



Syntactic structure assembly in human parsing: a computational model based on competitive inhibition and a lexicalist grammar

Theo Vosse, Gerard Kempen*

Experimental and Theoretical Psychology Unit, Department of Psychology, Leiden University, P.O. Box 9555, 2300 RB Leiden, The Netherlands

Received 23 September 1997; received in revised form 16 July 1999; accepted 15 January 2000

Abstract

We present the design, implementation and simulation results of a psycholinguistic model of human syntactic processing that meets major empirical criteria. The parser operates in conjunction with a lexicalist grammar and is driven by syntactic information associated with heads of phrases. The dynamics of the model are based on competition by lateral inhibition ('competitive inhibition'). Input words activate lexical frames (i.e. elementary trees anchored to input words) in the mental lexicon, and a network of candidate 'unification links' is set up between frame nodes. These links represent tentative attachments that are graded rather than all-or-none. Candidate links that, due to grammatical or 'treehood' constraints, are incompatible, compete for inclusion in the final syntactic tree by sending each other inhibitory signals that reduce the competitor's attachment strength. The outcome of these local and simultaneous competitions is controlled by dynamic parameters, in particular by the Entry Activation and the Activation Decay rate of syntactic nodes, and by the Strength and Strength Build-up rate of Unification links. In case of a successful parse, a single syntactic tree is returned that covers the whole input string and consists of lexical frames connected by winning Unification links. Simulations are reported of a significant range of psycholinguistic parsing phenomena in both normal and aphasic speakers of English: (i) various effects of linguistic complexity (single versus double, center versus right-hand self-embeddings of relative clauses; the difference between relative clauses with subject and object extraction; the contrast between a complement clause embedded within a relative clause versus a relative clause embedded within a complement clause); (ii) effects of local and global ambiguity, and of word-class and syntactic ambiguity (including recency and length effects); (iii) certain difficulty-of-reanalysis effects (contrasts between local ambiguities that are easy to resolve versus ones that lead to serious garden-path effects); (iv) effects of agrammatism on parsing performance,

* Corresponding author.

E-mail addresses: vosse@fsw.leidenuniv.nl (T. Vosse); kempen@fsw.leidenuniv.nl (G. Kempen)

in particular the performance of various groups of aphasic patients on several sentence types.
© 2000 Elsevier Science B.V. All rights reserved.

Keywords: Syntactic structure assembly; Computational model; Competitive inhibition; Lexicalist grammar

1. Introduction

Almost 30 years ago, Ron Kaplan (1972) and John Kimball (1973) set the stage for comprehensive psycholinguistic models of parsing, that is, models aiming to account for a broad range of human sentence processing phenomena rather than for isolated experimental effects. Kimball proposed a thorough and coherent collection of ‘parsing principles’. Kaplan implemented the first computational parser embodying several such principles. Since then, the validity of proposed parsing principles has been scrutinized in detailed experimental psychological research, and new parsing principles have emerged from this work. A prominent recent example is the Construal Principle of Frazier and Clifton (1996). Two important further developments in this period were sparked off by the rise of cognitive neuroscience. Neurocomputational models based on distributed and localist connectionism were applied to various syntactic aspects of language behavior (see the volumes edited by Reilly & Sharkey, 1992; Sharkey, 1992). Furthermore, functional brain imaging techniques and event-related brain potentials opened up an additional empirical window on the cognitive activity underlying human sentence processing (for detailed surveys see Friederici, 1999; Hagoort, Brown & Osterhout, 1999).

A variety of more or less comprehensive and detailed models of human sentence processing is presently available in the literature, and several of them have been implemented as computer simulation programs. On the basis of core concepts underlying them, these models may be divided – admittedly somewhat arbitrarily – into four groups.

- *Principle-based models*, which account for parsing preferences and complexity primarily on the basis of parsing principles such as Minimal Attachment, Late Closure, Construal, etc. (Berwick, Abney & Tenny, 1991; Crocker, 1992; Frazier & Fodor, 1978; McRoy & Hirst, 1990; Pritchett, 1992).
- *Resource-constrained models*, which emphasize the limited cognitive resources available to the human sentence processor and the ensuing upper bounds on acceptable linguistic complexity (Gibson, 1991, 1998; Haarmann & Kolk, 1991; Haarmann, Just & Carpenter, 1997; Henderson, 1994; Just & Carpenter, 1992; Lewis, 1993, 1996; Marcus, 1980).
- *Lexico-syntactic competition models*, focusing on the activation- or frequency-based competition between alternative attachment possibilities offered by syntactic building blocks retrievable from the mental lexicon (Jurafsky, 1996; Kempen

& Vosse, 1989; MacDonald, Pearlmutter & Seidenberg, 1994; Stevenson, 1993, 1994).

- *Distributed-connectionist models*, based on various types of neural network architectures, in particular on Simple Recurrent Nets (Chater & Christiansen, 1999; Christiansen & Chater, 1999; Elman, 1991; Tabor, Juliano & Tanenhaus, 1997).

Despite the extensive efforts that went into the development of these models, no single model has emerged to date that is effectively capable of parsing sentences and whose parsing performance simulates a broad range of phenomena of human sentence processing. The only models with a broad coverage of psycholinguistic phenomena are those developed by Gibson (1998) and Lewis (1993, 1996). The former model, although not implemented as a real parser, makes remarkably accurate quantitative predictions about the processing load effects of linguistic complexity. However, neither of these models has been applied to the parsing performance of language users other than normal adults, in particular, to aphasic patients suffering from agrammatism – the major neurological syndrome affecting syntactic performance.

In this paper we present the design, implementation and simulation results of a new model that, we claim, handles a significant range of normal and agrammatic parsing phenomena. It is a thoroughly revised version of the Unification Space model of Kempen and Vosse (1989). Key innovation is the introduction of lateral inhibition as the major force controlling the course of local competitions between mutually exclusive attachment options. The emphasis on competition is shared with the Competitive Attachment model of Stevenson (1993, 1994). However, in the latter model the competition is based on spreading activation rather than on lateral inhibition. Moreover, Stevenson only deals with the resolution of syntactic attachment ambiguities. Due to certain fundamental design assumptions, her model cannot handle effects of linguistic complexity, e.g. the contrast between center-embedded versus right-branching nestings of relative clauses.

Section 2 briefly lists the empirical and theoretical criteria the model aims to meet. In Section 3 we describe the model in detail. Sections 4 and 5 sketch the model's behavior and confront it with the empirical facts of normal and agrammatic sentence comprehension. Finally, in Sections 6 and 7 we present an overall evaluation of the model.

The description of the parsing model in this paper focuses on dynamic properties. The structures manipulated by the parser conform to a simplified version of the Performance Grammar developed by Kempen (2000).

2. Psycholinguistic criteria for parsing models

The central problem to be solved by any syntactic parser concerns lexical and syntactic ambiguity. Many lexical items belong to more than one word class or are associated with multiple sets of grammatical features within the same word class, and sentences often can be analyzed in several different ways even if every word is

unambiguous. In this section, we list the main types of empirical evidence and theoretical considerations bearing on how the human syntactic parser succeeds in dealing efficiently with massive disambiguation problems and on the circumstances that cause it to derail and fail. For details we refer to the various surveys of the literature on aspects of human syntactic parsing that have been published over the past few years (e.g. Frazier & Clifton, 1996; Garrod & Pickering, 1999; Gibson, 1998; Gorrell, 1995; Kempen, 1996; MacDonald et al., 1994; Mitchell, 1994; Pritchett, 1992). We provide many concrete examples in Sections 4 and 5 where the simulation runs are discussed.

It is generally assumed that parse tree assembly starts immediately upon arrival of the first word of the sentence, and that every new word is attached to the current – partial – tree as soon as possible (incremental processing). The human parser does not compute all possible parse trees for a given sentence, neither sequentially nor in parallel. Instead, it continually weighs up the viability of the competing analyses for the incoming sentence and spends most effort on the most promising ones. This strategy is risky, of course, because the low-ranking analyses may include the one intended by the speaker. In that case, the parser is forced to reconsider its viability rankings or even to re-parse the input completely.

Five groups of factors have been identified as the chief determinants of parsing success/failure rates in normal language users, of their parsing preferences, and of the time course of parsing processes.

1. *Complexity* of the linguistic input (center-embedding versus left branching versus right branching; sentence and constituent length; constituent order).
2. *Syntactic ambiguity* (various types of global and local attachment ambiguities; frequency differences between attachment alternatives).
3. *Lexical ambiguity* (multiple word senses and multiple part-of-speech membership of lexical items; also effects of frequency and contingent frequency on word sense and part-of-speech selection).
4. Lexical and syntactic *priming* (facilitation due to recent use of a lexical item or a syntactic structure).
5. *Semantic context* effects (*thematic fit*, subcategorization constraints; *pragmatic fit*, plausibility of a reading/analysis in the current linguistic and non-linguistic context).

Any quantitative model of human syntactic parsing should provide an account for all or most of the empirically established effects of these factors, or at least demonstrate the feasibility of such an account in extended versions of the model. The model we propose in this paper will not be concerned with the time course of the parsing process because as yet we have no straightforward analog of the processing resources consumed at various points in time. We do predict success/failure rates and preferences in parsing a wide range of sentence types under optimal and suboptimal processing conditions.

A source of syntactic processing evidence outside the normal range is provided by the comprehension performance of *aphasic* patients, in particular those affected by agrammatism. Because, with the exception of Haarmann et al. (1997), hardly any

computational model has taken agrammatic parsing phenomena into consideration, we discuss these here in some detail. Caplan, Baker and Dehaut (1985) (see also Caplan & Hildebrandt, 1988) present the data of an extensive study of sentence comprehension in 144 aphasics. These patients suffered from a variety of language disturbances, but their auditory comprehension of single words was relatively intact. They were presented with spoken sentences of different grammatical make-up and had to show their understanding by manipulating toy animals (enacting the thematic role assigned to each of the animals). Table 1 shows examples of the nine sentence types and the percentages of correct understanding. Successful models of sentence comprehension in normal language users should be able to simulate data like these by tuning some of their parameters.

3. The Unification Space

In this section we give a detailed description of the syntactic processing model that we have developed with the above criteria in mind. The model is an improved version of the Unification Space model published by Kempen and Vosse (1989). The most important innovation is the structure optimization mechanism, which is based on *lateral inhibition between incompatible attachment alternatives*. The model does not construct and manipulate syntactic trees directly. Instead, it creates and operates on a network of connections between nodes of ‘lexical frames’ (i.e. lexically anchored elementary syntactic trees) that are retrieved from the mental lexicon as ‘chunks’. Every connection in the network represents a potential attachment alternative. The *strength* of these connections varies continually, depending on the course of ongoing local competitions with other potential attachments. The dynamics of the model guarantee that the competitions will subside eventually and that the network enters a stable resting state. The parse is successful if the connections that emerge as winners of the competitions span a single grammatically correct syntactic tree covering the entire input sentence.

Table 1
Examples of stimulus sentences and comprehension scores in the Caplan et al. (1985) study^a

Active	A	The elephant hit the monkey	81.9
Cleft-subject	CS	It was the elephant that hit the monkey	80.2
Passive	P	The elephant was hit by the monkey	59.0
Dative	D	The elephant gave the monkey to the rabbit	58.7
Cleft-object	CO	It was the elephant that the monkey hit	51.1
Conjoined	C	The elephant hit the monkey and hugged the rabbit	45.0
Object-subject relative	OS	The elephant hit the monkey that hugged the rabbit	41.4
Dative passive	DP	The elephant was given to the monkey by the rabbit	37.7
Subject-object relative	SC	The elephant that the monkey hit hugged the rabbit	25.9

^a The sentences were taken from Table 1 in Caplan et al. (1985); the numerical data are percentages correct, averaged over the three sets of data reported in Tables 3, 9, and 15.

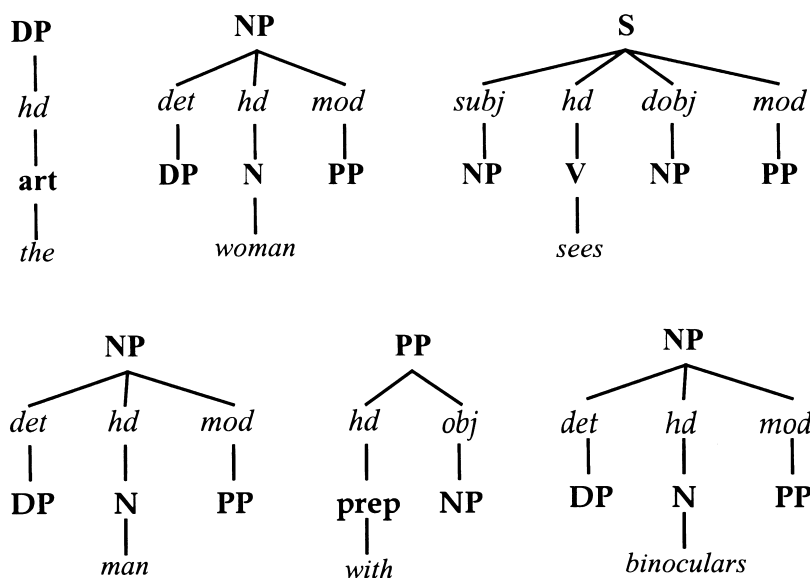


Fig. 1. Lexical frames associated with the words of sentence (1).

3.1. Grammar

The present version of the Unification Space utilizes a simplified version of the syntactic component of the Performance Grammar for English developed by Kempen (2000). This grammar is ‘lexicalist’ in the sense that all syntactic nodes are retrieved from the mental lexicon; there are no syntactic rules that introduce additional nodes. The lexical entries in this grammar are four-tiered mobiles called *lexical frames*. The examples in Fig. 1 correspond to the words of sentence (1). The top layer of a frame consists of a single *phrasal* node (called the ‘root’; e.g. S, NP, PP, DP¹), which is connected to one or more *functional* nodes in the second layer (e.g. SUBJect, HeaD, Direct OBJect, MODifier). At most one exemplar of a functional node is allowed in the same frame, except for MOD nodes which may occur several times. Every functional node dominates exactly one phrasal node in the third (‘foot’) layer, except for HD which immediately dominates a lexical (part-of-speech) node. Each lexical frame is ‘anchored’ to the lexical item which is printed below the HeaD node and constitutes the fourth layer. Every word of the language is head of a lexical frame (e.g. nouns head Noun Phrases, adverbs head Adverbial Phrases, a verb is the head of a Sentence, etc.).²

¹ DPs (Determiner Phrases) are part of NPs rather than – as in some other grammatical formalisms – the other way around.

² Kempen’s (2000) Performance Grammar does not use Verb Phrases (VPs) because the structural information encoded in VP nodes is covered by functional nodes (e.g. SUBJ, DOBJ) or by mappings between syntactic and conceptual relationships (e.g. ‘VP-modifiers’ and ‘S-modifiers’ map onto different types of conceptual relations).

The mobiles express important aspects of the subcategorization information of their lexical anchor. The left-to-right order of the branches is undefined. Constituent order rules are stated as simple linear precedence rules for branches of lexical frames. For example,

NP–det–DP < NP–hd–N

NP–det–DP < NP–mod–AP

NP–hd–N < NP–mod–PP

S–SUBJ–NP < S–HD–V

S–HD–V < S–DOBJ–NP

Checking the precedence rules is part of the parsing process. Not shown are the feature matrices associated with the phrasal and lexical nodes. They specify grammatical gender, person, number, case, etc., of the nodes.

(1) The woman sees the man with the binoculars.

The model's only structure assembly operation is called *unification*.³ A root node may become attached to an identically labeled foot node of another frame if two conditions are fulfilled. To begin with, their grammatical features must be compatible. For instance, The SUBJECT NP of the verb frame for *sees* carries the features ‘‘Case = nominative’’, ‘‘Person = third’’ and ‘‘Number = singular’’. These are compatible with those of the root NP node of *woman* (‘‘Person = third’’, ‘‘Number = singular’’ and ‘‘Case = nominative or accusative’’). The ‘*agreement check*’ between these NP nodes yields a positive result, so the first criterion for unification is met. The second requirement concerns *word order*. The root node of a lexical frame is allowed to unify with a foot node of another frame *only if this does not violate a precedence rule for the branch dominating the foot node*. Since *woman* precedes *sees* in the input string, and the above precedence rules order S–SUBJ–NP branches before S–HD–V branches, there is no word order impediment against attaching *woman* as the subject of *sees*. A ‘unification link’ (or U-link for short) between the NP node dominating *woman* and the SUBJECT branch of *sees* can now be established safely. Fig. 2 shows that until after processing the preposition *with* eight U-links are allowed by the present definition of the unification operation. Notice, in particular, that the first article is allowed to unify with both NPs (because neither of these unifications violates the above precedence rules for determiners; see the U-links numbered 1 and 4), and that the PP is U-linked to no less than three different nodes (U-links 6, 7 and 8).

We now turn to a systematic description of the competition mechanism whose

³ A detailed introduction to the notion of unification is given by Shieber (1986). The Unification Space model utilizes a simplified version: non-recursive feature unification.

duty it is to select a winner from among the sets of mutually incompatible U-links, in this case from {1, 4} and {6, 7, 8}.

3.2. Basic mode of operation

When a new word is read or heard, a copy of its lexical frame – or frames in case of lexical ambiguity – is entered into a workspace called *Unification Space*. Processing in the U-Space takes place in discrete time steps (cycles), and new input words are entered on a regular basis. The process of U-link formation between nodes of lexical frames starts immediately upon their arrival in the Unification Space. The structural preconditions on this process, in particular the agreement and word order checks, have been set out in Section 3.1. The dynamics of U-link formation, to be explained now, cause U-links to be formed gradually rather than instantaneously. A numeric value is associated with every U-link expressing the *strength* of the attachment (or quality, goodness of fit) between the pair of nodes at its extremes. The strength value of a U-link may vary between 0 (minimum) and 1 (maximum). It is initialized at 0 and, in the absence of counteracting forces, quickly rises to 1. Two nodes are considered to be attached if the strength of their U-link is at (or close to) 1. If the strength of their U-link is at or near 0, no attachment is supposed to exist. The dynamics of the system guarantee that intermediate strength values will ultimately be coerced towards either 0 or 1, so that no U-link remains in limbo.

The force that is capable of opposing the spontaneous increase of strength values over time is *lateral inhibition*. Two or more incompatible U-links (e.g. U-links 6, 7 and 8 in Fig. 2) mutually inhibit each other by emitting inhibitory force to their

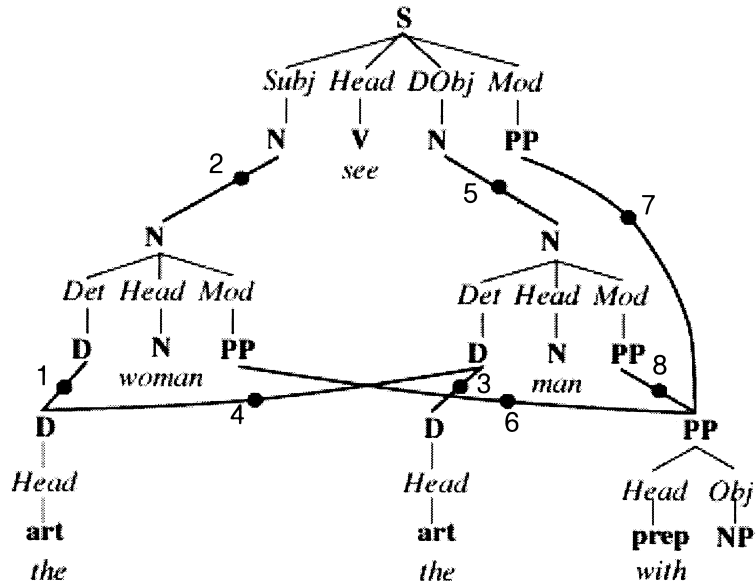


Fig. 2. Parsing sentence (1). Eight unification links ('U-links') proposed after processing PP with.

competitor(s), with the effect of reducing the strength of these competing U-link(s) by a certain amount. Whether or not two U-links are compatible depends primarily on grammar-independent conditions on ‘treehood’. For instance, since a root node is allowed to unify with just one foot node (and, vice-versa, one foot node with just one root node), two U-links trying to unify with the same node inhibit each other, and only the winner of this competition will belong to the final parse tree. Treehood conditions also forbid the formation of loops, that is, the creation of a U-link between the foot of a frame with the root of a frame that already occupies a higher position in the tree (or with the root of its own frame). In addition to treehood conditions there are *frame-based* conditions on compatibility. For example, certain lexical frames for verbs dictate the presence of either a direct object NP or a complement clause, but not both.

A parse is successful if, at the end of processing the last word of a sentence, a connected configuration of unified root–foot pairs exists that includes one lexical frame for every input word, and where every root node (except for the topmost one) is unified with exactly one foot node. Conventional syntactic trees can be derived from such configurations simply by picking the winning U-link landing on a root node, and merging this root node with the foot node at the U-link’s other extreme (so that they can be depicted as one single node).

This tree derivation procedure can be generalized slightly if there is a need to inspect the state of the parser before all local competitions have settled on a winner. In that case one picks, for every root node, the strongest U-link landing on that node, provided its strength is above threshold (e.g. higher than 0.50). Thus, tree derivation need not await the end of input. At any point in time during the parsing process for a sentence, it is possible to derive a single tree, or a set of trees, covering the input string up to and including the current word.

To sum up, the basic structure assembly operations take place in the U-Space proper, guided by two sources of constraints: information stored in lexical frames (syntactic categories and functions, grammatical features, word order constraints), and general, i.e. grammar-independent, conditions on treehood. The tree derivation component plays no role in the simulations; it merely serves as a visualization aid to external observers.

3.3. Step by step

When the lexical frame(s) of a new input word enter(s) the U-Space, Unification links are created from their root nodes to all matching foot nodes that are already present, and from all existing and matching root nodes to the newly entered foot nodes. Consider again sentence (1), whose analysis is illustrated in Figs. 1 and 2.

In Fig. 3a, the lexical frame of the article *the* has entered the Unification Space (left-hand panel). Then, a second word, the noun *woman*, is added. This frame contains an unfilled slot for the category DP. So, after entering, a U-link is proposed between the already present frame for *the* and the empty DP slot introduced by the frame for *woman* (right-hand panel). Fig. 3b shows that the situation is slightly more complex, in fact, because all lexical frames entering the U-Space are linked to an

ancillary top node called Apex. The U-links descending from Apex get involved in a struggle for the highest position in the syntactic structure. For example, if a sentence contains two verb forms interpretable as finite verbs (e.g. *raced* and *fell* in *The horse raced past the barn fell*), only one of them can become head of the topmost verb frame ('matrix clause'). In case of a successful parse, after the competitions have come to a halt, the U-link descending from Apex lands onto the root of the highest lexical frame in the hierarchy of frames that dominates the input sentence.

Multiple unification candidates presenting themselves to the same node are incompatible. For instance, when the PP frame for *with* enters the U-Space that is currently processing *The woman sees the man*, U-links are proposed between the PP and the NPs for both *the man* and *the woman* as well as for the S dominating *sees* (the links numbered 6, 7 and 8 in Fig. 2). Because these U-links are incompatible due to treehood constraints, they start inhibiting each other. Applying treehood constraints to the example in Fig. 1 yields the inhibition pattern depicted in Fig. 4. U-links 6, 7 and 8 inhibit each other because every root node can occupy at most one foot node as its attachment site. The inhibition between links 3 and 4 follows from the condition that foot nodes can accommodate no more than one root node. The competition between U-links 4 and 6 is due to a ban on crossing branches, as explained in Appendix A along with other treehood constraints. Every 'X' in the matrix of Fig. 4 corresponds to an *inhibitory link* between two U-links, as illustrated in Fig. 5.

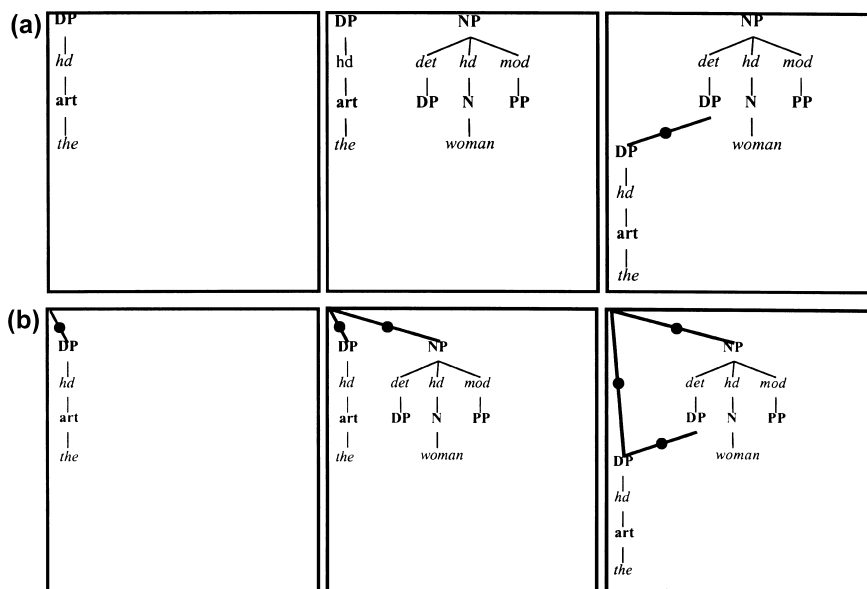


Fig. 3. (a) (Left) The U-Space containing the lexical frame of a single article, the first word of sentence (1). (Middle) The lexical frame for the second word is entered. (Right) A U-link between the two frames is proposed. (b) Unification links are proposed not only between nodes of lexical frames but also between the root node of every newly entered lexical frame and an ancillary top node ('Apex') represented by the upper left corner of the rectangle.

U-link	1	2	3	4	5	6	7	8
1				X				
2						X		
3				X				
4	X		X			X		
5								
6		X		X			X	X
7						X		X
8						X	X	

Fig. 4. Inhibition between the Unification links in Fig. 1.

3.4. Competitive inhibition

Upon entering the Unification Space, the nodes of lexical frames receive an *activation* value. The initial activation values are dictated by the mental lexicon

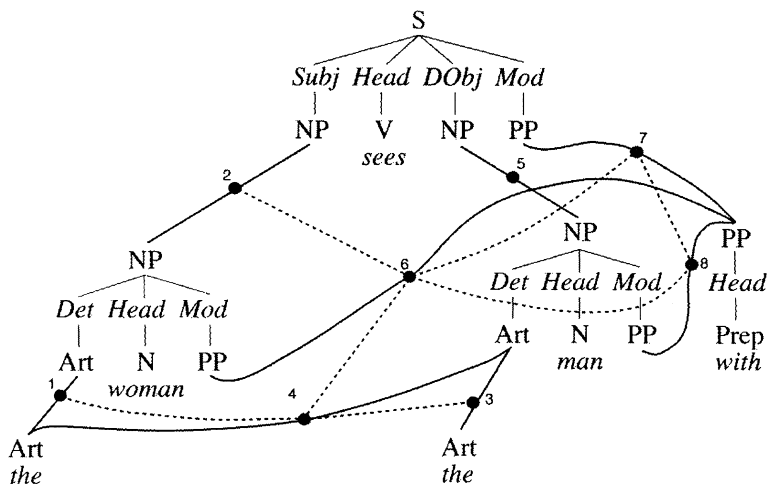


Fig. 5. Inhibitory links (dashed lines) between the U-links in Fig. 2. The preposition *with* is not yet attached. The option of attaching it to the first NP (*woman*, U-link 6) is inhibited from four different angles. Of the two remaining options, attachment of *with* to *man* via U-link 8 is the most likely outcome of this competition. NB: the articles are depicted here as individual nodes rather than as heads of DPs.

and depend, among other things, on the frequency of usage of the lexical frame. Activation gradually *decreases* towards the minimum value of 0. As said before, U-links between syntactic nodes are initially assigned the minimum *strength* value of 0, and strength *increases* spontaneously over time if unaffected by inhibitory force. Activation levels are not influenced by inhibition.

In every subsequent processing cycle, the strength value of a U-link is incremented by an amount proportional to the current activation of the root and foot nodes it connects (see Section 3.5 for mathematical details). However, this spontaneous strength increase is thwarted by the inhibitory forces emitted by competitors. A strong competitor (a U-link with a high strength value) emits a larger amount of inhibitory force than a weak competitor. Since a small random amount of *noise* is added to the summed inhibitory force emitted by the competitor(s), every competition is bound to produce a winner within a limited number of processing cycles. The winner will soon reach the maximum strength and push the strength value of its competitor(s) down to zero.

In terms of the example in Figs. 2 and 5, let us assume that, when the lexical frame for *with* is entered, the strength values of U-links 1, 2, 3, and 5 have already reached their maximum and U-link 4 has zero strength. At that point in time, U-links 6, 7, and 8 are created and start inhibiting each other because they share their root node (PP). Due to the gradual decay of activation, the activation level of the nodes in the ‘early’ NP frame for *woman* is lower than that in the verb frame for *sees*, while the nodes in the most recent NP frame for *man* have the highest activation level. Because the inhibitory potential exerted by a U-link depends on the activation levels of the root and foot nodes it connects, U-link 8 is capable of mounting a greater quantity of inhibitory potential than links 6 and 7, whereas U-link 6 is least capable of inhibiting its competitors. Therefore, U-link 8 will emerge as the winner of the competition, unless any of the others could profit from random noise. (Actually, the situation is slightly more complicated because U-link 6 is also inhibited by 2 and 4 due to the incompatibility constraint that prevents crossing branches. A branch from NP *woman* to PP *with* in the derived tree would cross the branch linking the second article with NP *man*; see also treehood constraints in Appendix A.)

3.5. Computation

The Unification Space contains two kinds of objects: syntactic nodes and Unification links.⁴ Syntactic nodes receive their initial activation value from the mental lexicon. From that level onwards, activation drops gradually with decay proportion d . The activation values are updated at regular intervals (every N cycles; see below). If syntactic node j is introduced at time t_j^0 with initial activation value v_j , its activation level at time t ($t \geq t_j^0$) is given by $n_{j,t} = v_j d^{t-t_j^0}$.

The arithmetic for computing the strength of a U-link is more complicated.

⁴ We will sometimes refer to the inhibitory effect of one U-link upon another as ‘inhibition link’ or ‘inhibitory link’. Such links might be considered as a third type of entity in U-Space.

- Let $r(i)$ and $f(i)$ denote the root node and the foot node that U-link i attempts to unify.
- Because we wish to confine the possible strength values of the U-link between a minimum and a maximum, we define a function C as follows

$$C(x) = \begin{cases} 0 & \text{when } x < 0 \\ 1 & \text{when } x > 1 \\ x & \text{otherwise} \end{cases}$$

- A (pseudo-)random function⁵ ρ returns a uniformly distributed value between -0.5 and $+0.5$.
- The cells of a matrix I specify the *inhibitory potential* between U-links i and j (cf. Fig. 4). The absence of inhibition between i and j is represented by $I_{ij} = 0$. The default value of the inhibitory potential of a U-link is $I_{ij} = 1$.
- The *strength* of a U-link is initialized at the minimum value 0. Then, during each processing cycle, the value is incremented by a proportion of the activation of the foot and root nodes it connects. The proportions are labeled $p_{\text{incr}}^{\text{foot}}$ and $p_{\text{incr}}^{\text{root}}$.
- Contrasting with the spontaneously rising strength of a U-link is its gradually diminishing *sensitivity* to inhibition. We assume that, from cycle to cycle, a U-link becomes less and less responsive to the inhibitory forces it is exposed to, so that its strength gradually consolidates. This decreasing sensitivity is – somewhat arbitrarily – supposed to depend on the activation level of the upper (foot)node of the U-link. It is therefore modeled by multiplying the incoming inhibition with the sum of a constant ($p_{\text{sens}}^{\text{const}}$) plus a proportion ($p_{\text{sens}}^{\text{foot}}$) of the current activation level of the foot node.
- During simulation runs with long sentences we discovered that the strength values of strong U-links that are exposed to high amounts of inhibitory force sometimes start oscillating. In order to damp these fluctuations, we decided to update U-link strength more frequently than activation. In fact, strength values are updated every cycle, but activation levels are updated every N th cycle, with $N = 5$. This is expressed in the following formula, which replaces the one at the beginning of this section: $n_{j,t} = v_j d^{\lfloor (t-t_j^0)/N \rfloor}$. Here, $\lfloor x \rfloor$ denotes x rounded down to the nearest integer.
- Together with the random noise that is added to the strength in every cycle, this gives the following formula for updating the strength of U-link i at time t

$$(u_{i,t} = C \left(u_{i,t-1} + \frac{(p_{\text{incr}}^{\text{root}} n_{r(i),t} + p_{\text{incr}}^{\text{foot}} n_{f(i),t}) + p_{\text{noise}} \rho - (p_{\text{sens}}^{\text{const}} + p_{\text{sens}}^{\text{foot}} n_{f(i),t}) \sum u_j I_{ij}}{N} \right))$$

The *inhibitory force* exerted by U-link i on its competitor j is defined as the product of its actual strength u_i and its inhibitory potential I_{ij} .

- So far we have assumed that the lexical frames of successive words of an input

⁵ The implementation utilizes a portable pseudo-random generator from Press et al. (1992).

sentence enter the U-Space at regular intervals. This is not entirely true. The model stipulates both a minimum (t_{\min}^{word}) and a maximum (t_{\max}^{word}) number of processing cycles between the entry of successive words. The maximum is used only when the total sum of inhibitory forces remains above a certain threshold, i.e. as long as

$$\sum_i \sum_j u_{i,t} u_{j,t} I_{i,j} > p_{\text{threshold}}$$

When this sum drops below the threshold, the next word is consumed without delay. A similar criterion holds for terminating the parsing process after the last input word. The process comes to a halt when the minimum (t_{\min}^{fin}) number of cycles has elapsed, given that the total sum of inhibitory forces drops below threshold, or otherwise after the maximum (t_{\max}^{fin}) number of cycles has elapsed.⁶

- Three additional model parameters are *co-inhib* (co-inhibition), n_{apex} and $\text{inhib}_{\text{apex}}$. (1) Co-inhibition enables two U-links that involve the same lexical frame to ‘conspire’ in order to enhance their common probability of survival. For example, consider a sentence like *Since the cat slept the dog bit her*, where an adverbial clause opens the matrix clause. The precedence rules for verb frames allow subordinating conjunction *since* to attach as a complementizer (CMPR) not only to verb frame *slept* but also to *bit*, as these rules merely require that a CMPR branch precede all other branches of the frame it belongs to – a condition which is fulfilled by both verb frames. Without special provisions, this would permit adverbial clauses to be stripped of their complementizer. What we need is a facility for promoting the joint survival of two U-links: in this example the U-link that connects *since* to verb frame *slept*, together with the one that connects this verb frame as an adverbial modifying clause to verb frame *bit*. The co-inhibition parameter is intended to give conspiring U-links a slight edge over their competitors. A small proportion of the strength of a conspirator is added to the strength of its conspiring partner (and vice-versa). This added strength enlarges the inhibitory force they can emit to their competitors and reduces the likelihood of the latter showing up in final parse tree. The co-inhibition parameter is listed in the lexical frames that are eligible for this facility (e.g. *since*), and from there it finds its way into inhibitory potential matrix I . (2) Parameter n_{apex} is the standard (low) activation level of the ancillary top node, Apex, which does not decay over time. (3) Similarly, $\text{inhib}_{\text{apex}}$ is the low inhibitory potential of U-links that descend from Apex.

⁶ The algorithm controlling the number of processing cycles for an incoming word does not aim to simulate the processing time required by novel input. The number of cycles may correlate with *syntactic* sources of processing difficulty, but one should realize that various non-syntactic (e.g. conceptual, referential) factors also contribute to the difficulty of processing new input words. See also the simulation of lexical ambiguity resolution discussed at the end of Section 4.

3.6. Parameter values

In order to estimate the 14 parameters⁷ of the model, we used simulated annealing as an optimization technique (Kirkpatrick, Gelatt & Vecchi, 1983). Monte Carlo simulations were run on eight ‘test’ sentences from the list in Table 3 below. For this set we defined an ‘optimal’ pattern of parsing results, that is, one reflecting parsing preferences and/or difficulties typically observed in human language users. We assume the following list can serve as an initial approximation.

- Complexity
 1. *The rat the cat chased escaped*: success rate 100%.
 5. *The rat the cat the dog bit chased escaped*: failure rate higher than 90%.
 7. *The cat chased the rat that escaped*: success rate 100%.
 8. *The dog bit the cat that chased the rat that escaped*: success rate 100%.
- Global syntactic ambiguity
 11. *John said he came yesterday*: success rate 100%, with low attachment of yesterday in more than 90% of the parses (i.e. yesterday modifying came rather than said).
- Local syntactic ambiguity
 16. *The man knew the woman slept*: success rate 100%.
 19. *Since the cat slept the dog bit her*: success rate 100%.
- Lexical ambiguity
 23. *They can fish*: success rate 100%, with either analysis (e.g. fish as noun or as verb) occurring between 10 and 90%.

Table 2 contains the set of parameter values delivered by the optimizer as conforming best to the basic pattern. These calibration values were entered in the U-Space program code and used throughout the simulation runs to be reported in Section 4, which address a wider range of human parsing phenomena in more detail.

3.6.1. Lexical fine-tunings

Lexical entries may contain information that overwrites the standard parameter values of Table 2. Certain types of U-links thus receive special values for certain parameters.

⁷ This number may seem rather big, considering that the simulation runs involve no more than 50-odd sample sentences (counting the 30 sentences in Table 3, the eight sentence types in Fig. 8, and several additional variants that we discuss *passim*). However, the skeptical reader should bear in mind that computing the desired parse trees for these sentences involves hundreds of competing attachments (Unification links), each carrying a strength value between 0 and 1. Even a simple sentence may require the instantiation of a dozen or more U-links. After the parsing process for an input sentence has come to a halt, the strength value of each of these U-links must be at (or close to) the maximum value if the U-link belongs to the correct parse tree, or at (or close to) the minimum value if it does not.

Table 2
Parameter values of the Unification Space^a

t_{\min}^{word}	7	d	0.98
t_{\max}^{word}	35	$p_{\text{noise}}^{\text{root}}$	0.18
t_{\min}^{fin}	9	$p_{\text{incr}}^{\text{root}}$	0.84
t_{\max}^{fin}	55	$p_{\text{incr}}^{\text{foot}}$	0.97
$p_{\text{threshold}}$	1.83	$p_{\text{sens}}^{\text{const}}$	0.65
$\text{inhib}_{\text{apex}}$	0.05	$p_{\text{sens}}^{\text{foot}}$	0.46
n_{apex}	0.09		
co-inhib	0.08		

^a The values in the first column were kept constant throughout the simulation runs described in Sections 4 and 5. The values in the third column were used only for the simulations reported in Section 4 (sentence comprehension in normals).

- U-links that descend from a foot node representing an obligatory constituent that always follows its head (e.g. a direct object NP after its head verb, or a prepositional object NP after its head preposition) summon up an above-average inhibitory potential. Thus, they are helped to fight off earlier or later lexical frames attempting to play these grammatical roles. In our current grammar of English, this value is 1.32.
- A similar measure helps the model to chose between alternative lexical frames that belong to the same ambiguous word, e.g. between *fish* as main verb and as noun. Two or more U-links that represent attachments of the root nodes of such alternative frames inhibit each other with an inhibitory potential of 1.11.
- The activation level of lexical frames is standardly initialized at the maximum level of 1. However, lexical entries may specify a lower initial activation level.

The values of the three types of lexical parameters did not result from systematic search of the parameter space but were chosen as convenient on the basis of informal simulation trials.

3.7. Tree derivation

Now that we have discussed the dynamic mechanism underlying the rise and fall of unification strength (u) in U-links, how can we derive syntactic trees from networks of Unification links? The following simple procedure enables us during any processing cycle, i.e. not necessarily at the end of parsing a complete sentence, to determine which tree, or which set of fragmentary trees, currently inhabits the Unification Space. Let us assume that root node $r(i)$ is unified with foot node $f(j)$ if the U-link connecting $r(i)$ to $f(j)$ has a strength value of at least half its maximum (that is, $u_{i,j} \geq 0.5$), and is the strongest one landing on $r(i)$. Tree derivation starts by establishing the unification partner (if any) of all root nodes. Any root node that either has no unification partner or is unified with Apex is the root of a syntactic tree. Such a tree is conveniently depicted without branches that end in a phrasal node rather than in a lexical item. These branches are pruned away and the remaining branches are linearized in accordance with input word order. However, one should

Table 3
Results of model simulations^a

Sentence	Score (%)
(1) The rat the cat chased escaped	S, 100
(2) The rat that chased the cat escaped	S, 100
(3) The rat that the cat chased escaped	S, 100
(4) The rat the cat with the binoculars chased escaped	S, 100
(5) The rat the cat the dog bit chased escaped	F, 100
(6) The rat the cat you bit chased escaped	F, 100
(7) The cat chased the rat that escaped	S, 100
(8) The dog bit the cat that chased the rat that escaped	S, 100
(9) The executive who the fact that the employee stole office-supplies worried hired the manager	F, 100
(10) The fact that the employee who the manager hired stole office-supplies worried the executive	S, 100
(11) John said he came yesterday	H, 1.2; L, 98.8
(12) The woman watches the man with the binoculars	H, 100; L, 0
(13) The horse raced past the barn fell	F, 100
(14) The horse raced past the barn yesterday	S, 100
(15) The stablehand groomed the horse raced past the barn	S, 100
(16) The man knew the woman slept	S, 100
(17) The man who knew the woman slept	H, 100
(18) Since Jay always jogs a mile seems a short distance to him	(see text)
(19) Since the woman slept the dog bit her	S, 100
(20) Since the horse kicked the dog bit her	S, 99; F, 1
(21) When the boys strike the dog kills	(see text)
(22) Before the boy kills the man the dog bites strikes	(see text)
(23) They can fish	VV, 54.4; VN, 45.6
(24) Without her contributions the funds are inadequate	S, 100
(25) Without her contributions are inadequate	F, 100
(26) I hate that	S, 100
(27) I hate that boy	S, 100
(28) I believe that John smokes annoys Mary	S, 100
(29) Before she knew that she went to the store	Pro, 51; CMPR, 49
(30) That coffee tastes terrible surprised John	(see text)

realize that the U-Space does not contain syntactic trees in the traditional sense but merely a network of Unification links between lexical frames. Trees are derived and drawn merely for the sake of notational and expository convenience.

4. Model simulations: basic phenomena of human syntactic parsing

In order to evaluate the model's performance, we tested it on a set of 30 sample sentences that, taken together, exhibit a large portion of the psycholinguistic phenomena reported in the experimental literature. Because of the probabilistic character of the simulations (due to the random noise parameter), each sentence

was run 1000 times.⁸ All simulation runs used the parameter settings listed in Table 2. Semantics played no role in the simulations. Unless stated otherwise, all activation levels were initialized at their maximum value of 1.

We present the simulation data in terms of the probability of parsing success or failure for cases where parsing difficulty is an issue. When ambiguity is at issue, proportions of alternative analyses are reported. Table 3 shows the results.

4.1. Complexity

Examples (1–8) show that self-embedded clauses present no problem up to one level of embedding. In case of a single nesting (see the snapshots in Fig. 6a,b), the number of U-links is relatively low, making it fairly easy for the second verb to force a solution upon the few ongoing local competitions – in particular the competition between the two attachment possibilities for *chased* allowed by the grammar: as head verb of a post-nominal object relative clause, or as matrix verb of the whole input sentence. Fig. 6c,d portray the essentially more complex situation arising in case of double self-embeddings. The number of inhibitory links between U-links rises steeply when the second verb enters the U-Space. This sparks off a large number of local competitions between competitors that are fairly well-matched in terms of inhibitory force: the lexical frames connected by these U-links were entered in brief succession and therefore have similar activation levels (e.g. *bit/chased* on one hand and *rat/cat* on the other). Consequently, such U-links can mount similar inhibitory forces against one another, thus making it hard for them to reach a settlement.

The right-branching counterparts of sentences (1) and (5) (examples (7) and (8)) are easy. The competitions now tend to be resolved quickly and correctly due to the imbalance between competitors. For instance, in terms of sentence (8), although the first NP (*the dog*) competes for the role of subject of the last verb (*escaped*), it will be defeated by the last NP (relative pronoun *that*) which is much more recent and therefore more active. Moreover, the unification of the first NP with the first verb's subject has gained considerable strength in the meantime and can only be undone by considerable inhibitory forces.

Sentence (4) shows that the all-or-none difference between single and double self-embeddings is not merely an effect of constituent length. Sentence (4) (a single self-embedding including a PP) has the same number of words as sentence (5) (double self-embedding) but is easy to analyze. This is due, among other things, to the fact that the preposition *with* and, a few words downstream, the NP *the binoculars* are immediately offered suitable attachment sites by their left-hand neighbors. The three corresponding input items in sentence (5) (*the dog bit*), on the other hand, find attachment sites only at a later point in time, viz. after the verb *bit* arrives.

In Section 6 we compare relative clauses with subject and object extraction –

⁸ This gives a 95% confidence interval ranging from $\pm 0.2\%$ (when P approaches either 0 or 1) to $\pm 3\%$ (when $P = 0.5$). So, when the simulation outcome is $P = 0$, the real P can be anywhere between 0 and 0.002 with a confidence of 0.95; when the simulation outcome is $P = 0.5$, the real P can be anywhere between 0.47 and 0.53 with the same confidence level.




Fig. 6. (a) Parsing sentence (1) (single nesting). Configuration of lexical frames and Unification links some time after the first verb (*chased*) was entered. The attachment of *the cat chased* as a relative clause modifying *the rat* is opposed by the U-link attaching *the rat* as the Subject of *chased*. NB: U-links descending from Apex have been removed for reasons of clarity. Furthermore, the lexical frame drawings have been simplified. First, the articles are depicted as single nodes rather than as heads of Determiner Phrases. Second, branches of a lexical frame that are not involved in a unification attempt are not shown. (b) Parsing sentence (1) (single nesting). Configuration immediately after the second verb was entered, yielding two more U-links (dashed lines). Because *escaped* is an intransitive verb, its S-frame cannot attach as an object relative clause. Therefore, it needs either of the two NPs as its subject, thus entering into competition with *chased*. The competition ends with *the rat* strongly attached to *escaped*, and *the cat* to *chased*. The number of inhibitory links between the eight U-links shown here is 21. (c) Parsing sentence (5) (double nesting). The configuration shown here exists sometime after entry of the first verb (cf. (a)). Five new U-links have been created (dashed lines). The Subject NP of *bit* attempts to unify with all three NPs. At the same time, the S-frame of *bit* tries to get unified with the two leading NPs as an object relative clause. (It cannot contract this relationship with the trailing NP because the grammar requires a subject NP preceding the finite verb of the relative clause.) (d) Parsing sentence (5) (double nesting). Configuration immediately after the second verb was entered (cf. (b)). The arrival of the transitive verb *chased* causes the addition of five U-links (dashed lines). The number of inhibitory links now rises to 80.

exemplified here by sentences (2) and (3), respectively – from the point of view of processing load.

Sentences (9) and (10) show that the model experiences great difficulty with a sentential complement clause embedded within a relative clause, whereas the reversed stacking, i.e. the relative clause embedded within the sentential complement, presents no problem. (See Gibson, 1998, for some human data concerning this contrast and for an explanation in terms of processing load.) The basic reason

why the U-Space performs so differently on these two sentences has to do with the attachment possibilities for the deepest clause (*that the employee stole office-supplies* in (9), and *who the manager hired* in (10)). In sentence (10) – the easier case – the verb *stole* has two important effects upon the competitions that are going on: it marks the end of the deepest clause, and it offers additional potential attachment sites (SUBJECT and Direct OBJECT) to the preceding NPs. The system rapidly resolves the competitions: *who the manager hired* is attached to *employee* as a relative clause, and *employee* becomes the SUBJECT of *stole*. *Employee* outwins the other NPs preceding *stole* in the competition for SUBJECT role, partly because its activation is higher (recency) and partly because its attachment to *stole* is not contested by a powerful competitor (*who* and *manager* are already strongly connected to *hired*). Moreover, the four remaining words of the sentence do not threaten this analysis – *worried* finds its SUBJECT in the sentence-initial NP.

Now consider the situation in difficult sentence (9) after the noun *office-supplies* has entered the arena. This word prompts an analysis with *that the employee stole office-supplies* as a finite clause. However, the attachment of this clause to *fact* is seriously challenged by a possible attachment to *executive* (as a post-nominal relative clause; cf. the legal NP *the executive that the employee stole office-supplies from*). This analysis is not precluded by precedence rules because these do not take intermediate constituents into account. Furthermore, the grammar allows more than one MODifier in an NP. Although *fact* is more recent than *executive*, it is also trying to become the SUBJECT of *stole*. This option is incompatible with its taking *that the executive stole office-supplies* as complement (because of treehood constraints) and, in effect, inhibits the attachment of the complement clause to *fact*.

In more general terms, sentences tend to become harder to analyze syntactically when a large number of U-links of roughly equal strength enter into competition with each other. Sentences are easier to parse when the number of U-links involved in the same local competition is small and of unequal strength, and when the solutions of different local competitions are consistent with each other.⁹

4.2. Global ambiguity

Sample sentence (11) demonstrates the strong preference for low attachment, caused by higher activation levels of more recent lexical frames. That no such recency effect becomes manifest in sentence (12) is caused by a strength manipulation we have applied. In the lexical frame for *watches* we have assigned a slightly higher strength value to instrumental modifiers (e.g. PPs headed by *with*), thereby indirectly increasing the inhibitory potential of the corresponding U-link link beyond its normal value. Due to this measure, the U-link between the

⁹ Note that these generalizations only concern syntactic parsing. They do not intend to rule out the possibility of non-syntactic (e.g. semantic and pragmatic) contributions to the processing difficulty of a sentence. See also footnote 11.

lexical frames for *with* and *watches* has an edge over the U-link connecting *with* and *man* (in terms of Fig. 2, U-link 7 tends to be stronger than U-link 8). This *verb bias* effect overrides the activation-based recency effect. Without the strength manipulation, the U-link between the most highly activated lexical frames is very likely to gain the upper hand (similar to what happens at sentence (11)).

4.3. Local ambiguity

The classic garden-path effect in sentence (13) is due to the fact that the lexical frame for *raced* as past participle is initialized at a lower activation value (0.8) than the frame for *raced* as finite main verb (1.0) – thus mimicking the frequency difference between these verb forms. In the absence of this frequency bias, the success rate rises to 100% because the two readings of *raced* are then able to inhibit each other's U-link with equal vigor, thereby preventing both these links gaining strength rapidly. This, in turn, enables the verb *fell* to dethrone *raced* as a potential matrix verb.

The model parses example (14) uniformly as a main clause as long as the entry level activation of *raced* as a past participle is around 0.8. However, when this value is raised to the maximum level of 1.0, the model sometimes (in 20% of the trials) delivers an NP analysis with the substring *raced past the barn yesterday* attached as a reduced relative clause.

Example (15) is unproblematic for the model, even if *raced* as past participle is initialized at low activation levels. This is at variance with informal observations by Stevenson and Merlo (1997), who noticed that the sentence is hard to comprehend. In order to reconcile the U-Space parsing architecture with such observations one needs to look for explanations external to the model. Stevenson and Merlo (1997) propose a detailed analysis – couched in the framework of the Competitive Attachment parser of Stevenson (1993, 1994) – departing from the special lexical properties of unergative manner-of-motion verbs like *race* that are used transitively.

In sentence (16), the NP following *knew* is temporarily functioning as this verb's direct object and is easily detached from there by the second verb, to become the latter's SUBJECT. Example (17) is analyzable not only as a main clause but also as an NP with a post-nominal relative clause. In the former reading, *slept* belongs to the main clause ('High Attachment', HA), while in the latter reading it is attached low as part of the relative clause. Recent experiments have revealed a strong preference for high attachment. Sturt, Pickering and Crocker (1999a,b) established this bias with materials like

(HA) *The detective who read the statement had contradicted himself and was a bit confused*

(LA) *The detective who read the statement had contradicted itself was a bit confused*

Viewing times for the reflexive pronoun were longer in the low attachment (*itself*)

than in the high attachment condition (*himself*). Schneider and Phillips (1999) obtained similar results with examples like

(HA) *The surprised woman who discovered the drunk man locked the front door herself with the spare key to keep him away*

(LA) *The surprised woman who discovered the drunk man locked the front door himself was amazed he could even walk*

Like human comprehenders, the U-Space model prefers High Attachment analyses. When parsing sentence (17), it never reanalyzes NP *the woman* and invariably produces the main clause parse. But then, why does it always apply reanalysis in case of sentence (16)? The answer: in (17), *slept* can satisfy its need for a Subject NP by unifying with clause-initial *the man* as its Subject NP. This unification is relatively easy because it is not contested by strong competing U-links. In (16), however, *slept* can only get hold of a Subject NP if it succeeds in detaching either *the man* or *the woman* from *knew*. Because the U-link that connects *the man* to *knew* as its Subject is older and therefore has gained more strength, it is capable of resisting the competition from *slept*. On the other hand, the U-link between *knew* and *the woman* is younger and easier to suppress. The result is reanalysis of the latter NP as the Subject of the complement clause.¹⁰

Sentence (18) is another well-known garden-path example (from Frazier & Rayner, 1982). The parsing problem is caused by the measure phrase *a mile*, which functions as the SUBJECT NP of the matrix clause but tends to become strongly attached to *jog* as an adverbial modifier. The model parses this sentence effortlessly, at least if the U-links corresponding to the two competing attachments of *a mile* are initialized at equal strength. However, parsing performance deteriorates dramatically when the ‘adhesion’ between *jogs* and the measure phrase modifier is raised. In the model, this requires enhancing the inhibitory potential of the corresponding U-link, as shown in Fig. 7. The graph shows that the likelihood of the model getting garden-pathed increases strongly with the adhesion between verb and modifier – an effect of verb bias (more specifically, of lexical frame preference; see also example (12)). The garden-path observed by Frazier and Rayner (1982) therefore may have been due, at least in part, to the effect of lexical frame preferences, e.g. a preference for *jog* to take a measure phrase, or, in several other experimental sentences, for the verb to take a direct object.

In line with this, sentence (19) with the strictly intransitive verb *slept* is unproblematic both for the U-Space and for human comprehenders (recent eye movement data collected by Adams, Clifton & Mitchell, 1998; see also Mitchell, 1987, and discussion by Corley, 1998). The verb *kicked* in example (20) was encoded as

¹⁰ The model’s behavior with respect to (16) should not come as a surprise because this sentence belongs to the test sample from which the parameter settings are derived (see Section 3.6). However, the simulation runs for sentence (17) were not anticipated. As a matter of fact, the U-Space version described here was finalized in November 1998, and the experimental results obtained by Schneider and Phillips (1999) and Sturt et al. (1999a,b) came to the authors’ attention during the CUNY Sentence Processing Conference, in March 1999.

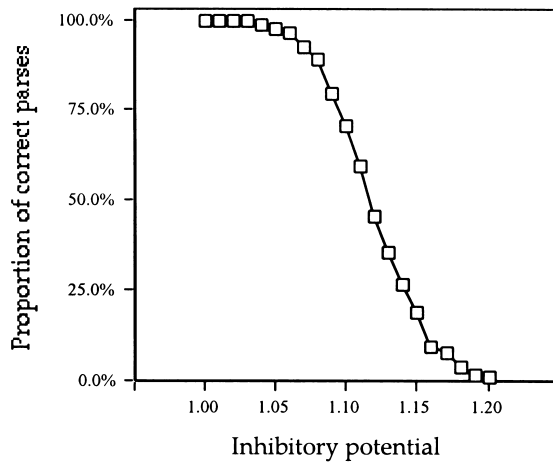


Fig. 7. Influence of a lexical frame preference (see text) on the success rate in parsing sentence (18) in Table 3.

a medium verb with the default bias ($I = 1.0$) towards attracting a Direct Object NP.

In recent experimental work, Sturt et al. (1999a,b) have compared the difficulty of reanalysis in structures like (16) and (20). Sentence (16) (*The man knew the woman slept*) requires reattachment of NP *the woman* from the Direct OBJECT of the Main clause to the SUBJECT of the Subordinate clause. In sentence (20) (*Since the horse kicked the dog bit her*), NP *the dog* is temporarily attached as the Direct OBJECT to *kicked* and is subsequently reattached as the SUBJECT of the Main clause. Sturt et al. (1999a,b) composed sentences of both these types such that the preference of the first verbs (*knew/kicked*) to take a Direct OBJECT NP was identical. Also, the plausibility of the two analyses was controlled for. They found that Main-to-Subordinate reanalysis (as in (16)) is easier than Subordinate-to-Main reanalysis (as in (20)). Fig. 8 shows that this also applies to the U-Space when the first verb is mildly biased towards taking an NP complement.

Examples (21) and (22) stem from a reading study by Warner and Glass (1987). The first member of this pair is grammatically similar to (20) and is easily parsed correctly in the U-Space. However, Warner and Glass observed that (22), the longer of the two, produced a garden-path effect: the readers tended to interpret *the man* as the Direct Object of *kills* rather than as the Subject of *strikes*. This tendency toward ‘late closure’ of the subordinate clause was absent in (21). In the U-Space model, the difficulty of (22) is due to the late arrival of the verb of the main clause (*strikes*). It is separated from the Subject noun (*man*) by three intervening words. Due to this time lag, the activation of the lexical frame for *man* has decayed substantially. Therefore, the U-link connecting this NP frame to the Subject NP of *strikes* can mount relatively little inhibitory force against its competitors, one of them being the older U-link between NP *man* and the Direct Object




Fig. 8. Influence of Verb-DOBJ 'adhesion' on probability of reanalysis (examples (16) and (20)). The horizontal axis represents the adhesion of the first verb (*knew/kicked*) to its Direct OBJECT: high inhibitory potential implies strong adhesion. The curves show the proportion of successful parses in 1000 trials for varying adhesion levels.

NP of *kills*. The inhibitory force exerted by the latter U-link is crucial in deciding the outcome of the competition. With the default value of I ($I = 1.0$), *kills* turns out to win the competition, so that *strike* cannot secure a Subject NP. Additional simulation runs show that *kills* starts losing the competition when the I values of this verb's Direct Object U-link drop below the level of about 0.85. In case of the shorter and easier sentence (21), on the other hand, the corresponding I value can rise until 1.10 without any garden-path effect manifesting itself.

4.4. Lexical ambiguity

Examples (23–30) illustrate the model's behavior in response to word-class ambiguities. *Can* is ambiguous between auxiliary and main verb, *fish* between noun and verb, and *her* between determiner and personal pronoun. No preference for either of these alternatives has been built into the lexical frames (equal activation values). Without such biases, the two legal analyses of (23) are expected to emerge about equally often. However, word-class ambiguity sometimes slows down the parsing process. The model tends to wait until one of the alternative lexical frames has taken a firm lead before digesting the next word. Removing one of *can*'s lexical frames therefore reduces the number of processing cycles consumed by this word. (After that, however, removing either of the lexical frames for *fish* does not affect processing time any further.)

This aspect of the model's behavior entails the prediction that lexical (and, for that matter, syntactic) ambiguity will slow down the parsing process in human comprehenders, especially when the alternative analyses are well-balanced in terms of frequency or off-line preference. This prediction has received support from eye fixation studies as far as the resolution of *lexical* ambiguities is concerned (Rayner & Duffy, 1986; see also Duffy, Morris & Rayner, 1988). The empirical picture is less

clear, however, in case of *syntactic* ambiguity, which sometimes delays (e.g. Spivey-Knowlton & Sedivy, 1995) and sometimes speeds up comprehension (e.g. Traxler, Pickering & Clifton, 1998; Van Gompel, Pickering, Liversedge & Traxler, 1999; Van Gompel, Scheepers & Pickering, 1999). Powerful conceptual or referential factors may be at work here, capable of overriding relatively small syntactic effects.¹¹

Examples (24) and (25) (from Pritchett, 1992) demonstrate that the U-Space strongly prefers to analyze *her contributions* as a single NP although, as a matter of fact, the entry activations of *her* as determiner and as personal pronoun are identical. The reason is that, initially, only the pronoun reading of *her* can attach as the prepositional object of *without*. Upon entry of *contributions*, this NP accommodates the determiner *her*, and together they conspire to gain the prepositional object role. Their combined inhibitory forces easily succeed in pushing the other *her* aside. The alternative analysis, with *her* and *contributions* as separate NPs, requires substantially different entry activation levels of the lexical frames corresponding to the two senses of *her*, in favor of the pronoun reading.

Examples (26–30) serve to explore treatment of the four-fold ambiguity of *that*. The model's lexicon contains four entries for this word: *that* as (1) determiner, (2) demonstrative pronoun, (3) relative pronoun, and (4) complementizer (subordinating conjunction). The lexical frames defined in these entries are all entered into the U-Space whenever *that* occurs in the input string. In sentences (2/3) and (9/10) above we already met *that* in its role as relative pronoun and complementizer, respectively. Sentences (26) and (27) have been adapted from Lewis (1993), who classifies them as unproblematic ambiguities. In (26), *that* is immediately attached as demonstrative pronoun because *hate* requires an NP as Direct OBJECT. The relative pronoun option, although also encoded as an NP frame, is eliminated by a word order rule stating that relative pronouns have to precede the verb in the clause they belong to. The same attachment briefly figures in sentence (27) but there it is soon contested by NP *boy*, which can accommodate the determiner *that* and easily win the competition (cf. the course of events elicited by *her* and *contributions* in (24/25)).

Sentence (28), classified as a garden-path by Gibson (1991), is nevertheless always analyzed correctly. The substring *that John smokes* is originally attached as *believe*'s complement and *annoys* easily succeeds in detaching it from there. Interestingly, the score drops dramatically when *annoys* is redefined as a verb that disprefers verb frames in the role SUBJECT. For (29), the U-Space delivers two alternative analyses in roughly equal proportions – one with *that* as demonstrative pronoun, the other one with *that* as complementizer. On the latter reading, the string

¹¹ On a speculative note, disambiguation by conceptual or referential mechanisms is perhaps better construed as a race-based than as a competition-based process. In mechanisms of this sort, alternative interpretations are computed in parallel without inhibiting each another, and the first one to be completed is chosen by the system. A race-based conceptual/referential disambiguator could operate in conjunction with a competition-based syntactic parser like the U-Space and modulate the strength development of U-links.

as a whole is a subordinate clause rather than a complete sentence. Finally, sentence (30) (from Gibson, 1991) receives a 100% correct score. This score drops when *surprised* is redefined as dispreferring a clause as its SUBJECT. Alternatively, we can assign a low activation value to *that* as a complementizer if it occurs in sentence-initial position. In Section 7 we will argue that the latter measure is less *ad hoc* than it may seem at first sight.

The explorations of the Unification Space described in this section attest to the model's capability to account for major effects of syntactic complexity and of word-class and syntactic ambiguity on human parsing performance. In the few cases where serious discrepancies arose between human and model performance, we have identified plausible external factors that may have caused the model to fail. These factors typically center around properties of individual lexical items – their internal structure (e.g. their preference profile for various types of complementation), their relationship with conceptual structures (cf. the manner-of-motion verbs), or their initial activation levels when entering the Unification arena.

In Section 7 we will outline briefly how the U-Space model could be interfaced with non-syntactic modules to accommodate lexical, semantic and pragmatic (discourse, referential) factors. But we first turn to syntactic parsing disorders in certain aphasic syndromes.

5. Model simulations: agrammatic sentence analysis in aphasia

In this section we report the results of our attempt to fit the U-Space model to the data by Caplan et al. (1985; in Section 1.1) (see Table 1). Remember that this study yielded comprehension scores for nine different sentence types. A total of 144 patients had participated in three different experiments. The authors divided this group into eight, five and seven subgroups, respectively, based on their performance on the various sentence types. The authors report comprehension scores for these 20 subgroups for each of the nine sentence types. The simulation of the data for *eight* of these sentence types – the U-Space cannot handle conjoined structures (coordination) – proceeded as follows.

First, we reasoned that patients who understand the words of the sentence but fail to take the grammatical structure into account have three options: to do nothing, to guess a response, or to apply some systematic response strategy that sometimes yields a correct response by chance (e.g. 'first NP is actor of first action'). If the model produces a certain percentage of correct parses for a sentence type, we can take that proportion as the model's 'true' score. This is equivalent to the 'do nothing' alternative in the patients. We can mimic a guessing strategy by adding to this proportion the percentage of correct responses that results from assigning the actors referred to in the sentences (toy animals) at random to the thematic roles offered by the verb(s). We assumed that 'doing nothing' in case of parsing failure is an unlikely patient reaction and that guessing is the better way of making patient and model data comparable. Correction for chance is also commendable on the assumption that the patients, rather than guessing, apply all sorts of *different* strategies. Therefore, we




Fig. 9. Simulation of agrammatic sentence analysis. Proportions of successful parses (corrected for chance) for eight target sentence types. (Patient data from the study by Caplan et al., 1985.)

applied chance correction to the patient scores as well as to the model scores.¹² The corrected average scores (proportions correct) for the eight sentence types are displayed in Fig. 9.

Second, because the 20 patient subgroups exhibited very different levels and patterns of sentence comprehension performance, we decided not to lump their scores together but to devise a model fitting procedure that takes inter-subgroup variation into account. The goal was to find a set of parameter settings yielding, for the eight target sentence types, a pattern of parsing successes and failures approximating the various patterns of the patient subgroups as closely as possible. The parameter settings representing sentence analysis in normal speakers (Table 2) can be viewed as a point in an n -dimensional space, where n equals the number of parameters. Let us call this point P . Impaired parsing performance in severely agrammatic patients can be thought of as another point, A , somewhere in the

¹² In order to make the model scores comparable to patient scores, we added the expected proportion of correct guesses, on the assumption that every NP is randomly assigned a thematic role (we did not assume a systematic response strategy). We used the formula $\text{obs} = \text{true} + (1 - \text{true})/n!$ to estimate the 'observed' proportion of correct responses for the model, given the 'true' proportion of correct parses obtained for a sentence. Here, n equals the number of actors (toy animals) mentioned in the input sentence. For example, if the model obtains a 'true' score of 60% correct trials on *The elephant hit the monkey*, this yields an 'observed' score of 80%. The true score of a patient, given his/her observed score, is approximated by $\text{true} = (n! \times \text{obs} - 1)/(n! - 1)$.

space surrounding *P*. The performance of less seriously handicapped (subgroups of) patients can then be represented as a point on the line between *A* and *P*. What we are looking for is the coordinates of *A*. In order to keep the search space within manageable proportions, we have made the following simplifying assumptions.

- Based on informal pilot studies of the model's behavior, we decided to restrict the search space to five dimensions. The parameters that were allowed to vary were the 'dynamic' ones that directly affect the strength of a U-link: activation decay rate, sensitivity to inhibition, and strength increase. They are listed at the right-hand side of Table 2.¹³ Additional simulation runs with a six-dimensional search space, where the noise rate was allowed to vary together with the five dynamic parameters, did not give a better fit with the empirical data than the runs reported below.
- We assumed that, in the five-dimensional subspace, the performance of the 20 patient subgroups would be located somewhere on the straight line connecting *P* and *A*.¹⁴ By moving step-by-step from *P* to *A*, at each step running the model with the eight target sentences as input, we get a specimen of parsing performance whose similarity to the performance of each of the subgroups can be determined by a least squares method. Each subgroup is then located on the line at the position where it produces the smallest sum of squares. The looked-for position of *A* is the one where the latter sums of squares are smaller on average than for any other position of *A*. The parameter values serving as the coordinates of *A*'s optimal position are returned as the final result of the search procedure.
- The resulting parameter settings imply a pattern of corrected parsing scores for each subgroup for the target sentences. Computing the average of these scores (weighted for the number of patients in the subgroup) over the 20 subgroups yields a set of simulation data directly comparable to the patient data.

Notice that throughout the described procedure the performance profiles of the 20 subgroups served as simulation targets: the performance profile of the group of patients as a whole (depicted in Fig. 9) was fitted only indirectly, namely, by adding subgroup error scores.

We developed a computer program embodying the search procedure and ran it for a large number of *A* positions in a subspace around *P*. Fig. 9 shows that the search indeed delivered a set of five parameter values generating a pattern of parsing performance that resembles the patient data quite well. Although the U-Space scores for three sentence types in the middle range are a little higher than the patients', the

¹³ The noise parameter listed in Table 2 was not allowed to vary. Although this parameter does affect parsing performance, in preliminary test runs it had failed to improve the model's performance pattern. Interestingly, Dell, Burger and Svec (1997) reach the same conclusion in exploring the behavior of their computational model of serial order in language production.

¹⁴ We realize that this assumption presupposes a one-dimensional view of the severity of agrammatic disorders. This methodological assumption may not be correct.

Table 4
Simulation of agrammatic sentence analysis^a

Parameter	Normal value	Agrammatic value
d	0.98	0.837
$p_{\text{incr}}^{\text{root}}$	0.84	0.50
$p_{\text{incr}}^{\text{foot}}$	0.97	0.078
$p_{\text{sens}}^{\text{const}}$	0.65	0.25
$p_{\text{sens}}^{\text{foot}}$	0.46	0.46

^a Comparison of normal and agrammatic values of dynamic model parameters; d , decay; incr(ement), spontaneous strength build-up as a function of time; sens, sensitivity to inhibition.

Spearman rank correlation between model and (corrected) patient data is very satisfactory ($\rho = 0.976$).

6. Competitive power as a computational resource

Table 4 shows that the agrammatic values of all free parameters are lower than normal, with the exception of the last one. Is it possible to lend a meaningful interpretation to these differences? Decay rate, speed of spontaneous strength build-up, and sensitivity to inhibition are the three¹⁵ critical variables that determine the inhibitory force mounted by a Unification link against its competitors (see the unification strength formula at the end of Section 3.5). Lowering these values reduces the model's *competitive power* (or 'competitiveness', 'competitive drive'). This reduction manifests itself in at least two ways. First, a newly entered lexical frame is harder to integrate into the existing structure, in particular when this integration requires the suppression of already strong U-links created earlier. To see this, remember that the activation levels of the nodes at the extremes of a U-link are important determinants of the inhibitory potential of that U-link: they 'fuel' the inhibition process. Second, if a new lexical frame cannot be integrated into the already existing structure and has to await the attachment opportunities offered by lexical frames further downstream, its own activation level may, in the meantime, have dropped to a level from where it is hard to pump sufficient 'energy' into the U-link(s) connecting it to the overall structure. These effects are amplified if U-links are relatively immune (insensitive) to inhibitory force, so that their strength values are harder to change – in other words, if assembled structures lose their capacity to respond to syntactic information coming in later. Finally, when unification strength builds up at too slow a rate, estab-

¹⁵ If the above simulation results could have been obtained with a solution that involves just one or two critical variables (e.g. only decay rate), the search procedure would definitely have found one. A simpler explanation of the agrammatism effects than one involving three parameters does not seem possible within the current framework.

lishing firm bonds between lexical frames will be difficult so that, when the competitions subside, the Unification Space is left with a network that includes many weak links and represents, not a single tree, but a set of disconnected partial trees.

How does ‘competitive power’ of the structure assembly mechanism compare to other explanations of the deteriorated parsing performance in agrammatic patients? Several aphasiologists consider problems with the activation of lexical and/or syntactic information as the prime source of agrammatism. At least three versions of this basic idea have been proposed (cf. Haarmann & Kolk, 1991; Haarmann et al., 1997; Zurif, 1998):

- accelerated decay of activation, i.e. the information is usable for too short a period of time
- delayed activation, i.e. the information becomes available too late, or
- insufficient activation, i.e. the incoming information cannot be supplied with enough activation to be used effectively.

The first version corresponds to one of the three U-Space parameters determining competitiveness. The second one – retarded activation – resembles our assumption of a slower build-up of unification strength. The third version has no parallel in a U-Space parameter. However, in a theoretical approach that views activation as the parser’s main source of energy, ‘insufficient activation’ fulfils the same role as ‘lack of processing capacity’ or, in our terms, ‘lack of competitive power’ or ‘poor competitiveness’.

The idea of competitive power as a computational resource provides us with a tool to obtain fine-grained estimates of the ‘processing load’ imposed by various types of syntactic constructions. Our simulations of parsing performance in normals tend to yield all-or-none scores, as shown by the many 100% successes or failures in Table 3. These scores, which are due to ceiling or floor effects, hide subtle differences in the amount of processing capacity they consume. In order to zoom in on them, we ran computer simulations with intermediate values of the four competitiveness parameters, that is, with a parser whose competitive power was gradually degraded from normal to agrammatic levels. We reasoned that parsing performance on a syntactic structure that consumes a larger amount of processing capacity (i.e. requires a highly competitive parser) will begin to deteriorate when degraded to a lesser extent than the performance on an easier structure.

We explored this issue by running the U-Space model with 20 different levels of degradation, that is, with the normal and the agrammatic parameter settings shown in Table 4 and 18 intermediate settings distributed evenly over the normal-to-agrammatic intervals at distances of 5 percentage points. The test case was the well-known contrast between relative clause with subject and object extraction, with the latter being harder to comprehend than the former (as confirmed, for example, by the patient data in Fig. 9 and Table 1: compare the Cleft-Subject and the Cleft-Object

sentences). Our sample material consisted of sentences (2) and (3) in Table 3 and two versions with a PP added.

Subject Relative: *The rat that chased the cat escaped*

Subject Relative with PP: *The rat that chased the cat with the binoculars escaped*

Object Relative: *The rat that the cat chased escaped*

Object Relative with PP: *The rat that the cat with the binoculars chased escaped*

Object Relatives stage a somewhat fiercer competition in U-Space than Subject Relatives. In the former type of constructions, when the subordinate verb arrives one more NP competes for the role of SUBJECT than in the latter. Word order rules only require that the SUBJECT precedes the verb. Moreover, Subject Relatives have the additional advantage that the NP following the subordinate verb experiences little competition in acquiring the role of Direct OBJECT.

The simulation runs yielded the expected result, shown in Fig. 10. The parses of the object relative sentences began to deteriorate at lower levels of degradation than

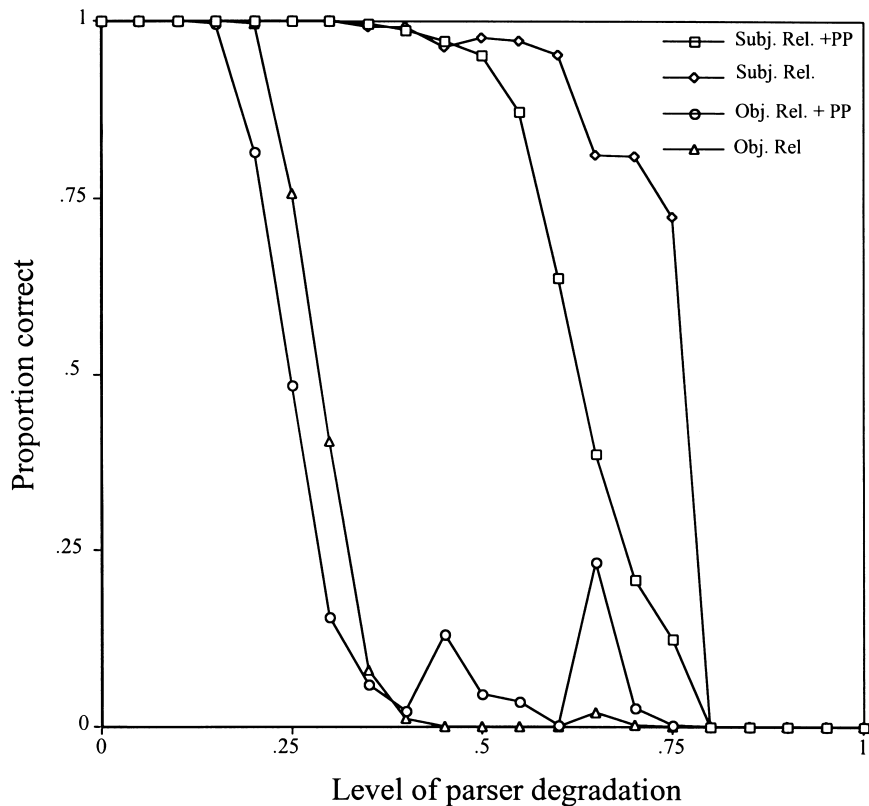


Fig. 10. Parsing success and failure for sentences containing a subject or an object relative clause, with or without an additional PP, at different levels of parser degradation.

those of subject relatives. Moreover, the insertion of a PP appears to consume some additional processing capacity, and the curves suggest an interaction between clause type and clause length.

7. Discussion: evaluating the Unification Space

The U-Space model of human syntactic parsing is based on the following theoretical assumptions.

- The underlying grammar is fully lexicalized: lexical frames retrievable from the Mental Lexicon are the elementary building blocks for assembling syntactic trees, and the only source of syntactic nodes.
- Nodes of lexical frames carry a gradually decaying activation value.
- Feature unification is the composition operation that combines lexical frames into larger syntactic trees.
- This operation is local, with the lexical frame defining the domain of locality.
- The resulting attachments are not all-or-none but have graded, dynamically varying strength.
- Mutually exclusive attachments compete with one another continually, and lateral inhibition is the dynamic mechanism controlling attachment strengths and thereby the current parse tree configuration.

The simulation results reported in Sections 4–6 show that the parsing behavior of the U-Space model meets the basic criteria put forward in Section 2. Parsing proceeds incrementally from left to right, delivers a single grammatically correct parse tree in case of success, and responds to syntactic complexity and to local and global syntactic ambiguities in ways similar to human understanders.

Although the present version of the model lacks conceptual and referential components, it is interactive in principle because semantics and pragmatics can influence parsing behavior. This may happen via two channels, one affecting the strength of U-links, the other the activation values of lexical frames. First, semantic factors are allowed to bias the inhibitory potential of U-links in a way similar to the treatment of ‘adhesion’ between a head and its dependents (cf. the *jogging a mile* example, Fig. 7). Syntactically proposed unifications that are supported by conceptual or referential information receive a higher inhibitory potential than unifications corresponding to implausible events or states-of-affairs. The second channel seizes upon the activation values of lexical frames and the syntactic nodes in them. If conceptual association and inference processes involve spreading of activation through a network of concepts, then it seems reasonable to hypothesize that this activation also affects the lexical frames corresponding to these concepts. So far we have assumed that the activation value of lexical frames is initialized at a level correlating with its frequency in language use (cf. the treatment of the *horse raced* example). This does not preclude other factors, semantic ones in particular, from modulating frame activation levels.

The activation level of a lexical frame upon entry in the U-Space may be

modulated, furthermore, by left–right contingencies between input items. These items may be concrete *words* (e.g. idioms), word *categories* (parts of speech), or words carrying certain grammatical *features* (e.g. case). A host of recent corpus-based studies in computational linguistics have demonstrated that bigram, trigram, etc., frequencies contain useful information for ‘tagging’ (determining actual word category membership of categorially ambiguous words) and even for predicting constituent boundaries. This information can guide the syntactic parsing process (cf. Charniak, 1993, 1997). There is no reason to assume that general-purpose pattern recognizers in our cognitive system could not pick up left–right contingencies. This opens up, at least in principle, the possibility for such statistical patterns to influence syntactic parsing, for example, by modulating the entry activation of lexical frames.

Experimental evidence for a mechanism of this sort is provided by Corley (1998). He conducted a self-paced reading study with materials such as the following:

- (NN) *The foreman knows that the warehouse prices are better value by far*
 (NV) *The foreman knows that the warehouse prices the beer very modestly*
 (VN) *The attendant discovered that the museum finds were destroyed last night*
 (VV) *The attendant discovered that the museum finds the public a nuisance*

The critical items in these sentences are ambiguous between Noun and Verb readings. However, *prices* occurs more frequently as a noun, *finds* as a verb. Corley (1998) predicted and observed a reading time increase on the words following *prices/finds* when the sentence context was incompatible with the most frequent reading (i.e. slower reading in conditions NV and VN than in NN and VV). Tabor and Richardson (1999) report a similar result with sentences containing a Verb–Noun pair that happens to constitute a frequent two-word noun compound.

- Its not wise to waste baskets needed by many people*
Its not wise to steal baskets needed by many people

The string following *waste basket* was read more slowly than the one after *steal basket*.

Tabor, Juliano and Tanenhaus (1997) collected data on the initially preferred interpretation of the function word *that* in different sentential environments. In a sentence like *The lawyer insisted that cheap hotels were clean*, the word *that* is taken to be a subordinating conjunction. However, in *That cheap hotels were clean and comfortable surprised us*, the word *that* is initially interpreted as a determiner. This conclusion was based on a comparison of self-paced reading times for these sentences and their counterparts with *hotel was* instead of *hotels were*. When the subordinate clause followed the main verb *insisted*, the passages *hotels were* and *hotel was*, were read at about the same speed. However, when the subordinate clause opened the sentence, *hotels were* turned out to take much more time than *hotel was*. Tabor et al. (1997) furthermore observed that these biases agree with the ‘contingent word frequencies’ of *that* in a corpus of written texts, i.e. with the frequencies of *that*

as a demonstrative pronoun or a subordinating conjunction in different grammatical positions.¹⁶

In separate simulation runs with these four sentences we have verified that the U-Space model yields a difficulty profile similar to that of the participants in the experiment by Tabor et al. (1997), if we vary the entry activation levels of *that* in accordance with contingent word frequencies in the corpus. (We already hinted at this possibility in the context of sentence (30) of Table 3; see the end of Section 4.4.)

If the effect of left–right contingencies can be substantiated in further research, this may have important bearings on the hotly debated issue of *head-driven parsing* (for a recent overview of the literature, with novel experimental data, see Kamide & Mitchell, 1999; Konieczny, Hemforth, Scheepers & Strube, 1997). The U-Space model exemplifies strictly head-driven parsing: phrasal nodes (S, NP, PP, etc.) are introduced by their heads (verb, noun, preposition), so that attachment of pre-head constituents is postponed until the head word has arrived. For instance, the initial NP of a declarative main clause is left unattached until it is offered a grammatical function (e.g. Subject) by the finite verb. Konieczny et al. (1997), however, argue that this view of parsing is incorrect. German-language data obtained in their laboratory (see Hemforth, Konieczny & Strube, 1993) indicate that the oft-reported preference to interpret clause-initial NPs as the subject is measurable before the verb has been encountered: if clause-initial NPs are explicitly marked as non-nominative they consume slightly more processing time than NPs that are unambiguously nominative or have ambiguous case marking. (Similar effects have recently been observed for Finnish: see Hyönä & Hujanen, 1997.) Strictly head-driven parsing indeed does not predict such results. However, the effect could also reflect the relatively low incidence of non-nominative case in (pro)nouns, determiners, and adjectives occupying the first few positions after a subordinating conjunction or after a sentence boundary. Rather than resorting to a syntactic processor that is not strictly head-driven, we

¹⁶ In a pilot study with native speakers of Dutch performing a self-paced reading task, we observed strong biasing effects of categorial trigram frequency on the syntactic analysis of the trigrams. The participants read Dutch sentences containing a categorially ambiguous word pair such as *stenen vormen*, which is ambiguous between Adjective + Noun ('stone shapes', i.e. shapes/forms made out of stone) and Noun + Verb ('stones form/constitute'). The definite article *de* (the) followed by this pair constitutes the string *de stenen vormen*, interpretable as an Article-Adjective-Noun trigram ('the stone shapes') or as an Article-Noun-Verb trigram ('the stones constitute...'). In Dutch, the frequency of the former categorial trigram is more than five times higher than that of the latter, presumably causing a bias in favor of the Adjective interpretation of *stenen* (frequency data based on a tagged corpus of one million Dutch words; Instituut voor Nederlandse Lexicologie, Leiden). Now consider the string *de blauwe stenen*, interpretable either as an Article-Adjective-Adjective trigram or as an Article-Adjective-Noun trigram (*blauwe* 'blue' is a categorially unambiguous adjective). Here, language frequencies favor the Noun interpretation of *stenen* by a 9:1 ratio. We observed that, as predicted on the basis of trigram frequency ratios, the Article-Adjective-Noun reading of *de stenen vormen* is strongly preferred (with *stenen* as an adjective and *vormen* as a noun), whereas in *de blauwe stenen vormen* there is a strong bias in favor of *stenen* interpreted as a noun and *vormen* as a verb. (If the mentioned trigram frequency ratios hold for English as well, the preferred reading of *iron forms* will switch from Adjective-Noun in *the iron forms* to Noun-Verb in *the blue iron forms*.)

prefer to explore the hypothesis that the entry activation levels of lexical frames can be fine-tuned in accordance with left–right contingencies.

Recently, Kamide and Mitchell (1999) leveled a similar criticism against data such as those cited above for German and Finnish, which many researchers interpret as evidence for the occurrence of pre-head attachments, i.e. against strictly head-driven parsing. They argue that these data may be attributed to the operation of a case-assignment mechanism outside of the parser proper. Such a pre-syntactic mechanism could assign case partly on a left-to-right basis, e.g. assign nominative case to the NP that opens the sentence. We agree with this argument, at least insofar as the proposed case-assignment mechanism can be viewed as a preprocessor that is sensitive to the frequencies of morphological cases in varying sentential environments (their contingent frequencies) and modulates the entry activation of lexical frames accordingly. Whether the remaining evidence in support of pre-head attachment, including the Japanese data reported by Kamide and Mitchell (1999), is also attributable to contingent frequency effects is open to debate.

In future projects we hope to work on extensions of the model in several directions:

- a more sophisticated and principled way of controlling U-link creation during the processing of sentences whose structure involves grammatical movement or discontinuous constituency;
- a conceptual component enabling the parser to compute conceptual and syntactic structures in synchrony and to adapt the strength value of Unification links to semantic plausibility;
- a *non-ad hoc* way of predicting on-line processing load, e.g. reading times for individual words or constituents;
- a computational architecture along the lines of the Unification Space that is capable of syntactic structure assembly not only in sentence comprehension but also in sentence production (cf. Kempen, 2000).

Acknowledgements


The work reported here has been partly supported by a grant from the Royal Academy of Sciences of The Netherlands. The authors are indebted to three anonymous reviewers for their very helpful comments.

Appendix A. Incompatibility conditions on Unification links

Two proposed Unification links are incompatible, and exert inhibitory force on each other, if merging the nodes they connect would lead to violation of general treehood constraints in the derived tree. This is the case if one or more of the following conditions are fulfilled:

- The U-links both try to attach to either the same root or the same foot node.

- The U-links unify two different lexical entries associated with the same (ambiguous) input word.
- Two root nodes try to get U-linked to two *different* foot nodes of the *same* lexical frame and violate each other's word order rules.
- The foot and root nodes of one U-link unify with, respectively, the *root* and the *foot* nodes of the other U-link (thus creating a loop; however, loops involving more than two lexical frames are not checked).
- The U-links lead to crossing branches in the derived tree. Checking this constraint is rather complicated: in the worst case, it requires complete inspection of all Unification links. Therefore, we adopt a simpler procedure at the risk of occasionally failing to eradicate a crossing branch and causing parsing failure. The scheme is easily explained in visual terms. First, type the input sentence twice, on two consecutive lines. This example corresponds to Fig. 2.



Assume that the words in the bottom line represent the root nodes of their lexical frames, while every word in the top line comprises all the foot nodes of their lexical frame. Now suppose we want to know which U-links cross with U-link 6, which in Fig. 2 corresponds to the connection between *woman* in the upper, and *with* in the lower line. Draw all U-links except for (a) those which start and end *outside* the range from *woman* up to and including *with*, and (b) those which are contained *within* this range. The drawing shows that U-link 6 crosses the U-links from *man* to the first occurrence of *the* (U-link 4) and from *sees* to *woman* (U-link 2). These links therefore inhibit U-link 6 (and vice-versa). In order to keep the computational complexity to $O(n^2)$ per step, only the branches originating from the immediate left-hand and right-hand context are checked, i.e. two words at either side of the current word.

References

- Adams, B. C., Clifton, C., & Mitchell, D. C. (1998). Lexical guidance in sentence processing? *Psychonomic Bulletin and Review: a Journal of the Psychonomic Society*, 5, 265–270.
- Berwick, R. C., Abney, S. P. & Tenny, C. (1991). *Principle-based parsing: computation and psycholinguistics*. Dordrecht: Kluwer.
- Caplan, D., & Hildebrandt, N. (1988). *Disorders of syntactic comprehension*. Cambridge, MA: MIT Press.
- Caplan, D., Baker, C., & Dehaut, F. (1985). Syntactic determinants of sentence comprehension in aphasia. *Cognition*, 21, 117–175.
- Charniak, E. (1993). *Statistical language learning*. Cambridge, MA: MIT Press.

- Charniak, E. (1997). Statistical techniques for natural language parsing. *AI Magazine*, 18, 33–44.
- Chater, N., & Christiansen, M. H. (1999). Connectionism and natural language processing. In S. Garrod, & M. J. Pickering. *Language processing*. Hove: Psychology Press.
- Christiansen, M. H., & Chater, N. (1999). Toward a connectionist model of recursion in human linguistic performance. *Cognitive Science*, 23, 157–205.
- Corley, S. (1998). *A statistical model of human lexical category disambiguation*. Ph.D. Dissertation. Edinburgh: University of Edinburgh.
- Crocker, M. W. (1992). *A logical model of competence and performance in the human sentence processor*. Ph.D. Thesis. Edinburgh: University of Edinburgh.
- Dell, G., Burger, L. K., & Svec, W. R. (1997). Language production and serial order: a functional analysis and a model. *Psychological Review*, 104, 123–147.
- Duffy, S. A., Morris, R. K., & Rayner, K. (1988). Lexical ambiguity and fixation times in reading. *Journal of Memory and Language*, 27, 429–446.
- Elman, J. L. (1991). Distributed representation, simple recurrent networks, and grammatical structure. *Machine Learning*, 7, 195–225.
- Frazier, L., & Clifton, C. (1996). *Construal*. Cambridge, MA: MIT Press.
- Frazier, L., & Fodor, J. D. (1978). The sausage machine: a new two-stage parsing model. *Cognition*, 6, 291–325.
- Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, 14, 178–221.
- Friederici, A. D. (1999). The neurobiology of language comprehension. In A. D. Friederici. *Language comprehension: a biological perspective*. Berlin: Springer.
- Garrod, S., & Pickering, M. (1999). *Language processing*. Hove: Psychology Press.
- Gibson, E. (1991). *A computational theory of human linguistic processing: memory limitations and processing breakdown*. Ph.D. Thesis. Pittsburgh, PA: Carnegie Mellon University.
- Gibson, E. (1998). Linguistic complexity: locality of syntactic dependencies. *Cognition*, 68, 1–76.
- Gorrell, P. (1995). *Syntax and parsing*. Cambridge: Cambridge University Press.
- Haarmann, H. J., & Kolk, H. H. J. (1991). A computer model of the temporal course of agrammatic sentence understanding: the effects of variation in severity and sentence complexity. *Cognitive Science*, 15, 49–87.
- Haarmann, H. J., Just, M. A., & Carpenter, P. A. (1997). Aphasic sentence comprehension as a resource deficit: a computational approach. *Brain and Language*, 59, 76–120.
- Hagoort, P., Brown, C., & Osterhout, L. (1999). The neurocognition of syntactic processing. In C. Brown, & P. Hagoort, *Neurocognition of language*. Oxford: Oxford University Press.
- Hemforth, B., Konieczny, L., & Strube, G. (1993). Incremental syntax processing and parsing strategies. *Proceedings of the 15th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Henderson, J. (1994). *Description based parsing in a connectionist network*. Ph.D. Thesis. Philadelphia, PA: University of Pennsylvania.
- Hyönä, J., & Hujanen, H. (1997). Effects of case marking and word order on sentence parsing in Finnish: an eye fixation analysis. *Quarterly Journal of Experimental Psychology*, 50A, 841–858.
- Jurafsky, D. (1996). A probabilistic model of lexical and syntactic access and disambiguation. *Cognitive Science*, 20, 137–194.
- Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: individual differences in working memory. *Psychological Review*, 99, 122–149.
- Kamide, Y., & Mitchell, D. C. (1999). Incremental pre-head attachment in Japanese parsing. *Language and Cognitive Processes*, 14, 631–662.
- Kaplan, R. M. (1972). Augmented transition networks as psychological models of sentence comprehension. *Artificial Intelligence*, 3, 77–100.
- Kempen, G. (1996). Computational models of syntactic processing in human language comprehension. In A. Dijkstra, & K. De Smedt, *Computational psycholinguistics: symbolic and subsymbolic models of language processing*. London: Taylor & Francis.
- Kempen, G. (2000). Human grammatical coding. Book manuscript, Leiden University, submitted.
- Kempen, G., & Vosse, T. (1989). Incremental syntactic tree formation in human sentence processing: a

- cognitive architecture based on activation decay and simulated annealing. *Connection Science*, 1, 273–290.
- Kimball, J. (1973). Seven principles of surface structure parsing in natural language. *Cognition*, 2, 15–47.
- Kirkpatrick, S., Gelatt Jr., C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220 (4598), 671–680.
- Konieczny, L., Hemforth, B., Scheepers, C., & Strube, G. (1997). The role of lexical heads in parsing: evidence from German. *Language and Cognitive Processes*, 12, 307–348.
- Lewis, R. (1993). *An architecturally-based theory of human sentence comprehension*. Ph.D. Thesis. Pittsburgh, PA: Carnegie Mellon University.
- Lewis, R. (1996). Interference in short-term memory: the magical number two (or three) in sentence processing. *Journal of Psycholinguistic Research*, 25, 93–115.
- MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). Lexical nature of syntactic ambiguity resolution. *Psychological Review*, 101, 676–703.
- Marcus, M. P. (1980). *A theory of syntactic recognition for natural language*. Cambridge, MA: MIT Press.
- McRoy, S. W., & Hirst, G. (1990). Race-based parsing and syntactic disambiguation. *Cognitive Science*, 14, 313–353.
- Mitchell, D. C. (1987). Lexical guidance in human parsing: locus an processing characteristics. In M. Coltheart, *Attention and performance XII: the psychology of reading*. Hillsdale, NJ: Erlbaum.
- Mitchell, D. C. (1994). Sentence parsing. In M. A. Gernsbacher, *Handbook of psycholinguistics*. San Diego, CA: Academic Press.
- Press, W. H., et al. (1992). *Numerical recipes in C*, (2nd ed.). Cambridge: Cambridge University Press.
- Pritchett, B. L. (1992). *Grammatical competence and parsing performance*. Chicago, IL: University of Chicago Press.
- Rayner, K., & Duffy, S. A. (1986). Lexical complexity and fixation times in reading: effects of word frequency, verb complexity, and lexical ambiguity. *Memory and Cognition*, 14, 191–201.
- Reilly, R. G., & Sharkey, N. E. (1992). *Connectionist approaches to natural language processing*. Hillsdale, NJ: Erlbaum.
- Schneider, D. A., & Phillips, C. (1999). Reanalysis as a last resort? In *Proceedings of the Twelfth Annual CUNY Conference on Human Sentence Processing*. New York.
- Sharkey, N. (1992). *Connectionist natural language processing. Readings from Connection*. Science Oxford: Intellect.
- Shieber, S. M. (1986). *An introduction to unification-based approaches to grammar*. Stanford, CA: Center for the Study of Language and Information, Stanford University.
- Spivey-Knowlton, M., & Sedivy, J. C. (1995). Resolving attachment ambiguities with multiple constraints. *Cognition*, 55, 227–267.
- Stevenson, S. (1993). A competition-based explanation of syntactic attachment preferences and garden path phenomena. In *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*. Columbus, OH.
- Stevenson, S. (1994). Competition and recency in a hybrid network model of syntactic disambiguation. *Journal of Psycholinguistic Research*, 23, 295–322.
- Stevenson, S., & Merlo, P. (1997). Lexical structure and parsing complexity. *Language and Cognitive Processes*, 12, 349–399.
- Sturt, P., Pickering, M. J. & Crocker, M. W. (1999). Exploring the ‘Reanalysis a last resort’ strategy. In *Proceedings of the Twelfth Annual CUNY Conference on Human Sentence Processing*. New York.
- Sturt, P., Pickering, M. J., & Crocker, M. W. (1999). Structural change and reanalysis difficulty in language comprehension. *Journal of Memory and Language*, 40, 136–150.
- Tabor, W. & Richardson, D. (1999). Ungrammatical influences in sentence processing. In: *Proceedings of the Twelfth CUNY Conference on Human Sentence processing*. New York.
- Tabor, W., Juliano, C., & Tanenhaus, M. K. (1997). Parsing in a dynamic system: an attractor-based account of the interaction of lexical and structural constraints in sentence processing. *Language and Cognitive Processes*, 12, 211–271.
- Traxler, M. J., Pickering, M. J., & Clifton, C. (1998). Adjunct attachment is not a form of lexical ambiguity resolution. *Journal of Memory and Language*, 39, 558–592.

- Van Gompel, R. P. G., Pickering, M. J., Liversedge, S. P., & Traxler, M. J. (1999). Further explorations of race-based parsers: evidence for a separate resource processor. In *Proceedings of AMLaP-99*. Edinburgh.
- Van Gompel, R. P. G., Scheepers, C. & Pickering, M. J. (1999). Race-based parsing: evidence from subject-object ambiguities in German. In *Proceedings of AMLaP-99*. Edinburgh.
- Warner, J., & Glass, A. L. (1987). Context and distance-to-disambiguation effects in ambiguity resolution: evidence from grammaticality judgments of garden path sentences. *Journal of Memory and Language*, 26, 714–738.
- Zurif, E. B. (1998). The neurological organization of some aspects of sentence comprehension. *Journal of Psycholinguistic Research*, 27, 181–190.