Designing Illustrated Texts:
How Language Production Is Influenced by Graphics Generation

Wolfgang Wahlster, Elisabeth André, Winfried Graf, Thomas Rist

February 1991

Deutsches Forschungszentrum
für
Künstliche Intelligenz

The German Research Center for Artificial Intelligence (Deutsches Forschungszentrum für Künstliche Intelligenz, DFKI) with sites in Kaiserslautern und Saarbrücken is a non-profit organization which was founded in 1988 by the shareholder companies ADV/Orga, AEG, IBM, Insiders, Fraunhofer Gesellschaft, GMD, Krupp-Atlas, Mannesmann-Kienzle, Nixdorf, Philips and Siemens. Research projects conducted at the DFKI are funded by the German Ministry for 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
- ❏ Intelligent Communication Networks
- ❏ Intelligent Cooperative Systems.

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.

Prof. Dr. Gerhard Barth
Director

Designing Illustrated Texts: How Language Production is Influenced by Graphics Generation

Wolfgang Wahlster, Elisabeth André, Winfried Graf, Thomas Rist

This paper will be published in the Proceedings of the 5th Conference of the European Chapter of the Association for Computational Linguistics, 9.-11. April 1991, Berlin, Germany.

Designing Illustrated Texts:
How Language Production Is Influenced by Graphics Generation

Wolfgang Wahlster, Elisabeth André, Winfried Graf, Thomas Rist

Authors' Abstract

Multimodal interfaces combining, e.g., natural language and graphics take advantage of both the individual strength of each communication mode and the fact that several modes can be employed in parallel, e.g., in the text-picture combinations of illustrated documents. It is an important goal of this research not simply to merge the verbalization results of a natural language generator and the visualization results of a knowledge-based graphics generator, but to carefully coordinate graphics and text in such a way that they complement each other. We describe the architecture of the knowledge-based presentation system WIP which guarantees a design process with a large degree of freedom that can be used to tailor the presentation to suit the specific context. In WIP, decisions of the language generator may influence graphics generation and graphical constraints may sometimes force decisions in the language production process. In this paper, we focus on the influence of graphical constraints on text generation. In particular, we describe the generation of cross-modal references, the revision of text due to graphical constraints and the clarification of graphics through text.

Table of Contents

1 INTRODUCTION

With increases in the amount and sophistication of information that must be communicated to the users of complex technical systems comes a corresponding need to find new ways to present that information flexibly and efficiently. Intelligent presentation systems are important building blocks of the next generation of user interfaces, as they translate from the narrow output channels provided by most of the current application systems into high-bandwidth communications tailored to the individual user. Since in many situations information is only presented efficiently through a particular combination of communication modes, the automatic generation of multimodal presentations is one of the tasks of such presentation systems. The task of the knowledge-based presentation system WIP is the generation of a variety of multimodal documents from an input consisting of a formal description of the communicative intent of a planned presentation. The generation process is controlled by a set of generation parameters such as target audience, presentation objective, resource limitations, and target language.

One of the basic principles underlying the WIP project is that the various constituents of a multimodal presentation should be generated from a common representation. This raises the question of how to divide a given communicative goal into subgoals to be realized by the various mode-specific generators, so that they complement each other. To address this problem, we have to explore computational models of the cognitive decision processes coping with questions such as what should go into text, what should go into graphics, and which kinds of links between the verbal and non-verbal fragments are necessary.

In the project WIP, we try to generate on the fly illustrated texts that are customized for the intended target audience and situation, flexibly presenting information whose content, in contrast to hypermedia systems, cannot be fully anticipated. The current testbed for WIP is the generation of instructions for the use of an espresso-machine. It is a rare instruction manual that does not contain illustrations. WIP's 2D display of 3D graphics of machine parts help the addressee of the synthesized multimodal presentation to develop a 3D mental model of the object that he can constantly match with his visual perceptions of the real machine in front of him. Fig. 1 shows a typical text-picture sequence which may be used to instruct a user in filling the watercontainer of an espresso-machine.

Fig. 1: Example Instruction

Currently, the technical knowledge to be presented by WIP is encoded in a hybrid knowledge representation language of the KL-ONE family including a terminological and assertional component (see Nebel 90). In addition to this propositional representation, which includes the relevant information about the structure, function, behavior, and use of the espresso-machine, WIP has access to an analogical representation of the geometry of the machine in the form of a wireframe model.

The automatic design of multimodal presentations has only recently received significant attention in artificial intelligence research (cf. the projects SAGE (Roth et al. 89), COMET (Feiner & McKeown 89), FN/ANDD (Marks & Reiter 90) and WIP (Wahlster et al. 89)). The WIP and COMET projects share a strong research interest in the coordination of text and graphics. They differ from systems such as SAGE and FN/ANDD in that they deal with physical objects (espresso-machine, radio vs. charts, diagrams) 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 to understand the operation of the machine. In spite of many similarities, there are major differences between COMET and WIP, e.g., in the systems' architecture. While during one of the final processing steps of COMET the layout component combines text and graphics fragments produced by mode-specific generators, in WIP a layout manager can interact with a presentation planner before

text and graphics are generated, so that layout considerations may influence the early stages of the planning process and constrain the mode-specific generators.

2 THE ARCHITECTURE OF WIP

The architecture of the WIP system guarantees a design process with a large degree of freedom that can be used to tailor the presentation to suit the specific context. During the design process a presentation planner and a layout manager orchestrate the mode-specific generators and the document history handler (see Fig. 2) provides information about intermediate results of the presentation design that is exploited in order to prevent disconcerting or incoherent output. This means that decisions of the language generator may influence graphics generation and that graphical constraints may sometimes force decisions in the language production process. In this paper, we focus on the influence of graphical constraints on text generation (see Wahlster et al. 91 for a discussion of the inverse influence).

Fig. 2 shows a sketch of WIP's current architecture used for the generation of illustrated documents. Note that WIP includes two parallel processing cascades for the incremental generation of text and graphics. In WIP, the design of a multimodal document is viewed as a non-monotonic process that includes various revisions of preliminary results, massive replanning or plan repairs, and many negotiations between the corresponding design and realization components in order to achieve a fine-grained and optimal division of work between the selected presentation modes.

2.1 THE PRESENTATION PLANNER

The presentation planner is responsible for contents and mode selection. A basic assumption behind the presentation planner is that not only the generation of text, but also the generation of multimodal documents can be considered as a sequence of communicative acts which aim to achieve certain goals (cf. André & Rist 90a). For the synthesis of illustrated texts, we have designed presentation strategies that refer to both text and picture production. To represent the strategies, we follow the approach proposed by Moore and colleagues (cf. Moore & Paris 89) to operationalize RST-theory (cf. Mann & Thompson 88) for text planning.

The strategies are represented by a name, a header, an effect, a set of applicability conditions and a specification of main and subsidiary acts. Whereas the header of a strategy indicates which communicative function the corresponding document part is to fill, its effect refers to an intentional goal. The applicability conditions specify when a strategy may be used and put restrictions on the variables to be instantiated. The main and subsidiary acts form the kernel of the strategies. E.g., the strategy below can be used to enable the identification of an object shown in a picture (for further details see André & Rist 90b). Whereas graphics is to be used to carry out the main act, the mode for the subsidiary acts is open.

Name:
Enable-Identification-by-Background
Header:
(Provide-Background P A ?x ?px ?pic GRAPHICS)
Effect:
(BMB P A (Identifiable A ?x ?px ?pic))
Applicability Conditions:
(AND   (Bel P (Perceptually-Accessible A ?x))
          (Bel P (Part-of ?x ?z)))
Main Acts:
(Depict P A (Background ?z) ?pz ?pic)
Subsidiary Acts:
(Achieve P (BMB P A (Identifiable A ?z ?pz ?pic)) ?mode)

For the automatic generation of illustrated documents, the presentation strategies are treated as operators of a planning system. During the planning process, presentation strategies are selected and instantiated according to the presentation task. After the selection of a strategy, the main and subsidiary acts are carried out unless the corresponding presentation goals are already satisfied. Elementary acts, such as **Depict** or **Assert**, are performed by the text and graphics generators.

2.2 THE LAYOUT MANAGER

The main task of the layout manager is to convey certain semantic and pragmatic relations specified by the planner by the arrangement of graphic and text fragments received from the mode-specific generators, i.e., to determine the size of the boxes and the exact coordinates for positioning them on the document page. We use a grid-based approach as an ordering system for efficiently designing functional (i.e., uniform, coherent and consistent) layouts (cf. Müller-Brockmann 81).

A central problem for automatic layout is the representation of design-relevant knowledge. Constraint networks seem to be a natural formalism to declaratively incorporate aesthetic knowledge into the layout process, e.g., perceptual criteria concerning the organization of boxes as sequential ordering, alignment, grouping, symmetry or similarity. Layout constraints can be classified as semantic, geometric, topological, and temporal. Semantic constraints essentially correspond to coherence relations, such as sequence and contrast, and can be easily reflected through specific design constraints. A powerful way of expressing such knowledge is to organize the constraints hierarchically by assigning a preference scale to the constraint network (cf. Borning et al. 89). We distinguish obligatory, optional and default constraints. The latter state default values, that remain fixed unless the corresponding constraint is removed by a stronger one. Since there are constraints that have only local effects, the incremental constraint solver must be able to change the constraint hierarchy dynamically (for further details see Graf 90).

2.3 THE TEXT GENERATOR

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 90 and Schauder 91). The grammar is divided into an LD (linear 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 (Finkler & Neumann 89).

The text generator uses a TAG parser in a local anticipation feedback loop (see Jameson & Wahlster 82). 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.

Since 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. Currently, we are exploring 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 (see McKeown & Feiner 90).

2.4 THE GRAPHICS GENERATOR

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 achieve a given presentation goal efficiently. E.g., when a picture of an object is requested, we have to determine an appropriate perspective in a context-sensitive way (cf. Rist&André 90). 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. E.g., 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 which map the 3D subject onto its depiction. E.g., 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. E.g., 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. 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.

3 THE GENERATION OF CROSS-MODAL REFERENCES

In a multimodal presentation, cross-modal expressions establish referential relationships of representations in one modality to representations in another modality.

The use of cross-modal deictic expressions such as (a) - (b) is essential for the efficient coordination of text and graphics in illustrated documents:

(a) *The left knob in the figure on the right is the on/off switch.*
(b) *The black square in Fig. 14 shows the watercontainer.*

In sentence (a) a spatial description is used to refer to a knob shown in a synthetic picture of the espresso-machine. Note that the multimodal referential act is only successful if the addressee is able to identify the intended knob of the real espresso-machine. It is clear that the visualization of the knob in the illustration cannot be used as an on/off switch, but only the physical object identified as the result of a two-level reference process, i.e., the cross-modal expression in the text refers to a specific part of the illustration which in turn refers to a real-world object[1].

---

[1] In the WIP system there exists yet another coreferentiality relation, namely between an individual constant, say `knob-1`, representing the particular knob in the knowledge representation language and an object in the wireframe model of the machine containing a description of the geometry of that knob.

Another subtlety illustrated by example (a) is the use of different frames of reference for the two spatial relations used in the cross-modal expression. The definite description *figure on the right* is based on a component generating absolute spatial descriptions for geometric objects displayed inside rectangular frames. In our example, the whole page designed by WIP's layout manager constitutes the frame of reference. One of the basic ideas behind this component is that such 'absolute' descriptions can be mapped on relative spatial predicates developed for the VITRA system (see Herzog et al. 90) through the use of a virtual reference object in the center of the frame (for more details see Wazinski 91). This means that the description of the location of the figure showing the on/off switch mentioned in sentence (a) is based on the literal **right-of(figure-A,center(page-1))** produced by WIP's localization component.

The definite description *the left knob* is based on the use of the region denoted by *figure on the right* as a frame of reference for another call of the localization component producing the literal **left-of(knob1,knob2))** as an appropriate spatial description. Note that all these descriptions are highly dependent on the viewing specification chosen by the graphics design component. That means that changes in the illustrations during a revision process must automatically be made available to the text design component.

Fig. 3: *The middle knob* in A is *the left knob* in the close-up projection B

Let's assume that the presentation planner has selected the relevant information for a particular presentation goal. This may cause the graphics designer to choose a close-up projection of the top part of the espresso-machine with a narrow field of view focusing on specific objects and eliminating unnecessary details from the graphics as shown in Fig. B (see Fig. 3). If the graphics designer chooses a wide field of view (see Fig. A in Fig. 3) for another presentation goal, **knob1** can no longer be described as *the left knob* since the `real-world' spatial location of another knob (e.g., **knob0**), which was not shown in the close-up projection, is now used to produce the adequate spatial description *the left knob* for **knob0**. Considering the row of three knobs in Fig. A, **knob1** is now described as *the middle knob*.

Note that the layout manager also needs to backtrack from time to time. This may result in different placement of the figure A, e.g., at the bottom of the page. This means that in the extreme, the cross-modal expression *the left knob in the figure on the right* will be changed into *the middle knob in the figure at the bottom*.

Due to various presentational constraints, the graphics design component cannot always show the wireframe object in a general position providing as much geometric information about the object as possible. For example, when a cube is viewed along the normal to a face it projects to a square, so that a loss of generality results (see Karp & Feiner 90). In example (b) the definite description *the black square* uses shape information extracted from the projection chosen by the graphics designer that is stored in the document history handler. It is obvious that even a slight change in the viewpoint for the graphics can result in a presentation situation where *the black cube* has to be used as a referential expression instead of *black square*. Note that the colour attribute *black* used in these descriptions may conflict with the addressee's visual perception of the real espresso-machine.

The difference between referring to attributes in the model and perceptual properties of the real-world object becomes more obvious in cases where the specific features of the display medium are used to highlight intended objects (e.g., blinking or inverse video) or when metagraphical objects are chosen as reference points (e.g., an arrow pointing to the intended object in the illustration). It is clear that a definite description like the *blinking square* or *the square that is highlighted by the bold arrow* cannot be generated before the corresponding decisions about illustration techniques are finalized by the graphics designer.

The text planning component of a multimodal presentation system such as WIP must be able to generate such cross-modal expressions not only for figure captions, but also for coherent text-picture combinations.

4 THE REVISION OF TEXT DUE TO GRAPHICAL CONSTRAINTS

Frequently, the author of a document faces formal restrictions, e.g., when document parts must not exceed a specific page size or column width. Such formatting constraints may influence the structure and contents of the document. A decisive question is, at which stage of the generation process such constraints should be evaluated. Some restrictions, such as page size, are known a priori, while others (e.g., that an illustration should be placed on the page where it is first discussed) arise during the generation process. In the WIP system, the problem is aggravated since restrictions can result from the processing of at least two generators (for text and graphics) working in parallel. A mode-specific generator is not able to anticipate all situations in which formatting problems might occur. Thus in WIP, the generators are launched to produce a first version of their planned output which may be revised if necessary. We illustrate this revision process by showing the coordination of WIP's components when object depictions are annotated with text strings.

Suppose the planner has decided to introduce the essential parts of the espresso-machine by classifying them. E.g., it wants the addressee to identify a switch which allows one to choose between two operating modes: producing espresso or producing steam. In the knowledge base, such a switch may be represented as shown in Fig. 4.

Since it is assumed that the discourse objects are visually accessible to the addressee, it is reasonable to refer to them by means of graphics, to describe them verbally and to show the connection between the depictions and the verbal descriptions. In instruction manuals this is usually accomplished by various annotation techniques. In the current WIP system, we have implemented three annotation techniques: annotating by placing the text string inside an object projection, close to it, or by using arrows starting at the text string and pointing to the intended object. Which annotation technique applies depends on syntactic criteria, (e.g., formatting restrictions) as well as semantic criteria to avoid confusion. E.g., the same annotation technique is to be used for all instances of the same basic concept (cf. Butz et al. 91).
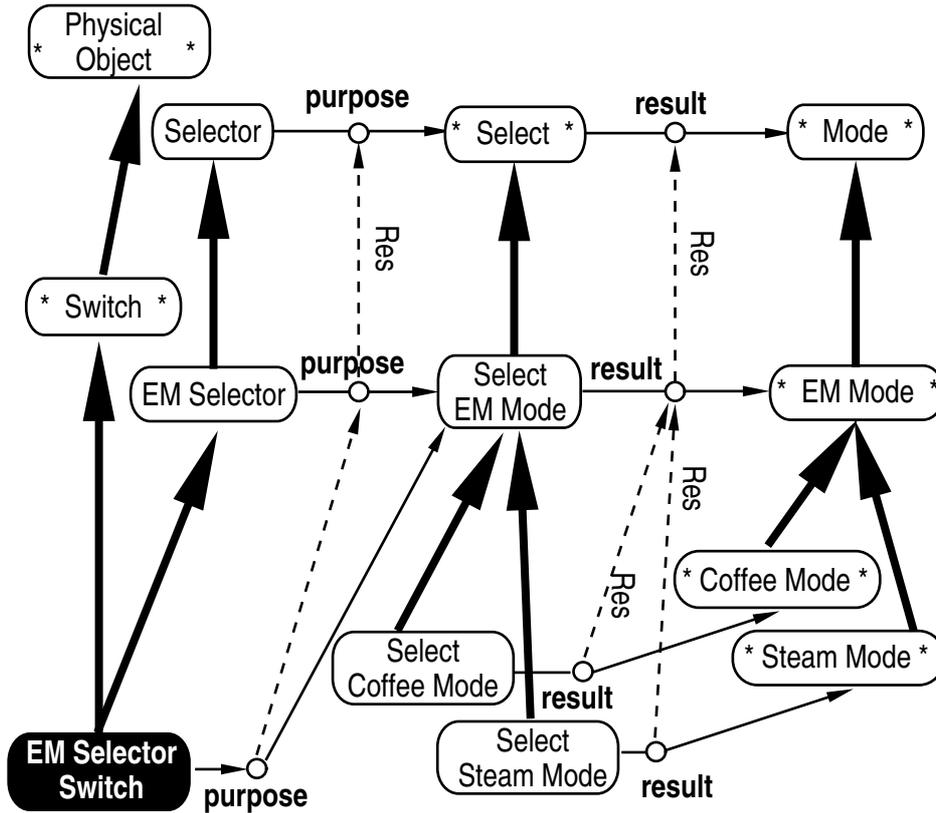
Fig. 4: Part of the Terminological Knowledge Base

Suppose that in our example, the text generator is asked to find a lexical realization for the concept **EM selector switch** and comes up with the description *selector switch for coffee and steam*. When trying to annotate the switch with this text string, the graphics generator finds out that none of the available annotation techniques apply. Placing the string close to the corresponding depiction causes ambiguities. The string also cannot be placed inside the projection of the object without occluding other parts of the picture. For the same reason, annotations with arrows fail. Therefore, the text generator is asked to produce a shorter formulation. Unfortunately, it is not able to do so without reducing the contents. Thus, the presentation planner is informed that the required task cannot be accomplished. The presentation planner then tries to reduce the contents by omitting attributes or by selecting more general concepts from the subsumption hierarchy encoded in terms of the terminological logic. Since **EM selector switch** is a compound description which inherits information from the concepts **switch** and **EM selector** (see Fig. 4), the planner has to decide which component of the contents specification should be reduced. Because the concept **switch** contains less discriminating information than the concept **EM selector** and the concept **switch** is at least partially inferrable from the picture, the planner first tries to reduce the component **switch** by replacing it by **physical object**. Thus, the text generator has to find a

sufficiently short definite description containing the components **physical object** and **EM selector**. Since this fails, the planner has to propose another reduction. It now tries to reduce the component **EM selector** by omitting the coffee/steam mode. The text generator then tries to construct a NP combining the concepts **switch** and **selector**. This time it succeeds and the annotation string can be placed. Fig. 5 is a hardcopy produced by WIP showing the rendered espresso-machine after the required annotations have been carried out.
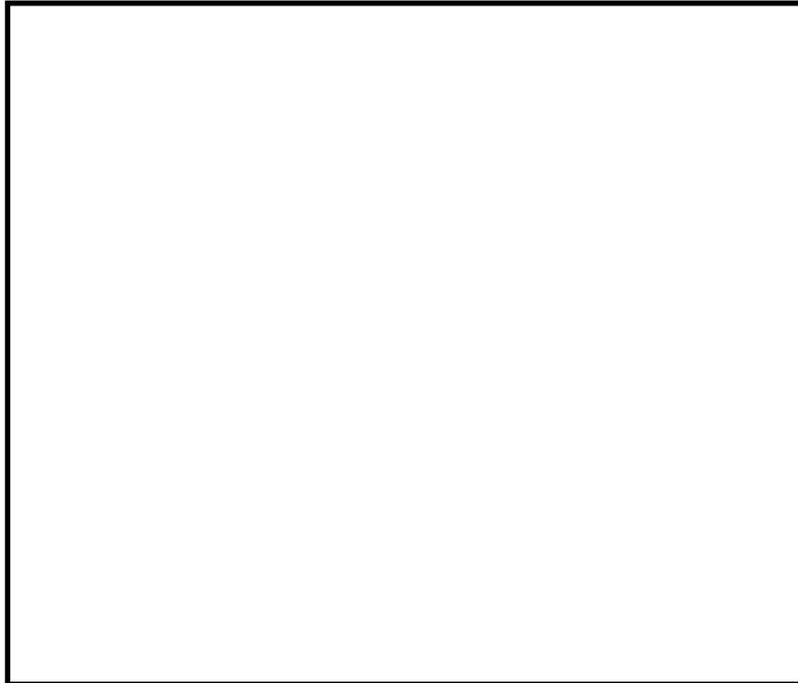
Fig. 5: Annotations after Text Revisions

5 THE CLARIFICATION OF GRAPHICS THROUGH TEXT

In the example above, the first version of a definite description produced by the text generator had to be shortened due to constraints resulting from picture design. However, there are also situations in which clarification information has to be added through text because the graphics generator on its own is not able to convey the information to be communicated.

Let's suppose the graphics designer is requested to show the location of **fitting-1** with respect to the **espresso-machine-1**. The graphics designer tries to design a picture that includes objects that can be identified as **fitting-1** and **espresso-machine-1**. To convey the

13

location of **fitting-1** the picture must provide essential information which enables the addressee to reconstruct the initial 3D object configuration (i.e., information concerning the topology, metric and orientation). To ensure that the addressee is able to identify the intended object, the graphics designer tries to present the object from a standard perspective, i.e., an object dependent perspective that satisfies standard presentation goals, such as showing the object's functionality, top-bottom orientation, or accessibility (see also Rist & André 90). In the case of a part-whole relationship, we assume that the location of the part with respect to the whole can be inferred from a picture if the whole is shown under a perspective such that both the part and further constituents of the whole are visible. In our example, **fitting-1** only becomes visible and identifiable as a part of the espresso-machine when showing the machine from the back. But this means that the espresso-machine must be presented from a non-standard perspective and thus we cannot assume that its depiction can be identified without further clarification.

Whenever the graphics designer discovers conflicting presentation goals that cannot be solved by using an alternative technique, the presentation planner must be informed about currently solved and unsolvable goals. In the example, the presentation planner has to ensure that the espresso-machine is identifiable. Since we assume that an addressee is able to identify an object's depiction if he knows from which perspective the object is shown, the conflict can be resolved by informing the addressee that the espresso-machine is depicted from the back. This means that the text generator has to produce a comment such as *This figure shows the fitting on the back of the machine*, which clarifies the graphics.

6 CONCLUSION

In this paper, we introduced the architecure of the knowledge-based presentation system WIP, which includes two parallel processing cascades for the incremental generation of text and graphics. We showed that in WIP the design of a multimodal document is viewed as a non-monotonic process that includes various revisions of preliminary results, massive replanning or plan repairs, and many negotiations between the corresponding design and realization components in order to achieve a fine-grained and optimal devision of work between the selected presentation modes. We described how the plan-based approach to presentation design can be exploited so that graphics generation influences the production of text. In particular, we showed how WIP can generate cross-modal references, revise text due to graphical constraints and clarify graphics through text.

REFERENCES

[André & Rist 90a] Elisabeth **André** and Thomas **Rist**. Towards a Plan-Based Synthesis of Illustrated Documents. In: 9th ECAI, 25-30, 1990.

[André & Rist 90b] Elisabeth **André** and Thomas **Rist**. Generating Illustrated Documents: A Plan-Based Approach. In: InfoJapan 90, Vol. 2, 163-170, 1990.

[Borning et al. 89] Alan **Borning**, Bjorn **Freeman-Benson**, and Molly **Wilson**. Constraint Hierarchies. Technical Report, Department of Computer Science and Engineering, University of Washington, 1989.

[Butz et al. 91] Andreas **Butz**, Bernd **Hermann**, Daniel **Kudenko**, and Detlev **Zimmermann**. ANNA: Ein System zur Annotation und Analyse automatisch erzeugter Bilder. Memo, DFKI, Saarbrücken, 1991.

[Feiner & McKeown 89] Steven **Feiner** and Kathleen **McKeown**. Coordinating Text and Graphics in Explanation Generation. In: DARPA Speech and Natural Language Workshop, 1989.

[Finkler & Neumann 89] Wolfgang **Finkler** and Günter **Neumann**. POPEL-HOW: A Distributed Parallel Model for Incremental Natural Language Production with Feedback. In: 11th IJCAI, 1518-1523, 1989.

[Graf 90] Winfried **Graf**. Spezielle Aspekte des automatischen Layout-Designs bei der koordinierten Generierung von multimodalen Dokumenten. GI-Workshop "Multimediale elektronische Dokumente", 1990.

[Harbusch 90] Karin **Harbusch**. Constraining Tree Adjoining Grammars by Unification. 13th COLING, 167-172, 1990.

[Herzog et al. 90] Gerd **Herzog**, Elisabeth **André**, and Thomas **Rist.** Sprache und Raum: Natürlichsprachlicher Zugang zu visuellen Daten. In: Christian Freksa and Christopher Habel (eds.). Repräsentation und Verarbeitung räumlichen Wissens. IFB 245, 207-220, Berlin: Springer-Verlag, 1990.

[Jameson & Wahlster 82] Anthony **Jameson** and Wolfgang **Wahlster**. User Modelling in Anaphora Generation: Ellipsis and Definite Description. In: 5th ECAI, 222-227, 1982

[Karp & Feiner 90] Peter **Karp** and Steven **Feiner**. Issues in the Automated Generation of Animated Presentations. In: Graphics Interface '90, 39-48, 1990.

[Mann & Thompson 88] William **Mann** and Sandra **Thompson**. Rhetorical Structure Theory: Towards a Functional Theory of Text Organization. In: TEXT, 8 (3), 1988.

[Marks & Reiter 90] Joseph **Marks** and Ehud **Reiter**. Avoiding Unwanted Conversational Implicatures in Text and Graphics. In: 8th AAAI, 450-455, 1990.

[McKeown & Feiner 90] Kathleen **McKeown** and Steven **Feiner**. Interactive Multimedia Explanation for Equipment Maintenance and Repair. In: DARPA Speech and Natural Language Workshop, 42-47, 1990.

[Moore & Paris 89] Johanna **Moore** and Cécile **Paris**. Planning Text for Advisory Dialogues. In: 27th ACL, 1989.

[Müller-Brockmann 81] Josef **Müller-Brockmann.** Grid Systems in Graphic Design. Stuttgart: Hatje, 1981.

[Nebel 90] Bernhard **Nebel**. Reasoning and Revision in Hybrid Representation Systems. Lecture Notes in AI, Vol. 422, Berlin: Springer-Verlag, 1990.

[Rist & André 90] Thomas **Rist** and Elisabeth **André**. Wissensbasierte Perspektivenwahl für die automatische Erzeugung von 3D-Objektdarstellungen. In: Klaus Kansy and Peter Wißkirchen (eds.). Graphik und KI. IFB 239, Berlin: Springer-Verlag, 48-57, 1990.

[Roth et al. 89] Steven **Roth**, Joe **Mattis**, and Xavier **Mesnard**. Graphics and Natural Language as Components of Automatic Explanation. In: Joseph Sullivan and Sherman Tyler (eds.). Architectures for Intelligent Interfaces: Elements and Prototypes. Reading, MA: Addison-Wesley, 1989.

[Schauder 90] Anne **Schauder**. Inkrementelle syntaktische Generierung natürlicher Sprache mit Tree Adjoining Grammars. MS thesis, Computer Science, University of Saarbrücken, 1990.

[Wahlster et al. 89] Wolfgang **Wahlster**, Elisabeth **André**, Matthias **Hecking**, and Thomas **Rist**. WIP: Knowledge-based Presentation of Information. Report WIP-1, DFKI, Saarbrücken, 1989.

[Wahlster et al. 91] Wolfgang **Wahlster**, Elisabeth **André**, Som **Bandyopadhyay**, Winfried **Graf**, and Thomas **Rist**. WIP: The Coordinated Generation of Multimodal Presentations from a Common Representation. In: Oliviero Stock, John Slack, Andrew Ortony (eds.). Computational Theories of Communication and their Applications. Berlin: Springer-Verlag, 1991.

[Wazinski 91] Peter **Wazinski**. Objektlokalisation in graphischen Darstellungen. MS thesis, Computer Science, University of Saarbrücken, forthcoming.
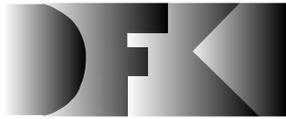
DFKI Publikationen

DFKI Research Reports

RR-92-43
Christoph Klauck, Jakob Mauss: A Heuristic driven Parser for Attributed Node Labeled Graph Grammars and its Application to Feature Recognition in CIM
17 pages

RR-92-44
Thomas Rist, Elisabeth André: Incorporating Graphics Design and Realization into the Multimodal Presentation System WIP
15 pages

RR-92-45
Elisabeth André, Thomas Rist: The Design of Illustrated Documents as a Planning Task
21 pages

RR-92-46
Elisabeth André, Wolfgang Finkler, Winfried Graf, Thomas Rist, Anne Schauder, Wolfgang Wahlster: WIP: The Automatic Synthesis of Multimodal Presentations
19 pages

RR-92-47
Frank Bomarius: A Multi-Agent Approach towards Modeling Urban Traffic Scenarios
24 pages

RR-92-48
Bernhard Nebel, Jana Koehler:
Plan Modifications versus Plan Generation:
A Complexity-Theoretic Perspective
15 pages

RR-92-49
Christoph Klauck, Ralf Legleitner, Ansgar Bernardi: Heuristic Classification for Automated CAPP
15 pages

DFKI Publications

RR-92-50
Stephan Busemann:
Generierung natürlicher Sprache
61 Seiten

RR-92-51
Hans-Jürgen Bürckert, Werner Nutt:
On Abduction and Answer Generation through Constrained Resolution
20 pages

RR-92-52
Mathias Bauer, Susanne Biundo, Dietmar Dengler, Jana Koehler, Gabriele Paul: PHI - A Logic-Based Tool for Intelligent Help Systems
14 pages

RR-92-53
Werner Stephan, Susanne Biundo:
A New Logical Framework for Deductive Planning
15 pages

RR-92-54
Harold Boley: A Direkt Semantic Characterization of RELFUN
30 pages

RR-92-55
John Nerbonne, Joachim Laubsch, Abdel Kader Diagne, Stephan Oepen: Natural Language Semantics and Compiler Technology
17 pages

RR-92-56
Armin Laux: Integrating a Modal Logic of Knowledge into Terminological Logics
34 pages

RR-92-58
Franz Baader, Bernhard Hollunder:
How to Prefer More Specific Defaults in Terminological Default Logic
31 pages

DFKI Documents

D-92-19
Stefan Dittrich, Rainer Hoch: Automatische, Deskriptor-basierte Unterstützung der Dokument-analyse zur Fokussierung und Klassifizierung von Geschäftsbriefen
107 Seiten

D-92-21
Anne Schauder: Incremental Syntactic Generation of Natural Language with Tree Adjoining Grammars
57 pages

D-92-22
Werner Stein: Indexing Principles for Relational Languages Applied to PROLOG Code Generation
80 pages

D-92-23
Michael Herfert: Parsen und Generieren der Prolog-artigen Syntax von RELFUN
51 Seiten

D-92-24
Jürgen Müller, Donald Steiner (Hrsg.):
Kooperierende Agenten
78 Seiten

D-92-25
Martin Buchheit: Klassische Kommunikations- und Koordinationsmodelle
31 Seiten

D-92-26
Enno Tolzmann:
Realisierung eines Werkzeugauswahlmoduls mit Hilfe des Constraint-Systems CONTAX
28 Seiten

D-92-27
Martin Harm, Knut Hinkelmann, Thomas Labisch: Integrating Top-down and Bottom-up Reasoning in COLAB
40 pages

D-92-28
Klaus-Peter Gores, Rainer Bleisinger: Ein Modell zur Repräsentation von Nachrichtentypen
56 Seiten

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

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

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

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

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-93-06
Jürgen Müller (Hrsg.):
Beiträge zum Gründungsworkshop der Fachgruppe Verteilte Künstliche Intelligenz Saarbrücken 29.-30. April 1993
235 Seiten
Note: This document is available only for a nominal charge of 25 DM (or 15 US-$).

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

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

D-93-09
Hans-Ulrich Krieger, Ulrich Schäfer:
TDL ExtraLight User's Guide
35 pages

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

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

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-93-14
Manfred Meyer (Ed.): Constraint Processing – Proceedings of the International Workshop at CSAM'93, July 20-21, 1993
264 pages
Note: This document is available only for a nominal charge of 25 DM (or 15 US-$).