# Task-Appropriate Hybrid Architectures for Explanation[1]

**Daniel D. Suthers**[2]

## COINS Technical Report 91-48

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

suthers@cs.umass.edu

**Abstract.** In this paper, I am concerned with making distinctions necessary for an explicit theory of explanation generation, and analyzing how a spectrum of architectural and structural ideas fit together to provide required functionality. Several information processing tasks involved in the choice and organization of the content of an explanation are identified. These are best modeled by distinct mechanisms; hence a particular class of hybrid planning architectures most clearly reflects the nature of the explanation task. I describe an implemented explanation planner to exemplify such an architecture. Various implications of the architecture are discussed, including a classification of structuring relations based on their sources and roles in planning; the elimination of goal-posting preconditions or satellites from plan operators; and the level at which nondeterminism is handled.

---

[1] This is an expanded version of a paper prepared for the AAAI Workshop on Comparative Analysis of Explanation Planning Architectures, held July 14, 1991 in Anaheim, CA.

[2] The author has been supported by the National Science Foundation under grant number MDR 8751362, and by Apple Computer, Inc., Cupertino, CA.

# Contents

1 Introduction 1

2 Roles and Limitations of Planning Mechanisms 2

2.1 Imposition of Structure . . . . . . . . . . . . . . . . . . . . . . . . . 2
2.2 Exploitation of Structure . . . . . . . . . . . . . . . . . . . . . . . . 3
2.3 Prescriptive Planning and Opportunism . . . . . . . . . . . . . . . . . 4
2.4 Perspective: Associative and Restrictive Tasks . . . . . . . . . . . . . 5
2.5 Summary of Analysis . . . . . . . . . . . . . . . . . . . . . . . . . . 6

3 A Hybrid Architecture 7

3.1 Goal Identification . . . . . . . . . . . . . . . . . . . . . . . . . . . 8
3.2 Prescriptive Refinement and Informativeness . . . . . . . . . . . . . . 8
3.3 Retrieval and Epistemic Functions . . . . . . . . . . . . . . . . . . . 12
3.4 Opportunistic Modification and Pedagogical Structure . . . . . . . . . 14
3.5 Exploitive Ordering and Epistemological Gradients . . . . . . . . . . 17
3.6 Realization . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 18

4 Discussion 21

4.1 Structural Abstractions for Generation . . . . . . . . . . . . . . . . . 21
4.2 Planning Without Preconditions . . . . . . . . . . . . . . . . . . . . . 22
4.3 Nondeterminism and Revision . . . . . . . . . . . . . . . . . . . . . . 23
4.4 Tradeoffs . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 23
4.5 Limitations of the Architecture . . . . . . . . . . . . . . . . . . . . . 24
4.6 Representation: The Neglected Architecture? . . . . . . . . . . . . . . 24

5 Conclusion 25

# 1 Introduction

The success of an explanation is primarily a function of choice and organization of content, and of its realization in media such as text and graphics. A variety of mechanisms for selecting and organizing the content of explanations have been utilized. This includes schema filling [McKeown 1985] and structure matching [Hovy 1988], graph traversal algorithms [Paris & McKeown 1986], and top-down goal expansion planning [Cawsey, 1989; Moore, 1989]. Yet, the literature lacks comparative analyses of the roles and limitations of these mechanisms for modeling explanation, apart from comments embedded in papers on other topics [e.g. Hovy 1990, Mooney *et al.* 1991, Moore 1989, Paris & McKeown 1986, Wolz 1990]

The purpose of this paper is twofold: to sketch such an analysis; and to offer a synthesis. The analysis of section 2 demonstrates that none of the above mechanisms taken alone can model certain explanation phenomena of interest (summarized in table 1 after the analysis). The task of choosing and organizing the content of an explanation involves associative, exploitive, opportunistic, prescriptive, and restrictive subtasks. Having distinct information processing requirements, these are best modeled by distinct computational mechanisms. Hence, a hybrid planning architecture that matches mechanisms to subtasks in a principled way most clearly reflects the nature of the explanation task. I illustrate this synthesis with the design of an implemented planner, presented in section 3. The architecture has implications for the debate over which relations best capture the structure of explanations for the purpose of generation. I make a distinction between intentional, pedagogical, conceptual, and epistemological relations, motivated by their different roles in the planning process.

My primary research focus is on choice of content from multiple perspective knowledge bases (those providing multiple models of a given phenomenon, using different predicates at different granularities and/or ontologies); inclusion of supporting material for pedagogical reasons; and organization of this content to facilitate the hearer's reconstruction [Suthers & Woolf 1990]. This aspect of my work is closest to that of Lester & Porter [1991], White & Frederiksen [1990] and Zukerman [1990, 1991]. A major design consideration in development of this architecture has been to provide a context for expressing a theory of explanation that makes the basis for such knowledge communication choices as explicit as possible. The hybrid architecture is more appropriate not only because it acknowledges the nature of the task in the mechanism, but also because it allows one to make explicit planning knowledge that remains hidden in plan operators for a uniform top-down architecture. In section 4, I propose elimination of preconditions and satellites from plan operators as part of the solution. This raises the nondeterministic aspect of planning above the level of selection between individual plan operators. The paper concludes with discussion of some tradeoffs, and limitations of the proposed architecture.

**Q1:** *How does a capacitor store charge?*
E1: A capacitor can be thought of as
E2: two flat metallic plates
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.

Figure 1: Example Explanation

## 2 Roles and Limitations of Planning Mechanisms

To initiate the analysis, I examine the roles and limitations of various mechanisms which have formed the basis for explanation planners. The analysis is not exhaustive, but provides the major motivations of the hybrid architecture described in section 3.

### 2.1 Imposition of Structure

Methods for structural imposition match abstract descriptions of the structure of explanations to a collection of propositions, the "relevant knowledge pool", in order to organize this material. This includes schema filling [McKeown 1985] and composition of rhetorical units [Hovy 1988]. The strength of such methods is the clarity with which they represent a descriptive theory of the structure of explanation, and their ability to impose organizations found to be effective in naturally occurring explanations.

Structure imposing mechanisms generally leave choice of the relevant knowledge pool to other mechanisms. This renders such methods inadequate for generating the pedagogical structure of certain explanations. The clearest example of this involves the **background** relation. Consider the explanation of Figure 1.[3] The process description of E5-E9 provides the informative response to the question. However, the

---

[3] The explanation is an edited version of a human-human protocol from sessions where participants communicated via pencil and paper. This medium better approximates human-computer interaction than spoken dialogues or textbook explanations, because it is interactive yet allows brief deliberation of responses, occurs in a nonvolatile physical medium, and allows both natural language and graphics.

explanation begins with a structural description in E1-E4. This functions as background for understanding the process description, by providing the context in which the process occurs and making terms such as "plate" meaningful. The background is *not* automatically part of the relevant knowledge pool for this question: E1-E4 might not be included if the interlocutor knew the basic structure of a capacitor. It is included *as background* by virtue of relationships between the interlocutor's assumed knowledge state and the content of the *previous* attempt at a relevant knowledge pool (initially, that chosen only for reasons of informativeness). Structurers that are ignorant of such dynamic relationships and see only the composite pool cannot identify the part that plays the role of background, so cannot account for placement of background before primary material.

Structure imposing techniques are inappropriate for exploiting existing structure as well [Paris & McKeown 1986]. This requires a different sort of mechanism, to which I turn.

## 2.2 Exploitation of Structure

Expert system rule traces [Davis, Buchanan, & Shortliffe 1977] and process strategies for object description [Paris & McKeown 1986, Paris 1987] exemplify graph traversal mechanisms for explanation. These selectively follow links in a knowledge base, with the path so traced out providing both the content and structure of the explanation. Hence, unlike structure matching, graph traversal methods need not assume a separate content selection mechanism.

Graph traversal mechanisms are appropriate for modeling how an explanation **exploits** the relational structure of knowledge. They implicitly embody the heuristic that it will be easier for the interlocutor to reconstruct a coherent model if each unit is introduced in relation to previous units. For example, parts of an object are introduced in relation to containing or adjacent parts, and process descriptions are organized to follow temporal and causal chains.

Graph traversal is limited to modeling local organization, or global organization that is a composite of local choices. Traversal methods don't naturally extend to global organization that 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 (as in Figure 1).

Lester & Porter [1991] use graph traversal to identify a relational network connecting target concepts to concepts believed to be known by the interlocutor. Their hybrid architecture is not subject to these limitations, as other processes perform initial content selection, editing and restructuring activities.

## 2.3 Prescriptive Planning and Opportunism

Top-down expansion of a discourse goal [Moore 1989, Moore & Paris 1991] is a stronger candidate for a uniform planning mechanism for explanation, integrating content selection and structuring. Hovy [1990] calls this **prescriptive** planning, as it handles goals that can be explicitly planned for, achieved by some act, and inactivated (though not forgotten). Like schematic approaches, top-down planning is a way of imposing structure (the difference between plan operators and schemata is a matter of degree). However, top-down planning records the intentional structure of the explanation at a finer granularity. As Moore [1989] has shown, the resulting record of discourse goals and of assumptions made in choosing between alternate strategies is required for certain functionality, such as recovery from failure to communicate and response to follow-up questions in context.

Cawsey [1989] handles the background problem with preconditions on plan operators. Planning for the precondition is invoked only if the specified background knowledge is not already understood (according to a user model). The planner knows what material was included as background, because it performed the selection. The example of Figure 1 is amenable to such treatment because a strategy for describing how a device works can state that the structure of a device must be understood first. For example, Cawsey's precondition for the goal to:

```
Know(Hearer, How-it-works(?device))
```

is:

```
Know(Hearer, Structure(?device))
```

However, one can only handle inclusion of background with expansion of abstract planning strategies if one can identify before-hand the kind of background knowledge required. This cannot always be guaranteed. Since explanation strategies are *abstract*, they don't specify exactly what content they will select. The retrieved content may contain unanticipated concepts that need to be explained as background. Hence, an **opportunistic** mechanism, such as data-driven "critics", is necessary in *any* computational model of explanation that plans with abstract strategies.

Various other functions require opportunistic mechanisms. Retrieved content must be examined in relation to a user model to edit material that has already been selected [Cohen & Jones 1987], and to prevent unintended false inferences [Zukerman 1990]. Independent expansion of multiple discourse goals to produce a single text may result in redundant or even conflicting content, necessitating filtering or merging of proposed content [Wolz 1990]. This is also an opportunistic task, as it can only proceed in response to redundancy observed in material proposed by other mechanisms. Finally, Mooney *et al.* [1991] show how extended explanations in certain domains exhibit content-determined organizational frameworks (e.g., time, space, events, factors) and consequent segmentation criteria (e.g., temporal, spatial, and causal event clusters) at the highest level of organization. They suggest using an opportunistic mechanism operating on initial "kernel" content specifications before applying top-down methods for selection and structuring of content within a segment.

## 2.4 Perspective: Associative and Restrictive Tasks

In many domains, explanations can differ on **perspective**, i.e., which properties and relationships are emphasized, and what concepts function as primitives, providing the basis for understanding the topic phenomenon [Suthers 1989]. Selection of an appropriate perspective 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 [McCoy 1989, McKeown *et al.* 1985, White & Frederiksen 1990]. I examine this complex phenomenon only to the extent needed to introduce two additional kinds of tasks.

In choosing perspective, an explainer strives for the (sometimes conflicting) desiderata of informative adequacy and comprehensibility [Suthers 1990]. Adequacy can be achieved in part by sensitivity to the interlocutor's goals [McKeown *et al.* 1985], as well as by fully utilizing what the application or knowledge base has to offer in response to the query. Comprehensibility is furthered in part by using concepts that are familiar to the interlocutor. The familiarity of a given concept may be evaluated based on its relations to concepts the interlocutor has used. For example, meaningful use of "electric potential" suggests that "force field" might also be understood, an inference one is less likely to make if the interlocutor never uses anything more specific than "electricity". This is an **associative** task, which could be handled by weighted links between concepts. My example deliberately suggests that certain semantic relations might mediate the association. McCoy [1989] provides an example of an associative influence on perspective for the sake of adequacy: responses to misconceptions emphasize attributes and objects associated with those mentioned in the previous dialogue. She uses collections of attribute weights representing domain-specific perspectives as filters when examining the user model and choosing possible responses.

Once chosen, the influence of perspective on content selection involves what Hovy [1990] calls **restrictive** planning. Perspective "goals", once generated, remain active throughout a span of discourse, and operate as *preferences* applied to choice points in content selection. Unlike prescriptive goals, such preferences are not necessarily satisfied after a single application, so cannot be inactivated once used, though they may change dynamically. The mechanism used for content retrieval should apply perspective preferences restrictively whenever a choice is available, and provide for conflict resolution, which Hovy points out is not well handled by prescriptive mechanisms.

| Phenomena | Task Type | Candidate Mechanisms |
|---|---|---|
| Providing requested information; Follow-up/recovery readdressing previous goal; Recurrence of strategies across situations and speakers. | Prescriptive | Top-down Goal Expansion with strategic operators; composition of Schemata (limited follow-up/recovery). |
| Impairment invalidation; Inclusion of background, illustrations, etc.; Omission of familiar material; Coordination of independently prescribed material; Content-based segmentation.[4] | Opportunistic | Data-driven Plan Critics for adding, removing, and reorganizing plan elements. |
| Ordering based on reconstruction gradients, including: Causal, Spatial, and Temporal organization of object and event descriptions; Background before foreground; Genus before differentia. | Exploitive | Graph Traversal (existing relations treated as constraints). |
| Influence of previous dialogue and user model on perspective and choice of terminology. | Associative | Weighted links, Perspective frames, Mediating abstractions. |
| | Restrictive | Filtering to satisfice Preferences. |

Table 1: Matching Mechanisms to Phenomena via Task Types

## 2.5 Summary of Analysis

Explanation involves a diversity of tasks with distinct information processing requirements, as shown in Table 1. A complete computational model of explanation will need to perform associative, exploitive, opportunistic, prescriptive, and restrictive tasks. Attempts to handle them with a single mechanism often resort to hiding mechanisms for the other task types in special purpose operators, resulting in inexplicit planning knowledge; and may result in an inelegant model due to theoretically irrelevant control operators that attempt to get the mechanism to handle tasks it is not well suited for. A uniform mechanism incorrectly implies that all information is equally available during each point of the planning process [McDonald & Pustejovsky 1985]. A principled match of computational mechanisms to each subtask of explanation planning provides a clearer expression of one's theory of what kind of process explanation is, and facilitates explicit representation of planning knowledge.

---

[4] The opportunistic aspect is choice of an organizational framework for segmentation. Ordering segments and their contents based on the chosen framework is Exploitive.

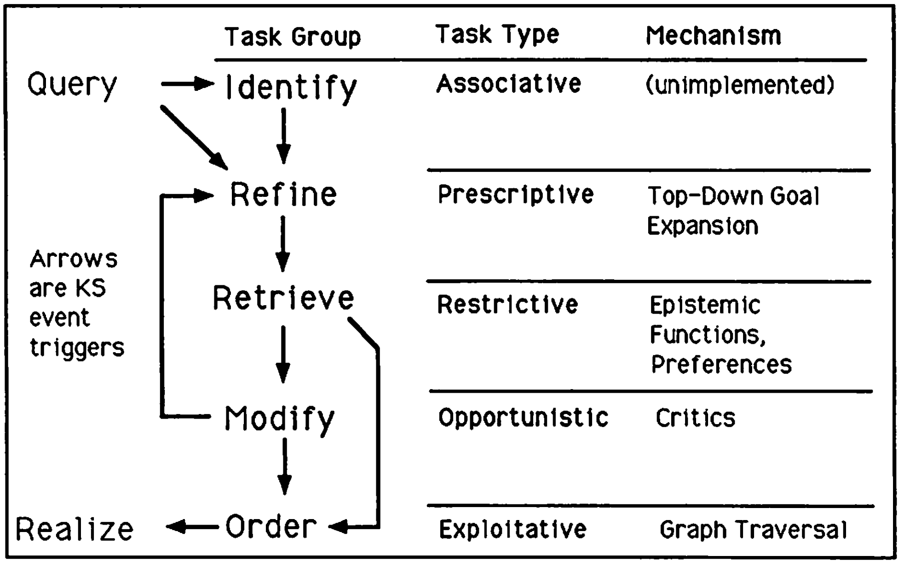

Figure 2: A Hybrid Architecture

## 3 A Hybrid Architecture

To provide a more concrete example of my architectural thesis, and to lead to further implications and conclusions, I describe the design of an implemented hybrid architecture. Bear in mind that the preceding analysis leaves open a variety of design choices, such as selection between candidate mechanisms and organization into modules. Hence I offer this particular design and implementation as an example and evidence of feasibility rather than as a necessity. Examples of other ways to modularize a hybrid architecture may be found in Lemaire [1991], Lester & Porter [1991], and (to a lesser extent due to lack of discussion of the top-down component) Mooney *et al.* [1991].

The intended application is an on-line educational resource which explains basic concepts in the physical sciences in response to user queries. Explanations are generated in a multimedia environment (a prototype is illustrated in figure 11) [Cornell *et al.* 1991]. Functional requirements include informativeness with respect to user queries; selection of perspective and simplifications and inclusion of prerequisite explanations based on evidence of user familiarity with the material; and handling of follow-up explanations. There is no need to plan extended explanations, as the dialogue is interactive and user controlled.

Figure 2 illustrates the architecture. Each task group is a collection of operators implemented in the indicated mechanism. The arrows indicate how the events generated by operators in one task group may invoke operators of another group. Operators from several task groups may be executing simultaneously, coordinated by an agenda mechanism and a common workspace, as in Nirenburg *et al.* [1989].

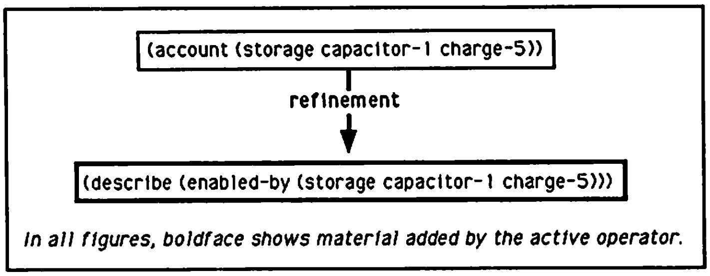

*In all figures, boldface shows material added by the active operator.*

Figure 3: Prescriptive Refinement of Example Query

A representation of the explanation that makes all theoretically motivated structure explicit is constructed in this workspace (an example will be given in figures 3-10).

## 3.1 Goal Identification

The initial *prescriptive* goal — which I call an **explanatory intention** — comes from interpretation of an interface event, such as a typed question, or menu selection applied to a previously displayed and mouse-selected object.[5] An explanatory intention is represented as an **explanatory illocution** (including Account, Define, Describe, Compare, and Illustrate) applied to one or more topics (atomic term, functional expression, or proposition).

Identification of *restrictive* perspective goals is under exploratory development at this writing. I am attempting to avoid weighted links or domain perspective frames, as there exists a potentially combinatorial number of associations between concepts, and hand encoded associations do not make much of a theoretical statement. Instead, I am investigating the use of abstractions of domain relations to mediate the association between concepts. In this approach, a concept will be considered to be familiar if it is associated with concepts the interlocutor has used via such abstractions. The corresponding restrictive goal is to prefer retrieval of material that minimizes the number of unfamiliar concepts used. If the approach succeeds, the associations will come "for free", without additional knowledge engineering, as it will be based on domain relations that need to be encoded anyway.

## 3.2 Prescriptive Refinement and Informativeness

Initial content selection is accomplished by a top-down goal expansion mechanism for refinement of explanatory intentions. The contents of the refinement operators

---

[5] I am using an expectation-based semantic parser, implemented by Keith Newstadt, which tolerates ungrammatical queries and can interpret a mixture of newly typed strings and references to previously displayed objects.

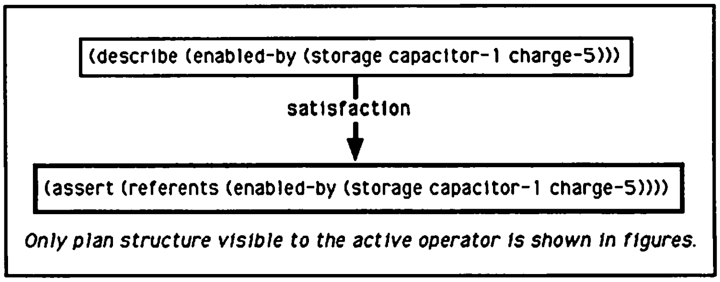

*Only plan structure visible to the active operator is shown in figures.*

Figure 4: Further Refinement to a Communication

are similar to those of Cawsey's [1989] content plans, but without preconditions. Some refinement operators say how a given explanatory illocution applied to a given topic can be fulfilled by application of another illocution to another topic related in a specified way. For example, one such operator says that one can account for something by describing what enables it:[6]

```
Account(?phenomenon) => Describe(enabled-by(?phenomenon)).
```

`Enabled-by` is a predefined abstract function for accessing domain-specific propositions. All references in plan operators to domain predicates, functions, and object classes are done via such abstract classes, which collectively constitute a nascent "upper model" [Bateman 1990] for explanation planning. As shown in figure 3, application of the above operator is the first refinement step taken in response to the query of Figure 1, which was represented as

```
Account(storage capacitor-1 charge-5).
```

Other operators take the final step, from an intention applied to a topic, to a speech act applied to a collection of propositions called a **view**. For example, if an intention to

```
Describe(structure(capacitor-1))
```

has been posted, then an act to

```
Assert(Referents(structure(capacitor-1)))
```

is set up, using the operator:[7]

```
Describe(?function(?object)) => Assert(Referents(?function(?object))).
```

`Assert` is a speech-act, and `Referents` is an "epistemic function", a device for referring to unretrieved content, discussed in section 3.3. Figure 4 illustrates a similar

---

[6] The notation of this and subsequent examples has been simplified for this paper, leaving out type specifiers, quantification, etc.

[7] `?Function` matches to any function in the *notational* sense regardless of whether the notational object has to do with purpose or functionality.

refinement of the derived informative goal of figure 3 to an intention to assert whatever enables a capacitor to store charge.

In addition to the refinement operators just described, there are currently operators embodying strategies for:

- Comparing, likening, and contrasting concepts, in a manner sensitive to whether or not the concepts are siblings or near cousins in the concept subsumption hierarchy.
- Defining terms by Aristotelian definition (genus and differentia), in terms of their definitional relationships to other concepts (see discussion of Skolem terms in section 3.4), or through description and illustration (below).
- Describing the structure of physical objects and the function and behavior of devices.
- Describing objects which have internal structure or behavior by drawing analogies. (Currently, the analogical mappings are stored as part of the domain knowledge, though a "structure mapping engine" [Falkenhainer *et al.* 1987] could be used without modifying the planner.)
- Illustrating class concepts by finding an example instance.
- Refining a generic intention to "explain" something, in a manner sensitive to its classification, to intentions that the above strategies can operate on.

Many other operators could be added by translating those of Cawsey [1989] and Moore [1989].

**Informative Satisfaction.** Refinement results in hierarchical dominance and satisfaction relations between explanatory intentions and communicative acts. This constitutes the **intentional structure** of the explanation [Grosz & Sidner 1986]. Creation of intentional structure requires examining unsatisfied goals, and posting subgoals or selecting content according to some theory of what it means to satisfy these goals. I maintain that this should be a theory that specifies, for any given explanatory intention, content that is **informative** with respect to that intention, in the sense that, if the intention is expressed in **content-question form** [Achinstein 1983, p. 49], the content answers the question. For example, a content-question form of the intention

```
(Account(storage capacitor-1 charge-5))
```

is

By what process does a capacitor store a charge?

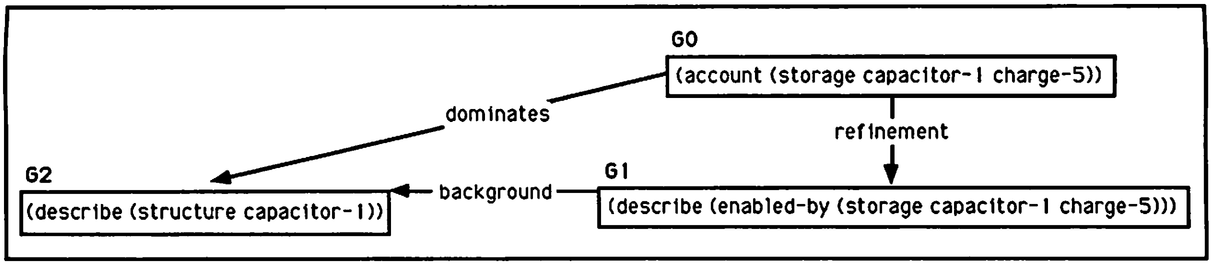

Figure 5: Refinement and Dominance

A content-giving sentence which answers this question is of the form:

> This is the process by which a capacitor stores a charge: *description of an appropriate process.*[8]

In order for an operator to specify content that satisfies the intention in the informative sense, it must select content which can be placed in the italicized position.

I further maintain that it is desirable to write refinement rules such that they express just the aspect of one's theory of explanation having to do with informative satisfaction, and nothing more. (One need not accept either of these positions to use the architecture.) In section 4.2, I will return to this point to explain why material included for other reasons (e.g., pedagogical augmentation) should not be selected by refinement operators. To preview, the appropriateness of such material is based on properties of the content selected to satisfy the intention, a condition which is best determined in a content-driven manner. The most we can *prescribe* in the form of refinement operators is that a given kind of question, by virtue of its *intrinsic* meaning, can be satisfied in the informative sense by a given kind of content. In summary, refinement is responsible *only* for intentional structure according to a theory of informative satisfaction.

**Dominance Relations.** The intentional hierarchy controls how perspective preferences and ordering constraints expressed at the granularity of intentions are inherited down the plan structure. Because of this, I have found it necessary to discriminate two variants of the dominance relation of Grosz & Sidner [1986]. My **refinement** and **satisfaction** relations (which differ from each other only in that the latter points to a communication object) have their semantics of dominance: hierarchical subordination plus informative satisfaction. However, a relation for inheritance of ordering constraints and preferences that does *not* imply satisfaction is also needed. Consider

---

[8] I use Achinstein's *this-is equivalent* of p. 36 rather than his content-giving form of p. 33 because the former is grammatical even when the inserted content can only be expressed as multiple sentences.

the case where a goal G2 is generated to provide Background for material selected by another goal Gl (figure 5). G2 does not satisfy Gl in the informative sense, so it would be wrong to install a Refinement relation between Gl and G2. The planner needs to know whether Gl has been refined into material that is informative respect to that goal. However, as an adjunct to Gl, G2 needs to inherit the same perspective and ordering constraints. The solution, shown in figure 10, is to install the (more restricted) **dominance** relation between Gl's dominator[9] G0 and G2, this relation allowing the inheritance of constraints without implying that G0 is satisfied by G2 in an informative sense.

## 3.3 Retrieval and Epistemic Functions

Epistemic functions may be thought of as abstract *knowledge patterns.* An expression such as

`Referents(structure(capacitor-1)),`

constructed with an epistemic function such as `Referents`, refers to certain subsets of the knowledge base that instantiate the abstract pattern. Epistemic functions serve two purposes: they are the means by which the planner refers to unretrieved content, and they are the prescriptive component of a retrieval specification. When implemented, perspective preferences inherited from the current context will provide the restrictive component.

The knowledge patterns retrieved by an epistemic function are defined *entirely* in terms of *epistemological level relations,* i.e. primitive relations required to structure concepts in a knowledge representation system [Brachman 1979]. Currently the relations used to define the patterns are:

- Structured inheritance between predicates, and between functions. This includes IS-A links between them, and role specialization links between their roles (figure 6).
- Specification of which function accesses propositions of a given predicate from fillers of a given role (figure 6).
- Subsumption relations between class concepts, and instantiation relations between class concepts and atomic terms. (I distinguish the subsumption predicate, which is used to make statements about domain concepts, from the representation system's IS-A, used to define predicates and functions.)

---

[9] The dominator of the top-level intention is a frame representing the dialogue context, not shown in the figure.

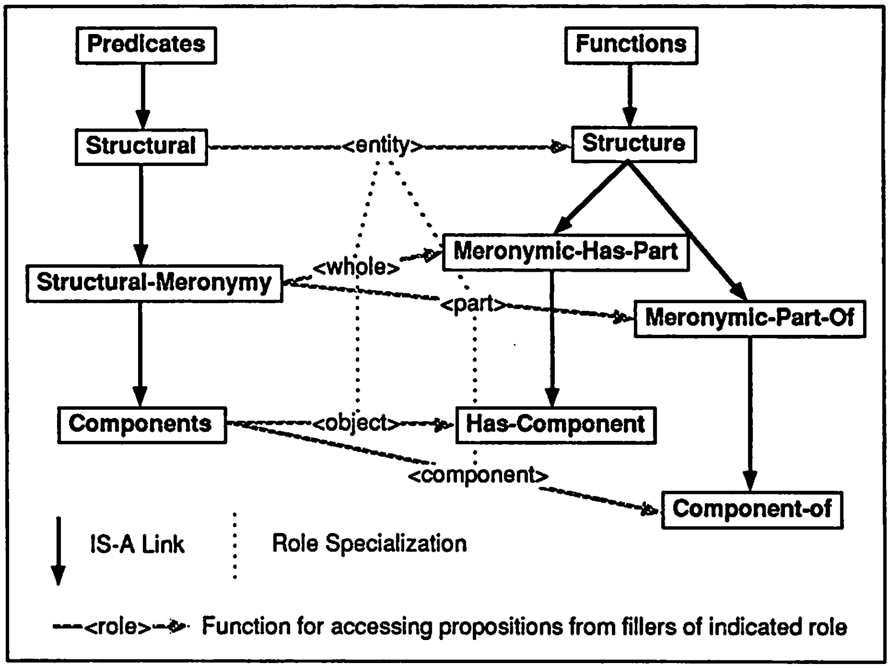

Figure 6: Definitional Relations between Predicates and Functions

- A Differentia relation which links a subsumption proposition to other propositions specifying what distinguishes the subclass from its siblings under the given superclass.

Because of this sole reliance on epistemological relations, refinement of explanatory intentions over domain concepts to speech acts over epistemic functions of those concepts constitutes a *reductionistic* step in the theory, from lesser-understood to better-understood entities. In contrast, "rhetorical predicates" [McKeown 1985] or "views" defined at the conceptual level [Suthers 1988, Souther *et al.* 1989] transfer the content selection problem to retrieval code without the theoretical commitment that reduction to a fixed set of epistemological patterns requires one to make.

The epistemic function **genus-and-differentia** is given one or more concepts, and returns a set of propositions, one identifying a common superconcept, and others differentiating the concepts from each other with respect to that superconcept. Wherever alternate propositions are found, perspective preferences apply. The epistemic function **referents** is given an ordinary functional expression, such as

`structure(capacitor-1),`

and returns a set of propositions providing the information those expressions refer to, such as

`(components capacitor-1 (set plate-1 plate-2)).`

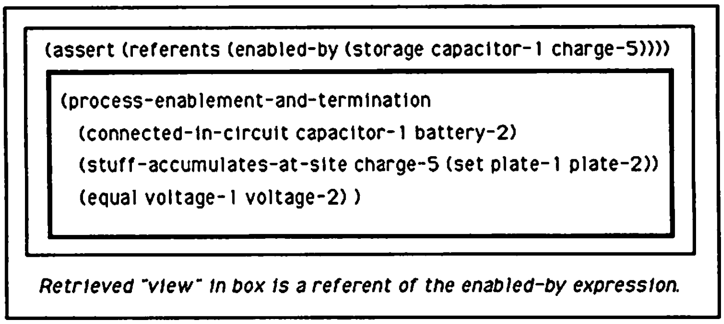

```
(assert (referents (enabled-by (storage capacitor-1 charge-5))))

(process-enablement-and-termination
  (connected-in-circuit capacitor-1 battery-2)
  (stuff-accumulates-at-site charge-5 (set plate-1 plate-2))
  (equal voltage-1 voltage-2) )
```

*Retrieved "view" in box is a referent of the enabled-by expression.*

Figure 7: Retrieval of Initial View

This relies on the *definitional* relations between an $N$-ary predicate (e.g. `components`) and an $N-1$-ary function that refers to the filler of the missing role (e.g. `has-component`, a specialization of `structure`), as shown in figure 6. The inner box of figure 7 illustrates the retrieved referent of

`(enabled-by (storage capacitor-1 charge-5)).`

`Enablement` is the predicate corresponding to `enabled-by`: in this case, the predicate `process-enablement-and-termination`, which specializes both `enablement` and `termination`, is found. There are currently four other epistemic functions: for finding an **instance** of a class concept, for finding **matches** and **mismatches** between the attributes of two concepts, and for finding **association** paths between two concepts using any specified relation class (e.g. `causality`).

## 3.4 Opportunistic Modification and Pedagogical Structure

Once a knowledge pool is available, data driven activities begin. Critics compare the knowledge pool and user model, **filtering** particular propositions likely to be familiar to the interlocutor, or **augmenting** the explanation by posting new intentions to illustrate abstract statements with examples or include background explanations.

**Representations and Reasoning for Posting a Background Goal.** For example, the background goal in figure 8 is posted as follows. An augmentation operator responds to the appearance of `plate-1` and `plate-2` by first examining internal evidence concerning whether the concepts are likely to be familiar to the interlocutor. These terms have not been used yet, and are not among the topics of active planning. The class concept `plate` is marked as a "common concept", normally assumed to be familiar. However, `plate-1` and `plate-2` are quantified with a Skolem function[10] of

---

[10] Skolem functions are used in standard logic to eliminate existential quantifiers in the scope of a universal, by rewriting $\forall(x)(\exists(y)P(x,y))$ as $\forall(x)P(x,f(x))$. That is, if the $y$ picked out depends

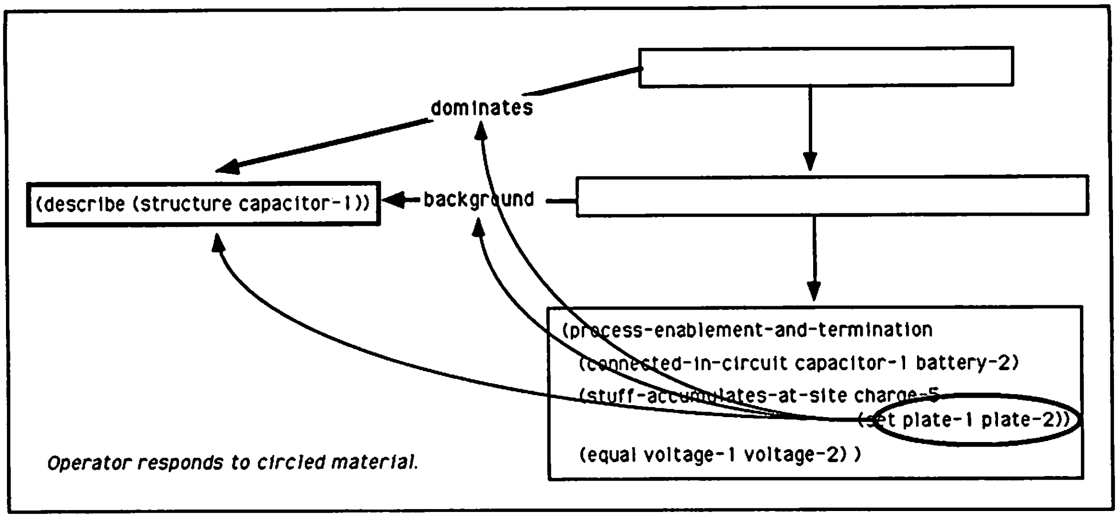

Figure 8: Opportunistic Posting of a Background Goal

`capacitor-1`, a situation that indicates an exception to the common concept rule, as the plates are dependent on a concept the interlocutor is *not* familiar with. The user is then asked (in a pop-up dialogue) whether she is familiar with "the plates of a capacitor", this expression being generated from the Skolem function. If the user responds "no", the augmentation operator must then construct and post an intention which provides the necessary background.

An intention to simply define `plate-1` and `plate-2` would be underconstrained, as we are interested in the plates in the context of the capacitor. The plates are Skolem of `capacitor-1` only because there is some relationship of interest between them. The planner needs this relationship to construct a more appropriate background intention. Hence, I require that the Skolem function be declared as part of the quantification. In this case, the function is `component-of`. This enables the planner to determine, via the function, the predicate that states the relevant relationship: that the plates are *components* of the capacitor (figure 6). Noting via `is-a` links that this relationship is a structural one, the augmentation operator can now post a background intention to describe the structure of `capacitor-1`. This last step, of abstracting to a high level predicate, allows the planner to access a complete model of the structure of a capacitor, shown in the collection of boxes in the left side of figure 9.

When other structural models of a capacitor are added to the knowledge base, I will need to ensure that the retrieved model includes the plates for whose sake the background explanation is being provided. To address this, I am experimenting with a version of the planner which does not take the last step of generalizing the predicates requested. Instead, it posts an intention to describe how the plates are associated

---

on which $x$ is picked out, then $y$ is a function of $x$. The $f$ is usually left unspecified. As discussed below, I need to know $f$.

with the capacitor via `component-of`. There is a cost: if additional descriptions of the composition of the plates, the insulator which is between them, etc. is desired, it must be planned for explicitly. I am currently examining heuristics for doing so. Further work is required to determine the best granularity at which to trade off explicit planning with pre-stored models.

**Planning Pedagogical Structure.** An explanation acquires pedagogical structure as material is included to facilitate the comprehension, reconstruction, acceptance, and retention of the material chosen for purposes of informative satisfaction. Planning this type of structure requires comparing proposed content to something functionally equivalent to a user model, and posting appropriate intentions to provide prerequisites, impairment invalidations, etc. Only the opportunistic task group need be responsible for this content-driven, context sensitive task. When augmentation occurs, an explicit record of the **pedagogical relation** (currently, Background and Illustration) between matched and requested material is created, important for ordering and realizing the explanation. I intend to add operators for other pedagogical augmentations based on Zukerman's [1990, 1991] work on Creative, Indicative, and Explanatory "rhetorical devices".

The architecture allows one to "compile" complex strategies into Refinement operators, e.g. to specify inclusion of Background and Illustrations prescriptively. (Initially, I modeled the inclusion of background in figure 8 with a precondition clause in a prescriptive operator, as Cawsey does.) However, for reasons to be discussed in section 4.2, I prefer a stricter division of responsibility between prescriptive and opportunistic operators, modeling inclusion of situation-dependent enabling and impairment invalidation material *uniformly* with situation-sensitive operators.

**Granularity of Background Selection.** Lester & Porter [1991] select background material by searching the knowledge base to find "integrative connections" from the target material as a whole to a user model, both being represented as a subgraph of the knowledge base. A more local background operator, such as that described above, should be able to select similar material by a form of backchaining in which the background operator is re-applied to material that was retrieved as background. The recursion ends on concepts that the planner decides is familiar to the interlocutor, possibly by asking her. Some tradeoffs are involved. A single search for integrative content may be more efficient, and allows search strategies to specify non-local constraints on the form the integrating content may take. On the other hand, finer granularity selection of background allows content selection to be more sensitive to the dialogue context, specifies why the content was included at a finer granularity, does not require that the model of the interlocutor's knowledge and beliefs be represented as an overlay to the knowledge base, and is less vulnerable to lack of information in this model.

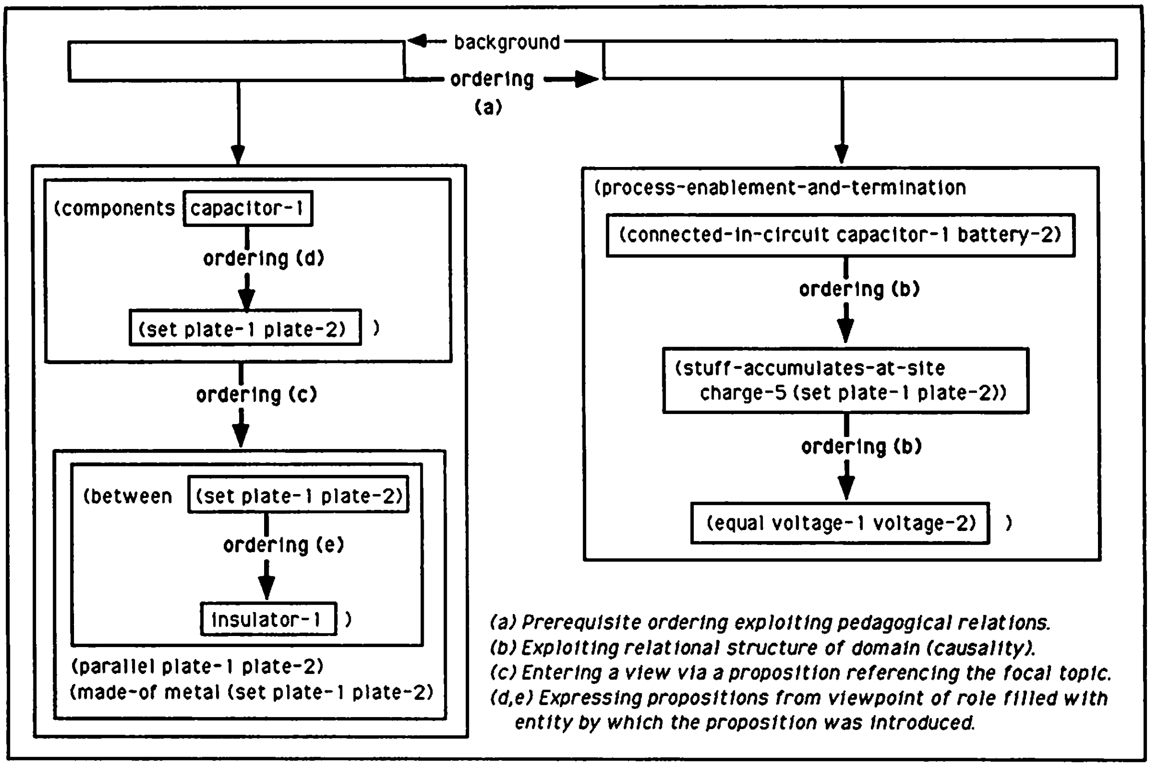

*(a) Prerequisite ordering exploiting pedagogical relations.*
*(b) Exploiting relational structure of domain (causality).*
*(c) Entering a view via a proposition referencing the focal topic.*
*(d,e) Expressing propositions from viewpoint of role filled with entity by which the proposition was introduced.*

Figure 9: Exploitive Installation of Ordering

## 3.5 Exploitive Ordering and Epistemological Gradients

Ordering operators respond to the installation of pedagogical, conceptual, and epistemological relations, and annotate the knowledge pool with a partial ordering. The ordering attempts to follow what I call **epistemological gradients**: the easiest ways a hearer can reconstruct the communicated concepts and propositions, and integrate it into her own knowledge [Goldstein 1979, Hewson 1981, Rissland 1978]. Lester & Porter's [1991] "integrative content selection" applies this intuition to identification of background and related knowledge.

**Heuristics for Derivation of Ordering.** Ordering annotations are at three granularities: between intentions (figure 9a), between propositions (figure 9c), and between roles of a proposition (figure 9b,d,e). Several types of relations installed by other processes are exploited:

- Two operators order based on the pedagogical relations Background and Illustration (9a). The ability to identify the pedagogical role of prerequisite or corequisite material independently of ordering decisions is an advantage over using an undifferentiated prerequisite relation.

- Some operators exploit domain relations: temporal and causal chains are followed in the forward direction (9b), and parts of an object are introduced in relation to other parts already mentioned.
- Another operator exploits the epistemological relation between a genus proposition and its differentia, introducing the genus first, as this provides the context in which the differentia is meaningful.
- Ordering is also derived from focus of attention: in each view, propositions involving the view's topic concepts (arguments of the epistemic function) are placed before other propositions, so the new content of a view is introduced via the focal concept (9c, where `capacitor-1` is the topic).
- Finally, propositions are expressed from the viewpoint of the role whose filler contains a topic concept (9d), or a concept by which the proposition was reached during the search that installs viewpoint markers (9e), providing smoother transitions between propositions. For example, this enables the generator to produce "A capacitor has two plates with an insulator between them" (merging two propositions by adjacent identical role fillers), rather than the more awkward "A capacitor has two plates. An insulator is between the plates."

**Ordering Relations.** I have found it necessary to use two ordering relations at the inter-propositional granularities, as well as the intra-propositional viewpoint annotation. **Prerequisite** ordering indicates material that must occur *sometime* before a given unit, and is used for Background and Illustration. **Juxtaposition** ordering, used for placement of Genus before Differentia, orders units in a *contiguous* sequence.

Grosz & Sidner present their ordering relation (satisfaction-precedence) as an intentional-level structure, albeit derived from domain-dependent considerations. While my ordering relations are also derived from other relations, it is best classified as pedagogical, as the derivation has strong pedagogical motivations and minimal implications for informativeness per se.

## 3.6 Realization

Currently I am using a simple template-based text generator[11] in an object-oriented, direct manipulation multimedia interface.[12] The screen dump of figure 11 shows one text the generator produces from the plan of figure 10, as well as the graphic media objects that were retrieved to support the explanation. While these components are not central to the architectural thesis of this paper, I discuss a few of their features for completeness.

---

[11] Modified from one initially developed by Bob Crites.

[12] Designed and implemented by Matthew Cornell.

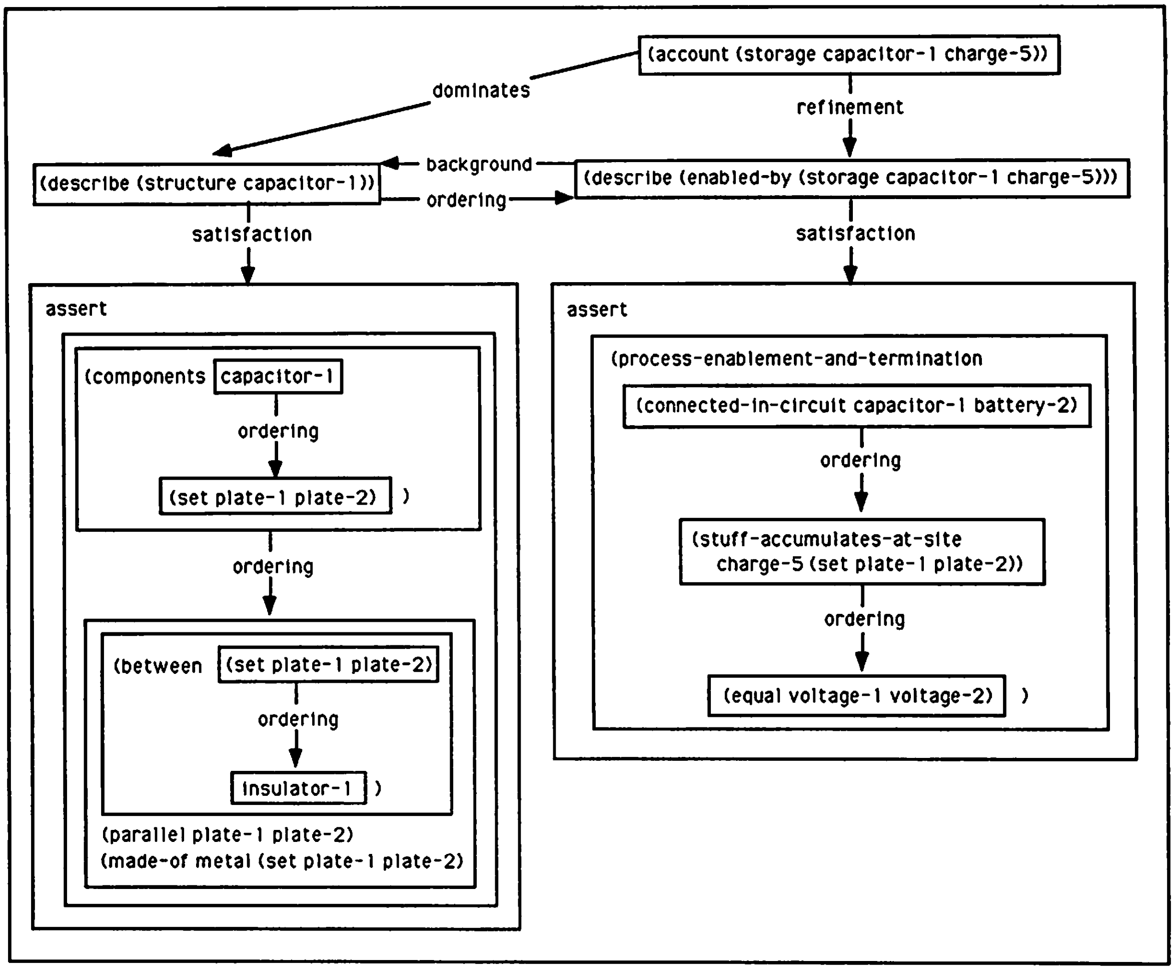

Figure 10: Resulting Plan for Example

**Text Generation.** Templates consist of a sequence of strings, role names, which are replaced with text generated recursively for the role fillers, and "variant lexemes", which are realized differently depending on characteristics of specified roles, such as plurality. Templates for active and passive sentences are indexed by the role which is active, and are chosen according to the viewpoint specified by intra-propositional ordering. When multiple templates are indexed, a random choice is made. The generator merges templates both sequentially and recursively to produce complex sentences. For example, the first sentence of figure 11 required three sequential merges, and the second three recursive merges. Sequential merging is allowed to violate prerequisite ordering, but not juxtaposition ordering, *within* a sentence to attach adjectives. Rhetorical transitions such as "however", "to provide some background", etc. are generated based on the pedagogical relations in the plan.

The generator is relatively fast, but is limited in several ways. The most serious limitation occurs when templates are merged: the generator is currently unable to

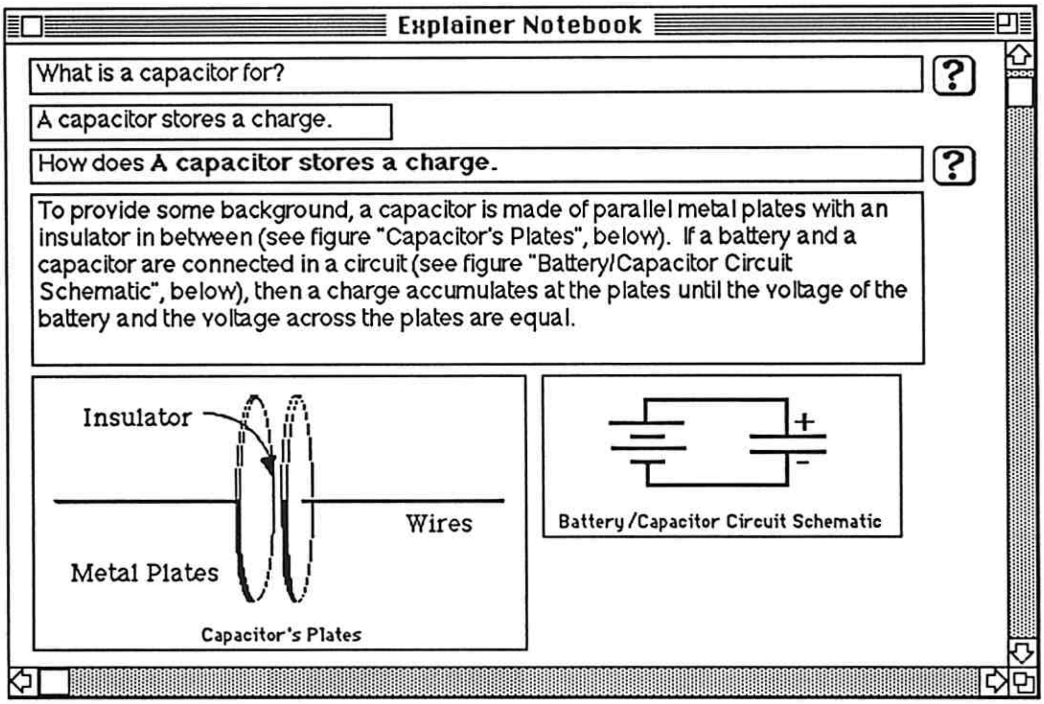

What is a capacitor for?

A capacitor stores a charge.

How does **A capacitor stores a charge.**

To provide some background, a capacitor is made of parallel metal plates with an insulator in between (see figure "Capacitor's Plates", below). If a battery and a capacitor are connected in a circuit (see figure "Battery/Capacitor Circuit Schematic", below), then a charge accumulates at the plates until the voltage of the battery and the voltage across the plates are equal.

Figure 11: Resulting Explanation in Prototype Multimedia Environment

ensure grammatical consistency such as in verb tense. Templates are currently written in the present perfect in an attempt to temporarily sidestep this problem.

**"Live Information" Interface.** The explanation is displayed in an interface in which displayed information retains its underlying semantics and functionality [Cornell *et al.* 1991]. In particular, all subsequences of the generated text that correspond to an underlying semantic object (proposition, term, class concept) retain pointers to the corresponding object, and can be mouse-selected to obtain a menu of questions, to reference the object in follow-up explanations, or for incorporation into a user-authored document. For example, the boldface text in figure 11 resulted when the previous response was dragged into the query box. This enables the explainer to directly access the intended topic of the query, by using the pointer to the underlying object rather than reparsing its string representation. I also use mouse selection to track focus of attention without complex focus heuristics: when the user selects an object, her focus is assumed to be on the dialogue context which encloses it. However, the current planner does not yet make use of the resulting ordering of focal contexts.

The explainer is being ported to a new version of this interface that enables users to pick up objects displayed in one way (e.g., as text or pictures) and drop them into displayers for other media modes to obtain an alternate view on them. This supports a "get it close and allow refinement" approach to multimedia presentation planning:

the explainer need only plan a reasonable initial presentation, which the user may then modify.

# 4 Discussion

I present several specific implications of the architecture for structural abstractions, plan operators, and nondeterminism; and conclude with discussion of tradeoffs that have been made, and limitations of the architecture.

## 4.1 Structural Abstractions for Generation

A number of research efforts emphasize rhetorical relations for the analysis of natural explanations and for expressing a generative theory of explanation [Hovy 1988, Mann & Thompson 1986, Maybury 1988, McKeown 1985, Moore 1989]. Rhetorical relations were an important development in explanation research, since they provided 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. However, because they were developed to describe *text*, i.e. the end product of some explanation generation process, rhetorical relations conflate the results of a variety of distinct knowledge sources operating on different levels of information. For example, an examination of rhetorical relations (such as in [McKeown 1985] and [Mann & Thompson 1986]) reveals that some are based on the structure of domain objects and processes (e.g. Constituency and Causality), some are derived relations between concepts that vary according to context (Comparison), and others are due to pedagogical considerations such as what the interlocutor knows and needs to know to better grasp a point (Amplification, Background, Evidence, and Illustration). A theory of explanation *generation* needs a vocabulary that makes distinctions relevant for describing the generation process, not just the end product.

My strategy has been to factor rhetorical structure into classes of relations that are derived from distinct knowledge sources. The foregoing task analysis and architectural design demonstrated that different subtasks are responsible for handling intentional, pedagogical, conceptual, and epistemological relations, and that they do so based on different components of a theory of explanation. The architecture provides a straightforward correspondence between the theoretical basis for each type of explanatory structure and the type of operator responsible for that structure, as summarized in table 2.

| Task Group | Mechanism | Structure Provided | Based on Theory of |
|---|---|---|---|
| Refinement | Prescriptive | Intentional | Informative Satisfaction |
| Retrieval | Restrictive | Conceptual | Application Domain |
| | | Epistemological | Knowledge Representation |
| Modification | Opportunistic | Pedagogical | Comprehension, Learning |
| | | Rhetorical | Persuasion |
| Ordering | Exploitive | Sequential | Reconstruction Gradients |

Table 2: Correspondence: Operators, Structure, and Theories

## 4.2 Planning Without Preconditions

All the major explanation planners based on a goal-refinement paradigm use either preconditions or satellites in their plan operators. These constructs specify how one might need to augment the explanation with additional material, in effect turning goal refinement operators into mini-schemata. The planning knowledge implicit in preconditions and satellites is most often knowledge of how the success of an explanatory intention depends on the interlocutor's knowledge or beliefs. Extra material is included only if some condition does not obtain, such as presence of the prerequisite concept in the user model. Inclusion of such material is *not* based on intrinsic properties of the primary intention. Rather, it is based on relationships between the *content* used to satisfy the intention and a model of the interlocutor's knowledge state. Therefore, knowledge of how to plan supporting material is knowledge about how to perform a context sensitive, data-driven task, best expressed as augmentation operators within an opportunistic mechanism rather than as preconditions on prescriptive operators.

The choice is not just a matter of design aesthetics: the functionality of the planner is also at stake. In a planner for explanation under multiple perspectives, one is unable to predict the content that will be used to satisfy any given intention, as choices are made during a given retrieval, not just in strategic refinement. When writing refinement operators, one doesn't know enough about what assumptions the model accessed will make about hearer knowledge to fully predict the appropriate auxiliary material. Hence, augmentation planning knowledge written as satellites or preconditions on strategic operators must either rely on contextual assumptions, sacrificing flexibility, or hide this planning knowledge in the interpreter.

For example, Cawsey can prescribe structural knowledge as a precondition for an account of how a device works (page 4) only because all her strategies for planning such accounts select content which *refer to* the device's structure. Inclusion of such background would become inappropriate if strategies and device models were added that gave the planner the option of accounting for a device's function with abstract equations, or by giving an analogy not involving structure. A planner that selects background using content-driven operators can generate the same prerequisites as Cawsey's operator when references to the device's structure are noticed, yet would not need to be modified when alternate strategies or process models were added.

When augmentation knowledge is made entirely explicit in opportunistic operators, top-down refinement operators are only responsible for prescribing which content is informatively relevant to an explanatory intention. They require no preconditions to do so, only applicability constraints. As a result, preconditionless operators can be written more abstractly. For example, Cawsey needs an operator that is restricted to intentions to account for how a device works, in order to write the structural precondition (page 4). When I removed this precondition from my version of this operator, I was able to abstract it to apply to intentions to account for any phenomenon (page 9).

## 4.3 Nondeterminism and Revision

Use of finer-grained, content prescribing refinement operators changes the nature of nondeterminism in planning. Conventional refinement operators are correctly treated as competitors for expansion of a given goal: as mini-schemata, they embody alternate, prepackaged ways of dealing with the goal. On the other hand, preconditionless informative operators for a given intention can be written to propose relatively disjoint kinds of information relevant to the intention. They are no longer clear competitors with each other. Hence, selection heuristics are not needed at the level of choosing one out of several refinement operators that match an intention. Instead, the problem is one of deciding which *subset* of expansions to allow, along with associated augmenting material.

The current planner takes an inclusive and least-commitment approach. All expansions are allowed to proceed. When an expansion fails to result in retrieved content (whether due to lack of an applicable plan operator or insufficient material in the knowledge base), failure is propagated up the refinement hierarchy, resulting in self-removal of the subtree and detachment of any pedagogical links involving that subtree. All successful expansions are currently included, regardless of what other expansions succeeded. As the size of my plan library and knowledge base grows, an inclusive least-commitment approach will be inadequate. I will need to add **revision** operators to the Modification task group for eliminating expansions based on *global* and *comparative* criteria such as redundancy, verbosity and the need for excessive supporting explanations [Lemaire & Safar 1991].

## 4.4 Tradeoffs

Theoretical clarity and flexible generation of explanations are compatible design goals: a theory that makes the basis for all decisions explicit is more likely to be implementable in a planner that tailors its explanations for the characteristics of the dialogue context. However, there is a tradeoff between these goals and others.

Simplicity has been sacrificed for the sake of diverse operator types and mechanisms. Acquisition of planning knowledge may be more difficult as a result; on the other hand, this should be compensated for by the greater modularity of the operators. For example, to improve the strategy for inclusion of Background, one need change only one or a few designated operators, rather than all refinement operators having preconditions.

Efficiency is lost whenever decisions that were previously made when writing schematic operators are made at runtime instead. The explanation of figure 11 required only several seconds on a stock Macintosh II, including question parsing and text generation. However, further increases in the number of plan operators and size of the knowledge base, and anticipated addition of inferences about the user [Zukerman 1990, 1991] and plan editing activities [Lester & Porter 1991, Lemaire & Safar 1991] will likely lead to less tolerable response times. Schematic plan operators (at whatever granularity) still have the advantage of being more efficient in their execution than approaches that derive content and structure dynamically.

## 4.5 Limitations of the Architecture

In the current architecture, the structure of extended dialogue can only emerge from the sensitivity of refinement and augmentation operators to the dialogue history. There is no higher-level knowledge source for curriculum planning or generating goal-oriented behavior, as in Cawsey [1991]. The architecture cannot model extended monologues that exhibit high-level content-derived organization, as in Mooney *et al.* [1991], because data driven operators operate only on the content of views proposed by informative refinement operators. It may be possible to add a curriculum knowledge source that affects refinement and modification choices via the preference inheritance mechanism. High-level content-derived organization may require more substantial changes to the architecture. However, generation of extended monologues is not required in the intended application.

## 4.6 Representation: The Neglected Architecture?

A limitation of the current implementation rather than the architecture is the small size of the knowledge base. The usual knowledge acquisition bottleneck has been exacerbated by repeated changes in representational requirements arising as I encountered a variety of demands that detailed run-time choice of content places on representation. Examples include classification of domain predicates into abstract classes; declaration of the relations between predicates and functions, and of role-specialization relations between predicates; and declaration of the functional basis for Skolem quantifications.

This brings up design of knowledge representation as an important aspect of architecture. The discrepancy between the large amount of my own time representational issues have consumed and the lack of discussion in the recent explanation planning literature has me concerned whether work on these issues within our community has declined since Clancey [1983] and Swartout [1983] provided us with some important insights. (Significant exceptions include Bateman [1990], Falkenhainer & Forbus [1991], and Neches *et al.* [1985].) I believe the community will need to determine which interactions between representation and functionality are due to essential requirements that explanation places on the knowledge base in order to support fully flexible pedagogical explanations under multiple perspectives.

# 5 Conclusion

I am primarily committed to the conclusion that explanation involves subtasks that are best modeled by distinct mechanisms, and hence that principled design of a hybrid planning architecture will most clearly reflect the nature of the explanation task, and allow us to make planning knowledge explicit for more flexible run-time choice and organization of content. The paper will have served its purpose if this has been made clear, even if you do not accept my other conclusions, such as the breakdown into associative, exploitive, opportunistic, prescriptive, and restrictive subtasks, with its concomitant elimination of preconditions on refinement operators; or my argument that structural relations should be distinguished according to their intentional (informative), pedagogical, conceptual, and epistemological sources.

# Acknowledgments

I am grateful to my advisors, Edwina Rissland and Beverly Woolf, and committee members Victor Lesser and Klaus Schultz, for helpful discussions; to Klaus Schultz for his expert explanations; to Matthew Cornell, Bob Crites, Sara Betz, and Keith Newstadt for providing assistance with the research; and to Johanna Moore for asking the right questions. Thanks also to Johanna and Alison Cawsey for understanding that when someone stands on your shoulders, it can appear a bit abusive. Development was facilitated by use of CMU's FrameKit, provided by Eric Nyberg, and especially by the generous support of Beverly Woolf and Apple Computer.

# References

**Achinstein, P. (1983).** *The Nature of Explanation.* New York: Oxford University Press.

**Bateman, J. (1990).** Upper modeling: organizing knowledge for natural language processing. *Fifth Int. Workshop on Natural Language Generation,* June 3-6, Dawson, Pennsylvania.

**Brachman, R. (1979).** On the epistemological status of semantic networks. In *Associative Networks: Representation and Use of Knowledge by Computers,* pp. 3-60, N. V. Findler, ed., New York: Academic Press.

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

**Cawsey, A. (1991).** Generating Interactive Explanations. To appear in *Proc. AAAI-91* .

**Clancey, W. J. (1983).** The epistemology of a rule-based expert system – a framework for explanation. *Artificial Intelligence,* vol. 20, no. 3, pp. 215-251.

**Cohen, R. & Jones, M. (1987).** Incorporating user models into an expert system for educational diagnosis. In *User Models in Dialog Systems.* W. Wahlster & A. Kobsa eds., Springer-Verlag.

**Cornell, M., Suthers, D. & Woolf, B. (1991).** Using "live information" in a multimedia framework. AAAI-91 Workshop on Intelligent Multimedia Interfaces.

**Falkenhainer, B. & Forbus, K. (1991).** Compositional Modeling: Finding the right model for the job. To be published in *Artificial Intelligence,* , special issue on qualitative physics, Fall 1991.

**Falkenhainer, B., Forbus, K., & Gentner, D. (1987).** The Structure-mapping Engine: Algorithm and Examples. Report No. UIUCDCS-R-87-1361, Department of Computer Science, Unviersity of Illinois at Urbana-Champaign.

**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.

**Lemaire, B. & Safar, B. (1991).** Some necessary features for explanation planning architectures: study and proposal. *AAAI-91 Workshop on Comparative Analysis of Explanation Planning Architectures, July 14, Anaheim, CA* (this volume).

**Lester, J., & Porter, B. (1991).** An architecture for planning multi-paragraph pedagogical explanations. *AAAI-91 Workshop on Comparative Analysis of Explanation Planning Architectures, July 14, Anaheim, CA* (this volume).

**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.

**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.

**Mooney, D., Carberry, S., & McCoy, K. (1991).** The need for bottom-up strategies for generating high-level structure. *AAAI-91 Workshop on Comparative Analysis of Explanation Planning Architectures, July 14, Anaheim, CA* (this volume).

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

**Moore, J. & Paris, C. (1991).** The EES explanation facility: its tasks and its architecture. *AAAI-91 Workshop on Comparative Analysis of Explanation Planning Architectures, July 14, Anaheim, CA* (this volume).

**Neches, R., Swartout, W. R., & Moore, J. D. (1985a).** Enhanced maintenance and explanation of expert systems through explicit models of their development; *IEEE Transactions on Software Engineering,* vol. SE-11, no. 11, pp. 1337-1351.

**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. (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. & McKeown, K. (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.

**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.

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

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

**Suthers, D. (1990).** The Epistemological Structure of Explanation. *AAAI Workshop on Explanation Working Notes,* July 30, 1990, Boston.

**Suthers, D., & Woolf, B. (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, University of Massachusetts, Amherst.

**Swartout, W. R. (1983).** XPLAIN: A system for creating and explaining expert consulting programs. *Artificial Intelligence,* vol. 21, no. 3, pp. 285-325.

**White, B. & Frederiksen, J. (1990).** Causal model progressions as a foundation for intelligent learning environments. *Artificial Intelligence,* vol. 42, no. 1, pp. 99-157.

**Wolz, U. (1990).** An object oriented approach to content planning for text generation. *Fifth Int. Workshop on Natural Language Generation,* June 3-6, Dawson, Pennsylvania.

**Zukerman, I. (1990).** Generating peripheral rhetorical devices by consulting a user model. *Fifth Int. Workshop on Natural Language Generation,* June 3-6, Dawson, Pennsylvania.

**Zukerman, I. (1991).** Failure prevention: an alternative approach to content planning. *AAAI-91 Workshop on Comparative Analysis of Explanation Planning Architectures, July 14, Anaheim, CA* (this volume).