The Acquisition of a Shared Task Model

Frances Brazier, Jan Treur and Niek Wijngaards

Artificial Intelligence Group,
Department of Mathematics and Computer Science
Vrije Universiteit Amsterdam
De Boelelaan 1081, 1081 HV Amsterdam, The Netherlands
E-mail: `{frances, treur, niek}@cs.vu.nl`

Keywords: shared task model, strategic interaction, cooperative agents, design support.

Abstract The process of the acquisition of an agreed, shared task model as a means to structure interaction between expert users and knowledge engineers is described. The role existing (generic) task models play in this process is illustrated for two domains of application, both domains requiring diagnostic reasoning. In both domains different levels of interaction between an expert user and a diagnostic reasoning system are distinguished.

1 Introduction

Decision support systems are most often designed to provide expert users with the information they need to solve a problem. More extensive support, however, is provided by knowledge-based systems that not only are capable of performing complex computation but that also are equipped with explicit knowledge of the decision process. The acquisition of such knowledge is not as trivial as it may seem. Although experts differ in their approaches to problems, in almost all situations different alternatives are thought through and compared. Decision support systems ideally support experts in this process. Not only the opportunity to influence the approach taken by systems (for example the sequence of tasks) is of importance, but also the opportunity to influence the more local levels of strategic reasoning involved in decision making processes.

User centered task analysis is essential to the design of such systems (Barnard, 1993; Brazier & Treur, 1994). The tasks users perform in specific decision making situations must be identified, in addition to relations between tasks. The designer of a system (in general a knowledge engineer) and one or more experts must reach a common understanding of the tasks involved in a specific decision making process. The types of decisions an expert user would prefer to make him/herself and the ways in which an expert user would like to be able to influence a system's reasoning, must be identified.

In this paper the role a shared task model can play as a means to acquire a common understanding of a task in interaction with expert users, is described. During an acquisition process different types of interaction between expert users and a system designed to support such users, can be identified. Three levels of interaction are introduced below in Section 2. In Section 3 existing (generic) task models used to structure the knowledge acquisition process are introduced. In Sections 4 and 5 the process of acquisition is described for two domains for which shared task models have been devised in practice, on the basis of which system architectures have been developed.

2 Knowledge Acquisition

To structure the exchange of knowledge between a knowledge engineer and an expert user often mediating representations are used (e.g., Ford, Bradshaw, Adams-Webber & Agnew, 1993). From our perspective, one of the results of knowledge acquisition (and task analysis) is a shared task model: a model which both the knowledge engineer and one or more expert user(s) agree to be an acceptable representation of the task structure for which support is to be provided.

2.1 Explicit Interaction

Within a shared model different types of tasks are distinguished: some of which may require interaction between the user and the system, and others which may not. Different types of information may be exchanged, depending on the subtask. These different types of information are used to define different levels of interaction.

Object level interaction is the interaction required to acquire specific facts about a current situation. This type of interaction in often modelled during the development and design of knowledge-based systems.

In addition, however, experts often reason about the approaches they take, comparing strategies and results. Systems designed to support expert users should therefore support such types of explicit meta-reasoning. Expert users should be able to influence the strategies employed by a system influencing factors such as specific goals, heuristics, preferences, assumptions, etc. Interaction at this level, is called *interaction at the level of strategic preferences* or *strategic interaction*.

Although a shared task model is the result of negotiation with one or more experts, it is not necessarily “the” correct model of a task for all problems in all domains. Expert users may wish to influence, for example, the sequencing or choice of subtasks in a particular situation. Interaction at this level, the level of *task model modification*, allows for individual expert users to adapt the task model to their own needs.

To model the knowledge required at these three levels of interaction within the task model, a task based framework for the design and development of compositional systems is required.

2.2 Declarative Compositional Approach

DESIRE (Langevelde, Philipsen & Treur, 1992; Brazier, Treur, Wijngaards & Willems, 1994, 1995) is a framework for the design and development of compositional systems. The framework provides support for the specification and implementation of compositional task models. These models include knowledge of the following types (comparable types of knowledge are distinguished in task analysis approaches such as KAT/KTS (Johnson & Johnson, 1991, 1993)):

1 knowledge of the task structure: task (de)composition,
2 knowledge of sequencing of tasks and goals: control (de)composition,
3 knowledge of knowledge structures,
4 knowledge of information exchange,
5 knowledge of task delegation.

Within the DESIRE framework different levels of abstraction are distinguished for each of these five types of knowledge. Tasks are defined at different levels of abstraction, resulting in a task (de)composition. Different levels of abstraction are also found within knowledge structures such as taxonomies, to which tasks refer. Sequencing of tasks and goals, and information exchange are defined not only at the level of primitive tasks, but also between composed tasks (and between composed tasks and primitive tasks), again providing a levelled structure of abstraction. Task delegation, the last of the five types of knowledge, can also be defined at all levels within a task model. More abstract tasks may be delegated to more than one party, whereas more specific tasks are often delegated to one particular party.

Within the DESIRE framework a distinction is made between the task dimension and the knowledge dimension. Together the knowledge structures define the knowledge dimension, related to the task dimension, but separately defined.

2.3 Task Models

A shared task model, as a mediating representation, is the result of negotiation between a knowledge engineer and one or more experts. The purpose of the negotiation is to acquire a common understanding of the task. An expert has extensive (often implicit) knowledge of a domain and of his/her task and strategies. A knowledge engineer has knowledge of existing models of related tasks which may or may not be applicable, and of ways to modify and combine such models for the domain at hand. Abstract task models are often used to structure the knowledge acquisition process.

Within the DESIRE framework (Langevelde *et al.*, 1992; Brazier *et al.*, 1995), a number of such abstract task models, generic task models, exist which are used for this purpose. These models have been defined on the basis of experience and logical analysis. The concept of a generic task, introduced by Chandrasekaran (1986, 1990) and Brown and Chandrasekaran (1989), is comparable to the notion of *generic task model* in that they are both generic with respect to domains. Generic task models within the DESIRE framework, however, are generic with respect to both tasks and domain: generic task models can be refined with respect to the task by *specialisation* (e.g., further decomposition of a subtask) and refined with respect to the domain by *instantiation* (e.g., addition of domain-specific knowledge). Moreover, the way a generic task model is specified in DESIRE is more declarative (with semantics based on temporal logic) than the way generic tasks are described in Chandrasekaran (1986, 1990) and Brown and Chandrasekaran (1989). The integral approach to levels of abstraction within the DESIRE framework supports the use of generic task models during knowledge acquisition. Different levels of abstraction and composition play a role during the negotiation phase.

2.4 The Common KADS Approach

The Common KADS model set (see de Hoog, Martil, Wielinga, Taylor, Bright and van de Velde, 1994) includes: an organisation model, a task model, an agent model, an expertise model, a communication model and a design model. An *organisation model* analyses the impact of a system in and on an organisation. A *task model* describes the tasks related to the realisation of a function in an organisation independent of the agent responsible for the performance of the tasks. A task model, however, when complete, relates each task to an agent. Agents are described in an *agent model*. The capabilities of an agent are described in an *expertise model*. Strategic knowledge is defined by inference and task aspects of the problem solving knowledge included in an expertise model. Communication tasks, defined in a *communication model*, are specified in terms of user models (defined in an agent model) and transfer tasks (defined in an expertise model). Important decisions made during the design of an application together form the *decision model*.

Not all of the three levels of interaction distinguished above in section 2.1 are easily distinguished within the Common KADS framework. Object level interaction is defined by transfer tasks. How interaction at the level of strategic preferences or task model modification can be modelled is less clear. One option is to use the task layer of the expertise model, another is to use the REFLECT principle (see van Harmelen, Wielinga, Bredeweg, Schreiber, Karbach, Reinders, Voß, Akkermans, Bartsch-Spörl & Vinkhuyzen, 1992). Using the task layer to model these levels of interaction may not be appropriate, as domain specific (strategic) knowledge is involved, which then would not be specified at the domain layer and inference layer of the expertise model. This is, therefore, not a very elegant solution. The REFLECT approach models an entire expertise-model in the domain layer of another expertise model. Explicit strategic reasoning can be modelled within this approach, but entails the (recursive) combination of two expertise models for this purpose.

Reasoning about states of different reasoning processes is quite common in, for example, multi-agent situations. The Common KADS framework does not include constructs or models which can be used for this purpose. The semantics of DESIRE, however, based on temporal logic (states and transitions between states) has been designed to model interaction between components (which may be tasks) by explicit specification of transitions between states.

From the above it follows that the way in which the three levels of interaction are incorporated into one knowledge based system is not as transparent in Common KADS as in DESIRE. In DESIRE the levels of abstraction and temporal semantics facilitate the modelling of these levels of interaction.

3 Generic Task Models for Diagnostic Reasoning

In most situations in which diagnosis is required not all relevant facts are known in advance. In practice, in fact, diagnosis is not often based on complete information. The acquisition of additional (test) information is an essential part of most diagnostic processes. In general, diagnosis includes a number of subprocesses such as: the determination of hypotheses, the choice of applicable tests, the performance of tests, and the interpretation of test results. Strategic considerations such as the suitability of a test, the likeliness of a hypothesis being true, and the cost and effect of a test, play an important role in these processes. A number of existing (generic) task models for diagnostic reasoning in which strategic knowledge is explicitly modelled are described in this section. These models are used in interaction with experts to structure acquisition of shared task models of diagnostic reasoning for specific domains of application. The process of knowledge acquisition is illustrated for two such domains in Sections 4 and 5.

In Section 3.1 a generic model for diagnostic reasoning is described. Two specialisations of this model are presented in Sections 3.2 and 3.3. The relation between the models and the different levels of interaction distinguished in Section 2.1 is discussed in Section 3.4.

3.1 The SIX Architecture for Diagnostic Reasoning

As described above shared task models are acquired in interaction with experts, using existing (generic) task models to structure the process of knowledge acquisition. A generic task model of diagnostic reasoning designed for this purpose (Treur, 1993; Brazier & Treur, 1994) is shown in Figure 1.

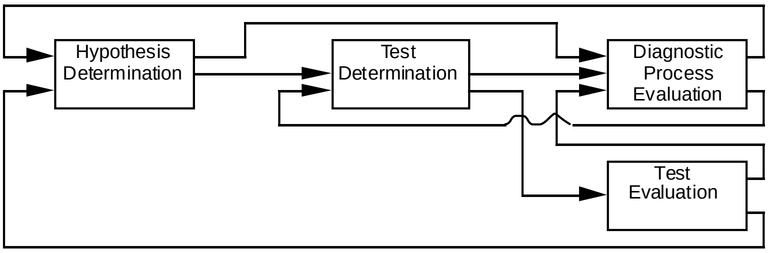

Fig. 1. SIX: a generic task model of diagnostic reasoning.

In this model four tasks are distinguished: *hypothesis determination*, *test determination*, *test evaluation* and *diagnostic process evaluation*. Hypothesis determination reasons about the appropriateness of possible hypotheses within a given state of the diagnostic process and determines which hypotheses are to be further investigated. Test determination analyses the current state of the diagnostic process with respect to test performance and determines which tests are most appropriate. Test evaluation performs the tests, and determines the relation between the test results and the current hypotheses. Diagnostic process evaluation analyses the implications of the test results for the hypotheses and determines which hypotheses are rejected and which are confirmed. On the basis of an analysis of the current overall state of the diagnostic process, the decision to conclude the diagnostic process may be made. If, however, the diagnostic process is continued, the required subsequent processes (for example, determination of hypotheses or tests) are identified.

Diagnostic reasoning processes can be based on *causal* or on *anti-causal* domain knowledge. In the first case derivations about the domain follow the direction of causality: the predicted observable consequences are derived from hypotheses (possible causes), after which (some of) the predicted observations are verified. For this type of reasoning causal knowledge is required that specifies how the causal consequences of hypotheses can be derived (e.g., represented by a causal network).

In the second case the domain knowledge is used to derive hypotheses from information on observables (symptoms). Here the direction of derivation is against the direction of causality: it proceeds from observable findings (in particular, those that actually were observed) to the causes. For this type of reasoning, knowledge is required that specifies how hypotheses can be derived from observable findings: this type of knowledge is called anti-causal knowledge.

In both cases strategic reasoning is required to determine the appropriate hypotheses on which to focus and the appropriate tests to be performed, as modelled by the generic task model for diagnostic reasoning SIX described above. This generic task model can be refined by specialisation to two slightly different models for diagnostic reasoning based on causal domain knowledge and anti-causal domain knowledge, respectively. These specialisations are described in Sections 3.2 and 3.3.

3.2 A Specialisation of the Generic Task Model for Anti-causal Diagnostic Reasoning

The specialisation for diagnostic reasoning based on anti-causal domain knowledge is obtained by decomposing the test evaluation task into two subtasks: *test performance* and *results interpretation*, as shown in Figure 2.

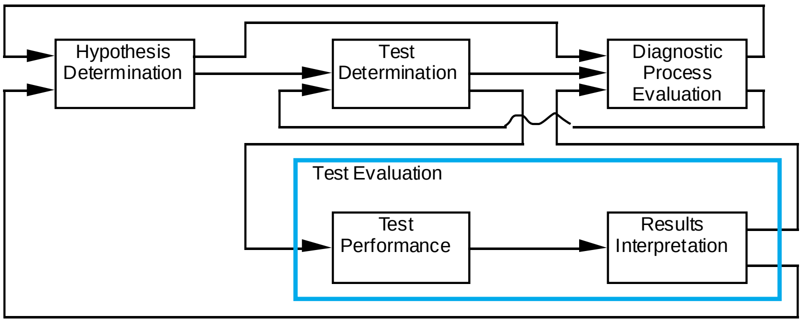

Fig. 2. A task model for diagnostic reasoning based on anti-causal knowledge.

Test performance is responsible for the “execution” of the tests selected by test determination. The results of the test may be acquired directly by object level interaction with an expert user, or may be acquired automatically from other systems. No further reasoning about the domain is performed in this task. The acquired test information is used by results interpretation to draw conclusions about the hypotheses, by means of the available anti-causal domain knowledge.

3.3 A specialisation of the Generic Task Model for Causal Diagnostic Reasoning

The specialisation for diagnostic reasoning based on causal domain knowledge is obtained by decomposing the test determination task into two subtasks: *test generation* and *test selection*, as shown in Figure 3.

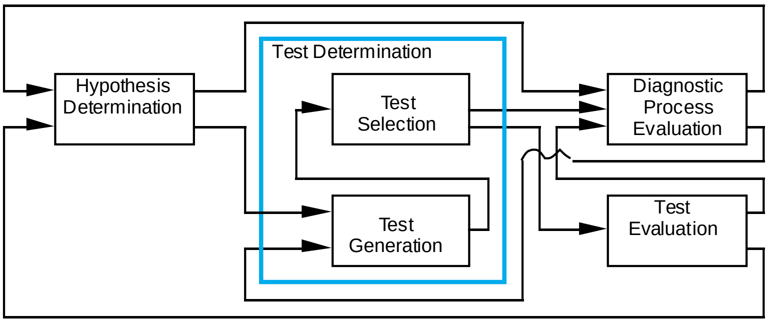

Fig. 3. A task model for diagnostic reasoning based on causal knowledge.

Test generation takes the hypothesis on which it is focussed as its input and using causal domain knowledge observable causal consequences are derived. These observable causal consequences are predictions of the findings that should be observed if the hypothesis hold. The predicted findings, influenced by the assumed hypothesis, are suitable candidates for tests to be selected. Test selection analyses the candidates and selects one or more tests on the basis of this analysis.

3.4 Interaction Levels in the Generic Task Models

In Section 2.1 three levels of interaction were distinguished. In this section the levels in the SIX model for diagnostic reasoning are discussed.

Object level interaction

The test evaluation subtask employs object level interaction with a user. The execution of tests may require users to provide (additional) information to the system.

Interaction at the level of strategic preferences

Strategic preferences are related to each subtask at the meta-level. Strategic preferences in the hypothesis determination subtask may influence the choice of hypotheses (e.g. frequency of occurrence, likelihood given a situation, preference of the user, …). Strategic preferences in the test determination subtask may influence the choice of tests (e.g. based on cost of the tests, duration, predictive power, order of execution, preference of the user, …). Strategic preferences in the diagnostic process evaluation subtask influence whether the system continues it search for (more) rejected and/or confirmed hypotheses or not.

Interaction at the level of task model modification

The user of the system influences which subtask is activated when. Knowledge of the sequence in which subtasks are activated may be overridden by users. For example, a user may look at the results of several settings for the strategic preferences for the test determination subtask before proceeding to the test evaluation subtask.

4 A Shared Task Model for Soil Sanitation

One domain in which a shared task model was developed in interaction with experts is the domain of soil sanitation (Boelens, 1991). During the acquisition process the generic task model of diagnostic reasoning (SIX) presented above played an important role. This model was used to structure interaction with the experts. In this section the domain of soil sanitation is introduced, an indication of the required functionality of a support system is given and finally the acquisition of the shared task model is described.

4.1 Soil Sanitation

Soil sanitation is a relatively young but fast-growing area of expertise. Polluted soil is found in many locations (in the Netherlands at least several thousand) and depending on the severity of the pollution the soil may need to be sanitized. At the level of provincial and local authorities the problem of soil sanitation usually is encountered during urban renewal. Pragmatic solutions are often chosen. Such solutions are based on two major decisions: how the site is to be sanitized and how the soil can be disposed.

Several procedures have been formulated concerning soil sanitation. Inventory research provides an indication of the different types of contaminations. Initial investigations aim to provide a global insight in the nature and concentrations of the contaminants. Further investigation concentrates on the nature, extent and concentrations of the contaminations as well as the spreading-probabilities. The goal of these investigations is to provide enough information for the sanitation procedure. The sanitation procedure consists of a comparison of the possible sanitation alternatives on environmental, technical and financial aspects. Sanitation is planned and executed.

The domain of sanitation consists of types of contaminations found (heavy metals, cyanide, aliphatic or hydrocarbons, aromatic compounds, and volatile helogenic hydrocarbons) and types of soil (sandy, loamy, loamy and clayey, peaty, and mixed). These types are only top-levels of taxonomies. Possible (general) sanitation techniques are: removing the contamination, prevent spreading of the contamination (isolation), or change of the function of the site. When removing the contamination either the soil is not removed (in situ techniques) or the soil is dug up. A soil sanitation alternative is a plan: one or more pollution remedial techniques are applied to the polluted site.

4.2 Acquisition of a Shared Task Model for Soil Sanitation

Experts working in the domain of soil sanitation were aware of the need for more support in choosing the best soil sanitation alternative. Although large bodies of knowledge are available the experts lacked support for flexible use of that knowledge. Ideally, the experts should be able to influence the use of the knowledge, when knowledge is used and what sanitation alternatives or tests may be investigated.

The following knowledge was readily available in pre-defined procedures and/or algorithms:

- How to choose remedial techniques based on their technological features and the situation at the polluted site.
- How to combine pollution remedial techniques into sanitation alternatives.
- How to predict the results of sanitation alternatives, based on the situation at the polluted site.
- How to compare sanitation alternatives to environmental standards or constraints.
- How to weight between various (groups of) evaluation criteria for sanitation alternatives.
- How to perform sensitivity analysis to determine which type of additional investigations is most effective with regard to selecting the best alternative.

Initial analysis of the experts’ reasoning process to find the best alternative for a specific situation given the option to collect additional information about the situation, is, in fact, a form of diagnosis. Experts agreed that this generic task model (the generic task model for diagnostic reasoning, described in section 3.1) provided a basis for subsequent discussion. The mapping of the terminology in the domain to the terminology in which the generic task model is presented, was relatively straightforward: sanitation alternatives in this domain are hypotheses, performance of “additional investigation” is the performance of tests, and acquired information corresponds to test results. The resulting task decomposition is shown in Figure 4.

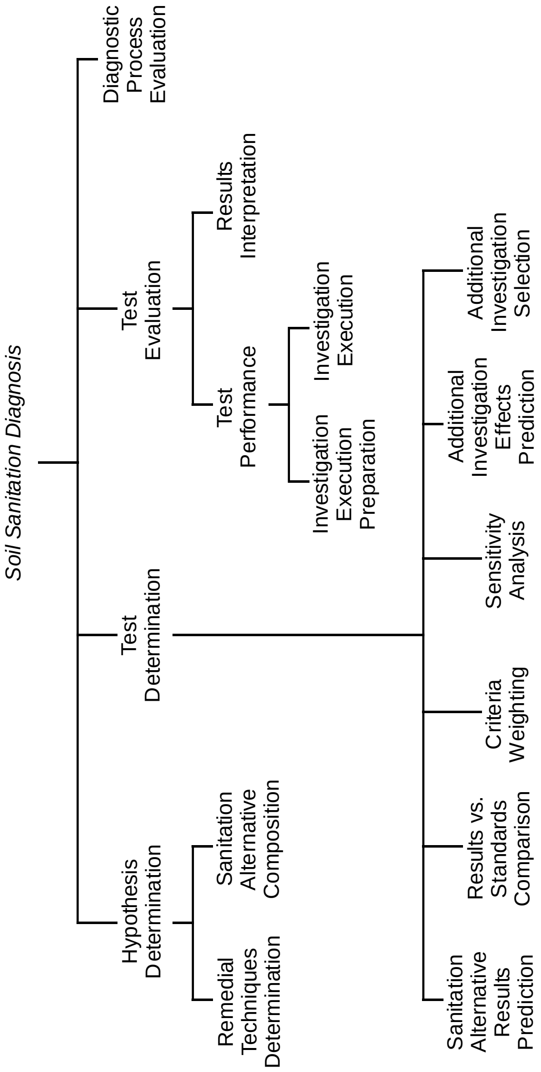

Fig. 4. Task decomposition of Soil Sanitation Diagnosis.

Hypothesis determination

During knowledge acquisition it became clear that the determination of the most appropriate sanitation alternatives, should be seen as two separate tasks. The first is the determination of the most appropriate technique for the reduction of one or more pollutants at a polluted site (*remedial techniques determination*). The second is the formulation of alternatives (i.e. hypotheses) on the basis of the available remedial techniques (*sanitation alternative composition*).

Test determination

The most extensive refinement of the generic task model was made with respect to test determination. Different subtasks were identified using different types of knowledge (including the knowledge mentioned above) to determine the most appropriate test.

Predictions about the (expected) reductions of pollutants at a particular site, are made using the available knowledge mentioned above (*sanitation alternative results prediction* task). Available knowledge is also used to determine the goodness of reductions, measured against directives on soil sanitation and construction materials (*results vs. standards comparison* task). Separate knowledge is available to decide how important different criteria are in the evaluation of sanitation alternatives (*criteria weighting* task).

Experts employ sensitivity analysis to determine which tests are most interesting in view of existing uncertainties (*sensitivity analysis* task). The knowledge experts have of models to predict the effect of tests on the criteria was also identified as a separate task (*additional investigation effects prediction* task). On the basis of the knowledge obtained by the performance of the above mentioned subtasks, a decision is made as to which additional investigations should be performed taking cost and duration (*additional investigation selection* task).

Test performance

Before actually performing tests experts reason about the information they expect to acquire and the way in which additional investigations should be performed (*investigation execution preparation*). This task is distinguished from the actual performance of the additional investigations (*investigation execution* task).

Results interpretation

Further decomposition of results interpretation was not necessary: the results of the tests are interpreted.

Diagnostic process evaluation

Experts recognized the appropriateness of a task for evaluation of the status of the process. On the basis of this analysis experts decide whether to pursue further analysis of a situation, or not.

Levels of interaction

Within the final version of the shared task model all three levels of interaction were modelled. Object level interaction is of importance in the test evaluation subtask, interaction at the level of strategic preferences in hypothesis determination and test determination subtasks, interaction at the level of task model modification in the task control process (using information provided by the diagnostic decision subtask).

5 Acquisition of Shared Task Models in Diagnosis of Chemical Processes

In a completely different domain, namely the domain of nylon production, the same SIX model was used during knowledge acquisition to structure discussions with an expert in this field. The expert involved identified the need for a system to support him in the diagnostic process, hopefully reducing the need for frequent on-site diagnosis. The nylon production process was described in detail and a few examples of types of problems with which the expert is confronted were discussed. As it was unclear how, in general, the expert structured his process of diagnosis, the two specialisations of the SIX task model described in Section 3: the causal and the anti-causal SIX task models, were introduced.

Initially, the process of nylon production in principle is based on causal knowledge in the domains of physics and chemistry, the knowledge engineers involved expected the diagnostic process to be based on causal reasoning. The SIX task model for causal diagnostic reasoning was introduced. Further discussion and analysis of cases of diagnostic processes, however, showed that during the diagnostic process in this domain hypotheses themselves could be confirmed or rejected on the basis of direct observation, i.e., no causal or anti-causal knowledge at all was required. In addition cases were identified in which hypotheses which could not be confirmed or rejected on the basis of direct observation, played an important role. In these cases the expert used anti-causal knowledge to derive hypotheses from observed findings. At this point the SIX task model for anti-causal diagnostic reasoning was introduced. The two models were compared, and the expert concluded that, in general, the anti-causal model was most applicable even though he realised that in some, more exceptional (and complicated) situations, the causal model would be more applicable (in which observable findings are derived from hypotheses). Hypothesis determination was further refined: a (limited) number of possible hypotheses are first identified, one of which is chosen for further examination. The first task is delegated to the system, the second to the expert user. By modelling the task in this way, the expert user explicitly and directly influences the reasoning process. The need for such strategic interaction was identified during the knowledge acquisition process.

The shared task model designed for diagnostic Nylon-6 production process is a specialisation of the generic task model for diagnostic reasoning based on anti-causal knowledge, presented in Section 3.3. The first version of a system for diagnosis of Nylon-6 production processes based on this model, has been implemented and is currently being evaluated. In other domains in the same chemical plant, the causal model has shown to be more applicable.

6 Discussion

To model a task in which an expert user and an intelligent decision support system collaborate, appropriate intermediate representations of the task at hand must be designed. The acquisition of a shared task model as an intermediate representation of the task (within which different levels of specificity are modelled), has been addressed in this paper.

The knowledge involved in a collaborative task, to the extent modelled in an agreed shared task model, includes the knowledge of different types of interaction involved within: (1) knowledge of the task structure, (2) knowledge of sequencing of (sub)tasks and goals, (3) knowledge of the knowledge structures, (4) knowledge of information exchange, and (5) knowledge of task delegation. These five types of knowledge are explicitly modelled in the declarative compositional framework for the design of complex reasoning tasks, DESIRE. Within the DESIRE framework existing abstract models of generic tasks, provide a means to structure initial interaction with the expert user during the acquisition of a shared task model. A number of agreed, shared task models have been used to develop applications (decision support systems) in different domains.

In this paper the principles behind the DESIRE approach to user-centered system design are presented and illustrated on the basis of the development of two applications of diagnostic decision support systems. In the first application, decision support in the domain of soil sanitation, one of the existing generic task models for diagnostic reasoning based on anti-causal knowledge provided a means to structure knowledge acquisition. The shared task model developed for this domain was, in the end, a specialisation of this existing generic task model.

In the second domain discussed in the paper, diagnosis of chemical processes, two existing generic task models for diagnostic reasoning were introduced: the first one based on causal knowledge, and later in the acquisition process the model based on anti-causal knowledge. In contrast to the knowledge engineers' expectations, the model based on causal knowledge was not in line with the expert's diagnostic approach. The anti-causal model, however, was useful: the acquisition process resulted in a shared task model for diagnostic reasoning of Nylon-6 production as a specialisation of this model.

The declarative nature of knowledge specification in DESIRE (for both examples), was of particular importance to modelling strategic preference interaction between the user and the decision support system. Explicit, declarative representation of strategic knowledge (for which modelling primitives exist within DESIRE) allows strategic knowledge itself to be subject of interaction, both from the user to the system (which preferences hold, which relations between preferences exist, etc. influencing the system's reasoning strategy), and from the system to the user (which preferences have been fulfilled, to which extent, etc.).

Not only the knowledge acquisition process (and task analysis) is structured on the basis of this shared model, but also the design of the interaction between the user and system. Three different levels of interaction between an expert user and an intelligent design/decision support system are distinguished in this paper: object level interaction, strategic preference interaction and interaction required for task model modification, each requiring specific modelling techniques. The role an agreed shared task model can play as the basis for modelling the necessary functionality of interaction between an expert user and the system, and thus as the basis for the design of an interface, is discussed in (Brazier & Ruttkay, 1993; Brazier, Treur & Wijngaards, 1996).

The role of shared task models in situations in which more than two parties (agents) are involved, is a current focus of research. A shared task model is an agreed model: in some situations agreement may be reached between more than two parties (resulting in a situation comparable to the situation described above for two parties), but in other situations different models of a task may exist between parties, thus requiring "attunement" between parties. Such collaborative tasks are currently being analysed, providing insight in the extensions required to the DESIRE framework.

Acknowledgements

This research has been (partially) supported by the Dutch Foundation for Knowledge-based Systems (SKBS), within the A3 project "An environment for modular knowledge-based systems (based on meta-knowledge) for design tasks" and NWO-SION within project 612-322-316, "Evolutionary design in knowledge-based systems".

References

P.J. Barnard (1993). Modelling users, systems and design spaces. In *Proceedings of HCI International '93,* Elsevier, Amsterdam, 1993, pp. 331-336.

J. Boelens (1991). *Soil sanitation and strategic interaction.* Masters thesis, Department of Mathematics and Computer Science, Vrije Universiteit Amsterdam, 1991.

F.M.T. Brazier and Zs. Ruttkay (1993). A Compositional, Knowledge-Based Architecture for Intelligent Query User Interfaces. In: S. Ashlund, K. Mullet, A. Henderson, E. Hollnagel and T. White (eds.), *Adjunct Proceedings of the INTERCHI '93* (INTERACT '93 + CHI '93), 1993, pp. 145-146.

F.M.T. Brazier and J. Treur (1994). User centered knowledge-based system design: a formal modelling approach. In: L. Steels, G. Schreiber and W. Van de Velde (eds.), "A future for knowledge acquisition," *Proceedings of the 8th European*

Knowledge Acquisition Workshop, EKAW '94. Springer-Verlag, Lecture Notes in Artificial Intelligence 867, 1994, pp. 283-300.

F.M.T. Brazier, J. Treur and N.J.E. Wijngaards (1996). Modelling Interaction with Experts: The Role of a Shared Task Model. In: W. Wahlster, ed. *Proceedings of the 12th European Conference on Artificial Intelligence, ECAI'96*. John Wiley & Sons, Ltd, 1996. To appear.

F.M.T. Brazier, J. Treur, N.J.E. Wijngaards and M. Willems (1994). A formalisation of hierarchical task decomposition. In: D. Fensel (ed.), *Proceedings of the ECAI '94 Workshop on Formal Specification Methods for Knowledge-Based Systems*, 1994, pp. 97-112.

F.M.T. Brazier, J. Treur, N.J.E. Wijngaards and M. Willems (1995). Formal specification of hierachically (de)composed tasks. In B.R. Gaines and M.A. Musen, eds. *Proceedings of the 9th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop KAW '95*, Calgary: SRDG Publications, Department of Computer Science, University of Calgary, 1995, Volume 2, pp. 25/1-25/20.

D.C. Brown and B. Chandrasekaran (1989). *Design Problem Solving; Knowledge Structures and Control Strategies*. Research Notes in Artificial Intelligence, London: Pitman, 1989.

B. Chandrasekaran (1986). Generic tasks in knowledge-based reasoning: high-level building blocks for expert system design. *IEEE Expert*, 1986, Vol. 1, pp. 23–30.

B. Chandrasekaran (1990). Design Problem Solving: a Task Analysis. *AI Magazine*, 11 (4), Winter 1990, pp. 59-71.

K.M. Ford, J.M. Bradshaw, J.R. Adams-Webber, and N.M. Agnew (1993). Knowledge Acquisition as a Constructive Modeling Activity. In: K.M. Ford and J.M. Bradshaw (eds.), *Knowledge Acquisition as Modeling, International Journal of Intelligent Systems*, Wiley and Sons, 1993, Vol. 8, Nr. 1, pp. 9-32.

F. van Harmelen, B. Wielinga, B. Bredeweg, G. Schreiber, W. Karbach, M. Reinders, A. Voß, H. Akkermans, B. Bartsch-Spörl, and E. Vinkhuyzen (1992). Knowledge-Level Reflection. In: B. Le Pape, and L. Steels (eds), *Enhancing the Knowledge Engineering Process -- Contributions from ESPRIT*, Elsevier Science, Amsterdam, The Netherlands, 1992, pp. 175-204.

R. de Hoog, R. Martil, B. Wielinga, R. Taylor, C. Bright and W. van de Velde (1994). *The Common KADS model set*. Deliverable DM1.1c of ESPRIT Project P5248 "KADS-II", 1994.

H. Johnson and J. Johnson (1991). Task Knowledge Structures: Psychological basis and integration into system design. *Acta Psychologica*, 1991, Vol. 78, pp. 3-26.

H. Johnson and P. Johnson (1993). Explanation facilities and interactive systems. In Gray, W., Hefley, W. & Murray, D. (eds.), *Proceedings of the 1993 International Workshop on Intelligent User Interfaces*, New York, ACM, 1993.

I.A. van Langevelde, A.W. Philipsen and J. Treur (1992). Formal Specification of Compositional Architectures. In: B. Neumann (ed.), *Proceedings of 10th European Conference on Artificial Intelligence, ECAI'92* ,Wiley and Sons, 1992, pp. 272-276.

J. Treur (1993). Heuristic reasoning and relative incompleteness. *Journal of approximate reasoning*, Vol. 8, 1993, pp. 51-87.