

Integrating Natural Language Components into Graphical Discourse

Stephan Dilley, John Bateman, Ulrich Thiel, Anne Tissen

GMD, Integrated Publication and Information Systems Institute (IPSI)
Dolivostraße 15, D-6100 Darmstadt, FRG
e-mail: { bateman,dilley,thiel,tissen }@darmstadt.gmd.de

Abstract

In our current research into the design of cognitively well-motivated interfaces relying primarily on the display of graphical information, we have observed that graphical information alone does not provide sufficient support to users — particularly when situations arise that do not simply conform to the users' expectations. This can occur due to too much information being requested, too little, information of the wrong kind, etc. To solve this problem, we are working towards the integration of natural language generation to augment the *interaction* functionalities of the interface. This is intended to support the generation of flexible natural language utterances which pinpoint possible problems with a user's request and which further go on to outline the user's most sensible courses of action away from the problem. In this paper, we describe our first prototype, where we combine the graphical and interaction planning capabilities of our graphical information system SIC! with the text generation capabilities of the Penman system. We illustrate the need for such a combined system, and also give examples of how a general natural language facility beneficially augments the user's ability to navigate a knowledge base graphically.

1 Introduction

Natural language interfaces to information systems allow the user to converse with the system in a natural way without being restricted by the syntax of a formal query language. However, users often have difficulties in stating their information need precisely (cf. Brooks et al., 1986) and so a user interface design emphasizing multimedia aspects of information systems may be more appropriate. An object oriented presentation meets this requirement because users are allowed to investigate information items directly. Thus, the distance between the users' intentions and the objects they involve is reduced to a minimum (cf. Hutchins et al., 1986). However, visual presentations of abstract information items have to be designed carefully to support the user in perceiving them correctly. Often the inherent limitations of available presentation forms, e.g. the number of displayable items, are an impediment to an adequate visualization. These difficulties which are well known in visu-

al languages can be attacked by combining the presentation of graphical objects to the user with cooperative system responses in natural languages.

Our approach to human-computer interaction provides the basis for an integration of different interaction styles, in our case natural language and graphics, in a *multi-modal information system*. The choice of an appropriate combination of modes is essential for a successful interface design. As there is not yet a complete theory of multi-modal interaction, this question has to be answered experimentally. Therefore, we propose a modular architecture for multi-modal interfaces composed of *interaction tools*. In this paper, we discuss the addition of natural language generation to the interaction tools provided. This capability is embodied in the prototype IGiNG (Integrating Graphics and Natural Language Generation) which combines

- the graphical information system SIC! (System for Information on Conferences) comprising a *dialog manager* which is able to track the user's goals and, consequently, to plan system reactions as cooperative responses (Tissen, 1991), a *database access module* providing relevant data (Kracker, 1991) and a *presentation manager* generating visualizations of the retrieved information items (Kerner and Thiel, 1991).
- the text generation component KOMET/Penman (Bateman et al., 1991; Penman, 1989), comprising extensive grammatical, semantic, and text organizational linguistic resources.

The paper is organized as follows: first, we describe related work and how our approach goes beyond this, second, we introduce the basic components of our system and explain how they can be appropriately interfaced. Then, we go on to discuss the benefits that including a natural language component can bring for augmenting the coherence and functionality of : user's interaction with the system. We offer some examples of helpful responses that the system can make which would be difficult, if not impossible, to represent graphically in a non ad hoc fashion. Finally, we describe some possible future directions.

2 Related Work

One body of related work has primarily intended to coordinate different modes of expression within a framework of *natural communication* (cf. Hayes, 1987; Neal and Shapiro, 1988; Cohen et al., 1989, Feiner and McKeown, 1990; Bandyopadhyay, 1990; Stock, 1991; Wahlster et al., 1991). The principle effort here is to ascertain factors that can *motivate* the distribution of information across different modes (e.g. Arens and Hovy, 1990). A further body of related work moves towards problems of *interaction* by exploring the potential of the combination of natural language and *deictic gestures* (cf. Allgayer et al., 1989; Moore and Swartout, 1990). In a similar vein, approaches to flexible graphical interaction based on the *conversational metaphor* (cf. Reichman, 1986, 1989; Thiel, 1990) treat user inputs such as mouse clicks, menu selections, etc. not as invocations of methods that can be executed without regarding the dialog context, but instead as dialog acts expressing a discourse goal of the user. The direct manipulation of an object then becomes itself a part of the dialog of the user with the system, meaning that the system can respond in a more flexible way by taking into account the illocutionary and semantic aspects of the user's input. Related work from generation includes the correction of misconceptions in the work of McCoy (1986) and the explicit representation of information about a system's own knowledge and planning activities that is found in the Explainable Expert Systems system of Swartout and Smoliar (1987). None of this work, however, addresses the problems of meta-dialog concerning graphically supported interaction.

In our approach we bring the kind of natural language capabilities required by the first body of related work (i.e., graphical and natural language information functioning together) to bear on the kinds of problems that arise in the second body of related work when the direct manipulation of objects by a user creates goals that the system cannot fulfil. Here we must not only respond to the user's attempt to manipulate an object or the user's deictic gesture as a dialog act, but also be able to en-

gage in a meta-interaction to debug that act if it creates a problematic situation. We show that natural language possesses properties that make it preferable over the graphical mode of expression for such meta-interaction and hence natural language generation needs to be supported even in graphics-oriented interfaces.

3 The SIC!-System

We demonstrate the combination of graphics and natural language output in the context of the SIC! system. SIC! is implemented using HyperNews (HyperNews, 1989), a hypermedia-like user-interface management system which can be controlled from a LISP-client. SIC! offers information on conferences; its domain consists of abstract information on conferences including workshops, tutorials, persons, institutes, and conference topics. A user who wants to obtain information from SIC! poses a query whereafter he can inspect the data retrieved. This is done by selecting a presentation form. In SIC! we use several cognitively motivated presentation forms (Kerner and Thiel; 1991).

One of these forms is the ring presentation form (cf. Figure 1) which surveys a structure of given data. All items are positioned on a virtual ring structure, the relations between the concepts being presented as single lines. In our example the ring presentation form contains the categories *workshop* and *topic*. The concepts of the workshops are on the left side of the figure, those of the topics on the right. However, this presentation form is clearly limited with respect to the quantity of data that can be presented simultaneously.

Consider, for example, the following situation: A user asks for a subset of the IJCAI89 workshops and their related topics. SIC! retrieves three workshops and several topics. If the user wants to get an overview of the data's structure, then for this goal the ring is the most adequate presentation form available in SIC! But this causes a problem when the ring form cannot display as many data as SIC! retrieved. There are several ways of solving this problem, depending on the user's main

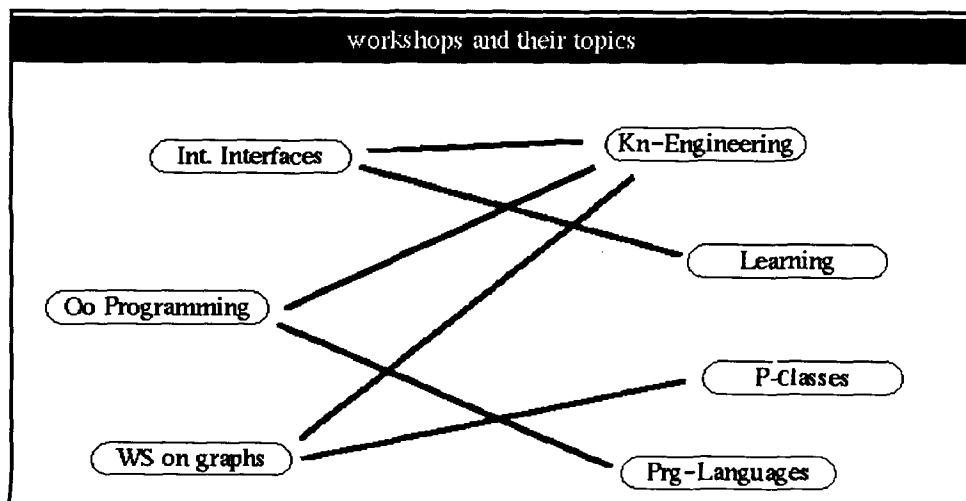


Figure 1:
Ring presentation form

point of interest. If he wants to see all the data he might select a presentation form that can show unlimited amounts of information, e.g. the table presentation form (cf. Figure 2). If he

Workshops and their topics:	
Workshops	Topics
OO Programming	Kn-Engineering, Prg-Languages,
Int. Interfaces	Learning, Kn-Engineering,
WS on graphs	Kn-Engineering, P-Classes,

Figure 2 : Table presentation form

is interested more in the structure of the data, he might change his initial query so that less data are retrieved. But, importantly, the user cannot be aware of the details of the current situation of retrieval and its implications simply from the graphical information displayed. Hence the system has to inform him of the situation and offer possible alternatives. This needs to be done in a way that enables users to grasp the situation and to choose the appropriate alternatives for their purpose. To achieve this, the system has to create a coherent informative act that is concise and yet unambiguous (in context), giving all the information necessary for the user to determine his future actions.

Based on work by Feiner and McKeown (1990) on the coordination of text and graphics in explanation generation and by Lombardi (1989), who examined the assignment of information to media, we assume that text is the appropriate medium for informative acts in meta-dialogs, since a well-constructed text is not only concise and easy to understand, but also guarantees the necessary flexibility to meet any situation that arises. Graphical 'acts' cannot be constructed compositionally to express possibly unforeseen complex circumstances: novel graphics must first be learnt by users — a situation avoided by the generation of situationally appropriate natural language. Thus, generating natural language text, particularly text involving controlled and appropriate deployment of text-forming resources such as rhetorical relations, enhances the total coherence of the user's dialog with the system. Our hypothesis is that the user's understanding of the situation and its implications is increased by the natural language output, which becomes an intermediary between the various possibilities for information presentation.

4 KOMET/Penman Text Generation System

We are using the KOMET/Penman¹ system (Mann and Matthiessen, 1983) for generating the natural language output our system requires. KOMET/Penman is a domain-independent text generation system based on systemic-functional grammar (Halliday, 1978). It consists of extensive grammars of English and German (Matthiessen, 1990; Teich, 1991), a linguistically motivated ontology, called the Upper Model (Bateman, Kasper, Moore and Whitney, 1989), a semantic interface that relates the categories of the conceptual ontology with their possible grammatical expressions in English and German (Matthiessen, 1990), and a basic lexicon containing English and German closed-class items and default lexical realizations for the concepts in the Upper Model ontology. The definition of the lexical items includes morphological information and sets of lexical features that determine the grammatical contexts in which items are to be selected.

The Upper Model is the component of the system that is primarily responsible for mediating between the knowledge specific to any given domain and the general lexical and grammatical expressions that are provided by a language. Because it is possible to state how any particular Upper Model concept is to be realized, subordinating domain concepts to particular Upper Model concepts causes those domain concepts to inherit appropriate forms of expression. For example, concepts from the *object*-class are usually realized as nominal phrases, while concepts from the *process*-class (e.g., mental-process, verbal-process, action-process, relation-process) are often realized by clauses². The relationship between Upper Model and domain model is diagrammed in the context of its application for SIC! in Figure 3.

Input to the KOMET/Penman text generation system is given in terms of the Sentence Plan Language (Kasper, 1989), of which we will see examples below. An SPL expression defines the semantic content of a sentence to be generated; it consists of a set of typed variables and relations defined between those variables. Both the types and the possible relations are defined either by the Upper Model directly or by concepts or relations in the domain model that have been subordinated to the Upper Model. In addition to this information, SPL expressions may also contain direct statements in terms of the grammar's semantic interface — in practical applications these latter are often abbreviated by use of macros (e.g. :tense present) or are defaulted.

¹ The original Penman system was developed at the Information Science Institute of the University of Southern California; the KOMET system of GMD/IPSI builds on this, working towards multilinguality and enhanced text planning capabilities.

² But not always: the existence of, for example, nominalizations motivates the maintenance of two distinct levels of representation.

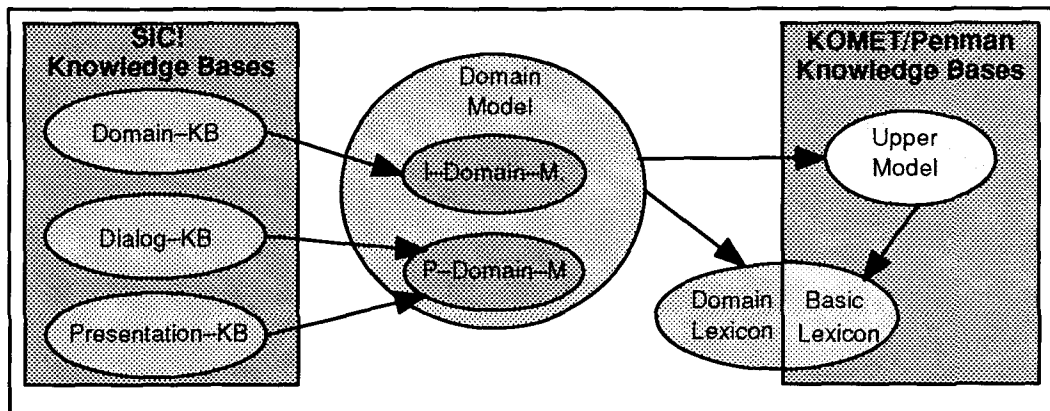


Figure 3 : Interfacing SIC! with KOMET/Penman on the Knowledge Level

5 Interfacing SIC! with the Text Generation System

To interface SIC! with KOMET/Penman we have to provide several types of knowledge (cf. Figure 3)

- A domain model, which is a taxonomy of knowledge specific to our application-domain. We split the domain into two parts: an Information-Domain (I-Domain), which contains concepts related to the information that is shown by SIC!, e.g. workshops and topics (cf. Figure 1), and a Presentation-Domain (P-Domain), which contains concepts related to the way this information is presented by SIC!, e.g. ring, table. By splitting the domain model we increase the adaptability in case of changes in the underlying application domain, e.g. replacing the conference knowledge base with a knowledge base on research projects. Every concept in the domain model has to be linked to some Upper Model concept from which it inherits attributes which enable KOMET/Penman to express the concept in a way that is grammatically correct. The I-Domain concepts can be generated automatically from the underlying SIC! knowledge bases (cf. Figure 3). Concepts can also be associated to lexical items.
- A domain lexicon, containing the definitions of lexical items of all the words that may appear in the application domain.

6 Creating the Natural Language Output

6.1 Planning Sentences

As stated in our example above, we want to produce text that, in this case, informs the user that not all the information that was requested can be shown because the current presentation-form's capacity is limited. Furthermore, we need to offer possible actions which solve this problem. In Figure 4, we show the semantic input to KOMET/Penman, expressed in SPL, that would cause KOMET/Penman to generate the first sentence

```
(a / ascription
 :domain (c / capacity
          :owned-by (p / presentation-form))
 :range (e / exceeded))
```

Fig. 4 : SPL-Plan for

"The presentation-form's capacity is exceeded."

that we require: i.e., *"The presentation-form's capacity is exceeded."* One type of abstract concept that the system requires is the status of a particular entity that may be displayed or used. Possible statuses are, for presentation-forms, *exceeded*, *incomplete*. These status concepts can then be attributed to objects by means of the Upper Model relation *ascription*, which has roles *'domain'* and *'range'*. They represent the concepts which are related, in our example the presentation-form's capacity and exceeded. In general, *'domain'* contains the essential concept of the relation while *'range'* contains additional information. The P-Domain concept *capacity* has been modeled as an object.

6.2 Using Rhetorical Relations

Figure 5 shows a more complex SPL-plan which demonstrates some of the more advanced possibilities given by KOMET/Penman. The most interesting aspect in this plan is the use of rhetorical relations based on Rhetorical-Structure-Theory (RST).

RST is a theory of the organization of natural language texts (Mann and Thompson, 1987). Mann and Thompson studied a wide variety of English texts and observed that there are approximately 25 relations that usually occur between coherent portions of English text. An RST relation consists of two parts, a *nucleus* and a *satellite*. The nucleus is that part that is most essential to the speaker's purpose, while the satellite contains additional information. The satellite is more easily replaced than the nucleus because of the nucleus' central role in the thematical progression of the discourse. Even though there are some critics questioning the use of rhetorical relations in discourse structure theory (Grosz and Sidner, 1986) we use

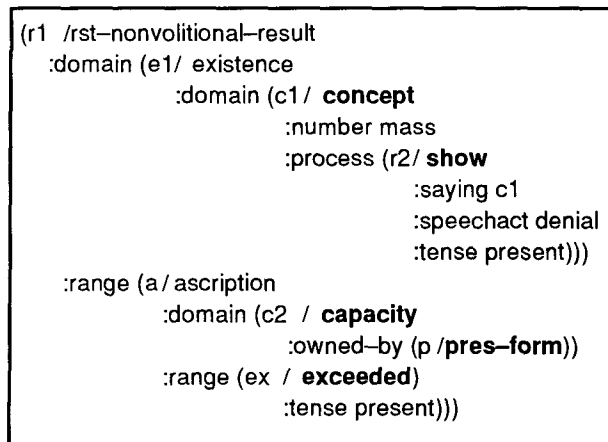


Figure 5 : SPL-Plan for "There are concepts that are not shown, because the presentation-form's capacity is exceeded".

RST relations because they proved to be quite useful when we link portions of information. In KOMET/Penman, RST-relations are treated the same way as other relations, e.g. *ascription* which we used in the plan shown in Figure 4 .

The SPL-plan shown in Figure 5 combines two relations: the *ascription*-relation, which we used in the SPL-plan in Figure 4 , and the *existence*-relation. *Existence* is a so called *one-place-relation*, because it contains only a :domain-role but no :range. It is usually realized as "There is ...", where :domain defines what exists. We link these two relations via an RST-relation called *rst-nonvolitional-result*. This RST-relation implies that the nucleus, which is defined in our :domain-role is a result of the satellite, defined in :range. One possible output is <domain>, because <range>, in our case "<There are concepts ...>, because <... capacity is exceeded>". Because what is defined in the :domain ("There are concepts that are not shown") is not volitional, we use *rst-nonvolitional-result* instead of *rst-volitional-result*. The fact that there is data that is not shown by the current presentation-form is essential to our informational purpose. Therefore this fact becomes the nucleus (represented by :domain) of our plan.

6.3 Supporting Coherence

In his work on coherence in multi-modal discourse, Bandyopadhyay (Bandyopadhyay, 1990) states that there are three levels of coherence: *syntactic coherence*, *semantic coherence* and *pragmatic coherence*. *Syntactic coherence* deals with the immediate connectivity among adjacent segments (in texts this is often called *text cohesion*). *Semantic coherence* ensures the wellformed thematic organization of a discourse. Discourse segments are connected by semantic ties (Hobbs, 1983). Bandyopadhyay defines a discourse to be *pragmatically coherent* if it is compatible with the addressees' interpretative ability. In our system syntactic coherence is enhanced by the way we present the natural language output in our graphical environment. Semantic coherence is supported by the use of

RST-relations which ensure the connectivity between our text segments. Pragmatic coherence is supported by the mere fact that we are using text as a medium for meta-dialogs, as these are difficult to understand on a graphical level.

7 Controlling Multimodal Discourse

The dialog manager is one of the main components of our interface system (cf. Figure 6). It chooses interaction modes (graphic or text) and controls the navigation or exploration in the information space.

In order to prevent the user from 'being lost in hyperspace', we guide the user by case-based dialog plans (Tissen, 1991). In a case-based planning system a new plan will be generated by retrieving the plan which is most appropriate to the user's goals and adapting it dynamically during the ongoing dialog. Two types of adaptations can be distinguished: first, system-driven modifications using domain dependent background knowledge, and second, *corrections of misconceptions*, handled interactively in meta-dialogs with the user.

The dialog manager detects misconceptions, i.e. situations in which an intended goal cannot be realized, e.g. *more items were retrieved than can be displayed in the current presentation form*. The corrector operates on knowledge bases of misconceptions and correction rules, e.g. "if there is a misconception like 'ring presentation: not all requested data can be presented in the ring' and there is no automatic plan modification possible then start a meta-dialog, which informs the user about the situation and offers alternatives." Because meta-dialogs will be handled in text mode, the dialog manager requests the SPL creator to produce SPL plans. Therefore, the dialog manager informs the SPL creator on the current misconception and possible alternatives the user has to choose from to resolve the situation. Then, the SPL creator produces the appropriate SPL plans by combining information on the misconception and possible alternatives with elements from the SPL library. The SPL plans are transformed into natural language text by the KOMET/Penman system. The resulting text is returned to the dialog manager which presents it to the user.

8 Controlling utterance selection

The IGiNG system intends to produce user adapted natural language output. It is an object oriented system consisting of several object classes (cf. Figure 7)

When IGiNG is requested to produce an utterance it calls the utterance's *express* method, which first builds a list of plan-objects starting from the initial plan-object given by the utterance object. Then, it is determined whether complex or short statements are desired. This information is kept in a user-stereotype and determines in which direction the list of plan-objects is to be traversed. Now IGiNG tests each plan-object's *select* conditions. If *all* conditions are satisfied, the plan-object is selected otherwise IGiNG tries the succeeding plan-object. Finally the plan defined by the plan-object is passed to the KOMET/Penman system which generates the utterance.

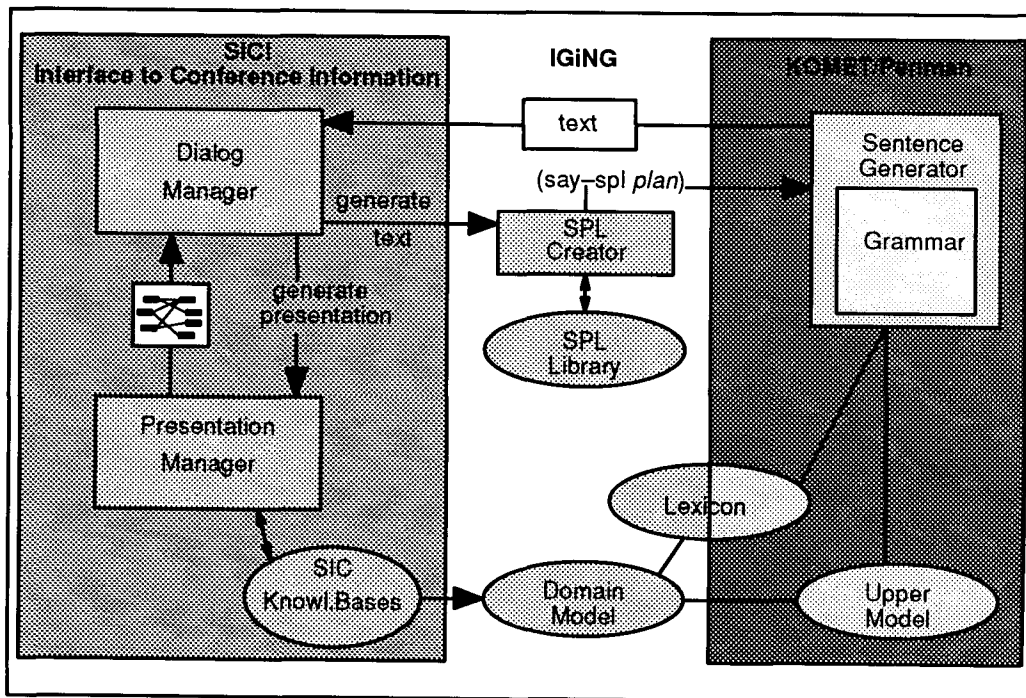


Figure 6 : Integrating SIC! and KOMET/Penman

Example: Let us consider the IGiNG-objects given in Figure 7. IGiNG is requested to express *mod-1*. It builds a list of possible plan-objects, which is (*plan-1 plan-2*). As concise statements are desired, *plan-1* is tested first. Because *con-1* is not satisfied (it demands that *user-level* is low while the current user level is advanced) *plan-1* is rejected. Next, *plan-2* is tested. As all conditions are satisfied, the plan given by *plan-2* is generated.

Adapting to new situations

When a new situations is to be included, the following steps have to be performed:

A new proposition-instance has to be defined as the successor of an existing proposition-class. If the new proposition is not a part of any of the existing proposition-classes, a new proposition-class should be defined first.

For any of the possible utterances a parameterizable partial sentence plan has to be written, which is stored in a plan-object, together with a reference to the plan-object's select-

condition-object. Of course, it is possible to use existing plan-objects, if they are suitable for the intended purpose.

Finally the plan-objects have to be linked to select-conditions. As these are domain independent, preexisting select-conditions can be reused.

These are all the steps necessary for defining new propositions. The new objects inherit from their ancestors all the functionality which is necessary for selection and expression.

9 Future Work

We are going to extend the use of natural language output to other situations which seem to be suited for textual information rather than graphics. We are working on a closer examination of what information has to be included in the user-stereotype and how this knowledge can be obtained. We intend to integrate the conversational roles model approach by Sitter and Stein (1991), so that we are able to track meta-dialogs, and to incorporate the current work on text planning to further improve the dynamic generation of plan-objects.

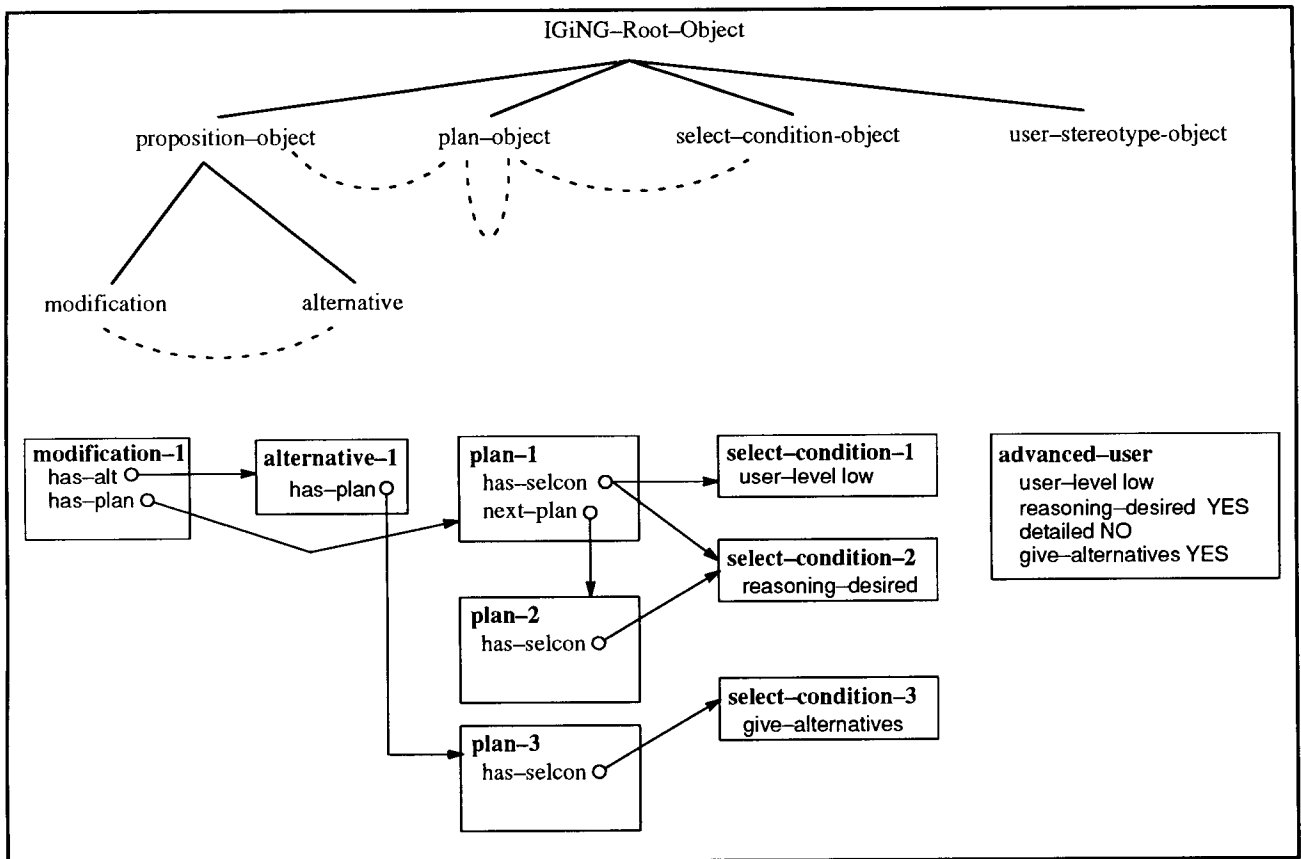


Figure 7 : IGiNG object hierarchy and sample instances

References

- Allgayer, J., Harbusch, K., Kobsa, A., Reddig, C., Reithinger, N., and Schmauks, D. 1989. *XTRA: A Natural Language Access System to Expert Systems*. In: *Int. J. Man-Machine Studies*, Vol. 31, No. 2, 1989, pp. 161-195
- Arens, Y. and Hovy, E.H. 1990. *How to describe what? Towards a theory of modality utilization*. In: *twelfth annual conference of the Cognitive Science Society*, pp. 487-94, Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1990, July 25-28, 1990, Cambridge, USA
- Bandyopadhyay, S. 1990. *Towards an Understanding of Coherence in Multi-Modal Discourse*, Technical Memo TM-90-01, Deutsches Forschungsinstitut fuer Kuenstliche Intelligenz GmbH, Saarbruecken, 1990
- Bateman, J., Kasper, B., Moore, J., and Whitney, R. 1989. *A general organization of knowledge for natural language processing: the Penman Upper Model*. USC/Information Sciences Institute, Technical Report, 1989
- Bateman, J. and Paris, C. 1989. *Phrasing text in terms the user can understand*. In: *Proceedings of IJCAI-89*
- Bateman, J., Maier, E., Teich, E., and Wanner, L. 1990. *Towards an Architecture for Situated Text Generation*, in *International Conference on Current Issues in Computational Linguistics*, Penang, Malaysia, 1991
- Brooks, H.M., Daniels, P.J. and Belkin, N.J. 1986. *A search on Information Interaction and Intelligent Information Provision Mechanisms*. In: *J. of Information Science* Vol. 12, 1986, pp. 37-44
- Cohen, P. R., Dalrymple, M., Moran, D., Pereira, F., Sullivan, J., Gargan Jr., R., Schlossberg, J., and Tyler, S. 1989. *Synergistic Use of Direct Manipulation and Natural Language*. In: Bice, K., and Lewis, C. (eds): *Proceedings CHI '89*, (Austin, Texas, April 30 - May 4, 1989), New York: ACM, 1989, pp. 227-233
- Feiner, S.K. and McKeown, K.R. 1990. *Coordinating Text and Graphics in Explanation Generation*. In: *AAAI-90 Proc. 8th Nat. Conf. on Artificial Intelligence*. July 1990 - Aug. 3, 1990. Vol. I. Menlo Park et al.: AAAI Press/The MIT Press, 1990, pp. 442-449.

- Grosz, B.J. and Sidner, C.L. 1986. *Attention, Intention, and the Structure of Discourse*, Computational Linguistics, Vol. 12, No. 3, pp. 175–204, 1986
- Halliday, 1978. *Language as Social Semiotic*. London: Edward Arnold.
- Hayes, Philip J. 1987. *Steps towards Integrating Natural Language and Graphical Interaction for Knowledge-Based Systems*. In: Boulay, B. du / Hogg, D. / Steels, L. (eds): *Advances in Artificial Intelligence — II*. (Proc. ECAI-86), Amsterdam et al.: North-Holland, 1987, pp. 543–552
- Hobbs, J. 1983. *Why is Discourse Coherent?* In: Neubauer (Editor), *Coherence in Natural Language Texts*, Buske, 1983
- Hutchins, E.L., Hollan, J.D. and Norman, D.A. 1986. *Direct Manipulation Interfaces*. In: Norman, D.A., and Draper, S.W. (eds): *User Centered System Design: New Perspectives on Human-Computer Interaction*. Hillsdale, NJ & London: Lawrence Erlbaum, 1986, pp. 87–124
- HyperNews 1989. *HyperNeWS User's Guide*, Hoff, Arthur van (Editor), The Turing Institute, Glasgow, UK, 1989
- Kasper, R. 1989. *A flexible interface for linking applications to Penman's sentence generator*. Proceedings of the DARPA Workshop on Speech and Natural Language.
- Kerner, A. and Thiel, U. 1991. *Graphical Support for Users' Inferences within Retrieval Dialogues*. In: *Proceedings of the IEEE Workshop on Visual Languages 1991*, pp.211–216.
- Kracker, M. 1991. *Unschärfes assoziatives Begriffswissen zur Unterstützung der Formulierung von Datenbankabfragen*, Dissertation, TU Wien, April 1991. Title in English: *Fuzzy Associative Conceptual Knowledge for Supporting Query Formulation*.
- Lombardi, C. 1989. *Experiments for Determining the Assignment of Information to Media in COMET*, Columbia University, New York, N.Y., USA, 1989
- Mann W.C. and Matthiessen C. 1983. *Nigel: A Systemic Grammar for Text Generation*, Technical Report RR-83-105, USC/Information Sciences Institute, 1983
- Mann, W.C. and Thompson, S.A. 1987. *Rhetorical Structure Theory: A Theory of Text Organization*. In: Polanyi, L. (Editor): *The Structure of Discourse*, Ablex Publishing Company, Norwood, N.J., 1987. Also available as USC/Information Sciences Institute Technical Report Number RS-87-190
- Matthiessen, C.M.I.M. 1990. *Lexicogrammatical Cartography: the systems of English*. Expanding ongoing draft. Department of Linguistics, University of Sydney, Australia.
- McCoy, K.F. 1986. *The ROMPER system: responding to object-related misconceptions using perspective*. Proceedings of the 24th. Annual Meeting of the Association for Computational Linguistics, New York.
- Moore, J.D. and Swartout, J.R. 1990. *Pointing: A Way Toward Explanation Dialogue*. In: AAI-90: Proc. 8th Nat. Conf. on Artificial Intelligence. July 29, 1990—Aug. 3, 1990. Vol. I. Menlo Park et al.: AAI Press / The MIT Press, 1990, pp. 457–464
- Neal, J.G. and Shapiro S.C. 1988. *Intelligent Multi-Media Interface Technology*. In: Sullivan, J.W. / Tyler, S.W. (eds): *Proceedings of the Workshop on Architectures for Intelligent Interfaces: Elements and Prototypes*. ACM/Addison-Wesley, 1989, pp. 69–91
- Penman 1989. *The Penman Project: The Penman Primer, User Guide and Nigel Manual*. USC / Information Sciences Institute, Marina del Rey, CA, 1989
- Reichman, R. 1986. *Communication Paradigms for a Window System*. In: Norman, D.A. / Draper, S.W. (eds): *User Centered System Design: New Perspectives on Human-Computer Interaction*. Hillsdale, NJ & London: Lawrence Erlbaum, 1986, pp. 285–313
- Reichman, R. 1989. *Integrated Interfaces Based on a Theory of Context and Goal Tracking*. In: Taylor, M.M. / Neel, F. / Bouwhuis, D.G. (eds): *The Structure of Multimodal Dialogue*. Amsterdam et al.: North-Holland, 1989, pp. 209–228
- Sitter, S. and Stein, A. 1991. *Modeling the Illocutionary Aspects of Information-Seeking Dialogues*, Arbeitspapiere der Gesellschaft fuer Mathematik und Datenverarbeitung 515, 1991
- Stock, O. 1991. *Natural Language and the Exploration of an Information Space: the ALFresco Interactive System*. IJCAI91, Sydney. pp972–978.
- Swartout, W.R. and Smoliar, S. 1987. *Explaining the link between causal reasoning and expert behavior*. In: *Proceedings of the Symposium on Computer Applications in Medical Care*. Washington, D.C.
- Teich, E. 1991. *A Systemic Grammar of German for Text Generation*, to appear in: Davies, M., L. Ravelli, *Advances in Systemic Linguistics: Recent Theory and Practice*, in progress.
- Thiel, U. 1990. *Konversationale graphische Interaktion mit Informationssystemen: Ein sprechakttheoretischer Ansatz*. Dissertation, Universität Konstanz. Title in English: *Conversational Graphical Interaction with Information-Systems. An approach based on Speechact Theory*.
- Tissen, A. 1991. *A Case-Based Architecture for a Dialogue Manager for Information-Seeking Processes*. In: *Proceedings of SIGIR '91, October 13–16, 1991, Chicago, USA*
- Wahlster, W., André, E., Graf, W., and Rist, T. 1991. *Designing illustrated texts: how language production is influenced by the graphics generation*. Proceedings of the 5th Conference of the European Chapter of the Association for Computational Linguistics.