

Explanation and Qualitative Reasoning

Anders Bouwer and Bert Bredeweg

Dept. of Social Science Informatics (SWI)

University of Amsterdam, Roetersstraat 15

1018 WB Amsterdam, The Netherlands

{anders,bert}@swi.psy.uva.nl

Abstract

Qualitative Reasoning is often seen as a powerful basis for generating explanations, because the behaviour of interest is explicitly modelled in terms of relevant components, processes, causality relations, quantity spaces, assumptions, states and transitions, while neglecting unnecessary details like quantitative values. However, the link between qualitative reasoning and explanation is often seen as a direct one-to-one mapping, whereas studies of human explanation indicate that this is a simplification. Explanation is an interactive process in which the context plays an important role. This position paper takes a closer look at the relation between qualitative reasoning, explanation generation and contextual factors such as the tasks and goals of the user, and the dialogue history.

1 Introduction

Explanation plays a major role in human communication, especially in learning, in both educational and business settings. When a computer system is intended to represent some knowledge, and this knowledge needs to be communicated to a user, facilities are necessary to generate appropriate explanations.

Qualitative reasoning techniques are often considered important for explanation generation, because of the articulateness of the simulation models these techniques deliver. Despite the potential of these techniques, actual use in real applications is limited (see e.g. Trave-Massuyes & Milne 1998). One reason is the poor 'interface' of qualitative reasoning problem solvers, both for supporting model building (input to the simulator) and for inspecting the simulation results (output of the simulator). In this paper we are concerned with the latter. The notion of interface we use is wider than 'just the screen layout of the program'. Interface actually refers to all the aspects relevant to an effective communication between artificial problem solver and user.

Experiments with real users highlight the need for more elaborate communication capabilities. For instance, when interacting with the CyclePad educational system, learners specifying and experimenting with several designs for thermodynamic cycles, seem to require additional coaching facilities, e.g., telling students how to proceed when they get stuck, and explaining the rationales behind design choices

(Forbus & Whalley 1998). As another example, when interacting with the qualitative reasoning shell GARP, using a text-based interface, the amount of information which is generated by running a simulation, or by inspecting a model can be overwhelming, making it hard to discern the important aspects of interest (de Koning 1997). The functionality of both of these simulation environments does not in all cases map directly onto the knowledge needs of their users.

In this paper we argue that the poor interaction capabilities of qualitative problem solvers can be addressed by augmenting these problem solvers with knowledge about their user's needs & goals and the dialogue history. This paper is a position paper by means of which we want to point out the need for the qualitative reasoning community to take into account techniques developed by adjacent communities such as natural language generation and intelligent tutoring systems. The content of this paper is as follows. The next section reviews qualitative reasoning techniques, arguing that these techniques in principle can generate parsimonious simulation models with sufficient articulateness to address user needs. However, these techniques are hampered in doing so, because they lack knowledge about these users. Section 3 discusses concepts and techniques from natural language generation and intelligent tutoring systems, which we believe are becoming increasingly more relevant to the qualitative reasoning community. In section 4 we elaborate on typical research topics that should be addressed in order for the qualitative reasoning community to profit from the insights developed in adjacent research areas. Finally, section 5 summarises and concludes this paper.

2 The Role of Qualitative Models

The role of qualitative models for intelligent systems has been discussed at large in the past (e.g. Forbus 1988; de Kleer 1990). Views range from 'the naive physics point of view' (e.g. Gentner & Stevens 1983), to 'reasoning about system behaviour in the absence of quantitative information' (e.g. de Kleer & Williams 1991). Whatever the specific view one takes, the model-based reasoning community is becoming more aware that one of the key roles for qualitative models is the *ontology* (vocabulary) it provides to have computers reason about system behaviour. This vocabulary is articulate and represents the distinct features of systems and their behaviour such that in

principle the behaviour analysis made by a computer program can be communicated with users in a way that is understandable by these users (Forbus 1997; de Koning 1997; Bredeweg & Winkels 1998). For example, one of the lessons learned from research projects that used 'less rich' knowledge representations (e.g. Sophie - Brown, Burton, & de Kleer 1982, and Steamer - Hollan, Hutchins, & Weizenbaum 1984) was that the ability to construct an ongoing dialogue with a user was severely limited, because the computer program (in this case, a quantitative simulation) lacked the vocabulary 'to talk' about the simulation results produced by the computer program. Particularly, generating explanations turned out to be a fundamental problem.

The qualitative reasoning community has produced many ideas in order to address this 'lack of articulation' problem (e.g. Weld & de Kleer 1990; Forbus & de Kleer 1993). Current techniques allow for the construction of artificial problem solvers that capture many important aspects of system behaviour, including knowledge about the structural composition of a system, behavioural features (particularly causal dependencies) and all kinds of time varying aspects. Of course, still many issues have to be solved, for instance reasoning about space or mechanics.

However, it is also realised that building articulate problem solvers, although required, introduces new problems. The 'information overflow' is probably the most outstanding. If many details are represented into a problem solver, many inferences can be made. This is both a computational problem and a problem for the user of the artefact. The computational problems have been alleviated by the development of truth maintenance systems (e.g. de Kleer 1986).¹ Other approaches have proposed the idea of assumptions (Falkenhainer & Forbus 1991; Rickel & Porter 1997), which address the problem 'of too much detail' both from a computational point of view as well as from the user's point of view. Using assumptions allows the artefact to neglect certain detail and as a result the user will also have to deal with less detail. Still other solutions have focused primarily on the idea of 'summarising' the details produced by a qualitative simulator and by doing so making the simulation results easier to understand for a user (de Koning 1997; Mallory, Porter, & Kuipers 1996). A problem with both approaches (summarising or using assumptions) is that although the simulation results may consist of less detail, there is no guarantee that this result is actually understandable by the user. In fact, most qualitative problem solvers have no idea whatsoever about what their users know, understand or don't know and are not able to understand.

In conclusion, it seems fair to argue that qualitative reasoning originated from problems encountered in trying to have humans (learners) interact with artificial problems solvers in order to discuss systems and their behaviour. It turns out that for such interactions to be effective specific aspects have to be dealt with. First, the qualitative problem solver should have access to an articulate vocabulary that captures all the important issues relevant to reasoning

about system behaviour. Second, because such articulate problem solvers may produce lots of detail, additional structuring and/or summarising mechanisms are required in order to reduce the amount of information detail delivered to the user. Finally, qualitative simulators have no knowledge of their users. This means that in principle the artefact is not able to adjust its problem solving capabilities to the knowledge needs of its users. We believe that this is one of the next boundaries to be tackled by the qualitative reasoning community.

3 Explanation Generation

This section gives an overview of research on explanation generation, with a focus on concepts and techniques from the fields of natural language generation and intelligent tutoring systems. In particular, attention is given to the notions of knowledge needs, text structure, and interaction.

Tasks, Goals, and Knowledge Needs

The main question that should guide explanation generation is: what do users need to know? This depends to a large extent on the tasks and goals of the user. In work by Khan, Brown, & Leitch (1997) on industrial simulation based activities, 'three types of user' are distinguished: those who want to (a) acquire knowledge of how the plant works; (b) develop skills of how to control the plant; and (c) receive advice on how to complete a task. To deal with these different demands for information, Khan et al. identified four fundamental types of explanations, each making primarily use of a certain type of knowledge: instructional (temporal knowledge), decision making (correlational knowledge), justification (causal knowledge), and theoretic (structural knowledge).

Winkels (1992) makes a distinction between *local* and *global needs*. If a user is engaged in a task and encounters a problem because he/she lacks some piece of knowledge or has a misconception, a local need occurs. Identifying a local need is done by analyzing not only the user's query (e.g., 'How do I ...?' or 'Why did that happen?'), but also the context in which it occurs, e.g., by monitoring the user's actions.² To resolve the local need, help can be given to supply the necessary operational knowledge; this is aimed at local repair or remediation, to ensure that correct task performance can continue. A local need may be related to more global needs, or educational goals. If this is the case, help can be expanded into a real teaching interaction, which may involve the introduction of new conceptual knowledge, giving examples and analogies, providing training exercises, and assessment and appropriate feedback.

When the computer system is largely in control of the interaction, the educational goals may be treated in a specific order derived from the way they are organised into a curriculum. When the user is more in control of the interaction, more opportunistic planning is necessary.

²The issue of diagnosing the user's knowledge state, missing knowledge and misconceptions, has received a lot of attention in the literature on intelligent tutoring systems (e.g., see Wenger 1987; de Koning 1997).

¹We will not discuss the computational problems any further in this paper.

Winkels (1992) discusses how the subject matter can be traversed by following didactic relations between concepts in the domain; they distinguish five relations (generalisation-specification, abstraction-concretion, inversion, analogy and identity) which may be automatically determined, if the domain is very structured.

Text Structure

Especially when explanations become longer than a single sentence, explicit means to structure the text are necessary. An important notion is the existence of relations between text segments, or text spans which may contain multiple segments. A very influential idea, from Rhetorical Structure Theory, or RST (Mann & Thompson 1988) is the assumption that a text is coherent when it can be analysed hierarchically in terms of rhetorical relations holding between consecutive text spans. Some research addresses the role of discourse markers, like 'because' or 'finally' in making these relations visible in the text (Knott & Sanders 1998). Starting from RST, Hovy (1993) describes what is necessary for automatic generation of coherent texts: text plans incorporating communicative intent; a collection of discourse relations; predefined structures (schemas); and control of focus shifts. He cites heuristics for sentence formation (p.368), describes aggregation rules (p.369), and also illustrates how text formatting can be integrated into text generation.

Interaction

Many researchers point out that explanation is an interactive process (e.g. Moore 1995; Cawsey 1991), and acknowledge the role dialogue plays in promoting learning, reflection, and the acquisition of scientific inquiry skills (e.g. Baker 1994). Dialogue can be analyzed on several levels of interaction, e.g., in terms of *moves* (which have a function), *turns* (which contain one or more moves by one participant, and have an initiating or responding mode), *exchanges* (which usually consist of a pair of turns – one for each participant) and *episodes* (which have an outcome, e.g., a conclusion agreed upon between both participants) (Pilkington 1997). An explanation consisting of a single turn is often not immediately successful, but requires follow-up interaction to make sure the hearer is satisfied and understands the knowledge communicated (Pilkington 1992).

To be able to deal with follow-up questions, the system should keep track of what has been said by the system and the user – the dialogue history. Moore (1995) argues for including in the dialogue history the intentions of the system together with the plan which resulted in a particular utterance, because this makes it possible to recover from a failed explanation by backtracking to the communicative goal and realising it in another way. An adequate representation of the dialogue history also allows linking new explanations to previous explanations, without repeating what has been said already.

Related to the notion of dialogue history is the user model, or student model, in the case of educational systems. Here, information is kept about the user's profile, goals, tasks, preferences, and knowledge. When this information is available, or can be derived from the dialogue history and/or dia-

gnostic processes, the interaction can be adapted to specific users. Hence, explanations can be generated at a level of detail appropriate to the expertise level of the user, avoiding concepts still unknown to the user. Because a user model represents a user's knowledge and characteristics, it should be updated when changes are appropriate.

Multimedia

Because interactivity is important, the interface between system and user plays a crucial role. Although the previous sections focused on language (in written form), there is also a need for other means of communication and ways of integrating them.

For instance, sometimes people are not able to state explicitly what they want to know, but they are able to identify the thing they don't understand, or want to know more about; in these cases direct manipulation, i.e., pointing at the point of interest (a part of text or a graphic), can be of great help (Moore 1995).

Graphics can be very helpful as part of effective explanation. Especially when structure needs to be conveyed, or large sets of data have to be visualised, a picture can be more illuminating than a verbal description. If both graphics and text were generated and coordinated by the same planning mechanism, even more flexible means of communicating would be possible, which combines the advantages of both media. Interesting work in this area has been done by Wahlster *et al.* (1993) and Kerpedjiev *et al.* (1997).

Animation is useful when a sense of realism is needed, but is often not important in qualitative systems, in which time is abstracted into qualitatively distinct states. So instead of real animation, a collection of distinct graphics, analogous to comic strips or film scenarios, may be better suited.

4 Discussion

After this short overview of some of the currently available techniques in qualitative reasoning, intelligent tutoring systems, and natural language generation, the question is: can we integrate these techniques to produce better explanation facilities, or are there still important pieces missing?

One first issue which may pose problems is the role of the knowledge represented in the system in relation to the task of the user interacting with the system. In qualitative reasoning, the main tasks which have received a lot of attention are prediction and diagnosis, especially in relation to physical systems.³ Most certainly, this has had an influence on the development of the ontological primitives in qualitative reasoning, since both of these tasks typically require knowledge about physical structure and causal relations. If we want to use qualitative reasoning in systems supporting (the learning of) other tasks, like design, planning or assessment, we will have to investigate whether the qualitative reasoning ontological primitives are still satisfactory. In these tasks, not only physical structure and behavioural properties are

³But even for these tasks, the domain knowledge necessary for the task is often not distinguished from the task knowledge which specifies how to reason with the domain knowledge.

important, but also human goals, actions, resources, requirements, preferences, etc.

A second issue is the mapping from communicative goals to concepts and relations in the domain. As we have seen, qualitative reasoning means specifying domain knowledge as an ontology in which components, processes and structural and causal relations between them are the key elements. If we want to *describe* a process or concept, this is no problem, but if we want to *identify* a concept, or *contrast* it to another, ways of fulfilling these communicative goals are less clear. This seems to require additional specification, or mechanisms for reasoning about, similarities and differences.

In a project which has started recently, the authors plan to develop a framework for generating explanations from articulate domain models which addresses these two points. As a short-term goal, a prototype will be developed which is capable of basic explanation generation on the basis of behavioural simulations in the context of an ecological domain, the Brazilian Cerrado (e.g. see Salles, Bredeweg, & Winkels 1997). In the long term, additional mechanisms will be added, like presenting counter-examples, analogies etc. The ultimate goal is to achieve thorough understanding of the way domain knowledge representation can interact with reasoning mechanisms and communication strategies to generate a successful explanatory interaction.

5 Conclusions

This paper has explored the relation between explanation generation and qualitative reasoning. Most importantly, what can be explained is dependent on the contents of the model in terms of structure, level of detail, ontological primitives, and vocabulary. If multiple levels of explanation are required, then multiple models, or views, and mechanisms for selecting appropriate ones are necessary. But in addition to these aspects of domain representation, explanation should also be seen as a communication process, which often requires planning and interaction to be successful. Principles underlying effective explanations were quoted from the literature in natural language generation; the main essence of these principles is that explanation should be seen in the context in which it is required, with special attention for the user's tasks, goals and intentions. Because of its focus on physical systems, causal chains and processes in time, textual explanations will not always be optimal; qualitative reasoning should also benefit from techniques in graphics and multimedia generation. For the future, integration of techniques developed in these various fields of research is an important goal to create explanation generation facilities which optimally support users of our systems.

References

- Baker, M. 1994. A model for negotiation in teaching-learning dialogues. *Journal of Artificial Intelligence in Education* 2(5):199-254.
- Bredeweg, B., and Winkels, R. 1998. Qualitative models in interactive learning environments: an introduction. *Interactive Learning Environments* 5:1-18.
- Brown, J. S.; Burton, R. R.; and de Kleer, J. 1982. Pedagogical, natural language and knowledge engineering techniques in SOPHIE I, II and III. In Sleeman, D., and Brown, J. S., eds., *Intelligent Tutoring Systems*, 227-282. New York: Academic Press.
- Cawsey, A. 1991. Generating interactive explanations. In *Proceedings of the ninth National Conference on Artificial Intelligence (AAAI-91)*, 86-91.
- de Kleer, J. H., and Williams, B. C. 1991. Qualitative reasoning about physical systems 2. *Artificial Intelligence* 51. (Special volume).
- de Kleer, J. H. 1986. An assumption-based TMS. *Artificial Intelligence* 28:127-162.
- de Kleer, J. H. 1990. Qualitative physics: A personal view. In Weld, D. S., and de Kleer, J. H., eds., *Readings in Qualitative Reasoning about Physical Systems*. San Mateo, California: Morgan Kaufmann. 1-8.
- de Koning, K. 1997. *Model-Based Reasoning about Learner Behaviour*. Amsterdam: IOS Press.
- Falkenhainer, B. C., and Forbus, K. D. 1991. Compositional modeling: Finding the right model for the job. *Artificial Intelligence* 51:95-143.
- Forbus, K., and de Kleer, J. 1993. *Building Problem Solvers*. Cambridge, Massachusetts: MIT Press.
- Forbus, K., and Whalley, P. 1998. Using qualitative physics to build articulate software for thermodynamics education: A preliminary report. *Interactive Learning Environments* 5:19-32.
- Forbus, K. D. 1988. Qualitative physics: Past, present and future. In Shrobe, H. E., ed., *Exploring Artificial Intelligence*. San Mateo, California: Morgan Kaufmann. 239-296.
- Forbus, K. D. 1997. Using qualitative physics to create articulate educational software. *IEEE Expert* 12(3):32-41.
- Gentner, D., and Stevens, A. L., eds. 1983. *Mental Models*. Hillsdale: Lawrence Erlbaum.
- Hollan, J. D.; Hutchins, E. L.; and Weizenbaum, L. 1984. STEAMER: an interactive inspectable simulation-based training system. *AI Magazine* 5:15-27, 2.
- Hovy, E. 1993. Discourse strategies for generating natural language text. *Artificial Intelligence* 63:341-385.
- Kerpedjiev, S.; Carenini, G.; Roth, S. F.; and Moore, J. D. 1997. Integrating planning and task-based design for multimedia presentation. In *Proceedings of the International Conference on Intelligent User Interfaces (IUI '97)*.
- Khan, T. M.; Brown, K. E.; and Leitch, R. R. 1997. Didactic and informational explanation in simulations with multiple models. In *Artificial Intelligence in Education: Knowledge and Media in Learning Systems. Proceedings of AI-ED 97 World Conference on Artificial Intelligence in Education*, 355-362.
- Knott, A., and Sanders, T. 1998. The classification of coherence relations and their linguistic markers: An exploration of two languages. *Journal of Pragmatics* 30.

- Mallory, R. S.; Porter, B. W.; and Kuipers, B. J. 1996. Comprehending complex behavior graphs through abstraction. In Iwasaki, Y., and Farquhar, A., eds., *Proceedings of the Tenth International Workshop on Qualitative Reasoning*, 137-146. Menlo Park, CA, USA: American Association for Artificial Intelligence Press. AAAI Technical Report WS-96-01.
- Mann, W. C., and Thompson, S. A. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *TEXT* 3(8):243-281.
- Moore, J. D. 1995. *Participating in Explanatory Dialogues. Interpreting and Responding to Questions in Context*. Cambridge, Massachusetts / London, England: MIT Press.
- Pilkington, R. M. 1992. Question answering for intelligent on-line help: The process of intelligent responding. *Cognitive Science* 16(4):455-491.
- Pilkington, R. M. 1997. Analyzing educational discourse: The DISCOUNT scheme. version 2, march 1997. Technical report, Computer-Based Learning Unit, University of Leeds.
- Rickel, J., and Porter, B. W. 1997. Automated modeling of complex systems to answer prediction questions. *Artificial Intelligence* 93:201-260.
- Salles, P.; Bredeweg, B.; and Winkels, R. 1997. Deriving explanations from qualitative models. In du Boulay, B., and Mizoguchi, R., eds., *Artificial Intelligence in Education: Knowledge and Media in Learning Systems*. Japan, Osaka: IOS-Press/Ohmsha. 474-481.
- Trave-Massuyes, L., and Milne, R. 1998. Gaps between research and industry related to model based and qualitative reasoning. In MONET., ed., *Proceedings of the European workshop on Model based systems and qualitative reasoning*, 54-57. Brighton, UK: ECAI 98 Workshop Programme.
- Wahlster, W.; Andre, E.; Profitlich, H.-J.; and Rist, T. 1993. Plan-based integration of natural language and graphics generation. *Artificial Intelligence* 63:387-327.
- Weld, D. S., and de Kleer, J. H. 1990. *Readings in Qualitative Reasoning about Physical Systems*. Morgan Kaufmann.
- Wenger, E. 1987. *Artificial intelligence and tutoring systems. Computational and cognitive approaches to the communication of knowledge*. Los Altos, California: Morgan Kaufmann.
- Winkels, R. G. F. 1992. *Explorations in Intelligent Tutoring and Help*. Amsterdam, The Netherlands: IOS.