Automated Modeling for Answering Prediction Questions: Exploiting Interaction Paths*

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

Abstract

This paper presents an algorithm for constructing differential equation models to answer prediction questions. A prediction question poses a hypothetical scenario and asks for the resulting behaviors of specific real-valued state variables ("quantities of interest"). A model for such a question should include only those aspects of the scenario relevant to the quantities of interest. Furthermore, it should include only as much detail as is necessary to answer the question. Some previous methods may build models that fail to relate the quantities of interest to the conditions of the scenario, while others have limited ability to exclude irrelevant aspects. Additionally, while many decisions on how much detail to include can be determined by an appropriate choice of the time scale of analysis, previous approaches require the question to provide this time scale. Our modeling algorithm overcomes these limitations by identifying the ways in which quantities of interest interact with conditions in the scenario. Only those aspects of the scenario that provide these interactions are included in the model. The algorithm also uses this knowledge to determine an appropriate time scale of analysis, which determines how much detail to include in the model.

1 Background: The Prediction Task and the Role of Modeling

Our long-term research objective is to develop knowledge bases with knowledge comparable to that of scientists and engineers and software that exploits this knowledge to answer questions. Such software will support intelligent tutoring, provide important information resources for scientists and engineers, and aid in designing and maintaining large engineered systems. While we expect the results of our research to be applicable to many fields of science and engineering, we have chosen to initially apply and evaluate our methods in the domain of plant physiology because a large knowledge base in this domain already exists [11].

*Support for this research is provided by grants from the National Science Foundation (IRI-8620052 and IRI-9120310) and NASA (NAG 9-512) and by donations from Digital Equipment Corporation and Lockheed AI Center.

Our current research focuses on prediction questions such as "How would a decreasing amount of soil water affect plant size and growth rate?" We define the **prediction task** as follows:[1]

Given:

- **domain knowledge** (e.g., knowledge of plant physiology)
- a **scenario**, consisting of the following:
  - **structural conditions**: physical individuals and relations among them representing physical structure (e.g., a plant in soil)
  - **behavioral conditions**: initial values or behaviors of state variables (e.g., the amount of water in the soil is decreasing)
- a set of **state variables of interest** whose behaviors the user wants to predict (e.g., size and growth rate of the plant)

Produce:

- predictions: the expected behavior of each state variable of interest (e.g., "size increases and eventually levels off; growth rate slows and eventually stops")
- explanations: justifications for the predictions (e.g., "water uptake decreases, causing a loss in turgor pressure, which is a driving influence on plant growth")

The scenario and state variables of interest collectively constitute the **prediction question.**

There are a variety of tools for prediction in science and engineering. State variables in such domains are typically real-valued functions of time (**quantities**). Analytic and simulation methods based on differential equations can predict the time-varying behavior of quantities. However, these **analysis** methods require a **model** of the scenario in terms of functional and differential relations among scenario quantities. The scientist or engineer must construct the model from the question, and few tools are available to help with this modeling task. The focus of our research is on constructing models from questions automatically.

Modeling is difficult because it requires identifying all, and only, the domain knowledge pertinent to the prediction question. If the model includes too much information, effort will be wasted during analysis and the predictions will be unnecessarily complicated and hard to explain. If it includes too little information, its predictions may be unreliable. To support a broad range of questions, the domain knowledge must include many perspectives and levels of detail, but most of them will be irrelevant to any particular question.

The domain knowledge provides the building blocks for models, called **model fragments**. Each model fragment specifies a set of functional or differential relations among quantities. It may also specify preconditions on the validity of these relations. For instance, the growth rate of a plant is a function of its turgor pressure (the hydraulic pressure in its cells), but only as long as the turgor pressure is above a threshold. We define a model fragment as a propositional implication whose antecedent is a conjunction of **operating conditions** and whose consequent is a conjunction of relations.[2] Operating conditions delimit the states in which the relations of the model fragment hold; they are equality and inequality constraints on quantity values. The relations are a set of functional and differential relations among quantities; while our modeling methods support many types of relations, such as quantitative and qualitative equations, this paper will restrict relations to be Qualitative Process Theory **influences** [4].

[1] Terms appear in boldface where first used.

[2] Falkenhainer and Forbus [3] define a model fragment as a first-order implication. When such a first-order implication is instantiated, the result is a propositional implication that we call a model fragment. We view the first-order implication as a schema for generating model fragments.

2 Fundamental Modeling Issues

For purposes of this paper, we assume that a previous step identifies all model fragments in the domain knowledge that describe some aspect of the scenario at some level of detail. Falkenhainer and Forbus [3] show how this can be done.

Of these model fragments, the modeler must select those that are pertinent to the quantities of interest. This poses two fundamental issues:

- Identify those model fragments that describe relevant aspects of the scenario (the **scope** of the model).
- Identify those model fragments that are at an appropriate **level of detail.**

Only model fragments that describe relevant aspects of the scenario (i.e., are within the chosen scope) at an appropriate level of detail should be included in the model.

2.1 Selecting Scope

The modeling algorithm of Falkenhainer and Forbus [3], which we refer to as FF, provides one of the best approaches to scope selection. FF chooses the scope of the model using three sources of knowledge:

- Each model fragment specifies the objects it models.
- Objects in the scenario must be decomposed into a structural partonomic hierarchy. FF assumes that each object in this hierarchy can be analyzed independent of other objects. This assumption holds to a large degree in many engineered devices, where components are designed to have specified functions largely independent of their context.
- The question specifies the quantities of interest.[3]

FF chooses the scope of the model by determining the relevant objects in the scenario. It identifies a minimal set of model fragments that could provide a model of the quantities of interest. Those objects modeled by these model fragments are relevant. To ensure that interactions among these objects are modeled, FF chooses a "minimal covering system," the lowest object down the partonomic hierarchy that subsumes the relevant objects; this object and its subobjects (down to the level of the relevant objects) are relevant. Any model fragment that models one of these relevant objects is within the scope of the model.[4] Importantly, no subsequent steps in the modeling algorithm can expand this scope, and model fragments that only model objects outside this scope cannot be included in the model.

There are two limitations in this approach. First, the assumption that scenario objects can be decomposed into a single partonomic hierarchy of objects that can be analyzed independently does not apply to natural systems such as plants and animals, and often it must be relaxed in engineering as well. Plants and animals can be organized into nearly-decomposable systems (e.g., the respiratory system and the circulatory system), although these are based on functional rather than partonomic breakdowns. However, system decompositions in natural systems are not rigid; they depend on operating conditions and the time scale and purposes of analysis. Even in engineering, designed

[3] Falkenhainer and Forbus allow questions to specify objects and relations of interest as well, but we focus on the guidance it receives from quantities of interest for comparison with our modeling approach.

[4] Other modeling decisions in FF might be viewed as further restricting the scope of the model, but these are irrelevant to our discussion.

system boundaries must be relaxed when considering faults or unintended interactions. Determining the appropriate system decomposition of the scenario is an important issue for the modeler; it cannot be specified a priori [2].

The second limitation of this approach is that it may fail to include important aspects of the scenario in the model. Consider the question "How would a decreasing amount of soil water affect plant size and growth rate?" The quantities of interest are plant size and growth rate. Since in general it is possible to reason about growth independent of the soil, the minimal covering system will exclude the soil, so soil water amount cannot be included in the model. Yet clearly a model for this question must include the interaction of soil water amount and growth processes. In this example, the inadequacy could be overcome if a query interpreter included soil water amount as a quantity of interest, forcing its use in selecting the minimal covering system; however, in general, including all quantities appearing in behavioral conditions (**given quantities**) as quantities of interest is a bad idea. The user may provide a complicated scenario with many irrelevant conditions. The modeler should not require users to know which aspects of their hypothetical scenario will affect the quantities of interest, and the query interpreter cannot easily determine which given quantities are relevant. If all given quantities are used in selecting the minimal covering system, an irrelevant given quantity may require consideration of an object far removed from the appropriate scope, resulting in an excessively large scope and hence an overly complex model. Thus, using all given quantities in selection of the scope of the model may result in unnecessarily large models, while ignoring them may result in an inadequate model.

FF cannot determine which given quantities are relevant because it is fundamentally insensitive to the ways in which behavioral conditions in the scenario interact with quantities of interest. Intuitively, a model should include the important **interaction paths** between given quantities and quantities of interest, where an interaction path is a set of functional and differential relations that connects two quantities, describing how they affect each other. While FF has contributed many useful ideas for automated modeling, it must be extended to exploit knowledge of interaction paths relevant to the question.

Several other researchers use interaction paths to guide modeling. Williams's "critical abstraction" method [16] constructs a model of how a device works with respect to a set of queries (which can be viewed as quantities of interest for purposes of this discussion). Williams assumes that the independent (exogenous) quantities of the device are given. Williams also assumes a most-detailed model of the device is given. Williams's method computes the "causal support structure" of the quantities of interest by tracing causally forward through the functional and differential relations of the detailed model from the independent quantities to the quantities of interest; those model fragments providing this support structure are deemed relevant and provide the scope of the model. Nayak et al. [9, 10] also provide a method for constructing a model of how a device works. In place of the queries used by Williams, their method requires an "expected behavior," a pair of quantities such that the device achieves a causal connection from the first to the second. Their objective is to find a model that provides such a connection.

Neither of these methods allows enough freedom in selecting the scope of the model; these methods are unable to select a subsystem of the device as the scope. For example, consider the question "How would decreasing turgor pressure affect the rates of potassium diffusion and active transport in the leaves?" It is not necessary to model the entire plant to answer this detailed question. Turgor pressure can be treated as an exogenous quantity and the potassium transport mechanisms in the plant can be examined in isolation. But Williams and Nayak et al. require the exogenous quantities of the plant as input; their methods cannot determine that turgor pressure, normally a quantity regulated by the plant, can be treated as the exogenous quantity for purposes of this question. It might seem possible to treat the given quantities of the question as the exogenous

quantities in the critical abstraction method, but that will not work because the method requires a complete set of exogenous quantities from which all other quantities can be determined. The methods of Williams and Nayak et al. assume that the exogenous quantities of the device are given and the objective is to provide a model of the whole device, whereas the ability to focus on a subsystem of a device, whose input quantities can be treated as exogenous for purposes of the question, is an important issue in selecting an appropriate scope for prediction questions.

2.2 Selecting Level of Detail

The modeler must determine the appropriate level of detail at which to describe each aspect of the scenario in the scope. This may involve many types of approximation and abstraction. In this paper, we focus on two types.

One important method of changing the level of detail in a model is by changing its accuracy [15]. Accuracy is a measure of how closely the predictions match reality. Although accuracy can be changed with a variety of mathematical approximations, we focus on approximation by ignoring insignificant influences. That is, a quantity or its rate of change may be a function of many other quantities; if the modeler can determine that some of these influences are insignificant for purposes of the question, the model can be simplified.

One useful criterion for significance is the time scale of analysis. For instance, diffusion of solutes through cell membranes is relatively slow. To model the changes in solute levels in a cell on a time scale of seconds, the influence of diffusion can be ignored because it cannot change the level significantly on that time scale. However, to model the changes on a time scale of hours, diffusion must be considered because it can cause significant change in that amount of time. Thus, the domain knowledge can associate an influence with the time scale at which it becomes significant.

Another related way of changing the level of detail in a model is by using **time scale abstraction**, a technique used in domains ranging from economics to engineering [5, 6, 7, 8, 13, 14]. Rather than ignoring influences that are significant only at time scales much slower than the time scale of interest, time scale abstraction allows the detailed dynamics of a process much faster than the time scale of interest to be abstracted out. In many dynamic systems, a set of detailed influences work to maintain an equilibrium functional relationship between quantities. If we are interested in the fast time scale and hence the transient dynamics, we must reason with the detailed influences. However, at a slower time scale we can simply model the behavior as an instantaneous functional relationship; that is, we assume that the system is always very close to equilibrium. For instance, at a detailed level, osmosis can be modeled by a set of influences that dynamically react to a gradient between the water potential in a cell and its environment and bring them into equilibrium. At a slower time scale, the model can ignore these dynamics and simply assert the equilibrium functional relationship (in this case, equality) between the water potentials. Because of the ubiquity of this type of abstraction in science and engineering, it is important to be able to represent these different levels of detail and the knowledge needed to choose among them.

There has been little work in choosing the level of detail in a model based on time scale considerations. Kuipers [8] shows how simulation can intermix models at different time scales, but he does not provide methods for constructing such models in response to questions. Iwasaki [6] presents a method that determines significant processes and appropriate time scale abstractions given the desired time scale and desired quantitative accuracy for each quantity of interest. Her method estimates the time scale at which each process will cause significant change and compares this to the specified time scale of interest; those processes acting at much faster time scales are modeled by their instantaneous functional relation, those processes acting at comparable time scales are modeled through their dynamics, and those processes acting at much slower time scales are treated

as insignificant. This is a useful method when the appropriate time scale of the model is given, but it will often be unreasonable to expect the person posing the prediction question (e.g., a student interacting with a tutoring system) to specify the time scale; rather, the modeling program should determine the time scale at which the behavioral conditions affect the quantities of interest.

2.2.1 Summary

Existing methods of selecting the scope of the model rely on assumptions inappropriate for the prediction task. Effective selection of scope requires a sensitivity to the interactions between behavioral conditions of the scenario and the quantities of interest. Identification of these interactions allows the modeler to determine how much of the scenario is relevant and which quantities can be treated as exogenous.

Existing approaches provide important ways of changing the level of detail of a model, but the criteria for making these decisions are too weak for the prediction task. Influences that are significant only at time scales slower than the chosen time scale of the model can be ignored. Detailed dynamics operating at time scales much faster than the chosen time scale can be modeled by time scale abstractions. However, existing methods require the user to specify the appropriate time scale of the model; instead, the modeler should determine the time scale at which the behavioral conditions affect the quantities of interest.

Our modeling algorithm, presented in the next section, identifies the interaction paths that relate behavioral conditions and quantities of interest. It uses this knowledge to select an appropriate scope for the model, and it chooses time scale abstractions and determines insignificant influences by selecting the time scale at which the behavioral conditions of the scenario interact with the quantities of interest.

3 Using Interaction Paths to Guide Approximation and Abstraction

3.1 Time Scale Conditions

Our modeling methods currently support two types of variation in level of detail that we have found to be among the most important. First, the accuracy of the model can be varied by ignoring insignificant influences whose significance depends on time scale. Second, our methods support time scale abstraction. Supporting these modeling dimensions requires two types of knowledge: the domain knowledge must include an ordered set of **time scales**, and model fragments must be qualified with **time scale conditions**.

Our approach to selecting level of detail requires that the domain knowledge include a set of the important time scales at which domain processes work. Each time scale represents an order of magnitude of time, such as seconds, minutes, or days. We further require a total ordering over these time scales such that a time scale t_1 is less than a time scale t_2 if t_1 is orders of magnitude smaller (faster). For instance, here is an ordering of several important time scales in plant physiology:

$$\text{cellular-water-balance} < \text{solute-balance} < \text{plant-water-balance} < \text{growth}$$

Cellular-water-balance represents the order of magnitude of time at which osmosis works to achieve an equilibrium between the water potentials inside and outside a cell; this is roughly seconds. Solute-balance represents the time scale at which solutes diffuse through the plant to achieve equilibriums; solutes diffuse much more slowly than water, with significant changes occurring on the order of minutes. Plant-water-balance represents the time scale at which the plant is working to balance

the water level in its tissues by taking water up from the soil, transporting it throughout the plant, and regulating transpiration (evaporation) from the leaves; significant changes occur on the order of hours. Finally, significant results of growth of the plant occur on the order of days. The domain knowledge must include such an ordering over all the important time scales in the domain.

Each model fragment in the domain knowledge can include in its operating conditions a set of time scale conditions. A time scale condition delimits the time scales of analysis for which the model fragment can be used. For instance, a model fragment with the time scale condition (> :time-scale plant-water-balance) can only be used in models that address time scales larger (slower) than plant-water-balance. As in the time scale ordering, the inequality relations in time scale conditions are order of magnitude comparators.

To illustrate the use of time scale conditions, consider the example question "How would a decreasing amount of soil water affect plant size and growth rate?" Figure 1 shows all the model fragments for this scenario. These are shown as an **influence graph** in which nodes are quantities and arcs are the influences of the model fragments. Operating conditions are not shown. Influences are labeled with their type (e.g., Q^+ or I^+) and with their time scale conditions.[5] A dotted line associates a time scale abstraction with its underlying detailed influences.

This figure demonstrates the three ways in which we use time scale conditions:

- Differential influences (I^+, I^-) cause changes in their influenced quantity. We qualify these influences with the minimum time scale needed for them to cause significant change. For instance, in the example, the influences of water uptake and transpiration on apoplast water amount are insignificant at time scales much faster than plant water balance, the effects of the ABA and K^+ processes are insignificant at time scales much faster than solute balance, the effects of osmosis are insignificant at time scales much faster than cellular water balance, and the effect of growth on the size of the plant is insignificant at time scales much faster than growth.

- Some functional influences represent time scale abstractions of dynamic processes. These influences are qualified with the minimum time scale at which a change in the independent quantity can be viewed as instantaneously causing a change in the dependent quantity (i.e., a time scale at least as large as the time it takes the underlying processes to reach equilibrium); at time scales faster than this, the functional relationship is invalid. For instance, the functional influence asserting the equilibrium functional relationship between the amounts of water in the apoplast and symplast only holds if the time scale is at least that of solute balance, since the dynamics cannot be ignored at faster time scales.

- We qualify an equilibrium process (influences on and from its rate quantity) with the time scale at which it reaches equilibrium; at this or slower time scales, the transient dynamics will have settled out and the time scale abstraction should be used. For instance, it only makes sense to reason about the dynamics of osmosis between the apoplast and symplast at time scales much faster than that of solute balance.

[5] In Qualitative Process Theory [4], there are two types of functional influences: Q^+ and Q^-. y Q^+ x means that y is a function of x (and perhaps other quantities) and the partial derivative is positive (i.e., if the other quantities on which y depends are held constant, y changes in the same direction as x); Q^- is similar but specifies a negative partial derivative. In the influence graph, the influence points from the independent to the dependent quantity. There are two types of differential influences: I^+ and I^-. These are used to specify that a process increases or decreases a quantity, respectively. For instance, the amount of water in a bucket is I^+ to the rate of inflow from a faucet and I^- to the rate of outflow through a hole. In the influence graph, the influence points from the process rate to the influenced quantity.

Those functional influences in the influence graph without time scale conditions are always treated as instantaneous functional relations.

Time scale conditions provide the criteria for selecting among alternative levels of detail in our representation. They ensure mutual exclusion of a time scale abstraction and its underlying details. They also provide the criteria for assessing significance of influences. Once a time scale for the model is chosen, these conditions determine the appropriate level of model accuracy and determine which processes should be modeled by their dynamics and which should be modeled through time scale abstractions. The next section shows how these conditions also provide the criteria for choosing the appropriate time scale of the model.

3.2 Selecting Level of Detail

Our approach to time scale selection ensures that the model includes the model fragments that relate the quantities of interest to behavioral conditions. This requires identifying the ways in which quantities of interest interact with given quantities (quantities in behavioral conditions).

To represent the way in which two quantities interact (i.e., through particular functional and differential influences), we first define an **interaction graph** as the undirected graph corresponding to an influence graph. Next, we define an **interaction path** between two quantities as a path in the interaction graph that may cross the same node (quantity) more than once but that may not cross the same edge (**interaction**) more than once. This definition allows loops in an interaction path, which is important for inclusion of feedback loops. Although relations in a model fragment may have causal direction, interactions are undirected. As such, interaction paths represent potential inference paths, but they need not correspond to causal paths. This is important because inferences need not follow the direction of causality; for example, if x is a function of y and z, and the values of x and y are known, it may be possible to infer the value of z. Two quantities are related by an interaction path when information about the values of one constrains the values of the other.

Given the interaction graph for the model fragments of a question, our modeling algorithm selects a time scale with the following steps:

1. Identify all the interaction paths linking any given quantity to any quantity of interest. For efficiency, avoid generating paths having incompatible interactions (i.e., interactions with mutually inconsistent time scale conditions).[6]

2. Associate each of these interaction paths with the set of time scales for which it is valid by finding all domain knowledge time scales that satisfy the conjunction of time scale conditions of interactions on the path.

3. Associate each quantity of interest with the time scales at which it interacts with given quantities by taking the union of the time scales of its interaction paths.

4. Form the intersection of the sets of time scales for the quantities of interest to find the time scales for which *every* quantity of interest is connected to *some* given quantities.

5. If the resulting set contains more than one possible time scale, choose among them using preference criteria. Our algorithm currently chooses the largest (slowest) time scale under the assumption that long-term trends are more important than transient behavior and that slower time scales correspond to time scale abstractions and hence simpler models.

[6] Ideally, the algorithm should also avoid generating paths containing interactions with mutually inconsistent operating conditions.

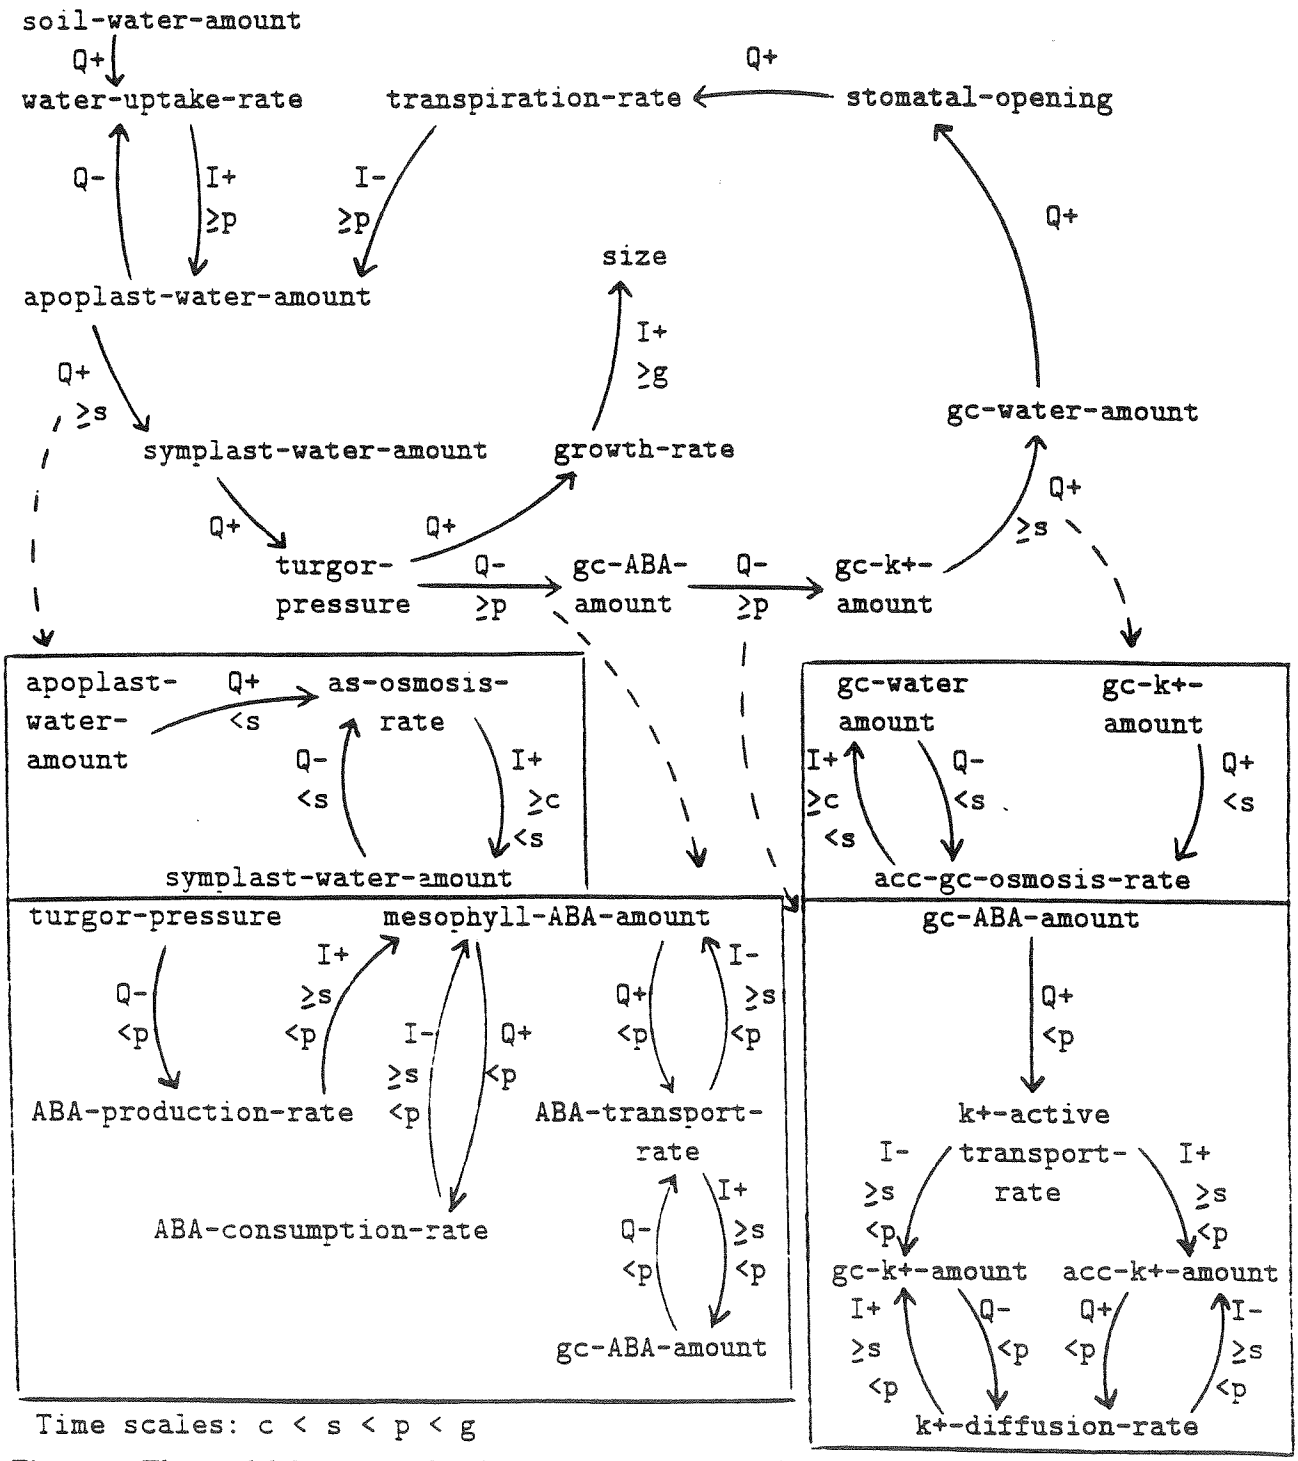

Figure 1: The model fragments for the example scenario. These are shown as a graph in which the nodes are quantities and the arcs are the influences of the model fragments. Operating conditions are not shown. Influences are labeled with their time scale conditions. A dotted line associates a time scale abstraction with its underlying detailed influences.

For example, suppose the prediction question is "How would a decreasing amount of soil water affect plant size and growth rate?" The quantities of interest in this question are size and growth rate and the given quantity is soil water amount. The only time scale at which there is a connection between soil water amount and size is that of growth, so this time scale is selected for the model.

Alternatively, suppose the prediction question is "How would decreasing turgor pressure affect the rates of K^+ diffusion and K^+ active transport?" The quantities of interest in this question are K^+ diffusion rate and K^+ active transport rate and the given quantity is turgor pressure. Figure 1 shows that connecting the given quantity to each quantity of interest requires a time scale at least as slow as solute balance for the ABA and K^+ processes, and reasoning about the dynamics of the K^+ processes requires a time scale faster than plant water balance, so the time scale of solute balance is selected for this model.

This algorithm selects a single time scale at which to model the scenario. This is sufficient for a large class of questions, but some questions require reasoning at different time scales and combining the results. We have not explored this issue; our program recognizes that no single time scale will suffice and informs the user. For relevant work on this issue, see Kuipers [8] or Iwasaki [6].

In our representation, the chosen time scale determines the appropriate level of detail of the model. Those model fragments whose time scale conditions are satisfied by the chosen time scale are all at the appropriate level of detail. The **time scale interaction graph** is the interaction graph corresponding to these model fragments. This interaction graph is used to select the model fragments to include in the model.

3.3 Selecting Scope

While the influences in the time scale interaction graph are all at an appropriate level of detail for the question, many of them will be irrelevant to the question. The most important criterion of relevance is that the model should include those influences that lie along interaction paths that relate quantities of interest to given quantities. However, this criterion alone is insufficient.

In addition to including interaction paths linking given quantities and quantities of interest, the modeler must also reason about the effects of given quantities on the operating conditions of model fragments in the model. Thus far we have ignored operating conditions. However, at each stage of simulation (or other analysis of the model), the analysis module must determine whether the operating conditions of each model fragment in the model are satisfied. To ensure that the model contains the information necessary to make this assessment, the model must include any interaction paths that connect a given quantity to a quantity in the operating conditions of a model fragment in the model. For example, if one of the model fragments relevant to the growth question has in its operating conditions a quantity q that is influenced by soil water amount, this influence must be included in the model; otherwise, as soil water amount changes, the model would erroneously predict that q (and the validity of the model fragment) remains unchanged.

The final criterion for relevance is the following: including one influence in the model may require including others. For instance, influences in Qualitative Process Theory represent pieces of an equation. If the model includes any functional influences on a quantity, it must include all functional influences on that quantity whose time scale conditions are satisfied by the selected time scale (similarly for differential influences). For example, a model cannot include the effect of water uptake on apoplast water without including the effect of transpiration, since these influences together determine changes in apoplast water. When one influence is relevant if and only if another influence is relevant, we will say that they are **relevance equivalent**.

We combine all these considerations into a definition of **relevance** of influences as follows:

1. An influence is relevant if it lies on an interaction path from a given quantity to a quantity of interest in the time scale interaction graph.

2. An influence in the time scale interaction graph is relevant if it is relevance equivalent to a relevant influence.

3. An influence is relevant if it lies on an interaction path (in the time scale interaction graph) from a given quantity to a quantity in the operating conditions of a model fragment containing a relevant influence.

4. No other influences are relevant.

Every model fragment containing a relevant influence is included in the model.

Figures 2 and 3 show the relevant influences for the prediction questions about growth and potassium transport. In these examples, all relevant influences come from clause 1 of the relevance definition. Figure 3, in particular, shows that the selection of scope eliminates irrelevant influences from the time scale interaction graph. The influences of soil water amount on water uptake rate, symplast water amount on turgor pressure, and stomatal opening on transpiration rate are all valid at the time scale of growth, but none are relevant to the question.

3.4 Current Status

The modeling algorithm presented in this section has been implemented and run on the examples in this section as well as other small examples. On the examples in this section, the program takes a few minutes on an Explorer I Lisp Machine. We have simulated each of the example models in this section using the Qualitative Process Compiler (QPC) [1]; this required us to provide, for each model, additional constraints on the ranges of quantities (e.g., amounts must be non-negative) and initial conditions that we would expect a query interpreter to select (based on defaults). QPC successfully simulated these examples, producing a small set of reasonable behaviors.[7]

4 Discussion

Our modeling algorithm extends existing modeling methods by exploiting the interaction paths that relate quantities of interest to the behavioral conditions of the scenario. The algorithm identifies a time scale at which quantities of interest and given quantities interact. This time scale determines which influences are insignificant, and it determines which processes can be modeled with time scale abstractions. Our relevance criteria determine a scope that includes those influences that relate quantities of interest to given quantities while excluding other influences. The relevance criteria also determine the quantities in the scenario that can be treated as exogenous for purposes of the question. This modeling algorithm demonstrates the benefits provided by identification of interaction paths.

Our method of selecting a time scale can be easily generalized to select among other types of modeling alternatives. Many modeling decisions, such as whether to model a process or its subprocesses, cannot be made based on the time scale of analysis. In another paper [12], we show how the methods described in this paper can be used to select among any modeling alternatives that can be encoded in the assumption class representation of Falkenhainer and Forbus [3]. We are currently applying these methods to automatic construction of models from the Botany Knowledge Base [11], which includes over 200 processes described at many different levels of detail.

[7]The experiments with QPC were run with the help of Adam Farquhar.

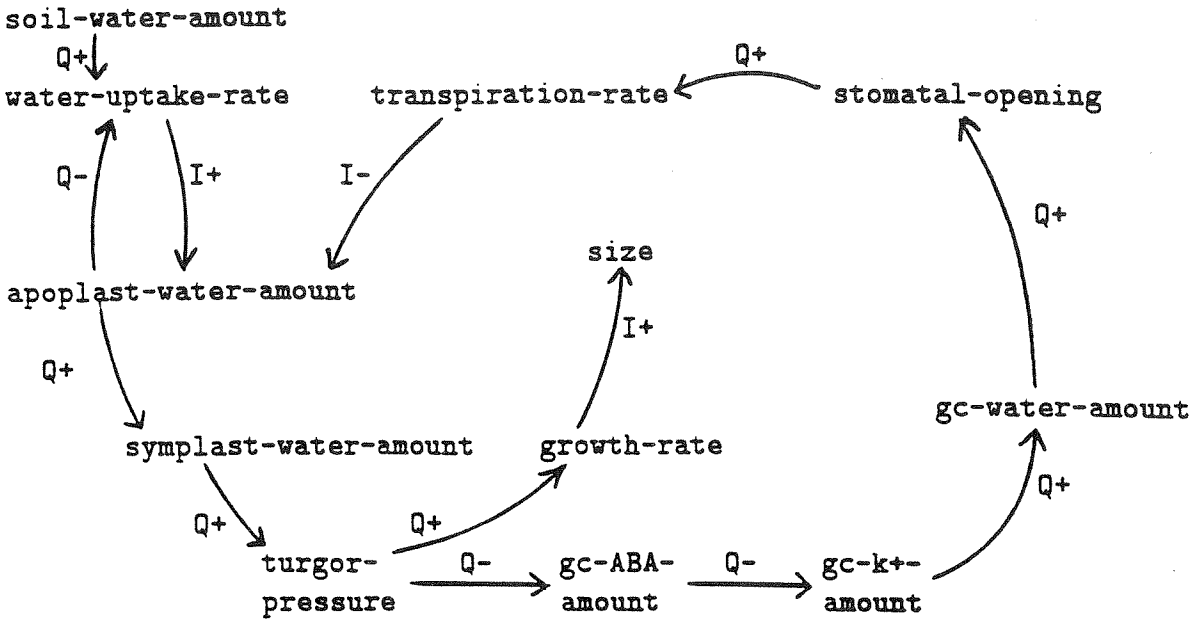

Figure 2: The relevant influences for the prediction question "How would a decreasing amount of soil water affect plant size and growth rate?"

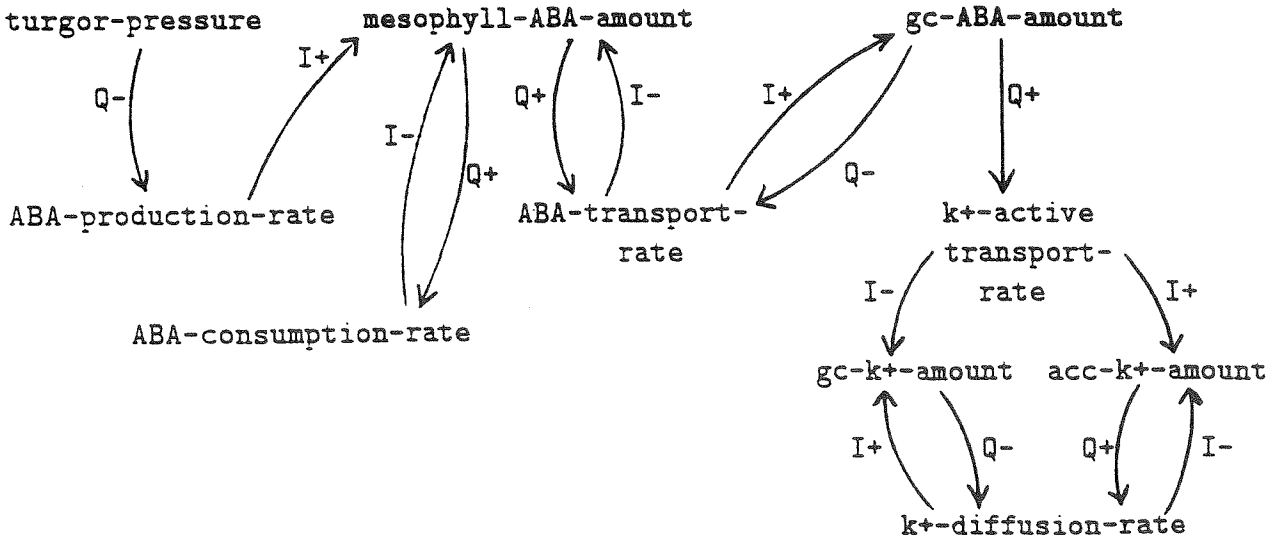

Figure 3: The relevant influences for the prediction question "How would decreasing turgor pressure affect the rates of K^+ diffusion and K^+ active transport?"

There are two important dimensions of research in automated modeling: modeling criteria and search efficiency. Research on modeling criteria attacks the question of what constitutes a good model; such research identifies the constraints and preferences a model should satisfy. Our research has been primarily aimed at this issue. However, within the space of models defined by these constraints and preferences, a modeler must search for the best model in an efficient manner. Our current, naive algorithm searches exhaustively through interaction paths. This approach will be unnecessarily expensive in a large knowledge base, so we are investigating ways of focusing the search for relevant interaction paths.

Finally, the modeling criteria discussed in this paper must be integrated with other modeling criteria. An effective modeler must exploit many sources of knowledge, including knowledge of the user's background, the context of the question, and feedback from analysis that can identify modeling errors that would be difficult to anticipate. The work described in this paper has focused on the important guidance that interaction paths can provide for modeling decisions.

5 Acknowledgements

Many people contributed to this paper. Art Souther helped with the examples. Adam Farquhar ran the experiments with QPC. Adam, Giorgio Brajnik, Brian Falkenhainer, Ben Kuipers, Ray Mooney, Gordon Novak, Art Souther, Liane Acker, James Lester, Ken Murray, and Rich Mallory gave excellent comments on earlier drafts. Rich's dedication to helping us make the original submission deadline was beyond generosity.

References

[1] Crawford, J., A. Farquhar, and B. Kuipers, "QPC: A Compiler from Physical Models into Qualitative Differential Equations," *Proc. of AAAI-90*, Los Altos, CA: Morgan Kaufmann, pp. 365-372, 1990.

[2] Davis, R., "Diagnostic Reasoning Based on Structure and Behavior," *Artificial Intelligence* **24**, pp. 347-410, 1984.

[3] Falkenhainer, B., and K.D. Forbus, "Compositional Modeling: Finding the Right Model for the Job," *Artificial Intelligence* **51**, pp. 95-143, 1991.

[4] Forbus, K.D., "Qualitative Process Theory," *Artificial Intelligence* **24**, pp. 85-168, 1984. Also appears in *Readings in Qualitative Reasoning about Physical Systems* (Eds. D.S. Weld and J. de Kleer), San Mateo, CA: Morgan Kaufmann, 1990, pp. 178-219.

[5] Iwasaki, Y., and I. Bhandari, "Formal Basis for Commonsense Abstraction of Dynamic Systems," *Proc. of AAAI-88*, San Mateo, CA: Morgan Kaufmann, pp. 307-312, 1988.

[6] Iwasaki, Y., "Reasoning with Multiple Abstraction Models," in Working Notes of the Workshop on Automatic Generation of Approximations and Abstractions (Eds. T. Ellman, R. Keller, and J. Mostow), 1990, pp. 122-134 (available through Tom Ellman at Rutgers University).

[7] Kokotovic, P.V., R.E. O'Malley, Jr., and P. Sannuti, "Singular Perturbations and Order Reduction in Control Theory – An Overview," *Automatica* **12**, pp. 123-132, 1976.

[8] Kuipers, B., "Abstraction by Time Scale in Qualitative Simulation," *Proc. of AAAI-87*, Seattle, WA, pp. 621-625, 1987.

[9] Nayak, P.P., L. Joskowicz, and S. Addanki, "Automated Model Selection Using Context-Dependent Behaviors," Working Papers of the Fifth International Workshop on Qualitative Reasoning about Physical Systems, pp. 10-24, Austin, Tx, 1991.

[10] Nayak, P. P., "Causal Approximations," to appear in AAAI-92.

[11] Porter, B., 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.

[12] Rickel, J.W., "Automated Modeling for Answering Prediction Questions: Exploiting Interaction Paths," thesis proposal, 1992.

[13] Saksena, V.R., J. O'Reilly, and P.V. Kokotovic, "Singular Perturbations and Time-Scale Methods in Control Theory: Survey 1976-1983," *Automatica* **20**(3), pp. 273-293, 1984.

[14] Simon, H.A., and A. Ando, "Aggregation of Variables in Dynamic Systems," *Econometrica* **29**, 1961.

[15] Weld, D.S., "Reasoning about Model Accuracy," to appear in *Artificial Intelligence*.

[16] Williams, B.C., "Critical Abstraction: Generating Simplest Models for Causal Explanation," Working Papers of the Fifth International Workshop on Qualitative Reasoning about Physical Systems, pp. 77-92, Austin, Tx, 1991.