Deutsches
Forschungszentrum
für Künstliche
Intelligenz GmbH

Document
D-93-05

PPP – Personalized Plan-Based Presenter

Elisabeth André, Winfried Graf, Jochen Heinsohn
Bernhard Nebel, Hans-Jürgen Profitlich
Thomas Rist, Wolfgang Wahlster

May 1993

Deutsches Forschungszentrum für Künstliche Intelligenz GmbH

Postfach 20 80
67608 Kaiserslautern, FRG
Tel.: + 49 (631) 205-3211
Fax: + 49 (631) 205-3210

Stuhlsatzenhausweg 3
66123 Saarbrücken, FRG
Tel.: + 49 (681) 302-5252
Fax: + 49 (681) 302-5341

Deutsches Forschungszentrum für Künstliche Intelligenz

The German Research Center for Artificial Intelligence (Deutsches Forschungszentrum für Künstliche Intelligenz, DFKI) with sites in Kaiserslautern and Saarbrücken is a non-profit organization which was founded in 1988. The shareholder companies are Atlas Elektronik, Daimler-Benz, Fraunhofer Gesellschaft, GMD, IBM, Insiders, Mannesmann-Kienzle, Sema Group, Siemens and Siemens-Nixdorf. Research projects conducted at the DFKI are funded by the German Ministry of Education, Science, Research and Technology, by the shareholder companies, or by other industrial contracts.

The DFKI conducts application-oriented basic research in the field of artificial intelligence and other related subfields of computer science. The overall goal is to construct systems with technical knowledge and common sense which - by using AI methods - implement a problem solution for a selected application area. Currently, there are the following research areas at the DFKI:

- ☐ Intelligent Engineering Systems
- ☐ Intelligent User Interfaces
- ☐ Computer Linguistics
- ☐ Programming Systems
- ☐ Deduction and Multiagent Systems
- ☐ Document Analysis and Office Automation.

The DFKI strives at making its research results available to the scientific community. There exist many contacts to domestic and foreign research institutions, both in academy and industry. The DFKI hosts technology transfer workshops for shareholders and other interested groups in order to inform about the current state of research.

From its beginning, the DFKI has provided an attractive working environment for AI researchers from Germany and from all over the world. The goal is to have a staff of about 100 researchers at the end of the building-up phase.

Dr. Dr. D. Ruland
Director

PPP - Personalized Plan-Based Presenter

Elisabeth André, Winfried Graf, Jochen Heinsohn, Bernhard Nebel, Hans-Jürgen Profitlich, Thomas Rist, Wolfgang Wahlster

DFKI-D-93-05

This work has been supported by a grant from The Federal Ministry for Research and Technology (FKZ ITW-8901 8).

$\mathcal{PPP}$

$\mathcal{P}$ersonalized $\mathcal{P}$lan-Based $\mathcal{P}$resenter

— Project Proposal

Dipl.-Inform. Elisabeth André
Dipl.-Inform. Winfried Graf
Dipl.-Inform. Jochen Heinsohn
Dipl.-Inform. Dr. Bernhard Nebel
Dipl.-Inform. Hans-Jürgen Profitlich
Dipl.-Inform. Thomas Rist
Prof. Dr. Wolfgang Wahlster

German Research Center for Artificial Intelligence (DFKI)
Stuhlsatzenhausweg 3, W-6600 Saarbrücken 11, Germany
e-mail: {last-name}@dfki.uni-sb.de

June 21, 1993

Contents

1 *PPP* Project Outline

The aim of the project 'Personalized Plan-Based Presenter' (*PPP*) is to explore and develop innovative presentation techniques for future intelligent user interfaces. The central issues of the project are:

- Planning Multimedia Presentation Acts
  A presentation system not only has to synthesize multimedia documents, but also has to plan how to present this material to various users. One objective of the *PPP* project is to emulate more natural and efficient presentations by using an animated character as a presenter who will show and explain the generated material.

- Interactive Multimedia Presentations
  Since it is impossible to anticipate the needs and requirements of each potential user, a presentation system should allow for user interaction. The *PPP* system responds to follow-up questions about the domain as well as to meta comments on the presentation act.

- Monitoring the Effectiveness of a Presentation
  In order to find out whether the user has really understood an instruction, a system must monitor the effects of its presentation. One way of getting feedback is using a data bus to physically connected technical devices, which are to be manipulated by the user, with the presentation system. Based on such a connection, the *PPP* system keeps track of the user's behavior and continuously adapts its presentations to the current situation.

- Providing a Firm Representational Foundation
  In order to allow for easy adaptations of new domains, representational techniques flexible and powerful enough to support a wide range of applications have to be employed. Further, these representation techniques should be accompanied by appropriate reasoning techniques that support the implementation of the multimedia presentation system.

Presentation design can be viewed as a relatively unexplored area of common-sense reasoning. Unlike most research on common-sense reasoning to date, the *PPP* project does not deal with metadomain research on general design principles, but focuses on formal methods capturing some of the reasoning in the design space of presentations for specific and realistic domains. The development of an interactive, multimedia presentation system requires efforts from various research areas such as planning, knowledge representation, constraint processing, natural language, and knowledge-based graphics generation.

2 *PPP* Project Description

2.1 The Need for Interactive Multimedia Presentation Systems

Rapid progress in technology for information processing, storing, distribution, and displaying is paving the way for the information society of the next century. It is one thing to have great potential in producing and accessing vast amounts of information, but quite another to make information available to human users in a profitable way. Since presentation of information is becoming more and more crucial in an expanding field of applications, intelligent presentation systems are needed as important building blocks for the next generation of user interfaces. Such presentation systems should be able to generate interactive multimedia presentations in order to account for:

- *Adaptivity*
  Interactive multimedia presentation systems translate from the narrow output channels provided by most of the current application systems into high-bandwidth communications tailored to the individual user. The need for adaptation is based on the fact that it is impossible to anticipate the needs and requirements of each potential user in an infinite number of presentation situations. In an intelligent presentation system like *PPP* design decisions concerning the presentation can be postponed until runtime and all parts of the presentation can be generated on the fly and so customized for the intended target audience and situation.

- *Effectiveness*
  In many situations, information is presented efficiently and effectively only through a particular combination of communication modes. For example, when explaining how to use a technical device, humans will often utilize a combination of language and graphics. It is a rare instruction manual that does not contain illustrations. Multimedia presentation systems take advantage of both the individual strength of each communication medium and the fact that several media can be employed in parallel, e.g., natural language and graphics to produce a flexible and efficient information presentation. Moreover, facilities of modern computer technology provide the potential to generate advanced presentations that go beyond the linear, static nature imposed by paper-printed documents. Examples are hyperdocuments, simultaneously commented animation, interactive graphics, and virtual realities. If carefully designed, these presentations will be much more effective than presentations based on traditional techniques, e.g., hardcopies, could ever be.

- *Reactivity*
  Fixed presentations such as paper printed documents provide only a one-way exchange of information. Since a user confronted with a non-interactive document is lost when he does not understand an instruction, it would be much better to allow for feedback from the user. Probably the greatest opportunity an interactive multimedia presentation provides lies in the generalization of methods, which generate cooperative responses to follow-up questions in natural language dialog systems, to the broader domain of multimodal communication. Moreover, in a sophisticated

presentation system the user could even criticize the ongoing presentation. Apart from direct user interaction, the presentation system could also obtain indirect feedback on the user's reaction to a presentation. This would make sense especially in a maintenance and repair application where the presentation system instructs a user. Based on an evaluation of the user's physical behavior after he has received instructions, the presentation system will be able to keep track of the relevant behavior of the user, monitor the effectiveness of the presentation and continuously adapt its presentations to the current situation.

- *Consistency*
  Intelligent presentation systems guarantee the consistency over several presentations. This is useful especially in technical documentation since companies will not have to waste time and money in designing similar instruction manuals again and again after small product changes.

Rapidly expanding activities in intelligent multimedia interfaces provide evidence that the importance of multimedia in human-computer communication has been well recognized world-wide. There are new funding programs currently in preparation, e.g., in USA, Japan, and France. Universities have founded multimedia groups (e.g., MIT Media Lab, Stanford University, UC Berkeley). Industrial interest and support have been shown by nearly all larger companies (e.g., Apple, IBM, Microsoft, SUN, Intel, NeXT, and Siemens). In Japan the Human Interface Laboratories at NTT and the FRIEND 21 project funded by all major companies are the driving force behind the research in this area. Specialized new conference series have been set up, e.g., IJCAI-89 Workshop on 'A New Generation of Intelligent Interfaces' (cf. [Arens *et al.*, 1989]), ACM Symposia on 'User Interface Software and Technology' (UIST, cf. confererence proceedings 1988-1992), International Workshop on 'Intelligent User Interfaces' (cf. [Sullivan and Tyler, 1991]), Workshop on 'Task Communication through Natural Language and Graphics' (cf. [Badler and Webber, 1990]), NATO Workshop on 'Computational Theories of Communication and their Applications' (cf. [Ortony *et al.*, 1992]), AAAI-91 Workshop on 'Intelligent Multimedia Interfaces' (cf. [Maybury, 1992]), International Workshop on 'Aspects of Automated Natural Language Generation' (cf. [Dale *et al.*, 1992]), and Advanced Visual Interfaces Workshop (AVI, cf. [Costabile *et al.*, 1992]). Furthermore a new ACL special interest group on Intelligent Multimedia Interfaces has been established and the first international book on 'Intelligent Multimedia Interfaces' will be published by AAAI Press ([Maybury, 1992]).

For the next International Joint Conference on Artificial Intelligence, IJCAI-93, a panel on 'Instructions and Language' has been organized by Prof. Webber (UPenn) and the *PPP* team has been invited to prepare a contribution on multimodal instructions for this panel discussion. Finally, in Saarbücken two spin-off companies have been founded by former members of Prof. Wahlster's research group, that develop, sell and deploy multimodal interfaces. The HQ company developed various multimodal information systems for Sony and multimodal entertainment systems for Philips and Ravensburger. The TransModul company sells an interactive multimodal interface to the DOS operating system (DOS-MAN) that integrates natural language and pull-down menus.

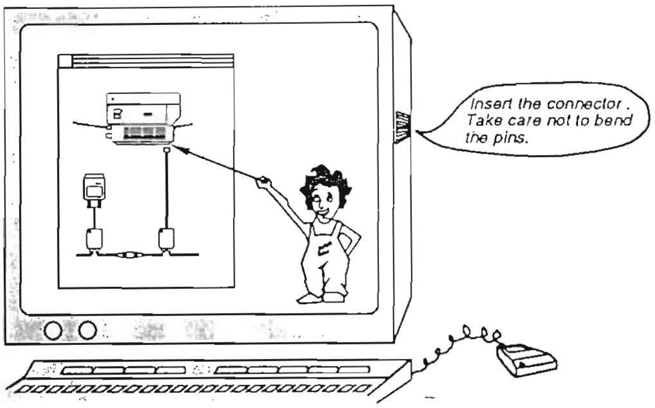

Figure 1: On-line Presentation Acts

2.2 Main Goals of the Project

2.2.1 Planning Presentation Acts

It cannot be denied that the success of human-human communication depends essentially on the rhetorical and didactical skills of the speaker or presenter. However, little attention has been paid to this aspect of computer-based presentation systems. Up to now, research has mainly concentrated on content selection and content encoding. Although multimedia documents which are synthesized by these systems may be coherent and even tailored to the user's specific needs, the presentation as a whole may fail because the generated material has not been presented in an appealing way. Such situations can often be observed when multimedia output is distributed on numerous windows, and the user himself must find out how to navigate through the presentation.

More efficient presentations are expected when using an animated character called PPP which will play the role of a presenter, showing, commenting and explaining the generated material. Thus the system should be able to plan presentations as well as presentation acts and their temporal coordination. This is exactly what speakers have to do when preparing a talk. They have to produce or select the material to be presented and to plan what to say at which point. Note that such a presentation system would support two operating modes. Firstly, it can be used for the generation of on-line screen presentations, e.g. in a maintenance and repair domain. Secondly, it can serve as a presentation tutor that assists a human speaker when preparing a talk. In this case, a generated presentation is considered as a proposal for presenting the synthesized text-picture combinations.

Figure 1 shows the conversational setting which is assumed in *PPP*. In the example, the system instructs the user in connecting the printer with his notebook. In order to accomplish this task, the system has to carry out the following presentation acts:

p-act-1: Show picture-1.
p-act-2: Say "Insert the connector".
p-act-3: Say: "Take care not to bend the pins".
p-act-4: Point to picture-object-2 in picture-1.

In contrast to existing presentation systems like WIP, *PPP* relies on an explicit ...

tation of temporal relationships between presentation acts. For example, we must express that a pointing gesture and speech output should start or finish at the same time. The following plan for the above presentation acts could be formulated by using, for example, Allen's interval-based temporal logic:

```
(AND (DURING p-act-3 p-act-1)
     (BEFORE p-act-2 p-act-3)
     (DURING p-act-4 p-act-1)
     (DURING p-act-2 p-act-1)
     (OR (OVERLAPS p-act-2 p-act-4)
         (DURING p-act-4 p-act-2)
         (FINISHES p-act-4 p-act-2)))
```

An important objective of *PPP* is to represent these temporal relationships in the framework of a terminological logics and to use the *PPP* knowledge representation formalism for representing domain plans as well as presentation plans.

2.2.2 Interactive Multimedia Presentations

It is obvious that a presenter cannot always have a detailed model of each individual conversational partner. Often the presenter's assumptions about the wants and beliefs of his audience are incomplete or even incorrect. Consequently, humans sometimes do not understand an instruction or they are rarely satisfied with the presentation. In such cases, it is quite natural to ask follow-up questions or to criticize the style of the presentation.

In order to emulate the multimodal interaction that occurs between humans *PPP* supports user interaction by taking advantage of hypermedia techniques. In particular, in the *PPP* system the user can interrupt the system and ask questions about the presentation already generated and change generation parameters during the presentation, e.g., by demanding the system to change the level of detail or the speed of the current presentation.

In Figure 2(a) the user clicks on a part of a generated hypergraphics. Thus, the presentation is interrupted to offer the user a menu of possible follow-up questions. The items on the menu are generated by the presentation system in a context-sensitive way. Figure 2(b) illustrates how the user criticizes the ongoing presentation. He clicks on the animated character and obtains options to change the presentation style.

Such situations require the system to revise the initial presentation plan. This includes inserting new subplans, reinstantiating variables or reordering goals. To enable immediate reactions to unexpected situations to take place, planning and execution must be interleaved.

In *PPP*, we see dialogues as resulting from the plans and the goals of the participants. The RST-based planner developed in the WIP project will provide the basis for extensions towards dynamic interactive presentations.

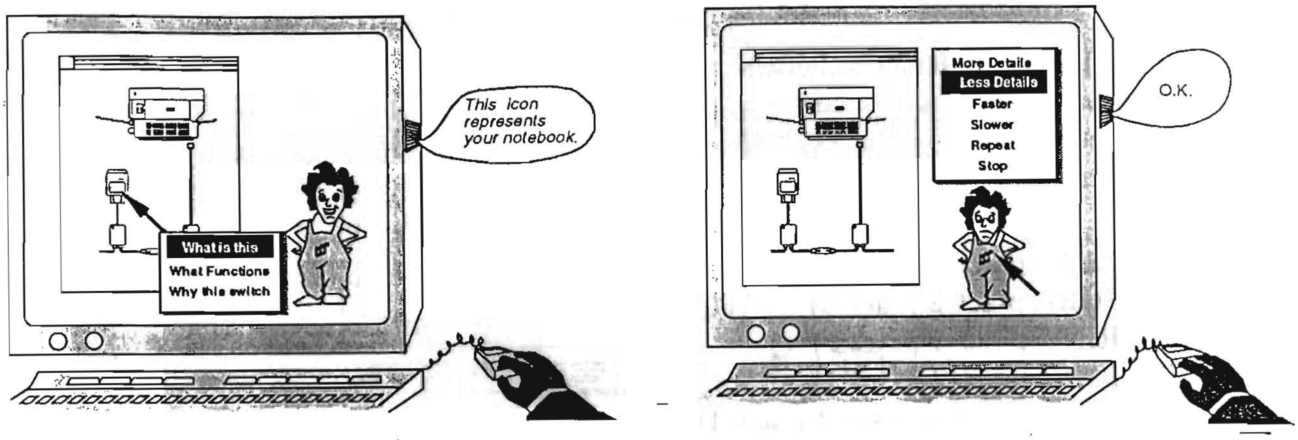

Figure 2: (a) Asking Follow-Up Questions, (b) Criticizing the Presentation Style

2.2.3 Monitoring the Effectiveness of a Presentation

Most approaches in the automatic presentation of information do not consider the user's reaction to a presentation. However, this is a severe limitation for most application scenarios where the presentation system has to communicate instructions, which must be carried out on-line by a user. In such an environment, problem solving becomes an iterative process involving the user, the application, and the presentation system (cf. Figure 3(a)). The presentation system must get feedback as to whether the user really understood the instructions in order to monitor the effectiveness of presentations and to continuously adapt these presentations to the current situation. A visual observation of the user's physical behavior after he has received instructions could provide the necessary information. However, this would require a sophisticated vision system which unfortunately is not available to date. An alternative is to physically connect the presentation system with the device to be manipulated via a data bus. This seems to be a more realistic alternative since data buses for technical devices are already available in many working areas. Such a situation is exemplified in Figure 3(b). In this case, the presentation system provides on-line help in maintaining a printer. By using the data bus, the presentation system receives information about whether instructions are carried out as intended. Since in $\mathcal{PPP}$ we will concentrate on presentation issues, we do not aim on diagnosis for troubleshouting in the application domain. In place of a diagnosis component, we will exploit status reports of the connected hardware in order to trigger predefined domain plans for problem solving. Of course, the integration of diagnosis components, e.g., an expert system for the diagnosis of printer problems such as the μ-UNIXPERT (cf. [Lessel and Boley, 1987]) would be a reasonable augmentation.

This application scenario also illustrates the need for reactive planning since existing plans have to be flexibly modified to adapt them to the new situation.

2.2.4 Providing a Firm Representational Foundation

Building a multimedia presentation system that can be used for more than one application domain requires the use of representational techniques which are powerful and flexible enough to cover a wide range of possible application domains. In addition, these techniques should be accompanied by appropriate reasoning techniques that support the

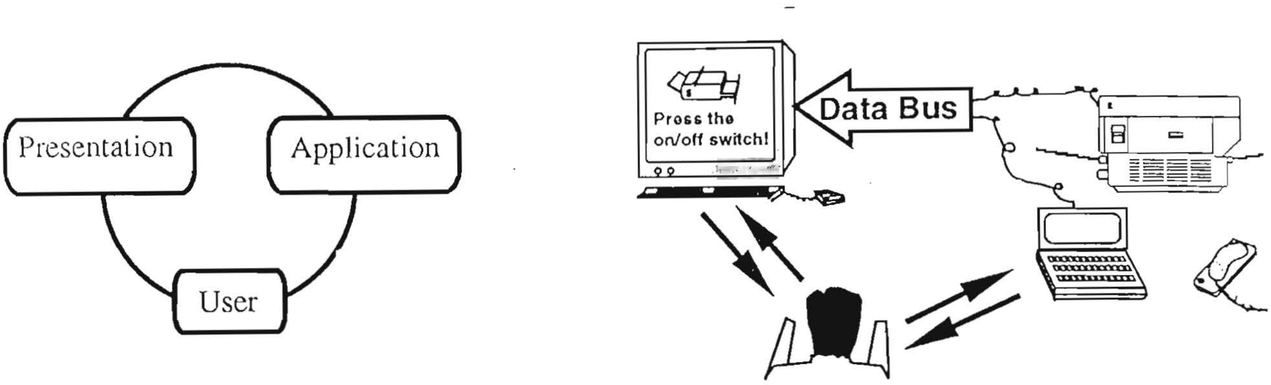

Figure 3: (a) Presentation Situation in $\mathcal{PPP}$, (b) $\mathcal{PPP}$ Application Domain

implementation of multimedia systems.

Since one of the design principles behind $\mathcal{PPP}$ is that the theoretical basis of all components should be sound enough to allow for scaling up, we will use and combine only relatively mature techniques, such as tree adjoining grammars for natural language generation. hierarchical planning and RST-theory for presentation planning, constraint-propagation techniques for layout design, and terminological logics for the representation of domain knowledge. The generation of interactive multimedia presentations creates challenges for the representational and reasoning subsystem that go beyond those usually encountered in generation systems. For instance, in order to deal with user interactions that express misunderstandings, the system must revise its beliefs about the user beliefs dynamically. Furthermore, in order to coordinate the presentation with the actions of a user, temporal reasoning must be incorporated into the presentation planning task. Finally, in order to allow for conceptually simple ways of representing and manipulating the knowledge. it seems desirable to provide a uniform representation for apparently different tasks that are structurally similar, such as domain and presentation knowledge.

Terminological Logics Terminological logics have been successfully applied in a number of different systems to represent important parts of the application domain. Basically. these logics provide a functionality offered by most semantic network formalisms extended by the facility to automatically *classify* new objects and concepts. This reasoning service can be exploited in the context of word-choice, when evaluating the specificity of presentation strategies and the applicability of such strategies as well as in the retrieval of presentation plans and canned multimedia units, such as video clips and graphical material. Terminological logics represent a mature technology which has even been subject to a standardization effort as part of the "DARPA Knowledge Sharing Effort." These logics will serve as the "representational backbone" around which all other representation and reasoning services are centered.

Reasoning about Action, Change, and Time Reasoning about actions and time takes place at the level of domain plans (representing operating instructions, for instance) and at the level of actual multimedia presentations. In order to represent and to reason about operating instructions, it is necessary to use some representational tools that are able to represent temporal orderings of actions and the causal relationships between actions and state changes. The RAT system (which is based on KRIS [Baader and Hollunder.

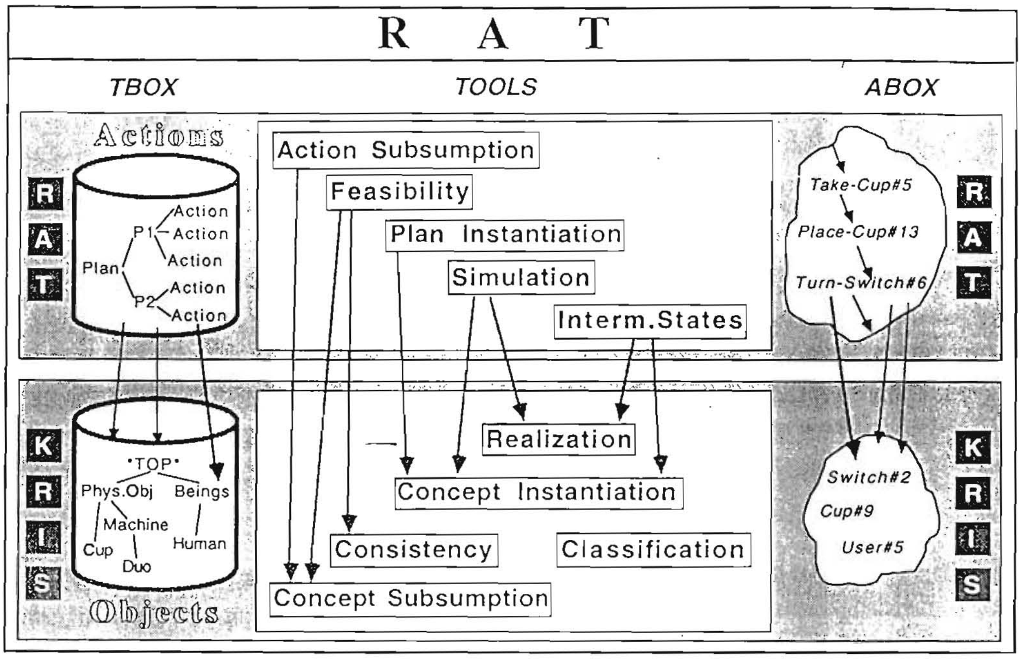

Figure 4: The Architecture of the RAT system

1991]) developed in the WIP project provides a firm base for the further development of such a tool (see Figure 4). In particular, we anticipate the need to represent more complex descriptions of the temporal ordering of actions and the need to deal with the interdependencies between simultaneously occurring actions. Such extensions are also necessary for dealing with the more complex task of planning interactive presentations as described in Section 2.2.1.

Incompleteness and Uncertainty A further prerequisite for a uniform knowledge representation and reasoning framework for the support of presentation planning concerns the handling of *ignorance* about domain and presentation knowledge. One particular type of ignorance, so-called *uncertainty*, is inherently present in interactive presentation planning and domain and user modeling and deals with those cases in which the current state of affairs is not completely determined but where we have to rely on preferences for the different possibilities. Planning a multimedia presentation in complex domains means that various possibly *graded* user characteristics (such as trained or untrained), *preferences* in the user's task, *rankings* in the effectiveness of presentation strategies. and *priorities* on more or less preferred modes are involved so that the representation formalism must be able to represent and combine *uncertain and incomplete* information sources of this kind to allow for decision making. This task also takes into account knowledge of categorical *exceptions*, like, "prefer the textual presentation unless the user is illiterate." Uncertain knowledge is also present in the automatic graphical layout of multimedia presentation [Graf, 1992]. While in the case of typical business letters there may be little doubt about the "best" structural layout, there may be lot in the case of instruction manuals or overhead slides for presentations. This aspect is closely related to the complex positioning problem for multimedia units, where, beside general consistency requirements and basic design principles, an "aesthetically pleasing layout" is also a requirement. These subjective criteria cover "typical" design decisions with regard to presentation task and user model and are *weighted* to allow the *ranking* of constraints in underconstrained cases.

Reasoning about Belief For personalized plan-based presentation, the most important issue is the *interactivity* of the planning process, that takes into account criticisms of a user and applies revision strategies. If revision of the (maybe partial) presentation plan is deemed necessary, the represented *knowledge about the user and user beliefs* must be taken into account. Therefore, representation and reasoning services must be provided, that allow for *reasoning about beliefs*. For example, to detect the more or less probable explanations for the interaction of the user that allow the revision of the presentation planning process, *(weighted) abductive techniques* may be appropriate.

Retrieval of Multimedia Units In the automatic layout of multimedia presentations in complex domains generally a large set of representational units such as text fragments, graphics, videos, animation, and virtual reality is employed. Thus the *retrieval of multimedia units* appears as an important task in knowledge representation and reasoning.

Closely related to information retrieval is the potential *reuse* of parts of this knowledge retrieved. In the framework of the *PPP* project, reuse can be viewed as the reapplicability of various kinds of knowledge including, e.g., design decisions previously made for parts of the document, knowledge of the document structure, and parts of the presentation plan. Taking presentation planning as an example, reuse of parts of the presentation plan means saving time and costs and reducing the risk of redundant and/or partially inconsistent knowledge.

Efficient Inference Mechanisms While the main problems in designing knowledge representation and reasoning (KR&R) systems to support multimedia presentation systems are of a conceptual nature, the efficiency of these systems cannot be completely ignored. First of all, the KR&R services must be efficient enough to allow for a reasonable overall performance of the system. Secondly, in order to permit scaling up of the system, algorithms must be provided, such that the runtime does not unreasonably increase in the size of the knowledge base.

Since most representation formalisms must be flexible enough to deal with a wide range of different situations, and since the reasoning services have to be powerful enough to support non-trivial tasks, usually it is not possible to guarantee that a KR&R system is efficient in all cases. Indeed, most reasoning services that are needed to support, for instance, multimedia presentations are computationally intractable in the worst case. Nevertheless, some level of performance must be guaranteed for the cases that occur in practice.

2.3 Application of Intelligent Presentation Systems

There is a growing application base for intelligent interactive multimedia presentation systems. Some interesting applications scenarios will be sketched below:

- **Multimedia Instructions**
  A good example for the *PPP* system are instructions for the maintenance, service and repair of technical devices. Computer-based presentation techniques provide more effective media for instructing people in task performance since they overcome

problems arising from the static and non-interactive nature of conventional technical documentation. Furthermore, interactive presentation systems provide a low-cost way of allocating personal trainers to learners. This has already been noted by industries, e.g., most car producers have begun introducing multimedia technology to train their mechanics (cf. [BDW, 1992]).

- **Adaptive Control Panels**
  With increases in the amount of information that must be communicated to the users of complex technical systems, a corresponding need arises to find new ways to present that information flexibly and efficiently. Siemens and Daimler-Benz are developing adaptive user interfaces for control panels in aircraft cockpits, cars, industrial plants and traffic control stations. They are already using multimodal systems, but are dissatisfied with the current level of media coordination and adaptability. The next generation of intelligent control panels must include the explicit planning of situated and tailored presentations. It is clear that PPP's approach to monitor the effect of presentations via a data bus is very attractive in the above mentioned applications. For the next generation of Mercedes cars a single data bus will collect information about the driver's behavior, the sensor measurements and all critical electronic and mechanical parts of the car, so that data fusion and multimodal presentation in a *hands-and-eyes-busy* situation will be an essential innovation. In particular, the combination of speech output coordinated with animated graphics is the wave of the future for the corresponding divisions of Siemens and Daimler-Benz.

- **Computer-Supported Collaborative Work (CSCW)**
  The concept of tailoring presentations for the user can be seen as an extended version of the view concept known from database technology. One step on the way to intelligent interfaces for computer-supported collaborative work (CSCW) is to use multimodal systems like PPP as presentation experts that map fragments of a shared knowledge-base onto a variety of presentations satisfying the information needs of the individual group members. It is clear that in a distributed setting various constraints for the individual members of a team supported by a groupware system have to be satisfied. Thus the same information should be displayed in different forms to the members of a team, e.g., in the setting of an international collaboration the information should be conveyed in the various mother tongues of the participants. At the same time the group members may have a diverse set of technical backgrounds, so that the presentation has to be tailored to various levels of expertise. Siemens has various strong CSCW groups that are interested in exploiting the techniques developed in PPP for their product development.

System	Media	Generation of Graphics	Communication between the Generators	Current Visual Domain	Project Team
XTRA	NL, graphics, pointing	manual	None	tax forms	Wahlster et al. (Saarbrücken)
CUBRICON	NL, graphics, pointing	partially automatic	None	geographic maps	Shapiro/Neal et al. (Buffalo)
ALFresco	NL, video, pointing, hypertext	manual	None	frescoes	Stock et al. (Trento)
MMI²	NL, graphics, pointing	automatic	None	computer networks, charts, tables	Wilson et al. (Oxon)
II	NL, graphics, menus	partially automatic	None	geographic maps	Arens et al. (Marina del Rey)
PEA	NL, hypertext	-	None	-	Moore (Marina del Rey)
IDAS	NL, hypertext	-	None	-	Reiter et al. (Edinburgh)
Weather Report System	NL, graphics	partially automatic	None	weather maps	Kerpedjiev (Sofia)
Map Display System	NL, graphics, pointing	manual	None	geographic maps	Maybury (Bedford)
SAGE	NL, graphics	automatic	None	business charts	Roth et al. (CMU)
FN/ANDD	NL, graphics	automatic	None	network diagrams	Marks/Reiter et al. (Harvard)
WIP	NL, graphics	automatic	between NL and graphics generator	espresso machine, mower, modem	Wahlster et al. (Saarbrücken)
COMET	NL, graphics	automatic	between NL and graphics generator	portable radio	Feiner/McKeown et al. (Columbia)
AnimNL	animated graphics	automatic	None	cooking devices	Badler et al. (Pennsylvania)

Figure 5: Current Research on Combining Natural Language, Graphics, Hypertext and Pointing

3 State of the Art

3.1 Presentation Planning and Design

3.1.1 Multimedia Presentation Systems

In the last few years, a number of projects have entered the area between natural language processing and multimodal communication, often focusing on a single specific functionality, such as the use of pointing gestures parallel to verbal descriptions for referent identification ([Kobsa *et al.*, 1986; Cohen *et al.*, 1989; Neal and Shapiro, 1991]). The automatic design of multimedia presentations has only recently received significant attention in artificial intelligence research. The most extensive discussion of active research in this field is documented in the proceedings of a series of workshops on intelligent multimedia interfaces (e.g., [Arens *et al.*, 1989; Sullivan and Tyler, 1991; AAAI-92, 1992; Costabile *et al.*, 1992]). Overviews on intelligent multimedia presentation and dialog management systems can be found in [Roth and Heffley, 1992; Edmonds and Murray, 1992]. Fig. 5 gives a survey of research activities in this area.

The first group of systems compared in Fig. 5 (XTRA, CUBRICON, ALFresco, MMI², I², PEA and IDAS) consists of multimodal dialog systems with an analysis and generation component. XTRA (cf. [Allgayer *et al.*, 1989]) provides multimodal access to an expert system that assists the user in filling out a tax form. CUBRICON (the CUBRC Intelligent CONversationalist, [Neal and Shapiro, 1991]) is an intelligent interface to a system for mission planning and situation assessment in a tactical air control domain. ALFresco (cf. [Stock, 1991]) displays short video sequences about Italian frescoes on a touchscreen and answers questions about details of the videos. Whereas the pointing actions and natural language utterances in these systems refer to visual presentations provided by the system builders, MMI^2 (A Multi-Modal Interface for Man Machine Interaction with Knowledge-Based Systems, [Wilson *et al.*, 1992]) also offers several graphical tools to assist the user in designing computer networks. In order to avoid many of the difficult referential problems in understanding natural language, I^2 (Integrated Interfaces, [Arens *et al.*, 1991]), PEA (Program Enhancement Advisor, [Moore and Swartout, 1990]) and IDAS (Intelligent Documentation Advisory System, [Reiter *et al.*, 1992]) do not have a natural language analysis component, but offer the user menus and forms or even a hypertext-style interface.

The second group of systems listed in Figure 5 focuses on the presentation task. They are designed more or less as presentation systems although the eventual application environment may also be that of an interactive system. SAGE (a System for Automatic and Graphical Explanation, [Roth *et al.*, 1991]) is a presentation system that uses text and graphics to explain the changes in the results generated by quantitative modeling systems. The ANDD (Automated Network-Diagram Designer) system automatically designs network diagrams from a list of relations and a basic network model whereas the FN system generates natural language expressions describing certain attributes of a particular object shown in the diagrams (see [Marks and Reiter, 1990]). Kerpedjiev has designed a system that transforms a dataset about a particular weather situation into a multimodal weather report consisting of a text illustrated by tables and weather maps with various icons and annotations (cf. [Kerpedjiev, 1992]). Maybury (cf. [Maybury, 1991]) is concerned with the planning of multimedia directions for a knowledge-based cartographic information system.

All the systems in Figure 5 *combine* natural language and graphics, but only systems that generate both forms of presentation from a common representation and allow for communication between the media-specific generators can address the problem of automatic *media choice and coordination*.

WIP (Knowledge-Based Presentation of Information, [Wahlster *et al.*, 1989]) and COMET (COordinated Multimedia Explanation Testbed, [Feiner and McKeown, 1990]) are the only systems in which the media-specific generators communicate with each other in order to achieve a fine-grained and optimal division of work between the selected presentation modes. Both systems deal with physical objects (espresso-machine, radio) that the user can access directly. For example, in the WIP project we assume that the user is looking at a real espresso-machine and uses the presentations generated by WIP in order to understand how the machine works. Likewise, COMET generates directions for the maintenance and repair of a portable radio, using text coordinated with 3D graphics. Although many similarities exist, there are also major differences between COMET and WIP, e.g., in the systems' architecture. To handle dependencies between content and mode selection, WIP selects the medium in which information should be presented during

	Informational Graphics	3D Graphics of Physical Objects
Static Media	Maps, Charts, Diagrams Example Systems: SAGE, FNN	Rendered Pictures Example Systems: WIP, COMET
Dynamic Media	Hypermedia Presentations Example Systems: AlFresco, IDAS	Animation Example Systems: VITRA-SOCCER, AnimNL

Figure 6: Combining Text Production with Four Types of Graphics Generation

content planning and not after as in COMET. Furthermore, WIP enables bidirectional communication to take place between the presentation planner and the layout manager. During one of the final processing steps of COMET the media layout component combines text and graphics fragments produced by media-specific generators, while in WIP a layout manager interacts with a presentation planner before text and graphics are generated, so that layout considerations can influence the early stages of the planning process and constrain the media-specific generators.

Whereas the majority of work has concentrated on combining static media, the VITRA-Soccer project (cf. [Herzog *et al.*, 1989]), for details of VITRA's animation component see [Schirra, 1992]), the AnimNL project (cf. [Badler *et al.*, 1991b]) and recent extensions of COMET (cf. [Feiner and McKeown, 1992; Feiner *et al.*, 1991]) and WIP in addition deal with dynamic media, such as animation. Systems like AlFresco (cf. [Stock, 1991]) and IDAS (cf. [Reiter *et al.*, 1992]) demonstrate that natural language generation can be enhanced by integration with hypermedia systems. In such systems the generated text may contain links to hypercards and canned text or images can be combined with generated text for a hypermedia presentation.

Figure 6 summarizes the various types of graphical presentations that have been combined with generated text in recent research prototypes. In all these projects the generation system is no longer merely the author of a text, but also plays the role of a desktop publisher, a hypertext designer, a multimodal interface designer or an animations commentator.

3.1.2 Automated Graphics Generation

Since graphics provides considerable potential in presenting information often more effectively than any other media can, graphic-based communication has provoked significant interest in research in intelligent user interfaces. Evidently it would not be feasible to handcraft and store graphics for each possible combination of relevant presentation parameters such as user characteristics, situation, and resource limitations. This leads to the question of how to automatically design and generate particular graphics for particular purposes on the fly in a context-sensitive way.

Previous work on the automatic design of graphics can be distinguished in view of the kind of graphics to be generated and the underlying design methods. The spectrum of graphics ranges from abstract presentation graphics such as pie- and bar charts (cf. [Gnanamgari, 1981], [Zdybel *et al.*, 1981], [Mackinlay, 1985], [Mertens, 1988], [Kansy, 1991], and [Casner, 1991]), node-link diagrams and networks (cf. [Kahn, 1979] [Marks, 1991]), symbol-based diagrams, e.g., for the visualization of process information in industrial control (cf. [Elzer *et al.*, 1988]), the presentation of electrical circuits (cf. [Geller and Shapiro, 1987]) and weather maps (cf. [Kerpedjiev, 1992]), schematic line drawings, e.g., to describe chemical apparatus (cf. [Strothotte, 1989]), up to 3D object depictions and environments (cf. [Friedell, 1984]) and illustrations of 3D objects (cf. [Feiner, 1985], [Seligmann and Feiner, 1991], [Rist and André, 1992a]). No less interesting, but somewhat unusual, are several approaches dealing with mental imagery (cf. [Kosslyn, 1980]). Instead of graphics being presented on a screen or printed on paper, they generate so called mental images as an analogical form of knowledge representation.

Whereas several approaches rely on a pure selection of predefined graphical presentations (cf. [Gnanamgari, 1981], [Zdybel *et al.*, 1981], [Mertens, 1988]) and thus do not address design issues, others provide techniques in order to select and combine graphical elements. Such "compositional" approaches can be further distinguished in view of the primitives they use. Several approaches rely on predefined icons that are stored in a database, either as bitmaps (e.g., see [Kerpedjiev, 1992], [Strothotte, 1989]) or as propositional descriptions (e.g., [Friedell, 1984] and [Geller and Shapiro, 1987]). Alternatively, following the approach of the graphics designer Bertin (cf. [Bertin, 1983]) a graphics can be described as an implantation of spots (either points, lines or areas) into an empty 2D drawsheet. With respect to perceptible variations of a spot, Bertin distinguishes between eight visual variables (x- and y-position in the plane, size, intensity, pattern, color, direction and shape). A particular piece of information is then encoded by certain variations of visual variables. Bertin's view of graphics has proven to be quite useful for the automated design of abstract presentation graphics (cf. Mackinlay's APT system described in [Mackinlay, 1985]). However, it is not clear how this approach can be transferred to graphics with illustrations of material 3D objects. The depiction of an object may be described as a configuration of spots together with their specific visual properties, but in general it is very difficult and costly to specify which information is encoded by which variation of spots. For example, to show an object, there is a choice of numerous perspectives. Each choice affects the arrangement of corresponding spots as well as the shape and size of the spots.

Important work on the generation of depictions of 3D objects, without relying on predefined icons, has been carried out by Feiner and Seligman (cf. [Seligmann and Feiner, 1991]). In their system IBIS, 3D objects are related to illustration objects on the picture level. When generating illustration objects, they consider both the underlying representation of the 3D object in the knowledge base and the purpose for which the illustration will be used. They use a generate and test approach in order to achieve a close relationship between the visual appearance of an object in the world and its appearance in the illustration. The graphics generator developed by Rist and André addresses both 3D and 2D graphics (cf. [Rist and André, 1992a] and [Rist and André, 1992b]). As in the

IBIS system, they start with a more or less complex presentation goal to be accomplished by graphics. Using a plan-based approach, presentation goals are refined in a top-down manner and eventually mapped onto realization operators which effect either 3D object models, functions for projecting 3D models, or the 2D constituents of a picture.

In most approaches the design process is driven by the information to be presented, and the communicative goal of a graphical presentation. Design decisions then frequently rely on heuristics and rules-of-thumb which are more or less empirically substantiated. An attempt towards a more analytical design approach has been made by Casner in his BOZ system (cf. [Casner, 1991]). Starting from an analysis of the task to be performed by the user, BOZ transforms a logical task description into a perceptual task description by substituting perceptual inferences in place of logical inferences. It then designs a graphics such that each perceptual inference is supported and visual search is minimized. On the one hand this seems to be a promising approach since it provides a means to characterize the effectiveness of a graphics by counting perceptual operations to be performed, whereas on the other hand, it is questionable whether the approach actually mirrors human perceptual behavior, e.g., in BOZ perceptual tasks are always modeled as a sequence of operations. Furthermore, it is less clear how to model perceptual tasks concerning the processing of complex 3D graphics.

An interesting approach for the synthesis of mental images has been taken by Kosslyn (cf. [Kosslyn, 1980]). Starting from a hierarchically structured propositional representation of domain objects, he instantiates a 2D cell matrix as a(n) (quasi)analogical representation of the object's shape, size and orientation. The instantiation process begins with a so-called skeletal image which will be recursively refined until all the available propositional information is mentally visualized. This visualization process is context-sensitive, e.g., if attention is focused on a specific part, a zoom operation is performed on the analogical representation of that part.

3.1.3 Automated Design of Animations

Animation as the computational control of images or objects over time is one of the most fascinating forms of presentations a computer system can support. Although animation is widely used in the entertainment industry and in scientific visualisation, it plays a subordinate role in research on intelligent user interfaces. Reasons for this are, among others, the fact that the fine-tuning of animation is a tedious and time-consuming task and that budgets of research projects are often overstrained by the costs of powerful high-speed graphics workstations which are indispensable in most applications including animation.

Previous work on animation concentrated on animation techniques and scripting systems. The spectrum of animation techniques includes key-framing (e.g., [Mezei and Zivian, 1971], [Reeves, 1981]), parametric interpolation (e.g., [Shelly, 1982], [Kochanek and Bartels, 1984], [Steketee and Badler, 1985]), tracking live action (e.g., [Ginsberg and Maxwell, 1983]), kinetics (e.g., [Thalmann and Thalmann, 1990]), inverse kinematics (e.g, [Badler *et al.*, 1980], [Korein, 1985], [Girard, 1987], dynamics (e.g., [Wilhelms, 1986]. [Isaacs and Cohen, 1987], [Wilhelms, 1987]), and constraints (e.g., [Badler, 1987], [Witkin *et al.*, 1987], [Barr and Barzel, 1988]). Animation techniques are usually embedded in a scripting system that provides an interface for a higher-level description of animation. Some scripting systems are conceived as imperative programming languages. There, animation scripts are written in special languages usually based on linear-list notations (e.g., [Catmull, 1972], [Gomez, 1984], [Strauss, 1988], [EXPLORE, 1992]), or in a general-purpose programming language with embedded animation directives (e.g., [Reynolds, 1982], [Thalmann and Thalmann, 1990]). Other scripting systems allow for a graphical specification of animation parameters (e.g., [Baecker, 1969], [Feiner *et al.*, 1982], [S-Dynamics, 1985]).

There are only a few projects in which the automated design of animation has been issued. Kahn's ANI system (ANImation, cf. [Kahn, 1979]) was one of the first attempts at the automatic scripting of animations. Starting from natural language story descriptions of physical actions the ANI system generates icon-based 2D animations. A similar approach, but for stories written in Japanese, has been taken for the system SDA (Story Driven Animation) (cf. [Takashima *et al.*, 1987]). The design of 3D animations has been investigated by Feiner and Karp ([Karp and Feiner, 1990]). They have implemented an expert system called ESPLANADE (Expert System for PLANing Animation Design and Editing) that uses a rule-based approach to automatically choose animation parameters such as camera trajectories. Important work on articulated human figure modelling, task performance assessment in a 3D environment, and animation has been done at UPenn ([Badler *et al.*, 1991b; Zeltzer, 1991]). They are concerned with both natural language-driven generation of animation scripts as well as the automatic synthesis of narrated animations (i.e., animations accompanied with natural language utterances) from propositional task descriptions.

Another context in which the problem of automatic design of animation has been addressed are help systems. There, animation is typically used to visualize a sequence of actions that must be carried out by the user of an application program. Neiman's system GAK (Graphical Animation from Knowledge) was one of the earliest attempts to extend an existing help system with animated help facilities (cf. [Neiman, 1982]). An approach towards domain independency has been made with the Cartoonist system (cf. [Sukaviriya and Foley, 1990]). Instead of developing an animated help system for a specific application programm, Cartoonist retrieves a specification of interaction techniques with the application programm and uses that knowledge to plan animated interaction examples. The animation component AniS$^+$ of the plan-based help system PLUS (cf. [Thies and Berger, 1992]) not only considers the screen context, - as Cartoonist does - but also takes into account the user's task when planning animated help. An extension of this approach to the animation of 3D interaction techniques has been proposed in [Graf and Thies, 1992].

An alternative to the generation of animations from scratch are approaches in which movies are assembled form video clips recorded, e.g., by human camera operators. In this case, the design of an animation is reduced to the selection and linearization of video clips stored in a database. This technique has been used in the Movie-Maps system that simulates driving a car freely through an assortment of US cities (cf. [Lippman, 1980]), and in a system by Rubin (cf. [Rubin, 1989]) that assembles a coherent visual narration from prerecorded video clips.

3.1.4 Automatic Layout

As graphics hardware becomes more and more sophisticated, computer-based graphical communication achieves a crucial role in intelligent user interfaces. While much research in this area has been focused on the automatic synthesis of graphics for either presenting relational information and realistic depictions of 3D objects (cf. Chap. 3.2), the automatic layout design of graphical presentations has remained unexplored. Beach (cf. [Beach, 1985]) has shown that the general layout problem formalized as a random packing problem, i.e., determining whether an unordered set of non-overlapping rectangular table entries can be arranged into a minimum space, is strongly *NP-complete* and thus, there is no general and efficient algorithm for solving it. So even the problem of finding an aesthetically pleasing layout for multimedia documents under certain outward restrictions seems to be intractable. Current work on layout design is essentially influenced by ideas and approaches known from general graphics design (e.g., [Müller-Brockmann, 1981; Lieberman, 1990]), computer graphics (e.g., [Foley *et al.*, 1990]), and psychology of visualization (e.g., [Arnheim, 1966; Csinger, 1991; Tufte, 1991]).

Layout of Static Presentations Some interesting early efforts in automating layout include Eastman's work on a *General Space Planner* that addressed the task of arranging objects (e.g., furniture) in a space subject to given constraints (cf. [Barr and Feigenbaum, 1981], Chap. III). Feiner's *GRIDS* (GRaphical Interface Design System, cf. [Feiner, 1988]) was constructed as an rule-based experimental system to investigate approaches in the automatic display layout of text and illustrations. The layout process is guided by the concept of a graphical design grid. The current version of the testbed system has been implemented using an OPS5-like production language. Other approaches using computer-based grids, modeled by a human designer, can be found in the system *VIEW* (cf. [Friedell, 1984]) for synthesizing graphical object depictions from high-level specifications and by [Beach, 1985] for low-level table layout, whose high-level topology was specified by the user as a matrix.

Recent approaches investigate the use of constraint-based and case-based reasoning methods for representing graphical design knowledge. So *Laylab*, WIP's knowledge-based layout manager (cf. [Graf, 1992] and Chap. 4.4), exploits advanced constraint formalisms, such as finite domains and constraint hierarchies for specifying graphical design principles as well as a technique for propagating prioritized constraints to position individual document fragments on an automatic generated design grid. WIP deals with page layout as a rhetorical force, influencing the intentional and attentional state of the reader. In WIP, layout is viewed as an important carrier of meaning. A system that combines both rule-based representation and case-based reasoning in a system that generates and adapts effective layouts of information is the *TYRO* graphics designer developed at MIT (cf. [MacNeil, 1990]). *LIGA* (Layout Intelligence for Graphics Automation, cf. [Colby, 1992]) is a prototype system that generates new layouts by modifying example layouts from its case library. Similar to the approach in WIP, the graphic design knowledge about so-called 'cases' is represented in the system using constraints.

Other systems have tried to drive the automatic generation of the layout of interfaces by incorporating recognized interface standards in the rule-based approach, e.g., the *ITS* system (cf. [Wiecha and Boies, 1990]) employs IBM's Common User Access. The importance of a deeper treatment of multimodal constraints in information presentation in order to address the ergonomic aspects of layout has also been stressed by [Dale, 1992]. An interval logic for reasoning about space, which is based on regions and connection is proposed in [Randell *et al.*, 1992].

Moreover, layout problems are inherent to most configuration tasks, e.g., the configuration of the passenger cabine of an AIRBUS A340 is addressed by constraint processing techniques (cf. [Kopisch and Günter, 1992]), and in [Paaß, 1992] associative methods are used to determine the geometrical arrangement of office furniture.

The importance of the layout dimension is also stressed by recent work at ISI that involves the generation of formatted text exploiting the communicative function of headings, enumerations and footnotes (cf. [Hovy and Arens, 1991]). A similar approach to text layout is followed by WIP's automatic typographer (cf. [Soetopo, 1992]). Here, high-level specifications of relations between textual devices are expressed by constraints which can be compiled into low-level text formatting routines. Other systems in the area of text layout exploit rule-based approaches for formatting text automatically (e.g., [Oemig *et al.*, 1991]).

Layout of Interactive Presentations Since constraint-satisfaction techniques have become more sophisticated during the last decade, and with the growing availability of advanced graphics hardware, there has been an upward trend in applying constraint techniques to user interface design. Thus, most of the related work on applications of constraint languages and systems has been done in the area of computer graphics and graphical interfaces, especially interactive geometric layout (e.g., [Borning and Duisberg, 1986; Kramer *et al.*, 1991]).

A pioneering system in both constraint-based languages and systems and interactive graphics was *Sketchpad* (cf. [Sutherland, 1963]) written by I. Sutherland at MIT in 1963. The Sketchpad system allowed a user to create complex objects by sketching primitive graphical entities and specifying constraints on them. Many of these ideas have been explored by Borning in the *ThingLab* system at Xerox PARC [Borning, 1979; Borning, 1981], a graphical constraint-oriented simulation laboratory implemented in Smalltalk-80. Later versions of ThingLab were concerned with extensions supporting constraint hierarchies, incremental compilation, and graphical facilities for defining new kinds of constraints (e.g. [Borning *et al.*, 1987; Freeman-Benson *et al.*, 1990]). Both systems exploit numeric techniques such as relaxation for solving constraint networks containing cycles, in contrast to symbolic techniques, e.g., used in Steele's constraint language (cf. [Sussman and Steele, 1980]). Further research activities in constraint-based graphics include the systems *Juno* [Nelson, 1985], *IDEAL* [van Wyk, 1982], *Magritte* [Gosling, 1983], *Bertrand* (cf. [Leler, 1988]), and the work of Cohen et al. on constraint-based tiled windows (cf. [Cohen *et al.*, 1986]).

An increasing number of interface-design systems mostly based on a graphical editor, have been developed during the last few years to make the interactive interface design process more efficient and comfortable than with conventional techniques. Here, constraints provide a means of stating layout requirements, e.g., the *Peridot* system deduces constraints automatically as the user demonstrates the desired behaviour (cf. [Myers, 1991b]). This approach has also been extended for text formatting by demonstration (cf. [Myers, 1991a]).

A similar system designed by Kurlander and Feiner (cf. [Kurlander and Feiner, 1991]) is able to infer constraints from multiple snapshots. The *Metamouse* system (cf. [Maulsby *et al.*, 1990]) is a demonstrational interface for graphical editing tasks within a drawing program. The user can specify a procedure by performing an example execution trace, manipulating objects directly on the screen and creating graphical tools. A grid-based approach to specifying simple number independent layouts by example is introduced in [Hudson and Hsi, 1992].

Further work has concentrated on methods for automating the layout of graphs. E.g., the article by Böhringer and Newbery [Böhringer and Paulisch, 1990] details a new way to achieve stability in automatic graph layout. This approach allows a continuum between manual and automatic layout by allowing the user to specify how stable the graph's layout should be through the use of layout constraints. Another approach in interactive graph layout proposes a novel methodology for viewing large graphs (cf. [Henry and Hudson. 1991]). The basic concept is to allow the user to interactively navigate through large graphs and learn about them in concise sections of appropriate size.

Layout of Animation Up to now only rudimentary work has been done in the areas animated layout and layout of presentations including animation. While currently, most animation is laboriously done by hand, *Animus* is one of the first systems that allows for easy construction of an animation with minimal concern for lower-level graphics programming (cf. [Borning and Duisberg, 1986; Duisberg, 1987]). Here temporal constraints are used to describe the appearance and structure of a picture, as well as as how those pictures evolve in time.

Other research in the area of animated layout was concerned with topics like *animation of programs*, and *visual programming* (e.g. [London and Duisberg, 1985; Duisberg, 1990]). In an application of the *Kaleidoscope* language temporal constraints are used to update the display of graphical objects which are manipulated by mouse actions interactively and maintain their consistency requirements (cf. [Freeman-Benson, 1990]).

Other representative research related to the area of automatic graphical layout has concentrated more on the theoretical background of constraint languages and systems (cf. [Leler. 1988]) including weak constraints (cf. [Borning *et al.*, 1987]), constraint logic programming (cf. [Jaffar and Lassez, 1987]) and new inference techniques (cf. [Hentenryck, 1989: Smolka, 1991]). An extended overview of recent work on constraint-based reasoning is given in a special volume of the AI Journal [Freuder and Mackworth, 1992].

3.2 Knowledge Representation and Reasoning

3.2.1 Terminological Logics

Most of today's work on taxonomic reasoning is based on the KL-ONE system, which was an implementation of Brachman's ideas of *structured inheritance networks* [Brachman. 1978]. KL-ONE was built to address the problem of making large knowledge bases comprehensible by forcing them to be constructed in terms of relatively few, well-understood representational operators, the so called epistemological primitives. Semantic networks up to then had suffered from ambiguities and misunderstandings caused by the unclear meaning of their primitives [Woods, 1975].

In order to better meet the characteristic properties, the name for KL-ONE-like systems has been changed from the original term Semantic Networks over Terminological Reasoning Systems, Term Subsumption Languages, Terminological Logics, Concept Languages to the current term *Description Logics.* The paradigm itself that is described by these names, however, remained unchanged: the structural description of classes of individuals–so called *concepts*—and binary relations between them—so called *roles.* The main advantage of these formal descriptions is that a well-defined formal semantics can be given for them and that they can be automatically classified into taxomomic hierarchies according to their generality. The classification process is based on the subsumption relationship and puts new concept descriptions automatically in the "right" place.

Most of the KL-ONE successors like KRYPTON [Brachman *et al.*, 1983; Brachman *et al.*, 1985], KANDOR [Patel-Schneider, 1984], NIKL [Moser, 1983; Schmolze. 1989; Schmolze and Mark, 1991], KL-TWO [Vilain, 1983; Vilain, 1985], LOOM [MacGregor, 1991a; MacGregor, 1991b], BACK [von Luck *et al.*, 1987; Nebel and von Luck. 1988; Peltason, 1991], MESON [Owsnicki-Klewe, 1988], KRIS [Baader and Hollunder. 1991], K-REP [Mays *et al.*, 1988; Mays *et al.*, 1991], SB-ONE [Profitlich, 1989; Profitlich. 1990; Kobsa, 1991a; Kobsa, 1991b], CLASSIC [Borgida *et al.*, 1989; Brachman *et al.*. 1991; Patel-Schneider *et al.*, 1991], or YAK [Cattoni and Franconi, 1990; Franconi, 1991: Franconi *et al.*, 1992], that have been developed up to now, have also comprised a second language to state assertions about instances of concepts and to reason about relations between instances and concepts, which led to yet another name: hybrid reasoning systems.

Brachman and Levesque [Brachman and Levesque, 1984] showed that the desired goal of sound, correct and tractable inferences (esp. subsumption) leads to a trade-off between the expressive power and the computational complexity because, even for small languages, subsumption can be intractable. The far ends of this discussion are taken by the KANDOR system on one side, which supported only a very small language and claimed to have complete algorithms (but see [Nebel, 1988]) and the LOOM system on the other side with a large variety of language constructs and inference mechanisms known to be incomplete.

During the late 80s more and more papers [Levesque and Brachman, 1987; Nebel. 1988: Patel-Schneider, 1987; Patel-Schneider, 1989; Schmidt-Schauß, 1989; Hollunder. 1989] showed that all reasonably expressive languages are intractable. In [Nebel, 1990b] it is shown that terminological reasoning is inherently intractable. The theoretical efforts in the area of pure terminological reasoning have come to an end as the sources of complexity now seem to be determined [Donini *et al.*, 1991a; Donini *et al.*, 1991b; Donini *et al.*. 1992]. This pushes the focus of attention on two old objections (see, e.g., [Doyle and Patil. 1991]). namely, whether classification is the central inference mechanism at all and whether these worst case results are really important for real applications (see, e.g., [FallSymp-92. 1992; MacGregor, 1992; Schaerf, 1992]). A very interesting contribution to the latter question has been published in [Heinsohn *et al.*, 1992a]. In this paper the results of an empirical analysis of six current terminological representation systems with respect to their "normal case" performance is documented.

The major interest of the KL-ONE developers now seems to be concentrated on the design of systems which meet the requirements of their applications by providing reasonably expressive languages, accepting the incompleteness of their algorithms in cases which–following the statements of the developers–do not occur in everyday use. The only exception to this trend is the design of the KRIS system [Baader and Hollunder, 1991], which provides complete subsumption algorithms in combination with a very expressive language (including, e.g., negation and disjunction of concepts). Although an initial empirical evaluation seemed to indicate that such an approach leads to a disappointing preformance [Heinsohn *et al.*, 1992a], a study of optimization techniques [Baader *et al.*, 1992b] showed that the completeness of the algorithms can not be blamed for the bad results. Now a kind of "pay-as-you-go" state seems to be achieved, i.e., the runtimes are roughly proportional to the complexity of the used language constructs.

Another direction on current research on description logics is the integration of other kinds of knowledge, e.g., temporal relations [Schmiedel, 1989; Schmiedel, 1990; Schild, 1991], actions and plans [Devanbu and Litman, 1991; Weida and Litman, 1992; Heinsohn *et al.*, 1992b], concrete domains [Baader and Hanschke, 1991], or defaults and nonmonotonic inferences [Baader and Hollunder, 1992; Quantz and Royer, 1992; Padgham and Nebel. 1992; Patel-Schneider, 1992].

Since description logics have reached a certain maturity, and since a number of systems have been implemented, the KRSS (Knowledge Representation System Specification) group [Neches *et al.*, 1991; Patil *et al.*, 1992] now aims at defining a standard for terminological representation systems.[1] The importance of these efforts has recently been confirmed in practice by observations during the empirical analysis described in [Heinsohn *et al.*, 1992a]. One result of this study was that sharing knowledge between several KL-ONE-alikes (which had been thought to be similar in that they are all based on the same paradigm) requires a surprising amount of effort caused more by differences in the design principles (like, e.g., allowing forward references) than by the differing languages.

A description of KL-ONE's different language constructs can be found in [Schmolze and Woods, 1990] together with the history of the KL-ONE-family starting from the origins of the KL-ONE system itself (see also [Nebel, 1990a]). A study of theoretical aspects of Description Logics is given in [Nebel and Smolka, 1991] and a good survey over the most recent of the various KL-ONE-like systems was provided at the AAAI Spring Symposium 1991 on Implemented Knowledge Representation and Reasoning Systems, where 10 KL-ONE-alikes were represented [Sigart Bulletin, 1991].

3.2.2 Reasoning about Action, Change and Time

Reasoning about action and change is one of the key topics in knowledge representation and AI in general. A large part of this work, however, is concerned with plan generation, a topic we will not address in this section. Instead, we will focus on less ambitious reasoning tasks such as *predicting* the outcome of the execution of a set of actions, *explaining* a result by hypothesizing that some actions have taken place, or *recognizing* a plan that is carried out by an agent.

The problem of predicting the result of the execution of actions, often called *temporal projection*, has been intensively studied, since it presents severe problems from a

[1]Part of this standard is a specification that has been developed by DFKI researchers [Baader *et al.*, 1990].

logical point of view. The *frame problem*, i.e., the problem to compute what is unchanged by an action, identified by McCarthy and Hayes [McCarthy and Hayes, 1969] played a key role in this context. In fact, most of the research is centered around the problem of how to "solve" the frame problem using nonmonotonic logics. While it was originally believed that nonmonotonic logics are suitable for solving this problem (see, e.g., [Reiter, 1980]), the paper by Hanks and McDermott [Hanks and McDermott, 1987] demonstrated that this is not the case. This negative result applies not only to the usually employed *situation calculus* but to all temporal represenation languages and logics and is independent of a particular nonmontonic logic used to specify the frame default. While there have been a number of proposals to account for the problem identified by Hanks and McDermott (e.g., [Shoham, 1986; Kautz, 1986; Morgenstern and Stein, 1988; Sandewall, 1989], more principled approaches to address the problem have, only recently, been developed [Lifschitz, 1991; Lin and Shoham, 1991; Sandewall, 1992a; Sandewall, 1992b].

Other approaches to cope with the frame-problem have been to use a procedural "update semantics" avoiding the problem altogether. Most prominently, the STRIPS-framework [Fikes and Nilsson, 1971] should be mentioned in this context. Although this approach avoids the frame problem, it has the disadvantage of not including time and the flow of actions into model, handling this only on the meta-level. Further, as pointed out by Lifschitz, one must be very careful when modeling actions and plans using STRIPS [Lifschitz, 1986].

Nevertheless, the STRIPS approach of modeling action and change has been very popular, in the area of planning in particular, because of its conceptual simplicity. Recently, the STRIPS model has been extended to allow for a richer modelling bringing back again the frame-problem, which is handled procedurally on the meta-level, though [Ginsberg and Smith, 1988a; Ginsberg and Smith, 1988b; Winslett, 1988; Katsuno and Mendelzon, 1991].

One elegant way to deal with the frame problem is provided by a logic programming approach to maintaining events and their effects over time—the *event calculus* [Kowalski and Sergot, 1986]. It should be noted, however, that the expressiveness of the event calculus and its nonmonotonic reasoning techniques is not comparable to the above mentioned work.

Even abstracting from the logical problems and assuming a simplified model of propositional STRIPS, there are considerable computational problems. If context-dependent effects are allowed or the ordering of the actions is only partial, temporal projection even for propositional STRIPS is intractable [Chapman, 1987; Dean and Boddy, 1988]. As shown in [Nebel and Bäckström, 1992a; Nebel and Bäckström, 1992b], however, projection over partially ordered, context-independent actions is tractable, provided a realistic execution model is assumed.

Most of the research in reasoning about action, including the work described above, assumes that actions do not occur simultaneously and there are only few approaches that go beyond this assumption (see e.g. [Große, 1992a]). In general, one can distinguish between approaches that require more or less independence of simultaneously executed actions (e.g., [Horz, 1992]), approaches that handle additional synergistic positive effects of simultaneously occuring actions (e.g., [Große and Waldinger, 1991;

Große, 1992b]), approaches that require simultaneous occurence for successful execution [Sandewall and Rönnquist, 1986; Bäckström, 1988; Allen, 1991], and approaches that permit the suppression of some effects in case of simultaneous execution [Lin and Shoham, 1992]. In short, the problem does not yet seem to be well understood. On the contrary, there seem to be a long way to go before the different perspectives on this problem will converge.

When representing simultaneously occuring event, one problem is to specify the exact or relative order (including overlap) of occurences. This, of course, is not possible in the *situation calculus* since actions happen in unit time [Gelfond *et al.*, 1991]. Hence, if actions and events are required to have duration, another formalism must be employed. One possibility is to use Allen's interval algebra [Allen, 1983] to specify the occurences of events (see also [Allen, 1991]), a formalism that also has attractive computational properties, at least for a reasonable subset [Nökel, 1989; van Beek, 1990; van Beek and Cohen, 1990] of the algebra, which is intractable in general [Vilain *et al.*, 1989], however. In addition, implemented systems are available that support reasoning with interval relations (e.g., the MATS system [Kautz and Ladkin, 1991]).

Until now only few attempts have been made to extend terminological representation systems in order to handle additional kinds of knowledge, e.g., temporal or causal relationships.

Three approaches to represent actions and plans that are similar in some aspects of their architecture are CLASP [Devanbu and Litman, 1991], T-REX [Weida and Litman, 1992], and RAT [Heinsohn *et al.*, 1992b]. They all use a terminological logic to represent the world states and atomic actions and add a second formalism to compose plans and reason about the temporal relationships. Caused by different requirements and objectives, however, the focus of the RAT system and with that the design of its language is different from that of CLASP and T-REX. Whereas in RAT the states to express preconditions and effects of actions can be *described* using a subset of the terminological language, actions and states in CLASP and T-REX are primitive, non-decomposable units. On the other hand, their language to compose plans goes beyond the linear sequences supported in RAT. CLASP provides regular expressions over actions (incl. conditionals, loops and disjunctions), T-REX uses Allen's temporal constraints [Allen, 1983] to construct plans. The inference services of the three systems are also determined by their applications: CLASP and T-REX both use the computed plan hierarchies mainly to support plan recognition tasks whereas RAT's services comprise consistency checks, simulated execution of plans and temporal projection of conditions.

Another approach has been made by Swartout and Neches [Swartout and Neches, 1986], who classified and retrieved plans according to their goal descriptions, which are formulated in a terminological logic. However, they made no attempt to represent plans in the terminological logic. Wellman [Wellman, 1990] also builds plans from actions represented in a terminological logic, and organizes plan classes into a hierarchy based on a notion of subsumption. His language, however, is completely atemporal and he does not reason about plan individuals but plan classes only.

A plan abstraction hierarchy is central to the plan recognition work of Kautz [Kautz, 1991]. However, in his taxonomy, plan nodes are still atomic rather than structural, and the suitability of the representation for computing terminological inferences is not of concern.

In the field of plan synthesis, Tenenberg [Tenenberg, 1989] uses a plan hierarchy to construct abstract plan solutions that restrict later search, where any abstract solution can always be specialized by choosing a specialization of each abstract plan step. Thus, while plans in Tenenberg's hierarchy are compositions of actions, plans must always be structurally isomorphic across abstraction levels.

Finally, there are approaches that try to integrate the notions of time with terminological logics by extending a terminological logic by a temporal logic [Schmiedel, 1990; Schild, 1991]. However, these approaches have been only theoretical efforts so far, and it is not clear in how far (reasonably efficient) systems based on these theoretical efforts can be built.

3.2.3 Reasoning about Uncertain Knowledge

For the application underlying the $\mathcal{PPP}$ project the *inherent uncertainty* of several kinds of knowledge (about the user, presentation task and plan, and layout planning, for instance) is one important feature that has to be taken into account. Apart from the question which uncertainty model is appropriate for which kind of uncertainty phenomenon, another important question is how to represent the uncertainty of knowledge in a uniform framework and how to perform inferences in the case of uncertainty. Such inferences may, for instance, allow for decision making if several non-conflicting solutions exist, support presentation planning in the case of uncertain knowledge, and allow to generate explanations in a (weighted) abductive manner, for instance, if the interaction with a user leads to an interruption of presentation planning.

Prompted by this application, we sketch below the state of the art concerning (i) numerical models for handling *uncertainty*, (ii) handling of *uncertainty in terminological logics*, and (iii) *planning and abduction under uncertainty*. While for the first item several models for handling uncertainty have been proposed and their foundations, advantages, and disadvantages are well explored now (see, e.g., [Kruse *et al.*, 1991; Shafer and Pearl, 1990] for overview and analyses) the examination of the influence of uncertain knowledge for neighbouring fields, such as knowledge representation in general and knowledge representation in (termino)logical formalisms in particular, planning, and abduction, etc. has started only recently.

Uncertainty Models The characteristic feature of *heuristic uncertainty models* is that their mathematical foundations are traced only partially or not at all to some sound theory, as given by probability theory, for instance. This is because heuristic approaches aim at avoiding certain "problems" arising from the use of, e.g., probability theory. The reasons that are often mentioned in this context are the amount of data needed (prior and conditional probabilities, joint probability distributions, etc.), the inability to distinguish between absence of belief and doubt, and the fact that it is impossible to represent ignorance. One of the most important (heuristic) uncertainty models that aim at solving these problems is the *certainty factor approach* developed by Shortliffe et al. [Shortliffe and Buchanan, 1975; Shortliffe, 1976; Buchanan and Shortliffe, 1984]. The certainty factor model has to be seen in relation to the development of the well known expert system

MYCIN which was built during the years 1972–1976 and is an expert system for advising physicians on how to treat patients suffering from bacteriogenous infectious diseases. Later systems related to MYCIN are EMYCIN [van Melle, 1980], a domain-independent system based on MYCIN's control mechanisms and data structures, and RMYCIN [Cendrowska and Bramer, 1984] which is a reconstruction of MYCIN. Since the certainty factor approach makes use of measures and algorithms that are heuristic and at most "syntactically similar" to probabilistic ones, it has often been criticized (e.g. in [Adams, 1976; Heckerman, 1986; Horvitz *et al.*, 1986; Horvitz and Heckerman, 1986]). Another heuristic model is based on the concept of *triangular norms* and *conorms* [Schweizer and Sklar, 1961; Schweizer and Sklar, 1983]. It is important to notice that, because of their generality, T-norms and T-conorms give an infinite number of different "calculi of uncertainty." The selection of a special pair describes a particular calculus uniquely and completely. An application of this general uncertainty model is the expert system RUM ("Reasoning with Uncertainty Module") [Bonissone and Wood, 1989; Bonissone, 1987; Bonissone *et al.*, 1987]. They are also discussed in [Smets and Magrez, 1987; Magrez and Smets, 1989]. The system INFERNO has been developed by Quinlan [Quinlan, 1983a; Quinlan, 1983b; Quinlan, 1985]. One characteristic feature of INFERNO's architecture is that the inference model is mainly based on *bounds propagation*: it can be used for both forward and backward inferences. Since the model does not make assumptions about (in)dependencies of data, all the propagation constraints can be proven to be correct. Because of this philosophy, and the resulting fact that computed bounds may sometimes be weak, INFERNO is also called a "cautious" approach to uncertain inference. Meanwhile, several modifications and improvements have been proposed ([Liu and Gammerman, 1988; Saunders, 1989]).

As often argued *probability theory* offers a theoretically sound model for representing uncertainty and for embedding it in reasoning techniques. Just in the last few years a revival of using the probability theory in representing uncertainty has taken place, giving considerable insight into the application of probability theory and pointing out some misconceptions about its applicability. Also, new theoretical results from statistics and probability theory present arguments for the utility of probabilities for reasoning. Simple early models, that may be viewed as *straightforward approaches* for making use of probabilities in rule-based uncertain reasoning, are introduced in [Ishizuka *et al.*, 1982] and in [Adams, 1976]. In the approach of "inference networks" [Duda *et al.*, 1976; Duda *et al.*, 1978; Duda *et al.*, 1981] expert rules are interpreted as directed links labelled with so-called *likelihood ratios* based on a probabilistic interpretation. A concrete expert system based on inference networks is the system PROSPECTOR. Discussions are mainly related to the restrictive independence assumptions [Glymour, 1985; Johnson, 1986; Steve, 1986]. A promising approach is that of "decomposable graphic models", also called *belief networks* (see, e.g., [Pearl, 1986; Pearl, 1988; Spiegelhalter, 1989]). A characteristic feature of this approach is that uncertainty and belief is propagated through a network by local operations only: each node of the network is viewed as a single processor which exchanges messages with its neighbor nodes. A prototype expert system based on this model is the system MUNIN [Andreassen *et al.*, 1987]. Based on this system the expert system shell HUGIN [Andersen *et al.*, 1989] has been developed. Also the system PATHFINDER [Heckerman, 1990] is based on the idea of belief networks. Analyses of decomposable graphic models and exhaustive references are also given in [Kruse *et al.*, 1991].

Like probabilistic approaches, the *Dempster-Shafer theory of evidence* aims to model and quantify uncertainty by *degrees of belief*. But, in contrast to probabilistic approaches, it permits assignment of degrees of belief to sets of hypotheses rather than to hypotheses in isolation. The underlying idea is that the process of narrowing the hypothesis set with the collection of evidence is better represented in terms of this theory than in terms of probabilistic approaches. For this reason the theory can be viewed as an alternative to probability theory. The classical approach to evidence theory has been proposed by Shafer [Shafer, 1976]. His mathematical model is essentially based on the notion of *belief functions* and *Dempster's rule of combination* [Dempster, 1967]. The ability to express (total) ignorance is one of the main features that has to be mentioned as an advantage of belief functions against the use of a single probability. The Dempster-Shafer theory of evidence has received wide attention since the "7th IJCAI" in 1981 where the three papers [Barnett, 1981; Friedman, 1981; Garvey *et al.*, 1981] considering aspects on the Dempster-Shafer theory were presented. In early proposals for performing uncertain reasoning the knowledge of experts is mainly represented in the form of explicit expert rules (see, e.g., [Ginsberg, 1984; Ishizuka *et al.*, 1982]). A characteristic feature of the Dempster-Shafer theory of evidence is that—because of being based on power sets—the hypothesis space can be *hierarchically ordered*. The elements of a power set can be related together by making use of subset and superset operators. In particular, Gordon and Shortliffe's extension to MYCIN's certainty factor approach [Gordon and Shortliffe, 1984; Gordon and Shortliffe, 1985] and Yen's "quasi-probabilistic" extension [Yen, 1989] to the Dempster-Shafer theory have to be mentioned in this framework. A characteristic feature of the proposals given in [Shafer *et al.*, 1987; Shenoy and Shafer, 1987] is that explicit expert rules do not appear. Instead, the knowledge on dependencies is represented by the *links of a network*—the underlying idea is similar to Pearl's "Bayesian networks". A concrete implementation of this is the system MacEvidence [Kruse *et al.*, 1991]. Notice that necessity (possibility) measures [Zadeh, 1978; Dubois and Prade, 1988; Yager, 1980] are special belief (plausibility) functions. Interesting discussions about the Dempster-Shafer theory can also be found in [Zadeh, 1984; Zadeh, 1986; Ruspini, 1987]. In [Halpern and Fagin, 1990] the relations between probability and belief are discussed and the difference between conditioning and update is made visible (see also [Kyburg, 1987]).

For more details and other uncertainty models the reader is referred to the books [Kruse *et al.*, 1991; Hajek *et al.*, 1992; Pearl, 1988] and the collections [Kruse and Siegel, 1991; Bonissone *et al.*, 1991], for instance.

Uncertainty and (Termino)Logical Approaches Apart from general models for handling uncertainty, recent work is also related to the role of uncertainty in *nonstandard logics* (see, e.g., [Smets *et al.*, 1988]) and to *extensions of logics* with respect to uncertainty in general. The best-considered approach in the second area is the integration of predicate logic with probability theory called *probabilistic logic* [Nilsson, 1986; Paaß, 1988]. The aim of such a combination is achieved by interpreting logical formulas as subsets of elementary events referred to as sets of possible worlds. Another kind of extension is related to the integration of possibility theory and predicate logic [Dubois and Prade, 1988; Dubois and Prade, 1991] in order to obtain a *possibilistic logic*. In this framework, the work of Bacchus [Bacchus, 1989; Bacchus, 1990; Bacchus *et al.*, 1992] is also important because he not only explores the question of how far one can go using *statistical* knowledge, but also presents LP, a logical formalism for representing and reasoning with statistical knowledge. However, in spite of providing a very powerful representational formalism, Bacchus does not offer a deep discussion of consistency requirements and inference mechanisms. There exist several other proposals which are, however, outside the scope of this section.

As *terminological formalisms* play an important role in knowledge representation and reasoning in general, as well as in the framework of this project in particular, these formalisms have to be extended w.r.t. handling of incomplete and uncertain knowledge. The importance of providing an integration of both term classification and uncertainty representation[2] was recently emphasized in some publications. Yen and Bonissone [Yen and Bonissone, 1990] consider this integration from a general point of view which, for instance, does not require a concrete uncertainty model (e.g., probabilistic, fuzzy, Dempster-Shafer). In [Yen, 1991] Yen proposes an extension of term subsumption languages to fuzzy logic [Zadeh, 1965] that aims at representing and handling vague concepts. His approach generalizes a subsumption test algorithm for dealing with the notion of vagueness and imprecision. Since our application is mainly influenced by the existence of uncertainty, our general objectives differ from those underlying Yen's proposal. Saffiotti [Saffiotti, 1990] presents a hybrid framework for representing epistemic uncertainty. His extension allows one to model uncertainty about categorical knowledge, e.g., to express one's belief on quantified statements such as "I am fairly (80%) sure that all birds fly". Note the difference from "I am sure that 80% of birds fly", which requires a completely different formalism. In [Heinsohn, 1991a; Heinsohn, 1992; Heinsohn, 1991b] a probabilistic extension of terminological logics is proposed that maintains the original performance of terminological logics of drawing inferences in a hierarchy of terminological definitions. It, however, enlarges the range of applicability to real world domains determined not only by definitional, but also by uncertain knowledge (arising with, e.g., "typical" properties) which can be modeled on the basis of the language construct "probabilistic implication". To guarantee (terminological and probabilistic) consistency, several requirements have to be met. Moreover, these requirements allow implicitly existent (probabilistic) relationships, including knowledge about exceptions, to be inferred.

Planning and Abduction under Uncertainty As argued in [André and Rist, 1992], for the automatic generation of illustrated documents a plan-based approach is adequate. However, because of the presence of several kinds of incompleteness and uncertainty also the inferences that allow to *reason about plans* have to cope with these phenomena. Below, we give an overview of the work that has recently been done in the areas of *planning under uncertainty* and *(weighted) abductive techniques*. In the framework of the $\mathcal{PPP}$ project *abduction* can be characterized as a method for finding, for-instance, the "best" explanation for an interaction that has been performed by a user during presentation.

For the task of *abduction in the presence of uncertain knowledge*, proposals have recently been made by Appelt and Pollack [Appelt and Pollack, 1990], Charniak and Shimony [Charniak and Shimony, 1990], Poole [Poole, 1991; Poole, 1988], and Peng and Reggia [Peng and Reggia, 1990], for instance. In their approach of "weighted abduction" Appelt

[2]Brachman [Brachman, 1990] considers "probability and statistics" as one of the "potential highlights" in knowledge representation.

and Pollack assign weighting factors to all literals in the premise of a rule for being able to single out the "best" hypotheses. These factors are used to compute the assumption cost of literals, and in the abductive procedure the assumption set with the lowest cost is preferred. Similarly, the model of Charniak and Shimony is based on a probabilistic semantics for cost-based abduction. The basis of Poole's work is the examination of (default) logics and non-monotonic formalisms in the framework of abduction. In [Poole, 1991] he presents a framework for Horn-clause abduction, including probabilities associated with hypotheses. His main contribution is in finding a relationship between logical and probabilistic notions of evidential reasoning. Peng and Reggia [Peng and Reggia, 1990] (see [Thagard, 1991] for a book review) consider abduction as the generation and comparative evaluation of explanations for a set of facts. Apart from analyses they provide a computational theory of abductive inference in medical diagnosis. It is important to mention that abductive techniques are also inherently present in some numerical methods for handling uncertainty such as belief networks introduced by Pearl. Also there exists a close relationship between the incompleteness and uncertainty of knowledge and the non-monotonicity in reasoning. For an overview about abduction in a *logical view* and respective references we recommend [Merziger, 1992]. Complexity analyses of abduction can be found in [Bylander *et al.*, 1991].

The central element of the book [Wellman, 1990] is the formulation of tradeoffs in *planning under uncertainty*. In particular, Wellman presents his SUDO-Planner, a program that formulates tradeoffs by constructing decision models from a multilevel knowledge-base of qualitative relations. A language for planning with statistics is provided by Martin and Allen. The paper [Martin and Allen, 1991] combines Allen's temporal interval reasoning with statistical inference to facilitate planning using inferences about probabilities. An overview about the recent state of the art in "reasoning about plans" is given in [Allen *et al.*, 1991b]. Shanahan [Shanahan, 1989] analyses the relations between deductive and abductive techniques.

An analysis of the proposals made in the area of planning and abduction under uncertainty visualizes that it is necessary to clarify in general the relationships and conceptual differences of numerical and logical approaches for both abduction and planning in order to provide an appropriate integrated formalism which aims at supporting presentation planning and explanation finding in the $\mathcal{PPP}$ project.

3.2.4 Efficient Inference Mechanisms

While the main problems in designing knowledge representation and reasoning systems to support multimedia presentation systems are of conceptual nature, considerations of efficiency cannot be completely ignored. Indeed, most reasoning services that are needed to support, for instance, multimedia presentations are computationally intractable in the worst case. Nevertheless, it is necessary to guarantee some level of performance for the cases that occur in practice. This issue has been recently recognized as important, as demonstrated by explicit sessions on computational complexity and tractability in AI conferences and by workshops on this topic, for instance, the AAAI'92 workshop on tractability [AAAI-WS, 1992] and the upcoming AAAI 93'Spring Symposium on NP-hard problems in AI.

A quite detailed investigation of reasoning with terminological logics has shown that this kind of reasoning is intractable for all reasonably powerful terminological logics [Donini *et al.*, 1991a; Donini *et al.*, 1991b]. Worse still, restricting the logic to be of "minimal" expressivity leads to a NP-hard inference problem provided we allow for the definitions of new concepts—something which is supported in all implemented representation systems supporting terminological logics [Nebel, 1990b].

Turning to temporal reasoning, a similar picture evolves. Planning is, of course, a kind of reasoning that is quite difficult as has been recently shown for a number of different planning models [Bylander, 1991a; Bylander, 1992; Erol *et al.*, 1992; Gupta and Nau, 1992] and severe restrictions on the quality of the solution and/or the allowable forms of action rules are necessary to guarantee tractability [Bylander, 1991a; Bäckström and Nebel, 1992a; Selman, 1992]. However, even less ambitious modes of temporal reasoning such as computing the implied ordering of events given a description in terms of interval relations following [Allen, 1983] is an NP-complete problem [Vilain *et al.*, 1989]. Also computing consequences of actions in a comparably simple setting or validating a given plan is a difficult problem [Chapman, 1987; Dean and Boddy, 1987; Dean and Boddy, 1988; Nebel and Bäckström, 1992a; Bäckström and Nebel, 1992b].

Finally, in considering reasoning about beliefs, it is a well-known fact that most propositional logics of beliefs have inference problems that are harder than reasoning in ordinary propositional calculus [Garey and Johnson, 1979; Halpern and Moses, 1992], which is already an NP-complete problem. Similarly, propositional abduction is a problem that is more difficult than ordinary propositional reasoning [Selman and Levesque, 1990].

Although the above results may be considered as very discouraging, from a more practical point of view they only indicate that it is impossible to come up with algorithms that are efficient in *all* cases. However, they do not rule out methods that give satisficing answers in almost all practical cases. While it would be, of course, desirable to have provably efficient reasoning methods [Brachman and Levesque, 1984], this goal is simply not achievable in most cases. For this reason, formal computational complexity results are often considered as non-informative and irrelevant for practical purposes. It should be noted, however, that complexity results provide us with insights into the computational structure of a problem that can guide us in developing efficient reasoning methods, for instance, by concentrating on special cases.

Efficient reasoning methods for worst-case intractable problems sometimes rely on the fact that in practice the full expressiveness of a representation formalism is not needed, leading to the situation that worst cases hardly occur. A prototypical example is the above mentioned problem of terminological reasoning in the presence of concept definitions [Nebel, 1990b]. Although the problem is worst-case intractable, algorithms that are usually employed in implemented systems do not encounter problems in practice.

This is however a rather unusual situation. Most of the time, some trade-offs along the line of reasoning accuracy or expressiveness have to be made—which may allow for a satisfying overall behavior. Although it is often difficult to restrict the expressiveness, sometimes non-trivial special cases that are expressive enough for a given application can be solved efficiently. For instance, while the temporal projection problem is intractable if events are conditional and are partially ordered, restricting events to be unconditional leads to a polynomial problem [Nebel and Bäckström, 1992a]. More generally, it is often possible to identify parameters of the problem that can be used in order to give a reasonable accurate characterization of the difficulty of given problem instances [Cheeseman *et al.*, 1991; Mitchell *et al.*, 1992].

Approaches that trade efficiency for accuracy include approximation methods (e.g., [Dean and Boddy, 1988; Selman and Kautz, 1991; Bylander, 1991b; Cadoli and Schaerf, 1992]) that compute upper and lower bounds on the result, and sound and efficient but in general incomplete methods (e.g., [Allen, 1983; van Beek, 1990; van Beek and Cohen, 1990]). The main problem with these two methods is a precise characterization of the inference capabilities of the system, a problem that can be solved by specifying, for example, a set of inference rules that lead to a provably polynomial inference problem [Givan and McAllester, 1992].

Further, probabilistic methods have been recently employed in dealing with problems that do not seem to be *apriori* well-suited to be solved by such methods. In particular, the probabilistic propositional satisfiability procedure GSAT [Selman *et al.*, 1992; Selman and Kautz, 1992] seems to be a promising tool in dealing with other similar NP-complete problems. For instance, terminological inferences in logics with an NP-complete inference problem seem to be an interesting candidate in this context. Also planning problems have been tried to solve using probabilistic methods with quite promising results [Langley, 1992; Kautz and Selman, 1992; Minton *et al.*, 1992].

Finally, it should be noted that the overall performance of a system does not only depend on reasoning methods that are efficient "in principle," i.e. that run in most practical cases in *polynomial time*. It is also necessary to use sophisticated implementation techniques and algorithms to achieve *practical efficiency*. An empirical investigation of terminological representation systems, for instance, revealed that there are drastic performance differences between different terminological representation systems [Heinsohn *et al.*, 1992a]. By carefully tuning one of the slowest systems using advanced algorithms and implementation techniques [Baader *et al.*, 1992b] it was shown that these differences were caused by parts of the system that were not worst-case intractable.

3.2.5 Retrieval of Multimedia Units

In addition to generating multimedia presentations from scratch, we also aim at retrieval and use of existing presentations and incoopertation of canned multimedia units. One prerequisite for this is a facility that supports storing, indexing, and retrieval of such multimedia units. For this purpose, we intend to use terminological logics since they provide a flexible description language and a powerful retrieval mechanism exploiting the classification inference [Nebel and Peltason, 1991; Beck *et al.*, 1989].

Similar approaches have been made in the area of managing information about multimedia units and software. In the multimedia system ALFresco [Stock, 1991], for instance, the terminological knowledge representation system YAK [Franconi *et al.*, 1992] is employed for representing domain knowledge as well as knowledge about video clips and pictures. This uniform representation allows to use the YAK system for the interpretation of user requests, for retrieval of canned multimedia units, and for explaining parts of pictures.

In the area of software information, the system that comes closest to our approach is LaSSIE (Large Software System Information Environment) [Devanbu *et al.*, 1991], which is based on the knowledge representation language KANDOR [Patel-Schneider, 1984]. LaSSIE incorporates a large knowledge base, a semantic retrieval algorithm based on the classification inference, and a powerful user interface incorporating a graphical browser and a natural language parser. LaSSIE primarily intends to process queries about actions and, on this basis, to help programmers to find useful information about a large software system. One basic observation underlying the development of LaSSIE is that a developer whose task it is to implement, modify, or add a special operation to the system often cannot determine if it has already been done. Given this difficulty, programmers, instead of reusing existing primitives, often reimplement. Thus, a library of reusable parts is required, along with a helpful access mechanism. Although the application domain of LaSSIE is software, the general principles used in this system seem to applicable to the task of multimedia management and retrieval.

The development of LaSSIE is partially based on the work that led to the ARGON system [Patel-Schneider *et al.*, 1984]. ARGON is an information retrieval system designed for use by non-expert users on heterogeneous knowledge bases. It assists users in retrieving information from its knowledge base by continually presenting a query and an example individual that satisfies the query. Similar to LaSSIE, ARGON stores information in the frame-based knowledge representation system KANDOR.

A model for retrieving software components for possible reuse that employes semantic nets or taxonomic knowledge representation is also proposed in [Prieto-Diaz and Freeman, 1987]. Prieto-Diaz and Freeman describe a taxonomic domain model for the set of data operations embodied in a library of software components, categorized along different facets. This domain model is used in query formulation/reformulation. There are some other semantic-net-based systems that are intended to be general-purpose software librarians: The AIRS system of Ostertag and Hendler [Ostertag and Hendler, 1987] employes a heuristic retrieval algorithm based on a numerical conceptual-distance measure that has to be specified by the user. Woods and Somerville [Woods and Somerville, 1988] use conceptual dependency diagrams, the associated query mechanism is based on a set of verbs. For a given verb the conceptual dependency graph is identified and by prompting for further information it is possible to further narrow the search for components. Beside LaSSIE and ARGON, none of these systems use a classification-style inference, however.

While the LaSSIE approach seems to be the most promising, LaSSIE also has some limititation. These, however, are mainly caused by the limitations of the underlying representation language KANDOR. For instance, *actions and plans* cannot be adequately expressed and handled within its representational framework. Further, due to certain expressive limitations made in KANDOR to make the classification algorithm faster and easier to implement, KANDOR is too weak to describe relationships between fillers of roles.

Retrieval, modification, and reuse of knowledge strutcures have been also considered at a recent AAAI Spring Symposium [SpringSymp-92, 1992]. Franke [Franke, 1992] argues that information regarding structure and behavior of a mechanism is readily captured in current design support systems and methodologies. While structure and behavior descriptions can be used to index design modifications, Franke claims that a more productive classification of design modifications for explanation and reuse is achieved via descriptions of purpose of these design modifications. A respective classification technique is introduced which partially orders the space of behavior abstractions. The overall goal of Johnson and Feather [Johnson and Feather, 1992] is to support the evolution of requirements and specifications for hardware and systems. In the proposed system ARIES the knowledge is organized into specialization hierarchies, folders and domains in order to facilitate reuse. The topic of reuse has also been discussed in the framework of *software engineering* (see, e.g., [Biggerstaff and Perlis, 1989a; Biggerstaff and Perlis, 1989b]).

In particular, the reuse of plans has been considered recently as an interesting research topic. Plan generation in complex domains is normally a resource and time consuming process. One way to improve the efficiency of planning systems is to avoid the repetition of planning effort whenever possible. For instance, in situations when the goal specification is changed during plan execution or when execution time failures happen, it seems more reasonable to *modify* the existing plan than to plan from scratch again. In the extreme, one might go as far as basing the entire planning process on plan modification, a method that could be called *planning from second principles*.

Instead of generating a plan from scratch, that method tries to exploit knowledge stored in previously generated plans. The current problem instance is used to find a plan in a plan library that—perhaps after some modifications—can be used to solve the problem instance at hand. Current approaches try to integrate methods from analogical or case-based reasoning to achieve a higher efficiency [Hammond, 1990; Veloso, 1992], integrate domain-dependent heuristics [Howe, 1992] or investigate reuse in the general context of deductive planning [Koehler, 1992; Biundo *et al.*, 1992]. The range of applicability for such techniques has not been investigated yet, though [Nebel and Koehler, 1992].

4 Previous Work of the Project Team

We have been engaged in work in the area of multimodal communication for several years now, starting with the HAM-ANS ([Wahlster *et al.*, 1983]) and VITRA systems ([André *et al.*, 1986], [Herzog *et al.*, 1989]), which automatically create natural language descriptions of pictures and image sequences shown on the screen. These projects resulted in a better understanding of how perception interacts with language production. Furthermore, we have been investigating ways of integrating tactile pointing with natural language understanding and generation in the XTRA project ([Kobsa *et al.*, 1986], [Wahlster, 1991]).

Our work on knowledge representation draws heavily on the experience gained in designing knowledge representation tools in the project HAM-ANS [Marburger and Nebel, 1983] and on designing of and working with terminological representation systems in the KIT-BACK project [Nebel and von Luck, 1988], in the JANUS project [Sondheimer and Nebel, 1986], in the XTRA project [Profitlich, 1990], and in the MESON project [Heinsohn and Owsnicki-Klewe, 1988]. In WIP, this experience was used to provide support in the area of knowledge representation in the form of adapting and enhancing existing tools and designing and implementing a system that supports representation of and reasoning about actions and plans.

Since 1989, we have been concerned with the coordination of text and graphics in the WIP project. Today, WIP is considered to be one of the leading projects in the area of multimodal presentation systems. This is reflected by numerous publications (among others two chapters in the first volume on intelligent multimedia interfaces and two articles in the AI Journal) and invited talks at major conferences and workshops. The knowledge representation work of the WIP group is internationally recognized as being very significant as can be seen from the group's two books, a number of publications at international conferences and in scientific journals, and invited talks at conferences and workshops.

In WIP, we have developed a computational model for the generation of multimodal communications (cf. [Wahlster *et al.*, 1991; Wahlster *et al.*, 1992a; André *et al.*, 1992; Wahlster *et al.*, 1992b]). The basic principles underlying the WIP project are that the generation of all constituents of a multimodal presentation should start from a common representation and that the design of a text-picture sequence can be modeled as a non-monotonic planning process.

4.1 Presentation Planning

For the automatic synthesis of illustrated documents, we have designed presentation strategies that refer to both text and picture production. In order to decide between several applicable presentation strategies, we have examined how the kind of information to be conveyed influences mode selection and which communicative functions single document parts play in text-picture combinations. In particular, we have shown that most semantic and pragmatic relationships which have been proposed for describing the structure of texts can be generalized in such a way that they are also appropriate for describing the structure of pictures and text-picture combinations. To represent the presentation strategies, we followed the approach proposed by Moore and colleagues (cf. [Moore and Paris, 1989]) to operationalize RST-theory for text planning.

For the automatic generation of illustrated documents, the presentation strategies have been treated as operators of a planning system (cf. [André and Rist, 1990b], [André and Rist, 1990a], [André and Rist, 1992]). The presentation planner receives as input a formal specification of a presentation goal. The result of the presentation planning process is a hierarchically structured plan of the document to be generated. This plan reflects the propositional contents of the potential document parts, the intentional goals behind the parts as well as rhetorical relationships between them. While the top of the presentation plan is a more or less complex presentation goal (e.g., introducing an object or explaining how to make coffee), the lowest level is formed by specifications of elementary presentation tasks (e.g., formulating a request or depicting an object) that are directly forwarded to the mode-specific design components.

4.2 Graphics Generation

When generating illustrations of physical objects, WIP does not rely on previously authored picture fragments or predefined icons stored in the knowledge base. Rather, we start from a hybrid object representation which includes a wireframe model for each object. Although these wireframe models, along with a specification of physical attributes, such as surface color or transparency, form the basic input of the graphics generator, the design of illustrations is regarded as a knowledge-intensive process that exploits various knowledge sources to efficiently achieve a given presentation goal. For example, when a picture of an object is requested, we have to determine an appropriate perspective in a context-sensitive way (cf. [Rist and André, 1990]). In our approach, we distinguish between three basic types of graphical techniques. First, there are techniques to create and manipulate a 3D object configuration that serves as the subject of the picture. For example, we have developed a technique to spatially separate the parts of an object in order to construct an exploded view. Second, we can choose among several techniques that map the 3D subject onto its depiction. For example, we can construct either a schematic line drawing or a more realistic looking picture using rendering techniques. The third kind of technique operates on the picture level. For example, an object depiction may be annotated with a label, or picture parts may be colored in order to emphasize them. The task of the graphics designer is then to select and combine these graphical techniques according to the presentation goal (cf. [Rist and André, 1992b], [Rist and André, 1992a]). The result is a so-called design plan which can be transformed into executable instructions of the graphics realization component. This component relies on the 3D graphics package S-Geometry and the 2D graphics software of the Symbolics window system.

4.3 Text Generation

WIP's text generator is based on the formalism of tree adjoining grammars (TAGs). In particular, lexicalized TAGs with unification are used for the incremental verbalization of logical forms produced by the presentation planner (cf. [Harbusch, 1990; Schauder, 1990; Harbusch *et al.*, 1991; Finkler and Schauder, 1992]). The grammar is divided into an LD (local dominance) and an LP (linear precedence) part so that the piecewise construction of syntactic constituents is separated from their linearization according to word order rules (cf. [Finkler and Neumann, 1989]).

The text generator uses a TAG parser in a local anticipation feedback loop (cf. [Jameson and Wahlster, 1982]). The generator and parser form a bidirectional system, i.e., both processes are based on the same TAG. By parsing a planned utterance, the generator makes sure that it does not contain unintended structural ambiguities.

As the TAG-based generator is used in designing illustrated documents, it has to generate not only complete sentences, but also sentence fragments such as NPs, PPs, or VPs, e.g., for figure captions, section headings, picture annotations, or itemized lists. Given that capability and the incrementality of the generation process, it becomes possible to interleave generation with parsing in order to check for ambiguities as soon as possible. We have explored different domains of locality for such feedback loops and trying to relate them to resource limitations specified in WIP's generation parameters. One parameter of the generation process in the current implementation is the number of adjoinings allowed in a sentence. This parameter can be used by the presentation planner to control the syntactic complexity of the generated utterances and sentence length. If the number of allowed adjoinings is small, a logical form that can be verbalized as a single complex sentence may lead to a sequence of simple sentences. The leeway created by this parameter can be exploited for mode coordination. For example, constraints set up by the graphics generator or layout manager can force delimitation of sentences, since in a good design, picture breaks should correspond to sentence breaks, and vice versa (cf. [McKeown and Feiner, 1990]).

4.4 Constraint-based Layout

In order to communicate generated information to the user in an adequate manner, we have integrated *LayLab*, an automatic layout manager, into the cascaded architecture of the WIP system (see also [Graf, 1991; Graf and Maaß, 1991; Graf, 1992]). In order to achieve a coherent output, this multimedia layout component is able to reflect certain semantic and pragmatic relations specified by a presentation planner to arrange the visual appearance of a mixture of text and graphics fragments delivered by the media-specific generators, i.e., to determine the size of the layout objects and the exact coordinates for positioning them on the document page.

WIP's presentation design process treats the layout problem as a constraint satisfaction problem. So, the design of an aesthetically pleasing layout is characterized as a combination of a general search problem in a finite discrete search space and an optimization problem. Therefore, we have integrated two dedicated constraint solvers, an incremental hierarchy solver and a finite domain solver, in a layered constraint solver model *CLAY*, which is triggered from a common metalevel by rules and defaults, in order to position the individual fragments on a graphic design grid. The underlying constraint language is able to encode graphical design knowledge expressed by semantic/pragmatic, geometrical/topological, and temporal relations. Furthermore, this mechanism allows us to prioritize the constraints as well as to handle constraint solving over finite domains. As graphical constraints frequently have only local effects, they are incrementally generated by the LayLab system on the fly. Beside the constraint-based positioning component (cf. [Maaß, 1992]), the architecture of the LayLab system includes an automatic grid generator, an intelligent typographer, and an interaction handler (cf. [Soetopo, 1992]).

4.5 Knowledge Representation and Reasoning

Most of the representation of domain knowledge in WIP is based on terminological logics. As far as only static knowledge about objects and their structure is concerned, "conventional" terminological representation systems are appropriate for this task. The KRIS system [Baader and Hollunder, 1991] which we employed in WIP, however, turned out to be inadequate in two respects. First, the interface to the application did not provide the required functionality. Second, the system was orders of magnitudes slower than other systems [Heinsohn *et al.*, 1992a; Heinsohn *et al.*, 1992c]. Both of these shortcomings stem from the fact that the system was only intended to be an experimental testbed for subsumption algorithms. As we were able to show, these limitations were not inherent to the general approach [Baader *et al.*, 1992b]. As a side effect of this work, a specification of a common terminological language has been developed jointly with the WINO project, which has become part of the KRSS standard effort for terminological representation systems, which is one of the projects of the "DARPA Knowledge Sharing Effort" [Patil *et al.*, 1992]. On the theoretical side, a number of existing features and possible extensions of terminological logics were explored and analyzed from a computational and logical point of view [Nebel, 1990b; Nebel, 1991; Nebel and Smolka, 1991; Baader *et al.*, 1992a; Heinsohn and Hollunder, 1992].

Since terminological representation formalisms are only aimed at representing categorical knowledge, but it is often also necessary to represent knowledge that is uncertain and/or vague, an integration of probabilistic approaches and terminological approaches appears to be desirable. Based on research that has been carried out in the area of representing uncertain and vague knowledge [Kruse *et al.*, 1991; Heinsohn and van Loon, 1988], we have designed an extension to terminological formalisms, the language ALCP, that allows the representation of and reasoning about uncertain knowledge in terminological representation systems [Heinsohn, 1991a; Heinsohn, 1992]. While this work has been purely theoretical up to now, we anticipate an implementation and application of this approach in the $\mathcal{PPP}$ project.

Another extension of terminological representation formalisms that proved necessary was an extension that supports the representation of actions and plans in order to adequately represent operating instructions. In order to support the presentation planning and generation task, new reasoning services such as computing the feasibility of a plan and the state of affairs after executing part of a plan have been implemented [Heinsohn *et al.*, 1992b; Heinsohn *et al.*, 1991]. Some of the theoretical problems associated with these reasoning services, such as computing the consequences of a plan, have been investigated in [Nebel and Bäckström, 1992a; Nebel and Bäckström, 1992b], showing that this problem is not as hard as other authors have claimed.

5 Research Plan

5.1 Presentation Planning and Design

In $\mathcal{PPP}$ we view the design of a multimedia document as a non-monotonic process that includes various revisions of preliminary results and negotiations between the system and the user.

Reactive Planning When planning presentations, unexpected situations may arise that require the system to revise the initial plan. Such revisions might be due to new high priority goals in the back-end system or the addressee's reaction to the output generated so far. In *PPP*, we will rely on work on encoding reactive behavior and represent reactive plans (i.e., plans that consist of reactions to possible situations) in the RAT formalism. However, we will not only store reactions for possible situations in advance, but also examine how reactions to unexpected situations can be computed at execution time. These investigations will provide the basis for the extension of WIP's RST-based planner to a reactive planner. To enable immediate reactions to unexpected situations, planning and execution will be interleaved.

Planning Presentation Acts *PPP* does not only synthesize documents, but also plans how to present this material to various users. This means, we have not only to define plan operators for generating the material to be presented, but also plan operators for planning presentation acts. In *PPP*, the processes for planning the presentation and planning the presentation acts will be interleaved to handle the dependencies between them. An interesting subproblem when planning presentation acts is to coordinate them temporally. To accomplish this task, we will extend the presentation planner developed in WIP by temporal constraints that will be propagated during the planning process.

Dialogue Planning In contrast to WIP, *PPP* supports user interaction during the multimodal presentation. Instead of relying on a sophisticated natural language analysis component, we offer the user a comfortable hypermedia interface. However, such an interface requires a system to understand the presentations it produces. For example, to interprete pointing gestures referring to picture parts, the system has to know what the picture presents. To avoid the generation of inconsistent or incoherent output, *PPP* has to record all its design decisions during the generation process. In addition to this, it will also have to keep track of the user's and the system's dialogue acts. To interpret the user's feedback, the system will rely on context-sensitive disambiguation heuristics.

In WIP, we planned the content and the structure of a multimodal document. In *PPP* we will go beyond this and plan both the content of a presentation and conversational moves in a multimodal environment. Consequently, we need not only plan operators to formalize content planning, but also discourse operators for conversational moves, such as interruptions, returning to a previously mentioned topic, checks, confirmations, etc. In *PPP* we will examine several kinds of dialogues in order to find out how they are structured. These studies will then provide a basis for defining discourse plan operators.

Hypergraphics The automatic generation of hypergraphics is a new issue arising from the need for interactive presentations in *PPP*. As illustrated in Section 2.2.2 it is quite useful to allow user interaction on a generated graphics, e.g, by clicking on an object depiction to obtain more information about that entity. In *PPP*, follow-up questions on a graphics can trigger passive and interactive presentation strategies. A passive presentation strategy to elaborate a picture part could be a textual explanation (as in an ordinary hypertext system), a further graphics, or a new text-picture combination presented by *PPP*. In contrast to this, *PPP*'s interactive presentation strategies involve the user in the explanation process. For example, in order to inform a user about certain object properties the user may be requested to zoom and pan on a particular part, to scroll a graphics in a window, or to simulate a walk-around by continously changing the viewing specification. To cope with interaction on graphics we have to maintain an explicit representation of the surface structure and the semantics of a graphics on display. In *PPP*, a propositional picture description will be built up during the graphics generation process.

There is also a technical dimension when producing graphical output. From our experience with commercial graphics software we know that there is no ideal graphics tool available on the market that meets all the requirements a system like *PPP* demands. Consequently, we also will have to address problems like adaptation, integration and augmentation of graphics tools at hand.

Controlling the Animated Presenter As elaborated in Section 2.2.1 the planning of presentation acts is one of the central research topics in *PPP*. In order to demonstrate and evaluate our results, we need a component that realizes planned presentation acts in a natural way. In *PPP*, we will use an animated character that plays the role of a presenter showing, commenting and explaining the generated material. It is clear that within the *PPP* project we cannot aim at sophisticated character animation; this is a hard and complex task in its own right and has been a hot topic for several years now in the computer graphics community. Rather, we will rely on a simple 2D icon-based character and concentrate on synchronizing some animated pointing gestures with natural language output.

Layout of Interactive Multimedia Presentations While our previous work in WIP has concentrated on automatically generated grids and constraint formalisms for supporting the layout design of static text-picture presentations, the *PPP* system will be enriched by further media including informational graphics (e.g., charts, diagrams) and dynamic as well as canned presentation parts (e.g., animation, hypermedia). So, a layout manager designed for *PPP* will be concerned with arranging the generated multimedia output as well as managing the interface to the user and the application. As we have proven in WIP, constraint processing techniques provide an elegant mechanism to specify layout requirements in graphical environments as well as to declaratively state design-relevant knowledge about heterogeneous geometrical relationships, characterizing properties between different kinds of multimedia items that can be maintained by the underlying system. Therefore, we will generalize the constraint-based approach used in WIP towards dynamic interactive layout design.

Editing of Incrementally Laid Out Presentations In *PPP* we will allow the user to tailor the interface to his needs by editing incrementally laid out presentations, changing default layout schemata interactively or working on virtual displays. We will address these goals through the extension of an existing incremental constraint hierarchy solver with regard to dynamic layout tasks. Here, we have to consider the fact, that in interactive graphical environments, not only the constraint hierarchy changes frequently, but the constraint solver must be capable of finding solutions without reducing the direct manipulation responsiveness. Another important topic will be concerned with the representation of layout stereotypes and the use of multimedia units retrieved from a 'case library' (see also Chap. 5.2).

Animated Layout Animating layout is an area of active research that is mostly based on experiences gained from algorithm visualization. In $\mathcal{PPP}$ animated multimedia presentations can enhance the effectiveness and expressiveness of both, the visualization of the incremental layout process and dynamic application scenarios, such as configuration tasks, process monitoring and viewing the dynamics of simulations. So, one of our efforts will be concerned with the evaluation of current work on animating layout of multimedia presentations in order to realize the flow layout technique for 2D graphics in our application domain. Constraints will be useful to describe the appearance and structure of multimedia items as well as how those items evolve over time. Thus, the layout of presentations including animation requires an extension of the exploited constraint language by introducing temporal constraints and mechanisms for satisfying them.

5.2 Knowledge Representation and Reasoning

Extensions of the RAT system and Plan Monitoring In the WIP project, terminological logics have been successfully employed for the purpose of representing knowledge about the domain and they have served as a base for extensions, such as modelling plans and actions. We intend to use terminological logics also in $\mathcal{PPP}$, in particular, we will further use the KRIS system and our extension RAT, which has been built on top of KRIS.

As the requirements of the $\mathcal{PPP}$ system with respect to knowledge representation and reasoning about knowledge will go beyond those of the WIP system, the RAT formalism developed and used in WIP must be extended in various ways. The lanuguage provided by RAT currently supports only atomic actions and linear sequences. In order to allow for the representation of more complex plans, it should include additional constructs as, e.g., partially ordered actions, simultaneous actions, conditionals, and complex temporal orderings of plan steps like those of Allen.

In order to achieve a firm design of an extended RAT formalism, existing approaches dealing with actions and plans will be studied. In the area of plan synthesis and recognition various attempts have been made to enrich the representation formalisms to handle plans more complex than linear sequences of actions. These formalisms, however, are not coupled with terminological logics and most of them add only one new aspect of plan compositions whereas the new RAT formalism must comprise all the above mentioned constructs. For example, when combining non-linear plans and the possibility of simultaneous actions new problems may arise. In order to provide a theoretically well-founded approach, we will study the existing approaches carefully and design an extended RAT-language.

Another aspect of the new RAT system is new inference services which are caused by the new requirements of $\mathcal{PPP}$. For instance, the validation of a plan execution with respect to the abstract description of a plan will be one of the tasks of RAT. This problem will arise especially when monitoring the user's execution of a presented domain plan. The user's actions and their effects must be mirrored by assertions in RAT's assertional knowledge base. Additionally, external events which may occur during the execution must be taken into account because they may effect the further execution of a plan.

Uniform Representation of Domain and Presentation Plans One of the objectives of $\mathcal{PPP}$ is to have a single formalism for the representation of knowledge about the domain plans as well as about the presentation plans. Therefore these two knowledge sources will be modeled using the RAT formalism. This includes on the one hand knowledge concerning the different application domains and temporal and causal relationships in these domains and on the other hand knowledge concerning strategies to present this knowledge to the user. Since the representation of presentation strategies requires a more powerful representation formalism than currently offered by KRIS, in particular, it is necessary to deal with modal belief operators, a necessary prerequisite for this task is an evaluation of the reasoning requirements in this context.

Reasoning about Beliefs Presentation planning in the $\mathcal{PPP}$ system is based on a theory of beliefs and intention similar to the one described in [Cohen and Levesque, 1990]. In particular, the planning process aims at satisfying goals of the form "the user should know ..." In the WIP project, reasoning about the user's beliefs has been already dealt with, however, in a limited form. We anticipate a number of necessary extensions and generalizations which are necessary to support this kind of reasoning in the broader context of $\mathcal{PPP}$. We will try to meet this goal by adapting existing solutions. One possibility may be to employ the modal extension of terminological logics developed by Ohlbach [Ohlbach, 1992]. However, it is not yet clear whether this extension will be efficient enough and, more importantly, will be ready to be used. For this reason, and because the representation of beliefs in the context of presentation planning does not seem to require a completely general form of reasoning about beliefs, we will also explore syntactic variants of reasoning about beliefs [Konolige, 1986], which promises to be more efficient than general methods as the ones proposed by Ohlbach.

A new problem that has to be addressed in the context of reactive presentation planning is the problem of *belief revision*. Since criticism or follow-up questions of the user indicate that some of the communicative goals of the presentation plan have not been achieved, it is necessary to revise the beliefs concerning the user's beliefs. Although the belief revision problem has been extensively studied in recent years [Gärdenfors, 1992], it is not yet clear whether it is possible to generalize the results to a context in a straightforward way, where beliefs about beliefs are revised. Starting with the results achieved in the area of belief revision, we will investigate this problem.

Uncertainty The language ALCP designed in the WIP project allows uncertainty in terminological logics to be dealt with and can be applied to model the typical (in the statistical interpretation) behavior of users of technical environments or to quantify preferences in choosing actions and plans, for instance. However, several questions concerning the application of such an extended terminological framework in the $\mathcal{PPP}$ project remain and must be considered in more detail: While the above approach allows the modeling of *generic* knowledge with a statistical interpretation and the drawing of inferences on the basis of terminological and statistical knowledge, several other uncertainty phenomena exist. An important one is related to the consideration of *individual beliefs* that require an extended assertional formalism.

Besides the task of implementing and extending ALCP, we have to note that uncertain and incomplete knowledge does not only exist in terms of *facts* but also in relation to *actions and plans* in the case of both domain modeling and presentation planning. Further, the automatic *constraint-based layout* is also influenced by design criteria that are "weighted" to allow for an optimal layout of multimedia documents in different situations—a process that has to be be supported by an appropriate uncertainty model.

Efficient Inference Algorithms Certain inference procedures that are needed in our applications are inherently worst-case intractable. Although good implementation techniques and clever algorithms can reduce the "average" runtime in most cases, there is the problem of designing methods that are provably fast or simply faster than other comparable methods. Our main emphasis will be on designing efficient methods for subsumption computation. Since approximation methods and probabilistic approaches have recently been shown to give quite satisfying results for a number of different problems, we will apply these methods in order to design new algorithms for subsumption/satisfiability computation in terminological logics. One particular point we will focus on is the computation of satisfiability for feature structures, which is of relevance because it is an important representation structure used in RAT and has relevance in the area of unification grammars. Due to the latter, we intend to cooperate with DISCO on this topic.

Retrieval of Multimedia Units From the viewpoint of information retrieval, terminological languages have several advantages: they allow the description of classes of objects with complex relational structure, they allow the handling of taxonomies, and, most important, they provide classification as key inference. By classifying descriptions of multimedia units, e.g., the most specific instances associated with the description can be retrieved. In order to support *interactive retrieval*, a "query by example" interface must be developed that will allow to process queries, to present sample individuals that satisfy the query, and, depending on the precision of the answer, to generalize or to specialize the query. In addition, in order to allow program interaction between the different modules of the $\mathcal{PPP}$ system, syntax and semantics of a *query language* for terminological information retrieval have to be developed.

References

[AAAI-86, 1986] *Proceedings of the 5th National Conference of the American Association for Artificial Intelligence*, Philadelphia, PA, August 1986.

[AAAI-88, 1988] *Proceedings of the 7th National Conference of the American Association for Artificial Intelligence*, Saint Paul, MI, August 1988.

[AAAI-90, 1990] *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence*, Boston, MA, August 1990. MIT Press.

[AAAI-91, 1991] *Proceedings of the 9th National Conference of the American Association for Artificial Intelligence*, Anaheim, CA, July 1991. MIT Press.

[AAAI-92, 1992] *Proceedings of the 10th National Conference of the American Association for Artificial Intelligence*, San Jose, CA, July 1992. MIT Press.

[AAAI-WS, 1992] *AAAI-92 Workshop on Tractable Reasoning*, San Jose, CA, July 1992.

[Adams, 1976] J. Adams. A probability model of medical reasoning and the MYCIN model. *Mathematical Biosciences*, 32:177–186, 1976.

[AIPS-92, 1992] *Proceedings of the 1st International Conference on Artificial Intelligence Planning Systems*, Washington, D.C., 1992. Morgan Kaufmann.

[Allen *et al.*, 1991a] J. A. Allen, R. Fikes, and E. Sandewall, editors. *Principles of Knowledge Representation and Reasoning: Proceedings of the 2nd International Conference*, Cambridge, MA, April 1991. Morgan Kaufmann.

[Allen *et al.*, 1991b] J. F. Allen, H. A. Kautz, R. N. Pelavin, and J. D. Tenenberg, editors. *Reasoning about Plans*. Morgan Kaufmann, San Mateo, CA, 1991.

[Allen, 1983] J. F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, November 1983.

[Allen, 1991] J. F. Allen. Planning as temporal reasoning. In Allen et al. [1991a], pages 3–14.

[Allgayer *et al.*, 1989] J. Allgayer, K. Harbusch, A. Kobsa, C. Reddig, N. Reithinger, and D. Schmauks. XTRA: A natural langugage access system to expert systems. *International Journal of Man-Machine Studies*, 31:161–195, 1989.

[Andersen *et al.*, 1989] S. Andersen, K. Olesen, F. Jensen, and F. Jensen. Hugin – a shell for building Bayesian belief universes for expert systems. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, Detroit, Mich., 1989.

[André and Rist, 1990a] E. André and T. Rist. Synthesizing illustrated documents: A plan-based approach. In *Proceedings of InfoJapan '90, Vol. 2*, pages 163–170, Tokyo, 1990. Also DFKI Research Report RR-91-06.

[André and Rist, 1990b] E. André and T. Rist. Towards a plan-based synthesis of illustrated documents. In *Proceedings of the 9th European Conference on Artificial Intelligence*, pages 25–30, Stockholm, Sweden, July 1990. Also DFKI Research Report RR-90-11.

[André and Rist, 1992] E. André and T. Rist. The design of illustrated documents as a planning task. In Maybury [1992]. Forthcoming.

[André *et al.*, 1986] E. André, G. Bosch, G. Herzog, and T. Rist. Characterizing Trajectories of Moving Objects Using Natural Language Path Descriptions. In *Proceedings ECAI-86*, pages Vol. 2, 1–8, Brighton, England, 1986.

[André *et al.*, 1992] E. André, W. Finkler, W. Graf, T. Rist, A. Schauder, and W. Wahlster. WIP: The automatic synthesis of multimodal presentations. In Maybury [1992]. Forthcoming.

[Andreassen *et al.*, 1987] S. Andreassen, M. Woldbye, B. Falck, and S. Andersen. Munin – a causal probabilistic network for interpretation of electromyographic findings. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, pages 366–372, Milan, Italy, 1987.

[Appelt and Pollack, 1990] D. Appelt and M. Pollack. Weighted abduction for plan ascription. Technical report, Artificial Intelligence Center and Center for the Study of Language and Information, SRI International, Menlo Park, California, 1990.

[Arens *et al.*, 1989] Y. Arens, S. Feiner, J. Hollan, and B. Neches, editors. *Proceedings of the IJCAI-89 Workshop 'A New Generation of Intelligent Interfaces', Detroit, MI*. 1989.

[Arens *et al.*, 1991] Y. Arens, L. Miller, and N. Sondheimer. Presentation design using an integrated knowledge base. In Sullivan and Tyler [1991], pages 241–258.

[Arnheim, 1966] R. Arnheim, editor. *Towards a Psychology of Art.* University of California Press, Berkeley, 1966.

[Baader and Hanschke, 1991] F. Baader and P. Hanschke. A scheme for integrating conrete domains into concept languages. In IJCAI-91 [1991].

[Baader and Hollunder, 1991] F. Baader and B. Hollunder. KRIS: Knowledge representation and inference system. *SIGART Bulletin*, 2(3):8–14, June 1991.

[Baader and Hollunder, 1992] F. Baader and B. Hollunder. Embedding defaults into terminological knowledge representation formalisms. In Nebel et al. [1992], pages 306–317.

[Baader *et al.*, 1990] F. Baader, H.-J. Bürckert, J. Heinsohn, B. Hollunder, J. Müller, B. Nebel, W. Nutt, and H.-J. Profitlich. Terminological Knowledge Representation: A proposal for a terminological logic. DFKI Technical Memo TM-90-04, Saarbrücken, 1990. A revised version has been published in [Nebel *et al.*, 1991].

[Baader *et al.*, 1992a] F. Baader, H.-J. Bürckert, B. Nebel, W. Nutt, and G. Smolka. On the expressivity of feature logics with negation, functional uncertainty, and sort equations. *Journal of Logic, Language and Information*, 1992. To appear. A preliminary version is available as DFKI Research Report RR-91-01.

[Baader *et al.*, 1992b] F. Baader, B. Hollunder, B. Nebel, H.-J. Profitlich, and E. Franconi. An empirical analysis of optimization techniques for terminological representation systems or "making KRIS get a move on". In Nebel et al. [1992], pages 270–281.

[Bacchus *et al.*, 1992] F. Bacchus, A. Grove, Y. Y. Halpern, and D. Koller. From statistics to beliefs. In *Proceedings of the 10th National Conference of the American Association for Artificial Intelligence*, pages 602–608, San Jose, Cal., 1992.

[Bacchus, 1989] F. Bacchus. A modest, but semantically well founded, inheritance reasoner. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 1104–1109, Detroit, Mich., 1989.

[Bacchus, 1990] F. Bacchus. Lp, a logic for representing and reasoning with statistical knowledge. *Computational Intelligence*, 6:209–231, 1990.

[Bäckström and Nebel, 1992a] C. Bäckström and B. Nebel. Complexity results for SAS^+ planning. Submitted for publication, 1992.

[Bäckström and Nebel, 1992b] C. Bäckström and B. Nebel. On the computational complexity of planning and story understanding. In Neumann [1992], pages 349–353.

[Bäckström, 1988] C. Bäckström. *Reasoning about Interdependent Actions.* Licentiate thesis, Linköping University, Linköping, Sweden, 1988.

[Badler and Webber, 1990] N. Badler and B. Webber. Task communication through natural language and graphics: A workshop report. Dept. of Computer and Information Science, University of Pennsylvania, June 1990.

[Badler *et al.*, 1980] N. Badler, I. O'Rourke, and B. Kaufman. Special Problems in Human Movement Simulation. *Computer Graphics*, 14(3):189–197, July 1980.

[Badler *et al.*, 1991a] N. Badler, B. Barsky, and D. Zeltzer, editors. *Making Them Move: Mechanics, Control, and Animation of Articulated Figures.* Morgan Kaufmann, Los Altos, CA, 1991.

[Badler *et al.*, 1991b] N. Badler, B. Webber, J. Kalita, and J. Esakov. Animation from instructions. In Badler et al. [1991a], pages 51–93.

[Badler, 1987] N. Badler. Computer animation techniques. In W. Brauer and W. Wahlster, editors, *Proceedings 2. Internationaler GI-Kongreß Wissensbasierte Systeme*, pages 22–34. Springer-Verlag, Berlin, Germany, October 1987.

[Baecker, 1969] R. Baecker. Picture Driven Animation. In *SJCC*, pages 273–288. AFIPS Press, Montvale, 1969.

[Barnett, 1981] J. Barnett. Computational methods for a mathematical theory of evidence. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, pages 868–875, Vancouver, Canada, 1981.

[Barr and Barzel, 1988] A. Barr and R. Barzel. A Modeling System Based on Dynamic Constraints. In *ACM SIGGRAPH '88*, pages 179–188, 1988.

[Barr and Feigenbaum, 1981] A. Barr and E. Feigenbaum, editors. *The Handbook of Artificial Intelligence, Vol. 1.* Pitman, London, England, 1981. Chap. III, C5: Direct (Analogical) Representations.

[BDW, 1992] Bild der Wissenschaft, Oktober 1992.

[Beach, 1985] R. Beach. *Setting Tables and Illustrations with Style.* PhD thesis, Dept. of Computer Science, University of Waterloo, Ontario, 1985.

[Beck *et al.*, 1989] H. W. Beck, S. K. Gala, and S. B. Navathe. Classification as a query processing technique in the CANDIDE semantic data model. In *Proceedings of the International Data Engineering Conference, IEEE*, pages 572–581, Los Angeles, CA, February 1989.

[Bertin, 1983] J. Bertin. *Semiology of Graphics.* The University of Wisconsin Press, Milwaulkee, Wisconsin, 1983. Translated by W. Berg.

[Biggerstaff and Perlis, 1989a] T. J. Biggerstaff and A. Perlis. *Software Reusability. Vol.1: Concepts and Models.* Addison-Wesley, Reading, Mass., 1989.

[Biggerstaff and Perlis, 1989b] T. J. Biggerstaff and A. Perlis. *Software Reusability. Vol.2: Applications and Experience.* Addison-Wesley, Reading, Mass., 1989.

[Biundo *et al.*, 1992] S. Biundo, D. Dengler, and J. Koehler. Deductive planning and plan reuse in a command language environment. In *Proceedings of the 10th European Conference on Artificial Intelligence, Vienna, Austria*, pages 628–632. John Wiley & Sons, Chichester, UK, 1992.

[Böhringer and Paulisch, 1990] K.-F. Böhringer and F. N. Paulisch. Using Constraints to Achieve Stability in Automatic Graph Layout Algoriths. In *Proceedings of CHI '90 (Human Factors in Computing Systems), Seattle, WA*, pages 43–51, 1990.

[Bonissone and Wood, 1989] P. P. Bonissone and N. C. Wood. T-norm based reasoning in situation assessment applications. In Kanal et al. [1989], pages 241–256.

[Bonissone *et al.*, 1987] P. Bonissone, S. Gans, and K. Decker. RUM: A layered architecture for reasoning with uncertainty. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, pages 891–898, Milan, Italy, 1987.

[Bonissone *et al.*, 1991] P. P. Bonissone, M. Henrion, L. Kanal, and J. Lemmer, editors. *Uncertainty in Artificial Intelligence 6.* North-Holland, Amsterdam, 1991.

[Bonissone, 1987] P. Bonissone. Summarizing and propagating uncertain information with triangular norms. *International Journal of Approximate Reasoning*, 1:71–101, 1987.

[Borgida *et al.*, 1989] A. Borgida, R. J. Brachman, D. L. McGuinness, and L. A. Resnick. CLASSIC: a structural data model for objects. In *Proceedings of the 1989 ACM SIGMOD International Conference on Mangement of Data*, pages 59–67, Portland, Oreg., June 1989.

[Borning and Duisberg, 1986] A. Borning and R. Duisberg. Constraint-based tools for building user interfaces. *ACM Transactions on Graphics*, 5(4):345–374, October 1986.

[Borning *et al.*, 1987] A. Borning, R. Duisberg, B. Freeman-Benson, A. Kramer, and M. Woolf. Constraint hierarchies. In *Proceedings of OOPSLA '87 (Object-Oriented Programming Systems, Languages, and Applications)*, pages 48–60, October 1987.

[Borning, 1979] A. Borning. *ThingLab - A Constraint-Oriented Simulation Laboratory*. PhD thesis, Dept. of Computer Science, Stanford University, Stanford, CA, 1979.

[Borning, 1981] A. Borning. The programming language aspects of ThingLab, a constraint-oriented simulation laboratory. *ACM Transactions on Programming Languages and Systems*, 3(4):353–387, October 1981.

[Brachman and Levesque, 1984] R. J. Brachman and H. J. Levesque. The tractability of subsumption in frame-based description languages. In *Proceedings of the 4th National Conference of the American Association for Artificial Intelligence*, pages 34–37, Austin, TX, 1984.

[Brachman *et al.*, 1983] R. J. Brachman, R. E. Fikes, and H. J. Levesque. KRYPTON: A functional approach to knowledge representation. *IEEE Computer*, 16(10):67–73, October 1983.

[Brachman *et al.*, 1985] R. J. Brachman, V. Pigman Gilbert, and H. J. Levesque. An essential hybrid reasoning system: Knowledge and symbol level accounts in KRYPTON. In IJCAI-85 [1985], pages 532–539.

[Brachman *et al.*, 1989] R. Brachman, H. J. Levesque, and R. Reiter, editors. *Principles of Knowledge Representation and Reasoning: Proceedings of the 1st International Conference*, Toronto, ON, May 1989. Morgan Kaufmann.

[Brachman *et al.*, 1991] R. J. Brachman, D. L. McGuinness, P. F. Patel-Schneider, L. A. Resnick, and A. Borgida. Living with CLASSIC: When and how to use a KL-ONE-like language. In Sowa [1991], pages 401–456.

[Brachman, 1978] R. J. Brachman. Structured inheritance networks. In W. A. Woods and R. J. Brachman, editors, *Research in Natural Language Understanding, Quarterly Progress Report No. 1, BBN Report No. 3742*, pages 36–78. Bolt, Beranek, and Newman Inc., Cambridge, MA, 1978.

[Brachman, 1990] R. J. Brachman. The future of knowledge representation. In *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence*, pages 1082–1092, Boston, Mass., 1990.

[Buchanan and Shortliffe, 1984] B. Buchanan and E. Shortliffe. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, Reading, Mass., 1984.

[Bylander *et al.*, 1991] T. Bylander, D. Allemang, M. Tanner, and J. Josephson. The computational complexity of abduction. *Artificial Intelligence*, 49:25–60, 1991.

[Bylander, 1991a] T. Bylander. Comnlexity results for planning. In IJCAI-91 [1991], pages 274–279.

[Bylander, 1991b] T. Bylander. The monotonic abduction problem: A functional characterization on the edge of tractability. In Allen et al. [1991a], pages 70–77.

[Bylander, 1992] T. Bylander. Complexity results for extended planning. In AIPS-92 [1992].

[Cadoli and Schaerf, 1992] M. Cadoli and M. Schaerf. Approximation in concept description languages. In Nebel et al. [1992], pages 342–353.

[Casner, 1991] S. Casner. A Task-Analytic Approach to the Automated Design of Graphic Presentations. *ACM Transactions on Graphics*, 10(2):111–151, April 1991.

[Catmull, 1972] E. Catmull. A System for Computer Generated Movies. In *ACM Annual Conference*, pages 422–431, ACM, 1972.

[Cattoni and Franconi, 1990] R. Cattoni and E. Franconi. Walking through the semantics of frame-based description languages: A case study. In *Proceedings of the Fifth International Symposium on Methodologies for Intelligent systems*, Knoxville, TN, October 1990. North-Holland.

[Cendrowska and Bramer, 1984] J. Cendrowska and M. Bramer. A rational reconstruction of the MYCIN consultation system. *International Journal of Man-Machine Studies*, 20:229–317, 1984.

[Chapman, 1987] D. Chapman. Planning for conjunctive goals. *Artificial Intelligence*, 32(3):333–377, July 1987.

[Charniak and Shimony, 1990] E. Charniak and S. Shimony. Probabilistic semantics for cost based abduction. In *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence*, pages 106–111, 1990.

[Cheeseman *et al.*, 1991] P. Cheeseman, B. Kanefsky, and W. M. Taylor. Where the *really* hard problems are. In IJCAI-91 [1991], pages 331–337.

[Cohen and Levesque, 1990] P. R. Cohen and H. J. Levesque. Rational interaction as the basis for communication. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*. MIT Press, Cambridge, MA, 1990.

[Cohen *et al.*, 1986] E. Cohen, E. Smith, and L. Iverson. Constraint-based tiled windows. *IEEE Computer Graphics and Applications*, 6(5):35–45, 1986.

[Cohen *et al.*, 1989] P. Cohen, J. Sullivan, M. Dalyrymple, R. Gargan, D. Moran, J. Schlossberg, F. Perreira, and S. Tyler. Synergistic Use of Direct Manipulation and Natural Language. In *Proceedings CHI-89*, pages 227–233, Austin, 1989.

[Colby, 1992] G. Colby. Maintaining legibility, structure, and style of information layout in dynamic display environments. Submitted to CHI'92 Monterey, CA, May 1992.

[Costabile *et al.*, 1992] M. F. Costabile, T. Catarci, and S. Levialdi, editors. *Advanced Visual Interfaces.* World Scientific Press, Singapore, 1992. Proceedings of the AVI '92 Workshop, Rome, Italy, May 27-29.

[Csinger, 1991] A. Csinger. The Psychology of Visualization. Psychology 579 Project, Department of Computer Science, University of British Columbia, Vancouver, Canada, April 1991.

[Dale *et al.*, 1992] R. Dale, E. Hovy, D. Rösner, and O. Stock, editors. *Aspects of Automated Natural Language Generation.* Springer-Verlag, Berlin, Germany, April 1992. Proceedings of the International Generation Workshop, Trento, Italy.

[Dale, 1992] R. Dale. Visible language: Multimedia constraints in information presentation. In Dale et al. [1992], pages 281–283. Panel: Extending Language Generation to Multiple Media.

[Dean and Boddy, 1987] T. L. Dean and M. Boddy. Incremental causal reasoning. In *Proceedings of the 6th National Conference of the American Association for Artificial Intelligence*, pages 196–201, Seattle, WA, July 1987.

[Dean and Boddy, 1988] T. L. Dean and M. Boddy. Reasoning about partially ordered events. *Artificial Intelligence*, 36(3):375–400, October 1988.

[Dempster, 1967] A. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Ann. Math. Statist.*, 38:325–339, 1967.

[Devanbu and Litman, 1991] P. T. Devanbu and D. J. Litman. Plan-based terminological reasoning. In Allen et al. [1991a], pages 128–138.

[Devanbu *et al.*, 1991] P. T. Devanbu, R. J. Brachman, P. G. Selfridge, and B. W. Ballard. LaSSIE: a knowledge-based software information system. *Communications of the ACM*, 34(5):35–49, May 1991.

[Donini *et al.*, 1991a] F. M. Donini, M. Lenzerini, D. Nardi, and W. Nutt. The complexity of concept languages. In Allen et al. [1991a], pages 151–162.

[Donini *et al.*, 1991b] F. M. Donini, M. Lenzerini, D. Nardi, and W. Nutt. Tractable concept languages. In IJCAI-91 [1991], pages 458–465.

[Donini *et al.*, 1992] F. M. Donini, M. Lenzerini, D. Nardi, B. Hollunder, W. Nutt, and A. M. Spacamella. The complexity of existential quantification in concept languages. *Artificial Intelligence*, 53(2-3):309–327, 1992.

[Doyle and Patil, 1991] J. Doyle and R. S. Patil. Two theses of knowledge representation: Language restrictions, taxonomic classification, and the utility of representation services. *Artificial Intelligence*, 48(3):261–298, April 1991.

[Dubois and Prade, 1988] D. Dubois and H. Prade. Default reasoning and possibility theory. *Artificial Intelligence*, 35:243–257, 1988.

[Dubois and Prade, 1991] D. Dubois and H. Prade. Possibilistic logic, preferential models, non-monotonicity and related issues. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, Sydney, Australia, August 1991.

[Duda *et al.*, 1976] R. Duda, P. Hart, and N. Nilsson. Subjective bayesian methods for rule-based inference systems. In *Proc. of National Computer Conference. AFIPS, Vol.45*, pages 1075 – 1082, 1976.

[Duda *et al.*, 1978] R. Duda, P. Hart, P. Barnett, J. Gasching, K. Konoloige, R. Reboh, and J. Slocom. Development of the prospector consultant system for mineral exploration. final report for sri projects 5821 and 6915. Technical report, SRI International Artificial Intelligence Center, 1978.

[Duda *et al.*, 1981] R. Duda, J. Gaschnig, and P. Hart. Model design in the prospector consultant system for mineral exploration. In D. Michie, editor, *Expert Systems in the Microelectronic Age*, pages 153–167. Edinburgh University Press, Edinburgh, Scotland, 1981.

[Duisberg, 1987] R. Duisberg. Animation Using Temporal Constraints: An Overview of the Animus System. In *Human-Computer Interaction*, pages 275–307. Lawrence Erlbaum Associates, Hillsdale, NJ, 1987.

[Duisberg, 1990] R. Duisberg. Visual Programming of Program Visualizations - A Gestural Interface for Animating Algorithms. In T. Ichikawa, E. Jungert, and R. Korfhage, editors, *Visual Languages and Applications*, pages 161–173. Plenum Press, New York, NY, 1990.

[Edmonds and Murray, 1992] E. Edmonds and B. Murray. Intelligent presentation and dialogue management. In *Proceedings Summer School Prague*, pages 465–476, 1992.

[Elzer *et al.*, 1988] P. Elzer, H. Siebert, and K. Zinser. New Possibilities for the Presentation of Process Information in Industrial Control. In *3rd IFAC/IFIP/IEA/IFORS Conference on Man-Machine Systems, Analysis, Design and Evaluation*, 1988.

[Erol *et al.*, 1992] K. Erol, D. S. Nau, and V. S. Subrahmanian. On the complexity of domain-independent planing. In AAAI-92 [1992], pages 381–386.

[EXPLORE, 1992] EXPLORE. Product Information. TDI Inc., 1992.

[FallSymp-92, 1992] *Working Notes of the AAAI Fall Symposium on Issues in Description Logics: Users meet Developers*, Royal Sonesta Hotel, Cambridge, MA, October 1992.

[Feiner and McKeown, 1990] S. Feiner and K. McKeown. Generating coordinated multimedia explanations. In *Proceedings CAIA90 (6th Conference on Artificial Intelligence Applications)*, pages 290–296, Santa Barbara, CA, March 1990.

[Feiner and McKeown, 1992] S. K. Feiner and K. R. McKeown. Automating the generation of coordinated multimedia explanations. In Maybury [1992]. Forthcoming.

[Feiner *et al.*, 1982] S. Feiner, D. Salesin, and T. Banchoff. DIAL: A Diagrammatic Animation Language. *IEEE Computer Graphics and Applications*, 2(7):17–25, 1982.

[Feiner *et al.*, 1991] S. Feiner, D. Litman, K. McKeown, and R. Passonneau. Towards coordinated temporal multimedia presentations. In *Proceedings of the AAAI-91 Workshop on Intelligent Multimedia Interfaces*, Anaheim, CA, July 1991.

[Feiner, 1985] S. Feiner. APEX: An experiment in the automated creation of pictorial explanations. *IEEE Computer Graphics and Applications*, 5(11):29–39, 1985.

[Feiner, 1988] S. Feiner. A grid-based approach to automating display layout. In *Proceedings of the Graphics Interface '88*, pages 192–197. Morgan Kaufmann, Los Altos, CA, June 1988.

[Fikes and Nilsson, 1971] R. E. Fikes and N. Nilsson. STRIPS: A new approach to the application of theorem proving as problem solving. *Artificial Intelligence*. 2:198–208, 1971.

[Finkler and Neumann, 1989] W. Finkler and G. Neumann. POPEL-HOW: A distributed parallel model for incremental natural language production with feedback. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 1518–1523, Detroit, MI, 1989.

[Finkler and Schauder, 1992] W. Finkler and A. Schauder. Effects of incremental output on incremental natural language generation. In *Proceedings of the 10th European Conference on Artificial Intelligence*, Vienna, Austria, August 1992.

[Foley *et al.*, 1990] J. Foley, A. van Dam, S. Feiner, and J. Hughes. *Computer Graphics: Principles and Practice, 2nd Edition*. The System Programming Series. Addison-Wesley, Reading, MA, 1990.

[Franconi *et al.*, 1992] E. Franconi, B. Magnini, and O. Stock. Prototypes in a hybrid language with primitive descriptions. *Computers & Mathematics with Applications*, 23(6-9), 1992. Special Issue on Semantic Nets.

[Franconi, 1991] E. Franconi. *The YAK Manual*. IRST, Trento, Italy. May 1991.

[Franke, 1992] D. Franke. Classifying and indexing design modifications via descriptions of purpose. In SpringSymp-92 [1992], pages 111–115.

[Freeman-Benson *et al.*, 1990] B. Freeman-Benson, J. Maloney, and A. Borning. An incremental constraint solver. *Communications of the ACM*, 33(1):54–63, 1990.

[Freeman-Benson, 1990] B. Freeman-Benson. Kaleidoscope: Mixing Objects, Constraints, and Imperative Programming. In N. Meyrowitz, editor. *Proceedings of ECOOP-OOPSLA '90*, pages 77–88, Ottawa, Canada, October 1990.

[Freuder and Mackworth, 1992] E. Freuder and A. Mackworth, editors. *Artificial Intelligence, Vol. 58, No. 1-3, Special Volume on Constraint-Based Reasoning*. Elsevier, Amsterdam, December 1992.

[Friedell, 1984] M. Friedell. Automatic synthesis of graphical object descriptions. *Computer Graphics*, 18(3):53–62, 1984.

[Friedman, 1981] L. Friedman. Extended plausible inference. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, pages 487–495, Vancouver, Canada, 1981.

[Gärdenfors, 1992] P. Gärdenfors, editor. *Belief Revision*, volume 29 of *Cambridge Tracts in Theoretical Computer Science.* Cambridge University Press, Cambridge, UK, 1992.

[Garey and Johnson, 1979] M. R. Garey and D. S. Johnson. *Computers and Intractability—A Guide to the Theory of NP-Completeness.* Freeman, San Francisco, CA, 1979.

[Garvey *et al.*, 1981] T. Garvey, J. Lowrance, and M. Fischler. An inference technique for integrating knowledge from disparate sources. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, pages 319–325, Vancouver, Canada, 1981.

[Gelfond *et al.*, 1991] M. Gelfond, V. Lifschitz, and A. Rabinov. What are the limitations of the situation calculus. In R. S. Boyer, editor, *Automated Reasoning. Essays in Honor of Woody Bledsoe*, pages 167–181. Kluwer, Dordrecht, Holland, 1991.

[Geller and Shapiro, 1987] J. Geller and C. Shapiro. Graphical Deep Knowledge for Intelligent Machine Drafting. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, pages 545–551, 1987.

[Ginsberg and Maxwell, 1983] C. Ginsberg and D. Maxwell. Graphical Marionette. In *SIGGRAPH/SIGART Interdisciplinary Workshop on Motion: Representation and Perception*, pages 172–179, April 1983.

[Ginsberg and Smith, 1988a] M. L. Ginsberg and D. E. Smith. Reasoning about action I: A possible worlds approach. *Artificial Intelligence*, 35:165–195, 1988.

[Ginsberg and Smith, 1988b] M. L. Ginsberg and D. E. Smith. Reasoning about action II: The qualification problem. *Artificial Intelligence*, 35:311–342, 1988.

[Ginsberg, 1984] M. Ginsberg. Non-monotonic reasoning using Dempster's rule. In *Proceedings of the 4th National Conference of the American Association for Artificial Intelligence*, pages 126–129, Austin, Tex., August 1984.

[Girard, 1987] M. Girard. Interactive Design of 3-D Computer-Animated Legged Animal Motion. *IEEE Computer Graphics and Applications*, 7(6):39–51, June 1987.

[Givan and McAllester, 1992] R. Givan and D. McAllester. New results in local inference relations. In Nebel et al. [1992], pages 403–412.

[Glymour, 1985] C. Glymour. Independence assumptions and Bayesian updating. *Artificial Intelligence*, 25:95–99, 1985.

[Gnanamgari, 1981] S. Gnanamgari. *Information Presentation Through Default Displays.* PhD thesis, University of Pennsylvania, 1981.

[Gomez, 1984] J. Gomez. Twixt: A 3D Animation System. In *Eurographics'84*, pages 121–133, July 1984.

[Gordon and Shortliffe, 1984] J. Gordon and E. Shortliffe. The Dempster-Shafer theory of evidence. In B. Buchanan and E. Shortlife, editors, *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, pages 272–292. Addison-Wesley, Reading, Mass., 1984.

[Gordon and Shortliffe, 1985] J. Gordon and E. Shortliffe. A method for managing evidential reasoning in a hierarchical hypothesis space. *Artificial Intelligence*, 26:323–357, 1985.

[Gosling, 1983] J. Gosling. *Algebraic Constraints.* PhD thesis, Dept. of Computer Science, Carnegie Mellon University, 1983.

[Graf and Maaß, 1991] W. Graf and W. Maaß. Constraint-basierte Verarbeitung graphischen Wissens. In W. Brauer and D. Hernández, editors, *Proceedings of 4. Internationaler GI-Kongreß Wissensbasierte Systeme - Verteilte KI und kooperatives Arbeiten*, pages 243–253. Springer-Verlag, Berlin, Germany, Oktober 1991. Also DFKI Research Report RR-91-35.

[Graf and Thies, 1992] W. Graf and M. A. Thies. Aspekte der kombination automatisch generierter animationen und planbasierter hilfe. *KI*, 6(4), 1992. Also DFKI Research Report RR-92-09.

[Graf, 1991] W. Graf. Constraint-Based Processing of Design Knowledge. In *Proceedings of the AAAI-91 Workshop on 'Intelligent Multimedia Interfaces'*, Anaheim, CA, July 1991.

[Graf, 1992] W. Graf. Constraint-based graphical layout of multimodal presentations. In Costabile et al. [1992]. Also DFKI Research Report RR-92-15.

[Große and Waldinger, 1991] G. Große and R. Waldinger. Towards a theory of simultaneous actions. In J. Hertzberg, editor, *Proceedings of the 1st European Workshop on Planning*, Sankt Augustin, Germany, August 1991.

[Große, 1992a] G. Große, editor. *Beyond Sequential Planning, Workshop Notes*. August 1992.

[Große, 1992b] G. Große. Simultaneous events and causation. In Große [1992a].

[Gupta and Nau, 1992] N. Gupta and D. S. Nau. On the complexity of blocks-world planning. *Artificial Intelligence*, 56(2):223–254, 1992.

[Hajek *et al.*, 1992] P. Hajek, T. Havranek, and R. Jirousek. *Uncertain Information Processing in Expert Systems.* CRC Press, Boca Raton, Ann Arbor, London, 1992.

[Halpern and Fagin, 1990] J. Y. Halpern and R. Fagin. Two views of belief: Belief as generalized probability and belief as evidence. In *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence*, pages 112–119, Boston, Mass., 1990.

[Halpern and Moses, 1992] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54(3):319–379, 1992.

[Hammond, 1990] K. J. Hammond. Explaining and reparing plans that fail. *Artificial Intelligence*, 45:173–228, 1990.

[Hanks and McDermott, 1987] S. Hanks and D. McDermott. Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33(3):379–412, November 1987.

[Harbusch *et al.*, 1991] K. Harbusch, W. Finkler, and A. Schauder. Incremental syntax generation with tree adjoining grammars. In W. Brauer and D. Hernández, editors, *Proceedings of 4. Internationaler GI-Kongreß Wissensbasierte Systeme - Verteilte KI und kooperatives Arbeiten*, pages 363–374. Springer-Verlag, Berlin, Germany, October 1991.

[Harbusch, 1990] K. Harbusch. Constraining tree adjoining grammars by unification. In *Proceedings of the 13th International Conference on Computational Linguistics*, pages 167–172, 1990.

[Heckerman, 1986] D. Heckerman. Probabilistic interpretations for mycin's certainty factors. In Kanal and Lemmer [1986], pages 167–196.

[Heckerman, 1990] D. E. Heckerman. *Probabilistic Similarity Networks*. MIT Press, Cambridge, Mass., 1990.

[Heinsohn and Hollunder, 1992] J. Heinsohn and B. Hollunder, editors. *Proceedings of the DFKI Workshop on Taxonomic Reasoning*. DFKI Document D-92-08, German Research Center for Artificial Intelligence, Saarbrücken, Germany, February 1992.

[Heinsohn and Owsnicki-Klewe, 1988] J. Heinsohn and B. Owsnicki-Klewe. Probabilistic inheritance and reasoning in hybrid knowledge representation systems. In W. Hoeppner, editor, *Künstliche Intelligenz. GWAI-88, 12. Jahrestagung*, pages 51–60, Eringerfeld, Germany, September 1988. Springer-Verlag.

[Heinsohn and van Loon, 1988] J. Heinsohn and J. van Loon. Numerical measures for handling uncertainty – looked at from a Bayesian perspective. Report MS-H 4795/88, Philips Research Laboratories Eindhoven/Hamburg, 1988. (55 pages).

[Heinsohn *et al.*, 1991] J. Heinsohn, D. Kudenko, B. Nebel, and H.-J. Profitlich. Integration of action representation in terminological logics. In C. Peltason, K. von Luck, and C. Kindermann, editors, *Terminological Logic Users Workshop – Proceedings*, page 117, Berlin, Germany, December 1991. Published as KIT Report 95.

[Heinsohn *et al.*, 1992a] J. Heinsohn, D. Kudenko, B. Nebel, and H.-J. Profitlich. An empirical analysis of terminological representation systems. In AAAI-92 [1992], pages 767–773.

[Heinsohn *et al.*, 1992b] J. Heinsohn, D. Kudenko, B. Nebel, and H.-J. Profitlich. Rat: Representation of actions using terminological logics. DFKI Research Report, German Research Center for Artificial Intelligence (DFKI), Saarbrücken, 1992. To appear.

[Heinsohn *et al.*, 1992c] J. Heinsohn, D. Kudenko, B. Nebel, and H.-J. Profitlich. An empirical analysis of terminological representation systems. DFKI Report RR-92-16, German Research Center for Artificial Intelligence, Saarbrücken, Germany, May 1992. (38 pages).

[Heinsohn, 1991a] J. Heinsohn. A hybrid approach for modeling uncertainty in terminological logics. In R. Kruse and P. Siegel, editors, *Symbolic and Quantitative Approaches to Uncertainty, Proceedings of the European Conference ECSQAU*, Lecture Notes in Computer Science 548, pages 198–205. Springer, Berlin, Germany, 1991. An extended version is available as DFKI Research Report RR-91-24.

[Heinsohn, 1991b] J. Heinsohn. A hybrid approach for modeling uncertainty in terminological logics. DFKI Report RR-91-24, German Research Center for Artificial Intelligence, Saarbrücken, Germany, August 1991. (22 pages).

[Heinsohn, 1992] J. Heinsohn. ALCP - an integrated framework to terminological and noncategorical knowledge. In *Proceedings of the 4th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU'92)*, pages 493–496, Palma de Mallorca, Spain, July 6–10 1992.

[Henry and Hudson, 1991] T. Henry and S. Hudson. Interactive Graph Layout. In *Proceedings of UIST '91 (ACM Symposium on User Interface Software and Technology), Hilton Head, SC*, pages 55–64, 1991.

[Hentenryck, 1989] P. V. Hentenryck. *Constraint Satisfaction in Logic Programming.* MIT Press, Cambridge, MA, 1989. Revision of Ph.D. thesis, University of Namur, 1987.

[Herzog *et al.*, 1989] G. Herzog, C. Sung, E. André, W. Enkelmann, H.-H. Nagel, T. Rist, W. Wahlster, and G. Zimmermann. Incremental Natural Language Description of Dynamic Imagery. In W. Brauer and C. Freksa, editors, *Proceedings of 3. Internationaler GI-Kongreß Wissensbasierte Systeme*, pages 153–162. Springer-Verlag, Berlin, Germany, October 1989.

[Hollunder, 1989] B. Hollunder. Subsumption algorithms for some attributive description languages. Master's thesis, Department of Computer Science, Universität Kaiserslautern, Kaiserslautern, Germany, August 1989.

[Horvitz and Heckerman, 1986] E. Horvitz and D. Heckerman. The inconsistent use of measures of certainty in artificial intelligence research. In Kanal and Lemmer [1986], pages 137–151.

[Horvitz *et al.*, 1986] E. Horvitz, D. Heckerman, and C. Langlotz. A framework for comparing alternative formalisms for plausible reasoning. In *Proceedings of the 5th National Conference of the American Association for Artificial Intelligence*, pages 210–214, Philadelphia, Pa., 1986.

[Horz, 1992] A. Horz. Simultaneous actions and classical planning. In Große [1992a].

[Hovy and Arens, 1991] E. Hovy and Y. Arens. Automatic generation of formatted text. In *Proceedings of the 9th National Conference of the American Association for Artificial Intelligence*, pages 92–97, Anaheim, CA, July 1991.

[Howe, 1992] A. E. Howe. Failure recovery analysis as a tool for plan debugging. In SpringSymp-92 [1992], pages 25–30.

[Hudson and Hsi, 1992] S. Hudson and C.-N. Hsi. A Grid-Based Approach to Specifying Simple Number Independent Layouts by Example. In *Proceedings of UIST '92 (ACM Symposium on User Interface Software and Technology), Monterey, CA*, 1992. Submitted to.

[IJCAI-85, 1985] *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, Los Angeles, CA, August 1985.

[IJCAI-91, 1991] *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, Sydney, Australia, August 1991. Morgan Kaufmann.

[Isaacs and Cohen, 1987] P. Isaacs and M. Cohen. Controlling Dynamic Simulation with Kinematic Constraints. *Computer Graphics*, 21(3), 1987.

[Ishizuka *et al.*, 1982] M. Ishizuka, K. Fu, and J. Yao. Speril: An expert system for damage assessment of existing structures. In *Proc. 6th Int. Conference on Pattern Recognition*, pages 932 – 937, 1982.

[Jaffar and Lassez, 1987] J. Jaffar and J.-L. Lassez. Constraint logic programming. In *Proceedings of the 14th ACM Symposium on Principles of Programming Languages*, pages 111–119, Munich, Germany, January 1987.

[Jameson and Wahlster, 1982] A. Jameson and W. Wahlster. User modelling in anaphora generation: Ellipsis and definite description. In *Proceedings of the 5th European Conference on Artificial Intelligence*, pages 222–227, 1982.

[Johnson and Feather, 1992] W. L. Johnson and M. S. Feather. Supporting modification and reuse of requirements knowledge. In SpringSymp-92 [1992], pages 59–62.

[Johnson, 1986] R. W. Johnson. Independence and bayesian updating methods. In Kanal and Lemmer [1986], pages 197–201.

[Kahn, 1979] K. Kahn. *Creation of Computer Animation from Story Descriptions*. PhD thesis, AI Lab, Massachusetts Institute of Technology, Cambridge, MA, January 1979.

[Kanal and Lemmer, 1986] L. Kanal and J. Lemmer, editors. *Uncertainty in Artificial Intelligence*. North-Holland, Amsterdam, 1986.

[Kanal *et al.*, 1989] L. Kanal, T. Levitt, and J. Lemmer, editors. *Uncertainty in Artificial Intelligence 3*. North-Holland, Amsterdam, 1989.

[Kansy, 1991] K. Kansy. Leitbeispiel Graphikdesigner. TASSO-Report 14, GMD, St. Augustin, 1991.

[Karp and Feiner, 1990] P. Karp and S. Feiner. Issues in the automated generation of animated presentations. In *Graphics Interface '90*, pages 39–48, Los Altos, CA, May 1990. Morgan Kaufmann.

[Katsuno and Mendelzon, 1991] H. Katsuno and A. O. Mendelzon. On the difference between updating a knowledge base and revising it. In Allen et al. [1991a], pages 387–394.

[Kautz and Ladkin, 1991] H. Kautz and P. B. Ladkin. Integrating metric and qualitative temporal reasoning. In AAAI-91 [1991], pages 241–246.

[Kautz and Selman, 1992] H. Kautz and B. Selman. Planning as satisfiability. In Neumann [1992], pages 359–363.

[Kautz, 1986] H. Kautz. The logic of persistence. In AAAI-86 [1986], pages 401–405.

[Kautz, 1991] H. A. Kautz. A formal theory of plan recognition and its implementation. In Allen et al. [1991b], chapter 2, pages 69–125.

[Kerpedjiev, 1992] S. M. Kerpedjiev. Automatic generation of multimodal weather reports from datasets. In *Proceedings of the 3rd ACL Conference on Applied Natural Language Processing (ANLP-92), Trento, Italy*, pages 48–55, 1992.

[Kobsa *et al.*, 1986] A. Kobsa, J. Allgayer, C. Reddig, N. Reithinger, D. Schmauks, K. Harbusch, and W. Wahlster. Combining Deictic Gestures and Natural Language for Reference Identification. In *Proceedings COLING-86*, pages 356–361, Bonn. FRG, 1986.

[Kobsa, 1991a] A. Kobsa. First experiences with the SB-ONE knowledge representation workbench in natural-language applications. *SIGART Bulletin*, 2(3):70–76, June 1991.

[Kobsa, 1991b] A. Kobsa. Utilizing knowledge: The components of the SB-ONE knowledge representation workbench. In Sowa [1991], pages 457–486.

[Kochanek and Bartels, 1984] D. Kochanek and R. Bartels. Interpolating Splines with Local Tension, Continuity, and Bias Control. *Computer Graphics*, 18(3):33–41. 1984.

[Koehler, 1992] J. Koehler. Towards a logical treatment of plan reuse. In AIPS-92 [1992], pages 285–286.

[Konolige, 1986] K. Konolige. *A Deduction Model of Belief.* Pitman, London, UK. 1986.

[Kopisch and Günter, 1992] M. Kopisch and A. Günter. Konfigurierung der Passagierkabine des AIRBUS A340 basierend auf einer Begriffshierarchie, einem Constraint-System und einer flexiblen Kontrolle. Prokon-bericht nr. 2, Universität Hamburg. 1992.

[Korein, 1985] J. Korein, editor. *A Geometric Investigation of Reach.* MIT press. Cambridge, MA, 1985.

[Kosslyn, 1980] S. Kosslyn. *Image and Mind.* Harvard University Press, Cambridge, Massachusetts, London, 1980.

[Kowalski and Sergot, 1986] R. Kowalski and M. Sergot. A logic-based calculus of events. *New Generation Computing*, 4:67–95, 1986.

[Kramer *et al.*, 1991] G. Kramer, J. Pabon, W. Keirouz, and R. Young. Geometric constraint satisfaction problems. In *Working Notes AAAI Spring Symposium 'Constraint-Based Reasoning'*, pages 242–251, Stanford University, CA, March 1991.

[Kruse and Siegel, 1991] R. Kruse and P. Siegel, editors. *Symbolic and Quantitative Approaches to Uncertainty, Proceedings of the European Conference ECSQAU.* Lecture Notes in Computer Science 548. Springer, Berlin, Germany, 1991.

[Kruse *et al.*, 1991] R. Kruse, E. Schwecke, and J. Heinsohn. *Uncertainty and Vagueness in Knowledge Based Systems: Numerical Methods.* Series Artificial Intelligence. Springer, Berlin, Germany, 1991. (491 pages).

[Kurlander and Feiner, 1991] D. Kurlander and S. Feiner. Inferring constraints from multiple snapshots. Technical report, Department of Computer Science, Columbia University, New York, NY, 1991.

[Kyburg, 1987] H. E. Kyburg, Jr. Bayesian and non-Bayesian evidential updating. *Artificial Intelligence*, 31:271–294, 1987.

[Langley, 1992] P. Langley. Systematic and nonsystematic search strategies. In AIPS-92 [1992], pages 145–152.

[Leler, 1988] W. Leler. *Constraint Programming Languages: Their Specification and Generation.* Addison-Wesley, Reading, MA, 1988.

[Lessel and Boley, 1987] M. Lessel and H. Boley. μ-UNIXPERT: Diagnosis of Printer Problems. In *Proceedings of the 12th Symposium on Operation Research*, Passau, 1987.

[Levesque and Brachman, 1987] H. J. Levesque and R. J. Brachman. Expressiveness and tractability in knowledge representation and reasoning. *Computational Intelligence*, 3:78–93, 1987.

[Lieberman, 1990] H. Lieberman. Communication of expert knowledge in graphic design. Visible Language Workshop, M.I.T. Media Laboratory, Cambridge, MA, 1990.

[Lifschitz, 1986] V. Lifschitz. On the semantics of STRIPS. In M. P. Georgeff and A. Lansky, editors, *Reasoning about Actions and Plans: Proceedings of the 1986 Workshop*, pages 1–9, Timberline, OR, June 1986. Morgan Kaufmann.

[Lifschitz, 1991] V. Lifschitz. Toward a metatheory of action. In Allen et al. [1991a], pages 376–386.

[Lin and Shoham, 1991] F. Lin and Y. Shoham. Provably correct theories of action. In AAAI-91 [1991].

[Lin and Shoham, 1992] F. Lin and Y. Shoham. Concurrent actions in the situation calculus. In AAAI-92 [1992].

[Lippman, 1980] A. Lippman. Movie-Maps: An Application of the Optical Video Disc to Computer Graphics. In *ACM SIGGRAPH'80*, 1980.

[Liu and Gammerman, 1988] X. Liu and A. Gammerman. A hybrid approach to deductive uncertain inference. *International Journal of Man-Machine Studies*, 28:671–681, 1988.

[London and Duisberg, 1985] R. London and R. Duisberg. Animating Programs Using Smalltalk. *Computer*, 18(8):61–71, August 1985.

[Maaß, 1992] W. Maaß. Constraint-basierte Plazierung in multimodalen Dokumenten am Beispiel des Layout-Managers in WIP. Master's thesis, Fachbereich Informatik, Universität Saarbrücken, Januar 1992.

[MacGregor, 1991a] R. MacGregor. The evolving technology of classification-based knowledge representation systems. In Sowa [1991], pages 385–400.

[MacGregor, 1991b] R. MacGregor. Inside the LOOM description classifier. *SIGART Bulletin*, 2(3):88–92, June 1991.

[MacGregor, 1992] R. MacGregor. What's needed to make a description logic a good KR citizen. In FallSymp-92 [1992], pages 53–55.

[Mackinlay, 1985] J. Mackinlay. *Automatic Design of Graphical Presentations*. PhD thesis, Dept. of Computer Science, Stanford University, Stanford, CA, 1985.

[MacNeil, 1990] R. MacNeil. Adaptive Persepectives: Case-based Reasoning with TYRO, the Graphics Designer's Apprentice. In *Proceedings of the IEEE Workshop on Visual Languages*, 1990.

[Magrez and Smets, 1989] P. Magrez and P. Smets. Epistemic necessity, possibility and truth. tools for dealing with imprecision and uncertainty in fuzzy knowledge-based systems. *International Journal of Approximate Reasoning*, 3:35–57, 1989.

[Marburger and Nebel, 1983] H. Marburger and B. Nebel. Natürlichsprachlicher Datenbankzugang mit HAM-ANS: Syntaktische Korrespondenz, natürlichsprachliche Quantifizierung und semantisches Modell des Diskursbereiches. In J. W. Schmidt, editor, *Sprachen für Datenbanken*, pages 26–41. Springer-Verlag, Berlin, Heidelberg, New York, 1983.

[Marks and Reiter, 1990] J. Marks and E. Reiter. Avoiding unwanted c onversational implicatures in text and graphics. In *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence*, pages 450–456, Boston, MA, July 1990.

[Marks, 1991] J. Marks. The competence of an automated graphic designer. In *Proceedings of the 1991 Long Island Conference on Artificial Intelligence and Computer Graphics, NYIT*, pages 53–61, New York, March 1991.

[Martin and Allen, 1991] N. G. Martin and J. F. Allen. A language for planning with statistics. In *Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence*, pages 220–227, UCLA, Los Angeles, CA, July 1991.

[Maulsby *et al.*, 1990] D. Maulsby, I. Witten, K. Kittlitz, and V. Franceschin. Inferring graphical procedures: The compleat metamouse. In *Proceedings of HCI (Human-Computer Interaction)*, 1990.

[Maybury, 1991] M. Maybury. Planning multimedia explanations using communicative acts. In *Proceedings of the 9th National Conference of the American Association for Artificial Intelligence*, pages 61–66, Anaheim, CA, July 1991.

[Maybury, 1992] M. Maybury, editor. *Intelligent Multimedia Interfaces.* AAAI Press, Menlo Park, CA, 1992. Forthcoming.

[Mays *et al.*, 1988] E. Mays, C. Apté, J. Griesmer, and J. Kastner. Experience with K-Rep: An object-centered knowledge representation language. In *Proceedings of IEEE CAIA-88*, pages 62–67, March 1988.

[Mays *et al.*, 1991] E. Mays, R. Dionne, and R. Weida. K-Rep system overview. *SIGART Bulletin*, 2(3):93–97, June 1991.

[McCarthy and Hayes, 1969] J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors. *Machine Intelligence*, volume 4, pages 463–502. Edinburgh University Press, Edinburgh. UK, 1969.

[McKeown and Feiner, 1990] K. McKeown and S. Feiner. Interactive multimedia explanation for equipment maintenance and repair. In *DARPA Speech and Language Workshop*. pages 42–47, 1990.

[Mertens, 1988] P. Mertens. Expertisesysteme als Erscheinungsform der betrieblichen Expertensysteme. Technical report, Universität Erlangen-Nürnberg, 1988. Arbeitspapiere Informatik-Forschungsgruppe VIII der Friedrich-Alexander-Universität Erlangen-Nürnberg.

[Merziger, 1992] G. Merziger. Approaches to abductive reasoning, an overview. Research Report RR-92-08, DFKI, Stuhlsatzenhausweg 3, 6600 Saarbrücken, February 1992.

[Mezei and Zivian, 1971] L. Mezei and A. Zivian. ARTA, an Interactive Animation System. In *IFIP Congress*, pages 429–434, 1971.

[Minton *et al.*, 1992] S. Minton, M. Drummond, J. L. Bresnia, and A. B. Philips. Total order *vs.* partial order planning: Factors influenciong performance. In Nebel et al. [1992], pages 83–92.

[Mitchell *et al.*, 1992] D. Mitchell, B. Selman, and H. Levesque. Hard and easy distributions of SAT problems. In AAAI-92 [1992], pages 459–465.

[Moore and Paris, 1989] J. Moore and C. Paris. Planning text for advisory dialogues. In *Proceedings of the 25th Annual Meeting of the ACL*, pages 203–211, 1989.

[Moore and Swartout, 1990] J. Moore and W. Swartout. Pointing: A Way toward Explanation Dialogue. In *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence*, pages 457–464, Boston, MA, July/August 1990.

[Morgenstern and Stein, 1988] L. Morgenstern and L. A. Stein. Why things go wrong: A formal theory of causal reasoning. In AAAI-88 [1988], pages 518–523.

[Moser, 1983] M. G. Moser. An overview of NIKL, the new implementation of KL-ONE. In *Research in Knowledge Representation and Natural Language Understanding, BBN Report No. 5421*, pages 7–26. Bolt, Beranek, and Newman Inc., Cambridge, MA, 1983.

[Müller-Brockmann, 1981] J. Müller-Brockmann, editor. *Grid Systems in Graphic Design.* Verlag Arthur Niggli, Niederteufen, Switzerland, 1981.

[Myers, 1991a] B. Myers. Text Formatting by Demonstration. In *Proceedings of CHI '91 (Human Factors in Computing Systems, Seattle, WA*, pages 251–256, 1991.

[Myers, 1991b] B. Myers. Using AI techniques to create user interfaces by example. In Sullivan and Tyler [1991], pages 385–402. Peridot.

[Neal and Shapiro, 1991] J. Neal and S. Shapiro. Intelligent multi-media interface technology. In Sullivan and Tyler [1991], pages 11–44. Cubricon.

[Nebel and Bäckström, 1992a] B. Nebel and C. Bäckström. On the computational complexity of temporal projection and plan validation. In AAAI-92 [1992], pages 748–753.

[Nebel and Bäckström, 1992b] B. Nebel and C. Bäckström. On the computational complexity of temporal projection, planning, and plan validation. *Artificial Intelligence*, 1992. To appear.

[Nebel and Koehler, 1992] B. Nebel and J. Koehler. Plan modification versus plan generation: A complexity-theoretic perspective. DFKI Research Report RR-92-48, German Research Center for Artificial Intelligence (DFKI), Saarbrücken, Germany, November 1992.

[Nebel and Peltason, 1991] B. Nebel and C. Peltason. Terminological reasoning and information management. In D. Karagianis, editor, *Information Systems and Artificial Intelligence: Integration Aspects*, pages 181–212. Springer-Verlag, Berlin, Heidelberg, New York, 1991.

[Nebel and Smolka, 1991] B. Nebel and G. Smolka. Attributive description formalisms ...and the rest of the world. In O. Herzog and C.-R. Rollinger, editors, *Text Understanding in LILOG*, volume 546 of *Lecture Notes in Artificial Intelligence*, pages 439–452. Springer-Verlag, Berlin, Heidelberg, New York, 1991.

[Nebel and von Luck, 1988] B. Nebel and K. von Luck. Hybrid reasoning in BACK. In Z. W. Ras and L. Saitta, editors, *Proceedings of the Third International Symposium on Methodologies for Intelligent systems*, pages 260–269, Torino, Italy, October 1988. North-Holland.

[Nebel *et al.*, 1991] B. Nebel, C. Peltason, and K. von Luck, editors. *International Workshop on Terminological Logics*, Schloß Dagstuhl, May 1991. Published as a KIT Report, TU Berlin, IWBS Report, IBM Germany, Stuttgart, and DFKI Document, DFKI, Kaiserlautern & Saarbrücken.

[Nebel *et al.*, 1992] B. Nebel, W. Swartout, and C. Rich, editors. *Principles of Knowledge Representation and Reasoning: Proceedings of the 3rd International Conference*. Cambridge, MA, October 1992. Morgan Kaufmann.

[Nebel, 1988] B. Nebel. Computational complexity of terminological reasoning in BACK. *Artificial Intelligence*, 34(3):371–383, April 1988.

[Nebel, 1990a] B. Nebel. *Reasoning and Revision in Hybrid Representation Systems*, volume 422 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, Heidelberg, New York, 1990.

[Nebel, 1990b] B. Nebel. Terminological reasoning is inherently intractable. *Artificial Intelligence*, 43:235–249, 1990.

[Nebel, 1991] B. Nebel. Terminological cycles: Semantics and computational properties. In Sowa [1991], pages 331–362.

[Neches *et al.*, 1991] R. Neches, R. Fikes, T. Finin, T. Gruber, R. Patil, T. Senator, and W. R. Swartout. Enabling technology for knowledge sharing. *The AI Magazine*, 12(3):36–56, 1991.

[Neiman, 1982] D. Neiman. Graphical animation from knowledge. In *Proceedings of the 2nd National Conference of the American Association for Artificial Intelligence*, pages 373–376, Pittsburgh, PA, July 1982.

[Nelson, 1985] G. Nelson. Juno, a constraint-based graphics system. *Proceedings of the SIGGRAPH '85*, 19(3):235–243, 1985.

[Neumann, 1992] B. Neumann, editor. *Proceedings of the 10th European Conference on Artificial Intelligence*, Vienna, Austria, August 1992. Wiley.

[Nilsson, 1986] N. J. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28:71–87, 1986.

[Nökel, 1989] K. Nökel. Convex relations between time intervals. In J. Rettie and K. Leidlmair, editors, *Proceedings der 5. Österreichischen Artificial Intelligence-Tagung*, pages 298–302. Springer-Verlag, Berlin, Heidelberg, New York, 1989.

[Oemig *et al.*, 1991] F. Oemig, A. B. Cremers, and G. Heyer. *Wissensbasierte Textverarbeitung: Schriftsatz und Typographie*. Deutscher Universitäts-Verlag, Wiesbaden, 1991.

[Ohlbach, 1992] H.-J. Ohlbach. Multi-dimensional $\mathcal{ALC}$. Submitted for publication, 1992.

[Ortony *et al.*, 1992] A. Ortony, J. Slack, and O. Stock, editors. *Communication from an Artificial Intelligence Perspective: Theoretical and Applied Issues*. Springer-Verlag, Berlin, Germany, 1992. In press.

[Ostertag and Hendler, 1987] E. Ostertag and J. Hendler. AIRS: An AI-based Ada reuse system. Tech. rep. cs-tr-2197, Computer Science Center, University of Maryland, 1987.

[Owsnicki-Klewe, 1988] B. Owsnicki-Klewe. Configuration as a consistency maintenance task. In W. Hoeppner, editor, *Proceedings of the 12th German Workshop on Artificial Intelligence (GWAI-88)*, pages 77–87. Springer, Berlin, Germany, 1988.

[Paaß, 1988] G. Paaß. Probabilistic logic. In Smets et al. [1988], pages 213–251.

[Paaß, 1992] G. Paaß. Vervollständigung von Layouts durch assoziative Verfahren. In K. Kansy and P. Wißkirchen, editors, *Proceedings GI-Workshop 'Innovative Programmiermethoden für Graphische Systeme'*, pages 30–44. Springer-Verlag, Berlin, Germany, Juni 1992.

[Padgham and Nebel, 1992] L. Padgham and B. Nebel. Combining classification and nonmonotonic inheritance reasoning: A first step. In FallSymp-92 [1992], pages 64–71.

[Patel-Schneider *et al.*, 1984] P. F. Patel-Schneider, R. J. Brachman, and H. J. Levesque. ARGON: Knowledge representation meets information retrieval. In *Proceedings of the 1st Conference on Artificial Intelligence Applications*, pages 280–286, Denver, Col., 1984.

[Patel-Schneider *et al.*, 1991] P. F. Patel-Schneider, D. L. McGuinness, R. J. Brachman, L. Alperin Resnick, and A. Borgida. The CLASSIC knowledge representation system: Guiding principles and implementation rational. *SIGART Bulletin*, 2(3):108–113, June 1991.

[Patel-Schneider, 1984] P. F. Patel-Schneider. Small can be beautiful in knowledge representation. In *Proceedings of the IEEE Workshop on Principles of Knowledge-Based Systems*, pages 11–16, Denver, Colo., 1984. An extended version including a KANDOR system description is available as AI Technical Report No. 37, Palo Alto, CA, Schlumberger Palo Alto Research, October 1984.

[Patel-Schneider, 1987] P. F. Patel-Schneider. *Decidable, Logic-Based Knowledge Representation.* PhD thesis, University of Toronto, Toronto, Ont., May 1987. Computer Science Department, Technical Report 201/87.

[Patel-Schneider, 1989] P. F. Patel-Schneider. Undecidability of subsumption in NIKL. *Artificial Intelligence*, 39(2):263–272, June 1989.

[Patel-Schneider, 1992] P. F. Patel-Schneider. Defaults and descriptions. In FallSymp-92 [1992], pages 72–73.

[Patil *et al.*, 1992] R. S. Patil, R. E. Fikes, P. F. Patel-Schneider, D. McKay, T. Finin, T. Gruber, and R. Neches. The DARPA knowledge sharing effort: Progress report. In Nebel et al. [1992], pages 777–788.

[Pearl, 1986] J. Pearl. Fusion, propagation and structuring in belief networks. *Artificial Intelligence*, 29:241–288, 1986.

[Pearl, 1988] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, Cal., 1988.

[Peltason, 1991] C. Peltason. The BACK system – an overview. *SIGART Bulletin*, 2(3):114–119, June 1991.

[Peng and Reggia, 1990] Y. Peng and J. A. Reggia. *Abductive Inference models for diagnostic problem-solving.* Artificial Intelligence. Springer, Berlin, Germany, 1990.

[Poole, 1988] D. Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36:27–47, 1988.

[Poole, 1991] D. Poole. Representing Bayesian networks within probabilistic Horn abduction. In *Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence*, pages 271–278, UCLA, Los Angeles, CA, July 1991.

[Prieto-Diaz and Freeman, 1987] R. Prieto-Diaz and P. Freeman. Classifying software for reusability. *IEEE Software*, 4:6–16, 1987.

[Profitlich, 1989] H.-J. Profitlich. *SB-ONE Benutzerhandbuch.* Department of Computer Science, Universität des Saarlandes, Saarbrücken, Germany, February 1989.

[Profitlich, 1990] H.-J. Profitlich. SB-ONE: Ein Wissensrepräsentationssystem basierend auf KL-ONE. Technical Memo 43, Department of Computer Science, Universität des Saarlandes, Saarbrücken, Germany, March 1990.

[Quantz and Royer, 1992] J. Quantz and V. Royer. A preference semantics for defaults in terminological logics. In Nebel et al. [1992], pages 294–305.

[Quinlan, 1983a] J. Quinlan. Consistency and plausible reasoning. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, pages 137–144, Karlsruhe, West Germany, 1983.

[Quinlan, 1983b] J. Quinlan. INFERNO, a cautious approach to uncertain inference. *Computer Journal*, 26(3):255–269, 1983.

[Quinlan, 1985] J. Quinlan. Internal consistency in plausible reasoning systems. *New Generation Computing*, 3:157–180, 1985.

[Randell *et al.*, 1992] D. Randell, Z. Cui, and A. Cohn. A spatial logic based on regions and connection. In *Proceedings of the KR'92*, pages 165–176, 1992.

[Reeves, 1981] W. Reeves. Inbetweening for Computer Animation Utilizing Moving Point Constraints. *Computer Graphics*, 15(3):263–269, August 1981.

[Reiter *et al.*, 1992] E. Reiter, C. Mellish, and J. Levine. Automatic Generation of On-line Documentation in the IDAS Project. In *Proceedings of the Third Conference on Applied Natural Language Processing*, pages 64–71, Trento, Italy, 1992.

[Reiter, 1980] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1):81–132, April 1980.

[Reynolds, 1982] C. Reynolds. Computer Animation with Scripts and Actors. *Computer Graphics*, 16(3):289-296, July 1982.

[Rist and André, 1990] T. Rist and E. André. Wissensbasierte Perspektivenwahl für die automatische Erzeugung von 3D-Objektdarstellungen. In K. Kansy and P. Wißkirchen, editors, *Graphik und KI, IFB 239*, pages 48–57. Springer, Berlin, 1990.

[Rist and André, 1992a] T. Rist and E. André. From Presentation Tasks to Pictures: Towards a Computational Approach to Automatic Graphics Design. In *Proceedings of the 10th European Conference on Artificial Intelligence*, Vienna, Austria, August 1992.

[Rist and André, 1992b] T. Rist and E. André. Incorporating graphics design and realization into the multimodal presentation system wip. In Costabile et al. [1992]. Proceedings of the AVI '92 Workshop, Rome, Italy, May 27-29.

[Roth and Heffley, 1992] S. Roth and W. Heffley. Intelligent multimedia presentation systems: Research and principles. In Maybury [1992]. Forthcoming.

[Roth *et al.*, 1991] S. Roth, J. Mattis, and X. Mesnard. Graphics and natural language as components of automatic explanation. In Sullivan and Tyler [1991], pages 207–240. SAGE.

[Rubin, 1989] B. Rubin. Constraint-Based Cinematic Editing. Master Thesis, MIT Media Arts and Science Section, Cambridge, MA, June 1989.

[Ruspini, 1987] E. Ruspini. Epistemic logics, probability and the calculus of evidence. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, pages 924–931, Milan, Italy, 1987.

[S-Dynamics, 1985] S-Dynamics. Reference Manual. Symbolics Inc., Cambridge, MA, 1985.

[Saffiotti, 1990] A. Saffiotti. A hybrid framework for representing uncertain knowledge. In *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence*, pages 653–658, Boston, Mass., 1990.

[Sandewall and Rönnquist, 1986] E. Sandewall and R. Rönnquist. A representation of action structures. In AAAI-86 [1986], pages 89–97.

[Sandewall, 1989] E. Sandewall. Filter preferential entailment for the logic of action in almost continuous worlds. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 894–899, Detroit, MI, August 1989. Morgan Kaufmann.

[Sandewall, 1992a] E. Sandewall. Features and fluents. Research Report LiTH-IDA-R-92-30, Department of Computer and Information Science. Linköping University, Linköping, Sweden, September 1992.

[Sandewall, 1992b] E. Sandewall. The range of applicability of nonmonotonic logics for the inertia problem. Submitted for publication, 1992.

[Saunders, 1989] D. Saunders. Improvements to INFERNO. In A. Cohn, editor, *Proc. of the 7th Conference of the Society for the Study of Artificial Intelligence and Simulation of Behaviour*, pages 105–112. Pitman and Kaufmann, Brighton, England, May 1989.

[Schaerf, 1992] A. Schaerf. On the role of subsumption algorithms in concept description languages. In FallSymp-92 [1992], pages 86–97.

[Schauder, 1990] A. Schauder. Inkrementelle syntaktische generierung natürlicher sprache mit tree adjoining grammars. Master's thesis, Dept. of Computer Science, University of Saarbrücken, 1990.

[Schild, 1991] K. Schild. A tense-logical extension of terminological logics. KIT Report 92, Department of Computer Science, Technische Universität Berlin, Berlin, Germany, September 1991.

[Schirra, 1992] J. Schirra. A Contribution to Reference Semantics of Spatial Prepositions: The Visualization Problem and its Solution in VITRA. Mouton de Gruyter, Berlin, Germany, 1992.

[Schmidt-Schauß, 1989] M. Schmidt-Schauß. Subsumption in KL-ONE is undecidable. In Brachman et al. [1989], pages 421–431.

[Schmiedel, 1989] A. Schmiedel. A temporal constraint handler for the back system. KIT Report 70, Department of Computer Science, Technische Universität Berlin. Berlin, Germany, November 1989.

[Schmiedel, 1990] A. Schmiedel. A temporal terminological logic. In AAAI-90 [1990], pages 640–645.

[Schmolze and Mark, 1991] J. G. Schmolze and W. S. Mark. The NIKL experience. *Computational Intelligence*, 6:48–69, 1991.

[Schmolze and Woods, 1990] J. G. Schmolze and W. A. Woods. The KL-ONE family. Technical Report TR-20-90, Aiken Computer Laboratory, Harvard University. Cambridge, MA, August 1990. To appear in *Computers & Mathematics with Applications*.

[Schmolze, 1989] J. G. Schmolze. The language and semantics of NIKL. Technical Report 89-4, Department of Computer Science, Tufts University, Medford, MA, September 1989.

[Schweizer and Sklar, 1961] B. Schweizer and A. Sklar. Associative functions and statistical triangle inequalities. *Publ. Math. Debrecen*, 8:169–186, 1961.

[Schweizer and Sklar, 1983] B. Schweizer and A. Sklar. *Probabilistic Metric Spaces.* North-Holland, Amsterdam, Holland, 1983.

[Seligmann and Feiner, 1991] D. Seligmann and S. Feiner. Automated generation of intent-based 3d illustrations. *Computer Graphics*, 25(3), July 1991.

[Selman and Kautz, 1991] B. Selman and H. Kautz. Knowledge compilation using horn approximations. In AAAI-91 [1991], pages 904–909.

[Selman and Kautz, 1992] B. Selman and H. Kautz. Domain-independent extensions to GSAT: Solving large structured satisfiability problems. AT&T Bell Labs., submitted for publication, 1992.

[Selman and Levesqüe, 1990] B. Selman and H. Levesque. Abductive and default reasoning: A computational core. In AAAI-90 [1990], pages 343–348.

[Selman *et al.*, 1992] B. Selman, H. Levesque, and D. Mitchell. A new method for solving hard satisfiability problems. In AAAI-92 [1992], pages 440–446.

[Selman, 1992] B. Selman. Near-optimal plans, tractability, and reactivity. AT&T Bell Labs., Submitted for publication, 1992.

[Shafer and Pearl, 1990] G. Shafer and J. Pearl, editors. *Readings in Uncertain Reasoning*. Morgan Kaufmann, Los Altos, Cal., 1990.

[Shafer *et al.*, 1987] G. Shafer, P. Shenoy, and K. Mellouli. Propagating belief functions in qualitative markov trees. *International Journal of Approximate Reasoning*, 1:349–400, 1987.

[Shafer, 1976] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, 1976.

[Shanahan, 1989] M. Shanahan. Prediction is deduction but explanation is abduction. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 1055–1060, 1989.

[Shelly, 1982] K. Shelly. Path Specification and Path Coherence. *Computer Graphics*, 16(3):157–166, July 1982.

[Shenoy and Shafer, 1987] P. Shenoy and G. Shafer. Propagating belief functions with local computations. *IEEE Expert*, 1:43–52, 1987.

[Shoham, 1986] Y. Shoham. Chronological ignorance: Time, nonmonotonicity, necessity, and causal theories. In AAAI-86 [1986], pages 389–393.

[Shortliffe and Buchanan, 1975] E. Shortliffe and B. Buchanan. A model for inexact reasoning in medicine. *Mathematical Biosciences*, 23:351–379, 1975.

[Shortliffe, 1976] E. Shortliffe. *Computer Based Medical Consultations: MYCIN*. Elsevier, New York, 1976.

[Sigart Bulletin, 1991] Special issue on implemented knowledge representation and reasoning systems, June 1991.

[Smets and Magrez, 1987] P. Smets and P. Magrez. Implication in fuzzy logic. *International Journal of Approximate Reasoning*, 1:327–347, 1987.

[Smets *et al.*, 1988] P. Smets, E. Mamdani, D. Dubois, and H. Prade, editors. *Non-Standard Logics for Automated Reasoning*. Academic Press, New York, N.Y., 1988.

[Smolka, 1991] G. Smolka. Residuation and guarded rules for constraint logic programming. Technical report, German Research Center for Artificial Intelligence, Saarbrücken, Germany, March 1991.

[Soetopo, 1992] D. Soetopo. "automatische verarbeitung typographischen wissens". Master's thesis, Fachbereich Informatik, Universität Saarbrücken, 1992.

[Sondheimer and Nebel, 1986] N. K. Sondheimer and B. Nebel. A logical-form and knowledge-base design for natural language generation. In AAAI-86 [1986], pages 612–618.

[Sowa, 1991] J. F. Sowa, editor. *Principles of Semantic Networks.* Morgan Kaufmann, San Mateo, CA, 1991.

[Spiegelhalter, 1989] D. J. Spiegelhalter. A unified approach to imprecision and sensitivity of beliefs in expert systems. In Kanal et al. [1989], pages 199–208.

[SpringSymp-92, 1992] Working notes of the AAAI Spring Symposium "computational considerations in supporting incremental modification and reuse". Stanford University, March 1992.

[Steketee and Badler, 1985] S. Steketee and N. Badler. Parametric Keyframe Interpolation Incorporating Kinetic Adjustment and Phrasing Control. *Computer Graphics*, 19(3):255–262, 1985.

[Steve, 1986] G. Steve. Probabilistic inferential engines in expert systems: How should the strength of rules be expressed? In *Proceedings of the 1st International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU'86)*, pages 451–455, Paris, France, 1986.

[Stock, 1991] O. Stock. Natural language and exploration of an information space: the ALFresco interactive system. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, pages 972–978, Sydney, Australia, August 1991.

[Strauss, 1988] P. Strauss. *BAGS: The Brown Animation Generation System.* PhD thesis, Dept. of Computer Science, Brown University, 1988.

[Strothotte, 1989] T. Strothotte. Pictures in Advice-Giving Dialog Systems: From Knowledge Representation to the User Interface. In *Graphics Interface 89*, pages 94–99. 1989.

[Sukaviriya and Foley, 1990] P. Sukaviriya and J. D. Foley. Coupling a ui framework with automatic generation of context-sensitive animated help. In *Proceedings of the ACM Symp. on User Interfaces Sooftware and Technology*, pages 152–166, Snowbird, UT, October 1990.

[Sullivan and Tyler, 1991] J. Sullivan and S. Tyler, editors. *Intelligent User Interfaces.* Frontier Series. ACM Press, New York, NY, 1991.

[Sussman and Steele, 1980] G. Sussman and G. Steele. Constraints – a language for expressing almost-hierachical descriptions. *Artificial Intelligence*, 14(1):1–39, 1980.

[Sutherland, 1963] I. Sutherland. Sketchpad: A man-machine graphical communication system. In *IFIPS Proceedings of the Spring Joint Computer Conference*, pages 329–345, 1963.

[Swartout and Neches, 1986] W. R. Swartout and R. Neches. The shifting terminological space: An impediment to evolvability. In AAAI-86 [1986], pages 936–941.

[Takashima *et al.*, 1987] Y. Takashima, H. Shimatzu, and M. Tomono. Story Driven Animation. In *ACM SIGCHI and Graphics Interface 87 Joint Conference*, 1987.

[Tenenberg, 1989] J. D. Tenenberg. Inheritance in automated planning. In Brachman et al. [1989], pages 475–485.

[Thagard, 1991] P. Thagard. Review of [Peng and Reggia, 1990]. *SIGART Newsletter*, 2(1):72–75, 1991.

[Thalmann and Thalmann, 1990] N. M. Thalmann and D. Thalmann. *Computer Animation. Theory and Practice. 2nd edition.* Computer Science Workbench. Springer-Verlag, Berlin, Germany, 1990.

[Thies and Berger, 1992] M. A. Thies and F. Berger. Plan-based graphical help in object-oriented user interfaces. In *Proceedings of AVI'92 (Advanced Visual Interfaces)*, Rome, Italy, May 1992.

[Tufte, 1991] E. Tufte. *Envisioning Information.* Graphics Press, Cheshire, CT, 1991.

[van Beek and Cohen, 1990] P. van Beek and R. Cohen. Exact and approximate reasoning about temporal relations. *Computational Intelligence*, 6:132–144, 1990.

[van Beek, 1990] P. van Beek. Reasoning about qualitative temporal information. In AAAI-90 [1990], pages 728–734.

[van Melle, 1980] W. J. van Melle. *System Aids in Constructing Consultation Programs.* UMI Research Press, Ann Arbor, Michigan, 1980.

[van Wyk, 1982] C. van Wyk. A high-level language for specifying pictures. *ACM Transactions on Graphics*, 1(2):163–182, 1982.

[Veloso, 1992] M. M. Veloso. Automatic storage, retrieval, and replay of multiple cases using derivational analogy in PRODIGY. In SpringSymp-92 [1992], pages 131–136.

[Vilain *et al.*, 1989] M. B. Vilain, H. A. Kautz, and P. G. van Beek. Constraint propagation algorithms for temporal reasoning: A revised report. In D. S. Weld and J. de Kleer, editors, *Readings in Qualitative Reasoning about Physical Systems*, pages 373–381. Morgan Kaufmann, San Mateo, CA, 1989.

[Vilain, 1983] M. B. Vilain. Assertions in NIKL. In *Research in Knowledge Representation and Natural Language Understanding, BBN Report, No. 5421*, pages 45–79. Bolt, Beranek, and Newman Inc., Cambridge, MA, 1983.

[Vilain, 1985] M. B. Vilain. The restricted language architecture of a hybrid representation system. In IJCAI-85 [1985], pages 547–551.

[von Luck *et al.*, 1987] K. von Luck, B. Nebel, C. Peltason, and A. Schmiedel. The anatomy of the BACK system. KIT Report 41, Department of Computer Science, Technische Universität Berlin, Berlin, Germany, January 1987.

[Wahlster *et al.*, 1983] W. Wahlster, H. Marburger, A. Jameson, and S. Busemann. Over-answering Yes-No questions: Extended Responses in a NL Interface to a Vision System. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, pages 643–646, Karlsruhe, FRG, 1983.

[Wahlster *et al.*, 1989] W. Wahlster, E. André, M. Hecking, and T. Rist. WIP: Knowledge-Based Presentation of Information. Project Proposal, German Research Center for Artificial Intelligence, Saarbrücken, Germany, November 1989.

[Wahlster *et al.*, 1991] W. Wahlster, E. André, W. Graf, and T. Rist. Designing illustrated texts: How language production is influenced by graphics generation. In *Proceedings of the 5th Conference of the European Chapter of the Association for Computational Linguistics*, pages 8–14. Springer-Verlag, Berlin, Germany, April 1991. Also DFKI Research Report RR-91-05.

[Wahlster *et al.*, 1992a] W. Wahlster, E. André, S. Bandyopadhyay, W. Graf, and T. Rist. WIP: The coordinated generation of multimodal presentations from a common representation. In Ortony et al. [1992], pages 121–144. Also DFKI Research Report RR-91-08.

[Wahlster *et al.*, 1992b] W. Wahlster, E. André, W. Finkler, H.-J. Profitlich, and T. Rist. Plan-based integration of natural language and graphics generation. *Artificial Intelligence, Special Issue on Natural Language Processing*, 1992. Forthcoming.

[Wahlster, 1991] W. Wahlster. User and Discourse Models for Multimodal Communication. In J. Sullivan and S. Tyler, editors, *Intelligent User Interfaces*, pages 45–67. Addison-Wesley, Reading, 1991.

[Weida and Litman, 1992] R. Weida and D. Litman. Terminological reasoning with constraint networks and an application to plan recognition. In Nebel et. al. [1992], pages 282–293.

[Wellman, 1990] M. P. Wellman. *Formulation of Tradeoffs in Planning Under Uncertainty.* Morgan Kaufmann, San Mateo, CA, 1990.

[Wiecha and Boies, 1990] C. Wiecha and S. Boies. Generating User Interfaces: Principles and Use of its Style Rules. In *Proceedings of UIST '90 (ACM Symposium on User Interface Software and Technology), Snowbird, UT*, pages 21–30, 1990.

[Wilhelms, 1986] J. Wilhelms. Virya - A Motion Editor for Kinematic and Dynamic Animation. In *Proc. of Graphics Interface '86*, pages 141–146, Vancouver, 1986.

[Wilhelms, 1987] J. Wilhelms. Towards automatic motion control. *IEEE Computer Graphics and Animation*, 7(4):11–22, April 1987.

[Wilson *et al.*, 1992] M. Wilson, D. Sedlock, J.-L. Binot, and P. Falzon. An Architecture For Multimodal Dialogue. In *Proceedings of the Second Vencona Workshop on Multi-Modal Dialogue*, Italy, September 1992.

[Winslett, 1988] M. S. Winslett. Reasoning about action using a possible models approach. In AAAI-88 [1988], pages 89–93.

[Witkin *et al.*, 1987] A. Witkin, K. Fleisher, and A. Barr. Energy Constraints on Parameterized Models. *Computer Graphics*, 21(3), 1987.

[Woods and Somerville, 1988] M. Woods and I. Somerville. An information system for software components. *Proc. of ACM SIGIR Forum*, 22(3), 1988.

[Woods, 1975] W. A. Woods. What's in a link: Foundations for semantic networks. In D. G. Bobrow and A. M. Collins, editors, *Representation and Understanding: Studies in Cognitive Science*, pages 35–82. Academic Press, New York, NY, 1975.

[Yager, 1980] R. Yager. Aspects of possibilistic uncertainty. *International Journal of Man-Machine Studies*, 12:283–298, 1980.

[Yen and Bonissone, 1990] J. Yen and P. Bonissone. Extending term subsumption systems for uncertainty management. In *Proceedings of the 6th Conference on Uncertainty in Artificial Intelligence*, Cambridge, Mass., July 1990.

[Yen, 1989] J. Yen. GERTIS: a Dempster-Shafer approach to diagnosing hierarchical hypotheses. *Communications of the ACM*, 32(5):573–585, 1989.

[Yen, 1991] J. Yen. Generalizing term subsumption languages to fuzzy logic. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, Sydney, Australia, 1991.

[Zadeh, 1965] L. Zadeh. Fuzzy sets. *Information and Control*, 8:338–353, 1965.

[Zadeh, 1978] L. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1:3–28, 1978.

[Zadeh, 1984] L. Zadeh. A mathematical theory of evidence. book review. *The AI Magazine*, 5:81–83, 1984.

[Zadeh, 1986] L. Zadeh. A simple view of the Dempster-Shafer theory of evidence and its implication for the rule of combination. *The AI Magazine*, 7:85–90, 1986.

[Zdybel *et al.*, 1981] F. Zdybel, N. Greenfeld, and M.Yonke. An Information Presentation System. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, pages 978–984, 1981.

[Zeltzer, 1991] D. Zeltzer. Task level graphical simulation: Abstraction, representation, and control. In Badler et al. [1991a], pages 3–33.

Deutsches
Forschungszentrum
für Künstliche
Intelligenz GmbH

-Bibliothek, Information
und Dokumentation (BID)-
PF 2080
67608 Kaiserslautern
FRG

Telefon (0631) 205-3506
Telefax (0631) 205-3210
e-mail
dfkibib@dfki.uni-kl.de
WWW
http://www.dfki.uni-sb.de/dfkibib

Veröffentlichungen des DFKI

Die folgenden DFKI Veröffentlichungen sowie die aktuelle Liste von allen bisher erschienenen Publikationen können von der oben angegebenen Adresse oder (so sie als per ftp erhaeltlich angemerkt sind) per anonymous ftp von ftp.dfki.uni-kl.de (131.246.241.100) im Verzeichnis pub/Publications bezogen werden. Die Berichte werden, wenn nicht anders gekennzeichnet, kostenlos abgegeben.

DFKI Publications

The following DFKI publications or the list of all published papers so far are obtainable from the above address or (if they are marked as obtainable by ftp) by anonymous ftp from ftp.dfki.uni-kl.de (131.246.241.100) in the directory pub/Publications.
The reports are distributed free of charge except where otherwise noted.

DFKI Research Reports

1995

RR-95-20
Hans-Ulrich Krieger
Typed Feature Structures, Definite Equivalences, Greatest Model Semantics, and Nonmonotonicity
27 pages

RR-95-19
Abdel Kader Diagne, Walter Kasper, Hans-Ulrich Krieger
Distributed Parsing With HPSG Grammar
20 pages

RR-95-18
Hans-Ulrich Krieger, Ulrich Schäfer
Efficient Parameterizable Type Expansion for Typed Feature Formalisms
19 pages

RR-95-17
Hans-Ulrich Krieger
Classification and Representation of Types in TDL
17 pages

RR-95-16
Martin Müller, Tobias Van Roy
Title not set
0 pages

Note: The author(s) were unable to deliver this document for printing before the end of the year. It will be printed next year.

RR-95-15
Joachim Niehren, Tobias Van Roy
Title not set
0 pages

Note: The author(s) were unable to deliver this document for printing before the end of the year. It will be printed next year.

RR-95-14
Joachim Niehren
Functional Computation as Concurrent Computation
50 pages

RR-95-13
Werner Stephan, Susanne Biundo
Deduction-based Refinement Planning
14 pages

RR-95-12
Walter Hower, Winfried H. Graf
Research in Constraint-Based Layout, Visualization. CAD, and Related Topics: A Bibliographical Survey
33 pages

RR-95-11
Anne Kilger, Wolgang Finkler
Incremental Generation for Real-Time Applications
47 pages

RR-95-10
Gert Smolka
The Oz Programming Model
23 pages

RR-95-09
M. Buchheit, F. M. Donini, W. Nutt, A. Schaerf
A Refined Architecture for Terminological Systems: Terminology = Schema + Views
71 pages

RR-95-08
Michael Mehl, Ralf Scheidhauer, Christian Schulte
An Abstract Machine for Oz
23 pages

RR-95-07
Francesco M. Donini, Maurizio Lenzerini, Daniele Nardi, Werner Nutt
The Complexity of Concept Languages
57 pages

RR-95-06
Bernd Kiefer, Thomas Fettig
FEGRAMED
An interactive Graphics Editor for Feature Structures
37 pages

RR-95-05
Rolf Backofen, James Rogers, K. Vijay-Shanker
A First-Order Axiomatization of the Theory of Finite Trees
35 pages

RR-95-04
M. Buchheit, H.-J. Bürckert, B. Hollunder, A. Laux, W. Nutt,
M. Wójcik
Task Acquisition with a Description Logic Reasoner
17 pages

RR-95-03
Stephan Baumann, Michael Malburg, Hans-Guenther Hein, Rainer Hoch,
Thomas Kieninger, Norbert Kuhn
Document Analysis at DFKI
Part 2: Information Extraction
40 pages

RR-95-02
Majdi Ben Hadj Ali, Frank Fein, Frank Hoenes, Thorsten Jaeger,
Achim Weigel
Document Analysis at DFKI
Part 1: Image Analysis and Text Recognition
69 pages

RR-95-01
Klaus Fischer, Jörg P. Müller, Markus Pischel
Cooperative Transportation Scheduling
an application Domain for DAI
31 pages

1994

RR-94-39
Hans-Ulrich Krieger
Typed Feature Formalisms as a Common Basis for Linguistic Specification.
21 pages

RR-94-38
Hans Uszkoreit, Rolf Backofen, Stephan Busemann, Abdel Kader Diagne,
Elizabeth A. Hinkelman, Walter Kasper, Bernd Kiefer, Hans-Ulrich Krieger,
Klaus Netter, Günter Neumann, Stephan Oepen, Stephen P. Spackman.
DISCO–An HPSG-based NLP System and its Application for Appointment Scheduling.
13 pages

RR-94-37
Hans-Ulrich Krieger, Ulrich Schäfer
TDL - A Type Description Language for HPSG, Part 1: Overview.
54 pages

RR-94-36
Manfred Meyer
Issues in Concurrent Knowledge Engineering. Knowledge Base and Knowledge Share Evolution.
17 pages

RR-94-35
Rolf Backofen
A Complete Axiomatization of a Theory with Feature and Arity Constraints
49 pages

RR-94-34
Stephan Busemann, Stephan Oepen, Elizabeth A. Hinkelman.
Günter Neumann, Hans Uszkoreit
COSMA – Multi-Participant NL Interaction for Appointment Scheduling
80 pages

RR-94-33
Franz Baader, Armin Laux
Terminological Logics with Modal Operators
29 pages

RR-94-31
Otto Kühn, Volker Becker, Georg Lohse, Philipp Neumann
Integrated Knowledge Utilization and Evolution for the Conservation of Corporate Know-How
17 pages

RR-94-23
Gert Smolka
The Definition of Kernel Oz
53 pages

RR-94-20
Christian Schulte, Gert Smolka, Jörg Würtz
Encapsulated Search and Constraint Programming in Oz
21 pages

RR-94-19
Rainer Hoch
Using IR Techniques for Text Classification in Document Analysis
16 pages

RR-94-18
Rolf Backofen, Ralf Treinen
How to Win a Game with Features
18 pages

RR-94-17
Georg Struth
Philosophical Logics—A Survey and a Bibliography
58 pages

RR-94-16
Gert Smolka
A Foundation for Higher-order Concurrent Constraint Programming
26 pages

RR-94-15
Winfried H. Graf, Stefan Neurohr
Using Graphical Style and Visibility Constraints for a Meaningful Layout in Visual Programming Interfaces
20 pages

RR-94-14
Harold Boley, Ulrich Buhrmann, Christof Kremer
Towards a Sharable Knowledge Base on Recyclable Plastics
14 pages

RR-94-13
Jana Koehler
Planning from Second Principles—A Logic-based Approach
49 pages

RR-94-12
Hubert Comon, Ralf Treinen
Ordering Constraints on Trees
34 pages

RR-94-11
Knut Hinkelmann
A Consequence Finding Approach for Feature Recognition in CAPP
18 pages

RR-94-10
Knut Hinkelmann, Helge Hintze
Computing Cost Estimates for Proof Strategies
22 pages

RR-94-08
Otto Kühn, Björn Höfling
Conserving Corporate Knowledge for Crankshaft Design
17 pages

RR-94-07
Harold Boley
Finite Domains and Exclusions as First-Class Citizens
25 pages

RR-94-06
Dietmar Dengler
An Adaptive Deductive Planning System
17 pages

RR-94-05
Franz Schmalhofer, J. Stuart Aitken, Lyle E. Bourne jr.
Beyond the Knowledge Level: Descriptions of Rational Behavior for Sharing and Reuse
81 pages

RR-94-03
Gert Smolka
A Calculus for Higher-Order Concurrent Constraint Programming with Deep Guards
34 pages

RR-94-02
Elisabeth André, Thomas Rist
Von Textgeneratoren zu Intellimedia-Präsentationssystemen
22 Seiten

RR-94-01
Elisabeth André, Thomas Rist
Multimedia Presentations: The Support of Passive and Active Viewing
15 pages

1993

RR-93-48
Franz Baader, Martin Buchheit, Bernhard Hollunder
Cardinality Restrictions on Concepts
20 pages

RR-93-46
Philipp Hanschke
A Declarative Integration of Terminological, Constraint-based, Data-driven, and Goal-directed Reasoning
81 pages

RR-93-45
Rainer Hoch
On Virtual Partitioning of Large Dictionaries for Contextual Post-Processing to Improve Character Recognition
21 pages

RR-93-44
Martin Buchheit, Manfred A. Jeusfeld, Werner Nutt, Martin Staudt
Subsumption between Queries to Object-Oriented Databases
36 pages

RR-93-43
M. Bauer, G. Paul
Logic-based Plan Recognition for Intelligent Help Systems
15 pages

RR-93-42
Hubert Comon, Ralf Treinen
The First-Order Theory of Lexicographic Path Orderings is Undecidable
9 pages

RR-93-41
Winfried H. Graf
LAYLAB: A Constraint-Based Layout Manager for Multimedia Presentations
9 pages

RR-93-40
Francesco M. Donini, Maurizio Lenzerini, Daniele Nardi, Werner Nutt, Andrea Schaerf
Queries, Rules and Definitions as Epistemic Statements in Concept Languages
23 pages

RR-93-38
Stephan Baumann
Document Recognition of Printed Scores and Transformation into MIDI
24 pages

RR-93-36
Michael M. Richter, Bernd Bachmann, Ansgar Bernardi, Christoph Klauck, Ralf Legleitner, Gabriele Schmidt
Von IDA bis IMCOD: Expertensysteme im CIM-Umfeld
13 Seiten

RR-93-35
Harold Boley, François Bry, Ulrich Geske (Eds.)
Neuere Entwicklungen der deklarativen KI-Programmierung — *Proceedings*
150 Seiten

Note: This document is available for a nominal charge of 25 DM (or 15 US-$).

RR-93-34
Wolfgang Wahlster
Verbmobil Translation of Face-To-Face Dialogs
10 pages

RR-93-33
Bernhard Nebel, Jana Koehler
Plan Reuse versus Plan Generation: A Theoretical and Empirical Analysis
33 pages

RR-93-32
David R. Traum, Elizabeth A. Hinkelman
Conversation Acts in Task-Oriented Spoken Dialogue
28 pages

RR-93-31
Elizabeth A. Hinkelman, Stephen P. Spackman
Abductive Speech Act Recognition, Corporate Agents and the COSMA System
34 pages

RR-93-30
Stephen P. Spackman, Elizabeth A. Hinkelman
Corporate Agents
14 pages

RR-93-29
Armin Laux
Representing Belief in Multi-Agent Worlds via Terminological Logics
35 pages

RR-93-28
Hans-Ulrich Krieger, John Nerbonne, Hannes Pirker
Feature-Based Allomorphy
8 pages

RR-93-27
Hans-Ulrich Krieger
Derivation Without Lexical Rules
33 pages

RR-93-26
Jörg P. Müller, Markus Pischel
The Agent Architecture InteRRaP: Concept and Application
99 pages

RR-93-25
Klaus Fischer, Norbert Kuhn
A DAI Approach to Modeling the Transportation Domain
93 pages

RR-93-24
Rainer Hoch, Andreas Dengel
Document Highlighting — Message Classification in Printed Business Letters
17 pages

RR-93-23
Andreas Dengel, Ottmar Lutzy
Comparative Study of Connectionist Simulators
20 pages

RR-93-22
Manfred Meyer, Jörg Müller
Weak Looking-Ahead and its Application in Computer-Aided Process Planning
17 pages

RR-93-20
Franz Baader, Bernhard Hollunder
Embedding Defaults into Terminological Knowledge Representation Formalisms
34 pages

RR-93-18
Klaus Schild
Terminological Cycles and the Propositional μ-Calculus
32 pages

RR-93-17
Rolf Backofen
Regular Path Expressions in Feature Logic
37 pages

RR-93-16
Gert Smolka, Martin Henz, Jörg Würtz
Object-Oriented Concurrent Constraint Programming in Oz
17 pages

RR-93-15
Frank Berger, Thomas Fehrle, Kristof Klöckner, Volker Schölles,
Markus A. Thies, Wolfgang Wahlster
PLUS - Plan-based User Support Final Project Report
33 pages

RR-93-14
Joachim Niehren, Andreas Podelski, Ralf Treinen
Equational and Membership Constraints for Infinite Trees
33 pages

RR-93-13
Franz Baader, Karl Schlechta
A Semantics for Open Normal Defaults via a Modified Preferential Approach
25 pages

RR-93-12
Pierre Sablayrolles
A Two-Level Semantics for French Expressions of Motion
51 pages

RR-93-11
Bernhard Nebel, Hans-Jürgen Bürckert
Reasoning about Temporal Relations: A Maximal Tractable Subclass of Allen's Interval Algebra
28 pages

RR-93-10
Martin Buchheit, Francesco M. Donini, Andrea Schaerf
Decidable Reasoning in Terminological Knowledge Representation Systems
35 pages

RR-93-09
Philipp Hanschke, Jörg Würtz
Satisfiability of the Smallest Binary Program
8 pages

RR-93-08
Harold Boley, Philipp Hanschke, Knut Hinkelmann, Manfred Meyer
COLAB: A Hybrid Knowledge Representation and Compilation Laboratory
64 pages

RR-93-07
Hans-Jürgen Bürckert, Bernhard Hollunder, Armin Laux
Concept Logics with Function Symbols
36 pages

RR-93-06
Hans-Jürgen Bürckert, Bernhard Hollunder, Armin Laux
On Skolemization in Constrained Logics
40 pages

RR-93-05
Franz Baader, Klaus Schulz
Combination Techniques and Decision Problems for Disunification
29 pages

RR-93-04
Christoph Klauck, Johannes Schwagereit
GGD: Graph Grammar Developer for features in CAD/CAM
13 pages

RR-93-03
Franz Baader, Bernhard Hollunder, Bernhard Nebel, Hans-Jürgen Profitlich, Enrico Franconi
An Empirical Analysis of Optimization Techniques for Terminological Representation Systems
28 pages

RR-93-02
Wolfgang Wahlster, Elisabeth André, Wolfgang Finkler, Hans-Jürgen Profitlich, Thomas Rist
Plan-based Integration of Natural Language and Graphics Generation
50 pages

RR-93-01
Bernhard Hollunder
An Alternative Proof Method for Possibilistic Logic and its Application to Terminological Logics
25 pages

DFKI Technical Memos

1995

TM-95-04
Klaus Schmid
Creative Problem Solving
and
Automated Discovery
— An Analysis of Psychological and AI Research —
152 pages

TM-95-03
Andreas Abecker, Harold Boley, Knut Hinkelmann, Holger Wache, Franz Schmalhofer
An Environment for Exploring and Validating Declarative Knowledge
11 pages

TM-95-02
Michael Sintek
FLIP: Functional-plus-Logic Programming on an Integrated Platform
106 pages

TM-95-01
Martin Buchheit, Rüdiger Klein, Werner Nutt
Constructive Problem Solving: A Model Construction Approach towards Configuration
34 pages

1994

TM-94-04
Cornelia Fischer
PAntUDE - An Anti-Unification Algorithm for Expressing Refined Generalizations
22 pages

TM-94-03
Victoria Hall
Uncertainty-Valued Horn Clauses
31 pages

TM-94-02
Rainer Bleisinger, Berthold Kröll
Representation of Non-Convex Time Intervals and Propagation of Non-Convex Relations
11 pages

TM-94-01
Rainer Bleisinger, Klaus-Peter Gores
Text Skimming as a Part in Paper Document Understanding
14 pages

1993

TM-93-05
Michael Sintek
Indexing PROLOG Procedures into DAGs by Heuristic Classification
64 pages

TM-93-04
Hans-Günther Hein
Propagation Techniques in WAM-based Architectures — The FIDO-III Approach
105 pages

TM-93-03
Harold Boley, Ulrich Buhrmann, Christof Kremer
Konzeption einer deklarativen Wissensbasis über recyclingrelevante Materialien
11 pages

TM-93-02
Pierre Sablayrolles, Achim Schupeta
Conflict Resolving Negotiation for COoperative Schedule Management Agents (COSMA)
21 pages

TM-93-01
Otto Kühn, Andreas Birk
Reconstructive Integrated Explanation of Lathe Production Plans
20 pages

DFKI Documents

1995

D-95-12
F. Baader, M. Buchheit, M. A. Jeusfeld, W. Nutt (Eds.)
Working Notes of the KI'95 Workshop:
KRDB-95 - Reasoning about Structured Objects:
Knowledge Representation Meets Databases
61 pages

D-95-11
Stephan Busemann, Iris Merget
Eine Untersuchung kommerzieller Terminverwaltungssoftware im Hinblick auf die Kopplung mit natürlichsprachlichen Systemen
32 Seiten

D-95-10
Volker Ehresmann
Integration ressourcen-orientierter Techniken in das wissensbasierte Konfigurierungssystem TOOCON
108 Seiten

D-95-09
Antonio Krüger
PROXIMA: Ein System zur Generierung graphischer Abstraktionen
120 Seiten

D-95-08
Technical Staff
DFKI Jahresbericht 1994
63 Seiten

Note: This document is no longer available in printed form.

D-95-07
Ottmar Lutzy
Morphic - Plus
Ein morphologisches Analyseprogramm für die deutsche Flexionsmorphologie und Komposita-Analyse
74 pages

D-95-06
Markus Steffens, Ansgar Bernardi
Integriertes Produktmodell für Behälter aus Faserverbundwerkstoffen
48 Seiten

D-95-05
Georg Schneider
Eine Werkbank zur Erzeugung von 3D-Illustrationen
157 Seiten

D-95-04
Victoria Hall
Integration von Sorten als ausgezeichnete taxonomische Prädikate in eine relational-funktionale Sprache
56 Seiten

D-95-03
Christoph Endres, Lars Klein, Markus Meyer
Implementierung und Erweiterung der Sprache $\mathcal{ALCP}$
110 Seiten

D-95-02
Andreas Butz
BETTY
Ein System zur Planung und Generierung informativer Animationssequenzen
95 Seiten

D-95-01
Susanne Biundo, Wolfgang Tank (Hrsg.)
PuK-95, Beiträge zum 9. Workshop „Planen und Konfigurieren", Februar 1995
169 Seiten

Note: This document is available for a nominal charge of 25 DM (or 15 US-$).

1994

D-94-15
Stephan Oepen
German Nominal Syntax in HPSG
— On Syntactic Categories and Syntagmatic Relations —
80 pages

D-94-14
Hans-Ulrich Krieger, Ulrich Schäfer
TDL - A Type Description Language for HPSG, Part 2: User Guide.
72 pages

D-94-12
Arthur Sehn, Serge Autexier (Hrsg.)
Proceedings des Studentenprogramms der 18. Deutschen Jahrestagung für Künstliche Intelligenz KI-94
69 Seiten

D-94-11
F. Baader, M. Buchheit, M. A. Jeusfeld, W. Nutt (Eds.)
Working Notes of the KI'94 Workshop: KRDB'94 - Reasoning about Structured Objects: Knowledge Representation Meets Databases
65 pages

Note: This document is no longer available in printed form.

D-94-10
F. Baader, M. Lenzerini, W. Nutt, P. F. Patel-Schneider (Eds.)
Working Notes of the 1994 International Workshop on Description Logics
118 pages

Note: This document is available for a nominal charge of 25 DM (or 15 US-$).

D-94-09
Technical Staff
DFKI Wissenschaftlich-Technischer Jahresbericht 1993
145 Seiten

D-94-08
Harald Feibel
IGLOO 1.0 - Eine grafikunterstützte Beweisentwicklungsumgebung
58 Seiten

D-94-07
Claudia Wenzel, Rainer Hoch
Eine Übersicht über Information Retrieval (IR) und NLP-Verfahren zur Klassifikation von Texten
25 Seiten

D-94-06
Ulrich Buhrmann
Erstellung einer deklarativen Wissensbasis über recyclingrelevante Materialien
117 Seiten

D-94-04
Franz Schmalhofer, Ludger van Elst
Entwicklung von Expertensystemen: Prototypen, Tiefenmodellierung und kooperative Wissensevolution
22 Seiten

D-94-03
Franz Schmalhofer
Maschinelles Lernen: Eine kognitionswissenschaftliche Betrachtung
54 Seiten

Note: This document is no longer available in printed form.

D-94-02
Markus Steffens
Wissenserhebung und Analyse zum Entwicklungsprozeß eines Druckbehälters aus Faserverbundstoff
90 pages

D-94-01
Josua Boon (Ed.)
DFKI-Publications: The First Four Years
1990 - 1993
75 pages

1993

D-93-27
Rolf Backofen, Hans-Ulrich Krieger, Stephen P. Spackman,
Hans Uszkoreit (Eds.)
Report of the EAGLES Workshop on Implemented Formalisms at DFKI, Saarbrücken
110 pages

Note: This document is no longer available in printed form.

D-93-26
Frank Peters
Unterstützung des Experten bei der Formalisierung von Textwissen INFOCOM - Eine interaktive Formalisierungskomponente
58 Seiten

D-93-25
Hans-Jürgen Bürckert, Werner Nutt (Eds.)
Modeling Epistemic Propositions
118 pages

Note: This document is available for a nominal charge of 25 DM (or 15 US-$).

D-93-24
Brigitte Krenn, Martin Volk
DiTo-Datenbank: Datendokumentation zu Funktionsverbgefügen und Relativsätzen
66 Seiten

D-93-22
Andreas Abecker
Implementierung graphischer Benutzungsoberflächen mit Tcl/Tk und Common Lisp
44 Seiten

Note: This document is no longer available in printed form.

D-93-21
Dennis Drollinger
Intelligentes Backtracking in Inferenzsystemen am Beispiel Terminologischer Logiken
53 Seiten

D-93-20
Bernhard Herbig
Eine homogene Implementierungsebene für einen hybriden Wissensrepräsentationsformalismus
97 Seiten

D-93-16
Bernd Bachmann, Ansgar Bernardi, Christoph Klauck, Gabriele Schmidt
Design & KI
74 Seiten

D-93-15
Robert Laux
Untersuchung maschineller Lernverfahren und heuristischer Methoden im Hinblick auf deren Kombination zur Unterstützung eines Chart-Parsers
86 Seiten

D-93-14
Manfred Meyer (Ed.)
Constraint Processing – Proceedings of the International Workshop at CSAM'93, St.Petersburg, July 20-21, 1993
264 pages

Note: This document is available for a nominal charge of 25 DM (or 15 US-$).

D-93-12
Harold Boley, Klaus Elsbernd, Michael Herfert, Michael Sintek,
Werner Stein
RELFUN Guide: Programming with Relations and Functions Made Easy
86 pages

D-93-11
Knut Hinkelmann, Armin Laux (Eds.)
DFKI Workshop on Knowledge Representation Techniques — Proceedings
88 pages

Note: This document is no longer available in printed form.

D-93-10
Elizabeth Hinkelman, Markus Vonerden, Christoph Jung
Natural Language Software Registry (Second Edition)
174 pages

D-93-09
Hans-Ulrich Krieger, Ulrich Schäfer
TDLExtraLight User's Guide
35 pages

D-93-08
Thomas Kieninger, Rainer Hoch
Ein Generator mit Anfragesystem für strukturierte Wörterbücher zur Unterstützung von Texterkennung und Textanalyse
125 Seiten

D-93-07
Klaus-Peter Gores, Rainer Bleisinger
Ein erwartungsgesteuerter Koordinator zur partiellen Textanalyse
53 Seiten

D-93-06
Jürgen Müller (Hrsg.)
Beiträge zum Gründungsworkshop der Fachgruppe Verteilte Künstliche Intelligenz, Saarbrücken, 29. - 30. April 1993
235 Seiten

Note: This document is available for a nominal charge of 25 DM (or 15 US-$).

D-93-05
Elisabeth André, Winfried Graf, Jochen Heinsohn, Bernhard Nebel, Hans-Jürgen Profitlich, Thomas Rist, Wolfgang Wahlster
PPP: Personalized Plan-Based Presenter
70 pages

D-93-04
Technical Staff
DFKI Wissenschaftlich-Technischer Jahresbericht 1992
194 Seiten

D-93-03
Stephan Busemann, Karin Harbusch (Eds.)
DFKI Workshop on Natural Language Systems: Reusability and Modularity - Proceedings
74 pages

D-93-02
Gabriele Schmidt, Frank Peters, Gernod Laufkötter
User Manual of COKAM+
23 pages

D-93-01
Philipp Hanschke, Thom Frühwirth
Terminological Reasoning with Constraint Handling Rules
12 pages

D-93-05

Document

PPP – Personalized Plan-Based Presenter

Elisabeth André, Winfried Graf, Jochen Heinsohn
Bernhard Nebel, Hans-Jürgen Profitlich
Thomas Rist, Wolfgang Wahlster