


C2

AAAI 88

The Seventh National
Conference on
Artificial Intelligence

August 21-26, 1988
Saint Paul, Minnesota



PROCEEDINGS • VOLUME 2

The Integration of Unification-based Syntax/Semantics and Memory-based Pragmatics for Real-Time Understanding of Noisy Continuous Speech Input*

Hideto Tomabechi and Masaru Tomita

Center for Machine Translation
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Real-time understanding of speech input is difficult especially because the input is often noisy and elliptic. Multiple morphophonemic and lexical hypotheses generated for a single input sentence cannot be resolved by local semantics alone. We have developed a system in which unification-based parsing of speech input is integrated with thematic memory-based spreading activation that supplies extra-sentential knowledge that can help to disambiguate noisy and elliptic real-time speaker-independent continuous speech input.

1 Introduction

The difficulty of parsing speech input is that unlike written text input, a parser receives a multiple number of hypotheses as input for a particular voice input. This is partly due to current limitations on speech recognition systems which are incapable of determining specific phonemes for each input and generally produce several possible segmentations of the hypothesized phonetic stream. It is not rare that a speech parser outputs 30 to 50 well-formed, semantically acceptable parse results for each independent sentence of a speech recognition device output. This paper describes our theory of integrating pragmatic (contextual and thematic) knowledge with the unification-based syntax/semantic parsing of real-time speaker-independent continuous speech input. This approach is adopted by our speech understanding system and implemented as a part of our speech translation system at the Center for Machine Translation (CMT) at Carnegie Mellon University.

Overview

Our speech understanding system consists of three parts:

- the speech recognition system (device hardware and control programs);
- a phoneme-based generalized LR parser (Φ GLR¹);

*The authors would like to thank members of the Center for Machine Translation for fruitful discussions and their efforts in implementing the subparts of the system. Hiroaki Saito has contributed significantly in implementing the phoneme-based Generalized-LR parser (Φ GLR). Teruko Mitamura developed a significant part of the Japanese LFG grammar. We would also like to thank Dr. Morii of Matsushita Research Institute for his generous contribution of the speech recognition hardware used by our project. Eric Nyberg was especially helpful in preparing the final version of this paper.

¹ Φ GLR parser is based on the generalized LR algorithm, augmented by a pseudo-unification package with semantic case-

- a memory-based recognition mechanism through spreading activation.

This paper focuses on our approach for integrating memory-based recognition (recognize-and-record paradigm) with unification-based (LFG²) syntax/semantic parsing (build-and-store paradigm) which has proven to be effective for understanding continuous noisy speech input. Readers may also wish to consult Morii, *et al*[1986] and Saito&Tomita[1988] respectively for discussions of the speech recognition hardware and the phoneme-based generalized LR parser. We will emphasize that our approach has truly integrated syntactic, semantic, and pragmatic (contextual and thematic) processing during real-time speech understanding where the use of each source of knowledge is interleaved and inter-dependent. This is in contrast to other approaches found in the text processing literature where parsing and contextual inferences are performed in sequence.

3 Need for contextual knowledge in parsing speech input

3.1 Difficulty of Speech Understanding

When compared to the understanding of text input, the added difficulty in the understanding of continuous speech input can be seen in:

- **phonemic segmentation:** The input stream can be segmented in many possible ways, and each hypothesis may be equally likely to succeed, even after applying semantic restriction checks.
- **added lexical ambiguity:** More than one lexical entry may be matched to a given phonemic segment, so the problem of lexical ambiguity is enlarged during continuous speech input understanding.
- **extra-grammaticality and incomplete words:** Due to noise, speaker variability, and other limitations of speech recognition devices, parser input is often incomplete and may be ungrammatical.

In response to these problems, the understanding of continuous speech input has been countered with mostly engineering improvements in recognition devices (including improved algorithms for probability measures, etc.). However, there have been a few recent efforts in trying to solve these problems from the NLP side. This includes work of Hayes, *et al*[1986], Tomita[1986], Poesio&Rullent[1987], and Saito&Tomita, in

frame restriction checks generated automatically by the Universal Parser/Compiler (Tomita&Carbonell[1987]).

²Lexical Functional Grammar (Kaplan&Bresnan[1982]).

which case-frame based semantic restriction checks are integrated into efficient parsing algorithms that score the word lattice and/or phoneme sequence. These efforts³, combined with the engineering improvements in speech recognition devices, effectively reduce some of the problems in the area of parsing moderately noisy and moderately incomplete sentences. However, issues of phonetic segmentation and lexical ambiguity are not solvable by enhancements in parsing techniques or recognition engineering, because there is an underlying problem of choosing the most appropriate hypothesis for grouping phonetic segments and choosing the correct lexical-sense from multiple hypotheses supplied by the voice recognition system (when each hypothesis passes the test of syntactic and local semantic (case-frame) constraints). This difficulty is aggravated when speech understanding is performed in real-time and the noise in the input is not easily controllable. We have seen cases where one input sentence is hypothesized into 30 to 50 syntactically well-formed and semantically (sentential level) acceptable candidates under noisy circumstances.

For example, testing the CMU-CMT speech understanding system, the Japanese input "atamagaitai" ("I have a headache") was spoken into a speech recognition system⁴ and accepted by the integrated⁵ parser with 57 ambiguous interpretations. Each of the ambiguous interpretations are semantically legitimate, meeting the local restrictions set forth by the case-frame instantiation restrictions. Below are some of the highly scored interpretations:

atamagaitai	(I have a headache.)
kazokuwaitai	((The) families want to stay.)
kazokuheitai	((My) family is soldier(s).)
kazokudeitai	(I want to stay as (a) family.)
koputoaitai	(I want to see (meet) (my) cup.)
asabanaisou	(Love (make love) (every) morning and night.)
askaraikou	(Go (come) (from) tomorrow morning.)
kazokuwa ikou	((The) families go.)
asamadeikou	(Go before morning, Come until morning.)
okosanaika	(Shall we wake (one) up?)
okosumaika	(Shall we not wake (one) up?)
kazokuheikou	((The) family is disappointed.)
kazokudeikou	(Go with the family.)
gohunasiou	(Love (make love) for five minutes.)
ugokumaika	(Shall I not move?)
atukunaika	(Is it not hot?)
dokoeikou	(Where shall we go?)
dokodeikou	(Where shall we come?)
koupumadeikou	(go to (the) cup.)

These are just some of the 57 disambiguations that were produced as acceptable readings by the speech understanding system given the input "atamagaitai" ("I have a headache"). As we can see from the examples, the sentences are perfectly acceptable, unless the system makes use of the discourse context

³And others such as Lee, *et al*[1987].

⁴Matsushita Research Institute's speech recognition hardware. The speech recognition system and the speech input enhanced LR parser is described in detail in Saito&Tomita. The experiment was conducted in an uncontrolled ordinary office environment.

⁵By 'integrated', we mean concurrent processing of syntax and semantics during parsing as opposed to some parsing methods where syntax and semantics are separately processed.

(doctor/patient communication) to further restrict the candidate set.

3.2 Need of contextual knowledge

Even with the semantic restrictions set forth by case-frame restriction checks, we suffer from the problem of ambiguities that are not possible with ordinary text inputs. This problem increases when the vocabulary of the speech understanding system enlarges and the variety of sentences that are accepted by the system expands. Although possible morphophonemic analyses of the speech input are ordered by the scores recorded during the speech recognition, the difference between candidate hypotheses as indicated by scoring is often within the tolerance of the system's error checking mechanism.

The CMU-CMT speech understanding system (without the contextual disambiguation mechanism we describe in this paper) attains about 85% accuracy on sentences in the doctor/patient dialogue domain (under a controlled, relatively noiseless environment). This is possible mainly because we restricted world knowledge (the semantic case-frame knowledge-base) to the doctor/patient dialogue domain. As a result, a sentence such as "asabanaisou" ((make) love (every) morning and night) was not accepted as candidate morphophonemic realization, simply because the knowledge-base was not large enough to accept such sentences. In a sense, this imposes an arbitrary contextual restriction on our system by restricting vocabulary to the limited domain. However, we are interested in expanding our vocabulary to cover sentences that are realistically possible in the actual use of a speech understanding system (such as aiding the hearing-impaired, translation for foreign language speakers, etc.), which inevitably includes vocabulary from a domain that is much larger than the target domain of the speech understanding system. As we have seen from our 57 well-formed and acceptable sentences (at the sentential level), once we enlarge the vocabulary (and world knowledge) to be a realistic size, we will suffer from the explosion of multiple ambiguities after the input is interpreted by the parser.

Local semantic restriction checks are not sufficient for disambiguating continuous speech input, since an interpretation can be totally legitimate semantically, but can mean something drastically different from what has been input into the speech recognition system (as well as being contextually inappropriate). The speech understanding system needs extra-sentential knowledge to choose an appropriate hypothesis for grouping phonetic segments and for selecting the appropriate word-sense of lexical entries. In other words, the need for contextual knowledge in speech understanding systems is even more urgent than in text input understanding systems; in a speech understanding system, the input can be interpreted in a way that is not possible in text input systems, and the input can still be acceptable to the local semantic restriction checks that integrated parsers perform within a sentence (such as slot-filler restriction checks of case-frame parsers).

Our belief that natural language understanding must be performed under concurrent processing of syntax, semantics and pragmatics (thematic and contextual) is effectively supported by the difficulty of understanding noisy continuous speech input without the integration of contextual and thematic knowledge.

4 Accessing contextual memory during parsing

First, we look briefly at our Φ GLR parser. Syntactic knowledge in the system is represented in an LFG formalism which is based on Kaplan&Bresnan[1982] using a notation similar to PATR-II (Shieber, *et al*[1983]). Below is an arbitrary example of an LFG (Japanese) syntactic representation⁶:

```
(<v> <--> (<v-IMP>) ((x0 = x1) ((x0 :mood) =
IMP)))
:
(<v-mizen2> <--> (<v-fsahen> @so) ((x0 = x1)))
(<v-renyol> <--> (<v-fsahen> @si) ((x0 = x1)))
:
(<v-IMP> <--> (<v-fsahen> @se @yo) ((x0 =
x1)))
(<v-fsahen> <--> ($sail)
((x0 root) = aisuru)
((x0 cat) = V)
((x0 subcat) = trans)))
($sail <--> (@a @i)) ; aisuru
```

In unification-based parsing, syntactic knowledge is used for a series of unifications which yield a sentential feature structure when the sentence is syntactically well-formed. Our speech (phoneme) parser, which is based on Tomita&Carbonell[1987]'s syntax/semantics parser, integrates syntactic unification with sentential (local) semantic restriction tests, using the case-frame-based syntax/semantics mapping rules. Below is an example of a rule from the doctor/patient dialogue domain⁷:

```
(f *HAVE-A-PAIN
(is-a (value *HAVE-A-SYMPOM))
(:symptom (sem *PAIN))
(:pain-spec (sem *PAIN-TYPE))
(:severity (sem *SEVERITY))
(:location (sem *BODY-PART)))
(j *HAVE-A-PAIN <==> (*OR* itai itamu)
(:symptom = (*PAIN))
(:severity <==> (advadjunct))
(:location <==> (obj))
(:pain-spec <==> (advadjunct))
(:freq <==> (advadjunct))
(:duration <==> (advadjunct))
)
```

As we can see, the mapping between semantic case-frame slot restrictions and syntactic feature structure paths are represented. The syntax/semantics parser utilizes these semantic restriction checks while performing unification. Also, note that these semantic restrictions are domain dependent, and therefore context independent. In our knowledge source, we have a multiple number of this type of mappings for lexical entries between syntax and semantics for each 'sense' of the words. Of course, in speech understanding, since a phonetic stream is segmented in multiple hypothetical ways, we will have an even greater number of concurrently active mapping-rules for each segment of a phonetic stream.

Now, we look at our integration of contextual (thematic) memory activity with this unification-based syntax/semantics parsing. In essence, we perform spreading activation in mem-

⁶See Mitamura, *et al*[1988] for details of the representational scheme.

⁷Consult Tomita&Carbonell, Tomita, *et al*[1987] for details of this representation.

ory every time a local semantic restriction test succeeds during the syntax/semantics unification. Our algorithm is as follows:
FOR each sentence in the speech input DO;

1. When unification of one feature-structure and another feature-structure succeeds (syntactically well-formed), and this unification accompanies the addition of one concept (semantic case-frame) to another concept as a part of the receiving concept's features (namely, succeeds in meeting case-frame slot filling restrictions), then:
2. Activate the concept that succeeded in the above unification and semantic test (receiving another concept as its feature⁸).
3. Activate the concepts that are abstractions of the activated concept in the memory-net.
4. If an activated concept is a thematic root concept, then send the thematic activation to the concepts that are thematic-children of the node.
5. If a thematically activated concept is a thematic root concept, then send the thematic activation to the thematic-children of the node.
6. When unifications build sentential case-frames, activate the sentential case-frame with the highest level of thematic-activation. Deactivate all other sentential case-frames and non-sentential case-frames. Perform upward activation as in 3 and thematic-activation (4, 5) for the chosen sentential case-frame.

END FOR;

To clarify, we are assuming a frame-based semantic-net as a representation of domain knowledge which is organized by inheritance links and also by relation (feature) links that are mapped with syntactic feature-attributes at some level of abstraction. We are also using links that group thematically related concepts. This is attained by having some nodes characterized as thematic root nodes packaging the thematic children nodes. This packaging can be thematic as well as episodic⁹.

As we stated in the description of the algorithm above, we have two kinds of activations: 1) unification triggered conceptual activations; and 2) thematic concept triggered thematic activations. Both are spreading activations but they are not so called 'dumb' spreading activations, because we do not spread activations everywhere. Instead, in the case of the first activation, we only activate upwards in the abstraction hierarchy; in the case of the second type, we only activate the thematic packaged nodes (and thematic packaged nodes of the packaged nodes). The thematic activation is analogous to DMTRANS (Tomabechi[1987])'s 'C-MARKER' marker passing, except that DMTRANS uses lexical activation of contextual mark-

⁸This reception of another concept as a specific feature of the concept is equivalent to 'concept refinement' that is central to parsers such as MOPTRANS (Lyinen[1984]), DMTRANS (Tomabechi[1987]) and DM-COMMAND (Tomabechi&Tomita[1988]).

⁹We do not distinguish episodic memory from thematic memory. Often memory representations tend to put emphasis on episodic memory (scriptal groupings); however, actual input may not accompany any scriptal episodic contents, instead, the input can be purely thematic. For example, input about configuration of a personal computer will have thematic grouping of concepts such as 'cpu', 'memory', 'key board'; but may not accompany any scriptal utterance. Thus, we treat both episodic and thematic grouping of concepts uniformly by categorizations under each thematic root nodes.

ers whereas we use unification-triggered activation of thematic root nodes as the source of thematic activations.

5 Discussion:

Context in a conceptual memory network can be represented as a grouping of concepts that are associated in a certain manner, i.e. an activation of one concept in memory triggers (or can potentially trigger) some other concepts in the network. To put it in another way, there is a relationship between concepts in which activation (recognition) of one concept reminds some other concept that it is related in a certain way. In our model, we have two types of activations, unification-based (local semantic) and thematic. The first type of activation represents the recognition of what is being said (can be multiple hypotheses) and the second type represents what is likely to be heard. The second type of activation is important because the context highlights some concepts during the understanding of input which is more than one sentence in length. Spreading activation through a network of concepts is our choice for performing such thematic activation. As claimed by the direct memory access literature (Riesbeck&Martin[1985]'s DMAP0, DMTRANS, etc.), such a scheme has an advantage of being able to perform memory-based inferences based upon knowledge existent in memory which is not possible by conventional build-and-store syntax/semantics parsers. We also believe that this memory-based activity (recognize-and-record understanding) needs to be integrated with the syntactic and local (domain-based) semantic analyses. We attain this by performing the spreading activation thematic recognition at each acceptance of LFG based unification with simultaneous acceptance of domain specific semantic tests. Our scheme is contrasted with recognize-and-record parsers¹⁰ such as DMAP0 and DMTRANS, where almost no syntactic analyses were performed (except linear ordering of concepts), and syntactically complex sentences were not recognized. Such schemes are not desirable, particularly with a speech understanding system, since without strong syntactic restrictions, the possible hypotheses of morphophonemic segmentation can grow exponentially with the increase in length and noise level of the continuous speech input. Norvig[1987] has a system which performs similar spreading activation inference for text input; however, in his system, syntax/semantic parsing and contextual inference modules are separate (performed sequentially); yet, we believe that these processings need to be integrated for the reason we stated above.

Our method of networking thematically related concepts in addition to an ordinary semantic network may resemble associative models that are researched by connectionists. However, we have not adopted connectionist associative architecture and back-propagation in our thematic conceptual clusters. Our spreading activations are guided and do not use weighted links. Since we are using an efficient generalized LR parser for syntactic analyses, combined with an unification-based information processing as a base for spreading activation memory activity, our model naturally solves problems of metonymy such as the example below (taken from Touretzky[1988]):

John cut an apple from the tree.

As Touretzky suggests, to correctly understand this, we need

¹⁰Also including recent efforts by Charniak&Santos[1987], Bookman[1987], and Berg[1987].

selectional restrictions created by combination of "cut" and "from the tree" and also the knowledge that apples are connected to trees by stems, etc.. This type of understanding, so far, is not possible under the connectionist paradigm. Also, under noisy continuous speech input, this sentence can also be hypothesized in multiply syntactic and local semantically acceptable ways and it is beyond the capacity of the current level of connectionist parsing. On the other hand, combination of unification-based approach with associative (thematic) memory handles this type of sentence naturally¹¹.

We believe our scheme of integration of unification-based syntax/semantics parsing with memory-based spreading activation recognition is equally viable for other types of unification formalisms such as HPSG (Pollard&Sag[1987]), GPSG (Gazdar, *et al*[1985]), and JPSG (Gunji[1986]). Since the method of unifying feature structures is shared among different unification-based grammar formalisms, our scheme guarantees that the memory-based activity is integrated at each unification that succeeded and passed the (domain/local) semantic-test for adding one concept as a part of another concept.

6 Conclusion

We have reported our scheme of integrating thematic (contextual) disambiguation with syntax/semantics unification-based parsing for understanding continuous speaker-independent noisy speech input. Our system represents the paradigm of integrating build-and-store type parsers (exemplified by our unification-based parser) and recognize-and-store memory-based activity during real-time processing. Our experimental results show that multiple hypotheses of morphophonemic and lexical segmentations and selections are effectively narrowed in our scheme. In most cases, our understander outputs a single semantic representation of the input speech for the same input that could be represented in over 50 possible ways that are syntactically well-formed and local semantically acceptable when the understander is run without the integrated contextual/thematic recognizer. Because our scheme of integrating memory-based activation for contextual recognition is based on feature-structure unifications and case-frame semantic knowledge representations, our paradigm is applicable regardless of the grammar formalism that is chosen (LFG, HPSG, JPSG, etc¹²) and therefore must be highly effective for understanding noisy continuous speech input when adopted to systems that utilize such formalisms as well.

Appendix: Implementation

As the speech recognition front-end to our speech understanding system, we have adopted a high-speed and speaker-independent speech recognition device built by Matsushita Research Institute (Morii, *et al*[1986]; Hiraoka, *et al*[1986]),

¹¹Local semantic restriction tests during the unification augmented by the contextual/thematic knowledge of concurrently activated associative memory attain the selectional restrictions based on syntax, semantics (case-frame restrictions) and pragmatics (contextual/thematic). This is not currently possible under the connectionist model especially due to the fact the connectionist model still lacks the complex compositionality and variable binding. In contrast, such a task is rather trivial in our scheme.

¹²Including categorial grammars utilizing unifications (Karttunen[1986], Zeevat, *et al*[1987], etc.).

which takes a Japanese speech utterance and produces a sequence of phonemes. The LFG phoneme parser (Φ GLR) is a generalized LR parser augmented by pseudo/full unification packages. The run-time grammar is precompiled for the LR parser for run-time efficiency. Semantic memory is represented using FRAMEKIT (Nyberg[1988]). Domain knowledge is coded as case-frames and syntax/semantic mappings are represented as mapping rules between case-frame slots and feature structure syntactic attributes. The parallelism of spreading activation is simulated using lazy evaluations in CommonLisp. The completed system runs on a 12Meg IBM-RT¹³ running CommonLisp. Currently, an implementation is underway using MULTILISP (Halsead[1985]), which is a parallel lisp developed at MIT for Concert multi-processors and is now implemented on Mach (Rashid, *et al*[1987]) at CMU.

References

- Berg, G. (1987) *A Parallel Natural Language Processing Architecture with Distributed Control*. In 'Proceedings of the CogSci-87'.
- Bookman, L.A. (1987) *A Microfeature Based Scheme for Modelling Semantics*. In 'Proceedings of the IJCAI-87'.
- Charniak, E. and Santos, E. (1987) *A Connectionist Context-Free Parser Which is not Context-Free, But Then It is Not Really Connectionist Either*. In 'Proceedings of the CogSci-87'.
- Gazdar, G., Pullum, G., and Sag, I. (1985) *Generalized Phrase Structure Grammar*. Harvard University Press.
- Gunji, T. (1986) *Japanese Phrase Structure Grammar*. Dordrecht: Reidel.
- Halstead, R. (1985) *Multilisp: A language for Concurrent Symbolic Computation*. In ACM Trans. on Prog. Languages and Systems.
- Hayes, P., Hauptmann, A., Carbonell, J. and Tomita M. (1986) *Parsing Spoken Language: A Semantic Case-frame Approach*. In 'Proceedings of Coling-86'.
- Hiraoka, S., Morii, S., Hoshimi, M. and Niyada, K. (1986) *Compact Isolated Word Recognition System for Large Vocabulary*. In 'Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing' (ICASSP86).
- Kaplan, R. and Bresnan, J. (1982) *Lexical Functional Grammar: A formal System for Grammatical Representation*. In 'The Mental Representation of Grammatical Relations', ed J. Bresnan. MIT Press.
- Karttunen, L. (1986) *Radical Lexicalism*. Report No. CSLI-86-68. Stanford: CSLI.
- Lee, L., Tseng, C., Chen, K.J., and Huang, J. (1987) *The Preliminary Results of A Mandarin Dictation Machine Based Upon Chinese Natural Language Analysis*. In 'Proceedings of the IJCAI-87'.
- Lytinen S. (1984) *The organization of knowledge in a multi-lingual, integrated parser*. Ph.D. thesis Yale University.
- Mitamura, T., Musha, H., Kcc, M. (1988) *The Generalized LR Parser/Compiler Version 8.1: User's Guide*, ed. Tomita, M.. CMU-CMT-88-MEMO. Carnegie Mellon University.
- Morii, S., Niyada, K., Fujii, S. and Hoshimi, M. (1986) *Large Vocabulary Speaker-independent Japanese Speech Recognition System*. In 'Proceedings of ICASSP85'.
- Nyberg, E. (1988) *The FrameKit User's Guide Version 2.0*. CMU-CMT-88-107. Carnegie Mellon University.
- Norvig, P. (1987) *Inference in Text Understanding*. In 'Proceedings of the AAAI-87'.
- Pollard, C. and Sag, A. (1987) *An Information-based Syntax and Semantics*. Vol 1. CSLI.
- Poesio, M. and Rullent, C. (1987) *Modified Caseframe Parsing for Speech Understanding Systems*. In 'Proceedings of the IJCAI-87'.
- Rashid, R., A. Tevanian, M. Younge, D. Youge, R. Baron, D. Black, W. Bolosky and J. Chew (1987) *Machine-Independent Virtual Memory Management for Paged Uniprocessor and Multiprocessor Architectures*. CMU-CS-87-140. Carnegie Mellon University.
- Riesbeck, C. and Martin, C. (1985) *Direct Memory Access Parsing*. Yale University Report 354.
- Saito, H. and Tomita, M. (1988) *Parsing Noisy Sentences*, In 'Proceedings of Coling-88'.
- Shieber, S., Uszkorcit, H., Robinson, J., and Tyson, M. (1983) *The Formalism and Implementation of PATR-II*. In Research on Interactive Acquisition and Use of Knowledge. Artificial Intelligence Center, SRI International.
- Tomabechi, H. (1987) *Direct Memory Access Translation*. In 'Proceedings of the IJCAI-87'.
- Tomabechi, H. and Tomita, M. (1988) *Application of the Direct Memory Access paradigm to natural language interfaces to knowledge-based systems* In 'Proceedings of COLING-88'.
- Tomita, M. (1986) *An Efficient Word Lattice Parsing Algorithm for Continuous Speech Recognition*. In 'Proceedings of ICASSP86'.
- Tomita, M. and Carbonell, J. (1987) *The Universal Parser Architecture for Knowledge-Based Machine Translation*. In 'Proceedings of the IJCAI-87'.
- Tomita, M., Kee, M., Mitamura, T. and Carbonell, J. (1987) *Linguistic and Domain Knowledge Sources for the Universal Parser Architecture*. In 'Terminology and Knowledge Engineering' Eds. H. Czap, and C. Galinski INDEKS Verlag.
- Touretzky, D (1988) *Beyond Associative Memory: Connectionists Must Search for Other Cognitive Primitives*. In 'Proceedings of the 1988 AAAI Spring Symposium Series. Parallel Models of Intelligence'.
- Zeevat, H., Klein E., Calder, J. (1987) *Unification Categorical Grammar*. In 'Categorical Grammar, Unification Grammar and Parsing'. Eds Haddock, Klein, and Morrill. Edinburgh Working Papers in Cognitive Science, Vol. 1. Edinburgh: Center for Cognitive Science, University of Edinburgh.

¹³A technical report by the authors from the CMU-CMT is forthcoming which contains the sample runs of the system on an IBM-RT.