Priming the attachment of relative clauses in German

Report on our experiment for the project seminar
Psycholinguistische Forschungspraxis

Dr. Christoph Scheepers

Winter semester 2000/01

Garance PARIS
8th April 2000

Computational Linguistics
Department 4.7 General Linguistics
University of the Saarland

Contents

1 Introduction

1.1 Syntactic priming

Priming is a widely used method in psycholinguistics, and has been used to investigate the way human beings process a variety of linguistic phenomena, from phonology and morphology to lexical semantics. The context of priming is activation theory, which assumes that concepts that have some relation to each other are connected in some mental network, so that if one concept is activated, then concepts related to it are also activated (Dawson and Medler).

Priming's most commonly mentioned field of application is lexical semantics: here the presentation of one stimulus, termed the "prime", speeds up response to another stimulus, called the target, if both stimuli are related in some way (Sekerina, 2000). For example, after being shown the prime "nurse", subjects will react faster to the word "doctor" than to the word "butter" (Dawson and Medler).

Lexical priming deals with the comprehension of language by humans, but the use of the priming paradigm is wider: in the production of language, it can also serve to explore the natural tendency of speakers to repeat certain words or structures.

In combination with a sentence completion or picture description task, it is thus possible to observe under which circumstances people will repeat certain syntactic structures: this is called syntactic priming, and may involve questionnaires or oral descriptions (Bock, 1986[1]), as well as on-line methods which additionally measure reaction times (Corley and Scheepers, 2000).

1.2 A model to explain priming effects

Pickering and Branigan (1998) propose an activation model to explain observed syntactic repetitions. This consists of three levels of processing: a conceptual stratum where messages are generated, a lemma stratum responsible for syntactic construction, and a word-form

[1] J.K. Bock. Syntactic persistence in language production. *Cognitive Psychology*, volume 18, pages 355-387, 1986. Cited in Pickering and Branigan (1998).

stratum where the correct phonological and morphological form of the utterance are finally retrieved.

The lemma stratum is so called because it contains the base form of a word, along with the corresponding syntactic and semantic information. This information is made up of three types of nodes representing the category (word class), feature information (such as number, person, tense, gender, case), and combinatorial information (the subcat and thematic roles). There is one node for each category and combination possibility, and these are shared by different lemmas.

Pickering and Branigan (1998) suggest that syntactic priming can be explained by the activation of nodes on the lemma stratum. They assume production of a word activates the nodes connected to its lemma and the links in-between, and that this activation decays gradually over time. Therefore recently activated nodes have more chance than normal of being selected again.

The sharing of the combinatorial nodes is evidenced by the data Pickering and Branigan (1998) collected to test their model. However, the model "leave[s] it open whether the combinatorial nodes encode [a lemma's] subcategorisation frames or its overall combinatorial potential (i.e. with respect to both argument and non-argument phrases)". Pickering and Branigan (1998) discuss a few possibilities as to the exact nature of the information represented in the combinatorial nodes, but actually, little data is available concerning the behaviour of modifiers (adjuncts) in priming.

This was thus the reason for our exploratory experiment: we wished to observe whether it is possible to prime the attachment of relative clauses, i.e., of optional modifiers of a noun. Should our data show priming effects, it would then help specify what kind of information the combinatorial nodes of the model should carry.

1.3 Attachment preferences for relative clauses in German

To be able to observe priming effects, we needed a structure offering some kind of alternative or ambiguity. This is the case with the following sentence:

Someone shot the daughter of the colonel who was on the balcony

in which the relative clause can be attached to either *daughter* or *colonel*. Such ambiguities can be constructed in many languages, and similar sentences have often been discussed in the literature (for a review see Hemforth et al., 2000) with respect to determining the comprehension preferences in different languages. The results seemed to show that in most languages (including German), the preferred attachment is „high" attachment (i.e. attachment to the first noun), whereas in English, „low" attachment is preferred.

Of course, the preferences for comprehension can not be generalised to production, and the preferences for production in German have not been established (a corpus study in the Frankfurter Rundschau yielded no sentences of this type). On top of this, there is the problem that the relative saliency of both nouns may have an influence on the attachment preference: binding a relative pronoun to its antecedent is an anaphoric process, and Hemforth et al. (2000) mention Sanford and Garrod[2] saying that anaphors tend to go for salient discourse referents.

For this reason, we used a variant of our primes that did not end in a relative clause to observe the attachment preference in the target independent of any priming effects. We then respectively compared the proportion of high/low attachment targets in the high/low attachment prime condition with the proportion of high/low attachment targets after this primes in this control condition.

[2]S.C. Garrod and A.J. Sanford. On real-time character of interpretation during reading. *Language and Cognitive Processes*, volume 1, pages 43-59, 1985.
A.J. Sanford and S.C. Garrod. *Understanding written language: Explorations of comprehension beyond the sentence.* John Wiley and Sons, Chichester, England, 1981

2 Method

2.1 Materials

A sentence completion task was used in combination with the priming paradigm. Our items were prime-target pairs of incomplete subject-verb-object sentences, in which the object was a complex NP with a genitive noun modifying the head noun, followed by a relative clause (or by some other subordinate clause in the baseline condition) and a blank.

We constructed 24 such pairs of sentences, building on the fact that German relative pronouns must agree in gender and number with the noun they refer to, but that in some cases, the same form of the pronoun is used for different genders and numbers.

There were three variants of each prime: baseline, low and high attachment. Primes in the low attachment and in the high attachment conditions used gender information on the relative pronoun to constrain the possible attachment points of the relative clause, allowing only attachment to one of both sites.

Thus in

> *Der Zeuge erkannte den Fahrer des Autos, der/das/als…*
> 'The witness recognised the driver of the car, who/that/when…'

the masculine relative pronoun *der* can only refer to the head noun, the neuter *das* can only refer to *Auto*, whereas *als* initiates a temporal subordinate clause.

In the target sentences on the other hand, there were always two possible attachment points. Here we made use of the ambiguous nominative/accusative relative pronoun *die*, which can be either a feminine singular, or the plural for all genders. In all target sentences, either the head noun or the genitive was a plural, the other noun a feminine singular, followed by *die*, thus rendering the sentence locally ambiguous, as in:

> *Die Lehrerin sprach mit der Mutter der Kinder, die…*
> 'The teacher spoke to the mother of the children who…'

In the choice of the nouns, we did not take into account possible biases due to the fact that sometimes one noun might be more easily subject to modification by a relative clause than the other (this could be improved in a follow-up experiment, see p. 12).

Since the form of German verbs differs in the singular and plural, we should thus be able to determine from the verb ending whether the participant thought the relative clause modified the head noun or its genitive modifier.

2.2 Procedure

Three files were created, each containing one version of each of our 24 items, equally distributed among the three conditions (three times eight items in each condition). Each item appeared in a different condition in each of the three files.

Five fillers were added at the beginning of each file, and two between each sentence pair, so that all in all, there were 99 sentences to be completed. The fillers were sentence fragments ending in a blank: most were main clauses followed by a subordinating conjunction, some were conjunctions of two sentences, others incomplete simple sentences.

In each file, the prime-target pairs and the fillers were then randomised to produce ten questionnaires. The questionnaires were then formatted so that each fragment was followed by a blank line for completing the sentence.

Instructions on the first page of the booklets provided an example and asked the subjects to complete the fragments by building grammatical sentences. The subjects were told there were no right or wrong answers, but that they should complete the sentences rather quickly, with the first thing that came to mind.

2.3 Participants

30 subjects were paid 5 DEM to fill in our questionnaire. Completing all 99 sentence fragments usually took between 25 and 45 minutes.

2.4 Coding

Of course, not all primes were completed by the subjects as we had planned: sometimes participants produced an ungrammatical sentence. We thus coded the primes as “correct” whenever the participants produced an acceptable sentence and “incorrect” otherwise.

In some cases, participants also built other structures than the ones we had expected. For example, since German declension of the relative pronoun has the same form, *der*, for both the nominative masculine and the dative feminine, in some cases when one of both nouns was masculine and the other feminine, participants used the relative pronoun *der* as the dative object of a verb, and this resulted in priming attachment to the alternative site. In such cases, the prime was also coded as "incorrect".

The targets were coded according to two criteria: target type and target attachment. First we classified the relative clauses into subject relative clauses, object relative clauses and indeterminate, depending on the function of the relative pronoun in the subordinate clause. "Indeterminate" was used in cases like

Franziska sammelt alle Werke der Schriftstellerin, die sie so bewundert

$$\text{'Franziska}_i\text{ collects the works}_j\text{ of the author[fem.]}_k \begin{cases}\text{which}_{j,k}\text{ she admires'}\\ \text{which}_k\text{ admires her}_i\end{cases}$$

in which both the relative pronoun and *sie* could be either nominative or accusative, leaving unclear whether *die* is subject or object in the relative clause. This problem could in part have been avoided by not using feminine NPs as subjects of the main clause in the targets.

The classification into target types was helpful for coding target attachment. Indeed, given that German verbs agree in number with their subjects, whenever the participants completed the sentence as a subject relative clause, we could easily determine the attachment point to be high or low according to the number of the verb.

Unfortunately, for object relative clauses coding was not so easy and we had to resort to soft, semantic criteria, to decide what the attachment point was.

Actually, in some cases, syntactic ambiguities could not be even be lifted with the help of semantic criteria, and we then coded the attachment as "X" for "ambiguous". This was primarily the case with object relative clauses, as in the following variant of the above example: *Hans sammelt alle Werke der Schriftstellerin, die er so bewundert* (here *die* can only be the object of the relative clause, but it remains unclear whether the referent is *Werke* or *Schriftstellerin*).

We also used “X” for ungrammatical completions, incomplete sentences (e.g. when the participant forgot the subject), or when syntax and semantics conflicted, as in

> *Der Hausmeister reparierte die Türklingel der Mieter, die zu leise waren*
> 'The janitor repaired the doorbell of the tenants which were not loud enough'

in which the plural verb spoke for low attachment to *Mieter*, but the semantics seem to indicate attachment to *Türklingel*.

2.5 Predictions and design

As we saw in chapter 1.2, the model proposed by Pickering and Branigan (1998) makes no predictions about the attachment of relative clauses, since is does not specify what the nature of the combinatorial information would be. Intuitively though, we expected, if a priming effect could be shown at all, that the frequency of high/low attachment in targets following corresponding primes would be higher than that of high/low attachment in the baseline condition.

By excluding items in which the primes had been completed in an unintended way (see p. 6), we turned a quasi-experimental design into an experimental one: the prime completion factor was the result of pure experimental manipulation, excluding any influence of the participants response.

Since frequencies of attachment are categorical data, we used hierarchical log-linear models for our analysis. Log-linear models, unlike ANOVAs, do not make the assumption of a normal distribution, nor of linear independence between factor levels (a condition not fulfilled here, since an increase in high attachment responses would correspond to a decrease in low attachments and vice versa). On the other hand, they offer the advantage over standard chi-square testing that they can test simultaneously main effects and interactions within a multi-factorial design (Corley and Scheepers).

3 Results

Two rather strange questionnaires were excluded from the analysis at coding time: the first participant built an unusual number of ungrammatical sentences (missing subject), and it appeared from the vocabulary of the second that he was obviously not taking the experiment seriously. These questionnaires were replaced and new data was gathered to maintain sample size.

Taking only items with “correct” primes into consideration, we were then left with 234 (98%) valid cases in the high attachment condition, 229 (95%) in the low attachment condition and 231 (96%) in the baseline condition. Within this, the bulk of the target completions were subject relative clauses(91.4%), followed by object relative clauses (7.2%) and “indeterminate” (1.4%, see p. 6) cases, lumped with cases in which the completion was not a relative clause.

As we had expected, “indeterminate” trials were relatively rare (1.7%) in subject relative clauses, but relatively high (44%) in object relative clauses, probably for the reason that we had only semantic criteria to rely on in this latter case.

Overall, the targets contained 35.3% high attachment, 58.7% low attachment und 6% indeterminate completions, showing a preference for low attachment. Figure 3.1 shows the distribution of the target attachments in relation to the prime condition (again considering only valid trials).

3.1 Descriptive statistics

A look at the descriptive data suggests a higher proportion of high attachment targets following high attachment primes, compared to both the low attachment and the baseline prime conditions, although a slight preference for low attachment remains. On the other hand, we observe only a small increase in the expected direction of the low attachment targets as opposed to the baseline condition.

Figure 3.1: Percentage Target Completions

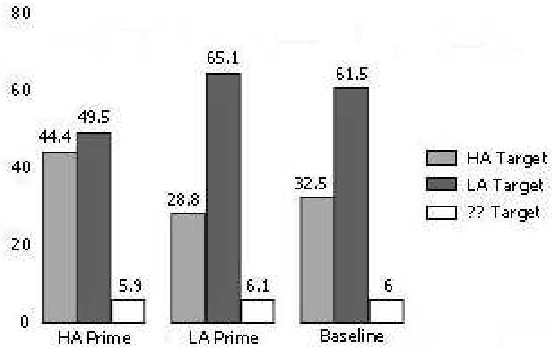

3.2 Inferential analysis

The factors we considered in the log-linear model analysis were prime condition, target attachment, and either participants (sample size N=30) or items (N=24). Results including participants as a random factor are reported as LRCS1, results including items as LRCS2 (where LRCS stands for "likelihood ratio chi square"). For each effect, we computed both partial and marginal associations: in the following we will report the most conservative of both measures.

Globally there was a significant prime type $\times$ target attachment interaction (LRCS1 = 12.012; df = 4; $p < .02$; LRCS2 = 12.347; df = 4; $p < .02$). However, the 3-way interactions Participants/Items $\times$ Prime Type $\times$ Target Attachment were not significant (ps $> .99$), allowing us to generalise over participants and items, and indicating a good fit between observed and expected frequencies in the saturated models.

Since the type of prime did not affect significantly the "indeterminate" attachments (ps $> .99$), these were left out in following analyses, high and low attachments being thus considered to be complementary.

Reflecting the descriptive analysis, the proportion of low attachment targets following high attachment primes was significantly smaller (and correspondingly the proportion of high attachments was higher) than those in the low attachment condition (LRCS1 = 10.913; df = 1; $p < .001$; LRCS2 = 11.216; df = 1; $p < .001$). The same effect was obtained for high attachment and baseline primes (LRCS1 = 6.369; df = 1; $p < .02$; LRCS2 = 6.539; df = 1; $p = .01$).

On the other hand, the slight numeric increase in low attachment targets after low attachment primes, compared to the baseline condition, did not reach significance (LRCS1, LRCS2 < 0.7; df $= 1$; $p > .40$).

3.3 Further analyses

We also tested whether the relative position of trials in the experiment (estimated by partitioning the trials into those occurring in the first half of a questionnaire and those in the second half) had an influence on target attachment. Nothing here reached significance (ps $> .24$), so that we can exclude "long term" priming (learning) as an explanation of our results.

Since we had not taken into account the modification preferences of the nouns when we generated our items (see p. 4), we looked again at descriptive patterns in the data after partitioning the items into those with a high attachment baseline preference (N $=$ 9) and those with low attachment baseline preference (N $=$ 15). Due to the small sample size, we could not do any inferential analysis, but had only the numeric tendencies to observe.

Considering only the items with a low attachment baseline preference, we observed a pattern that resembled very much that in figure 3.1. Interestingly though, the pattern of items with a high attachment baseline preference showed a greater contrast between the low attachment condition and the baseline condition. This could lead to suppose that the amplitude of the priming effect depends on the baseline preference, a possibility which will be considered in more detail in the discussion.

4 Discussion

4.1 The combinatorial nodes in Pickering and Branigan (1998)'s model

Let us now come back to the combinatorial nodes' issue mentioned in section 1.2. Pickering and Branigan (1998) discuss two possible representations of combinatorial information: subcategorisation frames or syntactic rules.

They rule out subcategorisation frames because this kind of representation could not explain why *The 747 was landing by the airport's control tower*, in which the PP is a locational adjunct, primes *The man is being stung by a bee*, in which the PP is an argument of the verb (Bock and Loebell, 1990)[1], since the subcat of a verb does not mention optional elements.

The syntactic rule hypothesis, however is retained, because some variants of phrase-structure grammars make use of the same rule for combining arguments and non-arguments with their verb. For example, in the Bock and Loebell sentence above, the rule that is applied to produce the prime and the passive target, $S \rightarrow NP\ V\ (PP)$, is different than the one that would be needed for the generation of the corresponding active target ($S \rightarrow NP\ V\ NP$), thus possibly explaining the priming effect.

Unfortunately, this does not offer a satisfying explanation of our data. In the case of our NP-NP-RC constructions, the rule that would be used to generate both high attachment and low attachment is $NP \rightarrow Det\ N\ (NP)\ (RC)$. To produce a high attachment, this rule would be applied twice as $NP \rightarrow Det\ N\ NP\ RC$ and $NP \rightarrow Det\ N$, for low attachment it would be applied as $NP \rightarrow Det\ N\ NP$ followed by $NP \rightarrow Det\ N\ RC$.

Thus after both high attachment and low attachment primes, the combinatorial node carrying $NP \rightarrow Det\ N\ (NP)\ (RC)$ would be activated, and this node would be selected again for the generation of both variants of the targets, so that higher activation of one node compared to some other node can not be used as an explanation of the priming effect. Consequently, some other representation of combinatorial information would be required if our data is to be explained within the model proposed in Pickering and Branigan (1998).

[1] J.K. Bock and H. Loebell. Framing sentences. *Cognition*, volume 35, pages 1-39, 1990.

4.2 The influence of markedness on priming

The analysis of our data in section 3.3 lead us to think the baseline preference could influence priming: when we considered only those items with a high attachment baseline preference, there was a numeric trend towards a greater contrast between the low attachment condition and the baseline condition.

This could mean that unpreferred attachments are more sensitive to priming than preferred ones. It should be noted that the term "preferences" mentioned throughout this paper does not pertain to personal preferences: rather preferences refer to the relative degree of markedness of linguistic elements. Many linguistic phenomena can be viewed as pairs, in which one element is typically more "usual", the other more "special". Linguists refer to both elements, respectively, as "unmarked" or "marked" elements. Typically in such a pair, the unmarked element is more frequent in texts, and cognitively easier to process (Waugh and Lafford, 1994).

The fact that a "marked" attachment should be more subject to priming would be relatively easy to explain: if marked elements are more difficult to process in general, and if previous activation due to the prime facilitates one variant of expression over another in the targets, then one could argue that the harder something is, the more it is subject to being facilitated – whereas things already easy cannot be made much more easier.

This could be the subject of a future experiment with more balanced items (half low attachment baseline preference and half high attachment baseline preference), in which the prediction would be that the priming effect in the low attachment condition should be as high as the priming effect in the high attachment condition.

Bibliography

Martin Corley and Christoph Scheepers. Log-linear analyses of response category frequencies. http://www.ed.ac.uk/~martinc/CSloglin.html - Consulted on 25th August 2001.

Martin Corley and Christoph Scheepers. Syntactic priming in english sentence production: Categorical and latency evidence from an internet-based study. *Psychonomic Bulletin & Review*, 2000.

Michael R.W. Dawson and David A. Medler. The university of Alberta's cognitive science dictionary. http://web.psych.ualberta.ca/~mike/Pearl_Street/Dictionary/contents/P/priming.html - Updated 24th February 1997, consulted on 1st September 2001.

Barbara Hemforth, Lars Konieczny, and Christoph Scheepers. Modifier attachment: relative clauses and coordinations. In Barbara Hemforth and Lars Konieczny, editors, *German sentence processing*, pages 159–183. Kluwer Academic Press, Dordrecht, NL, 2000.

Martin J. Pickering and Holly P. Branigan. The representation of verbs: Evidence from syntactic priming in language production. *Journal of Memory and Language*, volume 39, pages 633-651, 1998.

Irina Sekerina. Introduction to psycholinguistics - Mental lexicon and lexical processing. 7th EGG Summer School, Blagoevgrad - Course syllabus, 2nd August 2000.

L.R. Waugh and B.A. Lafford. Markedness. In R.E. Asher, editor, *The Encyclopedia of Language and Linguistics*, volume 5. Pergamon Press, Oxford, 1994.