

REPORT DOCUMENTATION PAGE

0674

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Project (0704-0188).

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE		3. REPORT TYPE AND DATES COVERED FINAL/15 APR 93 TO 14 APR 95	
4. TITLE AND SUBTITLE AN INTERACTIVE, INTELLIGENT TUTORING SYSTEMS FOR PREDICTION TASKS				5. FUNDING NUMBERS	
6. AUTHOR(S) BURCE PORTER				2304/GS F49620-93-1-0239	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) UNIVERSITY OF TEXAS DEPARTMENT OF COMPUTER SCIENCES AUSTIN, TEXAS 78712				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NM 110 DUNCAN AVE, SUTE B115 BOLLING AFB DC 20332-0001				10. SPONSORING / MONITORING AGENCY REPORT NUMBER F49620-93-1-0239	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE: DISTRIBUTION IS UNLIMITED				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) A major limitation of current advisory systems (e.g., intelligent tutoring systems and expert system) is their restricted ability to give explanations. The goal of our research is to develop and evaluate a flexible explanation facility, on that can dynamically generate responses to questions not anticipated by the system's designers and that can tailor these responses to individual users. To achieve this flexibility, we are developing a large knowledge base, viewpoint construction facility, and a modeling facility. In the long term we plan to build and evaluate advisory systems with flexible explanation facilities for scientists in numerous domains. In the short term, we are focusing on a single complex domain in biological science, and we are working toward two important milestones: 1) building and evaluating an advisory system with a flexible explanation facility for freshman-level students studying biology, and 2) developing general methods and tools for building similar explanation facilities in other domains.					
14. SUBJECT TERMS DTIC QUALITY INSPECTED 8				15. NUMBER OF PAGES	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED		18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	
20. LIMITATION OF ABSTRACT SAR(SAME AS REPORT)					



19951018 060



COLLEGE OF NATURAL SCIENCES
THE UNIVERSITY OF TEXAS AT AUSTIN

RIA 9/13/95

Department of Computer Sciences · Taylor Hall 2.124 · Austin, Texas 78712-1188 · (512) 471-7316

September 13, 1995

Marilyn McKee
Contracts Officer
Air Force Office of Scientific Research

Dear Ms. McKee,

The final report from my funded research (contract F49620-93-1-0239)
is attached. I apologize for the delay.

Sincerely,

A handwritten signature in cursive script that reads "Bruce Porter".

Bruce Porter,
Associate Professor

An Interactive, Intelligent Tutoring System for Prediction Tasks: Final Technical Report¹

Bruce Porter

Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

Abstract

A major limitation of current advisory systems (e.g., intelligent tutoring systems and expert systems) is their restricted ability to give explanations. The goal of our research is to develop and evaluate a *flexible* explanation facility, one that can dynamically generate responses to questions not anticipated by the system's designers and that can tailor these responses to individual users. To achieve this flexibility, we are developing a large knowledge base, a viewpoint construction facility, and a modeling facility.

In the long term we plan to build and evaluate advisory systems with flexible explanation facilities for scientists in numerous domains. In the short term, we are focusing on a single complex domain in biological science, and we are working toward two important milestones: 1) building and evaluating an advisory system with a flexible explanation facility for freshman-level students studying biology, and 2) developing general methods and tools for building similar explanation facilities in other domains.

¹Support for this research was provided by the Air Force Office of Scientific Research (contract number F49620-93-1-0239).

1 Research Objectives

The goal of our research is to develop and evaluate a *flexible* explanation facility that can dynamically generate responses to questions not anticipated by the system's designers and that can tailor these responses to individual users. Previous advisory systems have lacked these capabilities for a variety of reasons. In this section we will describe the problems of current advisory systems, the solutions to these problems that we propose, and our research activities for achieving those solutions.

Problems. The explanation facilities of current advisory systems are inflexible for two reasons:

- **Inadequate domain knowledge:** At least two factors limit the adequacy of the knowledge base as a source of "raw materials" for flexibly generating explanations: small size and task specificity. Although small size is an obvious limitation, few research projects have built a large-scale knowledge base as their "starting point" for research on explanation. Furthermore, because the knowledge for most advisory systems supports only a single task, most research on explanation has overlooked issues outside the task requirements, such as answering a range of questions, explaining terminology, and customizing explanations for specific users [12]. (For notable exceptions see work by Moore and Swartout [23, 13].)
- **Inability to reorganize knowledge:** Little work has been done to develop methods to select coherent packets of knowledge from a knowledge base, and even less on the reorganization of portions of the knowledge base to improve specific explanations. These issues have been avoided by "hardwiring" knowledge structures that are suitable for the limited explanations required by a particular advisory system. (For notable exceptions see work by McKeown [11] and Suthers [22].)

Solutions. We have developed a five-part solution to the problems of current advisory systems. Our solution comprises: (1) constructing a knowledge base which is large-scale and contains very fine-grained representations, (2) selecting and organizing knowledge with viewpoints and models, (3) generating new viewpoints on demand, (4) generating explanations which relate new information to what the user already knows, and (5) constructing and simulating models and using them to explain the behavior of mechanisms. We briefly describe each of these in turn, but we focus on the last one.

First, we have built an extensive knowledge base for one area of biology — college-level anatomy and physiology of plants [16]. Although it is under constant development, it is already one of the largest knowledge bases in existence. (Our knowledge base currently contains about 3,000 frames and over 28,000 facts.) Unlike knowledge bases built with instructional frames [8] or hypertext [2], our knowledge base consists of “atomic facts” that our explanation facility can combine in different ways to produce different explanations.

Second, we have developed methods for selecting information from the knowledge base and organizing it into a coherent bundle appropriate to the situation at hand. One organizing structure is that of *viewpoints*, which provide coherent descriptions of objects or processes. For instance, the viewpoint “photosynthesis as a production process” selects and organizes facts to explain how photosynthesis produces glucose from carbon dioxide and water. Another organizing structure is that of *models*, which are built from viewpoints and support computer simulation. For example, an energy flow model of the plant includes the viewpoints “photosynthesis as an energy transduction process” and “respiration as an energy transfer process,” and it allows an advisory system to predict and explain the effects of changes in light wavelength on a plant’s photosynthetic or respiratory rate under a variety of specific circumstances.

Third, we have developed methods to automatically generate *new* viewpoints. This ability is important because, as system designers, we cannot anticipate all the viewpoints necessary for effective explanations. For example, Table 1 lists several viewpoints on photosynthesis and the situations in which they might arise. Our question answering facility is able to construct these viewpoints by selecting and reorganizing the individual facts comprising existing viewpoints in the knowledge base (see [1]).

Forth, we have developed methods to automatically generate *integrative explanations*, which explicitly relate new information to what the user already knows. This is important to advisory systems because the coherence of an explanation depends upon the particular situation. Our system records the discourse with each user and explains new topics in ways that relate to that user’s knowledge and interests (see [10]).

Finally, we have developed methods for automatically constructing and simulating models and interpreting the consequences of simulations. These methods use existing methods of qualitative reasoning, but add two new capabilities: constructing models from large knowledge bases and generating explanations from these models. This allows our explanation facility to answer “what-if” questions that were unanticipated when the knowledge base was

<i>Viewpoint on Photosynthesis</i>	<i>Contextual Situation</i>
as a destructive process	To explain the effects of the first oxygen producing plants on other organisms during evolution.
as an essential process in ecosystem energy flow	To explain how almost all living things depend on photosynthesis for deriving energy from an abiotic source.
as a magnesium-utilizing process	To explain the effects of magnesium deficiency on the plant.
as an enabling process	To explain how photosynthesis is important for any processes which use glucose or oxygen.
as a constructive process	To explain how photosynthesis is vitally important to plant growth and reproduction.

Table 1: A few of the viewpoints on photosynthesis and the teaching situations in which they might be appropriate.

built (see [18]). Developing this capability has been our primary focus during the two years of AFOSR funding, and it is the focus of the remainder of this report.

2 Automated Modeling of Complex Systems to Answer Prediction Questions

The ability to answer prediction questions is crucial in reasoning about physical systems. The following question, from the plant physiology domain, illustrates the general form of a prediction question: "How would decreasing soil moisture affect a plant's transpiration² rate?" A prediction question poses a hypothetical *scenario* (e.g., a plant whose soil moisture is decreasing) and asks for the resulting behavior of specified *variables of interest* (e.g., the plant's transpiration rate). An answer to a prediction question includes the desired predictions and, perhaps more importantly, an explanation of the assumptions and principles that

²Transpiration is the process by which water evaporates from the leaves.

justify the predictions. In biology and ecology, such questions are important for predicting the consequences of natural conditions and management policies as well as for teaching biological and ecological principles. Because prediction is time consuming and error prone, and requires people with special knowledge, automation would be valuable.

A tool for answering prediction questions would be particularly useful for predicting the effects of global climate changes on plants and animals in specific regions. Answering these questions requires considerable knowledge: general principles of plant and animal physiology and species interactions as well as specific data on individual species, climatic events, and geologic formations. The central issue in answering prediction questions is constructing, from this wealth of information, a model that captures the important aspects of the scenario and their relationships to the variables of interest.

This section describes TRIPEL, a modeling program for answering prediction questions. Section 3 defines the modeling task. Section 4 presents TRIPEL's criteria for distinguishing relevant aspects of the scenario from irrelevant aspects. Section 5 describes the algorithm that uses these criteria to construct the simplest adequate model for answering a question.

While TRIPEL is designed to support a wide variety of domains, it has been extensively tested in the domain of plant physiology. Specifically, TRIPEL has been used to answer questions from the Botany Knowledge Base [17]. The BKB is a large (over 200,000 facts), multipurpose knowledge base covering plant anatomy, physiology, and development. It was developed by a domain expert. Section 6 discusses the results of evaluating TRIPEL using the BKB.

Because the BKB covers many different physical phenomena at many levels of detail, constructing simple yet adequate models from it is a difficult task. The techniques that allow TRIPEL to perform this task efficiently are applicable throughout science and engineering, but they are especially useful for biology and ecology.

3 The Modeling Task

TRIPEL's inputs are a prediction question and domain knowledge. The question has two parts: the scenario and the variables of interest. The scenario includes physical objects, spatial relations among them, and *driving conditions*. Driving conditions specify the behavior of selected variables (e.g., soil moisture is decreasing), their initial value (e.g., the temperature is above the freezing point), or both.

TRIPEL uses the compositional modeling approach [3], in which the modeler's job is to select those elements of domain knowledge that are needed to answer the question. Our research focuses on building differential equation models, so the elements of domain knowledge are the *influences* that pertain to the scenario.

An influence is a causal relation between two variables, as in Qualitative Process Theory [5]. The variables are real-valued, time-varying properties of the scenario (e.g., soil moisture or the plant's transpiration rate). Each influence specifies that a variable, or its rate of change, is a function of another variable.

Conceptually, each influence represents a physical phenomenon in the scenario at some level of detail. Typically, an influence represents the effect of a process (e.g., the amount of water in the plant is negatively influenced by the rate of transpiration) or a factor that affects a process's rate (e.g., the rate at which the plant absorbs water from the soil is positively influenced by the level of soil moisture). To emphasize their role in modeling, we call the set of all influences that pertain to the scenario the *candidate influences*.

TRIPEL's output, the *scenario model*, is the subset of candidate influences that are relevant to the question. Another program, the Qualitative Process Compiler [4], built on QSIM [9], simulates the scenario model starting from the initial state of the scenario. This simulation generates the predictions that are needed to answer the question. A colleague at the University of Texas is developing a program that will use the model and simulation results to answer the question and explain the answer.

4 Modeling Criteria

When the domain knowledge is extensive, as with plant physiology, it will describe many phenomena in the scenario, some at multiple levels of detail. Thus, there are two fundamental issues in modeling. First, the modeler must decide which phenomena are relevant to the question and which can be ignored. Second, for each relevant phenomenon, the modeler must choose a relevant level of detail. A candidate influence is relevant if it represents a relevant level of detail for a relevant phenomenon.

4.1 Scope

Of the many phenomena in any scenario, only a few are needed to answer any particular question. For example, of the many processes at work in a plant, the question about decreasing soil moisture only requires a model of the plant's water regulation processes. The *scope* of a model is the set of phenomena it covers.

There are two types of irrelevant phenomena. The first type, insignificant phenomena, can be ignored because they do not significantly influence the variables of interest. For instance, in our example, growth processes can be ignored because they do not significantly influence the transpiration rate.

The second type of irrelevant phenomena are those that can be treated as *exogenous*. For instance, in our example, the processes that regulate soil moisture (e.g., rain and evaporation from the soil) can be treated as exogenous. Although exogenous phenomena *do* significantly influence the variables of interest, they are nonetheless irrelevant to the question; they do not help predict the effects of the driving conditions (in our example, decreasing soil moisture) on the variables of interest.

To choose a suitable scope for the model, the modeler must eliminate both types of irrelevant phenomena. To eliminate insignificant phenomena, the modeler needs criteria for recognizing insignificant influences. By pruning insignificant influences, the modeler disconnects the model from all the insignificant phenomena in the scenario.

TRIPEL determines whether an influence is significant using time scale information. Processes cause significant change on widely disparate time scales. For example, in a plant, water flows through membranes on a time scale of seconds, solutes flow through membranes on a time scale of minutes, and growth requires hours or days. In TRIPEL, each influence that represents an effect of a process may have associated knowledge specifying the fastest time scale on which the effect is significant. Before constructing the scenario model, TRIPEL automatically determines a suitable *time scale of interest* for the question [20]. The time scale of interest allows TRIPEL to conclude that any candidate influence operating on a slower time scale is insignificant. This significance criterion is used by human modelers in many domains, including biology, ecology, and many branches of engineering [6, 15, 21].

To eliminate exogenous phenomena, the modeler needs criteria for choosing the exogenous variables of the model. Exogenous variables are those variables in the model whose behavior is determined by influences that are outside the scope of the model. All other variables in the model are *dependent*; their behavior is determined by influences in the model. Thus,

the exogenous variables constitute the boundary of the model, separating the model from exogenous phenomena in the scenario. For instance, in our example, by treating soil moisture as an exogenous variable, the processes that regulate soil moisture are excluded from the model.

To determine whether a variable in the model can be treated as exogenous, TRIPEL uses two criteria. First, by definition, the variable must not be significantly influenced, in the scenario, by any other variable in the model. One variable significantly influences another if there is a chain of candidate influences leading from the first variable to the second and every influence in the chain is significant. Second, note that the objective in a prediction question is to predict the effects of the *driving variables* on the variables of interest. A driving variable is one whose behavior or initial state is specified in the question (in our example, soil moisture). To meet that objective, the modeler must ensure that the exogenous variables do not separate the model from the driving variables of the question. Therefore, a variable in a model can be treated as exogenous only if it is not significantly influenced, in the scenario, by any driving variable of the question. TRIPEL tests these two criteria using a graph connectivity algorithm on the candidate influences [20].

In summary, TRIPEL eliminates irrelevant phenomena from the scope of the model by pruning insignificant influences (using time scale information) and by choosing suitable exogenous variables for the model. Phenomena that do not significantly influence the variables of interest, or that influence the variables of interest only through exogenous variables, are not included in the model (at any level of detail).

4.2 Level of Detail

The domain knowledge may provide multiple levels of detail for many phenomena in the scenario. For example, water in the plant can be treated as an aggregate, or the water in the roots, stem and leaves can be modeled individually. Similarly, processes can be aggregated. For example, the chemical formula for photosynthesis summarizes the net effects of its component reactions. Also, the dynamics of a process can often be summarized by its equilibrium results. For example, when the level of solutes in a plant cell changes, the process of osmosis adjusts the cell's water to a new equilibrium level. If the dynamics of this process are irrelevant, the modeler can simply treat the level of water as an instantaneous function of the level of solutes. Each of these types of alternatives arises in many areas of science and engineering.

For each relevant phenomenon in the scenario, the modeler must choose a suitable level of detail. Irrelevant details complicate simulation and make the resulting explanation less comprehensible, so the modeler must choose the simplest level of detail that is adequate for answering the question.

TRIPPEL has several criteria for choosing the level of detail. First, some approximations may be invalid in the context of the question. For example, process dynamics can only be summarized by their equilibrium result if the process reaches equilibrium very quickly relative to the time scale of interest. TRIPPEL includes a variety of general principles for recognizing that a level of detail is invalid or inadequate for a question.

Second, TRIPPEL includes coherence criteria. These ensure that the level of detail chosen for different phenomena in the model are compatible. The coherence criteria also ensure that the model does not include different levels of detail for any single phenomenon.

Finally, for those alternatives that are adequate for the question and coherent with other parts of the model, TRIPPEL chooses the one that leads to the simplest adequate model. While any simplicity criteria could be used, TRIPPEL defines one model as simpler than another if it has fewer variables. The number of variables in a model is a good heuristic measure of the complexity of simulation and of the model's comprehensibility.

In summary, the domain knowledge often provides alternative levels of detail for relevant phenomena, and the modeler must determine which level is relevant. In TRIPPEL, a level of detail is relevant if it is adequate for answering the question, coherent with other elements of the model, and it leads to the simplest adequate model.

5 Modeling Algorithm

Each candidate influence represents some phenomenon at some level of detail, so TRIPPEL's criteria for choosing scope and level of detail allow it to determine the influences that should be included in the scenario model. This section explains TRIPPEL's algorithm for selecting the relevant influences.

TRIPPEL conducts a best-first search for the simplest adequate scenario model for the question. Each state in the search space is a *partial model*, a model whose scope may not include all relevant phenomena. A partial model may contain *free variables* (variables not yet chosen as exogenous or dependent). The initial state in the search is a partial model consisting only of the variables of interest, all free. The successor function, described below,

extends the scope of a partial model to include any additional phenomena relevant to the free variables, possibly adding new free variables. A partial model has multiple successors when there are alternative levels of detail for the new phenomena. A partial model is pruned from the search if it is incoherent (i.e., violates the coherence criteria) or invalid (i.e., includes an invalid level of detail); any extension of an incoherent/invalid partial model is also incoherent/invalid. The search ends when an adequate model is found that is at least as simple as all remaining partial models; these partial models can only grow. The simplicity criterion (i.e., number of variables in the model) also serves as the evaluation function for the best-first search.

The successor function, *extend-model*, extends the scope of a partial model. *Extend-model* first determines whether all the free variables in the partial model can be exogenous, as described in Section 4.1. If so, it marks each one as exogenous and returns the resulting model, which is now complete. Otherwise, it chooses a free variable that must be dependent, and it determines all combinations of candidate influences on that variable that include every significant influencing phenomenon at some level of detail (multiple combinations arise from alternative levels of detail for these phenomena). *Extend-model* returns a set of new partial models, each the result of extending the original partial model with one of the combinations.

To extend the original partial model with a combination of candidate influences, *extend-model* adds the influences to the model, marks the chosen free variable as dependent, and adds any new free variables to the model. The new free variables are those variables referenced by the new influences but not already in the model (e.g., an influencing variable).

This algorithm is guaranteed to return the *simplest* adequate scenario model whenever an adequate scenario model exists. To see this, note that each partial model represents all its extensions. Thus, the initial partial model in the search represents all scenario models that include the variables of interest. Conceptually, the guarantee results from the following strategy:

- From the space of all possible scenario models, the algorithm repeatedly prunes away models until only a single scenario model (if any) remains.
- It never prunes a scenario model unless either (1) the model is inadequate for the question or (2) if the model is adequate, there is an adequate scenario model still under consideration (i.e., that is an extension of a partial model on the search agenda) that is at least as simple.

For the details of the proof, see [19].

6 Evaluation

To evaluate TRIPEL, we tested it on a variety of prediction questions concerning the physiology of a prototypical plant. The questions were generated by a domain expert. Each question specifies the qualitative behavior of one variable (e.g., soil moisture is decreasing) and asks for the resulting behavior of another (e.g., transpiration rate).

The domain knowledge was provided by the Botany Knowledge Base (BKB) [17]. The BKB is a large (over 200,000 facts), multipurpose knowledge base covering plant anatomy, physiology, and development. It was developed by a domain expert. The BKB provides 691 variables representing properties of a plant and its environment (soil and atmosphere), and it provides 1507 candidate influences among them. The candidate influences cover many different types of processes, including water regulation, metabolism, temperature regulation, and transportation of gasses and solutes. These processes operate on many different time scales. Many phenomena covered by the BKB are represented at multiple levels of detail, as described in Section 4.2.

The evaluation, described in detail in [19], suggests that TRIPEL is already an effective modeling program. Despite the size of the BKB, TRIPEL typically generates simple, adequate models, as judged by a domain expert. Models ranged in size from 3 variables to 93 variables, and more than half had fewer than 20 variables. Furthermore, the knowledge TRIPEL requires to construct these models is fundamental plant physiology knowledge that is natural for a domain expert to encode.

The evaluation also identified the most important limitation of TRIPEL: its criterion for determining whether one variable significantly influences another should be more sophisticated. Currently, TRIPEL concludes that one variable significantly influences another if there is a chain of influences connecting them and every influence in the chain is significant on the time scale of interest. The evaluation suggests that TRIPEL should also consider extra time lags due to the length of the chain or the spatial distance it covers. Due to this limitation, TRIPEL sometimes chooses a time scale of interest that is too fast, and it sometimes includes irrelevant elements in models. TRIPEL is designed to easily incorporate additional criteria for determining the significance of influences and chains of influences, so the main challenge for future research is simply to formulate the criteria.

7 Related Work

The modeling programs of Falkenhainer and Forbus [3], Nayak [14], and Iwasaki and Levy [7] are most similar to TRIPEL. The program of Falkenhainer and Forbus is notable for its contrasting method of selecting the scope, and Nayak's program is notable for its contrasting method of constructing a model (it builds an overly complex model and then repeatedly simplifies it). The modeling algorithm developed by Iwasaki and Levy is most similar to TRIPEL's algorithm, although their algorithm cannot automatically choose exogenous variables. For a detailed comparison between TRIPEL and these programs, see [20] and [19].

8 Conclusions

The primary results of our research are three-fold. First, we developed general methods for building intelligent tutoring systems that teach prediction. Second, we built a substantial tutoring system for the task of prediction and experimentally evaluated it. Third, we built an extensive knowledge base for college-level biology and developed prototype software for answering questions with coherent explanations. From this experience, we have learned how to structure large knowledge bases using viewpoints, and we have created a foundation on which to build tutoring system for a wide variety of prediction tasks.

References

- [1] L. Acker and B. Porter. Access methods for large knowledge bases. *Artificial Intelligence*, (submitted), 1992.
- [2] J. Conklin. Hypertext: An introduction and survey. *IEEE Computer*, 20(9):17-41, 1987.
- [3] Brian Falkenhainer and Kenneth D. Forbus. Compositional modeling: Finding the right model for the job. *Artificial Intelligence*, 51:95-143, 1991.
- [4] Adam Farquhar. A qualitative physics compiler. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 1168-1174, Menlo Park, CA, 1994. AAAI Press.

- [5] Kenneth D. Forbus. Qualitative process theory. *Artificial Intelligence*, 24:85-168, 1984.
- [6] H.J. Gold. *Mathematical Modeling of Biological Systems*. John Wiley and Sons, New York, 1977.
- [7] Yumi Iwasaki and Alon Y. Levy. Automated model selection for simulation. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 1183-1190, Menlo Park, CA, 1994. AAAI Press.
- [8] G. Kearsley. Intelligent computer-aided instruction. In Stuart C. Shapiro, editor, *Encyclopedia of Artificial Intelligence*, pages 154-159. John Wiley and Sons, New York, 1987.
- [9] Benjamin Kuipers. Qualitative simulation. *Artificial Intelligence*, 29:289-338, 1986.
- [10] J. Lester and B. Porter. A revision-based model of instructional multi-paragraph discourse production. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, August 1991.
- [11] K. McKeown. Discourse strategies for generating natural language text. *Artificial Intelligence*, 27:1-42, 1985.
- [12] K. McKeown and W. Swartout. Language generation and explanation. *Annual Review of Computing Science*, 2:401-409, 1987.
- [13] J. Moore and W. Swartout. A reactive approach to explanation. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 1504-1510, 1989.
- [14] P. Pandurang Nayak. Causal approximations. *Artificial Intelligence*, 70:277-331, 1994.
- [15] R.V. O'Neill, D.L. DeAngelis, J.B. Waide, and T.F.H. Allen. *A Hierarchical Concept of Ecosystems*. Princeton University Press, Princeton, NJ, 1986.
- [16] B. Porter, J. Lester, K. Murray, K. Pittman, A. Souther, L. Acker, and T. Jones. AI research in the context of a multifunctional knowledge base: The Botany Knowledge Base project. Technical Report AI88-88, University of Texas at Austin, 1988.

- [17] B. Porter, J. Lester, K. Murray, K. Pittman, A. Souther, L. Acker, and T. Jones. AI research in the context of a multifunctional knowledge base: The botany knowledge base project. Technical Report AI88-88, University of Texas at Austin, 1988.
- [18] J. Rickel and B. Porter. Using interaction paths for compositional modeling. In *Proceedings of the Sixth International Workshop on Qualitative Reasoning about Physical Systems*, pages 82-95, 1992.
- [19] Jeff Rickel. *Automated Modeling of Complex Systems to Answer Prediction Questions*. PhD thesis, Department of Computer Science, University of Texas at Austin, May 1995. Technical Report AI95-234.
- [20] Jeff Rickel and Bruce Porter. Automated modeling for answering prediction questions: Selecting the time scale and system boundary. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 1191-1198, Menlo Park, CA, 1994. AAAI Press.
- [21] V.R. Saksena, J. O'Reilly, and P.V. Kokotovic. Singular perturbations and time-scale methods in control theory: Survey 1976-1983. *Automatica*, 20(3):273-293, 1984.
- [22] D. Suthers. Providing multiple views of reasoning for explanation. In *Proceedings of the International Conference on Intelligent Tutoring Systems*, pages 435-442, 1988.
- [23] W. Swartout. Xplain: A system for creating and explaining expert consulting programs. *Artificial Intelligence*, 21(3):285-325, 1983.

Publications Resulting from this Research

1. J. Rickel and B. Porter, "Automated Modeling for Answering Prediction Questions: Selecting Relevant Influences", Fall Symposium on Relevance, American Association for Artificial Intelligence, 1994.
2. J. Rickel and B. Porter, "Automated Modeling for Answering Prediction Questions: Selecting the Time Scale and System Boundary", Proceedings of the National Conference on Artificial Intelligence, Seattle, Washington, 1994.
3. L. Acker and B. Porter, "Extracting Viewpoints from Knowledge Bases", Proceedings of the National Conference on Artificial Intelligence, Seattle, Washington, 1994.
4. J. Rickel and B. Porter, "Automated Modeling for Answering Prediction Questions: Selecting Relevant Influences", Fall Symposium on Relevance, American Association for Artificial Intelligence, 1994.
5. J. Rickel and B. Porter, "Automated Modeling for Answering Prediction Questions: Selecting the Time Scale and System Boundary", Proceedings of the National Conference on Artificial Intelligence, Seattle, Washington, 1994.
6. L. Acker and B. Porter, "Extracting Viewpoints from Knowledge Bases", Proceedings of the National Conference on Artificial Intelligence, Seattle, Washington, 1994.
7. C. Callaway and J. Lester. Robust Natural Language Generation from Large-Scale Knowledge Bases. *Proceedings of the Fourth Bar-Ilan Symposium on Foundations of Artificial Intelligence*, pp. 96-105, Jerusalem, Israel, 1995.
8. J. Lester and B. Porter. The KNIGHT Experiments: Empirically Evaluating an Explanation Generation System. *Symposium for Empirical Methods in Discourse Interpretation and Generation, AAAI Spring Symposium Series*, pp. 74-80, Stanford University, 1995.
9. J. Lester and B. Porter. "Interruption Handling:" Using Dynamic Replanning and Partially Refined Plans to Deal with Users' Interruptions. *AAAI Workshop on Planning for Interagent Communication*, pp. 71-77, AAAI-94, Seattle, Washington, 1994.

10. J. Lester and B. Porter. Designing Multi-Media Knowledge Delivery Systems: The Strong Representation Paradigm. *Symposium for Intelligent Multi-Media Multi-Modal Systems, AAAI Spring Symposium Series*, pp. 64-72, Stanford University, 1994.
11. B. Blumenthal and B. Porter, "Analysis and Empirical Studies of Derivational Analogy", *Artificial Intelligence Journal*, volume 67, number 2, pp. 287-328, June 1994.