INTEGRATING WORLD KNOWLEDGE WITH COGNITIVE PARSING

Harold Paredes-Frigolett and **Gerhard Strube**
Center for Cognitive Science
Institute for Computer Science and Social Research
University of Freiburg
D-79098 Freiburg i. Brg.
Germany
{paredes,strube}@cognition.iig.uni-freiburg.de

Abstract

The work presented in this article builds on the account of cognitive parsing given by the SOUL system (Konieczny & Strube, 1995), an object-oriented implementation of Parameterized Head Attachment (Konieczny *et al.*, 1991) based on Head-Driven Phrase-Structure Grammar (Pollard & Sag, 1994). We describe how the initial semantic representation proposed by the parser is translated into a logical form suitable for inference, thus making it possible to integrate world knowledge with cognitive parsing. As a semantic and knowledge representation system we use the most expressive implemented logic for natural language understanding, Episodic Logic (Hwang & Schubert, 1993), and its computational implementation, Epilog (Schaeffer *et al.*, 1991).

Introduction

The work reported in this article can be seen as a continuation of psycholinguistic research in sentence parsing (Strube *et al.*, 1990; Konieczny *et al.*; Hemforth *et al.*, 1993) and computational models of human parsing (Konieczny & Strube, 1995). This former work has resulted in a psycholinguistic theory of sentence parsing called Parameterized Head Attachment and in the SOUL parser as a computer model that implements the theory. SOUL makes use of typed feature formalisms to describe the syntax of natural language, especially Head-Driven Phrase-Structure Grammar, HPSG (Pollard & Sag, 1994). SOUL operates in a word-by-word incremental fashion, computing a HPSG structure according to human parsing preferences that have been identified in guiding the first syntactic analysis.

This model of human parsing has to be complemented by processes of semantic interpretation for two reasons. On the one hand, SOUL needs to judge the appropriateness of the analyses given by the parser based on world knowledge in order to trigger reanalysis whenever necessary. On the other hand, the preliminary logical form proposed by SOUL has to be transformed into a logical form that allows for inferences to be drawn in the knowledge base. Our model enables the integration of world knowledge with cognitive parsing and benefits from using Episodic Logic, the most expressive formal computational logic for general natural language understanding. The objective of the present article is to show how the output of the SOUL parser can be semantically interpreted into an initial representation, which is further transformed removing scope ambiguities and context-dependency. The final result is a logical form in Episodic Logic (Hwang, 1992; Hwang & Schubert, 1993) that is suitable for inference. We also show how this final logical form accounts for semantic biasing of syntactic analyses.

Parameterized Head Attachment

Cognitive parsing in the SOUL system, Semantics-Oriented Unification-Based Language Processing (Konieczny & Strube, 1995), is based on the Parameterized Head Attachment Principle of sentence processing (Konieczny *et al.*, 1991), PHA henceforth. PHA is a theory of sentence processing that originated in recent psycholinguistic investigations of PP-attachment and other phenomena in German, using self-paced reading and eye-movement experiments (Strube *et al.*, 1990; Konieczny *et al.*, 1991; Hemforth *et al.*, 1993; Konieczny *et al.*, 1994). Though a serial model in the tradition of the Garden-Path Theory (Frazier & Fodor, 1978), PHA is a model of sentence processing whose results differ from those predicted by Minimal Attachment and Late Closure, and by the principles put forth with the Construal Theory (Frazier & Clifton, 1996).

PHA consists of the following principles:

- **Head attachment**: Prefer to attach an item to a phrasal unit whose lexical head has already been read.
- **Preferred-role attachment**: Prefer to attach an item to a phrasal unit whose head preferentially subcategorizes for it.
- **Most recent head attachment**: Prefer to attach an item to a head that was read most recently.
- **Parameterized head attachment**: Attempt to apply **Head Attachment** before **Preferred Role Attachment** before **Most Recent Head Attachment**.

SOUL is a cognitive parser that functions complying with the principles described above and has been implemented using HPSG.[1] We will now concentrate on the syntax/semantics

[1] The interested reader is referred to Konieczny & Strube (1995) for a description of the SOUL system.

interface used in the system, called the preliminary logical form, and the translation process of this initial representation into a representation suitable for inference using Episodic Logic (Hwang, 1992, Hwang & Schubert, 1993).

Incremental semantic interpretation, scoping, and deindexing

Consider the natural language sentence (1), as presented below:

(1) Marion $_{\uparrow a}$ watched $_{\uparrow b}$ the $_{\uparrow c}$ horse $_{\uparrow d}$ with $_{\uparrow e}$ the $_{\uparrow f}$ white $_{\uparrow g}$ fleck $_{\uparrow h}$. $_{\uparrow i}$

Using this expression, we illustrate the process of incremental semantic interpretation, scoping, and deindexing by virtue of which the preliminary logical form is transformed into a so-called episodic logical form suitable for inference in Epilog (Schaeffer *et al.*, 1991), the computational system for Episodic Logic, EL henceforth.

The extended syntax/semantics interface

We consider the parse obtained by SOUL at point *f*, once the second definite article *the* has been absorbed from the input string. Using subcategorization information for the preferred lexical entry of *watch*, a transitive verb with an instrumental complement, SOUL proposes a parse in which the expected prepositional phrase attaches to the main verb of the sentence. The semantic information is given in HPSG under the feature *content*. Due to its restricted expressiveness we have defined a more expressive syntax/semantics interface in SOUL, which we have termed *preliminary logical form*, $\mathcal{PLF}$ for short, as shown below.

$\mathcal{PLF}_f$:

$$\begin{bmatrix} pred & watch \\ arg_{agent} & Marion \\ arg_{theme} & \begin{bmatrix} ind & ind_1 \\ spec & the \\ cond & \begin{bmatrix} pred & horse \\ arg_{inst} & \square_1 \end{bmatrix} \end{bmatrix} \\ arg_{instr} & \begin{bmatrix} ind & ind_2 \\ spec & the \\ cond & \begin{bmatrix} pred & P \\ arg_{inst} & \square_2 \end{bmatrix} \end{bmatrix} \\ a\text{-}mod & \begin{bmatrix} op & adv\text{-}a \\ pred & \begin{bmatrix} pred & with\text{-}instr \\ arg_{theme} & \square_2 \end{bmatrix} \end{bmatrix} \\ e\text{-}mod & \begin{bmatrix} op & past \end{bmatrix} \\ mood & \begin{bmatrix} op & decl \end{bmatrix} \end{bmatrix}$$

In $\mathcal{PLF}_f$, the action of Marion watching a horse and being modified so as to be performed with an instrument has been represented. This representation is based on the formalism used to represent the syntax/semantics interface in HPSG (Pollard & Sag, 1994). In $\mathcal{PLF}_f$, the predicate *watch* takes three arguments. The first argument corresponds to the agent of the event described, the second argument corresponds to the theme of the action being described, and the third argument corresponds to an instrument with which the action is performed. The format of the preliminary logical form is borrowed from Fenstad *et al.* (1987). The feature *cond* introduces a well-formed formula and the features *spec* and *ind* introduce the quantificational force of a determiner and the variable quantified over, respectively. In $\mathcal{PLF}_f$, the argument positions are filled with variables, but the scope of the quantifiers is left underspecified, much as in Schubert & Pelletier (1982), Fenstad *et al.* (1987) and Pollard & Sag (1994).

Pollard & Sag (1994) assume a situation-theoretic framework for semantic representation. Unfortunately, Situation Semantics (Barwise, 1987; Devlin, 1991) does not yet offer a framework for representing a variety of semantic phenomena. Note for example that in $\mathcal{PLF}_f$ the expected prepositional phrase is syntactically analyzed as an adverbial. The extension proposed here for the representation of adverbials distinguishes between those that operate on sentences and those that operate on monadic predicates. We take advantage of the representational framework put forth with EL (Hwang, 92; Hwang & Schubert, 1993) and adopt the functions *adv-e* and *adv-a*, which map predicates over episodes/actions into predicate modifiers. In EL, (*adv-e* π), with π a predicate over episodes, is an episode modifier, and (*adv-a* π), with π a monadic predicate over attributes/actions, is an action modifier. *with-instr* is a relational predicate taking in our example the monadic predicate P and mapping it into a monadic predicate over actions. In $\mathcal{PLF}_f$, the tense operator *past* and the mood operator *decl* have also been introduced under the features *e-mod* and *mood*, respectively.

Incremental semantic interpretation

After the parser has generated this initial representation, the preliminary attribute-value representation $\mathcal{PLF}_f$ is translated into a set of fact-schemata using recursive procedures that operate directly on the attribute-value matrices under the feature *cond*. Then, full schemata are constructed for the full sentence. For $\mathcal{PLF}_f$, we obtain the following set of fact schemata:

$$\begin{aligned} &\langle C_1 : horse, ind_1, 1\rangle \\ &\langle C_2 : P, ind_2, 1\rangle \\ &\langle C_3 : watch, Marion, ind_1, ind_2, 1\rangle \end{aligned} \quad (1)$$

Incremental scoping

In HPSG the feature *QSTORE* gives the scope disambiguation information. Our approach to scope disambiguation consists in generating a scoped logical form according to psycholinguistically plausible heuristics for scope disambiguation (Kurtzman & MacDonald, 1993). In our example, the surface speech act operator *decl* is assigned scope over the

whole sentence, the tense operator *past* is assigned scope within speech act operators and wider scope than all other operators. The two definites are assigned scope through all barriers only if they are salient in the current context, otherwise they are scoped within tense operators.

We obtain the set of scoped fact schemata in (2):

$$(The\ (\langle ind_1 \mid C_1 \rangle)\ (\langle ind_1 \mid The\ (\langle ind_2 \mid C_2 \rangle)(\langle ind_2 \mid C_3 \rangle)))) \quad (2)$$

Replacing C_1, C_2, and C_3 in (2) and incorporating the action-modifying operator *((adv-a (with-instr ind_2))*, the tense operator *past*, and the mood operator *decl*, we obtain the parameterized indexical logical form $\mathcal{PILF}_f$.[2]

$\mathcal{PILF}_f$:

```
(decl (past (The x:[x horse]
              (The y:[y P]
                   [Marion ((adv-a (with-instr y))
                            (watch x))]))))
```

Incremental deindexing

To be suitable for inference in Epilog, the computational system for EL, the final logical form has to be independent of context. Thus, the indexical information conveyed by the tense operator *past* and the mood operator *decl* has to be brought into the final representation. The incremental deindexer applies the compositional deindexing rules for translating the indexical logical forms put forth in Hwang & Schubert (1992). These rules transform the indexical logical form containing tense operators and adverbials into a so-called parameterized episodic logical form, $\mathcal{PELF}$ for short.[3] At point f, the result of this deindexing process is the parameterized episodic logical form $\mathcal{PELF}_f$.

$\mathcal{PELF}_f$:

```
(∃u1:[[u1 same-time Now1] ∧
      [u0 immediately-precedes u1]]
      [[Speaker tell Hearer (That
        (∃e1:[[e1 before u1] ∧ [e0 orients e1]]
             [[(The x:[x horse]
                  (The y:[y P]
                       [[Marion | e1]
                        ((with-instr y)(watch x))]))]
              ** e1]))]
      ** u1])
```

[2] We assume that the definites in question are not salient in the current context.

[3] Due to space limitations, we do not present the deindexing rules here. The interested reader is referred to Hwang (1992) and Hwang & Schubert (1992) for a detailed description.

Expressing knowledge in EL

Knowledge in Epilog takes the form of meaning postulates and world knowledge axioms. These are expressed as probabilistic conditionals of form $\phi \rightarrow_{p,\alpha_1,\ldots,\alpha_n} \psi$, where $\alpha_1, \ldots, \alpha_n$ are controlled variables and p is a statistical probability.

We now introduce the following meaning postulates and world knowledge axioms about people seeing objects.[4]

Meaning postulates about people seeing objects:

If a person watches a thing, then that person sees that thing.[5]

```
MP 1: (∃x:[x person]
         (∃y:[y thing]
            (∃e1:[[[[x | e1] P] ∧ [x watch y]] ** e1])))
      → (∃e2:[e1 ≼ e2]
            [[[[x | e2] P] ∧ [x see y]] ** e2])
```

If a person sees something with something, then she/he is seeing it with a viewing instrument.

```
MP 2: (∃x:[x person]
         (∃y:[y thing]
            (∃z:[z thing]
               (∃e1:[[[[x | e1] (with-instr z)] ∧
                      [x see y]] ** e1]))))
      → [z ((nn viewing) instrument)]
```

Meaning postulate about unlocated formulas:

$\mathcal{MP}$ 3: $(\forall e_1{:}[[[\phi \wedge \psi] ** e_1] \rightarrow [\phi \wedge (\exists e_2{:}[e_2 \preceq e_1][\psi ** e_2])]]]$

World knowledge axiom about people seeing things with viewing instruments:

If someone sees something with a viewing instrument, then she/he probably sees it clearly.

```
WK 1: (∃x:[x person]
         (∃y:[y thing]
            (∃z:[y ((nn viewing) instrument)]
               (∃e1:[[[[x | e1] (with-instr z)] ∧
                      [x see y]] ** e1]))))
      →0.85,e1,x,y (∃e2:[e1 ≼ e2]
                        [[[[x | e2](in-manner clear)] ∧
                         [x see y]] ** e2])
```

[4] In the meaning postulates above, the operator $\preceq$ is a metalogical operator that corresponds to the operator *coextensive-part-of* in EL. $e_1 \preceq e_2$ indicates that situation e_1 is coextensive with situation e_2, that is, e_1 and e_2 have the same spatiotemporal location. Finally, the modal operators $*$ and $**$ are introduced with the following intuitive meanings: $[\phi ** \eta] \equiv \phi$ *describes* η *as a whole* and $[\phi * \eta] \equiv \phi$ *describes some part of* η.

[5] Here, P is a parameter that stand for a monadic predicate over actions/attributes.

Triggering inferences in Epilog

After a formula is asserted in Epilog, a process of input-driven inference triggering via rule instantiation using meaning postulates and world knowledge axioms is started. This process consists of the following six steps: existentially quantified variables are skolemized, top-level conjuncts are split, simplification schemas are applied, the new formulas are checked for their consistency with the previously stored knowledge, the new formulas are classified and stored in the knowledge base, and a process of input-driven inference chaining is started using the meaning postulates and world knowledge axioms in the knowledge base.

For our example above we define the following simplification schema:

$\mathcal{SS}$ 1: For P a parameter: $(\forall x{:}[x\ P] \rightarrow [x\ \text{thing}])$

Skolemizing e_1/E_1, x/X, and y/Y in $\mathcal{PELF}_f$, splitting conjunctions, and applying the simplification schema above we obtain:[6]

F_1 $[E_1$ before $U_1]$

F_2 $[X$ horse$]$

F_3 $[Y$ thing$]$

F_4 [[Marion | E_1] (with-instr Y)]

F_5 [[[[Marion | E_1] (with-instr Y)] $\wedge$ [Marion watch X]] ** E_1]

The described episode E_1 is characterized by the action of Marion watching a horse and being modified so as to be performed with an instrument. F_5 matches the antecedent of $\mathcal{MP}_3$. We obtain the following additional formulas:

F_6 $[E_2 \preceq E_1]$

F_7 [[Marion watch X] ** E_2]

Episode E_2 is coextensive with episode E_1 and is characterized by the action of Marion watching a horse. Let us now assume the following facts:

F_8 [Marion woman]

F_9 [Marion person]

F_{10} $[X$ thing$]$

F_9 and F_{10} are obtained by type-hierarchical knowledge. F_3, F_5, F_9, and F_{10} match the antecedent of $\mathcal{MP}_1$,[7] a meaning postulate connecting watching with seeing events. Thus, the following formulas can be obtained:

F_{11} $[E_3 \preceq E_1]$

F_{12} [[[[Marion | E_3] (with-instr Y)] $\wedge$ [Marion see X]] ** E_3]

F_{12} matches the antecedent of $\mathcal{MP}_3$. This accounts for obtaining the following formulas:

F_{13} $[E_4 \preceq E_3]$

F_{14} [[Marion see X] ** E_4]

F_3, F_9, F_{10}, and F_{12} match the antecedent of $\mathcal{MP}_2$, a meaning postulate about seeing things with instruments. Thus, we obtain the following additional formula:

F_{15} $[Y$ ((nn viewing) instrument)]

At point f, we are able to make the prediction that Y is a viewing instrument.[8] Finally, by a process of input-driven inference chaining using $\mathcal{WK}_1$, a world knowledge axiom about seeing things with viewing instruments, we obtain from F_9, F_{10}, F_{12}, and F_{15} the following inferences:[9]

I_1 $[E_5 \preceq E_3]$

I_2 [[[[Marion | E_5] (in-manner clear)] $\wedge$ [Marion see X]]$_{0.85}$ ** E_5]

Incremental interpretation

Although inferences can be triggered at any point by the process described above, we contend that full inferences are not required to be triggered on a word-by-word incremental basis. Our claim is based on the assumption that the number of inferences made during on-line sentence comprehension is to be constrained by computational resources, and on the observation that the inferences drawn during incremental on-line sentence comprehension are a function of the rhetoric aspects used by the writer/speaker. To account for this in our model we introduce the notion of restricted inference.

Restricted inference

We define *restricted inference* as a process of hierarchy climbing, logical form pattern matching, and inference triggering using only meaning postulates in the knowledge base. For semantic biasing the parser uses this restricted form of inferences on a word-by-word incremental basis. We believe that the use of a restricted form of inference based on meaning postulate inference is quite plausible. On the one hand, there

[6] We neglect the relation orients that may relate the reported episode E_1 to a prior episode E_0.

[7] With P = (with-instr Y) in the antecedent of $\mathcal{MP}_1$

[8] In formula F_{15}, nn is a function introduced in EL to map 1-place nominal predicates into predicate modifiers.

[9] We write F for inferences drawn by meaning postulates and I for inferences drawn by world knowledge axioms. The probability attached to I_2 tells us that inference I_2, though uncertain, is quite likely. It can be paraphrased as "Marion probably sees the horse clearly."

is enough empirical evidence suggesting that lexical inferences based on both the meaning of open class words[10] and the link between the meaning of a verb and its syntactic characteristics become available to human comprehenders as soon as the open class word or verb in question has been absorbed from the input string. On the other hand, there is a natural, though not ultimate, distinction between knowledge about lexical meanings and knowledge about the world. These observations lead to making the distinction between two modes of inference: (i) restricted inference for semantic biasing via meaning postulates and (ii) full inference for text understanding based on world knowledge axioms and meaning postulates. While the former will be used in a word-by-word incremental fashion, the latter will be used only at certain points, essentially where there is a single parse and the need for full inferential activity arises. This may be the case when resolving anaphoric antecedents, establishing causal and explanatory relations among events, or when a prosodic cue arises during on-line speech comprehension.

Fine-grained, weak interaction

The interface between the parser and the incremental interpreter corresponds to a structure called *analysis* that contains information concerning the analysis currently being pursued by the parser. This structure contains three fields corresponding to the top node of the partial parse tree for the analysis in question, a field that indicates whether or not the analysis has just been repaired by the parser, and a field that indicates whether the set of inferences drawn for the interpretation obtained so far is consistent or not with the previously asserted inferences in the knowledge base. This interface is weak in that the information needed by the parser for semantic biasing corresponds to just one bit, namely, to whether or not the formulas stored in the knowledge base are consistent with the new formulas asserted by the incremental interpreter, or with the inferences derived from them.

Restricted inference for semantic biasing

We now illustrate how the process of incremental interpretation interacts with the parser for purposes of semantic biasing.

After incremental semantic interpretation, scoping, and deindexing of the analysis preferred by SOUL at point h we obtain $\mathcal{ELF}_h$, a shown below:

$\mathcal{ELF}_h$:

(∃u_1:[[u_1 same-time $Now1$] ∧
[u_0 immediately-precedes u_1]]
[[Speaker tell Hearer (That
(∃e_1:[[e_1 before u_1] ∧ [e_0 orients e_1]]
[[(The x:[x horse]
(The y:[[y fleck] ∧ [y white]]
[[Marion | e_1]((with-instr y)
(watch x))))]]
** e_1]))] ** u_1])

[10] Examples of open class words are adjectives, verbs, nouns, and adverbs. Examples of closed class words are prepositions and determiners.

After asserting $\mathcal{ELF}_h$ we obtain:[11]

F_{16} [Y fleck]

F_{17} [Y white]

By type-hierarchical knowledge, an inconsistency arises between F_{16} and F_{17} and the predicted formula F_{15}. The inconsistency is reported, and the conflicting formula is stored in the knowledge base. The incremental interpreter sets the field *analysis.consistent* to the value *false* and returns the structure *analysis* to SOUL. At point h, SOUL proposes as preferred continuation the analysis in which the prepositional phrase attaches to the noun "the horse," sets the field *analysis.just-repaired* to the value *true* and sends the repaired analysis to the incremental interpreter. The incremental interpreter drops the set of formulas obtained for the first analysis, sets the field *analysis.just-repaired* to the value *false*, obtains the new episodic logical form, and finally asserts it in Epilog.

The resulting episodic logical form for the repaired analysis at point h is shown below:

$\mathcal{ELF}_h$':

(∃u_1:[[u_1 same-time $Now1$] ∧
[u_0 immediately-precedes u_1]]
[[Speaker tell Hearer (That
(∃e_1:[[e_1 before u_1] ∧ [e_0 orients e_1]]
[[(The x:[[x horse] ∧ (The y:[[y fleck] ∧ [y white]]
[x with-part y])]
[Marion watch x])] ** e_1]))] ** u_1])

Skolemizing e_1/E_1, x/X, and y/Y in $\mathcal{ELF}_h$' and splitting conjunctions we obtain a new set of formulas at point h:

F_1 [E_1 before U_1]

F_2 [X horse]

F_3 [Y fleck]

F_4 [Y white]

F_5 [X with-part Y]

F_6 [[Marion watch X] ** E_1]

The described episode E_1 is characterized by the action of Marion watching a horse with a white fleck.

[11] In Epilog, the skolem constant introduced for an existentially quantified variable is kept on the property list of the variable. Thus, future references to the variable are replaced by the constant within future formulas.

Conclusions

The work reported in this article presents a formal computational framework for integrating world knowledge with cognitive parsing, the center piece of which is a computational model of incremental semantic interpretation, scoping, and deindexing. The model is incremental in that a partial parse tree is transformed into a complete preliminary logical form in which parameters construed as metalogical variables are introduced in the logical form for the missing constituents in the input sentence, thus enabling further semantic processing in a word-by-word incremental fashion. The model is compositional in that each syntactic rule comes equipped with a semantic annotation. Thus, the semantics of a sentence at any point during incremental semantic interpretation is a function of the semantics of the constituents absorbed so far plus the semantics of the parameters introduced for the missing constituents in the corresponding partial parse tree. The psycholinguistic plausibility of the model is not only grounded in SOUL, but also in the use of psycholinguistic well-founded heuristics for scope disambiguation. In our current implementation,[12] we have defined a set of heuristics based on recent empirical results reported in Kurtzman & MacDonald (1993) to transform the preliminary logical form into a single parameterized indexical logical form that is scopally unambiguous, but still context-dependent.

Finally, we have described how our model of incremental semantic interpretation, scoping, and deindexing can be used for fine-grained, weakly interactive incremental interpretation. We think that the distinction between two forms of inference, restricted inference for semantic biasing based on lexical inference through meaning postulates and full inference for text understanding through input-driven inference chaining using both meaning postulates and world knowledge axioms, gives a plausible account of the incremental inference process. In this view of sentence comprehension using a first-analysis parser like SOUL, incremental interpretation is construed as a question answering process (Sanford, 1990). Thus, reanalysis is not triggered on account of whether or not the first analysis reaches a certain plausibility threshold, but rather on whether or not the expectations that arise during incremental interpretation are confirmed or refuted as more information becomes available from the input string.

Acknowledgements

This research was supported by the German National Science Foundation (Deutsche Forschungsgemeinschaft) through the Graduate School of Human and Computational Intelligence at the University of Freiburg. We want to thank our colleagues Barbara Hemforth, Lars Konieczny, and Christoph Scheepers for their helpful comments. Special thanks to Lenhart Schubert for many fruitful discussions on EL.

[12] A model of incremental semantic interpretation, scoping, and deindexing has been implemented in Common Lisp for the SOUL system.

References

Barwise, J. (1987). *The Situation in Logic*. Stanford, CA: CSLI.

Devlin, K. (1991). *Logic and Information*. Cambridge, UK: Cambridge University Press.

Fenstad, J.E., Halvorsen, P.K., Langholm, T. & van Benthem, J. (1987). *Situations, Language and Logic*. Reidel.

Frazier, L. & Fodor, J. (1978). The sausage machine: a new two-stage parsing model. *Cognition*, 6, 291–325.

Frazier, L. & Clifton, C. (1996). *Construal*. Cambridge, MA: MIT Press.

Hemforth, B., Konieczny, L. & Strube, G. (1993). Incremental sentence processing and parsing strategies. *Proceedings of the 15th Annual Conference of the Cognitive Science Society*, (pp. 539–545).

Hwang, C.H. (1992). A Logical Approach to Narrative Understanding. Doctoral dissertation. Edmonton, Alberta: University of Alberta, Department of Computing Science.

Hwang, C.H. & Schubert, L.K. (1992). Tense trees as the fine structure of discourse. *Proceedings of the 30th Annual Meeting of the American Association for Computational Linguistics*, (pp. 239–240).

Hwang, C.H. & Schubert, L.K. (1993). Episodic Logic, a comprehensive, natural representation for language understanding. *Minds and Machines*, 3, (pp. 381–419).

Konieczny, L., Hemforth, B. & Strube, G. (1991). Psycholinguistisch fundierte Prinzipien der Satzverarbeitung jenseits von Minimal Attachment. *Kognitionswissenschaft*, 2.

Konieczny, L., Scheepers, C., Hemforth, B. & Strube, G. (1994). Semantikorientierte Syntaxverarbeitung. In C. Felix, C. Habel and G. Rickheit (Eds.), *Kognitive Linguistik: Repräsentationen und Prozesse*. Opladen: Westdeutscher Verlag.

Konieczny, L. & Strube, G. (1995). SOUL: A cognitive parser. *Proceedings of the 17th Annual Conference of the Cognitive Science Society*, (pp. 631–636).

Kurtzman, H. & MacDonald, M. (1993). Resolution of quantifier scope ambiguities. *Cognition*, 48, (pp. 243–279).

Pollard, C. & Sag, I (1994). *Head-Driven Phrase-Structure Grammar*. University of Chicago Press.

Sanford, A. (1990). On the nature of text-driven inference. In D. Balota, G. Flores d'Arcais & K. Rayner (Eds.), *Comprehension Processes in Reading* (pp. 515–535). Hillsdale, NJ: Lawrence Erlbaum Associates.

Schaeffer, S., Hwang, C.H., de Haan, J & Schubert, L.K (1991). EPILOG: The computational system for Episodic Logic. Technical Report, Department of Computing Science, Edmonton, Alberta: University of Alberta.

Schubert, L.K. & Pelletier, J. (1982). From English to Logic: Context-free computation of conventional logical translations. *American Journal of Computational Linguistics*, 10, (pp. 165–176).

Strube, G., Hemforth, B. & Wrobel, H. (1990). Resolution of structural ambiguities in sentence comprehension: On-line analysis of syntactic, lexical, and semantic effects. *Proceedings of the 12th Annual Conference of the Cognitive Science Society* (pp. 558–565).