The Knowledge Engineering Review, Vol. 21:4, 317–334. © 2006, Cambridge University Press
doi:10.1017/S026988890600097X Printed in the United Kingdom

Qualitative futures

C. J. PRICE[1], L. TRAVÉ-MASSUYÈS[2], R. MILNE, L. IRONI[3], K. FORBUS[4], B. BREDEWEG[5], M. H. LEE[1], P. STRUSS[6], N. SNOOKE[1], P. LUCAS[7], M. CAVAZZA[8] and G. M. COGHILL[9]

[1]*Department of Computer Science, University of Wales, Aberystwyth SY23 3DB, UK;*
email: cjp@aber.ac.uk
[2]*CNRS-LAAS, Toulouse*
[3]*IMATI-CNR, Pavia*
[4]*Northwestern University*
[5]*Universiteit van Amsterdam*
[6]*Technische Universität München*
[7]*Radboud Universiteit Nijmegen*
[8]*University of Teesside*
[9]*University of Aberdeen*

Abstract

This paper reviews the state of the art in model-based systems and qualitative reasoning, and considers where the field will be in 20 years time. It highlights six areas where developments in model-based systems in general, and in qualitative reasoning in particular, have the potential to provide significant computer-based help. The paper also examines where further technological developments might be needed in order to achieve these qualitative futures.

1 Introduction

Model-based systems and qualitative reasoning (MBS&QR) as a visible sub-field of Artificial Intelligence can be traced back some 20 years to the publication in 1984 of the seminal collection of papers in the field (Bobrow, 1985), although there was earlier work in this area (de Kleer, 1977; Brown *et al.*, 1982).

The 1984 collection of papers showcased early research, much of which has been developed further over the past 20 years, and resulted in real-world demonstrations of the potential of the technology.

This paper reviews where we are in being able to build systems using MBS&QR, and looks forward to the next 20 years, considering what MBS&QR will be recognized for in 20 years time.

The review of past work is done by giving a detailed example of the use of model-based reasoning for remote diagnosis and repair, and by considering more briefly the wide spectrum of applications of MBS&QR technology.

The exploration of the future of MBS&QR is done by selecting six areas where MBS&QR can make a significant difference, and detailing the nature of the technological challenge in each area. None of these challenges can be met completely with the MBS&QR technology that has been developed so far. For this reason, the paper also explores the areas where further research in MBS&QR is needed in order to fulfil the vision of model-based applications given here.

2 Key ideas of model-based systems

This section explores the key ideas behind Model-Based Systems (MBS), and explains why qualitative reasoning (QR) is foundational to effective reasoning with models.

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

- Separation of the system model from problem solving. MBS are based on a separation of the problem-solving algorithm from the model of the domain. Once a library of appropriate component models has been established, only a structural description of the respective device or system (e.g. obtained from design data) is required to automatically generate a system model and, based on it, problem-solving software dedicated to this device or system.
- Compositionality. Since systems are assembled from standard components and the behaviour (and misbehaviour in the case of a fault) of the system emerges from the behaviour of these components, establishing a model library is feasible and entails collecting models of (correct and faulty) behaviour of such standard components. This is important: this kind of model-based reasoning cannot be performed if the overall behaviour of the system cannot be composed from the behaviour of the components and the way in which they are linked. Where there is compositionality of models, a high degree of reuse is possible, as well as a prediction of what will happen in unexpected circumstances such as failure situations.
- Existence of different abstract problem solvers. Section 3 gives examples of a wide range of tasks that can be solved in a model-based way, where solving a new problem involves fitting a different appropriate kind of model into the generic problem solver. Combining compositional models of a domain with domain-independent problem solvers can produce MBS that are able to reason about different versions of a product, or about a process as its operating state changes.
- Modelling at different levels of abstraction. In diagnosis, for example, modelling different kinds of problems may well involve modelling different phenomena and at different levels of detail. However, while there may be a need for quantitative and/or semi-quantitative models, qualitative models provide a nice solution for representing phenomena at a higher level of abstraction.

Engineers generally prefer to work with numerical models, but in model-based reasoning the capability to exploit qualitative models turns out to be crucial for several fundamental reasons (Forbus, 1988).

- In particular in early design phases, only a partial specification of components and parameters is available, which prevents the use of numerical techniques.
- Many tasks, such as failure modes and effects analysis (FMEA) or the generation of diagnostic manuals, aim at statements about classes of (fault) behaviour and of symptoms rather than specific instances. For example, the effect of a leakage of unknown size has to be predicted, rather than just a leakage of a specified size.
- Faults can be defined as qualitative deviations from normal functioning (e.g. flow through a pipe is reduced), rather than arbitrary discrepancies with respect to precise values (e.g. flow is 4.12 gallons min^{-1}, but should be 6.73 gallons min^{-1}).
- Precise values often do not exist; for example, a vehicle is operated in a noisy and widely unmeasurable environment, and only incomplete data are available (e.g. about properties of the road surface).
- Qualitative models provide an appropriate level of abstraction for modelling complex systems and processes where standard mathematical models do not exist or are not tractable (e.g. modelling the combustion process, or the communication among the control units composing the system).
- They enable an intuitive representation and presentation of knowledge and information to the users.
- Where more detailed models are needed in order to produce precise answers, the qualitative models provide a way of focusing the numerical modelling on the answers that are needed, and of interpreting the results of numerical reasoning.

The qualities outlined above mean that qualitative models often provide appropriate answers for a wide range of systems where only incomplete knowledge is available. This enables automation of reasoning about the complex systems found in modern machines, early identification of safety and

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

reliability issues and generation of good diagnostics. This can be done for many different system variants with little extra effort.

A system's evolution can be explored in discrete terms, by defining states and the events that trigger transitions between states. This is generally the adopted point of view when continuous dynamics of behaviour are not relevant. The important contribution of QR is to provide an intermediate level between discrete event and continuous models, in which the state space is discretized into a number of finite states and transitions between those states obey continuity constraints.

The ultimate goal in modelling would be to represent a system in an integrated manner even where the details of the system may be unequally known. Such a goal could only be reached by a system able to make use of the various existing modelling paradigms. Indeed, complex systems are characterized by the heterogeneity of their components: a continuous behaviour component may be controlled by a valve including a commanding mechanism (driver), where the valve driver is implemented in digital logic and so such a mixture of levels of modelling is quite typical for complex systems.

Three types of alternative approaches have been proposed towards achieving such a goal. They have each been shown to be powerful and can be seen as contributions towards this unified modelling view.

- To use more qualitative and quantitative information, as is done in the semi-quantitative simulation approach. Attempts at such an approach include quantitative extensions of QSIM such as Q2 and Q3 which preserve the underlying qualitative model, and interval model-based simulation (Kuipers, 1986, 1994; Berleant & Kuipers, 1997).
- To utilize results in the area of systems theory, such as the qualitative phase space analysis approach.
- To integrate QR with traditional engineering modelling approaches such as numerical simulation or system identification. Self-explanatory simulation and hybrid systems domains exemplify the first type of integration, and the second is explored in the next paragraph.

In the system identification approach, QR complements numerical identification approaches and can play a crucial part in structural identification, i.e. in selecting the form of the equation within the model space. QR techniques naturally complement grey box and black box system identification techniques: in one case, they allow us either to supply the necessary knowledge or to emulate the expert's reasoning about structural identification; in the other, when the box is not completely black (as is quite often the case), they allow us to easily choose the proper equation complexity but above all to embed prior knowledge with a significant gain in model robustness (Travé-Massuyès *et al.*, 2003; Coghill *et al.*, 2004).

Another important fact is that QR models can complement engineering models by their conceptual nature and higher level of abstraction. They bring in new features, such as the underlying causality, that are crucial for supporting reasoning mechanisms.

Numerical models have a poor background for logical and causal reasoning. Hybrid MBS proposed recently have clearly shown that QR concepts establish a necessary link between numerical continuous models and the logical reasoning level (Williams & Nayak, 1996, 1997; Bénazéra & Travé-Massuyès, 2003; Bénazéra, 2003). Where the models used are only sets of differential equations in matrix form manipulated algebraically, they do not enable traceable predictions and do not include the necessary attached labels for allowing logical reasoning. The use of QR concepts, for example the underlying causal influences employed in Ca-En (Travé-Massuyès & Milne, 1997), makes these engineering models suitable for supporting logical reasoning.

A clear example of the need for qualitative models comes in the diagnosis domain. More often than not, even when numeric models of normal behaviour are available, fault models are not available, and neither are the data that would be necessary to derive numerical fault models. In this case, qualitative models are an obvious solution: they are good at capturing the uncertainty related to faults and they are generally sufficient for diagnostic purposes.

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

3 Range of QR and MBS applications

This section gives an idea of the variety of applications and domains in which model-based reasoning is currently being used. It can be seen that this is already a technology with a wide range of applicability. It also provides a good return on investment for those who apply it.

3.1 Fault detection by model-based prediction: numeric and non-numeric

If one knows what values the system parameters should have, then one can detect faults by seeing if the system is producing these values or not. However, for many systems, the behaviour of the components and subsystems is not well enough known to be used for a numerical simulation. In this case, QR and simulation can be used to produce a description of the overall expected system behaviour, thus enabling fault detection. For example, QR has been used in applications as diverse as spacecraft monitoring and diagnosis (Williams & Nayak, 1996), gas turbine monitoring (Travé-Massuyès & Milne, 1997) and ecological system observation (Heller & Struss, 1997) to detect and react to qualitative behavioural changes.

3.2 System simulation before the real system is built

This is particularly important in areas such as satellite design and virtual prototyping of vehicles, where design is iterative and prototyping is expensive. The developers want to understand what the system will be like, even though it will not be physically constructed for some time. Complex products involving discrete and process subsystems are very difficult to model with traditional simulation systems, but the qualitative nature of the behaviour of the system can be determined with qualitative simulation. Ward & Price (2001) give an example of this in the automotive sector, where virtual prototyping is used to explore the implications of a specific design for an automotive subsystem. Benazera & Travé-Massuyès (2003) discuss the use of models in the KOALA software prototype to produce efficient diagnostics before the actual system has been built.

3.3 Process understanding and monitoring

The operators of process plants need to be certain that the plant is reacting as it should; for example, they need to be certain that temperatures are increasing and decreasing when they should be. In such situations, numerical simulation creates a complex set of numbers when the user really wants to understand that the key system parameters are increasing or decreasing. QR can provide an appropriate level of reasoning (Trelease & Park, 1996; Adam & Grant, 2001). The CHEM project (`http://www.chem-dss.org`), funded by the European Commission, has successfully used model-based reasoning in conjunction with a range of other techniques to monitor and diagnose problems on chemical process plants.

3.4 Explanation of numerical simulations

Numerical simulators produce a battery of numbers, but not the easy to understand description of system behaviour the user is looking for. Forbus & Falkenheiner (1991) describe how qualitative and quantitative simulation can be coupled to provide simulators that are capable of explaining how they reached their results. Price *et al.* (2003) describe how QR can be used to extract the system's qualitative behaviours from the simulation output. This can be used with functional labelling to extract a description of system operation at the appropriate level for the user.

3.5 Cognitive applications

Qualitative models provide a useful level of representation for capturing many of the mental models (Gentner & Stevens, 1983) that people use in reasoning about the everyday world and about

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

complex systems. For example, qualitative models combined with analogical processing can be used to answer everyday physical reasoning questions (Klenk *et al*., 2005).

Qualitative models are also promising for formalizing some aspects of natural language semantics, both because the level of detail in which they carve up the world is very close to that of event descriptions commonly found in human language and because the compositionality of the modelling formalism enables knowledge of a situation to be built up incrementally, across multiple sentences (Forbus & Kuehne, 2005). In addition to handling many aspects of the semantics of the physical world, qualitative models appear to be a useful level of representation for many aspects of semantics which can be construed in terms of continuous parameters, such as metaphors and reasoning about the properties of other agents.

3.6 Compositional model-based diagnosis and state tracking

By linking together a collection of component descriptions, diagnosis can be performed on the whole system and the state of the system can be tracked over time (Dvorak & Kuipers, 1991; Miguel & Shen, 2005). Struss & Price (2003) give a range of examples of this happening in the automotive industry. The experience is that the use of such techniques is much faster and requires less effort than traditional manual design and analysis approaches. Bénazéra *et al*. (2002) describe the application of these techniques to uncertain hybrid concurrent systems. Lucas *et al*. (2004) describe an innovative use of these techniques to model what is happening to a patient who has an artificial heart pacemaker. This work has been employed in the production of the first fully digital pacemaker.

3.7 MBS provide many opportunities for re-usability

Once the model-based description of a component is created, it can be used in many system configurations. MBS build system descriptions from the composition of many sub-models. This strategy is the key to effective reusability. An excellent example of how effective this can be is given by the automotive industry where many variants of the same system might be deployed on individual vehicles because of the provision of customer choice. These variants on the basic design are desirable, but very expensive to support. Previous approaches to performing design analysis and generating diagnostics for individual variants have a high cost of development for the different versions of the variant subsystems. A MBS automatically generating the diagnosis makes this practical and hence opens up a whole new area of commercial opportunity (Struss & Price, 2003).

3.8 FMEA generated from the design description and component models

FMEA involves engineers in exploring the effects of every possible failure that can happen to a system. It involves the engineers in analysing the behaviour of a system under many different conditions. This analysis can be very effectively automated for an automobile's electrical systems by the use of model-based reasoning, based on the design description of the electrical systems (Price & Taylor, 2002; Struss & Price, 2003). This automation saves considerable man effort and is more accurate and complete than when FMEA is performed manually. Model-based FMEA is now a standard module offered by design company Mentor Graphics.

3.9 QR models in the educational context

QR can be used to simulate systems for students so that they can understand the errors they have made and how a system should function (de Koning *et al*., 2000; Bredeweg & Forbus, 2003).

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

3.10 QR to help decision making under uncertainty

QR can be used where numerical approaches are not applicable. For example, in marketing, knowledge is imprecise and often unknown. The evaluation of credit risk of companies and classifying the profiles of consumers can be achieved using qualitative descriptors (Flores *et al.*, 2001). A similar situation exists for ecological knowledge (Salles & Bredeweg, 2006).Another application of this kind of technology is in decision support systems used for crime investigation, where it is necessary to reason with uncertain and incomplete models. This can be achieved by instantiating and composing abstract component events of typical crime scenarios to uncover unanticipated cases (Shen *et al.*).

4 Visions for the future

This section presents a range of challenging potential applications which are more advanced than we are capable of achieving at present, and which depend on model-based reasoning for successful execution. The presentation of these as *20 year visions* is somewhat arbitrary – depending on the amount of money and effort invested, some of these visions may occur sooner or later than that. The visions highlighted in the section are as follows:

- the science-bot – automated education;
- virtual vehicles – from conception to recycling;
- understanding and managing complex natural systems;
- interpretation of 4D medical data;
- robust autonomous problem solvers in the face of uncertain situations;
- encyclopaedia of human mental models.

4.1 The science-bot: automated education

4.1.1 Scenario

A science-bot is an interactive agent that is knowledgeable about a set of topics in science. Each science-bot is specialized in its own area of expertise. It will have considerable amounts of domain knowledge and be able to assist learners in helping them to acquire knowledge, understanding and awareness. Science-bots will recognize and know the informational needs of their learners and users and adjust the communicative interaction so it is appropriate to the specific user. In addition, they will have their own teaching and communication goals depending on the circumstances in which they have been placed. Specifically, science-bots will be able to discuss topics from multiple perspectives, explain phenomena and criticize ideas and thoughts presented to them.

Tutoring and training was one of the earliest applications of model-based reasoning (e.g. Brown *et al.*, 1982; Hollan *et al.*, 1984; Wenger, 1987). At present, there are several types of model-based tools available for use in educational settings. Examples of these typically take the form of model-building environments (using the idea of ''learning by knowledge articulation'') and interactive simulations, and they deal with a variety of issues. For surveys of QR and education, see, for example, Forbus (1996), Bredeweg & Winkels (1998) and Bredeweg & Forbus (2003).

4.1.2 Applicability of MBS&QR

MBS&QR technology is of great importance for developing, strengthening and further improving education and training on topics dealing with systems and their behaviours. Educators and learners need the means to capture and share conceptual knowledge. That is, they need a means to formally represent (and automate reasoning with) knowledge that is qualitative, incomplete, fuzzy and uncertain, and in communicative interactions frequently expressed verbally and diagrammatically. Not being able to sufficiently represent this knowledge in a computer-processable format, preserving its unique characteristic, hampers the sharing and communication of insights and

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

theoretical developments. This is particularly a problem in education and training situations. QR technology can provide computer-based facilities to represent and reason with this kind of conceptual knowledge. However, MBS&QR technology is not well known to a wider audience and there are currently not many ready to use the products and tools available to exploit the capabilities of this technology. As result, the full potential of qualitative models as a key component of tutoring systems and interactive learning environments is still to be established.

We envision that the following products can and should be developed in order to address the need for educational software dealing with learning about systems and their behaviour. Interactive articulation devices are model-building environments that allow learners to articulate knowledge (conceptual models) and by doing so learn about a domain. Learning by modelling using traditional approaches has been shown to be effective for enhancing student understanding, but is often hampered by the mathematical complexity of knowledge representations and the lack of means to represent causal knowledge. QR has the capacity to overcome these hurdles. Based on MBS&QR technology, tools can be developed that will allow diagrammatic sketching of ideas and conceptual knowledge and have this automatically transformed into simulations. In order to be effective, such environments should also have the means to criticize models and simulations, and help learners with debugging them.

The concept of autonomous science-bots further advances the idea of individualized support, by the building of resources of previously defined models/model parts and coaching. Science-bots focus on knowledge transfer related to institution-defined goals (where the institution might be a university or school, etc.). Autonomous training-bots are a special class of science-bots. They operate side-by-side with workers (for instance, in factories or business-oriented environments) providing online help and also support for these workers with performing their tasks. MBS&QR technology can provide the basis for developing such tools.

4.2 The virtual vehicle: from conception to recycling

4.2.1 Scenario

Vehicle manufacturers and their suppliers face increasingly serious challenges. The complexity and sophistication of vehicles is growing, and so it is becoming harder to predict interactions between vehicle systems, especially when failures occur. Legal regulations and the demand for safety also impose strong requirements on the detection and identification of faults, and the prevention of their effects on the environment or dangerous situations for passengers and other people. Finally, customer satisfaction is important in order to remain competitive, and this means that the manufacturer must minimize break-downs and reduce maintenance time and the number of misdiagnoses.

The cost of meeting these challenges for a new vehicle model has increased over time, and is becoming overwhelming, both in terms of manpower and elapsed time. In response, vehicle manufacturers have gradually moved towards virtual prototyping and automated analysis. Virtual prototyping involves using software to construct a model of a system, and testing the model works correctly, thereby reducing the need for actual prototyping. This process can be significantly improved by automated analysis, having software performing analysis on the models – for example, failure modes and effects analysis – so that the engineers need to spend less time analysing the system.

The ideal end point of this activity would be the virtual vehicle – a model of the complete vehicle that can be developed and used throughout the lifetime of the vehicle. When it is first decided to make a new vehicle, then the requirements can be used to build a functional model of what the vehicle will be required to do. This might allow automatic specification of much of the complex equipment in the vehicle. As the design is fleshed out by the engineers, either stipulating physical components or specifying the aesthetic aspects of the vehicle (which will constrain design choices), then the extra information should be incorporated into the model of the vehicle from databases of

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

component models. When enough information becomes available, it will be possible to perform model-based tasks of the type described earlier: FMEA, system simulation, diagnosability analysis, production of diagnostics, generation of control software, etc. As variants of the new vehicle design are produced, all this work can be repeated with much less effort, reusing all information that can be used from the original model. When the vehicle is finally disposed of, the virtual vehicle can be used to plan disassembly and efficient disposal of materials.

4.2.2 Applicability of MBS&QR

The achievement of this scenario is far from trivial. At present, many of the models being created are useful for a single task, at a single point in the vehicle's lifecycle. Model-based reasoning is a vital technology for the virtual vehicle, and Struss & Price (2003) list a large number of vehicle manufacturers who are employing it. The use of compositional models makes it possible to automate the repeated reasoning on a design which is necessary for this kind of work. In particular, QR has an important contribution in enabling early analysis before all information is available, and also in focusing numerical reasoning to obtain more specific results. One issue that will become more important is the ability to reason as effectively as possible about a system where different subsystems are specified with different degrees of detail – perhaps only a qualitative model exists for one subsystem, a functional model for several others while one or two subsystems can provide detailed numerical models. Combining these different levels of information is not possible at present, but will become vital if the virtual vehicle is to be realized.

4.3 Understanding and managing complex natural systems

4.3.1 Scenario

We wish to understand the mechanisms and rates of change for complex natural systems where much data are available, but good models are not. Traditional machine learning techniques can produce models of such systems, but they do not provide models where the mechanisms in the domain are visible and capable of explanation.

For example, *Plasmopara viticola* is a fungal pathogen that develops on grape vines. We might be interested in the germination dynamics of the oospores of *Plasmopara viticola*, in response to both endogenous factors, either metabolic (e.g. the influence of the calcium ion) or genetic, and to exogenous factors due to the climate (e.g. water availability) and environment. A deep comprehension of such complex interactions is essential for a rational and optimized treatment planning of vine plants with a consequent benefit for the health of both consumers and operators, and for the impact on the ecosystem. The available pathophysiological knowledge on the endogenous mechanisms at work is qualitative but highly incomplete whereas the exogenous factors can be completely and quantitatively known. Moreover, the mechanisms involved may occur at different time scales. There is the need for the development of proper QR-based modelling methods that are capable of dealing with different levels of knowledge and, even more importantly, with different time scales.

4.3.2 Applicability of MBS&QR

Models of the dynamics of natural systems offer potential benefits to the deep comprehension of the system under study as well as to the performance of specific tasks. The dynamics of such systems result from complex interacting mechanisms, and are very often regulated by both endogenous and exogenous factors. Unfortunately, the available knowledge of the underlying mechanisms is very often highly incomplete, and identifying mechanisms with quantitative methods is a challenging prospect. This makes the modelling problem quite hard to solve, and even unsolvable when, as can occur for natural systems, the available observational data set is inadequate.

QR methods properly integrated with quantitative methods could overcome the identification problems outlined above. An example of a successful early application of a QR-based hybrid

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

method to solve serious identification problems deals with the identification of the intracellular Thiamine (vitamin B1) kinetics in intestinal tissue (Bellazzi *et al*., 2001). Understanding this system is quite important, as Thiamine is one of the basic micronutrients present in food and essential for health; it participates in carbohydrate metabolism, in the central and peripheral nerve cell function and in the myocardial function, and its deficiency causes beriberi with peripheral neurologic, cerebral and cardiovascular manifestations.

4.4 Interpretation of 4D medical data

4.4.1 Scenario

One of the most stimulating application domains where QR can fruitfully support traditional quantitative techniques in the investigation and comprehension of complex phenomena is Electrocardiology. In present clinical practice, information about the heart electrical activity is routinely gathered through Electrocardiographs (ECGs), which record electrical potential from just nine sites on the body surface. However, thanks to the latest technological advances, body surface potential maps are becoming available, as well as epicardial maps obtained non-invasively from body surface data through mathematical model-based reconstruction methods. This 3D data is gathered over time, giving a 4D data set. Electrocardiographic maps can capture a number of electrical conduction pathologies (e.g. arrythmias, Wolf–Parkinson–White syndrome) that can be missed by ECG analysis, but the interpretation of such maps requires skills that are possessed by very few experts.

4.4.2 Applicability of MBS&QR

An important role in the process of defining an interpretative rationale for ECG maps can be played by QR methodologies for spatial/temporal reasoning: (i) to support the expert in identifying salient features in the map; and (ii) to achieve the long-term goal of automating map interpretation to be used in a clinical context. QR approaches based on spatial aggregation (Bailey-Kellogg & Zhao, 2003) can be used to identify patterns and salient features in epicardial activation isochronal maps (Ironi & Tentoni, 2003a). In this kind of map, the time at which each point starts activating, derived from the electrical data of a whole heart beat, is visualized by means of isocurves. A great deal of information about the excitation wavefront structure and propagation can be summarized in a single such map, since isocurves represent subsequent snapshots of the travelling wavefront.

Breakthrough location, high- and low-velocity pathways and conduction block regions, for example, are salient features that characterize the heart electrical activity: they visually correspond to specific geometric patterns to be identified in the map, such as minima location, maximum and minimum elongation directions in the isocurve shapes (Ironi & Tentoni, 2005).

Spatial aggregation approaches, designed for the interpretation of numeric fields that are spatially represented, and capable of identifying global patterns and capturing structural information about the underlying events, exist in the literature. However, such methods only consider 2D geometrical domains that can be discretized by a uniform mesh, and given the complexity of the geometry of the heart (3D and non-uniform meshes), such methods are not directly applicable to the interpretation of cardiac maps: significant extensions are necessary to deal with non-uniform meshes, and maps have to be built on reference surfaces, usually epicardial/endocardial surfaces. Therefore, to properly and efficiently cope with the problem complexity, there is the need for the development of methods capable of dealing with 3D complex geometries over time.

Besides helping medical research in the important phase of the definition of interpretative rationales through models and their simulation, QR methods should lead to the automated interpretation of numerical fields in specific medical domains, and therefore to the development of tools that could eventually enter clinical practice.

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

4.5 *Robust autonomous problem solvers in the face of uncertain situations*

4.5.1 *Scenario 1*

Satellite systems need to make decisions no matter what information is available. A satellite system has constructed a plan of how to achieve its goals. However, the key infrared sensor is not responding. Using its MBS a new plan is constructed. It then uses a qualitative simulation to verify that the plan meets the goals. The simulation also generates expectations which can be used to monitor the execution of the plan to detect problems. The satellite executes its plan, matching available sensory data to expected measurements, and completes the mission.

4.5.2 *Scenario 2*

An autonomous planetary rover, comparing its limited sensory data to a prediction of the sensor readings, detects an inconsistency as it moves down the side of a shallow crater. It uses a MBS composed of models of each component to determine that a component has failed. Even if a sensor is lost, it needs to plan what it will do to complete the mission. It then reconfigures itself and re-plans the mission with its new system structure. Its sensory data now match the predictions of its MBS and it reaches the crater floor to continue its explorations.

These are two situations where autonomous decision making is needed by a system. There are others outside of the planetary exploration domain: robots in hostile environments, or building maintenance systems where a human supervisor is not continually present.

Autonomy requires a global *perception*→*state identification*→*action* loop. This is essential in order to provide the system with adaptable behaviour to face unknown events. Fault Detection Identification and Reconfiguration (FDIR) involves a set of functions, which are obviously crucial to adaptability.

4.5.3 *Applicability of MBS&QR*

Model-based diagnosis (MBD) techniques can undoubtedly benefit the overall spacecraft and rover design process. These tools provide an integrated development framework able to produce diagnostics that are the equivalent to the currently used on-board FDIR systems, and that also provide substantial additional benefits when going from the development step to the operation step:

- the FDIR design will be easier to build, reusable and more generic when based on MBD;
- MBD enables a global and unified management of the equipment and functional levels;
- the models can provide support for validation;
- MBD should lead to a decreased level of false alarms by making maximum use of redundancies and numerous non-telemeasured on-board observations;
- MBD should be able to handle more situations without human intervention than the current FDIR systems, enabling the satellite or rover to transit to *safe mode* and consequently increasing availability.

The Remote Agent eXperiment conducted in May 1999 by NASA demonstrated the applicability of MBS&QR by running the Livingstone software on board the Deep-Space-One satellite. State identification and reconfiguration were based on a description of the behaviour of the components in each of its operating modes in terms of simple qualitative constraints in propositional logic (Williams & Nayak, 1996, 1997). This work was extended later to hybrid modelling by Bénazéra & Travé-Massuyès (2003) and Bénazéra *et al.* (2002) in collaboration with the French Space Agency CNES. Complementing qualitative models with continuous models brought important benefits such as greater flexibility for the modelling and more precise fault detection when needed. Based on the hybrid formalism, a new set of algorithms interleaving the search for the most probable diagnoses with consistency checking were defined and implemented in the KOALA software prototype (Bénazéra, 2003).

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

Techniques for autonomy will offer new possibilities for the development of spacecraft missions by helping engineers automatically produce a large part of the ground and on-board software, as well as offering great help for the hardware specification and the design of the most useful on-board sensors and telemeasures. Space engineers should be able to produce more complex constellations and spacecrafts for difficult exploration or critical missions.

4.6 Encyclopedia of human mental models

4.6.1 Scenario

Understanding the nature of human knowledge is one of the central problems of cognitive science. Building up an encyclopaedia of the kinds of mental models people have, across a broad range of domains, levels of expertise and cultures, would be an invaluable resource. For example, one can view learning an area as a trajectory through the space of models that one might have for that area, so understanding the kinds of models that people have and how they change over time can shed light on the nature of human learning. In addition to supporting scientific research, such a compendium would be a valuable resource for systems that interact with people, including educational software and assistant systems, by knowing what kinds of models users with different backgrounds are likely to have.

4.6.2 Applicability of MBS&QR

Existing mental models research (cf. Gentner & Stevens, 1983; Collins & Gentner, 1987) suggests that QR plays an important role in human mental models. To be sure, human reasoning is often grounded in perception (Forbus, 1994), and quantitative answers are often required (Paritosh & Forbus, 2005). Nevertheless, given the paucity of data that people reason with regularly, compared to what is required for the quantitative modelling of physical phenomena, some form of QR appears necessary to explain everyday human reasoning. The analysis of protocols suggests that many of the representations developed by the QR community can be used to capture important aspects of human mental models. However, two kinds of questions remain open. The first category of questions concerns coverage: How far do today's QR models go in capturing human mental models reasoning? What other kinds of knowledge are needed for a complete account? Very few studies have focused on early stages of cognitive development, for example, so there could be some interesting surprises there. The second category of questions concerns the details of human QR. There are several aspects of todays qualitative simulation algorithms that appear to be psychologically unrealistic, including the amount of branching and the need for extremely detailed state descriptions (Forbus & Gentner, 1997). Developing psychologically realistic QR methods will be an important challenge in evaluating models, and hence important to the construction of such an encyclopedia.

Today, protocol analyses remain one of the best ways of uncovering the fine structure of mental models and their developmental trajectory. Unfortunately, such analyses are extremely time consuming and labour intensive. The creation of better tools is critical to the success of this enterprise. The first category of tools is scientist-friendly QR systems. That is, systems which enable scientists to rapidly formulate qualitative models and run them, to compare predictions of alternative models and to compare predictions against human data. In some ways this requirement echoes the need for friendlier environments for education and training emphasized earlier. However, the interface needs for scientists are different: as more technical users, they may require less scaffolding, but they will also demand more powerful tools. Simplifying the comparison of data with models will be important as well.

The second category of tools involves natural language processing, machine learning and data mining. To create a comprehensive encyclopedia by hand would be prohibitive. On the other hand, text-mining and automated analysis of protocols via natural language processing could fundamentally change the economics of constructing such an encyclopaedia. Interactive knowledge capture

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

tools that could interact with people and pose follow-up questions automatically, during an interview, would facilitate capture of expert mental models. Sketch understanding systems (cf. Klenk *et al.*, 2005) or some other kind of visual reasoning system will be needed to interactively capture spatial mental models. Developing standardized representation conventions will facilitate the interchange of models and the process of characterizing and cataloging them. Developing automatic classification and induction algorithms that could manage this data-intensive process would be another factor making such an enterprise practical.

5 Technological priorities

The previous section has presented a number of targets for the application of MBS&QR. This section will consider what improvements to available MBS&QR technology are needed in order to achieve those targets.

One issue which will not be addressed any further in this section is that of cross-discipline research. While we believe that MBS&QR have a key contribution to make to the realization of the visions in the previous section, they are not MBS&QR problems *per se*. The science-bot, for example, will also need advances in analogical reasoning and user modelling, in order to be able to work in the way that is outlined.

However, for all of the visions described, we consider that MBS&QR are central to the efficient production of workable systems, and this section explores the technological improvements in MBS&QR needed in order to realize the visions described above.

5.1 More powerful modelling formalisms/frameworks

Many of the processes that we are modelling evolve over time, happen in a particular space and are impossible to specify completely as not all relevant parameters can be determined (giving rise to uncertainty). In addition, the lack of precise data makes it impossible to describe the system quantitatively. Many real-world systems are very complex and, while the exact nature of the complexity varies from system to system, the contributors to the degree of complexity are as follows: nonlinearity, order, dimensionality, degree of coupling and non-determinism. Further research is needed in more powerful modelling languages, in coupling models at varying levels of abstraction and in developing spaces of models from which an appropriate model can be selected.

5.2 QR methods using more sophisticated mathematics

In many cases, methods from MBS&QR build upon existing mathematical methods from calculus (e.g. differential equations), algebra (equations, functions and sets) and logic (Travé-Massuyès *et al.*, 2003). The basic methods are geared towards the area of MBS&QR: (1) by restricting the domains and co-domains of functions to be discrete, possibly ordered, instead of being continuous, and the results are then still consistent with the underlying axioms; (2) by adding task-specific problem-solving methods, such as methods for diagnosis, which are able to act on particular representations in a particular fashion. There are many mathematical methods which are restricted in their practical usefulness because qualitative versions of them are as yet not available.

5.3 Integration of models from different domains

In many situations it is necessary to consider phenomena with different natures in order to reason about a system. In the field of continuous industrial processes, many devices, such as pumps, comprise phenomena related to hydraulics and mechanics. In the automotive industry cars comprise different inter-related subsystems such as hydraulic, electric and electronic.

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

Integration of models is an open problem, and further research on this topic is closely related with research on ontologies. Some work has been done using domain-independent ways of modelling such as bond graphs, although that work has not been as successful as might have been expected. One of the reasons may be that bond graphs are well suited to simulation, but less adapted for the other tasks performed by MBS.

It may be that a combination of appropriate methodologies for individual domains, plus the development of standards in an integrated manner for interfacing models in different domains, may finesse this problem, but at present it is still an open problem.

5.4 *Models of software*

The modelling of the action and influence of software is an issue for almost any advanced man-made device or system. For example, in the automotive domain, electronic control units (ECUs) containing many thousands of lines of software control the state of vehicle subsystems, and often perform monitoring, diagnosis and reconfiguration of systems. It is necessary to incorporate the actions performed by software in models in order to understand the state of the device and perform device-specific tasks. Similar issues occur in other domains, such as model-based reasoning about process control systems.

5.5 *Models to represent system specifications and requirements*

One of the major advantages of model-based reasoning for problem solving is that it can consider many more possible scenarios than a human could. One of the key concepts for qualitative MBS is that of an *envisionment* (Forbus, 1990). An envisionment is a map of all of the possible states that a given system can reach, and how the system moves from one state to another. It is generated by exhaustive simulation from all states to see what other states can be reached. For systems where safe operation is an issue, an envisionment provides important indications of the possibility of reaching unsafe states or situations. In other types of applications it might be possible to specify "interesting" states of a different sort. In order to identify interesting/unsafe states, two things are needed.

- Descriptions of what is interesting. This can involve capturing descriptions of the way in which the system should work, and might include issues of complex dynamic time varying and continuous systems, dealt with elsewhere in this section.
- Abstraction of state descriptions. It must be possible to abstract the results of an envisionment so that they can be compared with the descriptions of interesting states.

This area is in its infancy, but we would expect it to make a significant contribution to system safety and reliability as it becomes better developed.

5.6 *Hybrid modelling*

Different modelling techniques allow the capture of different aspects of the same phenomenon. Hence, in order to include different aspects of the same phenomenon or even different phenomena in one model, it is necessary to integrate models from different sources.

- Pure qualitative models allow one to focus on significant behaviours, while pure numerical models allow one to detail each one of these behaviours or even to solve ambiguities related to the QR.
- Causal models and models based on quantitative differential equations provide two different views of the same phenomenon.

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

Few systems are capable of combining different modelling approaches. Currently, the main research effort is devoted to the production and use of semi-qualitative models. In the future, MBS need to be able to combine different types of models to solve a given problem. The target to be achieved might be the kind of reasoning displayed by human experts, who seem able to combine information from different types of model seamlessly, and to combine information from partial models of each type.

5.7 Multi-level modelling

It is necessary to combine models at different levels of abstraction to solve a particular problem, usually to cope with complexity and multiple scales.

- In the food industry, an evaporation station can be modelled, at least, at three different levels: simple mass balances (product conservation), detailed balances (mass and energy conservation) and detailed dynamical model for control purposes.
- In the computer industry, a computer can be viewed at different levels, from high-level functional components to chips.

Currently, there are different theoretical proposals: automated handling of diagnosis hypotheses, multiple models considering available time for diagnosis and multiple levels of abstraction regarding the quality of the diagnosis. However, there is almost no application on industrial systems capable of handling models at different levels of abstraction, because there is no systematic way to share results from different models within the same task.

A real applicable methodology needs to be proposed to change smoothly from one level to another, exploiting results from different levels.

Eventually the reasoning system should be able to select the adequate level of abstraction automatically, and to switch from one to another as required.

5.8 Combining qualitative and functional models

Much qualitative research has concentrated solely on reasoning about the structure and behaviour of systems. For many applications, it is necessary to abstract the results in terms of the function or teleology of the system. That implies being able to represent teleological knowledge, to reason about it and to map behavioural knowledge to it. This has been done for systems with fairly static behaviour (Price, 1998). That work needs to be extended to cover complex, dynamic time-varying functions.

5.9 Automated model generation from simulation models

The exploitation of MBS in industry will greatly depend on the (additional) modelling efforts they require. This leads us to the attempt to reducing these efforts by automated conversion of existing simulation models into abstract models suited for model-based problem solvers.

Simulation models of a system are often created for control purposes. However, for diagnosis, for example, specific properties are needed from a model.

There is a need for methodologies that ensure that the models are built correctly in the first place for use in tasks other than simulation, and techniques that facilitate the process of converting those models to be appropriate for diagnosis and other tasks.

5.10 Derivation of qualitative models from requirements

During the design process, the correct operation of a system is often described at a high level, perhaps in terms of state charts. Such information is often very useful when performing

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

model-based reasoning. Better methods are needed to integrate such descriptions of requirements into the construction of MBS.

5.11 *Automated modelling*

Automated model building and model transformation need continued theoretical work and more effective and efficient algorithms (Ironi & Tentoni, 2003b). This is emphasized by application requirements. Much of the expected gain depends on the fast and economic creation of models from a library. Since different tasks may 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. QR needs to develop techniques to generate task-oriented models from generic models.

This also touches upon a more general goal, namely integrating QR results and techniques with standard engineering practice and tools. The lack of integration presents a major obstacle to transferring QR technologies into industrial practice. Deriving qualitative models from numerical models that have been developed (e.g. in the phase of design verification) is of high practical importance. However, it may require changes in current modelling practice towards modular, component-oriented models. The need to blend in with current practice also applies to other domains, such as medicine, economy, biology and ecology.

5.12 *Model-based system identification*

Model building is a difficult and time-consuming process. A much more efficient alternative to building models by hand would be to learn models from observed data. This is still a very difficult machine learning challenge for complex domains. QR can help with this in two ways.

- Learning qualitative models.In domains such as some areas of biology, where the underlying models may not be known, it will be possible to learn qualitative models from data. Early research in this area indicates that it is more useful to build qualitative models rather than numerical models at this stage, in order to facilitate understanding by domain experts. Good examples of this are the work of King *et al*. (2005), which identifies the dynamic qualitative model structure of biological systems from sparse qualitative or quantitative data, and that of Bratko & Suc (2003), which looks for qualitative patterns in numeric data.
- Deriving quantitative models from qualitative models. In domains where qualitative models are known, but are not executable, QR provides graphical ways of building executable models, and makes clear the assumptions behind the models, enabling domain experts to compare their models on a like-for-like basis. Where models are known and data are available, this type of work should make it possible to develop accurate numerical models with known assumptions and limitations (Bellazzi *et al*., 2000; Guglielmann & Ironi, 2005).

The main impact of these techniques may well be in science rather than in engineering, providing tools for scientists to understand the world better and having a dramatic impact on the way that we carry out scientific research.

5.13 *Conversion of qualitative models*

One issue concerns the development of better engineered and easy-to-use tools that facilitate the exchange of results among researchers and make QR techniques available to potential users in other areas and application work. The field, so far, 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

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

principles and systems. If the field can make progress on this, it will become easier to create and exchange libraries of reusable models.

6 Conclusions

MBS&QR has had a productive and useful childhood. As it enters its third decade, both the challenges for the technology and the promise it holds are greater than ever. The visions in this paper are intended as encouragements to researchers in the field. We are engaged in an enterprise with immense potential benefits. In order to succeed, we need developers who will drive work on the difficult applications described here, addressing problems of integration with other technologies and commercial issues, as well as researchers who will address the scientific issues that need to be resolved in order to build these visionary systems.

Acknowledgements

This paper is based on work carried out by the MONET European Network of Excellence in Model-based and Qualitative Reasoning (`http://monet.aber.ac.uk`). The authors listed were those involved in production of the final technological roadmap produced by the network. However, we also acknowledge the contribution of other members of the network, and of those who took part in the *Futures* session at the *18th International Workshop on Qualitative Reasoning* in Evanston, Illinois. A brief version of the paper was presented at the *19th International Workshop on Qualitative Reasoning* in Graz (Price *et al*., 2005).

References

Adam, G and Grant, E, 2001, Qualitative modelling of lime-mash flow processes and measurements. In Mastorakis and Pecorelli-Peres (eds.), *Advances in Systems Science: Measurements, Circuits and Control*. WSES Press, pp. 60–63.

Bailey-Kellogg, C and Zhao, F, 2003 Qualitative spatial reasoning: extracting and reasoning with spatial aggregates. *AI Magazine* (special issue on Qualitative Reasoning) **24**(4), 47–60.

Bellazzi, R, Guglielmann, R and Ironi, L, 2000, How to improve fuzzy-neural system modelling by means of qualitative simulation. *IEEE Transactions on Neural Networks* **11**(1), 249–253.

Bellazzi, R, Guglielmann, R, Ironi, L and Patrini, C, 2001, A hybrid input–output approach to model metabolic systems: an application to intracellular thiamine kinetics. *Journal of Biomedical Informatics* **34**(4), 221–248.

Bénazéra, E, 2003, Diagnosis and reconfiguration based on hybrid concurrent models – application to autonomous satellites. PhD thesis, Université Paul Sabatier, Toulouse, France.

Bénazéra, E and Travé-Massuyès, L, 2003, The consistency approach to the on-line prediction of hybrid system configurations. In *Proceedings of the 2003 IFAC Conference on Analysis and Design of Hybrid Systems (ADHS'03), Saint-Malo, France*.

Bénazéra, E, Travé-Massuyès, L and Dague, P, 2002, State tracking of uncertain hybrid concurrent systems. In *Proceedings of the International Workshop on Principles of Diagnosis DX02*, pp. 106–114.

Berleant, D and Kuipers, B, 1997, Qualitative and quantitative simulation: bridging the gap. *Artificial Intelligence Journal* **95**(2), 215–255.

Bobrow, D (ed.) 1985, *Qualitative Reasoning About Physical Systems*. Amsterdam: North-Holland (originally published as *Artificial Intelligence* **24**, 1984).

Bratko, I and Suc, D, 2003, Learning qualitative models. *AI Magazine* (special issue on Qualitative Reasoning) **24**(4), 107–119.

Bredeweg, B and Forbus, K, 2003, Qualitative modelling in education. *AI Magazine* (special issue on Qualitative Reasoning) **24**(4), 35–46.

Bredeweg, B and Winkels, R, 1998, Qualitative models in interactive learning environments: an introduction. *Interactive Learning Environments* **5**(1–2), 1–18.

Brown, JS, Burton, R and de Kleer, J, 1982, Pedagogical and knowledge engineering techniques in SOPHIE I, II and III. In Sleeman, D and Brown, JS (eds.), *Intelligent Tutoring Systems*. New York: Academic Press, pp. 227–282.

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

Coghill, G, Garrett, S and King, R, 2004, Learning qualitative metabolic models. In de Mantaras, RL and Saitta, L. (eds.), *Proceedings of the 16th European Conference on Artificial Intelligence*. Amsterdam: IOS Press, pp. 445–449.

Collins, A and Gentner, D, 1987, How people construct mental models. In Holland, D and Quinn, N. (eds.), *Cultural Models in Language and Thought*. Cambridge: Cambridge University Press, pp. 243–265.

de Kleer, J, 1977, Multiple representations of knowledge in a mechanics problem-solver. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence (IJCAI-77)*. San Francisco, CA: Morgan Kaufmann, pp. 299–304.

de Koning, K, Bredeweg, B, Breuker, J and Wielinga, B, 2000, Model-based reasoning about learner behaviour. *Artificial Intelligence* **117**(2), 173–229.

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

Flores, J, Flores, B and Gonzalez, F, 2001, Qualitative/quantitative financial analysis. In *Proceedings of the 15th International Workshop on Qualitative Reasoning (QR01)*.

Forbus, K, 1988, Intelligent computer-aided engineering. *AI Magazine* Fall, 23–36.

Forbus, K, 1990, The qualitative process engine. In Weld, D and de Kleer, J. (eds.), *Readings in Qualitative Reasoning about Physical Systems*. San Francisco, CA: Morgan Kaufmann, pp. 220–235.

Forbus, K, 1994, Qualitative spatial reasoning: framework and frontiers. In Glasgow, J, Narayanan, H and Chandrasekaran, B (eds.), *Diagrammatic Reasoning : Computational and Cognitive Perspectives*. Menlo Park, CA: AAAI Press, 1995.

Forbus, K and Gentner, D, 1997, Qualitative mental models: simulations or memories? In *Proceedings of the 11th International Workshop on Qualitative Reasoning (QR97), Cortona, Italy*.

Forbus, K, 1996, Qualitative reasoning. In Tucker, AB (ed.), *The Computer Science and Engineering Handbook*. Boca Raton, FL: CRC, pp. 715–733.

Forbus, K and Falkenheiner, B, 1991, Compositional modelling: finding the right model for the job. *Artificial Intelligence* **51**, 95–143.

Forbus, K and Kuehne, S, 2005, Towards a qualitative model of everyday political reasoning. In *Proceedings of the 19th International Workshop on Qualitative Reasoning (QR'05)*, pp. 157–162.

Gentner, D and Stevens, A (eds.), 1983, *Mental Models*. Hillsdale, NJ: Lawrence Erlbaum.

Guglielmann, R and Ironi, L, 2005, *Generating Fuzzy Models from Deep Knowledge: Robustness and Interpretability Issues (Lecture Notes in Artificial Intelligence, 3571)*, pp. 600–612.

Heller, U and Struss, P. 1997, Transformation of qualitative dynamic models – application in hydro-ecology. In *Proceedings of the 11th International Workshop on Qualitative Reasoning*, pp. 83–92.

Hollan, J, Hutchins, E and Weitzman, L, 1984, STEAMER: an interactive inspectable, simulation based training systems. *AI Magazine* **5**, 15–27.

Ironi, L and Tentoni, S, 2003a, *Towards Automated Electrocardiac Map Interpretation: An Intelligent Contouring Tool Based on Spatial Aggregation (Lecture Notes in Computer Science, 2810)*, pp. 397–408.

Ironi, L and Tentoni, S, 2003b, A model-based approach to the assessment of physicochemical properties of drug delivery materials. *Computers & Chemical Engineering* **27**(6), 803–812.

Ironi, L and Tentoni, S, 2005, *Electrocardographic Imaging: Towards Automated Interpretation of Activation Maps (Lecture Notes in Artificial Intelligence, 3581)*, pp. 323–332.

King, R, Garrett, S and Coghill, G, 2005, On the use of qualitative reasoning to simulate and identify metabolic pathways. *Bioinformatics* **21**(9), 2017–2026.

Klenk, M, Forbus, K, Tomai, E, Kim, H and Kyckelhahn, B, 2005, Solving everyday physical reasoning problems by analogy. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05), Pittsburgh, PA*, pp. 209–215.

Kuipers, B, 1986, Qualitative simulation. *Artificial Intelligence* **29**, 289–338.

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

Lucas, P, Kuipers, R and Feith, F, 2004, A system for pacemaker treatment advice. In *Proceedings ECAI 2004*, pp. 735–739.

Miguel, I and Shen, Q, 2005, Exhibiting the behaviour of time-delayed systems via an extension to qualitative simulation. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* **35**(2), 298–305.

Paritosh, P and Forbus, K, 2005, Analysis of strategic knowledge in back of the envelope reasoning. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05), Pittsburgh, PA*, pp. 651–656.

Price, C, 1998, Function directed electrical design analysis. *Artificial Intelligence in Engineering* **12**(4), 445–456.

Price, C, Snooke, N and Lewis, S, 2003, Adaptable modelling of electrical systems. In *Proceedings of the 17th International Workshop on Qualitative Reasoning (QR'03)*, pp. 147–153.

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press

Price, C and Taylor, N, 2002, Automated multiple failure FMEA. *Reliability Engineering and System Safety* **76**(1), 1–10.

Price, C. *et al*., 2005, Qualitative futures. In *Proceedings of the 19th International Workshop on Qualitative Reasoning (QR'05)*, pp. 29–37.

Salles, P and Bredeweg, B, 2006, Modelling population and community dynamics with qualitative reasoning. *Ecological Modelling* **195**(1–2), 114–128.

Shen, Q, Keppens, J, Aitken, C and Schafer, B. Towards scenario driven decision support systems for serious crime investigation. *Law, Probability and Risk*. Oxford: Oxford Press.

Struss, P and Price, C, 2003, Model-based systems in the automotive industry. *AI Magazine* (special issue on Qualitative Reasoning) **24**(4), 17–34.

Travé-Massuyès, L, Ironi, L and Dague, P, 2003, Mathematical foundations of qualitative reasoning. *AI Magazine* (special issue on Qualitative Reasoning) **24**(4), 91–106.

Travé-Massuyès, L and Milne, R, 1997, TIGER™: gas turbines condition monitoring using qualitative model based diagnosis. *IEEE Expert Intelligent Systems & Applications* **12**(3), 21–31.

Trelease, R and Park, J, 1996, Qualitative process modelling of cell–cell–pathogen interactions in the immune system. *Computer Methods and Programs in Biomedicine* **51**, 171–181.

Ward, D and Price, C, 2001, System functional safety through automated electrical design analysis. *SAE 2001, Transactions. Journal of Passenger Cars: Electronic and Electrical Systems* **110**(7), 341–347.

Wenger, E, 1987, *Artificial Intelligence and Tutoring Systems: Computational and Cognitive Approaches to the Communication of Knowledge*. San Francisco, CA: Morgan Kaufmann.

Williams, B and Nayak, P, 1996, A model-based approach to reactive self-configuring systems. In *Proceedings of AAAI-96*, pp. 971–978.

Williams, B and Nayak, P, 1997, A reactive planner for a model-based executive. In *Proceedings of IJCAI-97*, pp. 1178–1185.

https://doi.org/10.1017/S026988890600097X Published online by Cambridge University Press