

Artificial Intelligence for Nature

Why Knowledge Representation and Problem Solving Should Play a Key Role in Environmental Decision Support

Peter Struss¹

Abstract

Decision Support Systems for the environment have to incorporate and exploit knowledge about the phenomena and the interdependencies in the affected natural systems, i.e. a model. Using an ongoing project on mangrove forest management as an example, we discuss different tasks and requirements and propose to use methods and techniques developed in research on modeling and model-based systems in Artificial Intelligence for computer support to solving the problems.

1 Introduction

Every day, we can witness how effectively modern technology influences and changes the world and our lives. The most outstanding effects, indeed, are destructive. By "destructive", we mean that human activities destroy the existing balance of complex natural systems - irreversibly, at a large and even global scale. Or is it not "impressive" that it took only a few decades to eliminate the stability of self-organizing systems that have existed and developed over thousands and millions of years: rivers, rain forests, even oceans and the atmosphere, interrelationships of species?

Despite this "success", humans will not be able to destroy the system of life on earth itself. It will survive human impact as it survived the impact of a meteor 60 million years ago, move to a different balance, develop new species, and eliminate others (perhaps including the human species). However, the impact of human activities has started to change and threaten the living conditions and even the lives of the originators of the disturbances.

In general, no person intended these effects. The original goals are usually constructive, positive: generating energy, producing food and timber, improving

¹ Department of Computer Science, Technical University of Munich, Orleansstr. 34, D-81667 Munich, Germany, email: struss@in.tum.de, and OCC'M Software GmbH, Gleissentalstr. 22, D-82041 Deisenhofen, Germany

le we will return to later, the purpose
s to avoid flooding and/or to generate
e the downstream transportation of
of sediment deposits in the river delta,
ure stagnant tidal water with the effect
o a point where it causes the dying of
clones, etc. - far from the dam - these

logy: it seems to categorize interde-
actually, it characterizes human inten-
guish between primary and secondary
activities, the others are "side-effects"
cipated or unknown to us. We do not
Environment" - this is an even more
environment of human society and
complex gigantic system - reduced to
activities, or, even worse, to a set of
ig. 1a). Let us "manage" "resources",

It is this ridiculous worm's-eye view
us from solving our "environmental
standing their causes. What is required
individuals and populations of living
l conditions, including the embedding
. This view, certainly alien to politics
h and created scientific progress. We
ns we are interacting with that reflect




Fig. 1 1a The anthropocentric view 1b A "natural" view

2 The Role of Information Technology

This insight is a great achievement. However, it is only the starting point of the difficulties. It only poses the question. Finding the answers is the hard part. Why? Just because the systems we are looking at are so complex, incorporate so many participating substances and organisms, contain so many interactions, actual ones as well as potential and hidden ones. We are simply overwhelmed by the complexity when we try to model them, by the amount and structure of data when we try to identify behavioral patterns and to exploit them for explanation of existing phenomena or prediction of future ones, by unforeseen responses of the systems when we attempt to influence them according to our intentions. We know only bits and pieces. And even when we believe to have understood some fundamental mechanisms in one system, they almost certainly have to be revised when we try to apply them to a similar system, because some hidden preconditions are not satisfied, or new phenomena cancel out the expected effects.

Now, there exists also a technology that has been developed to cope with large amounts of data and complex information, to prevent them from overwhelming us. Computers and information technology have provided effective tools that can support research and decision making on "environmental problems". They enabled substantial results in acquiring information and in the analysis and prediction of the behavior of natural systems and their response to human impact. In particular, they help to

- acquire data, e.g. through remote sensing and image processing,
- store and retrieve data, e.g. in data bases and geographical information systems,
- transfer data, e.g. via the internet,
- process data, e.g. in statistical analysis and in simulating numerical models of ecosystems.

Although these computer systems are often extremely useful and even prerequisites for improving our understanding and decision making, they do not directly address these tasks that we have identified to be crucial. They provide and process **data**, rather than knowledge, and, in fact, using them requires substantial knowledge:

- A basic understanding is already necessary to determine what kind of data should be acquired or retrieved to provide insight.
- If we have a model as an input to a simulation system, the most difficult part of the work has already been done.
- And having obtained data, by means of sensing, retrieval, or simulation, they are of no use without an interpretation based on prior knowledge. Knowledge is what turns data into information.

Currently, all these knowledge-intensive tasks have to be carried out solely by human experts, mainly researchers, and computer support for them is almost non-existent. In a sense, the improved facilities for acquiring and generating data even creates

problems rather than solving them: huge amounts of data are buried on tapes and data bases, and many of them will never be excavated because interpreting them and drawing conclusions is by orders of magnitude more time-consuming than creating and storing them.

But is it really feasible to build computer systems that support such knowledge-intensive tasks? Does this not imply that they must be able to capture and manipulate at least a substantial part of the knowledge that their users need to solve problems? Yes, it does. Artificial Intelligence pursues the goal of providing theoretical foundations, methods, and systems for such capabilities. They have proven to be of use in other application domains, and, in particular, in modeling other kinds of physical systems, namely technical systems, artifacts, and solving problems related to them, such as configuration and diagnosis. We will discuss what can carry over to establish solutions to the problems discussed here.

To become more specific, we will use a particular project we are involved in to extract typical tasks, problems, and requirements from it.

3 Tasks and Requirements

We started a collaboration with the M.S. Swaminathan Research Foundation (MSSRF) and the Indian Institute of Technology (IIT) Madras. MSSRF is conducting research on the causes of degradation of mangrove forests in the coastal wetlands of South India. As in other regions of the world, this degradation has reached a stage that represents a serious threat to function of the mangrove forests as a shelter against cyclones, a habitat for various species, and a source of timber and food for the local population. While MSSRF can investigate the interdependencies in this ecosystem and the factors influencing its development and, based on these results, develop strategies for the restoration and management of the mangrove forests, ultimately decisions and actions have to be taken locally, by forest departments and the authorities of the local self-administration.

As the local conditions and relevant factors (in terms of geomorphology, hydrology, climate, species, exploitation etc.) vary considerably, it is impossible to produce general management plans that can be applied to all sites without further investigation, modification and specialization. On the other hand, MSSRF's experts are not available for continuous consultation at the various mangrove areas. In response to this situation, our collaboration aims at capturing their growing expertise in a knowledge-based system which can then support local investigations, problem solving, and decision making at the various sites in reflection of their specificity. Certainly, a geographic information system is a useful tool which can store the specific data of each site. But it is fairly obvious that it cannot provide the solution to the ambitious goals. Pure retrieval of existing data, although important to the work of the experts, offers very limited help to the kind of users involved in the management

of the mangrove areas. On the one hand, they would have to know what to inquire, and, on the other hand, for most of the problems, one cannot expect the solution to have been produced in advance and stored in the data base. Rather, it has to be constructed by reasoning about facts from the data base under the guidance of principled knowledge about the domain and task. For the non-experts, at least part of this knowledge and the appropriate reasoning mechanisms have to be provided by the knowledge-based system.

We informally discuss some of the problem solving tasks and their requirements. In section 4, we will revisit them and propose how to formalize them in order to provide a basis for automated reasoning systems.

Situation Assessment (What's Going on?)

Given a set of observations (measurements, descriptions of geomorphological or biological features, results of visual inspection, etc.) and general domain knowledge, determine the relevant phenomena and processes which cannot be or have not been observed directly. We might distinguish two different kinds of this tasks, although they frequently have to be solved together:

- **System Identification** which again has two aspects. One is **structure** identification, i.e. determining the (types of) constituents of the system and their interactions. Second, **parameter** identification, i.e. determining constants that characterize the particular instances of such constituents. The starting point is a description of the entities that are directly observable (the topography, present species, etc.) and measurable quantities (salinity of the water, current, etc.). The goal is to establish a behavior model of the target system. Obviously, this is a reasoning task which requires scientific knowledge and expertise, since it is about recognizing the processes that are caused by the visible configuration of entities, but happen behind them, often difficult or impossible to observe (e.g. deposit of sediments, chemical processes, etc.). Computer support to this task should satisfy a number of requirements. First, knowledge that has been gained in modeling similar systems and through generalization should be **re-used** directly where appropriate. Thus, we want to avoid having to build a model of each mangrove area from scratch. Second, the input to such a modeling system should be entered in a way that is natural to the user who may be a domain expert or a local semi-expert and whom we cannot force to use some mathematical formalism or computer language. Third, many of the available observations are inherently qualitative in nature, such as "trough-shaped area" or "increased degradation", but, nevertheless, carry crucial information that has to be taken into account. .

- **State Identification** interprets observation in order to infer internal states or tendencies in the system **at a particular time**. Obviously, this is based on, or combined with, system identification.

Both tasks can occur as "diagnostic" tasks if the given observations do not match the expectations or a given model.

Diagnosis (What's Going Wrong?)

This step tries to identify the causes of deviations of a system from what we consider normal or healthy. For instance, the threatening degradation of mangroves is a particular area of the Pichavaram forest mentioned earlier would require tracing the cause-effect chain sketched earlier back to the Cauvery dam. This is an example of

- **Identification of Ultimate Causes**, whereas
- **Identification of Controllable Causes** focuses on determining hooks for curing the symptoms in case we are not able or willing to remove the ultimate causes or if this does not suffice to also remove the undesirable effects. Since it is not easy to remove or open the dam, increasing the drainage rate of the water in the troughs is an options for reducing the salinity (whereas influencing the evaporation rate is not).

For diagnosis, the system needs, besides the objective description of the system under analysis, an appropriate representation of the desired states of evolution of the system. Note that this diagnostic reasoning may involve more than hypothesizing abnormal parameters or state values, but may involve the revision or extension of the original model in order to include the origins of the disturbances and all their effects. Again, the computer tool has to address the requirements stated above, in this task even more, since it should be part of the continuous use of the local authorities.

Therapy (What Can be Done?)

Given descriptions of a disturbed system (including the causes for the disturbance as discussed above) and of our goals, determine actions suited to re-establish a state that complies with our goals. In the simplest case, such actions are designed to shift a single quantity in the proper direction, e.g. increasing the flow from the trough to the sea by digging canals (which is what is actually being done). In general, the task can require planning a sequence of actions over time achieving a number of different sub-goals ("First, take steps to decrease salinity, then introduce species X to change the soil, finally, re-introduce species Y"). In addition to representing the ecosystem itself and the goals of management or conservation, computer support to this task needs to incorporate knowledge about the potential interventions, and, more specifically, this has to be done in a way that enables the analysis of their impact on the ecosystem.

Prediction (What's Going to Happen?)

Given a system model and some, potentially incomplete, description of its state (the identified actual one or a hypothesized one), anticipate the possible evolution of the system in the future. The purpose can be to check whether the actual state may lead to a critical one, or to explore the consequences of some action (of exploitation or therapy) before it is actually taken. Although this may sound like ordinary numerical simulation, one has to keep in mind that the initial conditions as well as the parameters may be specified only partially (e.g. qualitatively) and lead to ambiguous predictions. This is particularly true for the evaluation of potential actions which we cannot or would not want to describe numerically. Initially, we may be only interested in determining whether a certain class of measures could move the system in a proper or tolerable direction: "What will happen if a crab farm is established in this bay?", "What is the consequence of digging a canal from the trough-shaped areas to the sea?". Second, also the output, the predicted behavior, has to be stated in terms of concepts that are familiar to the user, rather than tables of numbers.

Explanation (Why Does It Happen?)

Describe the ecosystem and its response to disturbances and interventions in a way that is comprehensible to the user. Suppose the computer system had hypothesized the reduction of sediment deposit as a potential cause of mangrove degradation (a rather difficult task), then we would like it to deliver a "causal chain" of the kind presented in section 1. This is a crucial feature, because it enables the critical assessment of the performance of the system particularly for validation of the knowledge base. Furthermore, it is essential to the educational purpose of the system with regard to the local decision makers.

Discovery (What Can Possibly Happen, and under which Conditions?)

Given the existing theory about a domain (such as mangrove ecosystems), a description of a specific system (mangrove forest), and system observations that contradict the general theory or cannot be explained by it, propose revisions and extensions to the theory in order to account for the observations. This aims at supporting the research process itself and is certainly one of the most challenging tasks.

The easiest step is the mere detection of the contradiction. The following one would be to localize the elements of the theory that contribute to the contradiction (again, a "diagnostic" task!) and hence, are candidates for a revision. Beyond this,

hypotheses for certain revisions of these elements or, more difficult, for undiscovered interactions could be generated. In order to decide which of these potential revisions is appropriate, helpful observations or experiments should be proposed. All this may sound too futuristic. However, we will see what is required to provide such tools, and there exist not only theories, but also implemented tools.

4 Artificial Intelligence Approaches to Modeling and Model based Problem Solving

In our initial discussion, we emphasized in general terms the importance of explicitly representing general knowledge about the ecosystems under analysis in computer programs (as opposed to rules and algorithms special to particular goals and tasks). The examples in the previous section were meant to illustrate the crucial role of modeling and model-based reasoning.

Although traditional numerical modeling and simulation can and has to perform part of the work (e.g. in prediction when exact data are both available and necessary), the requirements discussed show that they cannot provide a solution to the central problems. For representation and presentation (in particular explanation) of existing and generated knowledge, conceptualization is essential.

- The system needs to maintain **a mapping to the phenomena in the physical world** as they are perceived by the user both for the model and the results obtained from it. In numerical modeling, this mapping has to exist in the heads of the experts when creating the model and when interpreting the generated data. But is lost in the mathematical model and in the simulation system. Differential equations and tables with numbers do not explain anything, and particularly for the non-expert users, they are not appropriate for conveying knowledge and information. This is why a weather report is presented (and, in fact, generated!) in terms of conceptual entities, such as low pressure systems, moving cold fronts, etc.
- Second, in our domain, much of the data, information, and knowledge can only be stated and should be communicated in **qualitative** terms. This is why qualitative and symbolic reasoning methods developed in Artificial Intelligence need to be exploited (Forbus 84), (Struss 97).

Both statements hold, especially, for spatial and temporal aspects that require an appropriate level of abstraction.

- The desired **re-use** of model elements also requires a structuring of the models in accordance with **conceptual entities** of basic phenomena and processes that occur in various systems in the domain.

This is why computer-supported conceptual modeling is needed (Heller-Struss 97). We will outline some theories, methods, and techniques that have been developed in the Artificial Intelligence area of knowledge representation and reasoning and that

offer means for tackling the tasks and requirements, without claiming to be comprehensive or to present ultimate solutions. Some of the aspects are discussed with more technical detail in another contribution to this volume (Heller-Struss 98) and in (Struss-Heller 98).

Process Modeling

Like to most other Artificial Intelligence approaches to modeling, compositionality is fundamental to process-oriented modeling (Forbus 84), (Struss-Heller 98): the behavior model of a specific system is derived through aggregation of re-usable model fragments that represent elementary processes in the domain. Model composition follows **structure-to-function reasoning**. This means, a description of the configuration of objects that constitute the system entails a set of relevant interacting behavior model fragments which are additionally controlled by certain conditions on the quantities involved.

Knowledge connected to each (generic) process comprises the following elements:

- The **conditions** under which the process is included in the model. This is stated in terms of structural preconditions, i.e. certain objects and relations between them, and constraints on parameters and variables involved in the process. For instance, an evaporation process requires a water body in touch with the air and a humidity of less than 100%.
- A set of **relations** among quantities local to the process. According to the nature of knowledge in the domain, they can represent relationships in a qualitative way, e.g. as monotonic dependencies. In the evaporation example, the evaporation rate may be stated to grow monotonically with the water temperature.
- A set of **influences** that capture the contributions of such a process to changes in the system. This concept reflects a requirement that arises from the compositional modeling scheme: we need to describe the effects local to a model element; but without knowing which other model elements will affect the same quantity, no definite constraint can be established for the influenced quantity. For instance, the evaporation rate influences the amount of the (liquid) water body negatively. But this may, nevertheless, grow due to other processes, such as precipitation. Influences can be positive or negative and act on a variable or its derivative. Roughly, an influence is a statement about the partial derivative of the influenced variable w.r.t to the influencing one.

Somewhat simplified, we can formalize the concept of a process in logic by stating that the process conditions imply certain constraints (i.e. mathematical relations), influences, and, potentially, other structural properties (since a process may generate new objects and relations):

$$(1) \quad \text{STRUCTURE-CONDITIONS}_i \wedge \text{QUANTITY-SPEC}_i$$

$$\Rightarrow \text{CONSTRAINTS}_i \wedge \text{INFLUENCES}_i \wedge \text{STRUCTURE}'_i$$

We can now categorize the contents of the knowledge base for our enterprise:

- The **domain theory** which represents the general knowledge about a particular type of ecosystem. Its core is the set of descriptions of generic processes as introduced above (a "library"). It has to be comprehensive in the sense that it contains all model elements required to model special ecological system of the respective type (of, course, it will never be complete).
- A **system structure** description which specifies the objects that constitute a particular system and their relations (e.g. spatial ones).
- **Quantity specifications** complete the specification of the ecosystem by determining parameters of objects involved. Secondly, they can represent observations about a particular **state** of this system.

We will now discuss how the various tasks can be formalized and implemented. It turns out that already the model composition step addresses some of the issues raised in section 3.

Situation Assessment, Step One: Model Composition

The idea underlying the process descriptions confines the input of a user to a "superficial" system description in terms of constellations of objects, their parameters, and observations about a particular state of the system. What the user does not have to enter or even know is what the relevant behavior models are. They are created automatically by the domain theory from the user input by applying process descriptions of the kind (1):

$$(2) \text{DOMAIN-THEORY} \cup \text{STRUCTURE-DESCRIPTION} \cup \text{PARAM-SPEC} \\ \vdash \text{CONSTRAINTS} \cup \text{INFLUENCES}:$$

At this point, for each variable, all existing influences on it have to be combined to form a constraint. The problem lies in "**all influences**". If the initial structure description missed some facts or was based on certain assumptions, then (2) may fail to generate all relevant influences. Speaking in logical terms, in order to turn influences into a constraint, the system has to make "closed-world assumptions" (CWA) meaning that there exist no additional objects or relations that might generate another influence on a variable, y :

$$(3) \text{INFLUENCES}(y) \wedge \text{CWA}_y \vdash \text{INFL-CONSTRAINT}(y).$$

Based on observations, the generated behavior model (a constraint network), can then complete the state description (as well as the parameter specification):

$$\text{CONSTRAINTS} \cup \text{QUANTITY-SPEC} \vdash \text{QUANTITY-SPEC}'.$$

Due to the lack of space, we can only mention that there are some non-trivial problems to solve because of the cyclic dependencies of behavior models (which can generate objects and derive quantity values) and process conditions on objects and quantities (which lead to behavior models) and refer to (Struss-Heller 98).

We also point out that (1), (2), and (3) enable the knowledge-based system to record which elements of the structure description, the quantity specification, the domain theory, and even the closed-world assumption lead to the existence of a particular constraint and, thus, to certain computed values. This is important, if some of these inputs are hypothetical or invalid in a given situation (e.g. due to disturbances in the system) and may have to be revised.

Model Revision - A Fundamental Task and Method

It turns out that several of the tasks of section 3 can be regarded and also implemented as a process of revising a given model. For instance, in situation assessment, we may start with a model of the normal state, but the given observations may contradict this model and force us to retract some normality assumptions to gain a picture of the actual situation.

Formally, the general revision process starts with a model that is inconsistent with the observations or some goals (both represented as sets of quantity specifications):

$$(4) \text{ MODEL}_0 \cup \text{OBSorGOALS} \vdash \perp .$$

Its result is a modified model (or several candidates) that removes the inconsistency with the observations (for situation assessment) or with the goals (for diagnosis and therapy):

$$(4') \text{ MODEL}_1 \cup \text{OBSorGOALS} \not\vdash \perp :$$

A stronger requirement would be that the revised model entails the observations (i.e. "explains" them) or the goals (i.e. accomplishes them):

$$(4'') \text{ MODEL}_1 \vdash \text{OBSorGOALS} :$$

Such a revision process is the core of model-based diagnosis systems that have been developed for technical systems and which currently enter industrial applications (Hamscher-Console-de Kleer 92), (Dressler-Struss 96). Their techniques also form the basis for solutions in our domain. The key to a focused proposal of model revisions is the following: the inconsistencies occur as conflicting values for variables derived from observations and the constraints of the behavior model. Techniques that record the dependencies of the constraints and values on specific elements of MODEL_0 identify candidates for a revision. Guidance is usually given by some minimality criterion (w.r.t. sets, number, or probability of revisions).

Usually, there is not a unique proposed revision. Since the model-based system can explore the consequences of different hypotheses and, hence, their distinctions, it

is able to propose measurements or tests that help to narrow down the set of candidates.

It is easy to see how these techniques can localize wrong assumptions about quantities or the existence of certain structural elements. The closed-world assumptions which underlie the INFL-CONSTRAINTS provide an entry point for hypothesizing objects that have **not** been anticipated initially. (Struss-Heller 98) discusses how a revision of a closed-world assumption can lead to a focused search for additional influences on a variable and the conditions that could generate it.

With this general background, we can formalize several tasks, especially by determining what is subject to potential revisions: the MODEL split into a fixed part which is considered to be fixed (true, unchangeable, not controllable), $MODEL_{fix}$, and the part which may be revised, $MODEL_{rev}$.

System Identification - Step Two: Observation-guided Revision

Here, assumptions about the system may be revised, and there are no goals taken into account:

$$\begin{aligned} MODEL_{rev} &= STRUCTURE_{rev} \cup PAR-SPEC_{rev} \cup CWAs, \\ GOALS &= \emptyset. \end{aligned}$$

State Identification - Step Two: Observation-guided Revision

In this case, it is assumed the system has been identified, but assumptions about the state may have to be retracted:

$$\begin{aligned} STRUCTURE \cup PAR-SPEC \cup CWAs &\subset MODEL_{fix}, \\ MODEL_{rev} &= VAR-SPEC_{rev}, \\ GOALS &= \emptyset. \end{aligned}$$

Diagnosis

Here, the task is to identify the origins of contradictions between given goals and the description of the system and its state as it has been derived from situation assessment. We make parts of this description revisable, not because they may be wrong, but in order to find modifications of the model that remove the conflict with the goals,

$$MODEL_1 \cup GOALS \not\vdash \perp,$$

hence, could have caused the conflict, and provide a starting point for therapy:

$$MODEL_{rev} = STRUCTURE_{rev} \cup PAR-SPEC_{rev} \cup VAR-SPEC_{rev} \cup CWAs,$$

Therapy

Therapy starts from a (not necessarily unique) result of the diagnostic step and has to identify actions that may implement the changes to the system proposed by this result:

$$\text{MODEL}_0 \cup \text{ACTIONS} \cup \text{GOALS}' \vdash \perp$$

This means the model knowledge base has to include actions. One way to achieve this is to represent the effects of an action as a process and the action itself as its (only) precondition. Thus, the revision process has to revise the closed-world assumption for variables that need to be influenced and then search for appropriate influences in the subset of the domain theory that represents potential actions.

Note that during therapy, the goal may be different from the ultimate ones, because it is in general impossible to satisfy them immediately. If the actual value of some variable is below what is specified by the goal, then an intermediate goal may be to increase the value, i.e. influence its derivative positively.

As stated above, the overall treatment may involve a real planning task which could exploit the model revision scheme presented here for generating steps of the plan and checking its compliance with the goals. Also, some goals may have to be violated temporarily, and by turning some goals into revisable options, the revision process can identify (minimal) sets of goals to be abandoned.

Prediction

This is simply the "model in action" over time, i.e. simulation:

$$\text{MODEL}(t=0) \vdash \text{MODEL}(t=t_1) .$$

But besides the computation of variable values over time (qualitative or numerical), a conceptual view of the evolution is derived in terms of activity of processes, creation and elimination of objects etc. This is important for the next task.

Explanation

This task now really benefits from the conceptual layer: rather than answering questions, such as "What happens?" or "Why does it happen?" by spitting out tables or plots of numerical values, it can present how certain processes were created, what their effects were, etc. Even questions why certain effects did **not** happen could be answered, based on the analysis of unsatisfied process conditions or influence combination.

Discovery

For all the previous tasks, the domain theory was not considered for revision. If we include elements of it (hypothetical processes or parts thereof) in the revisable part of the model, the general diagnostic technique could propose such elements for further inspection or modification. Finding missing processes could be supported by hypothesizing additional influences on variables that cannot be accounted for by existing process descriptions. Again, for alternative repairs of the theory, the system would be able to propose discriminating observations or experiments.

Of course, the ultimate driving force is the researcher. But a model-based system as his apprentice is a challenging perspective.

5 Conclusions

In this paper, we argued that

- solving our "environmental problems" requires to gain more insight in the complex interactions in ecosystems of which human activities are only a small part, i.e. **develop and use better models** of such systems,
- computer systems should be designed and implemented that support this task and, hence, have to support **conceptual modeling and problem solving** based on such models,
- theories, methods, and techniques developed in Artificial Intelligence research on knowledge representation and reasoning and, in particular, **qualitative modeling and model-based systems** are a promising starting point to pursue this goal.

This is why, although the goals are rather ambitious, there can be accomplishments, even in the short term. This is not the claim that we believe to have solved all fundamental problems. On the contrary, a lot of research and development of tools needs to be done. And we did not even touch some important research subjects, such as qualitative spatial and temporal reasoning. But there exists a strong basis in terms of theory, methods, **and implemented tools** and systems for tackling the tasks mentioned above in real applications.

What is required to progress? It is almost certain that wrong computer tools will be developed as long as the researchers in ecology do not pick up the offered solutions and spell out their requirements and Artificial Intelligence researchers do not expose their methods to these challenges. What is needed right now is closer collaboration between the domain experts and Artificial Intelligence researchers on developing principled solutions to specific problems.

Acknowledgments

Many thanks to the members of the MQM group at the Technical University of Munich for discussions and collaboration, and also to our project partners at the municipal department of water and sewage in Porto Alegre, Brazil, and the M.S. Swaminathan Research Foundation in Chennai, India, for supporting our work with their knowledge about the domain.

References

- Dressler, O., Struss, P. (1996): The Consistency-based Approach to Automated Diagnosis of Devices. In: Brewka, G. (ed.), Principles of Knowledge Representation, CSLI Publications, Stanford, ISBN 1-57586-057-0, pp. 267-311.
- Forbus, K. (1984): Qualitative Process Theory, in: Artificial Intelligence 24(1-3).
- Hamscher, W., Console, L., de Kleer, J. (eds.) (1992): Readings in Model-based Diagnosis, Morgan Kaufmann Publishers, San Mateo.
- Heller, U., Struss, P. (1997): Conceptual Modeling in the Environmental Domain. In: 15th IMACS World Congress on Scientific Computation, Modeling and Applied Mathematics, Berlin, August 1997, ISBN 3-89685-556-5, Vol. 6, pp. 147-152.
- Heller, U., Struss, P. (1998): Diagnosis and Therapy Recognition for Ecosystems - Usage of Model-based Diagnosis Techniques, this volume.
- Struss, P. (1997): Model-based and qualitative reasoning: An introduction. In: Annals of Mathematics and Artificial Intelligence 19 (1997) III-IV, Baltzer Science Publishers.
- Struss, P., Heller, U. (1998): Process-oriented Modeling and Diagnosis - Revising and Extending the Theory of Diagnosis from First Principles, in: Working Notes of the 9th International Workshop on Principles of Diagnosis (DX-98), Cape Cod.