Artificial Intelligence 59 (1993) 125–132
Elsevier

ARTINT 994

Reasoning with qualitative models *

Benjamin J. Kuipers

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

1. Introduction

Qualitative reasoning about physical systems has become one of the most productive areas in AI in recent years, due in part to the 1984 special issue of *Artificial Intelligence* on that topic. My contribution to that issue was a paper entitled "Commonsense reasoning about causality: deriving behavior from structure" [9]. From my perspective, that paper laid out a research program that has continued to be productive to this day, and promises to continue well into the future.

After establishing a framework for qualitative reasoning, the primary technical contribution of the paper was a simple, clear representation for qualitative structure and behavior, abstracted from ordinary differential equations. My subsequent *Artificial Intelligence* paper, "Qualitative simulation" [10], made that abstraction relation precise, presented the vastly improved QSIM algorithm for qualitative simulation, and used the abstraction relation to prove the soundness and incompleteness of QSIM. I will discuss developments in qualitative simulation in my retrospective on that paper [12], and concentrate here on the larger issue of reasoning with qualitative models.

Correspondence to: B.J. Kuipers, Computer Sciences Department, University of Texas at Austin, Austin, TX 78712, USA. E-mail: kuipers@cs.utexas.edu.
*This work has taken place in the Qualitative Reasoning Group at the Artificial Intelligence Laboratory, The University of Texas at Austin. Research of the Qualitative Reasoning Group is supported in part by NSF grants IRI-8905494 and IRI-8904454, by NASA grant NAG 2-507, and by the Texas Advanced Research Program under grant no. 003658-175.

2. Context

In 1978, thanks to Peter Szolovits, I started working on the problem of causal reasoning about physiological mechanisms by expert physicians. I was initially attracted to Rieger and Grinberg's causal link models, presented at IJCAI-77 [14], and I began working with Jerome P. Kassirer, an eminent nephrologist at Tufts Medical School. We applied Newell and Simon's protocol analysis methods to interviews with expert physicians, to extract clues about the cognitive representation of physiological mechanisms. This protocol analysis provided one of two essential constraints on the design of the qualitative representation: empirically, it should account for the behavior of the human subjects, and computationally, it should be capable of deriving the observed conclusions.

From analyzing the protocols, it quickly became clear that there was a cognitively meaningful distinction between the time-independent *structure* of a mechanism, and its time-dependent *behavior*. This distinction did not seem to map clearly onto causal networks, but fit better with Johan de Kleer's work on qualitative envisionment. Starting with this foundation, my applications and my intuitions led me away from the quasi-static equilibrium assumption, and toward monotonic function constraints and nonzero landmarks: essentially the QSIM representation, although the name came later. Kassirer and I published our protocol analysis and its explanation in terms of the qualitative representation in *Cognitive Science* [13], where it appeared at about the same time as the *Artificial Intelligence* special issue.

Meanwhile, Johan de Kleer, Ken Forbus, Dan Weld, Brian Williams and others were also developing and extending methods and applications for qualitative reasoning. While there were many fruitful discussions among the early participants in this research community, differences in outlook, assumptions, and notation often made it difficult for us to communicate clearly. Although it is sometimes said that differences in notation have acted as a barrier to unifying the different perspectives in qualitative reasoning, it now seems clear that the notational differences reflect genuine semantic distinctions among types of knowledge used in different types of reasoning: for example, model building versus model simulation, and dynamic simulation versus comparative statics.

3. Contributions and applications

An enormous amount of subsequent work has been inspired by the papers in the 1984 special issue, including international workshops on qualitative reasoning, model-based reasoning, and the principles of diagnosis. I will discuss developments specifically related to the qualitative simulation and

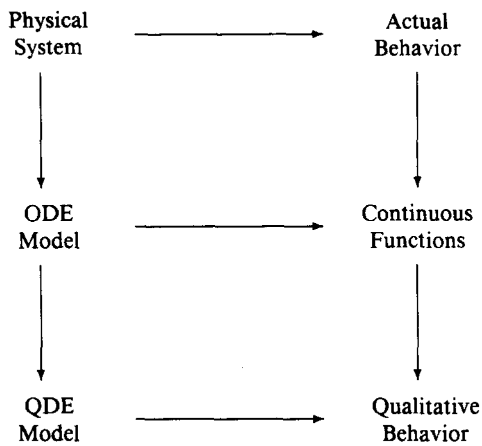

Fig. 1. All models are abstractions of the world. Qualitative models are related to ordinary differential equations, but are more expressive of incomplete knowledge.

the QDE representation in the retrospective [12] on the paper "Qualitative simulation", and focus here on applications of qualitative reasoning to diagnosis, monitoring, and design, and on several conceptual schemes that have been helpful to me in clarifying and factoring the problems and applications of qualitative reasoning.

3.1. Abstraction relations

Qualitative structure and behavior can be most clearly understood and analyzed as abstractions of ordinary differential equations and their solutions (Fig. 1). (Both types of models are, of course, abstractions of the physical world.) This abstraction relation (hypothesized in this paper, and proved rigorously in [10]) has been critical to making qualitative simulation mathematically tractable, and to communicating it successfully with the engineering and mathematics communities. It legitimizes the term *qualitative differential equation* or QDE for the qualitative structural description.

3.2. Structure, behavior, function, and design

There is an important distinction between three types of descriptions of a mechanism, and how they depend on each other:

structure → *behavior* → *function.*

A clear distinction among these is particularly important because function and behavior are often confused. The essence of a functional description is teleology or purpose, and hence the relation of the structure and behavior of a mechanism to its larger context. Now that the roles of structural and

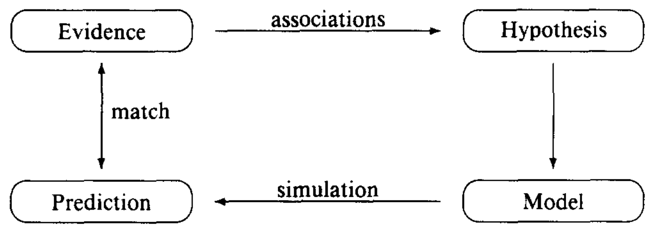

Fig. 2. Diagnosis viewed as a generate-and-test process, proposing models capable of explaining observations.

behavioral descriptions are better understood (Fig. 1), it has become clearer how to represent and reason with *functional* description of a device.

A motivating example from the 1984 paper is the assertion, "The function of the pressure release valve is to prevent explosions." Semantic analysis reveals that the verb *prevent* represents a relation between an element of structure (the pressure release valve) and a possible behavior (an explosion). However, simulation of the working device model does not include an explosion among its possible behaviors. The explosion referred to appears "upstream" in the *design process* for the device, prior to the addition of the pressure release valve.

David Franke [6,7] has defined a precise semantics for teleological relations such as *guarantee* and *prevent*, in terms of the incremental transformations to the device model that take place during design and a branching-time temporal logic over the set of behaviors predicted from each qualitative model. His system is associated with a CAD system, in order to acquire teleological descriptions when they are most available: during the design process. The teleological descriptions are then used to index design transformations for reuse, and to propose plausible candidates during diagnosis.

3.3. Diagnosis and monitoring

Diagnosis can be defined in several ways. The modeling and simulation perspective makes it natural to view diagnosis as an instance of the more general problem of causal explanation: a set of observations are explained by a set of general laws and specific initial conditions such that the observations can be predicted as consequences of the laws and initial conditions (Kuipers [11], Simmons and Davis [15]) (Fig. 2).

According to this view, the goal of diagnostic reasoning is to find a useful predictive model of a possibly faulty device, given a description of the working device and knowledge of possible fault modes, either of components or of the device as a whole. The task thus blends smoothly into monitoring, where the task is to maintain an accurate model of a mechanism and its state, even while faults occur and are repaired. This position contrasts with the "constraint suspension" view of diagnosis, where the task is to identify minimal sets of components whose correct behavior is inconsistent with observations. Constraint suspension seems most appropriate for devices where faults are relatively isolated, and where it is practical to shut the system down to replace components. Large-scale systems such as chemical plants, space vehicles, and the human body have many self-regulatory systems, and are expected to continue functioning even in the presence of numerous faults. These intuitions have led to the development of the MIMIC approach to monitoring and diagnosis (Dvorak and Kuipers [3,4]).

The basic idea behind MIMIC is very simple: track the observed state of a system against predictions derived from one or more models; discrepancies are used to refute some current models and suggest new ones. However, the success of this approach depends critically on the ability to cover a realistically large set of hypotheses with a tractably finite set of models. Traditional ODE models contain many assumptions of specific functional forms and numerical parameter values, often going beyond the knowledge available, particularly for fault models. Qualitative models can cover a wide range of possible ODEs with a single QDE, and can refine the qualitative predictions using numerical information when it is available.

3.4. Modeling and simulation

There is an important distinction between the tasks of model building and model simulation:

- *Model building*: starts with a description of a physical situation and builds an appropriate simplified model, in this case a QDE.
- *Model simulation*: starts with a model and predicts the possible behaviors consistent with the model.

The research issues involved in model building and model simulation are quite distinct, and the two tasks interact via the QDE representation. Therefore, as a research strategy, this factoring of the problem makes it possible for a qualitative reasoner to benefit from independent advances in the two areas.

Much of the work that my students and I have done within this framework has concentrated on the representation and tractable simulation of QDE models, with the view that QDE models could be constructed by a variety of different methods, including those pioneered in the work of de Kleer and Brown [2] and Forbus [5]. It has seemed to me that both of these pieces of work are clarified by separating their contributions to model building and model simulation. Accordingly, we have built two compilers that produce QDE models for simulation by QSIM.

- *CC* (Franke and Dvorak [8]) builds a QDE from a description of a physical system in terms of explicit connections among instances of components defined in a component library.
- *QPC* (Crawford et al. [1]) builds QDE models after the manner of qualitative process theory by identifying sets of active views and processes in a view library, and applying the Closed World Assumption to transform influences into constraints.

The two approaches to model building differ in the nature of the knowledge supplied by the modeler, and in the way the Closed World Assumption is applied. Specifically, by describing a device with a component-connection (CC) model, the *modeler* asserts that all relevant interactions between the components take place via explicit connections. In QPC, by contrast, the *system* is responsible for determining the set of relevant interactions and deciding when to apply the Closed World Assumption.

In the simplest, linear view, the modeler builds a model which the simulator uses to predict behaviors. More realistically, the process iterates, with the model builder responding to feedback from the simulator about the implications of the model and from the world about the success of the model to explain the physical phenomenon of interest.

Figure 3 illustrates the main linear path through model building and model simulation, separating the process into weakly interacting modules. Each module draws on certain information (e.g., the component library or asserted quantitative bounds), may make certain assumptions as needed (e.g., smoothness of monotonic functions or the CWA), and provides certain guarantees (i.e., that its conclusion follows soundly from its premises). One perspective on the field of qualitative reasoning is that its goal is to specify the modules and their intermediate representations so as to make this framework real.

4. Open problems

There are several "paths not taken" whose beginnings were visible in the 1984 paper, and which still seem very fruitful.

- The title of the 1984 paper refers to "causal" reasoning about physical mechanisms. In fact, we developed modeling and simulation methods based on constraint propagation and satisfaction. Any causality in these models is imposed from the outside by the viewer. The problem of describing and reasoning about causality seems important and has received a great deal of attention, but on careful examination both its content and its value are surprisingly hard to pin down, at least with the tools now available.

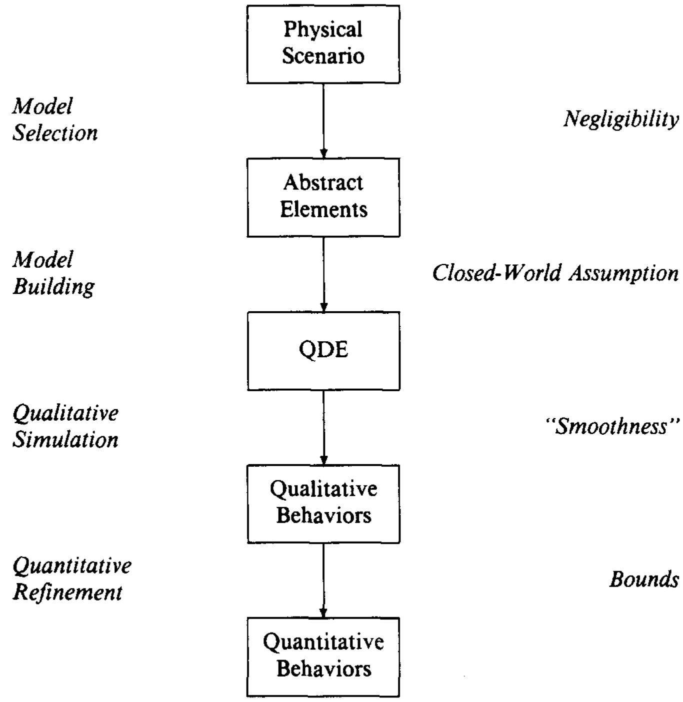

Fig. 3. Each step in modeling a physical system and predicting its behavior requires particular types of assumptions. Using qualitative models makes this set of assumptions tractable.

- A qualitative differential equation is a set of constraints, from which other constraints and equations can be derived by algebraic manipulation. One point of the 1984 paper was that certain conclusions that were intractable in the given model could be derived easily from an algebraically simplified abstraction of the model, automatically derived using the set of algebraic transformation rules in Appendix D. A good general-purpose algebraic manipulation utility applicable to QDE models would make it possible to search for Lyapunov functions or to symbolically evaluate the sign of a discriminant expression, and thus improve the power of qualitative simulation. Williams [16] made progress toward this goal, by clarifying the relationship between qualitative and real algebras, and demonstrating the use of MINIMA, an algebraic reasoner for qualitative models based on Macsyma.

References

[1] J.M. Crawford, A. Farquhar and B.J. Kuipers, QPC: a compiler from physical models into qualitative differential equations, in: *Proceedings AAAI-90*, Boston, MA (1990).

[2] J. de Kleer and J.S. Brown, A qualitative physics based on confluences, *Artif. Intell.* **24** (1984) 7–83.

[3] D. Dvorak and B.J. Kuipers, Model-based monitoring of dynamic systems, in: *Proceedings IJCAI-89*, Detroit, MI (1989).

[4] D. Dvorak and B.J. Kuipers, Process monitoring and diagnosis: a model-based approach, *IEEE Expert* **6** (3) (1991) 67–74.

[5] K.D. Forbus, Qualitative process theory, *Artif. Intell.* **24** (1984) 85–168.

[6] D.W. Franke, Representing and acquiring teleological descriptions, in: *Proceedings IJCAI-89, Model-Based Reasoning Workshop*, Detroit, MI (1989).

[7] D.W. Franke, Deriving and using descriptions of purpose, *IEEE Expert* **6** (1991) 41–47.

[8] D.W. Franke and D. Dvorak, Component-connection models, in: *Proceedings IJCAI-89, Model-Based Reasoning Workshop*, Detroit, MI (1989).

[9] B.J. Kuipers, Commonsense reasoning about causality: deriving behavior from structure, *Artif. Intell.* **24** (1984) 169–203.

[10] B.J. Kuipers, Qualitative simulation, *Artif. Intell.* **29** (1986) 289–338.

[11] B.J. Kuipers, Qualitative simulation as causal explanation, *IEEE Trans. Syst. Man Cybern.* **17** (3) (1987) 432–444.

[12] B.J. Kuipers, Qualitative simulation: then and now, *Artif. Intell.* **59** (1993) 133–140 (this volume).

[13] B.J. Kuipers and J.P. Kassirer, Causal reasoning in medicine: analysis of a protocol, *Cogn. Sci.* **8** (1984) 363–385.

[14] C. Rieger and M. Grinberg, The declarative representation and procedural simulation of causality in physical mechanisms, in: *Proceedings IJCAI-77*, Cambridge, MA (1977) 250–256.

[15] R. Simmons and R. Davis, Generate, test and debug: combining associational rules and causal models, in: *Proceedings IJCAI-87*, Milan, Italy (1987).

[16] B.C. Williams, MINIMA: a symbolic approach to qualitative algebraic reasoning, in: *Proceedings AAAI-88*, St. Paul, MN (1988) 264–269.