Artificial Intelligence 59 (1993) 115–123
Elsevier

ARTINT 991

Qualitative process theory: twelve years after

Kenneth D. Forbus

Qualitative Reasoning Group, The Institute for the Learning Sciences, Northwestern University, 1890 Maple Avenue, Evanston, IL 60201, USA

1. Introduction

The scientific goal of artificial intelligence is to understand minds by trying to build them. To make progress towards this goal requires careful decomposition. An especially productive strategy has been to focus on the knowledge and reasoning required for a particular class of domains or tasks. Some domains or tasks have special features which make them tractable, while still providing generalizable insights. If the domains and tasks have practical value then so much the better, for then scientific progress and economic benefits can go hand-in-hand. Qualitative physics is such an area. Reasoning about the physical world is clearly central to intelligences (human or machine). Moreover, it encompasses a variety of tasks and skills, thus providing a range of interesting problems. Physics and mathematics provide us with clues about what special constraints might help make such reasoning tractable. And the potential economic impact of intelligent computer-aided engineering systems is clearly enormous. Indeed, qualitative physics is one of the most exciting areas in AI today.

An important event in the growth of qualitative physics was the 1984 special issue of *Artificial Intelligence* edited by Bobrow. My paper in that issue, "Qualitative process theory", was one-half of my Ph.D. thesis; the other half described algorithms for implementing QP theory and results obtained with them. This essay summarizes the context of that work and the main contributions of QP theory. A cautionary point about basic research

Correspondence to: K.D. Forbus, Qualitative Reasoning Group, The Institute for the Learning Sciences, Northwestern University, 1890 Maple Avenue, Evanston, IL 60201, USA. E-mail: forbus@ils.nwu.edu.

0004-3702/93/$06.00

is made, and the essay ends by describing some current directions this work is taking.

2. Context

Johan de Kleer got me interested in qualitative physics when I was an undergraduate at MIT. I had been working at the AI Lab since 1973 with David Marr on vision, but had always wanted to tackle more central issues of cognition. Seeing SOPHIE in action was a revelation; building systems that smart seemed on a clear path towards the goal of building a humanlike intelligence. When I started graduate work at MIT in 1977 I joined Sussman's engineering problem solving group. I began to work on spatial reasoning, since it was clearly a rich, deep problem, and my earlier work on vision might provide a useful perspective on it. In 1978 draft copies of Pat Hayes' "Naive physics manifesto" [19] and liquids paper [20] galvanized me further: The notion of histories seemed crucial, for it decomposed reasoning about change into dynamics (to evolve possible behaviors) and spatio-temporal problem solving (to determine interactions). Ultimately this work led to a model of qualitative spatial reasoning, demonstrated by FROB, a system which used a diagram to reason about motion of point masses through space [8,10]. For my Ph.D. thesis I originally intended to extend this work into a system that could flexibly reason about a variety of more complex mechanical systems, including clocks. However, circumstances soon led me in a different direction. [1]

To help make ends meet I started working part-time at BBN with Al Stevens and Bruce Roberts on the STEAMER project [22,31]. The goal of STEAMER was to produce an intelligent tutoring system for training operators of oil-fired steam propulsion plants. Its core was a numerical simulation program, originally written to drive the steamship equivalent of a flight simulator, itself a warehouse-sized replica replete with gauges, noisy pumps, pipes, and so on. Such high-fidelity simulators are very useful for drilling students in operating procedures, but are very expensive. Our plan was to develop complementary desktop systems that could help trainees come to a global understanding of the system, including explanations of how physical and engineering principles applied to the safe operation of the plant. In the end, most of STEAMER's power came from providing a flexible, extensible graphical interface to the numerical simulator that a computer-naive trainee or instructor could use to explore the behavior of the steam plant. Today

[1] My work on spatial reasoning continued later, and the work of Faltings [7] and Nielsen [25] finally culminated in a system which could understand mechanical clocks (as well as other fixed-axis mechanical systems) in 1988 [16].

such interfaces are taken for granted, but in 1980 they were rather new, and in fact STEAMER was a major catalyst for the idea of direct manipulation interfaces [23].

My role in the project was to use ideas of qualitative physics to provide humanlike reasoning about the plant's principles and operation, in order to produce understandable explanations. The best qualitative physics available at the time (1979) was in de Kleer's Ph.D. thesis [2]. He had developed a device-centered ontology and the idea of incremental qualitative analysis, in which quantities were represented by the sign of their change from equilibrium in response to a disturbance. I used that physics in a demonstration system for teaching about feedback, using a spring-loaded reducing valve as an example. [2]

The Reducing Valve Demo was a great success [17]. It propagated a disturbance through a constraint network to construct a qualitative description of behavior. The resulting values and dependency structure were then used to automatically generate explanations, interleaving English with color animations of changes propagating through the system. Inspired by this initial success, I tried modeling other steam plant components and subsystems in the same way. These efforts were mostly failures. Several of the reasons were quite interesting:

(1) The notion of quantity used in incremental qualitative analysis was too weak. Often a quantity needed to be compared with several others, and knowledge about relative rates was often important. At minimum, ordinal information about amounts and derivatives seemed to be required.
(2) Identifying causality with the path of propagation in a constraint network sometimes led to unintuitive causal arguments. For example, the arguments "*The increased heat causes the temperature to rise.*" and "*The increased heat causes the amount of stuff to drop.*" can both be generated via propagation from a constraint defining the relationship between heat, temperature, and mass. The first is a legitimate causal argument but the second is not, and a qualitative physics that tries to capture human intuitions about causality should explain why.
(3) The device ontology which was so suitable for electronics seemed unnatural for steam plants. Explanations of phenomena in engineering thermodynamics typically included *physical processes*, which cut across component boundaries.

[2] The popular pressure-regulator example is actually a simplification of this valve: The original contains two stages, since the input stream is 1200 psi steam and would rip apart a single-stage regulator.

This analysis of the limitations of de Kleer's 1979 qualitative physics when applied to engineering thermodynamics led directly to the creation of QP theory.

The first paper about QP theory appeared in 1981 [9], with several other conference papers (AAAI [8] and Cognitive Science Society [17]) leading up to the paper in *Artificial Intelligence*. During this period (1980–1984) QP theory evolved considerably, thanks in part to a variety of stimulating interactions. While Johan de Kleer and John Seely Brown had since moved to Xerox PARC, our discussions continued electronically. QP theory's decidedly cognitive leaning, as well as my interest in cognitive modeling and simulation, grew out of working at BBN with Al Stevens, Dedre Gentner, and Allan Collins. And of course the MIT AI Lab remained a hotbed of engineering problem solving research, ranging from the work of Sussman's VLSI group to the work of Davis' Hardware Troubleshooting Group.

3. Contributions

QP theory introduced some key ideas of qualitative physics:

Physical processes as organizing principle

Ontology plays a central role in the organization of knowledge. A crucial observation in formulating QP theory was that concepts of physical processes (e.g., flows, motion, phase changes) seemed to play an important role in human reasoning about physical systems. Therefore it makes sense to organize theories of physical domains around a formalization of this intuitive notion of physical process.

Representing numerical values via ordinal relationships

Important qualitative distinctions are often tied to comparisons between parameters: Flows occur when pressures or temperatures differ, for instance, and phase changes occur when temperatures reach certain thresholds. If a parameter participates in only one comparison, sign values provide a satisfactory qualitative representations. However, for many circumstances representing values by a set of ordinal relationships (formalized in QP theory as the *quantity space* representation) is more natural.

Sole mechanism assumption

Physical processes are viewed as the mechanisms by which changes occur (excluding the actions of agents). Thus any changes must be explainable by the direct or indirect effects of some collection of physical processes. Causality in QP theory is thus grounded in ontology, rather than in constraint propagation as in de Kleer's account.

For instance, consider again the proposed causal arguments concerning heat, temperature, and mass above. In any reasonable model, heat and mass can each be directly affected by physical processes (e.g., mass or heat flows). Temperature is a causally dependent parameter, determined in terms of heat and mass. Thus the causal argument "*The increased heat causes the temperature to rise.*" is legitimate, since it follows the causality imposed by the physical processes in the domain. On the other hand, the argument "*The increased heat causes the amount of stuff to drop.*" is not legitimate since it reverses the direction of causality imposed by the physical processes.

Compositional qualitative mathematics

A missing ingredient in early attempts to formalize physical reasoning (cf. [20,21,29]) was the idea of *compositionality*. That is, a major factor in the flexibility of human reasoning about complex physical systems comes from the ability to use partial information and combine it as available.

The idea of *qualitative proportionality* captures one aspect of this productive use of partial information. For example, in modeling fluid resistance we might know that it depends on the area and the length of the path, i.e.,

$$\texttt{resistance(path22)} \propto_{Q+} \texttt{length(path22)},$$

$$\texttt{resistance(path22)} \propto_{Q-} \texttt{area(path22)}.$$

These qualitative relationships tell us that two potential ways to increase the resistance of `path22` is to increase its length or decrease its area.

Explicit representation and reasoning about modeling assumptions

De Kleer and Brown's notion of *class-wide assumptions* [3] was a valuable contribution in thinking about how knowledge about physical systems should be organized. However, their focus on electronics did not drive them to fully explore its consequences since the mapping from a circuit schematic to abstract electronic components is fairly direct. In thermodynamics the task of setting up a model is more difficult, and that factor, along with the goal of trying to capture different states of student knowledge, led me to focus on the task of model building as central to QP theory.

QP theory was the first system of qualitative physics to explicitly represent the conditions under which particular pieces of knowledge were applicable, and to make constructing the model of a specific system from a domain theory a central part of its computational account. QP theory can be viewed as a partial specification of a space of modeling languages for domains where physical processes are the appropriate ontology.

4. A cautionary note

The relationship between basic research and applied work is being debated strongly in the U.S. right now, with a strong tendency to push researchers towards applied work. I believe science progresses best when basic researchers remain connected to real problems, but are not obliged to solve them on any short-range timetable. The development of QP theory provides an illustration of this point.

In the summer of 1980 I tried to apply the existing theories of qualitative physics to the problems confronting us in the STEAMER project. The Reducing Valve Demo took two weeks to build. The next two months were spent mostly producing failures. The third month was spent reflecting on what was going wrong. Completing the first paper on QP theory [9] took roughly six months. In 1982 the *AI Journal* version was submitted, it appeared in 1984, the same year I received my Ph.D. Since then, my students and I have done basic research, motivated in part by the problems I could not solve then. Only now, ten years later, do we think we have enough ideas and technology to try developing some useful applications again.

I believe that we (and the field as a whole) have made substantial progress during those ten years. I also believe that our progress would have been impossible if we had been forced to field a new demo or working application every year. This lesson is of course an old one, deeply ingrained in the very division between basic and applied work. However, in today's troubled economic climate many seem to have forgotten it. Clearly practical problems and examples are crucial sources of inspiration and motivation—for instance, my group draws on engineering problems arising in space station design, civilian aviation, and propulsion plants in our research. However, the "demo or die" strategy often touted today seems to me to be a prescription for "demo *and* die", at least with respect to scientific progress.

5. Towards tutor compilers and learning machines

Research in qualitative physics is progressing well these days, and the future looks extremely bright. There is so much activity in qualitative physics these days that any brief summary cannot do it justice, so I will focus on my group's efforts.

One important problem is the need for intelligent tutoring systems and learning environments for science and engineering education and training. The kind of system that STEAMER was supposed to be, where a numerical simulation is integrated with an intuitive understanding of the artifact, is clearly an important component for such computer-based teaching systems. Developing a solid theoretical foundation for such systems has been a major motivation for our qualitative physics work. For example, what STEAMER should have been is what we now call a *self-explanatory simulator* [13,14], that is, a numerical simulation with an integrated qualitative understanding of the plant. Self-explanatory simulators can explain as well as reproduce the behavior of what they are modeling, and thus provide a basis for deeper reasoning about behavior. The qualitative model provides the information needed to compile such simulators automatically: It identifies what equations are appropriate under different conditions, and the causal account provides an order of computation in the numerical aspect of the simulation.

One of our medium-term goals is to develop a *tutor compiler* that can produce self-explanatory simulators that can be used either as stand-alone training systems or as modules in multimedia learning environments. In addition to self-explanatory simulators, two other ideas are key pieces to the puzzle of how qualitative physics can help build intelligent tutoring systems and learning environments. Reasoning about a complex system often requires formulating a model tuned towards that specific task, which in turn requires reasoning about modeling assumptions. This problem is especially acute in training situations, since the learner cannot be expected to know what models are appropriate. The *compositional modeling* strategy Falkenhainer and I developed [5] extends the modeling capabilities of QP theory to orchestrate the construction and use of domain theories that describe phenomena at multiple grain sizes and from different, often conflicting, perspectives. The other problem is that teaching correct operation of a system (or understanding how the operation of a system might be impaired by malfunctions) requires understanding the interaction of actions taken by an agent with the physical world. The idea of *action-augmented envisionments* [12] provides a simple conceptual framework for integrating action with dynamics, which should support the generation, verification, and teaching of operating procedures.

The tutor compiler work, like our work on monitoring [1,11], engineering analysis [30], and design, is heavily motivated by classes of applications. The perceived applications potential seems in fact to be a major reason for the popularity of qualitative physics. However, a side of the field which is just as important, but has been relatively neglected, is cognitive modeling. Certainly building engineering problem solvers under the sole constraint of optimum performance will reveal interesting properties of complex reasoning in complicated domains. However, providing a formalism for investigating human mental models of complex systems was, and should continue to be, another motivation for qualitative physics. [3]

[3]Even from an applications perspective such research should be interesting: Evidence suggests that qualitative physics representations provide a good conceptual account of human mental models, making qualitative physics a valuable tool for human/computer interaction.

So far the use of qualitative physics for modeling scientific discovery has been one of the few uses of qualitative physics in other parts of AI and Cognitive Science (cf. [4,26–28]). We believe there is a variety of exciting, productive possibilities for such research (cf. [15]). One approach we are exploring is integrating QP theory with other common-sense theories to see if we can develop programs that learn in a humanlike way from science books. We are focusing especially on analogical learning, based on Gentner's Structure-Mapping theory [6,18]. To provide the domain knowledge for analogizing we are implementing QP theory in CYC [24]. This is a long-term enterprise, but we hope to learn much on the way about how qualitative reasoning can be used in a much broader context.

References

[1] D. DeCoste, Dynamic across-time measurement interpretation, *Artif. Intell.* **51** (1991) 273–341.

[2] J. de Kleer, Causal and teleological reasoning in circuit recognition, MIT AI Lab, Tech. Report No. 529, Cambridge, MA (1979).

[3] J. de Kleer and J.S. Brown, Assumptions and ambiguities in mechanistic mental models, in: D. Gentner and A. Stevens, eds., *Mental Models* (Lawrence Erlbaum, Hillsdale, NJ, 1983).

[4] B. Falkenhainer, A unified approach to explanation and theory formation, in: J. Shrager and P. Langley, eds., *Computational Models of Scientific Discovery and Theory Formation* (Morgan Kaufmann, San Mateo, CA, 1990).

[5] B. Falkenhainer and K.D. Forbus, Compositional modeling: finding the right model for the job, *Artif. Intell.* **51** (1991) 95–143.

[6] B. Falkenhainer, K.D. Forbus and D. Gentner, The structure-mapping engine: algorithm and examples, *Artif. Intell.* **41** (1990) 1–63.

[7] B. Faltings, Qualitative kinematics in mechanisms, *Artif. Intell.* **44** (1990) 41–87.

[8] K.D. Forbus, Spatial and qualitative aspects of reasoning about motion, in: *Proceedings AAAI-80*, Stanford, CA (1980).

[9] K.D. Forbus, Qualitative reasoning about physical processes, in: *Proceedings IJCAI-81*, Vancouver, BC (1981).

[10] K.D. Forbus, Qualitative reasoning about space and motion, in: D. Gentner and A. Stevens, eds., *Mental Models* (Lawrence Erlbaum, Hillsdale, NJ, 1983).

[11] K.D. Forbus, Interpreting observations of physical systems, *IEEE Trans. Syst. Man Cybern.* **17** (3) (1987).

[12] K.D. Forbus, Introducing actions into qualitative simulation, in: *Proceedings IJCAI-89*, Detoit, MI (1989) 1273–1278.

[13] K.D. Forbus and B. Falkenhainer, Self-explanatory simulations: an integration of qualitative and quantitative knowledge, in: *Proceedings AAAI-90*, Boston, MA (1990).

[14] K.D. Forbus and B. Falkenhainer, Self-explanatory simulations: scaling up to large models, in: *Proceedings AAAI-92*, San Jose, CA (1992).

[15] K.D. Forbus and D. Gentner, Learning physical domains: towards a theoretical framework, in: R.S. Michalski, J.G. Carbonell and T.M. Mitchell, *Machine Learning: An Artificial Intelligence Approach*, Vol. 2 (Tioga Press, Palo Alto, CA, 1986).

[16] K.D. Forbus, P. Nielsen and B. Faltings, Qualitative spatial reasoning: the CLOCK project, *Artif. Intell.* **51** (1991) 417–471.

[17] K.D. Forbus and A. Stevens, Using qualitative simulation to generate explanations, in: *Proceedings Third Annual Conference of the Cognitive Science Society*, Berkeley, CA (1981).

[18] D. Gentner, Structure-mapping: a theoretical framework for analogy, *Cogn. Sci.* **7** (2) (1983).

[19] P.J. Hayes, The naive physics manifesto, in: D. Michie, ed., *Expert Systems in the Micro-Electronic Age* (Edinburgh University Press, Edinburgh, Scotland, 1979).

[20] P.J. Hayes, Naive physics 1: ontology for liquids, in: J.R. Hobbs and R. Moore, eds., *Formal Theories of the Commonsense World* (Ablex, Norwood, NJ, 1985).

[21] G.G. Hendrix, Modeling simultaneous actions and continuous processes, *Artif. Intell.* **4** (1973) 145–180.

[22] J.D. Hollan, E.L. Hutchins and L. Weitzman, STEAMER: an interactive inspectable simulation-based training system, *AI Mag.* **5** (2) (1984) 15–27.

[23] E.L. Hutchins, J.D. Hollan and D.A. Norman, Direct manipulation interfaces, *Human–Comput. Interaction* **1** (1985) 311–338.

[24] D.B. Lenat and R.V. Guha, *Building Large Knowledge-Based Systems* (Addison-Wesley, Reading, MA, 1990).

[25] P. Nielsen, A qualitative approach to mechanical constraint, in: *Proceedings AAAI-88*, St. Paul, MN (1988); also in: D.S. Weld and J. de Kleer, eds., *Qualitative Reasoning about Physical Systems* (Morgan Kaufman, San Mateo, CA, 1990).

[26] B. Nordhausen and P. Langley, An integrated approach to empirical discovery, in: J. Shrager and P. Langley, eds., *Computational Models of Scientific Discovery and Theory Formation* (Morgan Kaufmann, San Mateo, CA, 1990).

[27] P. O'Rorke, S. Morris and D. Schulenberg, Theory formation by abduction: a case study based on the chemical revolution, in: J. Shrager and P. Langley, eds., *Computational Models of Scientific Discovery and Theory Formation* (Morgan Kaufmann, San Mateo, CA, 1990).

[28] S. Rajamoney, A computational approach to theory revision, in: J. Shrager and P. Langley, eds., *Computational Models of Scientific Discovery and Theory Formation* (Morgan Kaufmann, San Mateo, CA, 1990).

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

[30] G. Skorstad and K.D. Forbus, Qualitative and quantitative reasoning about thermodynamics, in: *Proceedings Eleventh Annual Conference of the Cognitive Science Society*, Ann Arbor, MI (1989).

[31] A. Stevens, B. Roberts, L. Stead, K.D. Forbus, C. Steinberg and B. Smith, STEAMER: advanced computer-aided instruction in propulsion engineering, BBN Tech. Report, Cambridge, MA (1981).