To appear in *Proceedings of AAAI-94*

Automated Modeling for Answering Prediction Questions: Selecting the Time Scale and System Boundary*

Jeff Rickel and Bruce Porter
Department of Computer Science
University of Texas
Austin, Texas 78712
rickel@cs.utexas.edu, porter@cs.utexas.edu

Abstract

The ability to answer prediction questions is crucial to reasoning about physical systems. A prediction question poses a hypothetical scenario and asks for the resulting behavior of variables of interest. Prediction questions can be answered by simulating a model of the scenario. An appropriate system boundary, which separates aspects of the scenario that must be modeled from those that can be ignored, is critical to achieving a simple yet adequate model. This paper presents an efficient algorithm for system boundary selection, it shows the important role played by the model's time scale, and it provides a separate algorithm for selecting this time scale. Both algorithms have been implemented in a compositional modeling program called TRIPEL and evaluated in the plant physiology domain.

1 Introduction

The ability to answer prediction questions is crucial to reasoning about physical systems. 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 growth rate). Such questions are important in verifying designs, testing diagnostic hypotheses, and tutoring in science and engineering.

Prediction questions can be answered by simulating a model of the scenario. Simulation provides the desired predictions, and the model additionally supports subsequent explanation. The model must be sufficiently comprehensive to ensure reliable predictions yet simple so simulation is efficient and the explanation is comprehensible.

To balance these competing requirements, a modeler must choose a *system boundary* that separates aspects of the scenario that must be modeled from those that can be ignored. Despite the importance of choosing a suitable system boundary, current modeling programs for answering prediction questions shift responsibility for this issue to the people posing the question or representing the domain knowledge (see Section 8).

This paper presents an efficient algorithm for choosing system boundaries and explains its role in TRIPEL, a modeling program for answering prediction questions.[1] The paper shows that the system boundary can be chosen efficiently by first identifying the time scale on which the variables of interest are affected in the scenario. It presents a separate algorithm for determining this time scale. The correctness and efficiency of our methods have been evaluated on questions about plant physiology.

2 The Modeling Task

The input to the modeler consists of a prediction question and domain knowledge. The question has two parts: the scenario and the variables of interest. The scenario includes physical objects, relations among them, and *behavioral conditions*. Behavioral conditions specify the initial value of selected variables (e.g., the amount of soil water is above the permanent wilting percentage) and/or their behavior (e.g., the amount of soil water is decreasing).

TRIPEL uses the compositional modeling approach introduced by Falkenhainer and Forbus (1991), in which the domain knowledge provides a set of *model fragments*, the building blocks for models. Each model fragment describes some aspect of the scenario. (Falkenhainer and Forbus show how to generate the model fragments for a scenario from general domain knowledge.) The modeler constructs a model of the scenario by choosing a subset of the model fragments.

Model fragments specify relations among variables of the scenario. In plant physiology, *influences* are the most natural representation for such relations. An influence is a causal relation between two variables, as in Qualitative Process (QP) Theory (Forbus 1984),

*Support for this research is provided by a grant from the National Science Foundation (IRI-9120310), a contract from the Air Force Office of Scientific Research (F49620-93-1-0239), and donations from the Digital Equipment Corporation.

[1] TRIPEL is an acronym for "Tailoring Relevant Influences for Predictive and Explanatory Leverage." It is also a style of strong ale made by Trappist Monks in Belgium.

along with its *operating conditions* (behavioral conditions under which it holds) and associated *modeling assumptions*. The variables are real-valued, time-varying properties of the scenario. There are two types of influences: a *functional influence* specifies that one variable is a function of another (e.g., QP theory's indirect influences), and a *differential influence* specifies that the first derivative of one variable is a function of another variable (e.g., QP theory's direct influences).

In TRIPEL, each influence serves as a model fragment. This is natural, since each is an independent fact. It also allows the modeler flexibility to include or exclude any influence from the model. To emphasize their role in modeling, we call the influences (model fragments) for a scenario the *candidate influences*.

The output of the modeler, the *scenario model*, is a subset of the candidate influences. The variables referenced in this model are partitioned into *exogenous* variables, whose behavior is determined by influences external to the model, and *dependent* variables, whose behavior is determined by the model. To determine which combinations of candidate influences constitute an acceptable scenario model, the domain knowledge includes *coherence constraints*, which specify inconsistent combinations of modeling assumptions (e.g., assumption classes (Falkenhainer & Forbus 1991)).

Once constructed, the scenario model is simulated starting from the initial state, and the model and simulation results are used to answer the question and explain the answer.

3 Modeling Algorithm

The exogenous variables of a scenario model constitute its system boundary. To illustrate the role of system boundary decisions in compositional modeling, we briefly present our modeling algorithm.

TRIPEL conducts a best-first search for a scenario model for the question. Each state in the search space is a *partial model*, a model that 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 a partial model with alternative ways of modeling one of its free variables; this may add new free variables to the model. A partial model is pruned from the search if it is incoherent (i.e., violates the coherence constraints); any extension of an incoherent partial model is also incoherent. The goal of the search is to find the simplest adequate scenario model for the question.

A scenario model is adequate if it satisfies the following conditions: it includes all variables of interest, it satisfies all coherence constraints, its system boundary (set of exogenous variables) is adequate (discussed in Section 4), and each dependent variable has an adequate set of influences on it (i.e., the influences represent all significant influencing phenomena at some level of detail).

The adequate models are partially ordered by simplicity. While any simplicity criteria could be used, we define one model as simpler than another if it has fewer variables. With this criterion, 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. This criterion also serves as the evaluation function for the best-first search.

The successor function, extend-model, extends a partial model with alternative ways of modeling one of its free variables. Extend-model first determines whether all the free variables can be exogenous; if so, it marks each one as exogenous and returns the resulting model. Otherwise, it chooses a variable that must be dependent and determines all combinations of candidate influences on that variable that would provide an adequate model of it (multiple combinations arise from alternative ways of modeling some of the underlying influencing phenomena).[2] 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 one of the combinations of candidate influences, extend-model adds the influences to the model, marks the variable as dependent, and adds any new free variables to the model. These new free variables include any variable referenced by the new influences that was not already in the model (e.g., an influencing variable or a variable appearing in operating conditions).

System boundary decisions arise in the successor function extend-model. Given a partial model and one of its free variables, extend-model must determine whether the variable can be exogenous. Such decisions are important; if the variable is dependent, the model must be extended to include additional influences (on that variable) and variables (referenced by those influences). The next section describes TRIPEL's criteria and algorithm for determining if a variable can be classified as exogenous.

4 System Boundary Selection

Selection Criteria

An exogenous variable must satisfy two criteria. First, by definition, the variable must not be "significantly influenced" (defined below) by any other variable in the model. Second, the variable must not be significantly influenced by any *driving variable* (variable referenced in the question's behavioral conditions). The second criterion ensures that the system boundary doesn't disconnect the model from relevant behavioral conditions.

To determine whether one variable significantly influences another, TRIPEL uses the candidate influences.

[2] This step is not discussed in this paper. TRIPEL has a method for identifying these combinations, but any method will do; the algorithms in this paper do not depend on how the combinations are determined.

The candidate influences form a graph in which variables are nodes and the influences are directed edges from their influencing variable to their influenced variable. One variable *significantly influences* another variable if and only if there is an *influence path* (path in the graph) from the first variable to the second and every influence in the path is significant.

TRIPEL determines whether an individual 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, growth requires hours or days, and surrounding ecological processes may occur on a time scale of months or years. In TRIPEL, each differential influence, which specifies an effect of a process, has an associated modeling assumption that specifies the fastest time scale on which the effect is significant. Functional influences, being instantaneous, are significant on any time scale. After choosing an appropriate *time scale of interest* for the question (as discussed in Section 5), TRIPEL concludes that any candidate influence with a slower time scale is insignificant.

For example, consider the question "What happens to the amount of ABA in a plant's guard cells when the turgor pressure in its leaves decreases?" Turgor pressure is the hydraulic pressure in plant cells. ABA (abscisic acid) is a hormone that controls the plant's response to water stress. As will be discussed in Section 5, this question is best answered on a time scale of minutes.

Part (A) of Figure 1 shows some of the candidate influences for the example question. Leaf turgor pressure significantly influences guard cell ABA amount because there is an influence path from the former to the latter (along the top of the figure), and every influence in this path is significant on a time scale of minutes. However, the water uptake rate (lower left corner) does not significantly influence guard cell ABA amount, because the first influence on the influence path is significant only on a time scale of hours or longer.

Selection Algorithm

Using the notion of influence paths, extend-model *could* run a graph connectivity algorithm for each system boundary decision. A free variable in a partial model can be exogenous if the graph algorithm determines that the variable is not significantly influenced by any driving variable of the question or any other variable in the model.

However, this naive algorithm is inefficient. Each run of the graph algorithm will repeat much of the search that previous runs did. To avoid this problem, TRIPEL determines all variables and influences that *might* be relevant to the question and computes and caches connectivity relations among the variables *before* beginning the search for an adequate scenario model. These potentially relevant variables and influences constitute the search space that would be repeatedly searched by the naive algorithm.

To identify all the potentially relevant variables and influences, TRIPEL starts with the variables of interest and conducts a breadth-first search backwards through the candidate influences. If a variable is potentially relevant, so is any significant influence on it. If an influence is potentially relevant, so are its influencing variable and any variables appearing in its operating conditions. This search ends at variables that are not significantly influenced on the time scale of interest or variables that are significantly influenced only by previously-discovered relevant variables (i.e., through feedback loops).

In the example in Figure 1, the search for potentially relevant variables and influences begins with the influences on guard cell ABA amount. The influences of transpiration on leaf mesophyll water and water uptake on xylem water are insignificant on the time scale of interest (minutes); removing these two influences disconnects the potentially relevant variables from the remainder of the candidate influences, including the feedback loop through transpiration. Part (B) of Figure 1 shows the result, the potentially relevant variables and influences for the example.

As illustrated by the example, the search for potentially relevant variables and influences will typically have to traverse only a fraction of the variables and candidate influences of the scenario. In natural systems, like plants, animals, and ecosystems, modularity arises from the widely disparate time scales at which processes cause change (Allen & Starr 1982; Kuipers 1987; O'Neill *et al.* 1986; Rosswall, Woodmansee, & Risser 1988; Segal 1980). The result is a hierarchy of nearly decomposable subsystems; processes acting *within* a subsystem cause significant change quickly, while processes acting *across* subsystems cause change more slowly (Allen & Starr 1982; Kuipers 1987; O'Neill *et al.* 1986; Simon & Ando 1961). The time scale of interest filters out influences that are significant only on slower time scales, thus isolating the variables of interest in their own nearly decomposable subsystem. The search for potentially relevant variables and influences is confined to this subsystem because the influences from other subsystems are insignificant.

After determining the graph of potentially relevant variables and influences, TRIPEL constructs the adjacency matrix for the transitive closure of this graph. This two-dimensional, Boolean *connectivity array* records the connectivity between every pair of potentially relevant variables; thus, TRIPEL can tell whether any variable significantly influences any other variable by consulting a single cell of the array. This array can be computed efficiently; the Floyd-Warshall algorithm computes it in $\Theta(n^3)$ time, where n is the number of nodes (potentially relevant variables) in the

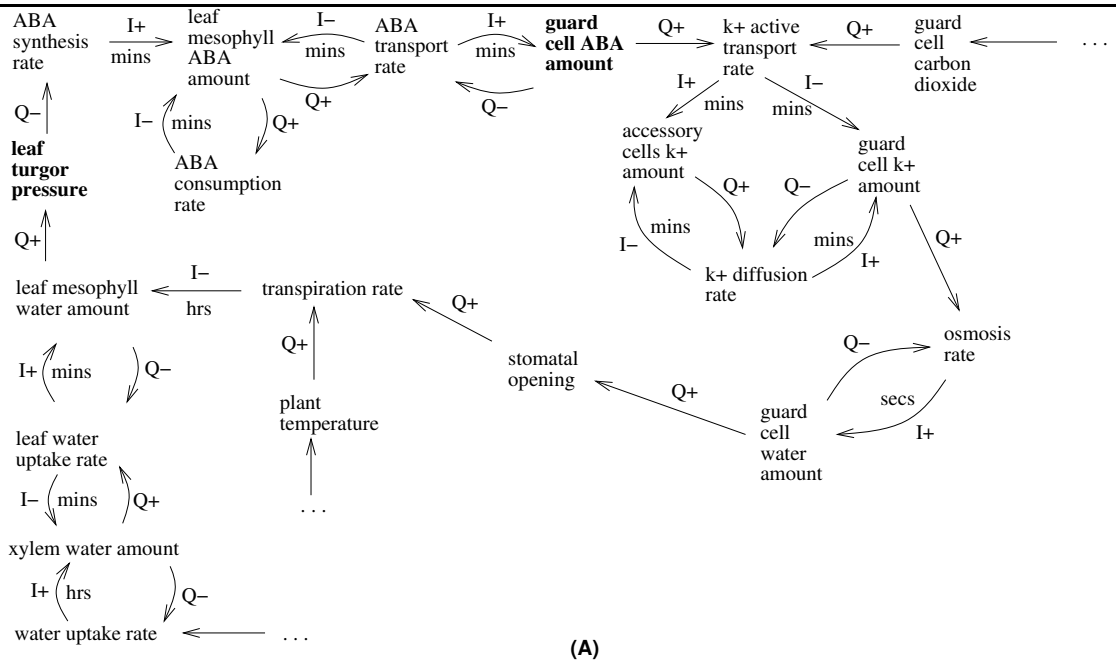

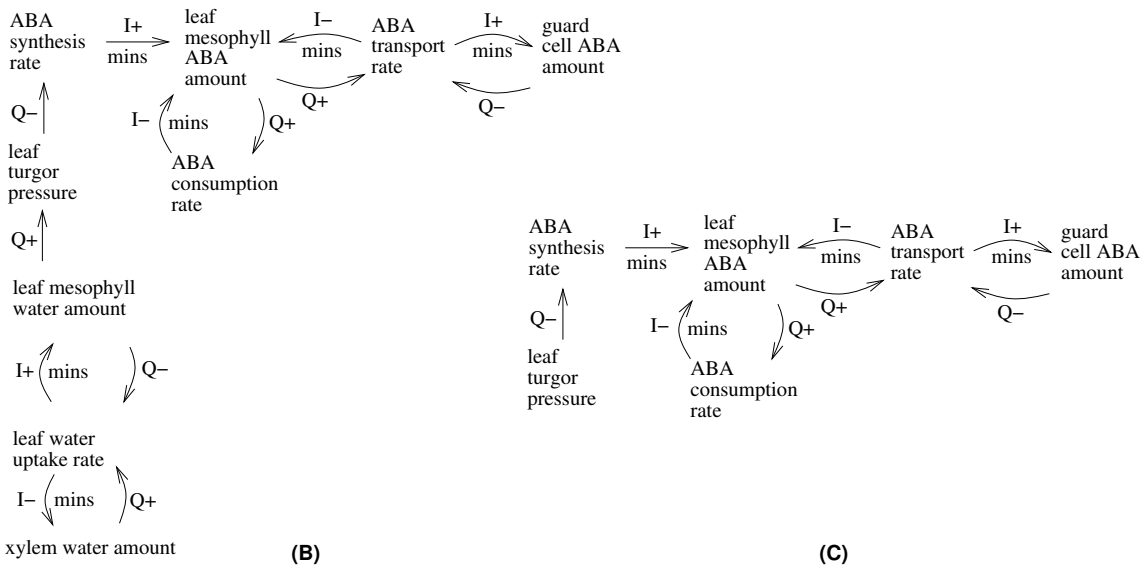

Figure 1:
(A) A subset of the candidate influences for the question "What happens to the amount of ABA in a plant's guard cells when the turgor pressure in its leaves decreases?" The driving variable, leaf turgor pressure, and the variable of interest, guard cell ABA amount, are shown in bold. Each influence is labeled with its type (Q+ and Q– are types of functional influences, and I+ and I– are types of differential influences) and the time scale on which it is significant (functional influences are significant on any time scale). Ellipses indicate connection to the remainder of the candidate influences. To focus on system boundary issues, the figure does not show alternative levels of detail. (B) The potentially relevant variables for the question. (C) An adequate scenario model for the question.

graph (Cormen, Leiserson, & Rivest 1989).

After computing the connectivity array, TRIPEL searches for an adequate scenario model as described in Section 3. Extend-model consults the array for each system boundary decision. A free variable v in a partial model m must be dependent in m (and any extension of m) if, in the connectivity array, v is influenced by any other variable in m or any driving variable of the question. If not, v can be exogenous in m but not necessarily in extensions of m, since they may contain additional variables that influence v. Therefore, as described in Section 3, extend-model doesn't mark variables in a partial model as exogenous until all remaining free variables in that model can be exogenous. At that point, the model is complete, so no other variables need to be added.

In the example, the search for an adequate scenario model begins with the partial model consisting only of guard cell ABA amount. TRIPEL incrementally extends this model until its contents match those shown in Part (C) of Figure 1. At this point, the free variable leaf turgor pressure is chosen as exogenous because it satisfies both criteria: it is not significantly influenced by any other variable in the partial model nor by any other driving variable. This model is the simplest adequate model for the question.

5 Time Scale Selection

A time scale of interest provides an important source of power in modeling. Besides providing the criteria for assessing the significance of influences, a time scale of interest also allows TRIPEL to use quasi-static approximations, in which fast processes are modeled through simple functional relations that summarize their equilibrium results (Iwasaki 1988; Kuipers 1987; Rickel & Porter 1992; Schaffer 1981; Simon & Ando 1961). Similarly, TRIPEL can model separate pools of substance or energy as a single aggregate compartment when they are kinetically distinguishable only on time scales much faster than the time scale of interest (Jacquez 1985; Simon & Ando 1961; Zeigler 1980). Thus, a time scale of interest allows many important model simplifications.

However, the person asking the question cannot be expected to provide the time scale of interest. Typically, this person will not even know which influences link the behavioral conditions to the variables of interest, much less their time scales. The modeler must choose, as the time scale of interest, a time scale that is adequate for answering the question. This section describes TRIPEL's criteria and algorithm for choosing a time scale of interest.

Selection Criteria

A prediction question asks for the effects of behavioral conditions on variables of interest. Therefore, a time scale is adequate for answering the question only if, on that time scale, every variable of interest is significantly influenced by some driving variable. Additionally, assuming that a prediction question asks for the behavior of the variables of interest *beyond* the initial state, the influence paths relating the driving variables to the variables of interest must be capable of causing changes in the variables of interest.

Through an individual influence, one variable can cause change in another variable in two ways: (1) with a differential influence, a specified value for the influencing variable (along with values for other influencing variables) provides the rate of change of the influenced variable; (2) in contrast, a functional influence can cause change only if the influencing variable is changing (Forbus 1984). This implies that a driving variable can cause change in a variable of interest only if the influence path connecting them contains a differential influence or the behavioral conditions specify that the driving variable is changing (in which case a path of functional influences will propagate the change). If either case is satisfied, the influence path is a *differential influence path*.

In our earlier example, since the question specifies that turgor pressure is decreasing, any influence path from turgor pressure to another variable is a differential influence path, capable of causing change. In contrast, if the question only specified that turgor pressure is above the "yield point" (above which the pressure causes cell growth), an influence path leading from turgor pressure is differential only if it contains a differential influence (as is the case with the influence of turgor pressure on cell growth).

Using this concept, the criterion for an adequate time scale is more concrete: A time scale is adequate for answering a prediction question only if, for every variable of interest, there is a differential influence path, consisting solely of candidate influences that are significant on that time scale, leading from some driving variable to that variable of interest. This criterion prevents TRIPEL from selecting a time scale on which simulation could only predict the initial state of the variables of interest resulting from the behavioral conditions.

Selection Algorithm

While the search for influence paths during system boundary selection is kept manageable by the time scale of interest, no such focus is available when choosing the time scale of interest. The complete set of candidate influences could be enormous, so generating that set and searching through it for influence paths could be prohibitively expensive. Efficient time scale selection requires the ability to generate and search through only a fraction of the candidate influences.

TRIPEL gains efficiency by starting with the fastest possible time scale and testing successively slower time scales until it finds one that is adequate. When TRIPEL tests a time scale, it can ignore all influences that are

significant only on slower time scales, so each test operates on a manageable fraction of the candidate influences. The set of significant influences grows monotonically as TRIPEL considers slower time scales, so TRIPEL performs the inexpensive tests before the more expensive ones. TRIPEL chooses the first adequate time scale it finds as the time scale of interest.

To determine whether a candidate time scale is adequate, TRIPEL conducts a breadth-first search, starting from the driving variables, for variables that are reachable via significant (on that time scale) influence paths. For each reachable variable, TRIPEL records whether it is reachable via a differential influence path or a functional one. The actual influence paths are not recorded. The search ends when every variable of interest is reachable by a differential influence path (in which case the time scale is adequate) or when the set of variables reachable at that time scale is exhausted (in which case the time scale is not adequate).

For the example question, TRIPEL first tests a time scale of seconds. Part (A) of Figure 1 illustrates that only the ABA synthesis rate is significantly influenced by leaf turgor pressure on this time scale. Next, TRIPEL tests a time scale of minutes. On this time scale, there is a differential influence path from leaf turgor pressure to guard cell ABA amount (along the top of the figure), so this time scale is chosen.

6 Evaluation

To evaluate our methods of time scale and system boundary selection, we tested TRIPEL on seven prediction questions concerning the physiology of a prototypical plant, including the example described above. Each question specifies the qualitative behavior of one variable and asks for the resulting behavior of another.

Our plant physiology knowledge base provides 77 variables and 155 candidate influences for this plant. Of the variables, 33 represent an amount of some substance or energy in a plant compartment, and 39 represent the rates of different processes. The candidate influences cover processes of water regulation, carbon dioxide regulation and carbohydrate regulation. The time scales of these processes range from seconds to hours. Many phenomena are represented at multiple levels of detail, based on the following:

- aggregation of pools (e.g., modeling water in the roots and stem as separate pools or as a single aggregate pool)
- aggregation of processes (e.g., modeling photosynthesis as an aggregate process or separately modeling its components, the light and dark reactions)
- quasi-static approximations (i.e., modeling the net equilibrium result of a set of processes or modeling their underlying dynamics)

For each question, TRIPEL chose the appropriate time scale and a reasonable system boundary, as judged by a domain expert. Consequently, the chosen scenario models included the variables and influences required for answering each question, and they excluded irrelevant ones. On average, the models contained 11 variables and 14 influences, substantially fewer than the number in the knowledge base. The largest scenario model contained only 15 variables and 20 influences. While the simplicity of these models is partially due to omitting unnecessary detail, their simplicity also reflects a well-chosen system boundary; each model excludes a number of plant subsystems.

Moreover, TRIPEL generated these models efficiently, requiring less than 15 seconds to find the time scale and the simplest adequate model for each question. In each case, connectivity analysis — the most expensive step in determining the system boundary — was performed on only a fraction of the influence graph. In the best case, connectivity analysis considered only 4 potentially relevant variables, and it considered 51 in the worst case. This shows how effectively the time scale of interest restricts the set of potentially relevant variables and influences; disregarding time scale, all the variables in the knowledge base are connected. We expect the fraction of potentially relevant variables to be even smaller for a knowledge base with a wider variety of time scales and a more extensive coverage of plant subsystems.

7 Future Work

Our method of time scale selection has several limitations. It assumes that a single time scale will suffice for answering the question, but some questions require multiple time scales (Iwasaki 1990; Kuipers 1987). Also, the criteria for an adequate time scale are necessary but not always sufficient; the most important connections between the behavioral conditions and the variables of interest may not lie at the fastest adequate time scale.

The algorithm for system boundary selection can be strengthened with additional methods for recognizing insignificant influences. Each such method further reduces the number of potentially relevant variables and tightens the resulting system boundary.

A more thorough evaluation requires more extensive domain knowledge. We are currently evaluating TRIPEL using the Botany Knowledge Base (Porter *et al.* 1988), which includes over 200 processes described at multiple levels of detail.

We expect our methods to apply to a wide variety of domains. In addition to biological and ecological domains, time scale knowledge appears useful in engineering domains as well. Kokotovic, O'Malley, and Sannuti (1976) and Saksena, O'Reilly, and Kokotovic (1984) survey hundreds of applications in many different engineering fields in which models are simplified using knowledge of the disparate time scales of processes.

8 Related Work

The modeling algorithm of Falkenhainer and Forbus (1991) requires, as input, a system decomposition for the scenario. In contrast, our algorithm determines system boundaries using only the model fragments. Falkenhainer and Forbus assume the system decomposition is based on partonomic structure; however, O'Neill *et al.* (1986) argue that approximate system boundaries in natural systems arise from differences in process rates and that these boundaries may not correspond to standard structural decompositions. Finally, as illustrated in our previous paper (Rickel & Porter 1992), Falkenhainer and Forbus's approach to selecting system boundaries is not sufficiently sensitive to the connection between behavioral conditions and variables of interest.

The modeling algorithm of Nayak *et al.* (Nayak 1992; Nayak, Joskowicz, & Addanki 1992) requires "model-as" constraints, provided in the domain knowledge, to identify potentially relevant model fragments. These constraints don't ensure that the model includes an influence path from the driving variable to the variable of interest, so a subsequent step adds model fragments until the model is adequate. In contrast, our method finds all potentially relevant variables and influences using only the candidate influences, and these variables and influences will include the significant influence paths from the driving variables to the variables of interest. Furthermore, their criteria for choosing exogenous variables are suitable for their task, explaining a specified causal relation, but are too weak for prediction questions. Their criteria only require some influence path from the driving variable to the variable of interest; the resulting scenario model may include an exogenous variable that, in reality, is significantly influenced by another variable in the model.

The modeling algorithms of Williams (1991) and Iwasaki and Levy (1993) require, as input, the variables that can be exogenous for the question. Although these algorithms can determine which exogenous variables must be included in the scenario model, neither algorithm can determine exogenous variables automatically.

Iwasaki (1990), Kuipers (1987), and Yip (1993) present modeling and simulation methods that exploit time scale information, but they do not provide methods for selecting the time scale of interest.

The time scale on which a differential influence is significant bundles two pieces of knowledge: the rates at which the influencing process operates and the level of change in the influenced variable that is considered significant. TRIPEL directly associates differential influences with their time scale of significance because this coarse level of knowledge is often more readily available than the underlying knowledge. However, it may be useful or necessary in some domains to infer the time scale from the underlying knowledge, especially if the level of significant change depends on the question. Iwasaki (1990) has explored this approach.

The work described in this paper builds on our previous methods for selecting a time scale and system boundary (Rickel & Porter 1992). The previous methods had to keep track of all the particular interaction paths, which can be prohibitively expensive when there are many candidate influences. In contrast, our current methods require only connectivity information, for two reasons: (1) the influences in a model are selected through incremental extension of partial models (Section 3) rather than a search for all interaction paths, and (2) a time scale of interest is selected by testing candidate time scales one by one (Section 5).

9 Conclusions

To provide a reliable, comprehensible answer to a prediction question, a modeler must choose an appropriate system boundary. The time scale of interest plays an important role in selecting the system boundary; TRIPEL uses this time scale to identify insignificant influences. To choose the time scale of interest and system boundary, TRIPEL searches for relevant influence paths. Time scale knowledge makes this search practical. Our evaluation indicates that TRIPEL efficiently selects appropriate time scales and system boundaries.

10 Acknowledgements

We appreciate the help of Susan Branting, Dan Clancy, Brian Falkenhainer, Bert Kay, and Rich Mallory, all of whom provided valuable comments on earlier drafts; and our botany experts Bassett Maguire and Art Souther, both of whom helped evaluate TRIPEL.

References

Allen, T., and Starr, T. 1982. *Hierarchy.* Chicago: University of Chicago Press.

Cormen, T. H.; Leiserson, C. E.; and Rivest, R. L. 1989. *Introduction to Algorithms.* New York: McGraw-Hill.

Falkenhainer, B., and Forbus, K. 1991. Compositional modeling: Finding the right model for the job. *Artificial Intelligence* 51:95–143.

Forbus, K. 1984. Qualitative process theory. *Artificial Intelligence* 24:85–168.

Iwasaki, Y., and Levy, A. Y. 1993. Automated model selection for simulation. In *Proceedings of the Seventh International Workshop on Qualitative Reasoning*, 108–116.

Iwasaki, Y. 1988. Causal ordering in a mixed structure. In *Proceedings of AAAI-88*, 313–318. San Mateo, CA: Morgan Kaufmann.

Iwasaki, Y. 1990. Reasoning with multiple abstraction models. In T. Ellman, R. K., and Mostow, J., eds., *Working Notes of the AAAI Workshop on Automatic Generation of Approximations and Abstractions*, 122–134.

Jacquez, J. 1985. *Compartmental Analysis in Biology and Medicine*. Ann Arbor, MI: University of Michigan Press.

Kokotovic, P.; O'Malley, Jr., R.; and Sannuti, P. 1976. Singular perturbations and order reduction in control theory – an overview. *Automatica* 12:123–132.

Kuipers, B. 1987. Abstraction by time scale in qualitative simulation. In *Proceedings of AAAI-87*, 621–625.

Man-kam Yip, K. 1993. Model simplification by asymptotic order of magnitude reasoning. In *Proceedings of AAAI-93*, 634–641. Menlo Park, CA: AAAI Press.

Nayak, P.; Joskowicz, L.; and Addanki, S. 1992. Automated model selection using context-dependent behaviors. In *Proceedings of AAAI-92*, 710–716. Menlo Park, CA: AAAI Press.

Nayak, P. P. 1992. Causal approximations. In *Proceedings of AAAI-92*, 703–709. Menlo Park, CA: AAAI Press.

O'Neill, R.; DeAngelis, D.; Waide, J.; and Allen, T. 1986. *A Hierarchical Concept of Ecosystems*. Princeton, NJ: Princeton University Press.

Porter, B.; Lester, J.; Murray, K.; Pittman, K.; Souther, A.; Acker, L.; and Jones, T. 1988. AI research in the context of a multifunctional knowledge base: The botany knowledge base project. Technical Report AI88-88, University of Texas at Austin.

Rickel, J., and Porter, B. 1992. Automated modeling for answering prediction questions: Exploiting interaction paths. In *Proceedings of the Sixth International Workshop on Qualitative Reasoning*, 82–95. Edinburgh, Scotland: Heriot-Watt University.

Rosswall, T.; Woodmansee, R.; and Risser, P., eds. 1988. *Scales and Global Change: Spatial and Temporal Variability in Biospheric Processes*. New York: John Wiley and Sons.

Saksena, V.; O'Reilly, J.; and Kokotovic, P. 1984. Singular perturbations and time-scale methods in control theory: Survey 1976-1983. *Automatica* 20(3):273–293.

Schaffer, W. 1981. Ecological abstraction: The consequences of reduced dimensionality in ecological models. *Ecological Monographs* 51(4):383–401.

Segal, L., ed. 1980. *Mathematical Models in Molecular and Cellular Biology*. Cambridge: Cambridge University Press. chapter 3.

Simon, H., and Ando, A. 1961. Aggregation of variables in dynamic systems. *Econometrica* 29:111–138.

Williams, B. 1991. Critical abstraction: Generating simplest models for causal explanation. In *Working Papers of the Fifth International Workshop on Qualitative Reasoning about Physical Systems*, 77–92. Austin, TX: University of Texas at Austin.

Zeigler, B. P. 1980. Simplification of biochemical reaction systems. In Segal, L., ed., *Mathematical Models in Molecular and Cellular Biology*. Cambridge: Cambridge University Press.