

The Failure to Use Gender Information in Parsing: A Comment on van Berkum, Brown, and Hagoort (1999)

Marc Brysbaert^{1,3} and Don C. Mitchell²

We critically review the empirical evidence published by van Berkum, Brown, and Hagoort (1999a, b) against syntax-first models of sentence parsing. According to van Berkum et al., discourse factors and word gender information are used instantaneously to guide the parser. First, we note that the density of the experimental trials (relative to fillers) and the slow presentation rate of the van Berkum et al. design seem likely to have elicited the use of tactics involving rapid reanalysis of the material. Second, we present new data from a questionnaire study showing that the grammatical gender information of a relative pronoun in Dutch is often completely ignored, even during the wrap-up phase at the end of the sentence.

KEY WORDS: sentence parsing; gender information; syntax-first models.

Recently, van Berkum, Brown, and Hagoort (1999a, b) published a collection of empirical studies that are likely to have an appreciable impact because of the claims made. In particular, van Berkum *et al.* argued (1) that discourse instantaneously contributes to the parsing decision taken and (2) that gender information also contributes within a word of its appearance, although at a slightly later stage than contextual information. Both arguments contradict positions previously defended by the present authors (Brysbaert & Mitchell, 1996; Mitchell, Brysbaert, Grondelaers, & Swanepoel, 2000; Mitchell, Cuetos, Corley, & Brysbaert, 1995). In our view, (1) the syntactic parser functions autonomously and discourse factors slip in during a second reanalysis stage and (2) gender information of the nouns is initially ignored by the syntactic parser, as well, and only taken into account during the second stage.

¹ Vakgroep Experimentele Psychologie, Universiteit Gent, H. Dunantlaan 2, B-9000 Gent, Belgium.

² Department of Psychology, University of Exeter, Exeter, Devon EX4 4 QG, United Kingdom.

³ To whom all correspondence should be addressed. email: Marc.Brysbaert@rug.ac.be

VAN BERKUM ET AL. FINDINGS

van Berkum *et al.* base their claims on the Dutch equivalent of the following ambiguous structure

- (1) *David told the girl that . . .*

In the English version of this structure the word *that* is temporarily ambiguous because it might begin a complement clause, as in “*David told the girl that there would be some visitors,*” or a relative clause, as in “*David told the girl that had been on the phone to hang up.*” Without context, the preferred interpretation of (1) is the complement clause; in a context with two possible referents (i.e., two girls have been introduced in the text before the critical sentence), there is a strong force toward a relative clause interpretation. The point at issue is when the discourse force intrudes, either instantaneously upon encountering the ambiguity (or even before the ambiguous word has been encountered; see below), or after the syntactic parser has made an initial proposal. We (e.g., Mitchell, 1987; Mitchell, Corley, & Garnham, 1992) have defended the latter view (see also other syntax-first models: e.g., Frazier, 1979, 1987). In contrast, van Berkum *et al.* argue they found evidence for the former (in line with other interactive and constraint satisfaction models; e.g., McClelland, 1987; Spivey-Knowlton & Sedivy, 1995; Trueswell, Tanenhaus, & Garnsey, 1994).

van Berkum *et al.* based their statements on the following quartets of sentences:

- | | | | |
|-------------------------------|-------------|--------------|-----------------------------------|
| (2) David vertelde het meisje | <i>dat</i> | er | visite kwam. |
| [<i>David told the girl</i> | <i>that</i> | there | <i>would be some visitors.</i>] |
| (3) David vertelde het meisje | <i>dat</i> | had | zitten bellen op te hangen. |
| [<i>David told the girl</i> | <i>that</i> | had | <i>been phoning to hang up.</i>] |
| (4) David vertelde de vrouw | <i>dat</i> | er | visite kwam. |
| [<i>David told the girl</i> | <i>that</i> | there | <i>would be some visitors.</i>] |
| (5) David vertelde de vrouw | <i>die</i> | had | zitten bellen op te hangen. |
| [<i>David told the girl</i> | <i>that</i> | had | <i>been phoning to hang up.</i>] |

These test sentences were preceded by one of two types of discourse context:

(6) one-referent context:

David had de jongen en het meisje gezegd hun kamer voor de lunch op te ruimen. Maar de jongen had de hele ochtend liggen slapen, en het meisje had voortdurend zitten bellen.

[*David had told the boy and the girl to clean up their room before lunch time. But the boy had stayed in bed all morning, and the girl had been on the phone all the time.*]

(7) two-referent context:

David had de twee meisjes gezegd hun kamer voor de lunch op te ruimen. Maar het ene meisje had de hele ochtend liggen slapen, en het andere had voortdurend zitten bellen.

[*David had told the two girls to clean up their room before lunch time. But one of the girls had stayed in bed all morning, and the other had been on the phone all the time.*]

van Berkum *et al.* presented the context sentences auditorily and recorded event-related brain potentials (ERPs) as the participants read the critical sentences in a serial visual presentation paradigm at a pace of 600 ms per word. van Berkum *et al.* used ERPs because earlier research had suggested that a syntactic violation elicits a distinct positive shift in the ERP, labeled the syntactic positive shift (SPS) or P600. van Berkum *et al.* reported a very reliable interaction between sentences (2) and (3) and the preceding context. When the sentences were preceded by a one-referent context, an SPS/P600 effect was elicited by the disambiguating word *had* in sentence (3) (forcing a relative-clause interpretation), whereas in the presence of a two-referent context, the SPS/P600 effect was confined to the disambiguating word *er* of sentence (2) (forcing a complement-clause interpretation). In addition, the effect started within about 500 ms after onset of the disambiguating word. On the basis of these results, van Berkum *et al.* ventured the proposal that discourse factors are used immediately to disambiguate potentially ambiguous syntactic structures and are not confined to a second reanalysis stage.

In sentences (4) and (5), the RC/Comp ambiguity is eliminated because, when used as a relative pronoun, the word *dat* can only be employed to refer to a neuter word (i.e., a word with the definite article *het*—as in *het meisje dat*). When a non-neuter, or common word is used (taking the definite article *de*), then the relative pronoun *die* is required (i.e., *de vrouw die*). Thus, when a common noun precedes the word *dat* (as in 4), this word can only act as a complementizer. Despite this, van Berkum *et al.* reported that there still was an SPS/P600 effect associated with the word *dat* in (4), when the sentence was

preceded by a two-referent context. No effect was found when the participants read the word *die* in (5), either after a one or after a two-referent context.

To explain this pattern of results due to grammatical gender, van Berkum *et al.* (1999b, p. 563) envisaged the following sequence of processing events:

1. When processed as part of “David vertelde de vrouw dat . . .,” the word form *dat* makes available two morpholexical entries: the generic complementizer *dat*_{COMPL} and the relative pronoun *dat*_{RELPRO(NEU)}.
2. The parser momentarily ignores gender and considers the two candidate analyses: a complement- and a relative-clause interpretation.
3. In a two-referent discourse context, the parser provisionally commits itself to the relative-clause analysis.
4. Before processing the next word, the preferred analysis is checked on gender agreement.
5. To the extent that a two-referent context has biased the parser to pursue the relative-clause analysis, the resulting gender agreement violation forces it to abandon its preferred analysis again, leading to an SPS/P600 in that context in the ERP waveform to *dat*. All this happens within 500 ms after the word *dat* has been encountered.

IMPACT OF VAN BERKUM ET AL. FOR SYNTAX-FIRST MODELS

Although the van Berkum *et al.* pattern of results seems to provide impressive support for the theoretical position they adopt, a closer examination reveals that their findings are not so distant from claims made by syntax-first models in the past. We will start with three comments concerning contextual effects and then, in the next section, discuss the role of gender information in parsing.

First, there is the problem of how to disentangle context- from syntax-driven models (i.e., the moment at which discourse factors exert their influence). It should be noted that this issue cannot be settled in a simple and symmetrical way (do we find context effects as soon as the ambiguous part is encountered or not?). The criteria for arguing for or against context-driven models are likely to be different than the criteria for arguing for or against syntax-driven models. Claims to have dismissed the notion of contextual influences can be compromised by methodological shortcomings like (1) a failure to demonstrate that the contextual manipulations are strong enough to exert at least some kind of influence (e.g., in an off-line questionnaire task) or (2) materials have been presented at faster than normal pace, so that words are arriving too fast for contextual influences to be implemented. Conversely,

claims to have demonstrated contextual effects may be compromised by a different set of methodological shortcomings, such as (1) contextual effects have been probed too far downstream (i.e., beyond the second word of the ambiguous region; see Mitchell *et al.*, 1992 for a discussion) or (2) materials were presented at a slower than normal reading rate, so that readers may get bored with waiting for new words to arrive and start using context to guess what is coming. As for the latter, it must be said that van Berkum *et al.*'s design contained a few characteristics that may have prompted just this strategy. For a start, sentences were presented word by word with a processing time of 600 ms per word. This is very long for the short, high-frequency, close-class words in the critical region (i.e., *dat er* or *dat had*; see Grondelaers & Brysbaert, 1996, for a discussion of the use of *er* in Dutch). Second, all possible combinations of sentences and contexts (for a total of 240) were presented in a single session, with only 40 filler sentences intermixed. Needless to say, it is likely that such a design prompts participants to pay particular attention to the manipulation made, not to mention the problem of sentence-to-sentence syntactic priming and short-term tuning (e.g., Pickering & Branigan, 1999; Brysbaert & Mitchell, 1996). Thus, while it appears to be the case that van Berkum *et al.* found early influences of discourse factors, they did so by using a paradigm that maximized the probability of such information being exploited as rapidly as possible.

Second, there is the question to what extent the SPS/P600 is a marker of the initial syntactic analysis, as conceived by syntax-first models. van Berkum *et al.* give the impression it is, but a closer look at the ERP literature shows that this is not a generally accepted assumption (see Coulson, King, & Kutas, 1998; Gunter, Stowe, & Mulder, 1997; Hahne & Friederici, 1999; Osterhout & Hagoort, 1999). In particular, the question has been raised whether the SPS/P600 is a distinct effect or an instantiation of the P300 family. The P300 is a positive peak, emerging some 300 ms after the critical information has been encountered and its amplitude is a function of the probability, salience, and informational content of the eliciting stimulus, with the most improbable, salient, and informative stimuli eliciting the largest-amplitude P300s. The peak latency of the P300 varies as a function of stimulus complexity and ranges from about 300 to 800 ms. From this, Coulson *et al.* (1998) concluded that the reported SPS/P600 effects are indicative of "the way that participants update the contextual models which govern their expectations" (p. 47). On such an interpretation, it could well be argued that SPS/P600 effects do not tap into the first stage of the *initial proposal*, but into the effects that, according to the syntax-first models, happen in the second, reanalysis phase.

Finally, the overlap of van Berkum *et al.* with positions we previously defended can be illustrated by comparing their five-step explanation of the

grammatical gender effect with the proposals Mitchell (1987, 1989) made concerning the use of verb information in sentences like “After the child sneezed the doctor . . .,” in which the intransitive verb *sneeze* indicates that *the doctor* cannot be the direct object of the verb, but has to be the subject of the main clause. Still, Mitchell observed a garden-path effect for this type of sentences because readers initially interpreted the phrase *the doctor* as the direct object of the verb *sneezed*. To account for this lack of lexical influences on sentence parsing, Mitchell proposed the following explanation:

1. When processed as part of “After the child sneezed the doctor . . .,” the word form *sneezed* makes available two lexico-syntactic entries: the transitive verb form and the intransitive verb form.
2. The parser momentarily ignores subcategorization information and considers two candidate analyses (taking the following NP as the direct object of the verb, versus closing the preposed adverbial clause and taking the following NP to be the subject of the matrix clause).
3. As a result of the principle of Late Closure, the parser provisionally commits itself to a transitive-verb analysis (with the following NP *the doctor* attached as direct object of *sneezed*, since it is the “phrase or clause currently being processed”).
4. Before proceeding beyond the following NP, the preferred analysis is checked against the subcategory information of the verb.
5. To the extent that Late Closure has biased the parser to pursue the transitive-verb analysis, the resulting subcategory violation forces it to abandon its preferred analysis again, leading to extended reading times (and very probably, associated P600 effects).

The only difference between Mitchell (1987, 1989) and van Berkum *et al.* (1999a, b) is that Mitchell placed steps 4 and 5 in the second, reanalysis stage, whereas van Berkum *et al.* consider them as part of the initial analysis. Their argument is mainly based on time constraints, but as we have made clear before, this may be less convincing than apparent at first sight. What the differential SPS/P600 effect shows, is that context in this particular experiment (with its less fortunate characteristics) is beginning to have some effect at the probe point. However, the effect could very well be due to reanalysis, so that we can learn less from the effect than the authors had hoped (and claimed). An analogy may be drawn with the Mitchell *et al.* (1992) study on which van Berkum *et al.* was based. Mitchell *et al.* (1992) paid virtually no attention to the one- versus two-referent differences in garden pathing in the relative-clause structures, because they acknowledged that these could very well be due to reanalysis effects. The finding they took to be more informative and noteworthy was the fact that they were

able to demonstrate a garden-path effect in the relative-clause condition, even when the continuation occurred in a supposedly favorable two-referent context. Since an early (discourse-driven) context effect would have been expected to direct the parser right from the start toward the relative clause rather than the complement-clause interpretation, they took this as evidence that the discourse context effect could not have exerted its influence early enough to prevent some competing (presumably syntax-first) mechanism from creating a bias in favor of the complement reading. It may be noted that van Berkum *et al.* did not look directly for a garden-path effect of a relative-clause interpretation in the two-referent context, relative to an unambiguous control.

In short, it is our contention that the recent van Berkum *et al.* findings do not provide a clear-cut resolution of the general question of whether discourse information is used to *direct* parsing or merely to introduce a check on decisions made on the basis of other considerations. On the question of the use of gender information, our reservations are even more marked. In this case, our contention is not just that the data are equivocal, but that van Berkum *et al.*'s gender findings may have little to do with syntactic parsing because they may be more in line with a semantic interpretation.

USE OF GENDER INFORMATION BY THE SYNTACTIC PARSER

According to the account outlined above, van Berkum *et al.* maintain that syntactic gender information is brought into play within the time occupied by the presentation of no more than a single word in their task. While this claim appears to be compatible with their own data, it is not consistent with data of our own that are discussed in their paper. Specifically, we (Brysbaert and Mitchell, 1996; Mitchell *et al.*, 2000) found no evidence at all that Dutch readers made use of grammatical gender information of relative pronouns, even after several words had elapsed since the delivery of this information.

We made use of the following structure (COM = common noun, NEU = neuter noun):

- | | | |
|-------------------------------------------------------|------------------------------------------------------------|-----------------------------------------------------------------------|
| (8) De zoon | van de actrice | die op het balkon zat . . . |
| [<i>The</i> _{COM} <i>son</i> _{COM} | <i>of the</i> _{COM} <i>actress</i> _{COM} | <i>who</i> _{COM} <i>was on the</i>
<i>balcony . . .</i>] |
| (9) Het zoontje | van de actrice | dat op het balkon zat. . . |
| [<i>The</i> _{NEU} <i>son</i> _{NEU} | <i>of the</i> _{COM} <i>actress</i> _{COM} | <i>who</i> _{NEU} <i>was on the</i>
<i>balcony . . .</i>] |
| (10) Het zoontje | van de actrice | die op het balkon zat . . . |
| [<i>The</i> _{NEU} <i>son</i> _{NEU} | <i>of the</i> _{COM} <i>actress</i> _{COM} | <i>who</i> _{COM} <i>was on the</i>
<i>balcony . . .</i>] |

Whereas in sentence (8) there is ambiguity as to which NP the relative clause should be attached (i.e., who was on the balcony?), no such ambiguity is present in sentences (9) and (10), where the grammatical gender of the relative pronoun unambiguously refers to NP1 or NP2. Using sentences like (8), we found that participants prefer to attach the relative clause to NP1 and require considerable extra reading time when the relative clause disambiguates toward NP2 attachment. Sentences (9) and (10) were added to see what would happen when the relative pronoun itself was disambiguating. Much to our surprise, however, we noticed that the gender information was completely ignored and that sentences (9) and (10) were processed very much like sentence (8). That is, there was a considerable time cost for sentence (10) relative to sentence (9), but this cost only arose when the disambiguating information within the relative clause was encountered. The effect was obtained in two eye-tracking studies (Brybaert & Mitchell, 1996, Experiment 3; Mitchell *et al.*, 2000) and in two self-paced reading studies (Brybaert & Mitchell, 1996, Experiment 2; unpublished data).

On the basis of these (and related) findings, we argued that the syntactic parser initially ignores fine-grained lexical information, including gender information (Mitchell *et al.*, 1995). Only in the second, reanalysis stage can this information be picked up and integrated with thematic and discourse information. In order to find out how prominent the information is at this stage, we ran a new questionnaire study. One hundred participants (first-year university students) got a printed list of 50 sentences each of which was followed by a question. Of the 50 sentences, 24 were test items of the types shown in (9) and (10). Thus, one-half of the participants would get a questionnaire with the trial:

- (11) Iemand schoot op **het zootje** van de actrice **dat** op de bank zat.
 Wie zat op de bank?
 [Someone shot *the son* of the actress *who* was sitting on the bench.]
 [Who was sitting on the bench?]

The other half of the participants would be given a questionnaire with the trial (according to a latin-square design):

- (12) Iemand schoot op het zootje van **de actrice die** op de bank zat.
 Wie zat op de bank?
 [Someone shot the son of *the actress who* was sitting on the bench.]
 [Who was sitting on the bench?]

Half of sentences had an attachment head that consisted of two human nouns (as in sentences 11 and 12); the other half of the sentences were of

the type nonhuman/human (e.g., *het boek van de jongen dat/die . . . [the book of the boy that . . .]*). If participants paid any attention to the gender of the relative pronoun, even if only during the wrap-up phase at the very end of the sentence, then they should have had no problem indicating the correct answer on the basis of the grammatical gender information. To elaborate on this, if van Berkum *et al.*'s account were correct, comprehension performance would be almost perfect. Although they argued for a delayed use of gender information, they maintained that the duration of the delay was no longer than the time it normally takes for a person to read one further word (beyond the gender-marked word). In all the sentences employed in the present study, this would have allowed gender information to exert its influence well before the end of the sentence. At the other extreme, if grammatical gender information was completely ignored, we could expect some 50% syntactically correct answers. Reality was in-between the two extremes with 79% syntactically correct answers (see Table I).

The mere existence of so many errors in this simple, off-line study is rather surprising, because we had every reason not to expect mistakes. The rule of grammatical correspondence between nouns and relative pronouns is intensively taught in primary school and high school and participants could easily pick up this information to make their performance flawless (as some ten of our participants, indeed, did). Still, the majority of the participants paid very little attention to the disambiguating grammatical information, even during the final wrap-up at the end of the sentence. This finding agrees with other recent evidence that readers, in many cases, do not make use of all syntactic cues to build a fully precise syntactic representation of a sentence, but content themselves with good enough representations (Ferreira & Henderson, 1999). An objection to our interpretation that participants often failed to pick up the syntactic information, might be that the gender information was available for use at some early point in sentence-processing but was overruled before the reader made a decision about the final analysis. There are, however, two reasons for doubting this interpretation. First, parsing studies typically show that in the absence of conflicting evidence, readers tend to stick to an initial or early analysis of the material rather than

Table I. Percentages Syntactically Correct Attachment of the Relative Clause

		relpro refers to NP1 (%)	relpro refers to NP2 (%)
h/h	<i>die</i>	84	55
	<i>dat</i>	96	76
nh/h	<i>die</i>	76	80
	<i>dat</i>	97	63

abandoning it in favor of an alternative. Given this tendency, the early use of gender information would have been expected to prevent the appearance of errors identified in this study. Second, and somewhat more directly, on-line studies using eye-tracking and self-paced reading techniques have failed to show any sign of *early* use of gender information (e.g. Brybaert and Mitchell, 1996). Thus, given that the evidence argues against both early and late use of gender information, there is little reason to suspect that this kind of information plays a full role in sentence interpretation at *any* point in the process (at least in the case of sentences of the kind under scrutiny here).

Further examination of the errors in our questionnaire study revealed that the mistakes were not completely randomly distributed, but were a function of three factors: (1) whether the relative clause referred to NP1 or to NP2, (2) whether the relative pronoun was *die* or *dat*, and (3) whether the head type was h/h or nh/h. Table I lists the percentages correct answers as a function of these three variables.

The overall NP1 advantage (88 vs. 68% correct; $F1(1,98) = 54.38$, $MSe = .147$, $p < .01$; $F2(1,23) = 38.02$, $MSe = .0174$, $p < .01$) agrees with the bias of the syntactic parser to attach relative clauses high when confronted with the structure NP1-van-NP2-RC in Dutch. The effects of the other two variables can readily be accounted for by an extraneous semantic variable, namely, the tendency of the participants to associate a common relative pronoun (*die*) with a human noun, irrespective of the grammatical gender of the noun, and a tendency to associate a neuter relative pronoun (*dat*) with a nonhuman noun. This explains why the error rate is so high for sentences such as:

- (13) De journalist ondervroeg HET dochtertje van DE kolonel DIE . . .
 [*The journalist heard the daughter of the colonel who . . .*]
 (h/h, relpro *die* refers to the human NP2: 36% syntactically correct attachment)
- (14) De oude vrouw keek naar DE teddybeer van HET kindje DAT . . .
 [*The old woman looked at the teddy bear of the child that . . .*]
 (nh/h, relpro *dat* refers to the human NP2: 58% correct attachment)

In (13), observers seemed to accept that the common relative pronoun *die* can refer equally well to “the daughter” as to “the colonel” (who both have a non-neuter sex), irrespective of the fact that all diminutives in Dutch are neuter and should be referred to by the relative pronoun *dat* according to the grammatical rule of gender agreement. Similarly, in (14) there is a conflict between the syntactic rule saying that the relative clause starting with the neuter pronoun should be attached to “the child,” and the semantic notion that the sex of a teddy bear is more likely to be neuter than the sex of a child.

The semantic component also explains why performance is better for sentences like:

- (15) De student staarde naar DE broer van HET kamermeisje DAT . . .
 [*The student stared at the brother of the servant girl who . . .*]
 (h/h, relpro *dat* refers to NP2: 86% correct attachment)
- (16) Een dief loerde naar HET koffertje van DE toerist DIE . . .
 [*A thief leered at the suitcase of the tourist that . . .*]
 (nh/h, relpro *die* refers to NP2: 92% correct attachment)

Following two human nouns with a non-neuter sex, the neuter relative pronoun *dat* in (15) is rather unexpected on semantic grounds and, therefore, is a stronger syntactic marker for correct attachment of the relative clause. In the same way, the common relative pronoun *die* in (16) semantically disagrees with the non-human nature of NP1 and, hence, is stronger force against the predominant tendency to attach high.

The important points to note from this questionnaire study are (1) in an unexpectedly high percentage of cases, the syntactic information included in the relative pronoun is ignored all together, even in an off-line questionnaire study and (2) when the information is not ignored, the attachment decision is based more on a semantic fit between the “sex” of the relative pronoun (neuter or not) and the nature of the two nouns (human or not). These two observations agree well with the view that grammatical gender information is not used by the syntactic parser for the proposal of a structure, but can be picked up by the thematic controller for the disposal of incorrect interpretations.

van Berkum *et al.* (1999b, p. 569) proposed two reasons why their results (showing a gender disambiguation effect) differed from ours (no gender disambiguation effect). The first is that our measures may not have been sensitive enough, and that we should have used ERPs to look for an effect. This explanation is not persuasive, however, because our argument is not based on the absence of an effect, but on the fact that supposedly informative gender information failed to direct readers away from a misinterpretation, even in an off-line questionnaire task. The second explanation van Berkum *et al.* offered, was that the difference may be due to the type of ambiguity being resolved: RC attachment vs. RC/comp interpretation. There is indeed some evidence for this interpretation. Several serial-processing theorists have suggested that there may be more of a role for parallel processing in RC attachment than there is in other forms of ambiguity resolution. Construal theory (Frazier & Clifton, 1996), for instance, proposes that both RC-attachment readings are kept in play until relevant disambiguating information comes along (see also Traxler, Pickering, & Clifton, 1998).

However, to these two explanations, we would like to offer a third one. Notice that sentence (15) has a marked resemblance to sentence (4) of van Berkum *et al.* In both sentences, there is not only a syntactic incompatibility between the gender of the relative pronoun and the preferred analysis, but, in addition, a semantic discrepancy between a neuter-relative pronoun and the fact that the preferred head noun refers to a person. So, the reason why van Berkum *et al.* thought they had found an effect of grammatical gender, may very well have been the fact that a relative pronoun interpretation of their complementizer *dat* was semantically incongruent with the preceding noun phrase, which always referred to a human (or a humanlike creature, such as a fairy or a gnome). If this explanation is accepted, then there is no incompatibility at all between van Berkum *et al.*'s findings and ours, but this still means that grammatical gender is not taken into account in parsing.

CONCLUSIONS

We argue that the reason van Berkum *et al.* (1999a) were able to demonstrate discourse effects near the beginning of an ambiguous region was that the methodological details of their experiment allowed or even encouraged this to happen. The density of experimental trials (relative to fillers) and the slow presentation rate (600 ms per word) seem likely to have introduced priming effects and to have allowed time for the use of tactics involving reanalysis of the material. We would expect that with presentation rates approaching those associated with normal reading speeds and more varied material sets, the results could well have been compatible with the predictions of syntax-first models.

With respect to the evidence concerning the use of gender information, we note the additional evidence for our claim that even when readily available gender information may not always be used immediately to guide parsing processing (cf. Brysbaert and Mitchell, 1996). We present new data which challenge van Berkum *et al.*'s (1999b) claim that the use of grammatical gender information is postponed no longer than for the time it takes to read a single word. The results of our questionnaire study strongly suggest that sentence processing data cannot be explained by a mere *delay* in using grammatical gender information of a relative pronoun. Rather, it seems that some readers at least failed to make any effective use of this information *at all* and relied on a semantic heuristic in their interpretation of the sentence (i.e., the relative pronoun *die* refers to humans and the relative pronoun *dat* to nonhumans). Of course, this finding does not rule out the possibility that grammatical gender information may be brought into play in resolving certain ambiguities. However, it undermines the claim that

such information is *always* used this rapidly. It is likely that a variety of ambiguities will have to be examined systematically before we can reach a clear understanding of the factors determining failures or delays in the use of gender information in parsing. Having said this, the common ground between the two perspectives is that, for some as yet unexplained reason, grammatical gender information does not appear to play as rapid and efficient a role in guiding syntactic processing as might have been expected from the formal constraints such cues place on the structures of sentences.

REFERENCES

- Brysbaert, M., & Mitchell, D. C. (1996). Modifier attachment in sentence parsing: Evidence from Dutch. *Quarterly Journal of Experimental Psychology*, *49A*, 664–695.
- Coulson, S., King, J., & Kutas, M. (1998). Expect the unexpected: Event-related brain responses to morphosyntactic violations. *Language and Cognitive Processes*, *13*, 493–513.
- Ferreira, F., & Henderson, J. M. (1999). Good enough representations in language and perception. Paper presented at the *Architectures and Mechanisms for Language Processing Conference in Edinburgh*, September 1999.
- Frazier, L. (1979). On comprehending sentences: Syntactic parsing strategies. PhD Dissertation, University of Connecticut. Bloomington: Indiana University Linguistics Club.
- Frazier, L. (1987). Sentence processing: A tutorial review. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 559–586). Hove UK: Erlbaum.
- Frazier, L., & Clifton, C., Jr. (1996). *Construal*. Cambridge, Massachusetts: MIT Press.
- Grondelaers, S., & Brysbaert, M. (1996). De distributie van het presentatieve er buiten de eerste zinsplaats. *Nederlandse Taalkunde*, *1*, 280–305.
- Gunter, T. C., Stowe, L. A., & Mulder, G. (1997). When syntax meets semantics. *Psychophysiology*, *34*, 660–676.
- Hahne, A., & Friederici, A. D. (1999). Electrophysiological evidence for two steps in syntactic analysis: Early automatic and late controlled processes. *Journal of Cognitive Neuroscience*, *11*, 194–205.
- McClelland, J. L. (1987). The case for interactionism in language processing. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 3–36). Hove UK: Erlbaum.
- Mitchell, D. C. (1987). Lexical guidance in human parsing: Locus and processing characteristics. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 601–618). Hove UK: Erlbaum.
- Mitchell, D. C. (1989). Verb guidance and other lexical effects in parsing. *Language and Cognitive Processes*, *4*, S1123–154.
- Mitchell, D. C., Brysbaert M., Grondelaers, S., & Swanepoel, P. (2000). Modifier attachment in Dutch: Testing aspects of Construal Theory. In A. Kennedy, R. Radach, D. Heller, & J. Pynte (Eds.), *Reading as a perceptual process* (pp. 493–516). Oxford: Elsevier.
- Mitchell, D. C., Corley, M. M. B., & Garnham, A. (1992). Effects of context in human sentence parsing: Evidence against a discourse-based proposal mechanism. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 69–88.
- Mitchell, D. C., Cuetos, F., Corley, M. M. B., & Brysbaert, M. (1995). Exposure-based models of human parsing: Evidence for the use of coarse-grained (non-lexical) statistical records. *Journal of Psycholinguistic Research*, *24*, 469–488.

- Osterhout, L. & Hagoort, P. (1999). A superficial resemblance does not necessarily mean you are part of the family: Counterarguments to Coulson, King and Kutas (1998) in the P600/SPS-P300 debate. *Language and Cognitive Processes*, *14*, 1–14.
- Pickering, M. J., & Branigan, H. P. (1999). Syntactic priming in language production. *Trends in Cognitive Sciences*, *3*, 136–141.
- Spivey-Knowlton, M. J., & Sedivy, J. C. (1995). Resolving attachment ambiguities with multiple constraints. *Cognition*, *55*, 227–267.
- Trueswell, J. C., Tanenhaus, M. K., & Garnsey, S. M. (1994). Semantic influences on parsing: Use of thematic role information in syntactic disambiguation. *Journal of Memory and Language*, *33*, 285–318.
- Traxler, M. J., Pickering, M. J., & Clifton, C., Jr. (1998). Adjunct attachment is not a form of lexical ambiguity resolution. *Journal of Memory and Language*, *39*, 558–592.
- van Berkum, J. J. A., Brown, C. M., & Hagoort, P. (1999a). Early referential context effects in sentence processing: Evidence from even-related potentials. *Journal of Memory and Language*, *41*, 147–182.
- van Berkum, J. J. A., Brown, C. M., & Hagoort, P. (1999b). When does gender constrain parsing? Evidence from ERPs. *Journal of Psycholinguistic Research*, *28*, 555–571.