Deep Knowledge and Domain Models

Jarmo J. Ahonen *

December 16, 1993

Abstract

An approach to the concept of deep knowledge is outlined. The approach is based on the assumption that the deepness of knowledge results from its explanatory powers. After considering some examples of deep and shallow knowledge and defining deep knowledge and robustness, an approach to the development of *quantitative* domain models based on deep knowledge is proposed. The proposed approach is based on the Salmonian concept of causal processes and it provides a uniform point of view to knowledge of physical domains and domain modeling.

1 Introduction

Since the Chandrasekaran and Mittal article (Chandrasekaran and Mittal, 1983) most of the writers concerned with so-called domain models have been using the term *deep knowledge*. Unfortunately the term has not been defined very well, and its use has been depending on vague notations of its nature. The vagueness of the meaning of the term invites us to try to give it more meaning, and one approach to the concept of deep knowledge and its nature and its use in domain modeling will be discussed later on. In addition to that some of the most obvious consequences of the use of deep knowledge to the overall usefulness and structure of domain models will be considered.

The use of the term deep knowledge seems to be closely related to the idea that we cannot create practical AI systems without turning our attention to the physical reality around us. Hayes proposed the use of models of the physical reality as a method of avoiding the problem of too simple and simplified problem domains (Hayes, 1979). In addition to the interest in physical domains Hayes said that reasoning is not enough, from which it may be concluded that an intelligent system would include both reasoning and knowledge of the physical reality. Such division

*Lappeenranta University of Technology, P.O.Box 20, 53851 Lappeenranta, Finland

leads us to adopt the idea of *model based reasoning*, according to which the descriptions of domain and reasoning are kept separate. The model based reasoning approach is more a methodology rather than a technique, because it does not determine the implementation (Saarenmaa, 1988) The model based approach seems to be especially useful in fault diagnosis and other fields that require the use of real-world models (see e.g. de Kleer (1987), Koton (1985), Nardi and Simons (1986), Xiang and Srihari (1986), Rich and Venkatasubramanian (1987), Cross (1984), and Adams (1986)). It is, however, worth stressing that our main interest will be in *quantitative* domain models, i.e. models which produce numeric results, and not so much in qualitative models normally considered when speaking about domain models (especially domain models in connection with the concept of expert systems). In that respect we have a clearly different emphasis than Clancey who has left quantitative models out of in his well covering considerations (Clancey, 1992).

In fault diagnosis and simulation there is one obvious requirement for the domain model, namely that it should be able to function as the modeled phenomenon in every case. Unfortunately the method by which this can be achieved is not clear, although some results and authors speak in favor of deep knowledge. The problem with deep knowledge is mainly that we do not know what it is although we know what it should do, i.e. provide better robustness for domain models.

An interesting possibility to approach deep knowledge is to consider the deepness of knowledge possessed by human experts. The human performance is especially interesting because human experts seem to be able to estimate the functioning of a physical system even in cases that have not been encountered earlier and because studies of cognitive processes of human experts seem to suggest that they, i.e. human experts, use some kind of domain models in their reasoning (e.g. Sweller's (1988) results support this belief).

2 Models

In our consideration the term *physical reality* means the concrete world. Similarly a *real-world phenomenon* or a *real-world object* is a phenomenon or an object which exists or can potentially exist in the physical world. A *model* is a representation of the physical reality. Such representations are simplifications of the physical reality, and unfortunately they will not produce reliable information on every aspect of the phenomenon being modeled (Lewis and Smith, 1979, 2). Our definition of models differs from normal definition in its clear nature as a representation of the physical reality (see e.g. (Futo 1990) for the normal approach).

Sometimes it may be the case that the model is intended to be used in a situation that has not been considered during the development of the model. Such a *new situation* is a case in which the model is intended to be used without prior knowledge of the possibility that such a case may be encountered during the use of the model. A model is *robust* if it behaves like its real counterpart would behave in a similar new situation. In other words, a robust model produces right values to the variables used to represent the modeled real-world phenomenon or object.

Models should be designed in a way that they do not lose their modeling ability even in new situations. The obvious way to achieve this goal is to develop models by using knowledge that is required to determine the behavior of the model in as many new situations as possible and do it despite the fact that models are simplifications (note that the models of human experts are that also). We will call such knowledge deep knowledge because it can be said to go beyond the surface of the knowledge on which our models are based. *Deep knowledge* is knowledge that defines the structure and behavior of the phenomenon or the object considered. Deep knowledge provides the basis on which robustness may be founded.

Our definition of deep knowledge is fairly obvious, although it leaves two important questions open. The first question is the actual nature of deep knowledge, and the other question is the structure of models that are based on deep knowledge.

3 Shallow models, explanation, and robustness

In the field of the philosophy of science, one of the most flourished topics of interest is the explanative nature of scientific knowledge. For our purposes we will consider the explanatory aspect of scientific knowledge, and suppose that scientific knowledge consists mainly of explanations. (This assumption seems to be a safe one because the explanatory nature of scientific knowledge tends to be so widely accepted proposition, see e.g. (Fetzer, 1981), (Niiniluoto 1983), (Salmon, 1971), and (Salmon, 1984), that it is not seriously questioned at all.) In order to avoid a multitude of arguments we will assume that concepts are considered to be defined in such explanations.

To clarify the connection between robustness and explanatory powers of models we will briefly consider so-called shallow models, which are often thought to be the opposite of deep models (see e.g. Chandrasekaran and Mittal (1983)). Although we will not define shallow models explicitly, we will briefly consider some of their features. For us the term shallow-model is interpreted in a way that makes clear that shallow models include primarily shallow knowledge acquired by e.g. empirical observations, which are then written into rules or equations that define the values of variables by relations between those values. If we consider shallow models to include only empirical knowledge acquired by observing the surface behavior of an object or a phenomenon, then we can assume that the features of such models differ from the features of the respective deep-knowledge models. In order to emphasize the distinction we will briefly consider some simple examples of models that provide desired robustness or do not provide it.

Consider, as an example, a situation where a human being cannot get enough vitamin A. It is well known that too little vitamin A causes illness, and it is similarly well known that by giving more vitamin A to the person in question, his/her health can be improved. From a limited amount of knowledge it can be concluded that by giving more and more vitamin A person's health will improve infinitely. It is, however, known that it is not the case. Knowledge of the biological low-level effects of the vitamin A would have enabled the avoidance of the wrong conclusion.

As an another example we can consider a model of the growth of pines. Having gathered a great number of observations, it is possible to generate a function, and say that

> 'According to our empirical observations we can give the statistical relationship between the height growth and the age of pines as
>
> $$h(t) = \frac{H}{1 + h_1 e^{-h_2 t}} \tag{1}$$
>
> where H (max. height), h_1 and h_2 are 21 m, 20.4 m and 0.064, respectively'

which is true. But by saying (as Oker-Blom et al (1988)) that

> 'The height growth was assumed to be independent of stand density and was modeled as a logistic curve (1) The values of H (maximum height), h_1 and h_2 were chosen as 21 m, 20.4 m and 0.064, giving a height development in accordance with existing growth and yield tables.'

and implicitly suppose that the equation explains something, we say something that is not true. If we say that the equation states the statistical relationship between h and t, we do not say anything wrong, but we do not provide any explanation, either.

Oker-Blom's equation (1) does not have great explanatory power. A question

> 'Why does the equation hold for the height growth and the age?' (2)

to which a human expert could easily provide an answer, cannot be answered by the model.

Although knowledge present in Oker-Blom's model can explicitly determine any model, it does not even implicitly answer the question (2). The fact that the model defined by Oker-Blom et al (1988) does not include the knowledge required for an adequate explanation does not mean that an explanation does not exist at all. The problem of inexistence of explanations in statistical relations has been noted earlier (Fetzer, 1981, 87; Salmon, 1984).

It is actually easy to find examples of cases in which obvious statistical relations have nothing to do with explanations or actual mechanisms of the case. We can, however, take Salmon's (1984, 268) example:

> '...there is a strict positive correlation between the amount of time required for clothes hung out on a line to dry and the distance required to get an airplane off the ground at the nearby airport. I take it as given

> that the higher the relative humidity, the longer the clothes will take to dry. Thus the phenomenon that requires explanation is the fact that increased relative humidity tends to make for greater takeoff distance, other things — such as the temperature, the altitude of the airport, the type of the plane and the load it is carrying — being equal...'

In the example chosen by Salmon there is no direct explanatory connection between the observed phenomena. The actual explanation gives the common reason for the phenomena.

Although statistical analysis of shallow knowledge are often very useful, the unfortunate fact is that sometimes even strict correlations or conditional probabilities of 1 or 0 do not mean that there is an explanatory relation. It can, at first, seem strange to say that $\mathrm{P}(A|B) = 1$ or $\mathrm{P}(A|B) = 0$ does not mean that there is any connection between A and B, but it is fairly easy to find examples which show that the claim is true. A great deal of this is implied by the fact that statistical correlation is a symmetric relation while causality is not (Irtzik and Meyer, 1987).

Considering Oker-Blom's equation and Salmon's airport-example and proper explanations for those cases, it is interesting to note that such explanations have to explain the observed phenomena by using concepts and entities which are not parts of the original description of the case. The same can be said about any physical phenomenon we are explaining by using the Newtonian mechanics, which surely are not present in the description of e.g. the famous apple often said to be connected with the late Newton. It seems to be the case that those explanations use deep knowledge, i.e. knowledge that goes beyond the observed surface features (i.e. the values of the actually interesting variables) of the case and the original description of the situation.

Problems with shallow models lead us to look for other approaches. As it was said, one of the most important aspect of explanation (and knowledge) is to answer why-questions, which often imply some kind of causality. Unfortunately the concepts of cause and causality are very difficult to define, but we will, however, briefly outline causality and explanation. We will continue according to the assumption that causal relations exist.

4 Causality and explanation

The use of causality in different forms is actually very tempting. The concept of causal relations may be very explanatory and cognitively plausible. Actually there seem to exist different levels of explanatory and probability in causal explanations (see e.g. Sober (1984)), which implies the versatile nature of causal explanation.

Probably the most useful approach to define the features of the physical reality spoken about is proposed by Salmon (1984), who proposes the concept of causal processes instead of causal events. Of causal processes he says that

> 'Causal processes propagate the structure of physical world and provide the connections among the happenings in the various parts of space time. Causal interactions produce the structure and modifications of structure that we find in the patterns exhibited by the physical world. Causal laws govern the causal processes and causal interactions, providing the regularities that characterize the evolution of causal processes and the modifications that result from causal interactions' (Salmon, 1984, 132)

The concept of causal processes is not, however, very clear; many features of them have to be defined with more precision. Those features include mark transmission, structure transmission, principle of causal influence, and causal interaction.

The most important criterion for a causal process is its ability to transmit a mark. By mark transmission it is meant that a causal process can transmit a mark from point A to point B (and every point between them) without further interactions. The *mark transmission* (MT) is defined, in a more explicit way (Salmon, 1984, 148) , as:

> Let P be a process that, in the absence of interactions with other processes, would remain uniform with respect to a characteristic Q, which it would manifest consistently over an interval that includes both the space-time points A and B ($A \neq B$). Then, a mark (consisting of a modification of Q into Q'), which has been introduced into process P by means of a single interaction at point A, is transmitted to point B if P manifests the modification Q' at B and all stages of the process between A and B without additional interventions.

Note that marks in a process are, actually, changes in the process itself. Therefore we can say that a transmission of a mark is a transmission of the changed structure of the process transmitting the mark, and a process can always be said to transmit its own structure, changed by a mark or not. The principle of *structure transmission* (ST) can be formulated as follows (Salmon, 1984, 154):

> If a process is capable of transmitting changes in structure due to marking interactions, then that process can be said to transmit its own structure.

The fact that a process does not transmit a particular type of mark does not mean that it is not a causal process. Consider, as an example, processes of a hard rubber ball and a particular beam of light (caused by a lamp and colored white). It is possible to paint a green mark on the surface of the ball but it is not possible to do the same to the beam of light, although it is possible to change the color of the beam to green by using a green filter. Marks must be consistent with the structure and properties of causal processes — a causal process cannot be marked by every method.

In accordance with the principle of structure transmission there must be a way to define how a causal process propagates causality from one space-time locale to another. The *principle of causal influence* (PCI) can be defined as (Salmon, 1984, 155):

> A process that transmits its own structure is capable of propagating a causal influence from one space-time locale to another.

Combined together the concepts MT, ST and PCI define what a causal process is. Although a causal process can be very effectively defined by them, no interactions between processes have been defined. Obviously processes interact and interactions constitutes the actual structure of causal relations.

As it can be deduced from the definitions of MT, ST and PCI, there exist, in addition to causal processes, a great deal of non-causal processes. The existence of non-causal processes makes the definition of causal interaction (CI) between processes quite difficult task, and one proposition is made by Salmon (1984, 171), and is as follows:

> Let P_1 and P_2 be two processes that intersect with one another at the space-time point S, which belongs to the histories of both. Let Q be a characteristic that process P_1 would exhibit throughout an interval (which includes subintervals on both sides of S in the history of P_1) if the intersection with P_2 did not occur; let R be a characteristic that process P_2 would exhibit throughout an interval (which includes subintervals on both sides of S in the history of P_2) if the intersection with P_1 did not occur. Then the intersection of P_1 and P_2 at S constitutes a causal interaction if:
>
> 1. P_1 exhibits the characteristic Q before S, but it exhibits a modified characteristic Q' throughout an interval immediately following S; and
> 2. P_2 exhibits the characteristic R before S, but it exhibits a modified characteristic R' throughout an interval immediately following S.

Salmon proposed that scientific explanations should use the concept of causal processes as the basis, and that actual explanations are linguistic descriptions of chains of causal interactions and mark transmissions. In the following chapter we will consider a modeling approach that is based on the concept of causal processes and briefly discuss the intuitively tempting claim that models able to provide explanatory knowledge are robust.

5 Modeling with explanatory concepts

In this section we will describe a domain modeling approach based on the Salmonian concepts of causality. Generally we will consider models that may be called explanatory models. An *explanatory model* is constructed according to the concept of causal processes and its methods to produce quantitative results, i.e. values to variables, are based on processing the descriptions of the modeled causal processes. Note that we are explicitly considering quantitative modeling, not qualitative in the meaning discussed in (Weld and de Kleer, 1990).

In order to prevent any confusion between explanatory inferences (i.e. the mechanism by which explanations are generated) and explanatory models (i.e. models of the causal structure of the reality) we will consider the sentence 'I see that there is snow outside and therefore I know that the temperature outside is under zero degrees'. The sentence is not an explanatory model, and even the sentence 'I see that there is snow outside and therefore I know that because water can be snow only if the temperature is below zero, I am able to deduce that the temperature outside is below zero' is not an explanatory model. Writing that 'When the temperature is below zero° C, water molecules arrange into a combination in which they do not constitute a liquid' we are much nearer to have an explanatory model. Note that a transmission of energy causes water molecules to rearrange – an obvious case of causal interaction and mark transmission.

5.1 The structure of explanatory models

An interesting feature of causal processes is the possibility that causal processes P_1 and P_2 can be in a continuous interaction. Consider, as an example of continuous causal interactions, a car. The motor of the car has always several different causal interactions with other parts of the car, and there are, in the motor, parts which can be thought to be causal processes continuously interacting with each others. But only certain characteristics or features of the motor can be modeled. Although some knowledge could have been omitted due to some restrictions, as much similarity as possible must be maintained in order to achieve the robustness requirement.

Another interesting feature of explanatory models is that they are divisible. Consider, again, a car. For a pedestrian injured by the car, the car constitutes one process, and for a mechanic the car consists of a multitude of different causal processes. The example of the car shows that we can divide a process into several processes depending on what we know of it. A causal process can sometimes be divided to several subprocesses, which have causal interactions and causal structures of their own.

Because some processes can be divided into subprocesses, a process should be clearly distinguishable from its subprocesses. One possible method for the separation is the introduction of the concept of levels; on lower levels there would be the subprocesses of process P_1 and P_1 itself would be on the upper level.

The number of levels of an explanatory model is not restricted in principle. Different models require different numbers of levels, and the number of levels of a model is restricted by the limitations dictated by the computer on which the model is used or by the fact that sometimes the required accuracy can be achieved by using fewer levels, or by the fact that human knowledge of the problem may not be complete enough.

In order to define the characteristics of the processes there must be some means to define the characteristics of any particular process. Therefore it is reasonable to consider the characteristic of a process to consist of the intensional name of the process (the name represents the intension of the intended real counterpart existing or potentially existing in the physical reality), and the name can be considered to mark the variables which constitute the features according which the process is identified and characterized.

A process is characterized by the values of the variables and the variables themselves which are used to represent the process in question. The name of the process is used to tell which variables with certain values are used to represent a certain ontological process, and names may be sometimes used to clarify the recognizable features of a process.

It must be required that every process of a model is distinguishable from other processes of the model. To achieve that we have to require that every process on any level must be characterized by at least one name all the time. This goal can be achieved by using special names. It is possible to name a causal processes with a special name like 'Fido'. This can be done by defining the characteristics of the process 'Fido' in a way that makes the extension of 'Fido' to always be the same process, and that the name 'Fido' always has an extension.

In some cases the causal interaction is constant, but sometimes it is only a possibility. The possibility of the causal interaction exists between any cup and any table, but it actually exists only if the causal processes in question fulfill certain characteristics (one example of that kind of characteristics is the space-time locations of the processes).

Obviously the movement of a car, which is a causal process, has an effect on the doors of the car — at least the effect of movement. Or consider a stone hitting one of the doors. Obviously that causal processes on different levels can have causal interactions with each others. It seems actually to be the case that there is no limit for the level number difference between processes which have causal interactions with each others.

It can be said that general interactions are intentions of interactions and actual interactions are their extensions. The general definition of an interaction could be thought to be a demon which causes a real interaction to happen when certain requirements are satisfied.

5.2 Remarks on explanatory models

In the subsection above we defined the structure of explanatory models and some other concepts. In this subsection we will consider some special semantic and syntactic features of our explanatory model structure, and try to clarify its intended interpretation.

In the explanatory model structure, a causal process is a set of variables. The set is named by a term, and the naming relation defines the characteristics of the process connected to the term. The mapping from the name to the set of variables defines the characteristics of a process. The mapping represents the intension of the name, and the set of the variables to which the mapping maps the name represents the extension of the name. The characteristics of a process often refers to the values of the variables representing the process, not only to the variables themselves.

One result of this characterization is that causal interactions and characterizations themselves can make the process to fall under different characterizations at different times. Consider, as an example, a steel plate. If the plate is crushed into another form, it is not a steel plate any more although it is still steel and still the same process. Actual processes do not disappear although they can change into forms which could be very different when compared to the original forms.

Physical relations are not static. A part of a machine can be removed, or a physical entity can be broken into several pieces. Causal interactions can change the physical relationships between processes.

Two models of two different chairs have very much common despite the fact that the chairs can be very different in some respects. A model of any physical entity can include knowledge of atoms and particles, and knowledge of atoms and particles is very similar in any case.

The relation between knowledge used in two models could be said to mean that causal processes on deeper levels are more and more alike conceptually, although physically different. Two processes may have very similar structure and definition. If models T_1 and T_2 are both models of pines, then they have very similar knowledge on their lower levels. In other words, knowledge used in a model is more common to other models on lower levels than it is on upper levels.

There is often no need to compute the actual causal interaction on every level of the model. The actual amount of causal interactions required to be computed depends on the model and the use of the model. Consider, for example, the familiar fact that a metal table will support a coffee cup. It is clear that the explanation of this fact involves the details of atomic structure and the appeal to conceptions as unfamiliar to everyday experience as the Pauli exclusion principle (the example is taken from (Salmon, 1984)). An explanatory model of the situation can be created, but a precise model would include the atomic structure, and the workings of the atomic structure would appear to be extremely difficult to be computed.

In order not to give the impression that the only useful type of explanatory models includes knowledge of even subatomic relations, we have to point out a quite obvious way to determine where to stop modeling. In every case we start modeling by using one or another scientific theory according to which we model. During the deepening of the model we will always end with a level which cannot be defined without using a scientific theory that is different than the theory according to which we started modeling. Such a change of theories provides a suitable point in which we may stop modelling and still be able to provide required robustness. If we are modeling a mechanical device, e.g. a clock, we may be able to stop modeling when we have reached a level on which we have exhausted the means of Newtonian mechanics. Explaining that level would require entirely another kind of theory, and in many cases such theories are not necessary in order to achieve the required robustness or explanatory power. Note, however, that the number of scientific theories used in a model depends on the pragmatic aspect of the modeling project, and those aspects depend on the case.

The computational problem can be avoided by hiding the atomic structure of the model from normal computations. Such hiding could be done by using the method described in (Ahonen, 1991). If computations are needed, the causal processes and interactions needed could be created when required and disposed after the need. The hiding of the more precise structure of a model in order to avoid computational problems seems to promise a solution to the question of computational effectiveness of causal models, see e.g. (D'Ambrosio et al, 1985).

5.3 The robustness of explanatory models

Although we would like to show that the structure of explanatory models do provide robustness, we have to admit that it is not easy. The best we are able to achieve is to discuss some features of explanatory models and the intuitively robust nature of those features.

A model should be able to function properly in new situations. From the nature of models it is clear that the robustness is actually provided by the knowledge already present in the model, and the same seems to be the case with the robustness present in estimations made by human experts. This makes us to adopt a view very similar to Hacking's. He says that *'a man can know how to use modus ponens, can know that the rule is valid, can know p, and can know $p \rightarrow q$, and yet not know q, simply because he has not thought of putting them together'* (Hacking, 1967, 319).

Assuming that the robustness of models is due to knowledge which is used in a new way we can understand the role of deep knowledge better. If our assumption of the growing similarity of deep knowledge is true, then it is tempting to conclude that new combinations of such basic knowledge, i.e. deep knowledge, provide the robustness. Since very deep knowledge of one model is very alike to similarly deep knowledge of another model, it may be the case that such knowledge can be easily arranged to combinations that provide the required robustness to the upper levels of the model.

We have considered the reasons of the robustness of models that use deep knowledge, but we have not paid any attention to the robustness of our explanatory models. We claim that explanatory models may provide a uniform approach to domain knowledge, explanatory power and robustness, because explanatory models employ only one concept — the concept of causal processes. The concept of causal processes enables the modeler to use uniform representations and an uniform way of though throughout the development of the model and maintain the explanatory power of scientific knowledge. Causal processes provide a method by which more or less deep knowledge may have similar conceptualization and representation undepending its deepness.

The concept of causal processes has originally been developed to provide a uniform approach to the study of scientific knowledge and scientific explanation, and there seems to be no reason to assume that the same approach could not be used in domain modeling. The causal process approach with a suitable formalization may provide a very promising alternative by defining what kind of knowledge should be used and what kind of interactions models include. Considering the nature of causal processes it is reasonable to believe that representations of processes enable very different and unanticipated causal interactions and mark transmissions to occur.

6 Discussion

It is plausible that the cognitive models used by humans are very robust due to the human ability to connect different items of knowledge in new ways. Assuming that it is easy to conclude that robustness of domain models may be based on the same method. The obvious way to introduce robustness to domain models seems, in that case, to depend on the conceptualization and representation of domain knowledge. Such conceptualization should provide versatile and uniform enough approach to domain knowledge and its representation.

The concept of causal processes provides a uniform view to domain knowledge, and it is not very difficult to develop a framework for modeling of causal processes and their features. This should enable us to follow a clear path in our modeling activities, and it is reasonable to assume that models with enough levels and accurate descriptions of processes and the features of processes are robust. Unfortunately the case is not obvious, although most of the writers that have discussed deep knowledge have adopted similar point of view to modeling with deep knowledge.

The uncertainty of the provided robustness is due to that we do not actually know how human experts generate new combinations of knowledge. If it is done by methods that work like normal logic or straightforward simulation, then there should not be any problems. But if the only method is induction, our attempts to develop models robust enough to be comparable to cognitive models used by human experts may be void.

Another hazard with the use of the concept of causal processes may be that the concept of causal processes seems to be extreamly difficult or even impossible to be formalized with descriptive methods. It may be possible that the nonlogical modalities and counterfactuals present in the definition of causal processes prevent exact formalization (Fetzer, 1987). This problem does not, however, make the use of the concept of causal processes impossible, it may only complicate it by giving more responsibility to the modeler.

One additional benefit of the concept of causal processes is that if quantitative models were developed according the approach, it could be possible to provide a natural connection between quantitative and qualitative models. If both quantitative and qualitative models were based on fairly similar principles (the concept of causal processes and the theory of qualitative processes, respectively), it should be much easier to provide the connection requested by Clancey (1992) and Kuipers (1993). That is, however, beyond the scope of this paper.

References

Adams, T. L. (1986). Model-Based Reasoning for Automated Fault Diagnosis and Recovery Planning in Space Power Systems. In *Proc. 21st Intersociety Energy Conversion Engineering Conference*. San Diego.

Ahonen, J. J. and Saarenmaa, H. (1991). Model-based reasoning about natural ecosystems: An algorithm to reduce the computational burden associated with simulating multiple biological agents. In *Proc. 6th Symposium on Computer Science for Environmental Protection*. 193-200.

Chandrasekaran, B. and Mittal, S. (1983). Deep Versus Compiled Knowledge Approaches to Diagnostic Problem Solving, *International Journal of Man-Machine Studies*, 22. 425-436.

Clancey, W. J. (1992). Model Construction Operators. *Artificial Intelligence*, 53. 1-115.

Cross, S. E. (1984). Model-based Reasoning in Expert Systems: An application to enroute air traffic control. In *Proc. AIAA/IEEE 6th Digital Avionics Systems Conference*. Baltimore.

D'Ambrosio, B., Fehling, M. and Adams, T. (1985). Levels of Abstraction in Model Based Reasoning. In *Proc. ROBEXS '85 (Robotics and Expert Systems)*. NASA. Johnson Space Center.

de Kleer, J. and Williams, B. C. (1987) Diagnosing Multiple Faults, *Artificial Intelligence*, 32. 97-130.

Fetzer, J. H. (1987). Critical Notice: Wesley Salmon's Scientific Explanation and the Causal Structure of the World, *Philosophy of Science*, 54. 597-610.

Fetzer, J. H. (1981). *Scientific Knowledge*. D Reidel Publishing Company. Boston.

Futo, I. and Gergely, T. (1990). *Artificial Intelligence in Simulation*. Ellis Horwood.

Chishester.

Hacking, I. (1967). Slightly More Realistic Personal Probability, *Philosophy of Science*, 34. 311-325.

Hayes, P. J. (1979). The Second Naive Physics Manifesto. In Michie, D. (ed) *Expert Systems in the Micro-Electronic Age.* Edinburgh University Press. Edinburgh.

Irtzik, G. and Meyer, E. (1987). Causal modeling: New Directions for Statistical Explanation, *Philosophy of Science*, 54. 495-514.

Koton, P. A. (1985). Empirical and Model-Based Reasoning in Expert Systems. In *IJCAI-85 Proceedings.* Los Angeles.

Kuipers, B. J. (1993). Qualitative simulation then and now. *Artificial Intelligence*, 59. 133-140.

Lewis, T. G. and Smith, B. J. (1979). *Computer Principles of modeling and Simulation.* Houghton Mifflin. Boston.

Nardi, B. A. and Simons, R. K. (1986). *Model-Based Reasoning and AI Problem Solving.* IntelliCorp Technical Article Reprint. IntelliCorp.

Niiniluoto, I. (1983). *Scientific Reasoning and Explanation* (in Finnish). Otava. Keuruu.

Oker-Blom, P., Kellomaki, S., Valtonen, E. and Vaisanen, H. (1988). Structural development of Pinus sylvestris Stands with Varying Initial Density: a Simulation Model, *Scandinavian Journal of Forest Research*, 3. 185-200.

Rich, S. H. and Venkatasubramanian, V. (1987). Model-Based Reasoning in Diagnostic Expert Systems for Chemical Process Plants, *Computers and Chemical Engineering*, 11. 111-122.

Saarenmaa, H. (1988). Model-Based Reasoning in Ecology and Natural Resource Management. In *Proc. Resource Technology 88: International Symposium on Advanced Technology in Natural Resource Management.* Fort Collins (Colorado).

Salmon, W. C. (ed) (1971). *Statistical Explanation and Statistical Relevance.* University of Pittsburg Press. Pittsburg.

Salmon, W. C. (1984). *Scientific Explanation and the Causal Structure of the World.*

Princeton University Press. Princeton.

Sober, E. (1984). Common Cause Explanation, *Philosophy of Science*, 51. 212-241.

Sweller, J. (1988). Cognitive Load During Problem Solving: Effects on Learning, *Cognitive Science* , 12. 257-285.

Weld, D. S. and de Kleer, J. (eds) (1990). *Qualitative Reasoning about Physical Systems*. Morgan Kaufmann Publishers, San Mateo, California.

Xiang, Z. and Srihari, S. N. (1986). Diagnosis Using Multi-level Reasoning. In *Proc. Expert Systems in Government Symposium*, McLean (VA USA).