

# Preferences for Model Selection in Explanation\*

Dan Suthers

Learning Research and Development Center  
3939 O'Hara Street  
University of Pittsburgh  
Pittsburgh, PA 15260  
suthers+@pitt.edu

## Abstract

The range of possible domain models on which an explanation can be based is often large, yet human explainers are able to choose models that address a questioner's informative needs without undue obscurity. However, few existing explanation systems use knowledge bases providing multiple models of their topics of explanation let alone account for the selection of a model for a given explanation. This paper demonstrates the utility of a preference-based mechanism for model selection. Selection heuristics are made explicit as preferences and can be added, modified, or removed without modifying the interpreter or other planning knowledge. The mechanism is more general than previous mechanisms for model selection or "perspective," and can apply previously identified model selection criteria.

## 1. Introduction

Consider a physics instructor describing the relationship between electric potential and current, answering a question about the charge-storage functionality of a capacitor, or explaining how a given circuit works. These phenomena can be modeled with mathematical descriptions of their behavior, with causal relations between underlying physical events, or by analogy to mechanical systems involving "flowing waters or teeming crowds" [Gentner and Gentner, 1983]. Idealized physical models can be used that ignore the resistance of a wire or inductive effects, or these factors can be taken into account. The instructor can identify the existence of a parameter such as voltage, point out the direction of change of that parameter, or give it a numeric value. In general, because the instructor knows a great deal about electricity and electrical devices, the range of possible models on which an explanation can be based is large. Yet the instructor is able to base the explanation on models that address the informative needs of the student without undue obscurity.

The success of an explanation can depend crucially on the domain model on which it is based [Paris, 1985; Stevens

and Collins, 1980; White and Frederiksen, 1990]. However, automated explanation facilities do not share our instructor's ability to select from among many domain models. Few existing explanation systems use knowledge bases providing multiple models of their topics of explanation, let alone account for the selection of a model for a given explanation. Notable exceptions include McCoy [1989], McKeown *et al.* [1985], and Paris [1987], who used mechanisms suited to their particular problems, mechanisms that do not generalize to all aspects of model selection. White and Frederiksen [1990] showed how to design a curriculum of models that supports a gradual transition from novice to expert in an instructional system, but did not automate model choice.

In this paper, the utility of a preference-based mechanism for model selection is demonstrated. Selection heuristics are made explicit as preferences, and can be added, modified, or removed without modifying the interpreter or other planning knowledge. Examples of preferences (discussed in detail later) include:

- say something you haven't said before,
- elaborate on previously discussed models,
- introduce curriculum topics when possible,
- minimize concepts likely to be unfamiliar to the questioner, and
- minimize the amount of information to be processed.

The independent and explicit treatment of model choice preferences enables model selection on more dimensions than have been previously attempted in automated explanation facilities. This in turn enables the use of knowledge bases that make available a greater diversity of models for each topic, increasing the range of users for which explanation facilities will be appropriate. The approach discussed in this paper has been implemented in an explanation system [Suthers *et al.*, 1992; Suthers, 1993].

We begin with definitions of "multiple models" and related terms. Then the control and information processing requirements of model selection are analyzed to motivate the choice of a preference-based mechanism for performing this selection. We describe the mechanism, list some preferences, and give an example of their application. The paper concludes with comments on related research.

## 2. Multiple Models, Perspectives, and Views

For our purposes, a knowledge base (or any resource from which domain models can be obtained) contains "multiple

\* Research conducted at the University of Massachusetts while supported by NSF grants MDR-8751362 and IRI-8908481 and by Apple Computer, and at the University of Pittsburgh while supported by NSF grant MDR-9155715.

models" if, for a non-negligible subset of its topics, it contains domain models that *compete* as the basis for constructing a response to a question about a topic. Our notion of "multiple models" is intended to include variation along the following (nonorthogonal) dimensions:

**Predication:** The category of predicates, relationships, or attributes used, for example, causal, evidential, functional, and structural.

**Granularity:** Level in a hierarchy of modeled entities, for example, predication over macroscopic, microscopic, atomic, or subatomic entities [Hobbs, 1985; Falkenhainer and Forbus, 1991].

**Ontology:** The type of primitive objects that the model is based on. For example, a physicist can model electromagnetic phenomena as particles or waves, and engineers can choose between Lagrangian and Eulerian methods of analysis [Falkenhainer and Forbus, 1991].

**Register:** The areas of knowledge or social communities in which the concept is primarily used [Bateman and Paris 1989]<sup>1</sup> For example, certain concepts are used primarily in mathematics, physics, or chemistry, while others pervade everyday life.

**Elaboration:** The number of factors or components taken into account by the model [White and Frederiksen, 1990].

**Order:** Whether parameters are characterized in terms of existence, direction of change, or quantitative values [White and Frederiksen, 1990].

Other taxonomies and examples of multiple models are provided by Souther *et al.* [1989], Stevens and Collins [1980], Stevens and Steinberg [1981], and White and Frederiksen [1990].

We use the term "perspective" to refer to a choice of values on the above dimensions independently of the topic it is applied to. A perspective is analogous to a "filter" through which a topic is viewed. A domain model is an application of a perspective to a topic. Previous writers have used the term "perspective" for a similar notion, although sometimes limited to mean choice of what subset of predicates, relationships, or attributes to consider [McKeown *et al.*, 1985; Falkenhainer and Forbus, 1991; McCoy, 1989; White and Frederiksen, 1990].

An explanation act often expresses only a fraction of what is properly called a complete model of a phenomenon. Indeed, an entire dialogue may be required to communicate a single domain model. We need to distinguish the portion of a model that is communicated in a given explanation from the model as a whole. The term "view" is used to refer to a portion of a domain model that is a candidate for communication in an explanation [Suthers, 1993]. In particular, a view is that portion of a domain model that is relevant to the rhetorical goal of an explanation in the sense that *communication of the contents of the view* (under the appropriate speech act) *satisfies the informative*

*requirements of the rhetorical goal* [Suthers, 1993]. A view can consist of one or more propositions.

### 3. The View Selection Task

When an explainer is asked a question about a topic, the explainer must choose a view that is appropriate for answering the question in the discourse situation. Appropriateness is a function of both the question and the context. We discuss why the contextual nature of view selection demands a distinct kind of mechanism.

#### 3.1 Goal Refinement: Helpful but Insufficient

It is often possible to specify the types of predicates that are potentially relevant to answering a given type of question based solely on the intrinsic meaning of the question. For example, structural and functional predicates are potentially relevant to a request for an object description and causal predicates are potentially relevant to a request for an account of why an event occurred. When a knowledge base contains only a single model for each topic within a given predicate class (e.g., only one structural decomposition, one causal etiology, etc.), then one can select content by considering the abstract rhetorical and semantic form of the goal alone (e.g., Describe(object) => structural predicates, Account-for(event) => causal predicates, etc.). For this reason, the goal refinement operators of various explanation planners specify the class of predicates to be used in formulating a segment of the response [Cawsey, 1989; Maybury, 1990; Moore, 1989; Suthers, 1991]. However, goal refinement can only specify relevant content at the level of abstraction at which goal refinement operators are written. When multiple models in the general sense outlined above are available, choices can remain within a specified predicate class as well as on the dimensions of granularity, ontology, register, elaboration and order. Additional constraints are required to choose an appropriate view. These constraints come from the interaction between the *content* of the proposed views and the *epistemic context* in which the explanation is taking place, namely:

- the organization of the explainer's domain knowledge [Rissland 1978],
- the explainer's beliefs about the questioner's domain knowledge [McCoy, 1989; Paris, 1987],
- the knowledge that has been shared in prior dialogue [McCoy, 1989; Moore, 1989], and
- the knowledge that will be expressed in other segments of the same explanation [Suthers, 1993].

Many of these interactions with the epistemic context cannot be identified until the contents of the candidate views are known. (Examples are given later.) Thus, it is necessary to propose potentially relevant candidate views before comparing these views to the epistemic context. Goal refinement is appropriate for proposing relevant candidate views (e.g., by allowing all refinements to proceed until they match to the knowledge base), but not for resolving the conflict between alternatives. Abstract goal refinement operators, being abstract, cannot resolve the conflict because they cannot prespecify all possible contextual constraints on all possible choices. Less abstract operators that fully

<sup>1</sup> In this paper, register is treated as a content selection problem. Bateman & Paris treat register as a lexical selection problem, *after* content selection is complete. The two approaches serve compatible purposes, although lexical selection cannot salvage an explanation that relies on obscure concepts.

specify the appropriate view for any situation would become numerous and cumbersome to manage. They would also cease to express a general theory of explanation and contain so much domain knowledge that they would compete with the knowledge base itself as the resource responsible for this knowledge. By *factoring* the constraints on view selection that derive from analysis of the question or rhetorical goal from those constraints that derive from the epistemic context, one can obtain a more compact and general theory of explanation and implement more flexible and portable explanation systems. (This point is made in detail in Suthers [1993].) Once factored out, the contextual constraints on perspective selection can be provided with an appropriate mechanism. This mechanism should make the contextual constraints explicit and modifiable, unlike the operator selection algorithms used by many goal refinement planners.

### 3.2 A Restrictive Planning Task

We now identify an appropriate mechanism for the view selection task based on its control and information processing requirements. First, the selection process will be invoked when a set of alternate views are identified (by goal refinement) as potentially relevant for an explanation. Second, there are a number of constraints and heuristics bearing on the choice between these views, some of which may be in conflict. For example, the explanation should be *relevant*, providing information that answers the question, yet be *comprehensible*, for example by minimizing the use of unfamiliar concepts. It should be *informative*, for example by saying something new, yet be *integratable*, for example by providing only information that can be related to the explanatory model under construction. Good explanations succeed in part by making appropriate tradeoffs between being relevant, comprehensible, informative, and integratable. Third, some of these heuristics require examination of the contents of the dialogue history and other evidence of the knowledge and beliefs of the questioner. In summary, we need a mechanism that is invoked when a choice between alternatives is available, is able to make tradeoffs between conflicting requirements, and does so in a context-sensitive manner.

The hybrid explanation planning architecture called HEP [Suthers, 1991, 1993] treats this context-sensitive selection between alternate views as a *restrictive planning task*. Hovy [1990] introduced "restrictive planning" to identify the use of selection criteria that are applied only when a choice is available. Here, we call these selection criteria "preferences." For example, the informativeness and comprehensibility constraints in the previous paragraph are cast as preferences.

Hovy showed that restrictive planning is necessary to make stylistic choices on dimensions such as formality, partiality, and simplicity under pragmatic constraints. Content choice on dimensions such as predication, granularity, ontology, register, and elaboration under informative and comprehensibility constraints is a similar problem. The primary difference is that Hovy's preferences mediate between text planning and natural language generation, while our preferences mediate between text planning and content retrieval. The similarities are in the

extent and handling of preferences. In both cases, preferences can be active across a segment of discourse in which multiple informative goals come and go. For example, the extent of a preference to minimize unfamiliar concepts is indefinite, while the extent of a goal to answer a particular question ends when the act of responding is completed. This orthogonality between the extents of informative goals and preferences is one reason Hovy distinguished restrictive planning from top-down goal refinement or "prescriptive" planning. Also, prescriptive goals can be planned for, while restrictive goals cannot. A goal to answer a question can be made more specific until action is possible, while criteria for context-sensitive selection between alternate views can only be applied when alternatives are available.

## 4. Preferences in HEP

In HEP, preferences are implemented as partial ordering predicates that take two views as arguments and return "<," "=", or ">" to indicate which view is preferred, if either. For example, a preference to minimize the number of unfamiliar concepts<sup>2</sup> used counts the number of unfamiliar concepts in each of two views and compares the sums, returning "<" if the first view has *more* unfamiliar concepts than the second (since the first is less preferred), ">" if the first view has fewer unfamiliar concepts, and "=" otherwise. A preference is applied to a *set* of alternate views by identifying the subset of views that are most preferred in the current context and restricting subsequent consideration to those preferred views. A *prioritized list* of preferences is applied to a set of alternate views by applying the preferences in priority order, successively restricting the set of views until only one view remains or there are no more preferences to apply.<sup>3</sup>

Preferences are activated by attaching them to the explanation plan itself. The plan data structure includes a hierarchy of dominance relations between contexts, rhetorical goals, and communicative acts. As illustrated in Figure 1, preferences are inherited down this dominance hierarchy, resulting in a priority-ordered list for any given view retrieval. Local preferences have priority over inherited ones. Default preferences are activated by placing them at a root node representing the dialogue context. This makes the default preferences active for the duration of the dialogue. A preference can be given priority for the duration of a rhetorical goal or communicative act by attaching it to the goal or act. When alternate views satisfy the relevance requirements for that act, the preferences inherited at the communicative act are applied to choose between the alternate views.

<sup>2</sup> Concept familiarity judgements are made by a user modeler that applies a number of familiarity heuristics to the dialogue history [Suthers 1993a].

<sup>3</sup> Alternately, one could use scalar evaluation functions instead of preferences, assign each evaluation function a weight and compute a total score for each view, thereby obtaining a more sensitive measure of the tradeoffs involved. However, this would require that most or all functions be evaluated for every view, whereas the present approach can stop evaluation as soon as a decision is possible.




Figure 1: Inheritance of Prioritized Preference Lists

Some of the preferences implemented in HEP are listed below in priority order. This prioritization results in a particular style of explanation. The ordering could be changed to model the styles of different explainers, to model strategic shifts taken by a single explainer, or to reflect further research concerning the effectiveness of these heuristics.

**Say Something New:** Ensures that an explanation is informative in the context of previous dialogue by preferring views that contain some new proposition over those views that do not.

**Elaborate Focal Views:** Helps the questioner relate new knowledge to knowledge active in his or her working memory by preferring views that elaborate on a view in the focus of attention (as opposed to those that take an entirely different perspective on the topic).

**Prefer Curriculum Goals:** Prefers views containing a topic that the explainer has a long term goal of introducing into the dialogue over those that do not. ("Curriculum goals" can be set by the user or by a client tutoring program.)

**Minimize Unfamiliar Concepts:** Promotes comprehensibility by counting the number of potentially unfamiliar concepts in each proposed view and preferring the view with the minimal count.

**Minimize New Propositions:** Reduces the amount of new information to be integrated by counting the number of propositions that have not been expressed previously and preferring the view with the minimal count.

**Maximize Active Concepts:** Keeps dialogue on the same conceptual basis for continuity by preferring views that contain more concepts that have been used in recent dialogue.

**Prefer General and Active Settings:** Increases the generality of the explanation by preferring universal over context-dependent statements, except when a view is dependent on a context currently in the focus of attention, in which case the preference is reversed.

**Random Preference:** Applied to all alternatives at once, this preference chooses randomly.

The above preferences can also be posted locally on a rhetorical goal to promote the priority of the preference

within the scope of the goal. In addition, other preferences are only posted locally on segments playing certain pedagogical roles [Suthers, 1993]. For example, Avoid-Circularity, posted on rhetorical goals for background definitions, prefers views that do not rely on terms which are the topics of explanation in the "foreground" segment.

## 5. An Example

We now turn to an example of the use of preferences expressing informative, comprehension, and integration constraints. This example has been run in HEP.

Various proportionality relationships concerning the capacitance of a capacitor are possible. Each relationship can be treated as a view on a simple model of capacitance. The most basic relationships are:

$C \propto A$ : capacitance is proportional to the area of the plates of the capacitor.

$C \propto K$ : capacitance is proportional to the dielectric constant.

$C \propto 1/D$ : capacitance is inversely proportional to the distance between the plates of the capacitor.

Combinations of these relationships are also available:  $C \propto AK$ ,  $C \propto A/D$ ,  $C \propto K/D$ , and  $C \propto AK/D$ . Views can also be more complex, containing multiple relationships of different types. We use a simple example here for expository purposes. In our implementation, these views are stored explicitly in a lattice-like "Graph Of Models" [Addanki *et al.*, 1991] with elaboration relations declared between them. This approach is not required by HEP: one could just as well generate the views dynamically, for example by "compositional modeling" [Falkenhainer and Forbus, 1991], provided that the elaboration relation is readily available where it applies.

In a dialogue we have recorded, a student asked how a capacitor is different from a break in a wire. The human explainer responded with the simplified view  $C \propto A$ : "The difference is in the areas of metal facing each other. The capacitance is proportional to the area of the plates." He could have also responded with  $C \propto A/D$ ,  $C \propto AK$ , or  $C \propto AK/D$ , as all of these are informative with respect to the question. When asked about the relationship between capacitance and area, HEP also prefers the simpler view.

The first inherited preference, Say Something New (an informativeness constraint), does not distinguish between the candidates, all of which contain a new proposition. If no related view is in focus and there are no curriculum goals involved in the candidate views, Elaborate Focal Views (an integration constraint) and Prefer Curriculum Goals (an informativeness constraint) do not restrict the candidates either. However, the fourth preference, Minimize Unfamiliar Concepts (a comprehension constraint), prefers  $C \propto A$  because this view contains fewer unfamiliar concepts (assuming K and D are not familiar).<sup>4</sup>

Suppose that in subsequent dialogue the explainer needed to express a constraint relationship between capacitance and the distance between the plates. As shown in Figure 2, the choice set would be  $C \propto 1/D$ ,  $C \propto A/D$ ,  $C \propto K/D$ , and  $C \propto AK/D$ . Say Something New is applied, and is again unable to restrict the set further. Next, Elaborate Focal Views applies. The view  $C \propto A$  is in focus. Views  $C \propto A/D$  and  $C \propto AK/D$  are elaborations of this view, while  $C \propto 1/D$  and  $C \propto K/D$  are not, so the latter are eliminated. Prefer Curriculum Goals is again unable to restrict the candidates. Next, Minimize Unfamiliar Concepts is applied. Concept A is assumed familiar because it has been used in previous dialogue. Concepts K and D are still assumed unfamiliar, so  $C \propto AK/D$  will be eliminated as containing two unfamiliar concepts ( $C \propto A/D$  contains only one). Hence the preferences choose  $C \propto A/D$ , resulting in a response which is not only informative with respect to the current question, but also integrates the new information with the previously expressed constraint and minimizes the number of new concepts that are introduced.

We briefly consider two variations of this example to further illustrate the functionality of the preference mechanism. For the first variation, suppose that K, dielectric constant, had been a curriculum goal when the explainer was responding to the query about the relationship between capacitance and area of the plates. Prefer Curriculum Goals would restrict the relevant views to  $C \propto AK$  and  $C \propto AK/D$  before Minimize Unfamiliar Concepts applied and chose  $C \propto AK$ . Thus the response would take the opportunity to introduce a goal topic while also answering the question. For the second variation, suppose that the user had asked "Tell me everything that affects capacitance." A preference to "Maximize Elaboration" (not on our list but easily implemented) could be posted at the explanation goal for this question. Then Maximize Elaboration would be at the front of the inherited priority list and  $C \propto AK/D$  would be preferred immediately. Maximize Elaboration would not affect retrievals outside of the scope of the goal to which it was attached. This is the utility of attaching preferences to the dominance hierarchy.

These examples illustrate how ordered preferences can select between relevant views in a manner that resolves a tradeoff between informativeness, comprehensibility, and integrability. The "views" in these examples are extremely




Figure 2: Choice Between Views in a Subsequent Exchange

simple but the mechanism and preferences apply just as well to the choice between more complex views (see for example Suthers *et al.* [1992]).

## 6. Related and Future Work

In the most closely related work, Paris [1987] observed that encyclopedia articles for novices tend to give functional descriptions while those for experts give structural descriptions. Paris partitioned her knowledge base into functional and structural information. Her planner switched between functional and structural explanation strategies within segments of an explanation according to whether the user was considered to be an expert in the portion of knowledge being covered within that segment. Restrictive planning using preferences is an improvement because the heuristics for model choice are made explicit, can be added, modified, or removed without modifying the interpreter, and can be expressed independently of other planning knowledge. In contrast, Paris' heuristic for selection between functional and structural explanation strategies was embedded in a decision algorithm that would have to be modified to add heuristics. The independent and explicit treatment of model choice preferences enables the use of more varied dimensions on which the selection is made than simply functional vs. structural.

Several other bodies of work address selection of "object perspective" using relevance as a primary criterion. For example, McKeown *et al.* [1985] select perspective by finding a high level goal in a goal lattice that generalizes goals implied by user utterances and then following an index from the identified goal to the class of predicates to be used in reasoning towards answering the question. Moore [1989] points out that some perspective choices in expert system explanation should be made based on how the object in question figures in the system's problem solving. McCoy [1989] represents perspectives as collections of related

<sup>4</sup> "Distance" is familiar to anyone, but D is "distance between the plates of a capacitor," a much more specific concept that requires correspondingly more specific evidence of familiarity.

attributes with numerical weights. A perspective is invoked when attributes mentioned in previous dialogue are highly weighted in the perspective. Once active, the weights in the perspective control the degree to which other attributes are visible to the response planner. These researchers have shown how the user's goals, the system's problem solving, and the contents of previous dialogue should influence choice of predicates or salient attributes. They did not generalize their techniques to include other influences on choice of perspective, nor to make the choice on other dimensions in addition to predication. However their techniques could be used to *generate* preferences that are applied using HEP's restrictive planning mechanism.

As presented in this paper, preferences only apply to choice between "views" that compete for fulfilling the informative requirements of a given rhetorical act. The HEP architecture is currently being redesigned to permit uniform use of preferences for *all* choices between alternatives, including the choice between sets of alternate rhetorical acts and whether or not to include optional supplemental material.

## 7. Summary

The success of an explanation can depend crucially on the domain model on which it is based. Good explanations provide relevant and adequate information while also being comprehensible. Choice between "views" on models occurs on a variety of dimensions such as predication, granularity, ontology, register, elaboration, and order. Given these choices, an abstract goal refinement strategy cannot fully specify the view to use in responding to a given question in a given situation. A restrictive planning mechanism is necessary to arbitrate between alternate views under sometimes-conflicting and context-sensitive criteria, here called "preferences." We have described a collection of preferences that can choose between alternate views in a manner that resolves the tradeoff between informativeness, comprehensibility, and integrability. The mechanism through which they are applied is sufficiently general to be used in a wide variety of discourse planning architectures.

## Acknowledgments

Thanks to David Skalak and Beverly Woolf and the anonymous reviewers for comments on earlier drafts, and to Edwina Rissland and Beverly Woolf for supporting this research.

## References

- [Addanki *et al.*, 1991] S. Addanki, R. Cremonini, and J. Penberthy. Graphs of models. *Artificial Intelligence*, 51:145-177.
- [Bateman and Paris, 1989]. J. Bateman and C. Paris. Phrasing a text in terms the user can understand. In *Proc. 11th IJCAI*, pp. 1511-1517, Detroit MI.
- [Cawsey, 1991] A. Cawsey. Generating Interactive Explanations. In *Proc. 9th NCAI (AAAI91)*, pp. 86-91, Anaheim CA.
- [Falkenhainer and Forbus, 1991] B. Falkenhainer and K. Forbus. Compositional Modeling: finding the right model for the job. *Artificial Intelligence*. 51:95-143.
- [Gentner and Gentner, 1983] D. Gentner and D. Gentner. Flowing waters or teeming crowds: Mental Models of Electricity. In *Mental Models*. D. Gentner and A. Stevens (Eds.), pp. 99-129. Hillsdale: LEA.
- [Hobbs, 1985] J. Hobbs. Granularity. In *Proc. 9th IJCAI*, pp. 432-435, Los Angeles, CA.
- [Hovy, 1990] E. Hovy. Pragmatics and natural language generation. *Artificial Intelligence*, 43:153-197.
- [Maybury, 1990] M. Maybury. Planning multisentential English text using communicative acts. Unpublished doctoral dissertation, Cambridge University.
- [McCoy, 1989] K. McCoy. Generating context-sensitive responses to object-related misconceptions. *Artificial Intelligence*, 41:157-195.
- [McKeown *et al.*, 1985] K. McKeown, M. Wish, and K. Matthews. Tailoring explanations for the user. In *Proc. 9th IJCAI*, pp. 794-798, Los Angeles CA.
- Moore [1989] J. Moore. A Reactive Approach to Explanation in Expert and Advice-giving Systems. Ph.D. Dissertation, University of California, Los Angeles.
- [Paris 1987] C. Paris. Combining discourse strategies to generate descriptions to users along a Naive/Expert spectrum. In *Proc. 10th IJCAI*, pp. 626-632, Milan, Italy.
- [Rissland, 1978] E. Rissland (formerly E. R. Michener). Epistemology, Representation, Understanding and Interactive Exploration of Mathematical Theories. Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, MA.
- [Souther *et al.*, 1989] A. Souther, L. Acker, J. Lester, and B. Porter. Using view types to generate explanations in intelligent tutoring systems. In *Proc. 11th Cognitive Science Conference*, Montreal.
- [Suthers, 1991] D. Suthers. A task-appropriate hybrid architecture for explanation. *Computational Intelligence*. 7: 315-333.
- [Suthers, 1993] D. Suthers. An Analysis of Explanation and its Implications for the Design of Explanation Planners. Ph.D. Dissertation, University of Massachusetts, Amherst.
- [Suthers *et al.*, 1992] D. Suthers, B. Woolf, and M. Cornell. Steps from explanation planning to model construction dialogues. In *Proc. 10th NCAI (AAAI92)*, pp. 24-30, San Jose, CA.
- [Stevens and Steinberg, 1981] A. Stevens and C. Steinberg. A typology of explanations and its application to intelligent computer aided instruction. Report No. 4626, Bolt Beranek and Newman Inc., Cambridge MA.
- [Stevens and Collins, 1980] A. Stevens and A. Collins. Multiple conceptual models of a complex system. In *Aptitude, Learning, and Instruction (Vol. 2)*. R. E. Snow, P. Federico, and W. E. Montague (Eds.), Hillsdale, NJ: Erlbaum, 1980. pp. 177-197.
- [White and Frederiksen, 1990] B. White and J. Frederiksen. Causal model progressions as a foundation for intelligent learning environments. *Artificial Intelligence*. 42:99-157.