

Current Topics in Qualitative Reasoning

Bert Bredeweg and Peter Struss

■ In this editorial introduction to this special issue of *AI Magazine* on qualitative reasoning, we briefly discuss the main motivations and characteristics of this branch of AI research. We also summarize the contributions in this issue and point out challenges for future research.

Qualitative reasoning automates reasoning and problem solving about the physical world around us, a topic that is fundamental to AI. Qualitative reasoning has also generated techniques and systems that enter significant application domains. Successful application areas include autonomous spacecraft support, failure analysis and on-board diagnosis of vehicle systems, automated generation of control software for photocopiers, and intelligent aids for learning about thermodynamic cycles. Qualitative reasoning is thus relevant for researchers who are interested in important AI issues as well as for managers, developers, and engineers who are looking for potential industrial benefits of AI.

A decade has passed since the publication of three collections of papers and a book covering the foundations of the field (Weld and de Kleer 1990), the status at that time (Faltings and Struss 1992; Williams and de Kleer 1991), and a comprehensive treatment of qualitative simulation (Kuipers 1994). One of the purposes of this special issue is to illustrate the significant progress that has been made since this time. We also point out open research problems, links to other AI fields, and future opportunities of industrial applications that should attract ambitious researchers and industrial developers. The goal was not to produce an updated survey of the foundations or provide a comprehensive overview of all the ongoing work (as, for example, in Forbus [2004]). In-

stead, we tried to select a number of recent pieces of work that are interesting and illustrative and clarify the objectives of the field, its methods, achievements, and application potential. However, many other excellent contributions could not be included.¹

Goal and Motivations

Reasoning about, and solving problems in, the physical world is one of the most fundamental capabilities of human intelligence and a fundamental subject for AI. However, what are the key research topics? There are the scientific disciplines such as physics and chemistry that develop theories, and there are engineering disciplines that develop solutions that change the physical world. Both use formal mathematical systems, as well as computer implementations, to derive conclusions about natural and artificial pieces of the world. Does this approach not provide a systematic and formal way to reason about the physical world? What remains to be done for AI research in this area?

An important motivation for qualitative research is the observation that scientific, mathematical, and engineering formalisms and models represent only a part of the intellectual activities of humans in their continuous interaction with the real world. Even stronger, their existence and application crucially depends on a more fundamental, conceptual, and qualitative way to perceive, analyze, understand, and model the world that has rarely been subject to formal treatment in the classical sciences.

The area where this phenomenon is most evident is the commonsense reasoning and problem solving humans perform continuously as they go about their tasks in the physical

Reasoning about, and solving problems in, the physical world is one of the most fundamental capabilities of human intelligence and a fundamental subject for AI.

world. Without qualitative reasoning, they would not survive. It allows humans to behave successfully in new situations, achieve goals, build and use tools, and so on. In addition, all this reasoning happens without possessing formal scientific theories that cover the relevant features in a systematic, detailed, and numeric way. Studying and formally representing this basic knowledge about the physical world is one of the motivations underlying research on qualitative reasoning. This motivation is illustrated by the Naïve Physics Manifesto (Hayes 1979, p. 2) that proposed “the construction of a formalization of a sizable portion of commonsense knowledge about the everyday physical world: about objects, shape, space, movement, substance (solids and liquids), time, etc.”

Although it appears natural that this kind of prescientific reasoning is qualitative in nature, it seems impossible that any new theory and nonnumeric formalism developed by AI could do better than the well-founded and elaborate theories and techniques developed over centuries in science and engineering. However, qualitative reasoning is not so much about solving the same problems differently or better. The goal is to solve a different set of problems. Mathematics and computer science have provided theories, calculus, and programs for implementing numeric simulations of physical situations. However, the creation of a model is beyond the scope of such instruments and so is a general formal account for the relevant concepts and inferences that could serve as the basis for implementing such reasoning in software.

To create the model, it is necessary to not only identify the relevant objects and interactions, their important properties and quantities, but also to determine what is irrelevant or negligible. In other words, the existence or creation of a conceptual and qualitative model of a situation for a certain task to be solved is a prerequisite for a scientific quantitative model, not its result, as sometimes suggested. An important task of qualitative reasoning is to provide formal languages for expressing such model elements, for developing inference schemes that compose them and enable conclusions to be drawn. This ability is essential to build more powerful and intelligent tools that support or automate tasks such as modeling, design, analysis, diagnosis, and explanation of natural or artificial systems (cf. de Kleer [1990]). It is also central to effective computer support for teaching science and training engineering skills. Interactive learning environments cannot be restricted to pure simulation

but have to interactively deal with queries, problems, assignments, and explanations at a conceptual and intuitive level.

Prominent Features of Qualitative Reasoning

Although the qualitative reasoning field has addressed diverse problem areas and developed a variety of theories and systems, there are a number of prominent features that are typical for many of the approaches. Some of the most important ones are as follows:

Ontologies: Qualitative reasoning provides explicit representations of the conceptual modeling layer rather than only an executable mathematical expression. This layer is crucial to any attempt to support model building and even more to automate it and has been one of the major issues of the field. The two main families of ontologies model physical systems in terms of interconnected components or interacting processes.

Causality: Analyzing and explaining the behavior of a system in terms of cause-effect relations is central to human intuitive reasoning and communication but is not present in the mathematical expressions describing system behavior. Formalizing this concept and exploiting it in automated reasoning is the basis for explanation facilities in model-based systems.

Compositional modeling: Most qualitative reasoning systems adopt a reductionist view of the world and aim at building libraries of elementary, independent model fragments (component behavior, processes, and so on). This approach provides the basis for automating the model composition and reusing models, a highly desirable feature for industrial applications.

Inference of behavior from structure: Many systems require as an input a purely structural description of a scenario (in terms of interrelated objects), build a behavior model automatically, and generate a behavior description either by qualitative simulation or as a comprehensive representation of all possible behaviors.

Qualitativeness: Qualitativeness means including only those distinctions in a behavior model that are essential for solving a particular task for a certain system. The goal is to obtain a finite representation that leads to coarse, intuitive representations of models and behavior (for example, as finite relations and a state graph, respectively), and efficient algorithms.

Contributions in This Issue

The eight papers selected for this special issue illustrate the features mentioned earlier in different domains and for different tasks. The contributions are self-contained, as much as possible, and can be read independently. We briefly highlight the main issues for each contribution.

Trave-Massuyes, Ironi, and Dague present some of the mathematical foundations of qualitative reasoning. Their contribution shows that representing system behaviors qualitatively does not mean resorting to heuristic or approximate knowledge. Rather, the qualitative calculi used have a formal basis and solid mathematical foundations. Also, their subsection on system identification shows that qualitative methods can be complementary or even superior to conventional numeric approaches in mathematics and engineering.

Despite the fact that qualitative reasoning research addresses very fundamental questions of knowledge representation and reasoning, its results have already taken the step toward real industrial applications on a broad scale, as documented by several contributions in this issue.

Discussing model-based programming of fault-aware systems, Williams, Ingham, Chung, Elliott, and Hofbauer demonstrate how their research for the National Aeronautics and Space Administration supports the goal of achieving autonomous behavior of designed systems, a feature highly important, but not confined, to spacecraft. In particular, an autonomous system needs to identify the current fault mode. Based on this result, it must plan and execute suitable actions to recover from this fault and achieve the mission, for example, by reconfiguring the system under exploitation of its structural redundancy. Qualitative models provide compact representations, and a compositional modeling approach is essential to automatically generate solutions rather than be limited to a set of enumerated preselected recovery plans.

Struss and Price present recent work on model-based systems in the automotive industry. They discuss various advanced applications during the entire life cycle of a vehicle. Being able to generate models and solutions for different work processes, rather than handcraft them for each variant or special purpose, is economically a highly attractive feature that exploits libraries of reusable model fragments and the compositionality in modeling. Because numeric models are often developed and used in the engineering design processes, the mathematical foundations that allow relating qualitative models to, and even generating them from, the numeric ones is an important topic.

Fromherz, Bobrow, and de Kleer describe model-based computing for design and control of reconfigurable systems. Their system automatically assembles control software for Xerox copiers from the specific configuration of a copier and even adapts the software when new physical modules are added in the field. Obviously, the component models exploited in the system have to be compositional to provide the required flexibility and coverage, but the granularity of the behavior description can be less coarse compared to other qualitative reasoning applications.

Bredeweg and Forbus discuss qualitative modeling in education. Following current theory on education that advocates learning by doing, they use the ontologies provided by qualitative reasoning as the basis for knowledge capture tools. The underlying hypothesis is that modeling is an important skill for learners to acquire and that the articulation of ideas in a formal representation is an important means to induce learning. Learners can formulate their insights on how systems behave in an appropriate qualitative and causal way. Using such external representations, they can test, debug, and refine ideas (for example, by running what-if scenarios) and share them with peers to foster learning in a collaborative setting.

Traditionally, qualitative reasoning research has focused on physics and engineering domains, which might suggest that the qualitative approach cannot extend to other areas. However, there are many other areas of research and application that can benefit from using qualitative reasoning. Salles and Bredeweg discuss qualitative reasoning about population and community ecology and show how qualitative reasoning allows formulating formal models in an area that has an inherent lack of precise numeric information. The models they present implement a compositional approach, allowing the overall system behavior to be explained in terms of basic processes, and can generate qualitatively different evolutions of populations, revealing their respective preconditions.

Bailey-Kellogg and Zhao apply qualitative reasoning to spatial reasoning. Their work on extracting and reasoning with spatial aggregates extracts abstract, qualitative spatial entities from numeric fields. Such aggregate concepts are essential for interpreting and processing the data. For example, the interpretation of numeric weather data requires the identification of higher-level concepts, such as high-pressure systems and cold fronts. The contribution also includes a brief discussion of other approaches to spatial reasoning.

Traditionally, qualitative reasoning research has focused on physics and engineering domains, which might suggest that the qualitative approach cannot extend to other areas. However, there are many other areas of research and application that can benefit from using qualitative reasoning.

Bratko and Suc discuss learning qualitative models. Their system generates qualitative models from data about human or algorithmic controllers. They show how qualitative reasoning provides the representational basis for explicitly encoding control laws and strategies that are either hidden in undocumented control software or difficult to obtain from skilled operators by introspection.

Challenges

The qualitative reasoning research field and its application domains have progressed fairly well. However, the field faces several open issues and challenges, many of which have been identified or emphasized by work on real applications. We highlight some of them that we currently find of high importance.

One issue concerns the development of better engineered and easy-to-use tools that facilitate the exchange of results among researchers and make qualitative reasoning techniques available to potential users in other areas and application work. The field, to date, has developed a variety of theories, formalisms, and techniques with different degrees of generality and is still far from delivering a small set of uniform principles and systems. If the field can make progress on these issues, it will become easier to create and exchange libraries of reusable models.

Automated model building and model transformation needs continued theoretical work and more effective and efficient algorithms, which is emphasized by application requirements. Much of the expected gain depends on fast and economic creation of models from a library. Because different tasks can require models at different levels of abstraction, there is a tension between the desired compositionality and generality (and, hence, reusability) of models and the necessity of task-oriented models. Qualitative reasoning needs to develop techniques to generate task-oriented models from generic ones.

This issue also touches on a more general goal, namely, integrating qualitative reasoning results and tech-

niques with standard engineering practice and tools. The lack of integration presents a major obstacle to transferring qualitative reasoning technologies into industrial practice. Deriving qualitative models from numeric ones that have been developed, for example, in the phase of design verification, is of high practical importance. However, it might require changes in current modeling practice toward modular, component-oriented models. The need to blend in with current practice also applies to other domains, such as medicine, economy, biology, and ecology.

To achieve the goals mentioned earlier, qualitative reasoning has to broaden its scope. To date, it has mainly focused on modeling physical systems and reasoning about them. To become more relevant and have a serious economic impact, the application of qualitative reasoning technology must be interwoven with human activities in real-life settings (such as planning and carrying out actions to localize a fault and repair it) and also develop formal representation of these activities and their interrelationships. On the technical side, there is the challenge to promote the integration of AI planning techniques and model-based problem solving in principled ways.

Acknowledgments

We appreciate the support provided by the *AI Magazine* staff, particularly by David Leake and Mike Hamilton. We thank the authors for providing their contributions and for being patient in dealing with our critical requests. In addition, we would especially like to thank the following persons for helping us realize this special issue: Gautam Biswas, Joost Breuker, George Coghill, Tony Cohn, Ken Forbus, Francois Guerrin, Johan de Kleer, Ben Kuipers, Chris Price, Julie-Ann Sime, and Feng Zhao.

Note

1. Visit the Monet MIR database for recent publications on qualitative reasoning: monet.aber.ac.uk/.

References

de Kleer, J. 1990. Qualitative Physics: A Personal View. In *Readings in Qualitative Rea-*

soning about Physical Systems, eds. D. S. Weld and J. de Kleer, 1–8. San Francisco, Calif.: Morgan Kaufmann.

Faltings, B., and Struss, P., eds. 1992. *Recent Advances in Qualitative Physics*. Cambridge, Mass.: MIT Press.

Forbus, K. 2004. Qualitative Reasoning. CRC Handbook of Computer Science and Engineering. 2d ed. Boca Raton, Fla.: CRC. Forthcoming.

Hayes, P. 1979. The Naive Physics Manifest. In *Expert Systems in the Micro-Electronic Age*, ed. D. Michie, 242–270. Edinburgh, United Kingdom: Edinburgh University Press.

Kuipers, B. 1994. *Qualitative Reasoning: Modeling and Simulation with Incomplete Knowledge*. Cambridge, Mass.: MIT Press.

Weld, D., and de Kleer, J., eds. 1990. *Readings in Qualitative Reasoning about Physical Systems*. San Francisco, Calif.: Morgan Kaufmann.

Williams, B. C., and de Kleer, J., eds. 1991. *Artificial Intelligence* (Special issue on Qualitative Reasoning about Physical Systems) 51(1–3).



Bert Bredeweg's research (University of Amsterdam, Department of Social Science Informatics) is driven by fundamental questions about computational intelligence. It includes issues related to interactive learning environments, human-computer interaction, and the use of AI technology to enable communicative interactions between humans and computers. His e-mail address is bert@swi.psy.uva.nl.



Peter Struss is a professor of computer science at the Technical University of Munich and a managing director of OCC'M Software GmbH. He obtained his master degree in mathematics at the University

of Goettingen, a Ph.D. in computer science at the University of Kaiserslautern, and his habilitation at the Technical University of Munich. The focus of his work is on qualitative modeling and model-based systems and the transfer of technology into industrial applications. His e-mail address is struss@in.tum.de.