Reassessing Rhetorical Abstractions and Planning Mechanisms[1]

Daniel D. Suthers

Department of Computer and Information Science
University of Massachusetts
Amherst, Massachusetts 01003

suthers@cs.umass.edu

Abstract

The utility of rhetorical abstractions and certain text planning mechanisms were assessed from the standpoint of accounting for how an explainer chooses and structures content under multiple perspectives to meet knowledge communication goals. This paper discusses ways in which they were found to be inadequate, argues for greater emphasis on an epistemological level of analysis, and proposes a mixed architecture matching computational mechanisms to the explanation planning subtasks they are suited for.

Introduction

Our research is concerned with explanation in its broad sense, as "the act or process of making plain or comprehensible; elucidation; clarification" (*American Heritage Dictionary*). In many physical science domains, an explanation can be based on a variety of models of the phenomenon being explained. These models differ on the type of properties and relationships emphasized; what ontology is used and whether the phenomenon is described at a macroscopic, microscopic, or atomic granularity; what factors are ignored or assumed to be constant; the use of general statements vs. concrete examples; and in general on what concepts function as primitives, providing the basis for understanding the topic phenomenon. Selection of an appropriate model of the topic is an important aspect of content selection which impacts on the interlocutor's comprehension of the explanation and on its appropriateness for his or her purposes. We use the term **perspective** to refer to an abstract characterization of the kind of knowledge provided by a class of models. Our notion of perspective is a composite of distinctions made by Falkenhainer & Forbus (submitted), Stevens & Collins (1980) Stevens & Steinberg (1981) and White & Frederiksen (1989).

Most existing work in explanation and text planning has been directed at other problems, and hence has utilized single-perspective knowledge bases to simplify the research. (Notable exceptions include McCoy, 1989 and McKeown *et al.* 1985.) A number of such research efforts emphasize rhetorical abstractions for the analysis of natural explanations and for expressing a theory of explanation (Hovy, 1988; Mann & Thompson, 1986; Maybury, 1988; McKeown, 1985; Moore, 1989). A variety of mechanisms for selecting and organizing the content of explanations have also been explored. This includes schema filling (McKeown 1985) and structure matching (Hovy 1988), graph traversal algorithms (Paris & McKeown 1986), and top-down expansion planning (Cawsey, 1989; Moore, 1989).

Our own work on choosing explanatory content from multiple-perspective knowledge bases has uncovered some limitations of rhetorical abstractions, and led us to question previous applications of the computational mechanisms listed above. The purpose of this paper is to present our perception of the roles and limitations of these items, and suggest some alternatives. (We do not emphasize our work on perspective and content organization here: see Suthers & Woolf, 1990.) We begin with an example, used to illustrate some of the problems.

An Example

The following example explanation will be used to illustrate our points. The domain is elementary electricity. We emphasize the communication of an understanding of concepts such as "charge", "current", and "capacitance" within the context of qualitative reasoning about the behavior of simple circuits and their components. The explanation is an edited version of a human protocol.

Q1: *How does a capacitor store charge?*

E1: A capacitor

E2: can be thought of as two flat metallic plates

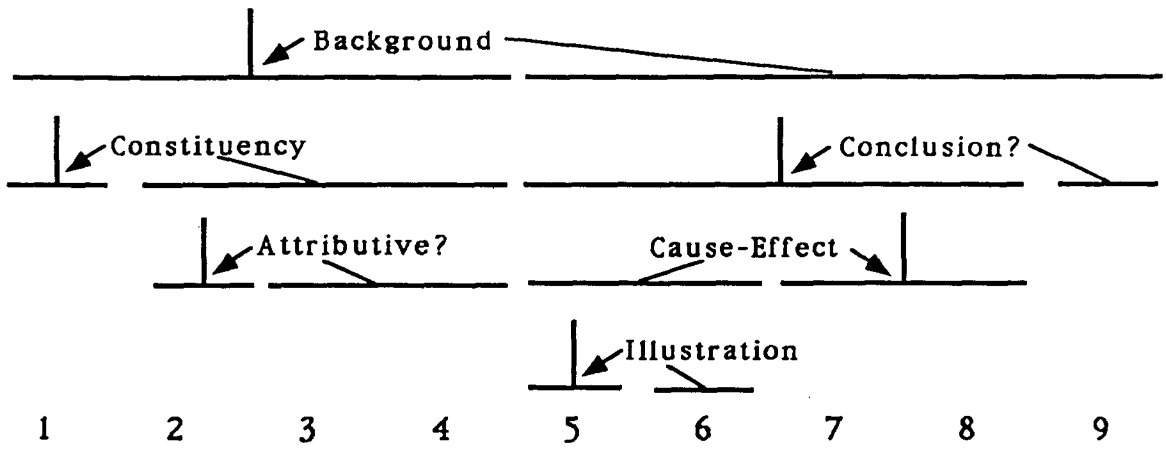

Figure 1: One possible RST analysis (partial)

E3: situated close to each other and parallel to each other,

E4: with air in between.

E5: If you connect a capacitor to a battery,

E6: as follows:

```
   ---------------------
   |                  |   +
-------             -----
  ---               -----
   |                  |   -
   ---------------------
```

E7: then positive charge builds up on one plate and negative charge on the other,

E8: until the voltage across the capacitor is equal to the voltage of the battery.

E9: The charge stays there when you disconnect the battery.

One possible rhetorical analysis of the first explanation, using relations from McKeown (1985) and Mann & Thompson (1986), is given in Figure 1. (Whether or not this is an optimal analysis is not the point; rather we are concerned with the role of rhetorical abstractions and various characterizations of the content planning task in accounting for the explanation.) The analysis does point out some features of interest. For example, we note that the explanation starts with Background material (E1-E4) before proceeding to the primary explanation (E5-E9). The Background relation in this case describes the high level organization of the explanation. The question is how this relation should be manifest in the mechanism for generating the explanation. Also of interest are the explainer's determination that a process account of how the capacitor carries out its function is appropriate; use of an abstract structural model of the capacitor, simplified to what is needed to support the process account; and use of a concrete situation for the enablement condition of the process.

Problems with Rhetorical Abstractions

Rhetorical relations were an important development in explanation research, since they provided a first pass at abstractions for a general description of explanatory structure, and in bringing various roles of the parts of explanation into the foreground, pointed out phenomena in need of further study. They are also useful via their "relational propositions" (Mann & Thompson 1983), for conveying propositional information implicitly in the structure of the text, hence reducing its length and redundancy. However, we claim that rhetorical abstractions fail to make the necessary distinctions for further advances in a theoretical understanding of explanation.

Potpourri. Rhetorical abstractions are descriptive of explanatory *text*, i.e. the end product of some explanation generation process, and so describe with one device structure due to a variety of distinct knowledge sources bearing on such a process. These knowledge sources operate on different levels of information, and hence need to be separated in a theory of explanation generation. For example, (drawing on relations in McKeown, 1985 and Mann & Thompson, 1986) some correspond directly to the fundamental structure of domain objects and processes (e.g. Constituency and Causality), while others are about derived relations between concepts which may vary according to the context (e.g. Comparison). Relations such as Amplification, Background, Evidence, and Illustration are primarily about relationships between propositions which arise in part out of consideration of what the interlocutor knows and needs to know to better grasp a point. Illocutionary acts are involved in the relations as well (most blatantly, Concession; others are not themselves illocutionary acts but only make sense in the context of certain such acts). Finally, relations such as Topic and Conclusion appear to be due to conventions governing writing style which direct focus of attention. Grosz & Sidner (1986) made similar criticisms from the standpoint of characterizing discourse coherence. They suggested that each rhetorical relation combines domain information with certain general relations between propositions, between actions, and between intentions.

Implicit Features Unaccounted For. Rhetorical abstractions are also inappropriate for a theory of content selection because, in describing the final text, they fail to identify important relations between the chosen content and external material, such as what is known about the user, or information left *out* of the explanation. For example, E5-E6 is more specific than is necessary and includes a concrete example. A more general and accurate way to state the condition for initiation of the charging process would be "If a voltage is applied to the plates ...". However, the explainer has opted to replace this with one of the many particular configurations which meet the condition. A rhetorical analysis of the text *cannot even tell us that this has happened*, let alone why, because it does not describe relations between the contents of the text and what is *not* included, viz., other models of the process being described. It can only report that an illustration is being used. Another example is provided by E9. Retention of charge when the voltage is removed is what is meant by "storage" in this case, so the fact expressed in E9 is essential to answering the question. Rhetorically, we can only identify relationships E9 has to the rest of the text, e.g. that it is a Conclusion. This does not illuminate the relationship between its content and the goal of the explanation.

Epistemological Analysis. The success of an explanation is primarily a function of choice and organization of knowledge. Hence, to account for how these choices further knowledge communication goals, one must examine explanation in part from an epistemological standpoint. Such an analysis examines how explanations are guided by:

- the types of knowledge in a given domain, and its logical and etiological structure (Rissland, 1978);
- the types of knowledge which, in principle, could fulfill a given request for information;
- the role of an individual's knowledge in understanding new concepts and situations, and hence in understanding a given explanation (Paris, 1987); and
- the ways in which individuals are willing or able to undertake conceptual change (Goldstein, 1979; Hewson, 1981; White & Frederiksen, 1989; vanLehn, 1987).

As discussed in Suthers (1989), most previous work has offered solutions to the subproblems of explanation in the form of mechanisms and data structures which are in part the result of, rather than the expression of, epistemological considerations. Epistemological problems have been avoided through direct selection of content based on well-formulated queries; the simplicity of the knowledge bases used, which only permit one way of discussing each topic; and through implicit conflation of epistemological constraints on organization of the explanation with those of rhetorical and linguistic origin. A major goal of our research (Suthers & Woolf, 1990) is an explicit theory of the epistemological structure of the activity of explaining.

Roles of Computational Mechanisms

In this section we illustrate how inclusion of background material and the use of multiple perspectives pose problems for various mechanisms in the literature, and suggest a mixed architecture solution.

Structure Matching. By "structure matching" we mean methods where abstract descriptions of the structure of explanations are matched to a collection of propositions (or other content to be expressed) in order to organize this material. This includes bottom-up composition of structural units (Hovy, 1988) and schema filling (McKeown, 1985). We question their adequacy for accounting for the prerequisite structure of explanations, such as the Background relation of the example. In our view, the explanation is organized this way because the explainer recognized in his process model (expressed in E5-E9) concepts the interlocutor may not be familiar with (the parts of a capacitor), and then added material prerequisite to understanding the process explanation (the structural description expressed in E1-E4). The background material is not automatically part of the relevant knowledge pool for this question, and structure matching methods leave choice of content and perspective to other mechanisms. Suppose, then, that some other mechanism accounts for inclusion of the background in the pool. It is included *as background* by virtue of relationships between the interlocutor's assumed knowledge state and the conceptualizations contained in the first attempt at a knowledge pool. Pattern matching techniques which are ignorant of such relationships and see only the composite pool would be unable to identify the part of the knowledge pool which plays the role of "background", and account for placement of background before primary material.

Graph Traversal. These are algorithms for selectively following links in a knowledge base, with the path so traced out providing the content and structure of the explanation (Paris & McKeown 1986). Such methods model how an explanation exploits the structure of knowledge. They implicitly embody the heuristic that it will be easier for the interlocutor to reconstruct the knowledge if it is presented such that each unit is introduced in relation to the previous unit. For example, parts of an object are introduced in relation to containing or adjacent parts, and process descriptions organized to follow temporal and causal relations in the forward direction. However, graph traversal is limited to modeling local organization, or global organization which is a composite of local choices. Traversal methods don't naturally extend to global organization which occurs at a higher level of abstraction than the links followed, e.g. the presentation of a coherent structural model before a process account begins.

Top-Down Goal Expansion. Top-down expansion of a discourse goal (Cawsey, 1989; Moore, 1989) is a stronger candidate for a uniform mechanism for explanation, integrating content selection and structuring. The background problem can be handled with preconditions on plan operators, as Cawsey does. It simplifies modeling explanations such as our example if one can specify the kind of background knowledge required in preconditions at the highest level of plan operators. To illustrate, consider a rhetorical plan operator which contains an optional satellite for Background but does not specify what constitutes background knowledge. Expansion of the satellite would have to be predicated on a comparison of the content selected by expansion of the nucleus with the user model. This decision could not be made at the time the operator is selected by the planning system, since the knowledge pool for the nucleus would not have been selected yet. Instead, one would have to place the satellite decision on hold, expand the nucleus, collect together the knowledge selected at the leaves of its expansion, and perform the comparison before deciding on the satellite. (One could make the decisions concerning the need for background locally to each leaf, avoiding the need for a high level decision depending on knowledge selected at many localities. But then one could not model the structure of explanations such as this one, where the need for prerequisite material is anticipated and provided in advance as a coherent model, rather than as an interruption to the flow of the process description.) Then, if it was decided that some background was required, expansion of the satellite would have to occur under a binding of some variable in the satellite to the concepts for which background is required.

With regards to perspective, some problems emerge. Content selection in top-down expansion can be influenced by the current perspective if some mechanism for sharing perspective decisions across the expansion tree is provided. However, neither the choice of perspective nor the actual mechanism by which it influences content selection are modeled appropriately by top-down expansion. Choice of perspective tries to balance the (sometimes conflicting) constraints of adequacy and comprehensibility. Adequacy constraints come from examination of the informative goals (McKeown *et al.* 1985). Comprehensibility requires answering the question: given the concepts which have been used in the dialogue so far, and/or which the explainer has evidence the interlocutor is familiar with, what other concepts are also likely to be familiar? This suggests a strength-of-association mechanism (McCoy 1989), which we comment on further in the next section. As Hovy (1990) points out, top-down planning is *prescriptive*, and does not handle conflicting goals easily. The influence of perspective on content selection involves what Hovy calls *restrictive* planning: some perspective goals, once generated, remain active throughout a span of discourse, and operate as preferences applied to choice points in content selection. They cannot be erased once satisfied, and they may change dynamically. Finally, McDonald & Pustejovsky (1985) point out that a uniform mechanism may not be desirable, as it incorrectly implies that all information is equally available during each point of the text planning process. A principled match of distinct computational mechanisms to each subtask of the

```
(Query) --> Identify
               |
               V
            Refine <----
               |       |
               V       |
            Retrieve   |
               |       |
               V       |
            Modify -----
               |
               V
             Order --> (Realize)
```

Figure 2: Content planning tasks

explanation planning task is one way to express one's theory of what kind of process explanation is. For these reasons, we postulate a separate level of planning for coordinating perspective decisions.

A Mixed Architecture. In our current approach, planning occurs at two granularities: selection and retrieval of coherent packages of knowledge similar to Suthers' (1988a,b) "views" or Souther, Acker, Lester, & Porter's (1989) "viewpoints"; and editing and ordering the propositions and examples which make up these views. Planning at the granularity of views is concerned with purely epistemological constraints on identification of appropriate content, including choice of perspective and prerequisite explanations. At the finer granularity, further epistemological constraints governing the comprehensibility of the explanation are applied, and rhetorical and linguistic constraints play a role as well.

The question at hand is: *what sort of task is content planning?* We postulate several, and comment on possible computational mechanisms for each. Figure 2 gives the rough relationships between the different tasks. The process is reminiscent of case-based design, where one identifies and refines design specifications, and retrieves and reconfigures a previous solution to fit the circumstances.

An explainer must first **identify** at least some of its goals. These are of two types, as discussed above: prescriptive informative goals, and restrictive goals such as comprehensibility. We mentioned that the latter suggests a concept association mechanism. This could be done by activating weighted links between concepts, or though intermediate frames of attribute weights, as in McCoy (1989). However, one would have to install a potentially combinatorial number of associations between concepts. Because we wish to make an epistemological theory explicit and implement it with an abstract interface to the knowledge base, we are investigating use of a mechanism akin to prototype induction, to generate an abstract description of the desired perspective from the user model and dialogue history.

At the view granularity, content selection may be seen as **refinement** of explanatory goals into a specification of the appropriate addition to the relevant knowledge pool. This refinement includes consideration of what type of knowledge, in principle, could fulfill the informative goal; of the concepts the interlocutor is likely to understand; and of the models which have already been shared in the dialogue. We are attempting to treat this refinement task as top-down planning, though our work has not progressed far enough to comment further on this.

Retrieval requires a knowledge-base specific mechanism: its only theoretical importance to us is that it correctly operationalize the dimensions used to describe the desired view (Suthers 1988a,b). We are examining a variant of compositional modeling (Falkenhainer & Forbus, submitted) for this purpose.

Once a knowledge pool is available, some data driven activities occur at a finer granularity, resulting in **modification** of the relevant knowledge pool. This includes filtering activities, such as removing particular propositions likely to be familiar to the interlocutor; and augmenting activities, such as illustrating abstract statements with examples. These are *opportunistic* planning tasks, that can be approached with *critics* which match to the knowledge pool and user model, and specify replacements or deletions to be made. As shown in figure 2, data driven operators may also reinvoke content planning at the refinement level to access material in a different type of model. For example, we have seen how process propositions may involve use of structural concepts the interlocutor is not likely to understand, causing the explainer to plan a prerequisite explanation. An explicit record of the prerequisite relations between views is created, important for ordering the explanation.

Ordering is sensitive to prerequisite and illustration links installed by the modification processes. Hence figure 2 should not be interpreted as a claim that ordering considerations do not arise during content selection. The ordering task involves two kinds of processes. One embodies epistemological constraints by *exploiting* existing structure in the knowledge pool. As discussed previously, techniques for traversing links in the knowledge pool apply here. The result will likely be a partial ordering. Further ordering requires *imposition* of structure for rhetorical, linguistic, and pictorial reasons during realization (generation of text and graphics). Matching of rhetorical patterns to the knowledge pool may be more appropriate at this later stage.

In summary, we have argued for a mixed architecture matching prescriptive, restrictive, and opportunistic mechanisms to explanation subtasks. Coordination of these diverse processes may, at the implementation level, require an agenda control mechanism as in Nirenburg, Lesser, & Nyberg (1989).

Conclusions

Single-perspective knowledge bases, i.e. those which provide only a single conceptual basis for a given description or explanation, have dominated existing work in planning expository text. This research has overemphasized rhetorical abstractions as the basis for theories of content planning, and used computational formalisms such as top-down goal expansion, traversal algorithms, and opportunistic structure matching which, taken alone, fail to fully account for the search for appropriate conceptualizations during content selection. We suggest a greater emphasis on epistemological abstractions, and viewing content selection in terms of identification, refinement, retrieval, modification, and ordering tasks. Rhetorical abstractions retain a place in initial analyses of explanations, to point out features in need of further study, and in the later stages of text planning, to further constrain a partial ordering of content and implicitly convey content via textual structure. Finally, the devices of top-down goal expansion, traversal algorithms, and structure matching retain potential utility for high level content planning, exploiting the structure of knowledge when ordering an explanation, and further ordering an explanation on a rhetorical basis, respectively.

Acknowledgments

The author has been supported by the National Science Foundation under grant number MDR 8751362, and by Apple Computer, Inc., Cupertino, CA. Partial support was also received from the Office of Naval Research under a University Research Initiative Grant, contract no. N00014-86-K-0764. Thanks are due to my advisors, Edwina Rissland and Beverly Woolf; to Klaus Schultz for his expert explanations; to Matthew Cornell for assistance with the research, and to the reviewers for comments.

References

Cawsey, A. (1989). *Generating Explanatory Discourse: A Plan-Based, Interactive Approach.* Ph.D. Dissertation, University of Edinburgh.

Falkenhainer & Forbus (submitted). Compositional Modeling: Finding the right model for the job. Submitted to *Artificial Intelligence.*

Goldstein, I. (1979). The genetic graph: a representation for the evolution of procedural knowledge. *International Journal of Man-Machine Studies,* vol. 11, pp. 51-77.

Grosz, B. J. & Sidner, C. L. (1986). Attention, intention, and the structure of discourse. *Computational Linguistics,* vol. 12, no. 3, pp. 175-204.

Hewson, P. W. (1981). A conceptual change approach to learning science. *European Journal of Science Education,* vol. 3, no. 4, pp. 383-396.

Hovy, E. H. (1988). Planning coherent multisentential text. *Proc. 26th Meeting of the ACL, Buffalo, New York;* Reprinted as ISI/RS-88-208, Information Sciences Institute, Marina del Rey, California.

Hovy, E. H. (1990). Pragmatics and natural language generation. *Artificial Intelligence,* vol. 43, no. 2, pp. 153-197.

Mann, W. C., & Thompson, S. A. (1983). Relational Propositions in Discourse. ISI/RR-83-115, Information Sciences Institute, University of Southern California, Marina del Rey, California.

Mann, W. C., & Thompson, S. A. (1986). Rhetorical structure theory: Description and construction of text structures; *Proceedings of the NATO Advanced Research Workshop on Natural Language Generation,* Nijmegen, The Netherlands, August 19-23, 1986.

Maybury, M. (1988). Explanation rhetoric: The rhetorical progression of justifications. *Proc. AAAI-88 Workshop on Explanation,* St. Paul, August 22, 1988, pp. 16-20.

McCoy, K. F. (1989). Generating context-sensitive responses to object-related misconceptions. *Artificial Intelligence,* vol. 41, no. 2, pp. 157-195.

McDonald, D. D. & Pustejovsky, J. D. (1985). Description-directed natural language generation. *Proc. 9th Int. Joint Conf. on Artificial Intelligence,* Los Angeles, August 18-23, 1985, pp. 799-805.

McKeown, K. R. (1985). Discourse strategies for generating natural language text. *Artificial Intelligence,* vol. 27, no. 1, pp. 1-41.

McKeown, K. R., Wish, M., & Matthews, K. (1985). Tailoring explanations for the user. *Proc. 9th Int. Joint Conf. on Artificial Intelligence,* Los Angeles, August 18-23, 1985, pp. 794-798.

Moore, J. D. (1989). *A Reactive Approach to Explanation in Expert and Advice-giving Systems.* Ph.D. Dissertation, University of California, Los Angeles.

Nirenburg, S., Lesser, V., & Nyberg, E. (1989). Controlling a language generation planner. *Proc. 11th Int. Joint Conf. on Artificial Intelligence,* Detroit, pp. 1524-1530.

Paris, C. L. (1987). Combining discourse strategies to generate descriptions to users along a Naive/Expert spectrum. *Proc. 10th Int. Joint Conf. on Artificial Intelligence,* August 1987, Milan, Italy, pp. 626-632.

Paris, C. L. & McKeown, K. R. (1986). Discourse strategies for describing complex physical objects; *Proceedings of the NATO Advanced Research Workshop on Natural Language Generation,* Nijmegen, The Netherlands, August 19-23, 1986. Also published by Martinus Nijhoff, Dordrecht, 1987.

Rissland, E. L. (1978). (Formerly Michener.) Understanding Understanding Mathematics. *Cognitive Science,* vol. 2, no. 4.

Souther, A. Acker, L., Lester, J., & Porter, B. (1989). Using view types to generate explanations in intelligent tutoring systems. *Proc. Cognitive Science Conf.,* Montreal, 1989.

Stevens, A. L. & Collins, A. (1980). Multiple conceptual models of a complex system. In R. E. Snow, P. Federico, and W. E. Montague (Eds.), *Aptitude, Learning, and Instruction (Vol. 2).* Hillsdale, NJ: Erlbaum, 1980. pp. 177-197.

Stevens, A. L. & Steinberg, C. (1981). A typology of explanations and its application to intelligent computer aided instruction. Report No. 4626, Bolt Beranek and Newman Inc., Cambridge, MA.

Suthers, D. D. (1988a). Providing multiple views of reasoning for explanation. *Proc. Int. Conf. on Intelligent Tutoring Systems,* Montreal, June 1988, pp. 435-442.

Suthers, D. D. (1988b). Providing multiple views for explanation. *Proc. AAAI-88 Workshop on Explanation,* St. Paul, August 22, 1988, pp. 12-15.

Suthers, D. D. (1989). Perspectives in Explanation; COINS Technical Report 89-24, Computer and Information Science, University of Massachusetts, Amherst.

Suthers, D. D., & Woolf, B. P. (1990). Accounting for the Epistemological Structure of Explanation. *Spring Symposium on Knowledge Based Environments for Learning and Teaching,* March 27-29, Stanford. Available as COINS Technical Report 90-36, Computer and Information Science, University of Massachusetts, Amherst.

vanLehn, K. (1987). Learning one subprocedure per lesson. *Artificial Intelligence,* vol. 31, no. 1, pp. 1-40.

White, B. Y. & Frederiksen, J. R. (1990). Causal Model Progressions as a Foundation for Intelligent Learning Environments. *Artificial Intelligence,* vol. 42, no. 1, pp. 99-157.