Chapter 8

Conclusions and Outlook

In this thesis we have presented an integrated approach to qualitative prediction of behaviour, both as a knowledge level description and in terms of a detailed account of an implemented system that performs state of the art qualitative reasoning. Protocol analysis of human problem solving behaviour supports the cognitive plausibility hypothesis of the conceptual model and thereby its utility for knowledge acquisition. In addition, the approach has been augmented with an initial step towards reflective competence assessment and improvement. In the following sections the contributions of our research to these issues are further discussed. The last section points out a number of issues that are worthwhile for further investigation.

8.1 A Unified Framework for Prediction of Behaviour

The conceptual model presented in this thesis extends previous approaches to qualitative prediction of behaviour by distinguishing between domain, inference, task and strategic knowledge. This unified approach presents a frame of reference for comparing the original approaches on how they use these different types of knowledge.

8.1.1 Domain and Inference Knowledge

At the domain layer the original approaches provide ontological primitives for modelling the domain specific knowledge. The knowledge at the inference layer abstracts from these modelling primitives by describing the *canonical* inferences used in the reasoning process, and pointing out the *role* the domain knowledge plays in this reasoning process. The three original approaches to qualitative reasoning can be viewed as using parts of the inference layer described in this thesis for qualitative prediction of behaviour.

8.1.1.1 Roles Played by the Domain Knowledge

In many ways the approach presented in this thesis advances each of the original approaches as a result of the integrated view.

- The notion of system model description, which is used in different ways in all three approaches, is extended to include a full description of the elements in the physical system, the partial behaviour models, the parameters, the parameter values and the parameter relations. This allows the use of powerful modelling techniques and heuristics for interpreting a physical situation, and thereby greatly reduces the number of states that need to be considered.

- The notions of view, qualitative state and process are unified and extended to a generic domain ontology for partial behaviour models that discriminates between static, process and agent models. Static models represent general properties of system elements. They can be further divided into single description, composition, and decomposition models, referring to modelling properties of a single system element, a collection of system elements or to how a system element can be decomposed into its sub-structure. Processes describe changes that are based on inequalities between interacting quantities of different system elements. Agent models are used for modelling changes that are caused by agents. As shown in a number of examples, each of these modelling primitives is an essential part of a qualitative model.

  The *no-function-in-structure* principle must be redefined as follows: the behaviour description of a partial model may refer only to properties of system elements which are defined in the conditions for applying that partial behaviour model. Recall that the *no-function-in-structure* principle was defined in order to develop a library of component models that can be used independently of the specific configuration in which the components appear. Obeying this principle, as redefined above, guarantees that each partial behaviour model represents the behaviour of a system element, or of a configuration of system elements, independent of the behaviour of *other* system elements in the same context, i.e not mentioned in the conditions of the partial behaviour model.

- A third extension beyond earlier approaches concerns parameter values and parameter relations. Parameter relations can be used for a number of purposes (table 4.4). The most important ones being: relating quantity spaces, relating parameters and relating derivatives. The set of parameter relations provided by the integrated framework enables a broader functionality for specifying dependencies between parameters. In particular, we can use both directed (causal) and undirected (non-causal) dependencies between derivatives and between parameter values in a single behaviour model. In addition, the notions of directed and undirected quantity space correspondences and directed value correspondences are new. The introduction of different types of directed relations allows explicit representation of different kinds of causal dependencies between variables.

8.1.1.2 Knowledge Sources

- The specification inference uses a depth-first search algorithm for generating a full state of behaviour (system model description), which provides feedback on the applicability of partial behaviour models. In addition, the inference step determines explicitly which (knowledge level) assumptions are necessary and which sets of assumable partial behaviour models mutually exclude each other.

- The use of inequality reasoning has several advantages. It is more general, but still generates a manageable search space. The qualitative calculus, as originally proposed in [57; 93], exhibits spurious behaviour, especially with regard to conservation of quantities such as energy, flow, and force. We severely reduced the generation of spurious behaviours by applying reasoning about inequalities and arithmetic summations. Our method combines the axioms for reasoning about transitivity and arithmetic summations and thereby avoids the problems associated with the approaches of [120; 73]. Furthermore, it enables the specification of quantity spaces containing any number of values.

- The transformation inference can make use of domain independent transformation rules similar to those used in earlier approaches, but may also use rules that refer to domain specific knowledge. Of particular interest is the conceptual distinction made between *termination*, *precedence* and *continuity* rules. Each of these rules refers to a different type of knowledge used in the transformation inference. In qualitative reasoning the transformation step, in particular, tends to cause unmanageable branching of possible states of behaviour. In our approach this ambiguity is reduced by representing precedence knowledge for merging related transitions or filtering out undesired transitions. The distinction between finding terminations and then explicitly ordering them was not present in earlier approaches.

8.1.2 Task and Strategic Knowledge

The task layer is used for representing typical chains of inferences that experts make in solving a particular, well-known task. In our framework, the task knowledge is organised such that all the states of behaviour are found that apply to, or follow from, a certain input system. When the input system specifies only a configuration of system elements, this leads to generating all possible states of behaviour (total envisionment). When additional parameter values and parameter relations are added to this input system, then the behaviour prediction will be more directed, resulting in a more specific trace of behaviour (attainable envisionment).

Strategic knowledge, in the sense of the four-layer model, is not present in the original approaches to qualitative reasoning. The approaches always execute the same task structure, are not able to monitor their own inference process, and as such are not able to modify or change their own reasoning process. In the integrated framework two solutions to this problem have been realised. Firstly, the task layer is accessible by the user. In this way the user can decide to give control to the prediction engine (resulting in full autonomous execution of the task), or the user can decide to control the prediction by him- or herself and manipulate the prediction engine such that a required trace of behaviour is derived. Secondly, in our research on reflection we have started to explore the knowledge needed for operationalising the strategic layer (see section 8.4).

8.2 Operationalising an Interpretation Model for *KADS*

The framework for qualitative prediction of behaviour presented in this thesis was partly developed simultaneously with the *KADS* methodology for building knowledge based systems. The research on qualitative reasoning served as one of the domains used for further enhancement and consolidation of the theoretical insights developed within the methodology (cf. [14]). In particular, the research on qualitative reasoning concentrated on operationalising a specific problem solving task according to the principles and guidelines provided by the *KADS* methodology. However, the approach taken to this study was contrary to a regular knowledge acquisition process. Instead of analysing the problem solving behaviour of a human expert, we focused on existing artificial intelligence programs that performed the same problem solving tasks. The *KADS* methodology, and more specifically the four-layer model, was applied as a method for pointing out similarities and differences in the problem solving capabilities offered by the different approaches to qualitative reasoning. In addition, the distinction between the knowledge level model of expertise and its computation specific counterpart (design and implementation), provided strong guideliness for discriminating between computational and conceptual issues.

Being part of a model driven methodology for knowledge based systems development, the framework for qualitative prediction of behaviour serves as a method for interpreting and analysing expertise for prediction tasks in new domains. The usability of the framework for knowledge acquisition is supported by a protocol analysis of humans solving prediction problems, which showed that the conceptual model represents expertise at a level of abstraction that closely matches the human conceptualisation of the prediction problems.

The design and implementation of the model of expertise resulted in a domain independent reasoning shell, called *GARP*, which allows a knowledge engineer to quickly prototype prediction models. After providing *GARP* with the domain specific knowledge an engineer can run the program and analyse the predicted behaviour. The solutions generated by *GARP* provide guidelines for focusing the knowledge acquisition process on gathering the additional knowledge needed for optimising and finalising the prediction model. Moreover, the implementation of the task and strategic layer allows the user to interact directly with the prediction process. This makes it possible to tune the behaviour predicted by *GARP* in such a way that it traverses those paths that are of most interest for the knowledge acquisition process. The choices that have to be made by a knowledge engineer provide a focus for where additional strategic knowledge is required in the prediction model.

8.3 Cognitive Plausibility

Cognitive plausibility is relevant for deciding which parts of the problem solving potential should be described in the design model and which by the analysis model.

The conceptual framework underlying the design and implementation of *GARP* appears to be useful for describing and interpreting the reasoning processes observed by students who predicted the behaviour of complex balance problems. Both the different viewpoints subjects have on the domain knowledge as well as their reasoning process can be modelled by the framework. The canonical inferences and the meta-classes defined in the model provide a strong means for interpreting the steps of the reasoning process in the protocols. The notion of strategic reasoning explains disruptions, and changes in the order in which the new states are determined.

It is therefore fair to conclude that the conceptual model presented for qualitative prediction of behaviour does describe this problem solving task at the right level of abstraction, i.e. it constitutes a psychologically plausible knowledge level model of this problem solving expertise.

8.4 Reflective Control and Improvement

Based on a framework for knowledge level reflection [112], we gave a classification of reflective behaviour and investigated a knowledge level theory of reflection. In particular, we analysed how the notion of knowledge conflicts can be used as a means for reasoning about the competence assessment and improvement of knowledge based systems. It turned out that impasses in the problem solving process of $GARP$ can be identified and described with one of the three basic conflicts: *inconsistent*, *missing*, and *irrelevant* knowledge. In addition, we discussed how remedies can be used to aid competence improvement. In particular, we presented algorithms for improving the problem solving behaviour of $GARP$ by removing irrelevant parameters and irrelevant states of behaviour from the behaviour description.

8.5 Outlook and Further Research

Despite the advances resulting from the research presented in this thesis there are still many topics that are insufficiently addressed by the study. In the following sections we discuss a number of directions for further research.

8.5.1 Support Knowledge for the Modelling Process

Building models of real-world systems is an important problem. Even the smallest prediction model soon takes a few hours before it 'runs'. In particular, modelling the domain knowledge is a major bottleneck. Research on qualitative reasoning would greatly benefit from a theory that supports the domain modelling process. The conceptual model described in this thesis provides a starting point for the knowledge elements that such a theory should reason about. The support knowledge should specify the characteristics of the available modelling primitives and compare these to the specific features of the domains that must be modelled in the framework. Based on this comparison the support knowledge should determine to what extent specific modelling primitives are appropriate for representing certain parts of the domain knowledge. The approach to competence assessment (chapter 7) provides important handles for developing such a theory. However, more research is needed for operationalising these ideas.

8.5.2 Cognitive Modelling

The protocol analysis, carried out in the course of the research presented here, provides material for studying further the cognitive validity of the specific model of the balance problems represented in $GARP$, as well as for further research on how humans perform common sense reasoning about the everyday physical world. From the protocol analysis it became clear that further research should focus on the learning aspects concerning the knowledge structuring principles that people use for developing their domain knowledge. Interesting work on cognitive modelling can be found in [77; 74].

8.5.3 Learning and Domain Knowledge Structuring

A research issue related to both supporting the knowledge acquisition process and cognitive modelling is that of machine learning. Machine learning techniques can in this respect be used for two purposes:

- to automate the knowledge acquisition process, and
- to improve domain models over a number of behaviour prediction sessions (see chapter 6).

Research in this direction can, for example, be found in [66; 107; 108]. In particular, questions concerning the level of abstraction at which domain knowledge must be modelled, and the set of parameters that is needed for analysing the behaviour of a certain system, are essential problems that must be dealt with in further research.

8.5.4 Reflective Control and Strategic Reasoning

Although some typical examples have been tried out in experiments (cf. [5]), further research is needed for operationalising the ideas concerning reflective improvement of problem solving behaviour presented in chapter 7. The crucial topic to be addressed here concerns the realisation of a strategic layer 'on top of' an artifact such as *GARP*. There is some work going on as an extension of the constraint centred approach (cf. [93; 95]) aiming at filtering, and thereby reducing, the number of generated states of behaviour. However, the framework underlying *GARP* realises a broader functionality than provided by *QSIM* and consequently requires a different control mechanism. Moreover, it is essential that the reflective control is active during the behaviour prediction, instead of filtering the output after the problem solver is finished, otherwise, an explosive growth of predicted states, for example, cannot be prevented. The approach presented in this thesis for competence improvement provides initial ideas on how this reflective control can be realised.

8.5.5 Prediction Based Diagnosis

A possible extension of our research on qualitative prediction of behaviour is to use the domain ontology for partial behaviour models in order to represent different behaviour models for prediction based diagnosis.

A crucial aspect in prediction based diagnosis (cf. [45; 59; 79]) is the use of partial models that represent behaviours of entities in the real-world. These partial behaviour models are used to determine the overall behaviour of some system that is the object of the diagnostic reasoning process. In the original *GDE* approach [59] each behaviour model represents the behaviour of a single component. The aggregation of these models into an overall behaviour description is realised by directly combining the specific behaviour models that are available for each component.

Extensions of the initial *GDE* approach, however, require an enhanced definition of behaviour models. The behaviour models must be able to represent fault models, multiple correct models and structural decompositions. In addition, a behaviour model does not always apply to a single component, but may refer to a functional abstraction or even to processes [70].[1]

In order to address the advanced requirements for the behaviour models the integrated framework for qualitative prediction, described in this thesis, can be used (cf. [19]). In particular, the ontology for partial behaviour models can be used to control the search for applicable behaviour models in the extended *GDE* paradigm. Fault models [129], for example, can be represented as an alternative class of behaviours next to correct behaviours. The modelling primitives needed for representing these fault models do not differ from those needed for representing correct behaviour, i.e. both static, process and agent models can be used.

The notion of simplification, abstraction, and approximation (cf. [128]) can be represented by using the super-type relation, modelling the more specific models lower in the hierarchy with respect to the more abstract and simplified models. Instead of using the complete hierarchy the diagnostic problem solver should take care only to take into account as much level of detail as is required for establishing a sufficient diagnosis. In addition, we can use more than one model at a single level of detail for representing the different views that may apply to a certain device or a subpart of the device. Multiple correct behaviours of a single component can, for example, be represented as different single description models that all apply to the same system element, but which have different conditions with respect to the parameters, parameter values, parameter relations and other partial behaviour models. The diagnostic problem solver should decide which of these condition sets is valid and therefore which of the available behaviour models should be used.

Structural decomposition, as used in *SiDia* [78] and in *MuDia* [10], can be represented by the decomposition models. The notion of an *unknown behaviour mode* as used in [60] should not be represented by a partial behaviour model (although this could be a way of representing it), but is better realised by the specific way in which the diagnostic problem solver applies the use of partial models.

In conclusion, the ontology for partial behaviour models developed for qualitative prediction of behaviour can be used for the integration of the different behaviour models used in the extended *GDE* paradigm. The integrated set of modeling primitives provides handles for controlling the search for applicable behaviour models.

8.5.6 Intelligent Tutoring Systems

A further extension of the presented research is the use of *GARP* as a shell that represents and simulates the (physical) world at the conceptual level (instead of at the mathematical level) and which is therefore better suited for use within an intelligent tutoring environment [22], both for:

Domain modelling Representing the domain knowledge that has to be acquired.

Student modelling Representing alternative notions held by the students.

[1] To our knowledge the latter has not yet been studied in practice, but only briefly discussed in [128]. However, such an extension would be essential for the *GDE* paradigm to be applicable in environments where there are no active components, but only physical objects that interact on the basis of some inequality.

A growing number of computer based learning environments is simulation based (see for instance [49] for an overview). Usually these environments are based on the *discovery learning* educational philosophy (cf. [29]) and are aptly called discovery worlds. Research shows, however, that such simulations are effective only when the actions of the students are monitored by a teacher (human or computer based) and guidance is provided (e.g. [98; 63]). One solution is to embed the simulation in an intelligent tutoring system. However, most of the simulations are based on complex mathematical procedures that calculate how the specific aspects within the simulation are to be manipulated (cf. *SOPHIE* I/II [28], *STEAMER*, [84], *RBT* [147]). These mathematical procedures, although efficient in simulating the physical world, provide no conceptual access to the objects and their behaviour in the simulation. This makes it hard, if not impossible, for the intelligent tutoring system to use the simulator for explanations or student modelling, because it has no means for relating the mathematical calculations within the simulation to the conceptual framework (that represents the knowledge the student possesses, or has to acquire). It is impossible to derive causal explanations of the behaviour of the particular device or system from the mathematical model, so these would have to be added *by hand*. Therefore we are looking for different models to use for simulations, models that explicitly represent the objects and their behaviours that play a role in the system. These models should allow for causal reasoning, and preferably comply with the way humans reason about the system.

The only well known intelligent tutoring system that uses a qualitative model for its simulation is *SOPHIE* III [28], a reactive learning environment for teaching troubleshooting in electronic circuits, of which the coaching module was never realised. (Note that although *SOPHIE* III is based on a qualitative approach, it does not have true causal models of the circuits, but uses circuit specific rules and links). More recently White et al. [138] describe their work on *QUEST*, an instructional system for the same task, based on a progression of mental models, from a qualitative to a quantitative one.

In order to make the use of qualitative reasoning techniques for teaching purposes effective, the following requirements should be met:

- The simulation should be based on a single conceptual framework, which allows the representation, and the use, of different types of models of the physical world in a single intelligent tutoring system.
- The conceptual framework must be represented explicitly in the implemented artifact, in order to facilitate access to the entities from the simulation in terms of the conceptual framework.
- The framework must be cognitively plausible in order to ensure that the access to the entities represented by the simulation model closely matches the human interpretation of that simulation.

GARP is a qualitative reasoning shell that meets these requirements. It can be used for both domain simulation and student modelling. The conceptual framework on which *GARP* is based is explicitly represented in the implemented artifact and maps well onto the conceptions that humans seem to use. In addition, *GARP* can reason with different types of qualitative models (component, process, or constraint centered) and with different conceptualisations of a particular system (including wrong ones). This makes the domain model suitable for providing explanations for student responses (i.e. be used for diagnosis) or for predicting them. This extends to selecting critical problems for diagnostic or teaching purposes, in particular, for stimulating learning or awareness of misconceptions on the part of the student (cf. [124]).

Bibliography

[1] J. Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23:123–154, 1984.

[2] S. Barthèlemy, G. Edin, E. Toutain, and S. Becker. Requirements Analysis in KBS Development. ESPRIT Project P1098 Deliverable D3 (task A2), Cap Sogeti Innovation, 1987.

[3] S. Barthèlemy, P. Frot, and N. Simonin. Analysis document experiment F4. ESPRIT Project P1098, Deliverable E4.1, Cap Sogeti Innovation, 1988.

[4] S. Barthélemy and N. Simonin. Experiment F4: design and implementation. ESPRIT Project P1098, Deliverable E4.2, Cap Sogeti Innovation, 1988.

[5] B. Bartsch-Spörl, B. Bredeweg, C. Coulon, U. Drouven, F. van Harmelen, W. Karbach, M. Reinders, E. Vinkhuyzen, and A. Voß. Studies and experiments with reflective problem solvers. ESPRIT Basic Research Action P3178 REFLECT, Report IR.3.1,2 RFL/BSR-UvA/II.2/1, REFLECT Consortium, August 1991.

[6] B. Bartsch-Spörl, M. Reinders, H. Akkermans, B. Bredeweg, T. Christaller, U. Drouven, F. van Harmelen, W. Karbach, G. Schreiber, A. Voß, and B. Wielinga. A tentative framework for knowledge-level reflection. ESPRIT Basic Research Action P3178 REFLECT, Deliverable IR.2 RFL/BSR-ECN/I.3/1, BSR Consulting and Netherlands Energy Research Foundation ECN, August 1990.

[7] J. Batali. Computational introspection. Technical Report 701, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, Cambridge, Massachusetts, 1983.

[8] R. Berndsen and H. Daniels. Qualitative economics in prolog. Technical Report IKT Research Report No. 8, University of Tilburg, November 1989.

[9] D. G. Bobrow, editor. *Qualitative Reasoning about Physical Systems.* Elsevier Science, Amsterdam, The Netherlands, 1984.

[10] C. Böttcher. Die verwendung multipler modelle in der modellbasierten diagnose. Technical report, University of Karlsruhe, January 1991. (Diplomarbeit, in German).

[11] R. J. Brachman. On the epistemological status of semantic networks. In N. V. Findler, editor, *Associative Networks*, New York, 1979. Academic Press.

[12] R. J. Brachman and J. G. Schmolze. An overview of the KL-ONE knowledge representation system. *Cognitive Science*, 9:171–216, 1985.

[13] I. Bratko. *Prolog Programming for Artificial Intelligence.* Addison-Wesley, Reading, Massachusetts, 1986.

[14] B. Bredeweg. GARP - K4 mini study. Esprit project p1098 - task k4, SWI, University Of Amsterdam, January 1989.

[15] B. Bredeweg. Introducing meta-levels to qualitative reasoning. *Applied Artificial Intelligence*, 3–2:85–100, 1989.

[16] B. Bredeweg and Z. Karssen. Domain representation of heart diseases and cover-and-differentiate. ESPRIT Basic Research Action P3178 REFLECT, Technical Report RFL/UvA/I.2/2, University of Amsterdam, May 1990. Describes intermediate software deliverable IS.1.

[17] B. Bredeweg, M. Reinders, and B.J. Wielinga. Garp: A unified approach to qualitative reasoning. Technical Report VF-memo 117, University of Amsterdam, January 1990.

[18] B. Bredeweg and C. Schut. Cognitive plausibility of a conceptual framework for modeling problem solving expertise. In *Proceedings of the 13th Conference of Cognitive Science Society*, pages 473–479, Hillsdale, New Jersey, August 1991. Lawrence Erlbaum.

[19] B. Bredeweg and H. Voß. A KADS conceptual model of prediction based diagnosis. Technical report, National german Research Centre for Computer Science (GMD), February 1992.

[20] B. Bredeweg and B. J. Wielinga. Integrating qualitative reasoning approaches. In *Proceedings of ECAI–88*, pages 195–201, Munich, Germany, 1988.

[21] B. Bredeweg and B. J. Wielinga. Reasoning about physical systems: a unifying perspective. In *Proceedings of IMACS–88, Paris*, pages 287–290, 1988.

[22] B. Bredeweg and R. G. F. Winkels. Teaching according to GARP. In *Proceedings of the first International Conference on the Learning Sciences*, pages 52–58, Charlottesville, Virginia, August 1991. Lawrence Birnbaum.

[23] J. A. Breuker and B. J. Wielinga. KADS: Structured knowledge acquisition for expert systems. In *Proc. of the Fifth International Workshop on Expert Systems and their Applications*, Avignon, France, 1985.

[24] J. A. Breuker and B. J. Wielinga. Use of models in the interpretation of verbal data. In A. L. Kidd, editor, *Knowledge Acquisition for Expert Systems, a practical handbook*, New York, 1987. Plenum Press.

[25] J. A. Breuker and B. J. Wielinga. Model Driven Knowledge Acquisition. In P. Guida and G. Tasso, editors, *Topics in the Design of Expert Systems*, pages 265–296, Amsterdam, The Netherlands, 1989. North-Holland.

[26] J. A. Breuker, B. J. Wielinga, M. van Someren, R. de Hoog, A. Th. Schreiber, P. de Greef, B. Bredeweg, J. Wielemaker, J. P. Billault, M. Davoodi, and S. A. Hayward. Model Driven Knowledge Acquisition: Interpretation Models. ESPRIT Project P1098 Deliverable D1 (task A1), University of Amsterdam and STL Ltd, 1987.

[27] D.C. Brown and B. Chandrasekaran. *Design Problem Solving.* Pitman Publishers, London, 1989.

[28] J. S. Brown, R. R. Burton, and J. de Kleer. Pedagogical, natural language and knowledge engineering techniques in SOPHIE I, II and III. In D. Sleeman and J. S. Brown, editors, *Intelligent Tutoring Systems*, pages 227–282, New York, 1982. Academic Press.

[29] J.S. Bruner. The act of discovery. *Harvard Educational Review*, 31:21–31, 1961.

[30] B. G. Buchanan and E. H. Shortliffe. *Rulebased Expert Systems: The Mycin Experiments of the Stanford Heuristic Programming Project.* Addison Wesley, Reading, Massachusetts, 1984.

[31] A. Bundy, editor. *A Catalogue of AI techniques.* Springer-Verlag, New York, 1985.

[32] A. Bundy, L. Byrd, G. Luger, and C. Mellish M. Palmer. Solving mechanics problems using meta-level inference. In *Proceedings of the 6th International Joint Conference on Artificial Inteligence, Tokyo*, 1979.

[33] T. Bylander. Some causal models are deeper than others. In *Proceedings of the 9th Conference of Cognitive Science Society*, pages 916–921, Seattle, 1987.

[34] T. Bylander and B. Chandrasekaran. Generic tasks in knowledge-based reasoning: The right level of abstraction for knowledge acquisition. In B. Gaines and J. Boose, editors, *Knowledge Acquisition for Knowledge Based Systems*, volume 1, pages 65–77. Academic Press, London, 1988.

[35] B. Chandrasekaran. Generic tasks in knowledge based reasoning: High level building blocks for expert system design. *IEEE Expert*, 1(3):23–30, 1986.

[36] B. Chandrasekaran. Towards a functional architecture for intelligence based on generic information processing tasks. In *Proceedings of the 10th IJCAI*, pages 1183–1192, Milan, Italy, 1987.

[37] B. Chandrasekaran. Generic tasks as building blocks for knowledge-based systems: The diagnosis and routine design examples. *The Knowledge Engineering Review*, 3(3):183–210, 1988.

[38] B. Chandrasekaran and S. Mittal. Deep versus compiled knowledge approaches to diagnostic problem solving. *International Journal of Man-Machine Studies*, 19:425–436, 1983.

[39] E. Charniak and D. McDermott. *Introduction to Artificial Intelligence.* Addison Wesley, Mass., 1985.

[40] W. J. Clancey. The epistemology of a rule based system -a framework for explanation. *Artificial Intelligence*, 20:215–251, 1983.

[41] W. J. Clancey. Heuristic classification. *Artificial Intelligence*, 27:289–350, 1985.

[42] A. G. Cohn. Qualitative reasoning. In R. T. Nossum, editor, *Advanced Topics in Artificial Intelligence (ACAI–87)*, pages 96–124. Spinger-Verlag, Berlin, Germany, 1988.

[43] E.W. Coiera. Monitoring diseases with emperical and model-generated histories. *Artificial Intelligence in Medicine*, 2:135–147, 1990.

[44] E. Davis. *Representation of Common Sense Knowledge.* Morgan Kaufmann, San Mateo, California, 1990.

[45] R. Davis. Diagnostic reasoning based on structure and behavior. *Artificial Intelligence*, 24:347–410, 1984.

[46] R. Davis and B. Buchanan. Meta-level knowledge. In B. G. Buchanan and E. H. Shortliffe, editors, *Rule-Based Expert Systems*. Addison-Wesley, 1984.

[47] R. Davis and W. C. Hamscher. Model-based reasoning: Troubleshooting. In H. E. Shrobe, editor, *Exploring Artificial Intelligence*, pages 297–346. Morgan Kaufmann, San Mateo, California, 1988.

[48] P. de Greef, J. A. Breuker, and T. de Jong. Modality: An analysis of functions, user control and communication in knowledge-based systems. ESPRIT Project P1098, Deliverable D6 (task A4), University of Amsterdam, 1988. Available from: University of Amsterdam, Social Science Informatics, Roetersstraat 15, NL-1018 WB, Amsterdam, The Netherlands.

[49] T. de Jong (editor). Computer simulations in an instructional context. *Education and Computing (Special issue)*, 6, 1991.

[50] J. de Kleer. Multiple representations of knowledge in a mechanics problem solver. In *Proceedings of IJCAI-77, Cambridge, MA*, pages 299–304, 1977.

[51] J. de Kleer. Qualitative pysics: A personal view. In D.S. Weld and J. de Kleer, editors, *Readings in Qualitative Reasoning about Physical Systems*, pages 1–8. Morgan Kaufmann, San Mateo, California, 1990.

[52] J. de Kleer and J.S Brown. Assumptions and ambiguities in mechanistic mental models. In D. Gentner and A.L. Stevens, editors, *Mental Models*, pages 155–190. Lawrence Erlbaum, Hillsdale, 1983.

[53] J de Kleer and B.C. Williams. Qualitative reasoning about physical systems 2. *Artificial Intelligence*, 51, October 1991. (Special volume).

[54] J. H. de Kleer. An assumption-based TMS. *Artificial Intelligence*, 28:127–162, 1986.

[55] J. H. de Kleer. Extending the ATMS. *Artificial Intelligence*, 28:163–196, 1986.

[56] J. H. de Kleer. Problem solving with the ATMS. *Artificial Intelligence*, 28:197–224, 1986.

[57] J. H. de Kleer and J. S. Brown. A qualitative physics based on confluences. *Artificial Intelligence*, 24:7–83, 1984.

[58] J. H. de Kleer and J. S. Brown. Theories of causal ordering. *Artificial Intelligence*, 29:33–61, 1986.

[59] J. H. de Kleer and B. C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32:97–130, 1987.

[60] J. H. de Kleer and B. C. Williams. Diagnosis with behavioral modes. In *Proc 11th. IJCAI*, pages 1324–1330, Detroit, 1989.

[61] D. de Klerk. Qualitatief redeneren over een complex apparaat. Technical Report Doctoraal skriptie, University of Amsterdam, Amsterdam, The Netherlands, 1991. (in Dutch).

[62] A.A. Disessa. Phenomenology and the evolution of intuition. In D. Gentner and A.L. Stevens, editors, *Mental Models*, pages 15–. Lawrence Erlbaum, Hillsdale, 1983.

[63] M. Elsom-Cook, editor. *Guided Discovery Tutoring*. Paul Chapman Publishing Ltd., London, 1990.

[64] L. Eshelman. MOLE: A knowledge-acquisition tool for cover-and-differentiate systems. In S. Marcus, editor, *Automating Knowledge Acquisition for Expert Systems*, pages 37–80. Kluwer, Boston, 1988.

[65] L. Eshelman, D. Ehret, J. McDermott, and M. Tan. MOLE: a tenacious knowledge acquisition tool. In J. H. Boose and B. R. Gaines, editors, *Knowledge Based Systems, Volume 2: Knowledge Acquisition Tools for Expert Systems*, pages 95–108, London, 1988. Academic Press.

[66] B. Falkenhainer. Scientific theory formation through analogical inference. In *Proceedings of the 4th International Workshop on Machine Learning*, pages 218–229, Irvine, 1987. University of California.

[67] B. Falkenhainer and K.D. Forbus. Setting up large-scale models. In *Proceedings of the AAAI*, pages 301–306, San Mateo, California, August 1988. Morgan Kaufmann.

[68] B. Falkenhainer and K.D. Forbus. Compositional modeling of physical systems. In *Proceedings of the 4th international Workshop on Qualitative Physics*, pages 1–14, Lugano, Switserland, July 1990.

[69] A.M. Farley. Qualitative modeling of economic systems. In *Proceedings of the IFAC Economics and Artificial Intelligence*, Aix-en-Provence, Fance, 1986.

[70] K. D. Forbus. Qualitative process theory. *Artificial Intelligence*, 24:85–168, 1984.

[71] K. D. Forbus. Qualitative process theory. Technical report, MIT Artificial Intelligence Laboratory, Massachusetts, 1984.

[72] K. D. Forbus. Introducing actions to qualitative reasoning. In *Proceedings of IJCAI, Detroit*, pages 1273–1278. Morgan Kaufmann, 1989.

[73] K. D. Forbus. The qualitative process engine. In D. S. Weld and J. H. de Kleer, editors, *Readings in qualitative reasoning about physical systems.* Morgan Kaufmann, San Mateo, California, 1990.

[74] K. D. Forbus and D. Gentner. Learning physical domains : Toward A theoretical framework. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, editors, *Machine Learning : An Artificial Intelligence Approach*, volume 2, pages 311–348. Morgan Kaufmann, Los Altos, California, 1986.

[75] K.D. Forbus. Modeling motion with qualitative process theory. In *Proceedings of the AAAI*, pages 205–208, San Mateo, California, 1982. Morgan Kaufmann.

[76] K.D. Forbus. Qualitative physics: Past, present and future. In H.E. Shrobe, editor, *Exploring Artificial Intelligence*, pages 239–296. Morgan Kaufmann, San Mateo, California, 1988.

[77] D. Gentner and A.L. Stevens, editors. *Mental Models.* Lawrence Erlbaum, Hillsdale, 1983.

[78] T. Guckenbiehl and G. Schäfer-Richter. Sidia: Extending prediction based diagnosis to dynamic models. In *Expert Systems in Engineering, Principles and applications*, pages 53–68. Springer Verlag, Berlin, New York, 1990. (Proceedings of workshop on Diagnosis, Vienna, Austria).

[79] W.C. Hamscher. Model-based troubleshouting of digital systems. Technical Report AI-TR 1074, MIT Press, Cambridge, Massachussetts, 1988. PhD thesis.

[80] P. J. Hayes. Naive physics 1: ontology for liquids. Technical report, University van Essex, Essex, 1978.

[81] P. J. Hayes. The naive physics manifesto. In D. Michie, editor, *Expert Systems in the Microelectronic Age.* Edinburgh University Press, 1979.

[82] P.J. Hayes. The second naive physics manifesto. In J.R. Hobbs and R.C. Moore, editors, *Formal Theories of the Common Sense World*, pages 1–36. Ablex Publishing Corporation, Norwood, 1985.

[83] S. A. Hayward, B. J. Wielinga, and J. A. Breuker. Structured analysis of knowledge. *International Journal of Man-Machine Studies*, 26:487–498, 1987.

[84] J.D. Hollan, E.L. Hutchins, and L. Weizenbaum. Steamer: an interactive inspectable simulation-based training system. *AI Magazine*, 5:15–27, 2, 1984.

[85] I. Bratko and I. Mozetič and N. Lavrač. *KARDIO: A Study in Deep and Qualitative Reasoning for Expert Systems.* MIT Press, Cambridge, MA, 1989.

[86] L. Ironi, G. Lanzola, and M. Stefanelli. A qualitative model of iron metabolism. In *Proceedings of the European Conference on Artificial Intelligence and Medicine*, pages 147–166. Springer-Verlag, 1989.

[87] Y. Iwasaki and H. A. Simon. Causality in device behaviour. *Artificial Intelligence*, 29:3–32, 1986.

[88] Y. Iwasaki and H.A. Simon. Theories of causal ordering: reply to de kleer and brown. *Artificial Intelligence*, 29:63–67, 1986.

[89] W.N.H. Jansweijer. *PDP: An artificial Intelligence Approach to Problem Solving and Learning by Doing in a Semantically Rich Domain.* PhD thesis, University of Amsterdam, Amsterdam, 1988. (in Dutch).

[90] W. Karbach, M. Linster, and A. Voß. OFFICE-PLAN: Tackling the synthesis frontier. In D. Metzing, editor, *GWAI–89: 13th German Workshop on Artificial Intelligence, Informatik Fachberichte 216*, pages 379–387, Berlin, Germany, 1989. Springer-Verlag.

[91] D. Klein and T. Finin. Whats in a deep model?; A characterization of of knowledge depth in intelligent safety systems. In *Proc. of IJCAI–87*, pages 559–562, 1987.

[92] B. Kuipers. Commonsense reasoning about causality: Deriving behaviour from structure. *Artificial Intelligence*, 24:169–203, 1984.

[93] B. Kuipers. Qualitative simulation. *Artificial Intelligence*, 29:289–388, 1986.

[94] B. Kuipers. New reasoning methods for artificial intelligence in medicine. *International Journal of Man-Machine Studies*, 26, 1987.

[95] B. Kuipers and C. Chiu. Taming intractible branching in qualitative simulation. In *Proceedings of IJCAI-87, Milan, Italy*, pages 1079–1085, 1987.

[96] B. Kuipers and J. P. Kassirer. Knowledge acquisition by analysis of verbatim protocols. In A. L. Kidd, editor, *Knowledge Acquisition for Expert Systems: A Practical Handbook*, pages 43–71. Plenum Press, New York, 1987.

[97] J.H. Larkin. The role of problem representation in physics. In D. Gentner A.L. Stevens, editor, *Mental Models*, pages 27–39. Lawrence Erlbaum, Hillsdale, 1983.

[98] D.R. Lavoie and R. Good. The nature and use of prediction skills in a biological computer simulation. *Journal of Research in Science Teaching*, 25:335–360,5, 1988.

[99] P. Maes. Introspection in knowledge representation. In *Proceedings of the European Conference on Atrificial Intelligence*, Brighton, UK, 1986.

[100] P. Maes. Computational reflection. Technical report 87-2, Free University of Brussels, AI Lab, 1987.

[101] P. Maes and D. Nardi, editors. *Meta-Level Architectures and Reflection*, Amsterdam, The Netherlands, 1988. North-Holland.

[102] D. McDermott. A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6:101–155, 1982.

[103] J. McDermott. The world would be a better place if non-programmers could program. *Machine Learning*, 4(3/4):337–338, 1989.

[104] C.T.C.W. Mettes and A. Pilot. Over het leren oplossen van natuurweteschappelijke problemen. Technical Report CDO-rapport 42, Technical University of Twente, 1980. (in Dutch).

[105] C.T.C.W. Mettes and H.J. Roossink. Linking factual and procedural knowledge in solving science problems: A case study in a thermodynamics course. *Instructional Science*, 10:333–361, 1981.

[106] R. S. Michalski. A theory and methodology of inductive learning. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, editors, *Machine learning: an Artificial Intelligence Approach*. Springer-Verlag, 1983.

[107] I. Mozetiĉ. The role of abstractions in learning qualitative models. In *Proceedings of the 4th International Workshop on Machine Learning*, pages 242–255, Irvine, 1987. University of California.

[108] I. Mozetiĉ, I. Bratko, and T. Urbanĉiĉ. Varying levels of abstraction in qualitative modelling. In J.E. Hayes, D. Michie, and E. Tyugu, editors, *Machine Intelligence 12*, pages 259–280. Clarendon Press, Oxford, 1990.

[109] A. Newell. The knowledge level. *Artificial Intelligence*, 18:87–127, 1982.

[110] O. Kühn and M. Linster and G. Schmidt. CLAMPING, COKAM, KADS and OMOS: The construction and operationalisation of a KADS conceptual model. In J. Boose, B. Gaines, and D. Smeed, editors, *Proceedings of the fifth European Knowledge Acquisiton Workshop, EKAW91*, Crieff Hydro, Scotland, May 1991. University of Strathclyde.

[111] O. Raiman. Order of magnitude reasoning. In *Proceedings of the AAAI*, pages 100–104, San Mateo, California, 1986. Morgan Kaufmann.

[112] M. Reinders, E. Vinkhuyzen, A. Voß, J. M. Akkermans, J. R. Balder, B. Bartsch-Spörl, B. Bredeweg, U. Drouven, F. van Harmelen, W. Karbach, Z. Karssen, A. Th. Schreiber, and B. J. Wielinga. A conceptual modelling framework for knowledge-level reflection. *AI Communications*, 4(2-3):74–87, 1991.

[113] E. Rich. *Artificial Intelligence*. McGraw-Hill, Singapore, London, 1983.

[114] E.P. Sacks and Doyle J.D. Prelegomena to any future qualitative physics. *Computational Intelligence*, 1991. (to appear).

[115] E. Sandewall. Combining logic and differential equations for describing real-world systems. In *Proceedings Knowledge Representation*, pages 412–420, Toronto, Ontario, 1989.

[116] M. Schachter and D. Wermser. A sales assistant for chemical measurement equipment. In *Proceedings of ECAI–88*, pages 191–193, London, 1988. Pitman.

[117] G. Schreiber, B. Bartsch-Spörl, B. Bredeweg, F. van Harmelen, W. Karbach, M. Reinders, E. Vinkhuyzen, and A. Voß. Designing architectures for knowledge-level reflection. ESPRIT Basic Research Action P3178 REFLECT, Deliverable IR.4 RFL/UvA/III.1/4, REFLECT Consortium, August 1991.

[118] G. Schreiber, B. Bredeweg, M. Davoodi, and B. Wielinga. Towards a design methodology for KBS. ESPRIT Project P1098, Deliverable D8 (task B2), University of Amsterdam and STL Ltd,, 1987.

[119] G. Schreiber, B. Bredeweg, P. de Greef, P. Terpstra, B. Wielinga, E. Brumet, N. Simonin, and A. Wallyn. A KADS approach to KBS design. ESPRIT Project 1098, deliverable B6 UvA-B6-PR-010, University of Amsterdam & Cap Sogeti Innovation, December 1989.

[120] R. Simmons. Commonsense arithmetic reasoning. In *Proceedings of the AAAI–86*, pages 118–124, 1986.

[121] B. C. Smith. Reflection and semantics in a procedural language. Technical Report TR–272, MIT, Computer Science Lab., Cambridge, Massachussetts, 1982.

[122] L. Steels. The deepening of expert systems. In *Proceedings of IMACS-88, Paris*, pages 323–326, 1988.

[123] L. Steels. Components of expertise. *AI Magazine*, Summer 1990.

[124] A.L. Stevens and A. Collins. The goal structure of a socratic tutor. In *Proceedings of ACM*, pages 256–263, 1977.

[125] J. Sticklen. Problem solving architecture at the knowledge level. *Journal of Experimental and Theoretical Artificial Intelligence*, 1989.

[126] J. Sticklen and B. Chandrasekaran. Integrating classification-based compiled level reasoning with function-based deep level. In W. Horn, editor, *Causal AI Models: Steps Towards Applications*, pages 191–220. Hemisphere Publishing Corporation, New York, 1989.

[127] P. Struss. Problems of interval-based qualitative reasoning. In D. S. Weld and J. H. de Kleer, editors, *Readings in Qualitative Reasoning about Physical Systems.* Morgan Kaufmann, San Mateo, California, 1990.

[128] P. Struss. What's in SD? towards a theory of modeling for diagnosis. In *Second international workshop on principles of diagnosis*, pages 41–51, Milano, 1991. Working notes.

[129] P. Struss and O. Dressler. Physical negation – integrating fault models into the general diagnostic engine. In *Proc 11th. IJCAI*, pages 1318–1323, Detroit, 1989.

[130] P. Toal and J. Hunter. A qualitative model of cardiac haemodynamics. Technical Report AUCS/TR9001, University of Aberdeen, Aberdeen, Scotland, 1990.

[131] J.L. Top and J.M. Akkermans. Processes as components: on the primitives of a qualitative scientific physics. In *Proceedings ECAI-90, Stockholm*, pages 643–648, London, 1990. Pitman.

[132] J.L. Top and J.M. Akkermans. Computational and physical causality. In *Proceedings of the IJCAI*, August 1991.

[133] B. Ueberreiter and A. Voß, editors. *Materials KADS User Meeting, February 14/15 1991.* Siemens AG ZFE IS INF 32, Munich Perlach, Germany, 1991. In German.

[134] F. van Harmelen. Classification of meta-level architectures. In P. Jackson, H. Reichgelt, and F. van Harmelen, editors, *Logic-Based Knowledge Representation*, chapter 2, pages 13–36. MIT Press, 1989. Also in: *Meta-Programming in Logic Programming (META88)*, Abramson, H. and Rogers, M. H. (eds.), MIT Press, 1989, pp. 103–122.

[135] F. van Harmelen. *Meta-level Inference Systems.* Research Notes in AI. Pitmann, Morgan Kaufmann, London, San Mateo California, 1991.

[136] A. Voß, W. Karbach, U. Drouven, and D. Lorek. Competence assessment in configuration tasks. In L. Aiello, editor, *Proceedings of the 9th European Conference on AI, ECAI–90*, pages 676–681, London, August 1990. ECAI, Pitman.

[137] D. S. Weld and J. H. de Kleer. *Readings in Qualitative Reasoning about Physical Systems.* Morgan Kaufmann, 1990.

[138] B.Y. White and J.R. Frederiksen. Qualitative models and intelligent learning environments. In *Artificial Intelligence and Education, volume 1*, pages 281–305. Ablex Publishing, Norwood, 1987.

[139] J. Wielemaker. SWI-Prolog 1.4 Reference Manual. Technical report, SWI, University of Amsterdam, Amsterdam, The Netherlands, January 1990.

[140] B. J. Wielinga and B. Bredeweg. Knowledge and expertise in expert systems. In G. C. van der Veer and G. Mulder, editors, *Human-Computer Interaction: Psychonomics Aspects*, pages 290–297, Berlin, Germany, 1988. Springer-Verlag.

[141] B. J. Wielinga, B. Bredeweg, and J. A. Breuker. Knowledge acquisition for expert systems. In R. T. Nossum, editor, *Advanced Topics in Artificial Intelligence (ACAI–87)*, pages 96–124. Spinger-Verlag, Berlin, Germany, 1988.

[142] B. J. Wielinga and J. A. Breuker. Interpretation of verbal data for knowledge acquisition. In T. O'Shea, editor, *Advances in Artificial Intelligence*, pages 41–50, Amsterdam, The Netherlands, 1984. ECAI, Elsevier Science. Also as: Report 1.4, ESPRIT Project 12, University of Amsterdam.

[143] B. J. Wielinga and J. A. Breuker. Models of expertise. In *Proceedings ECAI–86*, pages 306–318, 1986.

[144] B. J. Wielinga, A. Th. Schreiber, and J. A. Breuker. KADS: A modelling approach to knowledge engineering. ESPRIT Project P5248 KADS-II/T1.1/PP/UvA/008/2.0, University of Amsterdam, 1991. To appear in *Knowledge Acquisition*.

[145] B.J. Wielinga and J.A. Breuker. Interpretation of verbal data for knowledge acquisition. ESPRIT Project P12 Report 1.4, University of Amsterdam, 1984.

[146] D. M. Williams, J. D. Hollan, and A. L. Stevens. Human reasoning about a simple physical system. In D. Gentner A. L. Stevens, editor, *Mental Models*, pages 40–51. Lawrence Erlbaum, Hillsdale, 1983.

[147] B. Woolf and P.A. Cunningham. Multiple knowledge sources in intelligent tutoring systems. In *IEEE Expert*, pages 41–54, 1987.

[148] E. Yourdon and L. Constantine. *Structured Design: A discipline of Computer Program and System Design*. Yourdon Press, 1978.

Acknowledgements

The research reported here was partially funded by the ESPRIT Programme of the Commission of the European Communities as project numbers 1098 (Kads), 3178 (Reflect, Basic Research) and 5248 (Kads-II).

The partners in KADS are STC Technology Ltd. (UK), SD-SCICON Ltd. (UK), KBSC of TRMC (UK), NTE Neutech (D), Cap Sogeti Innovation (F), and the University of Amsterdam (NL).

The partners in REFLECT are the University of Amsterdam (NL), the Netherlands Energy Research Foundation ECN (NL), the National German Research Centre for Computer Science GMD (D) and BSR Consulting (D).

The partners in KADS-II are Cap Gemini Innovation (F), Cap Gemini Logic (S), Netherlands Energy Research Foundation ECN (NL), ENTEL SA (ESP), IBM France (F), Lloyd's Register (UK), Swedisch Institute of Computer Science (S), Siemens AG (D), Touche Ross MC (UK), University of Amsterdam (NL) and the Free University of Brussels (B).