

Parsing Verb-Final Clauses in German: Garden-path and ERP Effects Modeled by a Parallel Dynamic Parser

Theo G. Vosse (theo.vosse@fcdonders.ru.nl)

Cognitive Psychology Unit, University of Leiden
Wassenaarseweg 52, 2333 AK Leiden, The Netherlands

FC Donders Centre for Cognitive Neuroimaging
Kapittelweg 29, 6525 ER Nijmegen, The Netherlands

Gerard A.M. Kempen (gerard.kempen@mpi.nl)

Max Planck Institute for Psycholinguistics
Wundtlaan 1, 6525 XD Nijmegen, The Netherlands

Cognitive Psychology Unit, University of Leiden
Wassenaarseweg 52, 2333 AK Leiden, The Netherlands

Abstract

Experimental sentence comprehension studies have shown that superficially similar German clauses with verb-final word order elicit very different garden-path and ERP effects. We show that a computer implementation of the Unification Space parser (Vosse & Kempen, 2000) in the form of a localist-connectionist network can model the observed differences, at least qualitatively. The model embodies a parallel dynamic parser that, in contrast with existing models, does not distinguish between consecutive first-pass and reanalysis stages, and does not use semantic or thematic roles. It does use structural frequency data and animacy information.

Keywords: Unification-Space; word-order scrambling; German; ERP effects; parsing; connectionist modeling

Introduction: Word Order Variation and Grammatical Function Assignment

German is a language with a relatively free word order. The grammatical function of syntactic phrases is often indicated by morphological markings. This allows language users to produce these phrases in varying orders without confusing the comprehender. However, not all phrases have morphological marks. Certain noun phrases (NPs), in particular proper names and “bare” plural NPs (i.e., without determiners or adjectives), are highly ambiguous with respect to morphological case. Hence, speakers and comprehenders also use heuristics or preferences such as “subject before object” for NPs without case-marking. Many German verbs have their own case pattern. E.g., the verbs *danken* ‘thank’ and *sehen* ‘see’ both take a Subject and an Object—like their English translation equivalents; but while *sehen* governs a Direct Object with accusative case, *danken* takes an Indirect Object with dative case. Consequently, a simple default interpretation of an ambiguous NP will not always work and may cause “garden-path” effects and other comprehension difficulties.

A special word order phenomenon concerns the position of verbs in main and subordinate clauses. In main clauses, the finite verb occupies a relatively early position, often preceding one or more of its arguments (Subject and Objects). In subordinate clauses, however, the finite verb follows all its arguments. This raises two questions: how does the language comprehension system in speakers of German assign grammatical functions to pre-verbal NPs, in particu-

lar to case-ambiguous NPs; and what is the time-course of this process?

There is solid empirical evidence that in verb-final clauses grammatical functions are assigned before the head verb is processed (e.g. Konieczny, 1997; Kamide & Mitchell, 1999). These assignments must be preliminary because the actual head and its argument structure are still unknown. For example, consider sentence fragment (1)—a subordinating conjunction followed by two case-ambiguous NPs.

(1) ... *dass Richard Künstlerinnen* ...
... that Richard artists ...

Richard and *Künstlerinnen* can play any grammatical role: nominative Subject (Subj), accusative Direct Object (DObj), or dative Indirect Object (IObj). Based on the available empirical evidence, we can assume that the system prefers *Richard* to be Subj, and *Künstlerinnen* DObj, or perhaps IObj. This preference is so strong that it even influences processing of unambiguously case-marked sentence fragments, such as

(2a) ... *dass der Bischof dem Kardinal* ...
... that the bishop_{SUBJ} the cardinal_{IOBJ} ...
(2b) ... *dass dem Bischof der Kardinal* ...
... that the bishop_{IOBJ} the cardinal_{SUBJ} ...

Although (2a) and (2b) are both fully grammatical, the non-canonical order in (2b) tends to be read somewhat slower than canonical order (2a) (e.g. Scheepers & Vasishth, 2007).

Numerous empirical studies with a variety of experimental methods, including grammaticality judgments, reaction times, self-paced reading, eye-movement tracking, ERPs, and fMRI have unearthed aspects of how native speakers of German parse and interpret pre-verbal NPs in subordinate clauses (e.g., see Hemforth & Konieczny, 2000; Bader & Bayer, 2006; and Bornkessel & Schlesewsky, 2006). Given the space limitations of this paper, we focus on the empirical phenomena observed in studies on a special group of subordinate clauses: those beginning with the conjunction *dass* ‘that’ followed by two maximally case-ambiguous animate NPs, and ending with a single finite verb, as in (3a-f) below. For these target structures, we will explore the comprehension effects that manifest themselves at or after the onset of the clause-final head verb.

In this article, we take issue with the central assumptions in two theories that have been proposed to account for the key phenomena: The Linking and Checking Model (L&CM) developed by Bader & Bayer (2006), and the Extended Argument Dependency Model (eADM) by Bornkessel &

Schlesewsky (2006). Both models assume a SERIAL parser. We present a new PARALLEL model of parsing, implemented as a localist-connectionist network. The model, nicknamed SINUS, is the successor to the Unification Space parser by Vosse & Kempen (2000). We will show that the dynamic behavior resulting from the interplay of spreading activation and lateral inhibition in the neural network can simulate the essentials of the targeted comprehension phenomena. Hence, the assumption of a serial parser with discrete first-pass and reanalysis stages is not warranted by the data.

The Empirical Domain

The comprehension phenomena to be simulated by SINUS have been obtained with sentences of the following types. For overviews of the experimental data, see Bader & Bayer (2006) and Bornkessel & Schlesewsky (2006). The NPs in the target clauses are a singular animate (human) proper noun (NP1) and a bare plural animate noun (NP2), as exemplified by the well-formed examples (3a-f).

- (3a) ... *dass Richard*_{SUBJ} *Künstlerinnen*_{DOBJ} *sah*_{SING}
... that Richard saw artists
- (3b) ... *dass Richard*_{DOBJ} *Künstlerinnen*_{SUBJ} *sahen*_{PLUR}
... that artists saw Richard
- (3c) ... *dass Richard*_{SUBJ} *Künstlerinnen*_{IOBJ} *dankte*_{SING}
... that Richard thanked artists
- (3d) ... *dass Richard*_{IOBJ} *Künstlerinnen*_{SUBJ} *dankten*_{PLUR}
... that artists thanked Richard
- (3e) ... *dass Richard*_{SUBJ} *Künstlerinnen*_{IOBJ} *gefiel*_{SING}
... that Richard pleased artists
- (3f) ... *dass Richard*_{IOBJ} *Künstlerinnen*_{SUBJ} *gefielen*_{PLUR}
... that artists pleased Richard

Notice that the grammatical functions indicated by the subscripts are the ones required by the grammatical structure of the clause, more specifically, by agreement between number of the Subject and the SINGular or PLURal verb. In isolation, the nouns themselves can be Subj, IObj or DObj.

The verbs in (3a-f) exemplify three syntactically different classes: accusative verbs (*sehen*, governing DObj), dative-active verbs (*danken*, with IObj), and dative-experiencer verbs (*gefallen*, also with IObj). The former two classes are known to prefer animate Subjects; the third group prefers their Subj to be inanimate, their IObj animate. The empirical phenomena to be simulated by SINUS can be summarized as follows:

1. Strength of preferential function assignment to an animate NP1: Subj > IObj > DObj.
2. When, due to the requirements of subject-verb agreement, NP1 cannot be Subj, then the verb yields a garden-path effect (3b/d/f) in comparison with their counterparts (3a/c/e). As shown in reading time studies, the strength of this effect is strongest for accusative verbs, intermediate for dative-accusative verbs, and mildest for dative-experiencer verbs.
3. These garden-path effects are accompanied by three different ERP patterns at the verb. Compared to clauses where NP1 is Subj (3a/c/e), the clauses with NP2 as Subj yield an ERP wave with a positive deflection if the verb is accusative. However, the deflection is negative for both classes of dative-taking verbs—more pro-

nounced for dative-active than for dative experiencer verbs.

Serial Parsing Architectures Both the EADM and the L&CM parsers are serial in the sense that, when confronted with ambiguity, they pursue only one possible analysis, and consider further options only after earlier options have failed. Both models assume a first-pass parsing preference that assigns the function of Subj to NP1 and that of DObj to NP2. As soon as a plural clause-final verb is processed, or when the verb requires an IObj instead of a DObj, the initial function assignment causes a problem. This triggers an operation in another phase whose nature and complexity differ between verb types. The two models diverge widely with respect to the hypothesized phases and processes. We cannot and need not review the theoretical assumptions of the models here because our chief aim is to demonstrate that the above pattern of results does not necessitate a serial parser and can be simulated by a parallel, single-stage parsing mechanism.

In the following Section, we describe the essential properties of SINUS, a parallel self-organizing dynamic parser that we claim can simulate the target (and many related) phenomena at least qualitatively.

The Unification Space

Unification Space 2000 & Performance Grammar

The Unification Space model by Vosse & Kempen (2000), a single-pass model of human syntactic structure building, accounts for a considerable range of parsing preferences and garden-path phenomena. The model (henceforth called US2000) is based on the Performance Grammar formalism (PG; Kempen & Harbusch, 2002; Harbusch & Kempen, 2002). PG is a strongly “lexicalized” grammar—one assuming that the information needed to build grammatically correct sentences is associated with the individual lexical items. In PG, the grammatical information—word class & subcategorization in particular—is represented in treelets called LEXICAL FRAMES (somewhat similar to “elementary trees” in Tree Adjoining Grammar). Figure 1a shows two such treelets—one represents the personal pronoun *er* ‘he’, the other the finite verb *schliefe* ‘slept’. The branches of the latter lexical frame indicate that *schliefe* requires a Subj, and possibly one or more Modifiers. Other verbal lexical frames may also contain branches for DObj, IObj, and clausal complements. The node in the top layer of a lexical frame is called the




Figure 1. LEFT: Binding (dashed line) of lexical frames associated with *er* and *schliefe*. RIGHT: Derived tree.

ROOT node, the phrasal nodes in the third layer are FOOT nodes.

Syntactic trees for word groups and sentences are constructed by UNIFYING the root node of one treelet with a foot node of another treelet (the dashed line in Figure 1), provided that conditions on feature agreement and linear order are satisfied. A dynamic node represents the fact that two nodes are unified at a certain point in time. Whenever the activation level of this node is above a certain minimum value, root and foot node are supposed to be unified.

The lexical frames associated with the individual words of a sentence often allow more unifications than the actually correct one. In such cases, the model represents all these possibilities simultaneously and invokes lateral inhibition to make a selection. Parse trees thus emerge as the result of a dynamic self-organizing process. This architecture enables small revisions on the fly, without the need for discrete stages of parsing.

A localist-connectionist neural network

The US2000 model embodies a neurocognitive implausibility: It assumes that syntactic nodes and the connections between them can be created on-line. To address this issue, we started the development of SINUS—a parsing mechanism that is based on a “fixed” neural network, and where syntactic structure is represented in patterns of activation and inhibition, without the need for dynamic modifications to the network’s connectivity. This project also provided an opportunity to improve the linguistic plausibility of the parser by equipping it with a more sophisticated linear order component based on TOPOLOGIES. Our third goal was to remove a shortcoming for which US2000 has been criticized (Kamide & Mitchell, 1999): its lack of predictive parsing.

A topology is a one-dimensional array consisting of a small number of slots. Every phrase has a topology associated with it, but in this paper we only address clausal topologies. In our current implementation, they have nine slots, distributed over three “fields”: Forefield, Midfield and Endfield. The notion of “topological fields” stems from traditional German linguistics. Every slot serves as a destination for one or more types of syntactic functions and constituents; and many constituent types have several optional destinations. In terms of our present target structures, the subordinating conjunction *dass* ‘that’ and the clause-final verb fill the first and the last Midfield slots, whereas the Subj-, IObj- and DObj-NPs each have several placement options in between (thus giving rise to linear order flexibility and “scrambling” phenomena).

The SINUS network consists of a fixed number (eight in the current implementation) of identical columns, each capable of representing properties of any incoming word, whose connectivity derives from a symbolic grammar. The columns are divided into six layers, each representing a specific grammatical resource. From bottom to top (Figure 2):

1. **Input** This layer represents the current input word. During a few processing cycles, it activates the morphosyntactic information from “its” lexical entry.
2. **Word Category** This layer represents the lexical head of the frame(s) associated with the input word. Features

such as number, person and animacy are also found here. Several categories can be active simultaneously.

3. **Lexical Frame** This layer represents the lexical frame associated with the input word. Frames have features, some of which stemming straight from the lexicon, others resulting from unifications. Multiple frames may be active at the same time.
4. **Unification** This layer contains Unification nodes (U-nodes) that represent virtual unification (the dashed line in Figure 1). U-nodes receive activation from a root and a foot node of two different frames, and possibly from feature nodes. E.g., the activation level of a Subject U-node is influenced by the activation levels of the number and person features involved. The influence of the root node decreases slightly with distance.
5. **Linear Order** Nodes in this layer attempt to link U-nodes to a free and legal topology slot, thereby guarding the integrity of constituent order.
6. **Topology** Each node in this layer represents a slot in a topology. Activation of these nodes indicates that the slot is filled and signals to a potential new filler constituent to look for an empty slot further to the right.




Figure 2. Activation spreading through the SINUS network. The Linear Order layer is not shown.

From each layer, activation spreads to the adjacent higher and lower layers. Within each layer, incompatible elements inhibit each other. During each processing cycle, the activation level of all nodes is updated by adding feed-forward activation from the layer above; inhibition is subtracted, and activation decay is added. Incoming activation is transformed via the non-linear function in Appendix A.

Figure 2 shows SINUS at work for the simple clause *Er schlief* ‘He slept.’ Node 1 in the Input layer (representing *er*) is activated first. Activation spreads to node 2 in the Category layer (Pronoun), which in turn activates node 3 in the Frame layer (NP). This activation spreads further to U-node 4 (Subj-NP), thus offering partial support for the Subject role of *er*. The activation level of node 4 does not reach its maximum level since it only has support from a foot node: No root node that might serve as unification partner is seen yet. Node 4 partially activates node 5 in the topology layer

(Forefield slot). In the mean time, input node 1 has ceased to be active, which however doesn't mean that the other nodes cannot remain active: Since activation flows down from 5 to 4, to 3, to 2, and back upward again, the nodes sustain each other's activation. This goes on for some time, after which the next word enters and node 6 (*schlieft*) becomes active. This activates node 7 (V) and node 8 (S). At that moment, there is sufficient support for node 4 (Subject), which now can reach its full activation level and also fully activate node 5 (Forefield). At that moment, a parse can be extracted: a clause headed by *schlieft* with *er* heading the Subj-NP and placed in the Forefield slot.

In order to extract the parse tree that SINUS represents at a given moment: Collect the most active lexical frames per column (two in the example: an NP- and an S-frame), and unify those root-foot pairs whose U-node have the highest activation in the footnote's column.

Delayed Unification The activation of a Unification Node does not only depend on the activation of its root and foot nodes, but also on the outcome of a feature agreement check. (The implementation of the checking mechanism need not concern us here.) In the example *Richard schlieft* 'Richard slept', the person and number features of *Richard* and *schlieft* agree, so there is no penalty for the Subject U-Node in the form of inhibition. However, had we changed the verb to the plural form *schlafen*, then the outcome would have been different. Figure 3 shows the consequence. When the verb is introduced, the Subject U-node has low activation. When NP and S agree (the continuous line), the activation of Subj quickly rises to maximum, but when they do not agree (the dashed line), inhibition from other U-nodes suppresses the Subject U-node. It does continue to receive input activation from the verb, though, and after some time it manages to overcome the inhibition (also helped by feedback from the Unification to the Lexical Frame layer; not discussed here). The resulting analysis is identical to that of the correct sentence, but took more time.

Delayed binding could somehow be related to the appearance of a P600 component in the ERP signal elicited by grammatical errors in input sentences (Hagoort, 2003). Below we will come back to this suggestion, since low unification speed is also one of the factors involved in the strength of the garden-path effects elicited by our target sentences (3b/d/f).

Predictive parsing SINUS offers the possibility of provisionally assigning a grammatical function to a root node before a suitable unification partner becomes available. E.g., although *Richard* is three-way case-ambiguous, in *Richard schlieft*, NP *Richard* is immediately analyzed as Subject, before the verb has entered. This is a consequence of activation coming in from NP *Richard* and feedback from the topology. However, a U-node with only one active input reaches a much lower activation level than one with two comparable inputs. This makes these early assignments relatively easy to overcome.

Once activated, a node in one column can spread activation to nodes in the next-higher or next-lower layer of the next column. For instance, the topological layer can feed




Figure 3. Delayed unification due to missing subject-verb agreement.

feed into the Unification layer of the next column and pre-activate one or more grammatical functions; and the Lexical Frame layer can do the same. At the same time, inhibition from active U-nodes will suppress already active functions.

This dynamic interplay between partly converging and partly opposing forces may be illustrated by the example in Figure 2. After *schlieft* 'slept' has entered column 2 and its activation has reached the Topological layer, activated nodes in that area of the network will spread activation to DObj, IObj, and Mod U-nodes in column 3. Simultaneously, the lexical frame of *schlieft* will activate Subj and Mod there, whereas the already active Subj function in column 1 will inhibit its counterpart in column 3. Consequently, Mod will gain most activation (albeit not very much) and slightly activate Category nodes PP and AdvP. Thus, the network may be said to "expect" a modifier after *Er schlieft*. Had the sentence been *Er kaufte* 'he bought' instead, the activation pattern in column 3 would have been somewhat different: The transitive verb *kaufen* also pre-activates the DObj U-node. This would cause this U-node to be the highest activated U-node in column 3, and the system would "expect" a DObj and an NP. Below we will see that predictive parsing plays an important role in parsing preverbal NPs.

Computer simulations

Settings The model (which is fully deterministic; no random factors) consists of about 3000 PDP-like nodes, with different parameters for each layer, and a small number of control parameters (e.g., number of cycles for each input word). In total, there are 32 free parameters. The structure of the network was derived from a (simple) PG-grammar that specifies the hierarchical structures and topologies for many possible single-clause structures in German, with six types of verbs, different types of NPs, determiners, adjectival, adverbial and prepositional phrases. Lexical Frames activate U-nodes for each of the functions they contain; e.g., an NP frame with a noun head activates Determiner and Modifier functions in all active columns; a frame for an intransitive verb activates Subj nodes, an accusative-taking verb Subj and DObj nodes, and a dative-taking verb Subj and IObj nodes.

We trained the model using simulated annealing on a set of eight sentences, two of which are part of the simulations to be presented below: (3a) and (3c). The other sentences trained the system for different orders (SVO, OVS, VSO)

and/or contained case-unambiguous NPs with definite determiners. As the training set of eight sentences cannot claim to be representative for the language input of an average language learner, we introduced two linguistic preferences by hand-coding a number of connection weights:

- The first Midfield slot after the subordinating conjunction has a relatively strong preference to receive the Subj-NP, a weaker preference for the IObj, and a dispreference for DObj. No (dis)preferences were built into the other slots. These measures reflect the outcome of a corpus frequency study (Kempen & Harbusch, 2005).
- Dative-experiencer verbs send slightly less activation to the Subj U-nodes, and slightly more to the IObj U-nodes, in comparison with dative-active verbs. This is a somewhat *ad hoc* measure to reduce the Subj-before-IObj preference. (We are currently exploring a more principled way, based on an animate-before-inanimate preference that reduces the Subject-first preference. Subj-NPs of dative-experiencer verbs tend to be inanimate.)

The trained model is capable of parsing common single-verb main and subordinate clauses of German.

Simulation results The resulting model quickly assigns the first case-ambiguous NP (Richard) the provisional role of Subject. The U-node representing this assignment inhibits its counterpart in the next column. Similarly, the IObj receives some activation and inhibits the IObj U-node in the next column. Then, if the next input word is also a case-ambiguous noun, it will become the DObj—this being the only function that is not suppressed from the first column. (We disregard here the very weak influences of U-nodes that represent other grammatical functions, e.g. Prepositional Object.) The upshot is a strong preference for an initial Subj-before-DObj analysis of two fully case-ambiguous NPs in a subordinate clause.

When the verb arrives, it sends activation to those U-nodes in the preceding columns that correspond to their argument structure. This has the effect of boosting the activation levels of those nodes, provided that the agreement checks succeed. For the accusative *sah* in (3a), this means that the initial Subj-before-DObj analysis is reinforced. The analyses of the five other clauses proceed less smoothly. The activation that the two singular dative verbs *dankte* and *gefiel* send to the IObj U-nodes in the preceding columns, is sufficient for these U-nodes to win the competition with the DObj U-node. Because of feature agreement, only the Subj node for NP1 receives activation. Hence, *Künstlerinnen* will end up as Indirect Object in sentences (3c) and (3e).




Figure 4. Activation transmitted by *sahen* and *dankten* to U-nodes in NP1 and NP2 columns (thick lines). Size of squares reflects activation level at verb onset.

More complications arise when the verb is plural. In all three clauses, due to the requirements of feature agreement, the activation for Subject role does not reach NP1 (S1 in in Figure 4) but NP2 (S2). But whereas *sahen* conveys activation to both DObj NPs (D1, D2), *dankten* and *gefielen* send it to the IObj NPs (I1, I2). Consequently, in the case of *sahen*, two nodes, D1 and S2, need to overcome the inhibition from S1 and D2 in order to reach the correct analysis; and this problem is even enlarged by activation received by D2 at the same time. Hence, for *sahen* it is much harder to reach the correct analysis than for the dative verbs, where two freshly activated functions (I1, I2) each receive less inhibition than D1, three nodes (I1, I2 and S2) conspire to suppress S1 and D2, and D2 does not receive activation. In fact, (3d) and (3f) reach a correct analysis, but (3b) sticks to the initial Subj-DObj parse.

The resulting time course of the parsing process for clauses (3a-f) is summarized in Figure 5. The light grey timelines denote periods of unstable analyses, starting with the first fully correct one. Dark grey bars mark the period of stable correct analyses—except for *sahen*, which from the 15th cycle onward yields a stable but incorrect Subj-before-DObj analysis. The patterns agree with the size of the garden-path effects obtained in psycholinguistic studies.

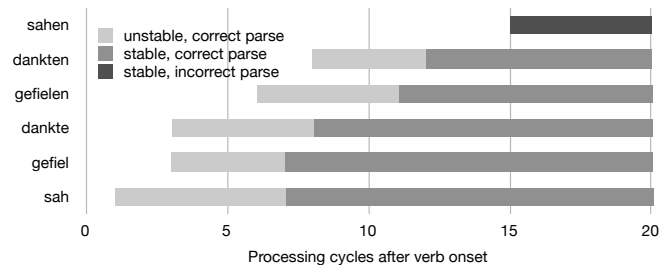


Figure 5. Time course of the parsing process.

Useful information about the dynamic interplay of activation and inhibition is provided by the SUM of the activation levels of all U-node at a given processing cycle. Figure 6 shows the development of this sum over time, onwards from the cycle in which the verb enters. We plotted the INCREASE of the summed activation. The time course of this sum score for the three singular verbs rises relatively fast at virtually identical pace, mainly because the Subject assignment need not change. The sum score rises more slowly for the plural




Figure 6. Increase of summed U-node activation after verb onset, for sentences (3a-f).

dative-taking verbs, and then overshoots the curve for the singular verbs—*dankten* for a somewhat longer time span than *gefielen*—before finally stabilizing at the same level as the singular verbs. When *sahen* ends the clause, the sum score undershoots the singular-verbs curves for a considerable number of cycles (and fails to reach the correct parse).

This aspect of SINUS' dynamics converges remarkably with the qualitatively different ERP effects observed for the clauses (3b/d/f): a positivity for *sahen* relative to *sah*; a negativity for *dankten* relative to *dankte*, and a somewhat weaker negativity for *gefielen* relative to *gefiel*. Given this result, SINUS may be said to reflect an important aspect of the ERP effects, at least in a qualitative sense.

It has been suggested before that parsing dynamics are involved in ERP effects. Hagoort (2003) hypothesized that a slow unification is a possible source of the SPS/P600 effect. Our simulations support this hypothesis, with the modification that it is the summed activation speed of all unification nodes rather than that of a single node. However, they do not exclude other causes, such as a well-formedness check. They also hint at the possibility that certain negativities might find their origin in parser dynamics as well.

The limited space available prevents us from discussing other aspects of the model here. But we can report that it also shows the desired behavior for comparable all-singular and all-plural clauses, with verbs in declarative main clauses, and in direct questions with a fronted non-nominative *wh*-phrase.

Discussion and future work

SINUS is a constraint-based, dynamic, self-organizing, parallel parsing system that can assign grammatical functions before the lexical head is encountered, and predict properties of upcoming words and phrases. We have demonstrated its capability of modeling differential garden-path and ERP effects between superficially very similar German sentences, at least in a qualitative manner, without the need for discrete first-pass parsing and reanalysis stages.

In future work, we have to broaden SINUS' empirical coverage, enabling it to model a wider range of psycholinguistic phenomena (in German and English), not only qualitatively but also quantitatively. This will require considerable extension of the system's syntactic sophistication. As for parsing German clauses with canonical and scrambled constituent orders, we presently do not foresee a need to invoke semantic (thematic) roles in explanations of garden-path effects such as the ones discussed here. The only meaning-related notion, we estimate now, is animacy.

SINUS generates an unexpected prediction for sentences of the types featuring in the present article: Agreement-based reversal of Subject and Object functions should be possible for accusative verbs that prefer an animate object and an inanimate subject. Such verbs will behave like *dankten*. We hope to put predictions like this one to experimental test.

Note The software used for the simulations is available for research purposes. Contact one of the authors in order to obtain the code or a demonstration version.

Appendix A

The exponential mapping of inputs I_i is $f = (1 - e^{\alpha \sum I_i}) / (1 - e^{\alpha M})$, where M is the sum of the maximum values of all inputs; $f(0, 0, 0, \dots) = 0$, and $f(M_1, M_2, M_3, \dots) = 1$. The shape of f is controlled by α , where $\alpha > 0$ gives a function with a positive first and second derivative.

Acknowledgements

The work reported here was supported by the Dutch Organization for Scientific Research (NWO) as part of Cognition Program (PLUS project). We thank two anonymous reviewers for their valuable comments.

References

- Bader, M., & Bayer, J. (2006). *Case and Linking in Language Comprehension*. Dordrecht: Springer.
- Bornkessel, I., & Schlesewsky, M. (2006). The Extended Argument Dependency Model: A neurocognitive approach to sentence comprehension across languages. *Psychological Review*, 113, 787-821.
- Hagoort, P. (2003). How the brain solves the binding problem for language: a neurocomputational model of syntactic processing. *NeuroImage*, 20, S18-S29.
- Harbusch, K. & Kempen, G. (2002). A quantitative model of word order and movement in English, Dutch and German complement constructions. In: *Proceedings of the 19th International Conference on Computational Linguistics (COLING-2002)*, Taipei (Taiwan). San Francisco: Morgan Kaufmann.
- Hemforth, B., & Konieczny, L. (2000). *German sentence processing*. Dordrecht: Kluwer.
- Kamide, Y., & Mitchell, D. C. (1999). Incremental pre-head attachment in Japanese parsing. *Language and Cognitive Processes*, 14, 631-662.
- Kempen, G., & Harbusch, K. (2002). Performance Grammar: A declarative definition. In: A. Nijholt, M. Theune & H. Hondorp (Eds.), *Computational Linguistics in The Netherlands 2001*. Amsterdam: Rodopi.
- Kempen, G. & Harbusch, K. (2005). The relationship between grammaticality ratings and corpus frequencies: A case study into word order variability in the midfield of German clauses. In: S. Kepser & M. Reis (Eds.), *Linguistic Evidence—Empirical, Theoretical, and Computational Perspectives*. Berlin: Mouton De Gruyter.
- 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.
- Scheepers, C., & Vasishth, S. (2007). What behavioural measures reveal about thematic reanalysis and the use of structural frequency in verb-final sentence processing. Paper submitted for publication, Department of Psychology, University of Glasgow.
- Vosse, T., & Kempen, G. (2000). Syntactic structure assembly in human parsing: A computational model based on competitive inhibition and lexicalist grammar. *Cognition*, 75, 105-143.