

# Upper Modeling: organizing knowledge for natural language processing

John A. Bateman  
USC/Information Sciences Institute  
4676 Admiralty Way  
Marina del Rey, CA 90292-6695, U.S.A.  
*e-mail:* bateman@isi.edu

## Abstract

A general, reusable computational resource has been developed within the Penman text generation project for organizing domain knowledge appropriately for linguistic realization. This resource, called the **upper model**, provides a domain- and task-independent classification system that supports sophisticated natural language processing while significantly simplifying the interface between domain-specific knowledge and general linguistic resources. This paper presents the results of our experiences in designing and using the upper model in a variety of applications over the past 5 years. In particular, we present our conclusions concerning the appropriate organization of an upper model, its domain-independence, and the types of interrelationships that need to be supported between upper model and grammar and semantics.

## Introduction: interfacing with a text generation system

Consider the task of interfacing a domain-independent, reusable, general text generation system with a particular application domain, in order to allow that application to express system-internal information in one or more natural languages. Internal information needs to be related to strategies for expressing it. This could be done in a domain-specific way by coding how the application domain requires its information to appear. This is clearly problematic, however: it requires detailed knowledge on the part of the system builder both of how the generator controls its output forms and the kinds of information that the application domain contains. A more general solution to the interfacing problem is thus desirable.

We have found that the definition of a mapping between knowledge and its linguistic expression is facilitated if it is possible to classify any particular instances of facts, states of affairs, situations, etc. that occur in terms of a set of general objects and relations of specified types that behave systematically with respect to their possible linguistic realizations. This approach has been followed within the PENMAN

text generation system [Mann and Matthiessen, 1985; The Penman Project, 1989] where, over the past 5 years, we have been developing and using an extensive, *domain- and task-independent* organization of knowledge that supports natural language generation: this level of organization is called the **upper model** [Bateman *et al.*, 1990; Mann, 1985; Moore and Arens, 1985].

The majority of natural language processing systems currently planned or under development are now recognizing the necessity of some level of abstract 'semantic' organization similar to the upper model that classifies knowledge so that it may be more readily expressed linguistically.<sup>1</sup> However, they mostly suffer from either a lack of theoretical constraint concerning their internal contents and organization and the necessary mappings between them and surface realization, or a lack of abstraction which binds them too closely with linguistic form. It is important both that the contents of such a level of abstraction be motivated on good theoretical grounds and that the mapping between that level and linguistic form is specifiable.

Our extensive experiences with the implementation and use of a level of semantic organization of this kind within the PENMAN system now permit us to state some clear design criteria and a well-developed set of necessary functionalities.

## The Upper Model's Contribution to the Solution to the Interface Problem: Domain independence and reusability

The upper model decomposes the mapping problem by establishing a level of *linguistically motivated* knowledge organization specifically constructed as a response

<sup>1</sup>Including, for example: the Functional Sentence Structure of XTRA: [Allgayer *et al.*, 1989]; [Chen and Cha, 1988]; [Dahlgren *et al.*, 1989]; POLYGLOSS: [Emele *et al.*, 1990]; certain of the Domain and Text Structure Objects of SPOKESMAN: [Meteer, 1989]; TRANSLATOR: [Nirenberg *et al.*, 1987]; the Semantic Relations of EUROTRA-D: [Steiner *et al.*, 1987]; JANUS: [Weischedel, 1989]. Space naturally precludes detailed comparisons here: see [Bateman, 1990] for further discussion.

to the task of constraining linguistic realizations<sup>2</sup>; generally we refer to this level of organization as *meaning* rather than as *knowledge* in order to distinguish it from language-independent knowledge and to emphasize its tight connection with linguistic forms (cf. [Matthiessen, 1987:259-260]). While it may not be reasonable to insist that *application domains* organize their knowledge in terms that respect linguistic realizations — as this may not provide suitable organizations for, e.g., domain-internal reasoning — we have found that it is reasonable, indeed essential, that domain knowledge be so organized if it is also to support expression in natural language relying on general natural language processing capabilities.

The general types constructed within the upper model necessarily respect generalizations concerning how distinct semantic types can be realized. We then achieve the necessary link between particular domain knowledge and the upper model by having an application *classify* its knowledge organization in terms of the general semantic categories that the upper model provides. This does not require any expertise in grammar or in the mapping between upper model and grammar. An application needs only to concern itself with the 'meaning' of its own knowledge, and not with fine details of linguistic form. This classification functions solely as an interface between domain knowledge and upper model; it does not interfere with domain-internal organization. The text generation system is then responsible for realizing the semantic types of the level of meaning with appropriate grammatical forms.<sup>3</sup> Further, when this classification has been established for a given application, application concepts can be used freely in input specifications since their possibilities for linguistic realization are then known. This supports two significant functionalities:

- interfacing with a natural language system is radically simplified since much of the information specific to language processing is factored out of the input specifications required and into the relationship between upper model and linguistic resources;
- the need for domain-specific linguistic processing rules is greatly reduced since the upper model provides a domain-independent, general and reusable conceptual organization that may be used to classify all domain-specific knowledge when linguistic processing is to be performed.

<sup>2</sup>Although my discussion here is oriented towards text generation, our current research aims at fully bi-directional linguistic resources [Kasper, 1988; Kasper, 1989]; the mapping is therefore to be understood as a *bi-directional* mapping throughout.

<sup>3</sup>This is handled in the PENMAN system by the grammar's *inquiry semantics*, which has been described and illustrated extensively elsewhere (e.g., [Bateman, 1988; Mann, 1983; Matthiessen, 1988]).

An example of the simplification that use of the upper model offers for a text generation system interface language can be seen by contrasting the input specification required for a generator such as MUMBLE-86 [Meteer *et al.*, 1987] — which employs realization classes considerably less abstract than those provided by the upper model — with the input required for Penman.<sup>4</sup> Figure 1 shows corresponding inputs for the generation of the simple clause: *Fluffy is chasing little mice*. The appropriate classification of domain knowledge concepts such as *chase*, *cat*, *mouse*, and *little* in terms of the general semantic types of the upper model (in this case, *directed-action*, *object*, *object*, and *size* respectively — for definitions see: [Bateman *et al.*, 1990]) automatically provides information about syntactic realization that needs to be explicitly stated in the MUMBLE-86 input (e.g., *S-V-O\_two-explicit-args*, *np-common-noun*, *restrictive-modifier*, *adjective*). Thus, for example, the classification of a concept *mouse* as an *object* in the upper model is sufficient for the grammar to consider a realization such as, in MUMBLE-86 terms, a *general-np* with a particular *np-common-noun* and *accessories* of gender *neuter*. Similarly, the classification of *chase* as a *directed-action* opens up linguistic realization possibilities including clauses with a certain class of transitive verbs and characteristic possibilities for participants, corresponding nominalizations, etc. Such low-level syntactic information is redundant for the PENMAN input.

The further domain-independence of the upper model is shown in the following example of text generation control. Consider two rather different domains: a navy database of ships and an expert system for digital circuit diagnosis.<sup>5</sup> The navy data base contains information concerning ships, submarines, ports, geographical regions, etc. and the kinds of activities that ships, submarines, etc. can take part in. The digital circuit diagnosis expert system contains information about sub-components of digital circuits, the kinds of connections between those sub-components, their possible functions, etc. A typical sentence from each domain might be:

*circuit domain*: The faulty system is connected to the input

*navy domain*: The ship which was inoperative is sailing to Sasebo

The input specifications for both of these sentences are shown in Figure 2. These specifications freely intermix upper model roles and concepts (e.g., *domain*,

<sup>4</sup>Note that this is not intended to single out MUMBLE-86: the problem is quite general; cf. unification-based frameworks such as [McKeown and Paris, 1987], or the Lexical Functional Grammar (LFG)-based approach of [Momma and Dörre, 1987]. As mentioned above, the current developments within most such approaches are now considering extensions similar to that covered by the upper model.

<sup>5</sup>These are, in fact, two domains with which we have had experience generating texts using the upper model.

---

```

/general-clause
:head (CHASES/S-V-0_two-explicit-args
  (general-np
    :head (np-proper-name "Fluffy")
    :accessories (:number singular
      :gender masculine
      :person third
      :determiner-policy
        no-determiner))
    (general-np
      :head (np-common-noun "mouse")
      :accessories (:number plural
        :gender neuter
        :person third
        :determiner-policy
          initially-indefinite)
      :further-specifications
        ((:attachment-function
          restrictive-modifier
          :specification
            (predication-to-be *self*
              (adjective "little")))
          ))
      :accessories (:tense-modal present :progressive
        :unmarked) )
    Input to MUMBLE-86 for the clause:
      Fluffy is chasing little mice
    from: Meter, McDonald, Anderson, Forster, Gay,
      Huettner, and Sibun (1987)
  )

```

---

```

(e / chase
:actor (e / cat :name Fluffy)
:actee (n / mouse
  :size-ascription (s / little)
  :multiplicity-q multiple
  :singularity-q nonsingular)
:tense present-progressive)
  Corresponding input to PENMAN

```

Figure 1: Comparison of input requirements for MUMBLE-86 and PENMAN

*range, property-ascription*) and the respective domain roles and concepts (e.g., *system, faulty, input, destination, sail, ship, inoperative*). Both forms are rendered interpretable by the subordination of the domain concepts to the single generalized hierarchy of the upper model. This is illustrated graphically in Figure 3. Here we see the single hierarchy of the upper model being used to subordinate concepts from the two domains. The domain concept *system*, for example, is subordinated to the upper model concept *object*, domain concept *inoperative* to upper model concept *quality*, etc. By virtue of these subordinations, the grammar and semantics of the generator can interpret the input specifications in order to produce appropriate linguistic realizations: the upper model concept *object* licenses a particular set of realizations, as do the concepts *quality, material-process*, etc.

Our present upper model contains approximately 200

---

```

(v1 / connects
:domain (v2 / system
  :relations
    (v3 / property-ascription
      :domain v2
      :range (v4 / faulty)))
:range (v5 / input)
:tense present)

```

Input for digital circuit example sentence:  
*The faulty system is connected to the input*

---

```

(v1 / sail
:actor (v2 / ship
  :relations
    (v3 / property-ascription
      :domain v2
      :range (v4 / inoperative)
      :tense past)
  :destination (sasebo / port)
  :tense present-progressive)

```

Input for navy example sentence:  
*The ship which was inoperative is sailing to Sasebo*

Figure 2: Input specifications from navy and digital circuit domains

such categories, as motivated by the requirements of the grammar, and is organized as a structured inheritance lattice represented in the LOOM knowledge representation language [MacGregor and Bates, 1987]. Generally, the upper model represents the speaker's experience in terms of generalized linguistically-motivated 'ontological' categories. More specifically, the following information is required (with example categories drawn from the current PENMAN upper model):

- abstract specifications of process-type/relations and configurations of participants and circumstances (e.g., NONDIRECTED-ACTION, ADDRESSEE-ORIENTED-VERBAL-PROCESS, ACTOR, SENSER, RECIPIENT, SPATIO-TEMPORAL, CAUSAL-RELATION, GENERALIZED-MEANS),
- abstract specifications of object types, for, e.g., semantic selection restrictions (e.g., DECOMPOSABLE-OBJECT, ABSTRACTION, PERSON, SPATIAL-TEMPORAL),
- abstract specifications of quality types, and the types of entities which they may relate (e.g., BEHAVIORAL-QUALITY, SENSE-AND-MEASURE-QUALITY, STATUS-QUALITY),
- abstract specifications of combinations of events (e.g., DISJUNCTION, EXEMPLIFICATION, RESTATEMENT).

These are described in full in [Bateman *et al.*, 1990].

Appropriate linguistic realizations are not in a one-to-one correspondence with upper model concepts, however. The relationship needs to be rather more complex and so the question of justification of upper model concepts and organization becomes crucial.




Figure 3: Upper model organization reuse with differing domains

### Degree of Abstraction vs. Linguistic Responsibility

The general semantic types defined by a level of meaning such as the upper model need to be 'linguistically responsible', in that mappings between them and linguistic form may be constructed. In addition, to be usable by an application, they must also be sufficiently operationalizable so as to support *consistent* coding of application knowledge. Both of these requirements have tended to push the level of organization defined closer towards linguistic form. However, it is also crucial for this organization to be *sufficiently abstract*, i.e., removed from linguistic form, so that it is possible for an application to achieve its classification purely on grounds of meaning. It is thus inadequate to rely on form-oriented criteria for upper model construction because grammatical classifications are often non-isomorphic to semantic classifications: they therefore need to deviate from semantic organization in order to respect the syntactic criteria that define them. Reliance on details of linguistic realization also compromises the design aim that the applications should not be burdened with *grammatical* knowledge.<sup>6</sup>

<sup>6</sup>This is also resonant with the design aim in text generation that higher level processes — e.g., text planners — should not need direct access to low level information such as the grammar [Hovy *et al.*, 1988]. For descriptions of all these

Thus, the level of abstraction of an upper model must be sufficiently high that it generalizes across syntactic alternations, without being so high that the mapping between it and surface form is impossible to state. This tension between the requirements of abstractness and linguistic responsibility presents perhaps the major point of general theoretical difficulty and interest for future developments of upper model-like levels of meaning. Without a resolution, substantive progress that goes beyond revisions of what the PENMAN upper model already contains is unlikely to be achieved. *It is essential for constraints to be found for what an upper model should contain and how it should be organized so that an appropriate level of abstraction may be constructed.*

### Constraining the Organization of an Upper Model

Figure 4 sets several methodologies have been pursued for uncovering the organization and contents of a level of meaning such as an upper model, with examples of approaches that have adopted them, along the continuum of abstraction from linguistic form to abstract ontology. While the problem of being too bound to linguistic form has been mentioned, there are also severe problems with attempts to construct an upper model independent of form and motivated by other criteria, e.g., a logical theory of the organization of knowledge *per se*. Without a strong theoretical connection to the linguistic system the criteria for organizing an abstraction hierarchy remain ill-specified; there is very little guarantee that such systems will organize themselves in a way appropriate for interfacing well with the linguistic system.<sup>7</sup>

An alternative route is offered by the approaches in the middle of the continuum, i.e., those which abstract beyond linguistic form but which still maintain a commitment to language as a motivating force. This is further strengthened by the notion, now resurgent within current linguistics, that the organization of language informs us about the organization of 'knowledge' (e.g., [Halliday, 1978; Jackendoff, 1983; Langacker, 1987; Matthiessen, 1987; Talmy, 1987]): that is, *the relation between grammar and semantics/meaning is not arbitrary*. Detailed theories of grammar can then be expected to provide us with insights concerning the organization that is required for the level of meaning.

We have found that the range of meanings required to support one particular generalized functional region of

distinctions in detail, see the PENMAN documentation [The Penman Project, 1989].

<sup>7</sup>Furthermore, the experience of the JANUS project (e.g., [Weischedel, 1989]) has been that the cost of using a sufficiently rich logic to permit axiomatization of the complex phenomenon required is very high, motivating augmentation by an abstraction hierarchy very similar to that of the upper model and facing the same problem of definitional criteria.

nonlinguistic	reality	ontological — 'logical'	Weischedel (1989)
	knowledge	cognitive — 'psychological' situational — 'socio/psycho-logical'	Langacker (1987) Steiner (fc)
linguistic	meaning	grammatical semantics inquiry semantics clause-based	Halliday & Matthiessen (fc) PENMAN UPPER MODEL
		lexical semantics word senses word-based	Jackendoff (1983), LFG Mel'čuk & Žholkovskij (1970)
	form	syntactic realization classes syntax	Steiner <i>et al.</i> (1987) LFG

Figure 4: Sources of motivations for upper model development

the grammar developed within the PENMAN system provides a powerful set of organizing constraints concerning what an upper model should contain. It provides for the representation of 'conceptual' meanings at a high level of abstraction while still maintaining a mapping to linguistic form. This functional region corresponds with the Systemic Functional Linguistic notion of the *experiential metafunction* [Matthiessen, 1987], one of four generalized meaning types which are simultaneously and necessarily made whenever language is used. Any sentence must contain contributions to its function from all four 'metafunctions' — each metafunction providing a distinct type of constraint. The value of this factorization of distinct meaning types as far as the design of an upper model is concerned can best be seen by examining briefly what it *excludes* from consideration for inclusion within an upper model: i.e., all information that is controlled by the remaining three metafunctions should not be represented.

The *logical metafunction* is responsible for the construction of composite semantic entities using the resources of interdependency; it is manifested in grammar by dependency relationships such as those that hold between the head of a phrase and its dependents and the association of concepts to be expressed with particular heads in the sentence structure. The removal of this kind of information permits upper model specifications to be independent of grammatical constituents and grammatical dominance relations.

This relaxes, for example, the mapping between objects and processes at the upper model level and nominals and verbals at the grammatical level, enabling generalizations to be captured concerning the existence of verbal participants in nominalizations, and permits the largely textual variations shown in (1) and (2)<sup>8</sup> to be removed from the upper model coding.

- (1) It will probably rain tomorrow  
It is likely that it will rain tomorrow

<sup>8</sup>Example taken from [Meteer, 1988].

There is a high probability that it will rain tomorrow

- (2) independently  
in a way that is independent

No change in upper model representation or classification is required to represent these variations.

This can be seen more specifically by considering the following PENMAN input specification that uses only upper model terms:

```
((c0 / cause-effect
:domain discharge
:range breakdown)
(discharge / directed-action
:actee (electricity / substance))
(breakdown / nondirected-action
:actor (system / object)))
```

This states that there are two configurations of processes and participants — one classified as an upper model *directed-action*, the other as a *nondirected-action* — which are related by the upper model relationship *cause-effect*. Now, the assignment of concepts to differently 'ranked' heads in the grammar governs realization variants including the following:

*Electricity being discharged resulted in the system breaking down.*  
*Because electricity was discharged, the system broke down.*  
*Because of electricity being discharged the system broke down.*  
*... the breakdown of the system due to an electrical discharge...*  
*Electricity was discharged causing the system to break down.*  
*... an electrical discharge causing the breakdown of the system...*  
 etc.

Many such 'paraphrase' issues are currently of concern within the text generation community (e.g., [Meteer, 1988; Iordanskaja *et al.*, 1988; Bateman and Paris, 1989; Bateman, 1989]).

The *textual metafunction* is responsible for the creation and presentation of text in context, i.e., for estab-

lishing textual cohesion, thematic development, rhetorical organization, information salience, etc. The removal of this kind of information allows upper model specifications to be invariant with respect to their particular occasions of use in texts and the adoption of textually motivated perspectives, such as, e.g., theme/rheme selections, definiteness, anaphora, etc. Thus, with the same input specification as above, the following variations are supported by varying the textual constraints:

*It was the electricity being discharged that resulted in the system breaking down.*

*The discharge of electricity resulted in the system breaking down.*

*The system breaking down — the electricity being discharged did it!*

etc.

These textual variations are controlled during the construction of text (cf. [Matthiessen, 1987; Dale, 1989; Hovy and McCoy, 1989; Meteer, 1989; Bateman and Matthiessen, 1990]) and, again, are factored out of the upper model.

The *interpersonal metafunction* is responsible for the speaker's interaction with the listener, for the speech act type of an utterance, the force with which it is expressed, etc. Thus, again with the same input specification, the following variants are possible:

*Did electricity being discharged result in the system breaking down?*

*Electricity being discharged resulted surprisingly in the whole damn thing breaking down.*

*I rather suspect that electricity being discharged may have resulted in the system breaking down.*

etc.

The metafunctional factorization thus permits the upper model to specify experiential meanings that are invariant with respect to the linguistic alternations driven by the other metafunctions. That is, a specification in upper model terms is consistent with a set of linguistic realizations that may be regarded as 'experiential paraphrases': the specification expresses the 'semantic' content that is shared across those paraphrases and often provides just the level of linguistically decommitted representation required for nonlinguistically oriented applications. Generation of any unique surface realization is achieved by additionally respecting the functional constraints that the other metafunctions bring to bear; particular surface forms are only specifiable when a complete set of constraints from each of the four metafunctions are combined. The application of these constraints is directly represented in the PENMAN grammar, which provides for the perspicuous and modular integration of many disparate sources of information. The interdependencies between these constraints and their conditions of applicability are also directly represented in the grammar. This organization of the grammar allows us to construct a rather abstract upper model while still preserving the necessary mapping to linguistic form. The value of achieving the abstract specification of meaning supported by the upper model

is then that it permits a genuinely form-independent, but nevertheless *form-constraining*, 'conceptual' representation that can be used both as a statement of the semantic contents of an utterance and as an abstract specification of content for application domains that require linguistic output.

## Summary and Conclusions

A computational resource has been developed within the PENMAN text generation project that significantly simplifies control of a text generator. This resource, called the *upper model*, is a hierarchy of concepts that captures semantic distinctions necessary for generating natural language. Although similar levels of abstract semantic organization are now being sought in many natural language systems, they are often built anew for each project, are to an unnecessary extent domain or theory specific, are required to fulfill an ill-determined set of functionalities, and lack criteria for their design. This paper has presented the results of our experiences in designing and using the upper model in a variety of applications; in particular, it presented our conclusions concerning the appropriate source of constraints concerning the organization of an upper model. We have found that restricting the information contained in an upper model to *experiential* meaning has significantly improved our understanding of how a semantic hierarchy should be organized and how it needs to relate to the rest of the linguistic system. We strongly feel, therefore, that subsequently constructed semantic organizations should follow the guidelines set out by the metafunctional hypothesis; the factorization that it provides concerning what should, and should not, be represented in an 'abstract semantic knowledge' hierarchy supports functionalities well beyond those envisioned in current text generation/understanding systems.

## Acknowledgments

The upper model has been under development for several years, and many have and continue to contribute to it. The ideas I have reported on here would not have been possible without that development. Those responsible for the present form of the upper model include: William Mann, Christian Matthiessen, Robert Kasper, Richard Whitney, Johanna Moore, Eduard Hovy, Yigal Arens, and myself. Thanks also to Cécile Paris and Eduard Hovy for improving the paper's organization. Financial support was provided by AFOSR contract F49620-87-C-0005, and in part by DARPA contract MDA903-87-C-641. The opinions in this report are solely those of the author.

## References

[Allgayer *et al.*, 1989] Jürgen Allgayer, Karin Harbusch, Alfred Kobsa, Carola Reddig, Norbert Reithinger, and Dagmar Schmauks. Xtra: a natural-language access system to expert systems. *Inter-*

- national Journal of Man-Machine Communication*, 1989.
- [Bateman and Matthiessen, 1990] John A. Bateman and Christian M.I.M. Matthiessen. Uncovering the text base. In Hermann Bluhme and Hao Keqi, editors, *Selected Papers from the International Conference on Research in Text and Language, Xi'an Jiaotong University, Xi'an, P.R. China, 29-31 March 1989*. 1990.
- [Bateman and Paris, 1989] John A. Bateman and Cécile L. Paris. Phrasing a text in terms the user can understand. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Michigan, 1989. IJCAI-89.
- [Bateman et al., 1990] John A. Bateman, Robert T. Kasper, Johanna D. Moore, and Richard A. Whitney. A general organization of knowledge for natural language processing: the penman upper model. Technical report, USC/Information Sciences Institute, Marina del Rey, California, 1990.
- [Bateman, 1988] John A. Bateman. Aspects of clause politeness in Japanese: an extended inquiry semantics treatment. In *Proceedings of the 26th International Conference on Computational Linguistics*, pages 147-154, Buffalo, New York, 1988. Association for Computational Linguistics. Also available as ISI Reprint Series report RS-88-211, USC/Information Sciences Institute, Marina del Rey, California.
- [Bateman, 1989] John A. Bateman. Upper modelling for machine translation: a level of abstraction for preserving meaning. Technical Report EUROTRAD Working Papers, No. 12, Institut für Angewandte Informationsforschung, Saarbrücken, West Germany, 1989.
- [Bateman, 1990] John A. Bateman. Upper modeling: current states of theory and practise, 1990. PENMAN Development Note, USC/Information Sciences Institute.
- [Chen and Cha, 1988] Keh-Jiann Chen and Chuan-Shu Cha. The design of a conceptual structure and its relation to the parsing of chinese sentences. In *Proceedings of the 1988 International Conference on Computer Processing of Chinese and Oriental Languages*, Toronto, Canada, August 29 - September 1 1988.
- [Dahlgren et al., 1989] Kathleen Dahlgren, Joyce McDowell, and Edward P. Stabler. Knowledge representation for commonsense reasoning with text. *Computational Linguistics*, 15(3):149-170, 1989.
- [Dale, 1989] Robert Dale. Cooking up referring expressions. In *Proceedings of the Twenty-Seventh Annual Meeting of the Association for Computational Linguistics*, Vancouver, British Columbia, June 1989. Association for Computational Linguistics.
- [Emele et al., 1990] Martin Emele, Ulrich Heid, Walter Kehl, Stefan Momma, and Rémi Zajac. Organizing linguistic knowledge for multilingual generation. In *COLING-90*, 1990. Project Polygloss Paper, University of Stuttgart, West Germany.
- [Halliday and Matthiessen, forthcoming] Michael A.K. Halliday and Christian M.I.M. Matthiessen. The Bloomington Lattice. Technical Report in preparation, University of Sydney, Linguistics Department, Sydney, Australia, forthcoming.
- [Halliday, 1978] Michael A. K. Halliday. *Language as social semiotic*. Edward Arnold, London, 1978.
- [Hovy and McCoy, 1989] Eduard H. Hovy and Kathy F. McCoy. Focusing your RST: A step towards generating coherent multisentential text. In *Proceedings of the 11th. Annual Conference of the Cognitive Science Society*, pages p667-674, University of Michigan, Ann Arbor, Michigan, August 16-19 1989. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- [Hovy et al., 1988] Eduard H. Hovy, Douglas Appelt, and David D. McDonald. Workshop on text planning and natural language generation, August 1988. Sponsored by American Association for Artificial Intelligence.
- [Iordanskaja et al., 1988] Lidija Iordanskaja, Richard Kittredge, and Polguère Alain. Lexical selection and paraphrase in a meaning-text generation model, July 1988. Presented at the Fourth International Workshop on Natural Language Generation. Also appears in selected papers from the workshop: Paris, Swartout and Mann (eds.)(1990)(*op. cit.*).
- [Jackendoff, 1983] Ray Jackendoff. *Semantics and Cognition*. MIT Press, Cambridge, MA, 1983.
- [Kasper, 1988] Robert T. Kasper. An Experimental Parser for Systemic Grammars. In *Proceedings of the 12th International Conference on Computational Linguistics, August 1988*, Budapest, Hungary, 1988. Association for Computational Linguistics. Also available as Information Sciences Institute Technical Report No. ISI/RS-88-212, Marina del Rey, CA.
- [Kasper, 1989] Robert T. Kasper. Unification and classification: an experiment in information-based parsing. In *Proceedings of the International Workshop on Parsing Technologies*, pages 1-7, 1989. 28-31 August, 1989, Carnegie-Mellon University, Pittsburgh, Pennsylvania.
- [Langacker, 1987] Ronald W. Langacker. *Foundations in Cognitive Grammar*. Stanford University Press, Stanford, California, 1987.
- [MacGregor and Bates, 1987] Robert MacGregor and Raymond Bates. The LOOM knowledge representation language. In *Proceedings of the Knowledge-Based Systems Workshop*, 1987. Held in St. Louis,

- Missouri, April 21-23, 1987. Also available as ISI reprint series report, RS-87-188, USC/Information Sciences Institute, Marina del Rey, CA.
- [Mann and Matthiessen, 1985] William C. Mann and Christian M.I.M. Matthiessen. Demonstration of the nigel text generation computer program. In J. Benson and W. Greaves, editors, *Systemic Perspectives on Discourse, Volume 1*. Ablex, Norwood, New Jersey, 1985.
- [Mann, 1983] William C. Mann. The anatomy of a systemic choice. *Discourse Processes*, 1983. Also available as USC/Information Sciences Institute, Research Report ISI/RR-82-104, 1982.
- [Mann, 1985] William C. Mann. Janus abstraction structure – draft 1, 1985. An informal project technical memo of the Janus project at ISI.
- [Matthiessen, 1987] Christian M.I.M. Matthiessen. Notes on the organization of the environment of a text generation grammar. In G. Kempen, editor, *Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics*. Kluwer Academic Publishers, Boston/Dordrecht, 1987. Paper presented at the Third International Workshop on Natural Language Generation, August 1986, Nijmegen, The Netherlands.
- [Matthiessen, 1988] Christian M.I.M. Matthiessen. Semantics for a systemic grammar: the chooser and inquiry framework. In James Benson, Michael Cummings, and William Greaves, editors, *Linguistics in a Systemic Perspective*. Benjamins, Amsterdam, 1988. Also available as USC/Information Sciences Institute, Reprint Series Report ISI/RS-87-189, Marina del Rey, CA.
- [McKeown and Paris, 1987] Kathleen R. McKeown and Cécile L. Paris. Functional unification grammar revisited. In *Proceedings of the 25th Annual Meeting of the ACL*, Palo Alto, California, 1987. Association of Computational Linguistics.
- [Mel'čuk and Žholkovskij, 1970] A. Mel'čuk, Igor and A.K. Žholkovskij. Towards a functioning "meaning-text" model of language. *Linguistics*, 57:10-47, 1970.
- [Meteer et al., 1987] Marie W. Meteer, David D. McDonald, S.D. Anderson, D. Forster, L.S. Gay, A.K. Huettner, and P. Sibun. MUMBLE-86: Design and implementation. Technical Report 87-87, COINS, University of Massachusetts, 1987.
- [Meteer, 1988] Marie W. Meteer. Defining a vocabulary for text planning, August 1988. Presented at the AAAI-88 Workshop on Text Planning and Realization, organized by Eduard H. Hovy, Doug Appelt, David McDonald and Sheryl Young.
- [Meteer, 1989] W Meteer, Marie. The SPOKESMAN natural language generation system. Technical Report BBN Report No. 7090, BBN Systems and Technologies Corporation, Cambridge, MA., 1989.
- [Momma and Dörre, 1987] Stefan Momma and Jochen Dörre. Generation from f-structures. In Ewan Klein and Johann Van Bentham, editors, *Categories, Polymorphism and Unification*. Cognitive Science Centre, University of Edinburgh, Edinburgh, Scotland, 1987.
- [Moore and Arens, 1985] Johanna D. Moore and Yigal Arens. A hierarchy for entities, 1985. USC/Information Sciences Institute, Internal Draft.
- [Nirenberg et al., 1987] Sergei Nirenberg, V. Raskin, and A. Tucker. The structure of interlingua in TRANSLATOR. In Sergei Nirenberg, editor, *Machine Translation: Theoretical and Methodological Issues*. Cambridge University Press, Cambridge, 1987.
- [Paris et al., 1990] Cécile L. Paris, William R. Swartout, and William C. Mann, editors. *Natural Language Generation in Artificial Intelligence and Computational Linguistics*. Kluwer Academic Publishers, 1990.
- [Steiner et al., 1987] Erich H. Steiner, Ursula Eckert, Birgit Weck, and Jutta Winter. The development of the EUROTRA-D system of semantic relations. Technical Report Eurotra-D Working Papers, No. 2, Institut der angewandten Informationsforschung, Universität des Saarlandes, Saarbrücken, West Germany, 1987.
- [Steiner, 1990] Erich H. Steiner. *A model of goal-directed-action as a structuring principle for the context of situation in systemic linguistics*. Mouton and de Gruyter, Berlin, 1990.
- [Talmy, 1987] Leonard Talmy. The relation of grammar to cognition. In B. Rudzka-Ostyn, editor, *Topics in Cognitive Linguistics*. John Benjamins, 1987.
- [The Penman Project, 1989] The Penman Project. The PENMAN documentation: User guide, primer, reference manual, and nigel manual. Technical report, USC/Information Sciences Institute, Marina del Rey, CA, 1989.
- [Weischedel, 1989] Ralph M. Weischedel. A hybrid approach to representation in the janus natural language processor. In *27th Annual Meeting of the Association for Computational Linguistics*, pages 193-202, Vancouver, British Columbia, 1989.