

Structure and Rules in Automated Multimedia Presentation Planning*

Yigal Arens, Eduard Hovy, and Susanne van Mulken

USC/Information Sciences Institute

4676 Admiralty Way

Marina del Rey, CA 90292-6695

U.S.A.

email: {ARENS, HOVY}@ISI.EDU, MULKEN@PSYCH.KUN.NL

Abstract

During the planning of multimedia presentations, at least two distinct processes are required: planning the underlying discourse structure (that is, ordering and interrelating the information to be presented) and allocating the media (that is, delimiting the portions to be displayed by each individual medium). The former process has been the topic of several studies in the area of text planning, but numerous questions remain for the latter, including: What information does the allocation process depend on and what does it produce? How should the two processes be performed — sequentially, interleaved, or simultaneously? In this paper, we define Discourse Structure and Presentation Structure and outline the kinds of information that media allocation rules must depend on, including, centrally, information about the discourse structure. We describe a prototype information-to-media planning system, arguing that since media allocation rules depend on the characteristics of the information to be presented, they can only be applied once the overall discourse structure has been essentially planned out and the individual portions of information have become apparent.

1 Central Problems of Multimedia Presentation Planning

The automated creation of multimedia presentations contains a number of challenging and interesting prob-

lems. Given a collection of information to display, what is the nature of the process of designing a multimedia presentation? What information does the process start with and what does it produce? On what information are its decision criteria based?

Borrowing from the conceptually related field of automated text generation, one might consider starting with a text planner that performs the tasks of information selection and organization, and then add to that a process that allocates the information to specific media for presentation. One may even hope that a single planning mechanism rather like a text planner will perform all three these tasks simultaneously; after all, it is dealing with the same information in all cases. Unfortunately, this appealing idea is not feasible. We argue in this paper that such integration cannot work because since the information-to-medium allocation process depends upon the overall structure of the discourse, this structure has to be determined before allocation can take place.

This fact may seem obvious when stated, but it is a problem that bedevils many current presentation planning systems, and has caused them to make significant simplifying assumptions in some cases or to finesse the issue in others. Although the timing of application of multimedia rules has serious consequences for the construction of multimedia presentation planners, we do not know of any work that solves the problem adequately. The process and timing of presentation rule application in WIP has only been described verbally (André, personal communication, April 1992), while the COMET reports we have seen are not specific on this point.

In this paper we argue the case as follows: first, we describe briefly the kind of background information required to perform flexible media allocation. Then we outline the need for the separation of planning and allocation processes. Next, given this separation, we define *Discourse Structure* and *Presentation Structure* and outline the nature of the transformations from the former (which is pre-media) to the latter (which is media-dependent), which can be quite significant. Finally, we provide two examples produced by a prototype media allocation planning system to illustrate the points made.

*The first author was supported in part by Rome Laboratory of the Air Force Systems Command and the Defense Advanced Research Projects Agency under contract no. F30602-91-C-0081. The second author was supported in part by Rome Laboratory of the Air Force Systems Command under RL contract no. FQ7619-89-03326-0001. The third author, a graduate student at the University of Nijmegen, Nijmegen, The Netherlands, spent a research visit of six months at USC/ISI working on this project for her Master's degree. Views and conclusions contained in this report are the authors' and should not be interpreted as representing the official opinion or policy of DARPA, RL, the U.S. Government, or any person or agency connected with them.

2 Background

2.1 Presentation Planning: Art or Science?

If humans and computers are to collaborate easily and well, computers will have to master the art of felicitous multimedia communication. Just like people, they will have to know when and how most appropriately to use words, pictures, and diagrams, as well as how to integrate different media into a coherent communication, since it is well established that different types of information are most clearly presented in different media [Dwyer 78, Tufte 83, Twyman 85].

Several research projects suggest that multimedia communication is a science, not an art: it can be captured in rules. The nature of such rules is an important question. If a system's rules are simply collections of states and actions that are not theoretically well motivated, the system may well work for some particular task, but its operation will be uninformative and, eventually, probably collapse under its own weight. On the other hand, if the rules are based on the principles of communication using accurate representations of pertinent aspects of media, data, and interlocutors, then the rules will probably continue to be useful, in some form or another, in future multimedia communication systems.

In this paper we argue for specific levels and types of representation as the most promising basis for rules for multimedia presentation planning. In particular, we argue that the planning process should convert a *discourse structure* — represented using intentional and rhetorical terms — into a *presentation structure*. We claim that the rules, to be appropriately general and powerful, should be based on very *fine-grained features* of *four kinds of information* that characterize the communication, the interlocutors, and the situation. Although we make no arguments here about the nature of the processes involved, we describe a prototype presentation planner that varies its output depending on delicate variations in features of the data to be communicated. This planner uses the kinds of representation features and other structures we believe are required for flexible, powerful, and theoretically interesting multimedia communication.

2.2 The Knowledge upon Which Rules Depend

Before describing the data structures and transformations, we outline the individual features of data upon which presentation building rules should be based.

In the simplest, most direct approach to multimedia presentation planning, the question of clustering and allocating information to appropriate media is handled simply by fixed rules that specify exactly what medium is to be used for each particular data type, as in the following rule from the Integrated Interfaces system [Arens et al. 88]:

1. Ships' locations are presented on maps.

For generality, portability, and flexibility, this is clearly not a satisfactory solution. For this reason, most approaches generalize this type of rule by using features, instead of data types and media, in both sides of the rule. One possible generalization is:

1'. Data duples (of which ships' locations are an example) are presented on two-dimensional media (such as maps, graphs, and tables).

Another possible generalization is:

1". Data with specific denotations (such as spatial location) are presented on media which can convey the same denotations.

Determining appropriate features upon which to base the rules is unfortunately not a simple matter. In an attempt to systematize the bewildering amount of relevant information, we surveyed several relevant fields, including some of the more theoretical literature on good presentation design [Tufte 90, Bertin 83, Tufte 83], recent studies in Psychology [Hartley 85, Fleming & Levie 78], Cognitive Science [Roth & Mattis 90, Mayer 89, Larkin & Simon 86], and Computer Science [Mackinlay 86, Wahlster et al. 92, Feiner & McKeown 90, Neal 90]. It was clear that the large number of parameters playing a role and their complex interrelationships, makes it imperative that features be defined with care at the proper level of generality to ensure their successful interaction.

As a result of our reading, and after analysis of multimedia aspects in business reports, instruction manuals, cookbooks, etc., we have identified and taxonomized four distinct collections of information features that play a role in the design of multimedia presentations [Vossers 91, Hovy & Arens 91, Hovy & Arens 90, Arens & Hovy 90]. These four classes of information are: characteristics of the information to be conveyed, characteristics of the media at hand, the presenter's goals, and the perceiver's interests and abilities. In [Arens et al. 92] we describe a feature-based formalism for this information, organized in an And-Or dependency network, and provide some examples of presentation rules based on these features. Some example features appear in the prototype system rules below. One rather satisfying finding is that many of the rules governing presentations operate cross-modally; that is, the same rule can be used to control the analysis or generation of some aspect of both a diagram and a piece of text, when the right representation formalism is employed, resulting in a reduction of the complexity and number of rules [Arens et al. 92].

Computational attempts to plan multimedia presentations make clear that describing the information to be presented, the goals of the presenter, and the nature of the media and interlocutor merely as lists of features is too restrictive and simple to permit realistic displays to be created. The underlying intentional structure of the communication as a whole plays an important role. Different portions of the information presented support, elaborate upon, or contrast with other portions, and such relationships have to be taken into account during presentation planning. Consequently, the presentation planning rules cannot apply solely to the features we identified above; they also have to take into account a more global view of the structure of the discourse. We next describe the nature of presentation planning based upon the discourse structure.

3 Discourse & Presentation Structures

In recognition of the importance of the structure of the communication, almost all multimedia presentation systems construct some sort of internal tree-like representation of the presentation before actually producing any output. The WIP system [Wahlster et al. 92, André & Rist 92, Graf 92] plans a text/graphics description of the use of an espresso machine, starting with a database of facts about the machine and appropriate communicative goals, and using text and presentation plans to build a presentation tree. The COMET system [Feiner 88, Feiner & McKeown 90] plans text/graphic presentations of a military radio using text schemas and pictorial perspective presentation rules. The AIMI system [Burger & Marshall 91] builds a tree to produce presentations using text, maps, and tables about database information of military operations and hardware.

Considerable difference of opinion still exists about the nature of the presentation tree and the changes it undergoes during planning. Is there one tree, which once built is read and possibly annotated by the media specialists but never fundamentally altered (as in COMET)? Or should it be dynamically constructed in collaboration with the media specialists, depending on what information they can handle — in which case some complex coordination is required to ensure that everything is indeed output (as in WIP)? Or is it simpler to plan the communicative structure and then allow local reorganization by media specialists?

While we believe that the answer depends on the particularities of the presentation task, we find it probable that the simplest and most useful answer is to organize overall communication planning and media-specific presentation planning as a pipeline by maintaining a clear distinction between the intentional structure (which we call here discourse structure) and the presentational structure, and by viewing presentation planning as a process that transforms the former structure into the latter using rules based on the features of data, media, and the interlocutors¹. In this architecture, the information placed into the Discourse Structure by the Communications Planner is the *minimum amount* of information that has to be contained in the communication. The Presentation Planner is free to add to the Presentation Structure (a small amount of) additional material when its rules require so, usually the kind of material that supports the presentation itself, such as captions for figures. We do not anticipate the need for feedback from the Presentation Structure to the Discourse Structure.

This view results in the general conceptual architecture of a presentation system shown in Figure 1.

We define a Discourse Structure as follows:

¹Not surprisingly, the same simplification can be found in the text planning literature, where several researchers have argued for a distinction between what can be called text planning (using communicative intentions to construct a discourse structure) and sentence planning (using linguistically oriented knowledge about focus, theme, pronominalization, sentence length and organization, etc. to create input for the sentence generator) [Rambow & Korelsky 92].




Figure 1: Conceptual architecture of multimedia planner-presenter.

Discourse segment: A discourse segment is a tuple (*name*, *purpose/relation*, *content*), where:

- The *name* is a unique identifier for the segment.
- The *purpose/relation* is a communicative goal the presenter has WRT the perceiver's mental state. For example, the *EQUATIVE* relation represents the goal to indicate to the perceiver that the two discourse subsegments express the same information in different form, while the *JOIN* purpose/relation has no communicative semantics (no linguistic connotations are implied in a discourse structure).
- The *content* is either:
 - the semantic material to be communicated (usually "small" enough to be generated in language using a single clause). This material is usually a set of knowledge base assertions or a database fact; or
 - a single discourse segment; or
 - an ordered list of discourse segments.

Discourse structure: A discourse structure is a discourse segment which is not contained in any discourse segment and all of whose leaves (the innermost segments) contain semantic material to be communicated.

In our formulation, the discourse structure simply provides the basic organization of the information to be presented. It is neutral with respect to the media and layout of the eventual presentation.

The discourse structure relations are taken from our work expanding the relations of Rhetorical Structure Theory [Mann & Thompson 88], which holds that textual discourses are coherent only if adjacent segments are related according to the constraints imposed by a small set of relations such as *PURPOSE* (signalled in text by "in order to" or "so that"), *SEQUENCE* (signalled by "then", "next", etc.), *SOLUTIONHOOD* (no English cue phrase), and so on. We provide a taxonomy of approximately 100 relations in [Maier & Hovy 91] and an expanded taxonomy of over 170 in [Hovy et al. 92]. In the work reported here, however, the relations are used differently — to indicate communicative constraints among portions of the discourse, not necessarily linguistic realizations (if any) of such constraints.

By applying its media allocation rules, the presentation planner transforms the discourse structure into a presentation structure, which we define as follows:

Presentation segment: A presentation segment is a tuple (*name*, *purpose*, *medium*, *position*, *content*), where:

A: Total Sales (absolute amount in Ktons): 120.1
 B: Total Sales (relative % over previous year): +10%
 C: Dom. Sales (absolute amount in Ktons): 87.8
 D: Dom. Sales (relative % over previous year): +3%
 E: Exp. Sales of Edam in 86 (abs. Ktons): 11
 F: Exp. Sales of Gruyzer in 86 (abs. Ktons): 10.9
 G: Exp. Sales of Sbrinz in 86 (abs. Ktons): 8.6

Figure 2: Portion of data from cheese domain.

- The *name* is a unique identifier for the segment.
- The *purpose* is one or more communicative goals the presenter has with respect to the perceiver's mental state.
- The *medium* denotes the medium of presentation of the segment.
- The *position* specifies where in the overall presentational layout the medium should appear in the presentation (see for example the layout algorithm of [Graf 92]).
- The *content* is a list of the semantic material to be communicated, taken from the leaves or branch nodes of the discourse structure.

Presentation structure: A presentation structure is a list of presentation segments in order of communication.

Once the presentation structure is complete, each presentation segment is sent to the appropriate medium generator, whose results are then composed by the layout specialist and presented on the screen.

4 A Prototype Media Allocation Planner and Presenter

In this section we describe examples of presentation planning in order to elucidate the nature and role of the various structures, the underlying features, the media allocation rules, and the process of applying them. We describe here one possible way to address the issue of timing media allocation rules. Our prototype system applies the rules to an existing predefined discourse structure and then creates a multimedia presentation of the information using graphs, bar charts, pie charts, and text. We further discuss the nature of the rules and describe how they encapsulate the information we outlined in [Arens et al. 92] and [Van Mulken 92].

This prototype media allocation planner operates over a collection of data regarding the domestic and export sales of cheese in Switzerland during 1986, taken from an annual banking report. An English paraphrase of some of the system's data appears in Figure 2.

4.1 The Prototype in Action: Example 1

In our design, the discourse structure would be built by a communication planner on being activated by the user or host system with one or more communicative goals of the form *Describe entity X to the perceiver*, or *Motivate the perceiver to perform action Y using resource Z*, where X, Y, and Z are database or expert system domain entities or processes. For more details on text planning see [Hovy 88, Moore 89, Maybury 90, Cawsey 90,




Figure 3: Discourse structure: input to media allocation planner (letters refer to data items in Figure 2).

Reithinger 91, Hovy et al. 92].

Since however we have not yet built a communications planner and we are here interested primarily in the problem of media allocation, we activate our prototype media allocation planner with a hand-crafted discourse structure representing the communication to be made. An example discourse structure appears in Figure 3.

The ELABORATION relation holds between a nuclear concept and material providing more details about its parts, instances, specific details, etc. The EQUATIVE relation holds between two equivalent descriptions of the same underlying material. The JOIN relation holds among two or more blocks of material among whom no more detailed relation is found.

Since the allocation rules must take into account characteristics of the information to be presented, the planner has to begin by applying them to the leaves of the discourse structure. Since, as we argue, the rules must also take into account the discourse relations and overall shape or balance of the discourse structure, they have to be applied while traversing the discourse structure. Combining these two requirements, one necessarily arrives at a rule application procedure that traverses the discourse structure bottom-up, creating presentation structure nodes at the leaves of the discourse structure and then gradually filtering them upward until finally the entire presentation structure emerges at the root node. At any point during the discourse structure traversal, presentation information has been determined by the media application rules for the subtree beneath it, and the rules are applied to the node at that point and to the presentation nodes for the subtree. Earlier decisions may be, and often are, overruled by rules applied at higher nodes of the tree — i.e., when a more global view of the situation is available. Thus media allocation rules depend crucially on the nature of the discourse structure relations. For example, it is possible to merge information related by EQUATIVE into a single medium display, but not so when the relation is PURPOSE.

To make this clearer, consider the two EQUATIVE nodes of the discourse structure in Figure 3: though they are separated by several intermediate nodes and one level in the structure, their contents may be intermixed in the presentation, since the two EQUATIVE left hand sides both contain absolute numbers (of cheese sales, in Ktons) while the two right hand sides contain the corresponding relative changes (as a percentage of the previous year's sales). Given the similar nature of the information to be presented in the two cases, appro-




Figure 4: Presentation structure: output of media allocation planner (letters refer to data items in Figure 2).

appropriate media allocation rules reorganize the information into the presentation segments P1 and P3 in Figure 4.

The system contains approximately a dozen media allocation rules. Each rule may refer only to specific features of the information to be presented by leaves of the current subtree, to the current discourse segment's purpose/relation, and to the system's communicative goals. For example, two rules governing the abovementioned reorganization can be paraphrased as:

```
; Display two instances of Equatively similar
; info using different instances of same medium.
R1: IF AND (relation == EQUATIVE)
      (contents (leaf1) <> contents (leaf2))
THEN (create presentation-node-1
      loc: current-relation-node
      contents: leaf1
      constraint: medium(presentation-node-1) ==
                  medium(presentation-node-2)
      (create presentation-node-2
      loc: current-relation-node
      contents: leaf2
      constraint: medium(presentation-node-1) ==
                  medium(presentation-node-2)

; Display Elaborated or Contrastive info using same
; medium if info is of same type, merging leaf nodes
; from both subtrees. Apply separately to each leaf.
R2: IF AND (OR (relation == ELABORATION)
              (relation == CONTRAST))
      (type (contents (leaf1)) ==
       type (contents (leaf2)))
      (medium (leaf1) == medium (leaf2))
THEN (create presentation-node
      loc: current-relation-node
      contents: (leaf1 leaf2)
      medium: medium(leaf1))
```

In the example, as we move up the tree, Rule R1 applies to both EQUATIVE subtrees (since the leaf content is not identical), and hence presentation segments are created at both EQUATIVE relations. Subsequently, the presentation segment for the rightmost EQUATIVE relation is filtered upwards to the top JOIN relation. Finally, when the traversal process arrives at the ELABORATION node, Rule R2 applies, and corresponding portions of the two lists of presentation segments are merged to form P1 and P3, which then give rise to the two bar charts in Figure 5.

Whenever two or more presentation segments are bound by a relation and are not merged, a default rule creates a new presentation segment between them to state in a sentence the relation binding them. At




Figure 5: (Annotated) Presentation produced from presentation structure.

present, the linking sentence is canned; it will eventually be generated by ISI's PENMAN sentence generator [Mann & Matthiessen 83].

4.2 Altering the Data: Example 2

If the media allocation rules are to prove their worth, they must effect appropriate changes even with subtle changes in the data characteristics. In order to test our rules, we performed several trials by changing just one or two basic characteristics of the data in order to see what the resulting presentations would be. In general, we found the results satisfactory.

As mentioned above, in order to ensure the clearest presentation of information, the media allocation rules must be sensitive to the discourse structure as well as to particularities of the media and information itself. We illustrate this in the following example. Starting with the same data as in the previous section, we changed one data characteristic, namely the value of node B, from a relative percentage (+10%) to an absolute amount (+10.9 Ktons). Given the sensitivity of the allocation rules to such characteristics, the merging of node B with node D cannot occur, resulting in two distinct presentation nodes. These two nodes remain unmerged throughout the media planning process; they percolate up the discourse structure, until eventually they appear separately in the presentation structure Figure 6 (Cf. Figure 4). The final presentation is shown in Figure 7.

(name: P1 medium: barchart content: (A C)	name: P2 medium: text content: (Equative Relation)	name: P3 medium: text content: (B')
name: P4 medium: text content: (D)	name: P5 medium: text content: (Join Relation)	name: P6 medium: piechart content: (E F G)

Figure 6: Presentation structure: output of media allocation planner (letters refer to Figure 2, except B').




Figure 7: Presentation produced from presentation structure after one data item, B', is altered.

5 Conclusion

In this paper, we recall that multimedia presentation planners require information to medium allocation rules based on very specific kinds of knowledge: the characteristics of the data (including its organization into a discourse structure), the characteristics of the media at hand, the communicative intentions of the presenter, and the characteristics of the perceiver. We argue that the medium allocation planning process should transform the discourse structure into a presentation-oriented construct we call the *presentation structure* by performing a bottom-up filtering traversal of the former. To illustrate the ideas, we outline a prototype planner that changes its output depending on subtle differences in the data to be presented. Two examples show how rules based on single data characteristics and discourse structure relations successfully handle display requirements which

fixed data-to-medium mappings could not.

The operation of the media application rules is clear from the example. In all cases, their actions move presentation segments from one level in the discourse structure to the next higher level, and eventually out of the root in the form of a presentation structure, which is linear. Operating (in the example) on characteristics of the information (the data) itself and on the discourse structure relations, they perform three types of action:

1. create new presentation segments containing data,
2. combine presentation segments into new ones that then contain lists of presentation segments, and
3. simply copy lower-level presentation segments up the discourse structure.

Specifically, we note that the data characteristics that media allocation rules must consider include the discourse structure, which embeds data and the presenter's goals using purpose/reasons, as defined in the paper. Without knowledge of the discourse structure, it is impossible to formulate rules that correctly distribute information among the available media.

Finally, we note that the explicit separation of Communication Planning from Presentation Planning (and the accompanying separation of Discourse Structure from Presentation Structure), where the Presentation Planner is allowed to add material to the Presentation Structure that is not present in the Communication Structure, facilitates the separation of the various kinds of knowledge that play roles at various stages of the planning of multimedia presentations, and may help solve some of the difficult problems inherent to this task.

References

- [André & Rist 92] André, E. and Rist, T. 1992. The Design of Illustrated Documents as a Planning Task. German Research Center for AI (DFKI) Research Report.
- [Arens et al. 88] Arens, Y., Miller, L., Shapiro, S.C. and Sondheimer, N.K. 1988. Automatic Construction of User-Interface Displays. In *Proceedings of the 7th AAAI Conference*, St. Paul, MN (808-813). Also available as USC/Information Sciences Institute Research Report RR-88-218.
- [Arens & Hovy 90] Arens, Y. and Hovy, E.H. 1990. How to Describe What? Towards a Theory of Modality Utilization. In *Proceedings of the 12th Cognitive Science Conference*, Cambridge, MA.
- [Arens et al. 92] Arens, Y., Hovy, E.H. and Vossers, M. 1992. Describing the Presentational Knowledge Underlying Multimedia Instruction Manuals. In volume edited by M. Maybury, forthcoming.
- [Bertin 83] Bertin, J. 1983. *Semiology of Graphics*, (trans. by J. Berg). Madison: University of Wisconsin Press.
- [Burger & Marshall 91] Burger, J. and Marshall, R. 1991. AIMI: An Intelligent Multimedia Interface. In *Notes of the AAAI-91 Workshop on Intelligent Multimedia Interfaces*, Anaheim, CA (23-28).

- [Cawsey 90] Cawsey, A. 1990. Generating Communicative Discourse. In *Current Research in Natural Language Generation*, Dale, R., Mellish, C., and Zock, M. (eds), Academic Press, Boston (75-102).
- [Dwyer 78] Dwyer, F.M. 1978. *Strategies for improving visual learning*. State College, PA: Learning Services.
- [Feiner 88] Feiner, S. 1988. An Architecture for Knowledge-Based Graphical Interfaces. ACM/SIGCHI Workshop on *Architectures for Intelligent Interfaces: Elements and Prototypes*, Monterey, CA.
- [Feiner & McKeown 90] Feiner, S. and McKeown, K.R. 1990. Coordinating Text and Graphics in Explanation Generation. In *Proceedings of the 8th AAAI Conference* (442-449).
- [Fleming & Levie 78] Fleming, M. and Levie, H.W. 1978. *Instructional Message Design: Principles from the Behavioral Sciences*. New Jersey: Educational Technology Publications.
- [Graf 92] Graf, W. 1992. Constraint-Based Graphical Layout of Multimodal Presentations. German Research Center for AI (DFKI) Research Report.
- [Hartley 85] Hartley, J. 1985. *Designing Instructional Text*. (2nd edition). Great Britain: Kogan Page Ltd.
- [Hovy 88] Hovy, E.H. 1988. Planning Coherent Multisentential Text. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, Buffalo, NY.
- [Hovy & Arens 90] Hovy, E.H. and Arens, Y. 1990. When is a Picture Worth a Thousand Words? — Allocation of Modalities in Multimedia Communication. Presented at the *AAAI Symposium on Human-Computer Interaction*, Stanford University.
- [Hovy & Arens 91] Hovy, E.H. and Arens, Y. 1991. Automatic Generation of Formatted Text. In *Proceedings of the 10th AAAI Conference*, Anaheim, CA.
- [Hovy et al. 92] Hovy, E.H., Lavid, J., Maier, E., Mital, V., and Paris, C.L. 1992. Employing Knowledge Resources in a New Text Planning Architecture. In *Aspects of Automated Natural Language Generation*, R. Dale, E. Hovy, D. Rösner, O. Stock (eds). Heidelberg: Springer Verlag Lecture Notes in AI no. 587 (57-72).
- [Larkin & Simon 86] Larkin, J.H. and Simon, H.A. 1986. Why a Diagram is (Sometimes) Worth Ten Thousand Words. *Cognitive Science* 11 (65-99).
- [Mackinlay 86] Mackinlay, J. 1986. *Automatic Design of Graphical Presentations*. Ph.D. dissertation, Stanford University.
- [Maier & Hovy 91] Maier, E. and Hovy, E.H. 1991. A Metafunctionally Motivated Taxonomy for Discourse Structure Relations. In *Proceedings of the 3rd European Workshop on Language Generation*, Judenstein, Austria (38-45).
- [Mann & Matthiessen 83] Mann, W.C. and Matthiessen, C.M.I.M. 1985. Nigel: A Systemic Grammar for Text Generation. In *Systemic Perspectives on Discourse: Selected Papers from the 9th International Systemics Workshop*, R. Benson and J. Greaves (eds), Ablex: London, England. Also available as USC/ISI Research Report RR-83-105.
- [Mann & Thompson 88] Mann, W.C. and Thompson, S.A. 1988. Rhetorical Structure Theory: Toward a Functional Theory of Text Organization. *Text* 8(3) (243-281).
- [Maybury 90] Maybury, M.T. 1990. *Planning Multisentential English Text Using Communicative Acts*. Ph.D. dissertation, Cambridge University. Also available as RADC Technical Report 90-411.
- [Mayer 89] Mayer, R.E. 1989. Systematic Thinking Fostered by Illustrations in Scientific Text. *Journal of Educational Psychology* 81 (240-246).
- [Moore 89] Moore, J.D. 1989. A Reactive Approach to Explanation in Expert and Advice-Giving Systems. Ph.D. dissertation, University of California at Los Angeles.
- [Neal 90] Neal, J.G. 1990. Intelligent Multi-Media Integrated Interface Project. SUNY Buffalo. RADC Technical Report TR-90-128.
- [Rambow & Korelsky 92] Rambow, O. and Korelsky, T. 1992. Applied Text Generation. *Proceedings of the Applied Natural Language Processing Conference*, Trento (40-47).
- [Reithinger 91] Reithinger, N. 1991. *Eine Parallele Architektur zur Inkrementellen Generierung Multimodaler Dialogbeiträge*, Ph.D. dissertation, University of the Saarland.
- [Roth & Mattis 90] Roth, S.F. and Mattis, J. 1990. Data Characterization for Intelligent Graphics Presentation. *CHI'90 Proceedings* (193-200).
- [Tufte 83] Tufte, Edward R. 1983. *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press.
- [Tufte 90] Tufte, Edward R. 1990. *Envisioning Information*. Cheshire, CT: Graphics Press.
- [Twyman 85] Twyman, M. 1985. Using Pictorial Language: A Discussion of the Dimensions of the Problem. In Duffy, T.M. & Waller, R. (Eds.), *Designing Usable Texts*. Florida: Academic Press (245-312).
- [Van Mulken 92] Van Mulken, S. 1992. Master's thesis. Master's thesis, University of Nijmegen, The Netherlands (forthcoming).
- [Vossers 91] Vossers, M. 1991. Automatic Generation of Formatted text and Line Drawings. Master's thesis, University of Nijmegen, The Netherlands.
- [Wahlster et al. 92] Wahlster, W., André, E., Bandyopadhyay, S., Graf, W., Rist, T. 1992. WIP: The Coordinated Generation of Multimodal Presentations from a Common Representation. In *Computational Theories of Communication and their Applications*, A. Ortony, J. Slack, and O. Stock (eds), Berlin: Springer Verlag.