Presented at: CSAPC’99, Villeneuve d’Ascq, France 21-24 September 1999.

MAKING SENSE OF THE ABSTRACTION HIERARCHY

Morten Lind
Department of Automation
Technical University of Denmark
DK-2800 Lyngby Denmark
Email: ml@iau.dtu.dk

ABSTRACT

The paper discuss the abstraction hierarchy proposed by Rasmussen(1986) for design of human machine interfaces for supervisory control. The abstraction hierarchy represent the domain of human work by multiple levels of means-end and part-whole abstractions. It is argued in the paper that the abstraction hierarchy suffer from both methodological and conceptual problems. The problems are illustrated by concrete examples from the power plant domain. It is concluded that the semantics of the means-end levels and their relations should be improved by making more distinctions. Furthermore, the commitment to a fixed number of levels of means-end abstractions should be abandoned and more attention given to the problem of level identification in the modeling building process.

Keywords: Cognitive Engineering, Human-Machine Interfaces, Abstraction Hierarchies, Modeling

INTRODUCTION

The use of means-end and part-whole abstractions in plant representations for human supervisory control has been the subject of both research and application for more than a decade. Rasmussen (1986) promoted the idea in the form of the abstraction hierarchy (AH). The AH is part of a cognitive engineering approach to human-machine systems design and has been adopted mainly by the nuclear power industry, presumably because of its appeal to system engineers designing display systems for plant supervision. Furthermore, the AH is supported by empirical studies of operators fault finding strategies, and offer therefore, together with the overarching cognitive engineering framework, a promising basis for design of supervisory displays. The AH was used by American nuclear power industries after the Three Mile Island incident in their efforts to improve the reliability of the human-machine interaction and was later adopted by the Japanese nuclear industry in the conceptual development of a new generation of control rooms.

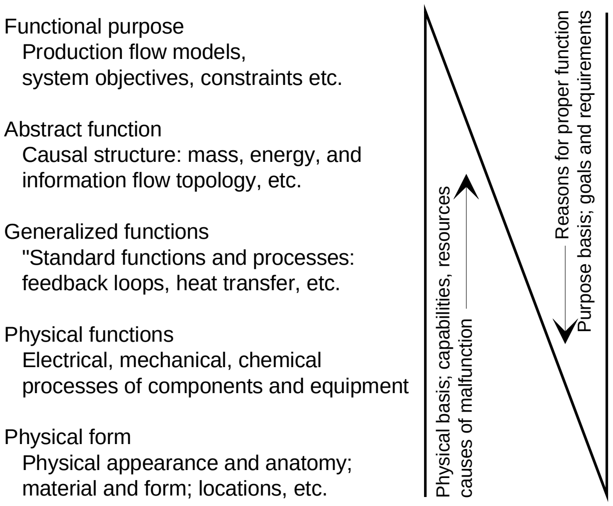

Fig.1. The abstraction hierarchy (Rasmussen, 1986)

In parallel with and partly motivated by the industrial interest, there has also been efforts within academia to develop the foundations and applications of means-end and part-whole abstractions. An example of academic research is the application of the AH for design of so-called ecological interfaces by Vicente and coworkers (Bisanz & Vicente, 1994). One of the aims of this research is to use the AH for the design of human-machine interfaces. This group has also demonstrated the application of the AH through concrete examples. Another example of academic research is the development of Multilevel Flow Modelling by Lind(1994) and his research group. The main objectives of MFM research are to develop concepts and methods for modeling of complex industrial artifacts and to use the models in conceptual design of industrial automation systems, including intelligent controls and supervisory functions for the operator (Lind, 1996). At the early stages of development MFM was presented as an articulation of the AH. However, this tight coupling of MFM to the AH proved later to be a hindrance to its development. An analysis of these restrictions led to the recognizion of the cluster of AH problems discussed below. Current MFM research (Lind, 1999) is addressing these problems and MFM is therefore no longer considered to be an integrated part of or a formalisation of the AH.

The industrials and the academics seems accordingly to share the conviction that the AH or its underlying idea of combining means-end and part-whole abstractions makes sense. In spite of these beliefs, it has actually proven to be difficult to apply the AH. However, systematic analyses of these difficulties are lacking. It is therefore not yet clear whether the very idea of using means-end and part-whole abstractions for modeling complex plants is flawed or the problems are specific to the AH. In the present paper we will identify a cluster of problems that the author has found to be sources of difficulties when building plant models using the AH.

PROBLEMS WITH THE ABSTRACTION HIERARCHY

The main principle of the AH is to describe a system on several levels of abstraction. Two types of abstraction are at work simultaneously. The *means-end* abstraction as shown in figure 1 decribe how physical resources and system functions can be organised into five levels so that each level define the means for the next upper level and define ends that are accomplished using items on the level below as the means. *The part-whole* abstractions (not in Fig. 1) decompose or aggregate items on each level of means-end abstraction. The original formulation of the AH by Rasmussen(1986) was based on studies of fault finding in electronic workshops and supervisory control of power plants. The principles of the AH were also applied by Rasmussen, Pejtersen & Goodstein(1994) to other non engineering domains leading to a generalized version. The present analysis of the AH is mainly based on the early version and shows that the AH suffer from problems even when used to represent engineering systems like power plants.

A general problem with the AH is actually its immediate appeal to common sense i.e. the fact that engineers knowledgeable especially in power systems control seems to recognize the idea of the AH without great difficulty, hence the success within the nuclear domain. This success could of course be simply explained by the fact that the AH originally was developed in the context of supervisory control of power plants. However, we also believe that the AH, more or less sucessfully, reflect tacit knowledge about engineering practices in the process control domain that also has some validity outside the narrow context of power plants. The problem is however, that the AH is lacking the expressive power to make this knowledge explicit. The common sense appeal of the AH is therefore a problem. It could lead some to the conclusion that the AH is without problems and ready to use. The analysis below shows that such a conclusion is quite premature.

The cluster of specific AH problems discussed below has been divided into two major groups; methodological and conceptual problems. The problems in the two groups are partly interdependent and the division is mainly done for convenience.

Methodological Problems

A user trying to use the AH for a particular modelling problem will be met with problems of methodological nature. One problem is that no procedure or guideline exist (to the authors knowledge) that can help in the acquisition of the background knowledge required for a modelling task. Another problem is that there is no process for model building or for revising, modifying and validating a model. These deficiencies are amplified by the circumstance, that the meanings of the different levels of means-end and part-whole abstraction in the AH are only defined in terms of prototypical incomplete examples from a few selected domains. It is therefore very difficult to apply the AH to a new or sligthly different modelling problem. A related problem is the lack of convincing arguments for the number of means-end abstraction levels. This problem will be discussed in more detail after the analysis of the conceptual problems.

The monolitical nature of the AH is another source of methodological problems. Thus, it is not clear how to combine a multiple of AH instances each representing a contextual unit of purposeful activities or how to break down a complex system in such units. A paradigm example illustrating the nature of this problem is the representation of the couplings of a production process and its control system. Production processes are physical transformations of materials or energy whereas control systems are information processing systems. It is suggested that the AH can be used to model both types of system but it is not clear how to express the fact that the information processed by the control system are representations of the physical phenomena in the production system. This coupling is established through the plant instrumentation but cannot be dealt with in an obvious way with the AH. A related problem, handled by MFM, is the representation of the couplings between a system and its supporting subsystems. These problems, indicate that the AH offer no distinctions between system categories and lacks a concept of system boundary. The problems involved in the modeling of the coupling of the plant and its control system through the instrumentation are discussed by Lind(1993) in the context of MFM.

Conceptual Problems

The hetereogeneous repository of concepts used to characterize the content of the five means-end levels is a source of confusion. According to the definitions (Rasmussen et. al.,1994), the levels includes activities, things, information, money, people, properties, constraints and prority measures. These concepts clearly refer to a variety of ideas and phenomena each belonging to different contexts of analysis, and cannot be combined without more fine grained distinctions.

Here we will mainly discuss conceptual problems that are caused by the generality of the means-end and the part–whole relations. This generality makes the AH open to interpretation, a quality that may seem as an advantage. It is actually a serious weakness because the unclear semantics disguises the real nature of the modelling problem and is a hindrance to model validation. The issues discussed in the following are; the five levels of means and ends, the concept of function, the means-end relations and the concepts of whole and parts.

The Five Levels of Means and Ends

The AH make distinctions between five different levels of means and ends. The levels of physical form and physical functions quite clearly refer to the spatial extension and behavioural characteristics of physical objects that in production plants have natural interpretations like pumps and heat exchangers. However, it is unclear whether the definition of these levels leave room for other types of entities like water and heat that according to common intuition also would count as physical means in a production plant. However, flows of water (mass) and energy are both included on the level of abstract function and not on the level of physical function. Actually, the concept of heat flow is necessary in order to describe the behaviour of a heat exchanger represented on the level of physical function and water is often used as as a means for storing energy. Of the same reason, it is not obvious how the AH handle modeling problems where a substance is both an agent and an object for action at the same time. Take for example a cooling loop. The cooling pump (agent) is transporting the water (object) and the water (agent) is transporting the energy(object). This example has three means-end levels but they do not fit naturally into the AH. The AH cannot handle this problem because substances seems to be on the level of abstract function and because it has no distinctions between agents and objects.

The inclusion of actions on the level of physical function in the AH (Rasmussen et. al., 1994) is also problematic. Most people would regard actions as genuine means (consider e.g. the following sentence ”the turning of the valve by 30 degrees is a means to increase the flow of water”) but actions does not to fit naturally in the same category as material objects like pumps and valves. Actions are events that take place in space and include objects and agents with physical extension but they have themselves no physical form or configuration. The relation implied in the AH between physical function and physical form is accordingly only valid for material objects. Furthermore, the intrinsic logical (i.e. not causal) relations between an action and its attributes (agent, object, means, manner, cause, result, intention etc.) considered by theories of action (Rescher, 1966) or verb semantics (Fillmore, 1968) cannot be represented fully in the AH.

Other problems occur on the level of generalized functions. The entities mentioned on his level are not physical material objects but refer to desirable outcomes of the dynamic interactions in the system and are therefore more of a process nature (e.g. cooling and control). According to the AH these processes are implemented by means of the entities described on the level of physical function. This follows from the logic of the means end relation connecting the levels of physical and generic function. Since the existence of processes then are assumed to be dependent on the existence of physical functions that again are assumed to have a spatial extension, it is seen that physical functions are considered ontologically primary to generalized functions. This implicit commitment to a thing’s ontology can create conceptual problems when modeling phenomena where it is more convenient to consider processes as primary entities i.e. to adopt a process ontology (Rescher, 1996). An illustrating example from the power plant domain is the problem of representing the burning of fuel in a boiler. The combustion of fuel is a process that interacts with the air gas mixture flowing through the boiler. But this process can also be described as a flame i.e. as an entity with a physical location and form. It should accordingly be represented in the AH at the levels of physical form and function. However, even if the flame is a means for heating the boiler it is not meaningful to describe it only as a material object or thing because its existence is conditional on the burning process. It is accordingly necessary in some way to allow both material objects and processes to be at the bottom of the hierarchy. It is not clear how this should be done in the AH. A solution could be to extend the hierarchy at the bottom with a new level. But there are no principles that prescribe how such an extension should be done in a systematic way.

The Concept of Function

The distinctions made in the AH between different types of function (physical, generalized and abstract) are troublesome because the concept of function has several meanings. Proper treatment of this problem will not be attempted here because it involves a complex analysis of a whole cluster of related concepts (goal, objective, function and disposition and action). We will only point out that the AH seems to confuse functions that are ascribed to entities by convention and function that are grounded in a dispositional property of the entity in question. As an example of an entity that is ascribed a function by convention could be mentioned a coin whose function is to exchange value. An example of an entity whose function is based on a dispositional property could be mentioned a pump. The ability of a pump to serve its function (to move water) is only dependent on its physical conditions and independent on conventions. The AH invite to a confusion of the two meanings because money flow in (Rasmussen et. al. , 1994) is mentioned as belonging to the level of abstract function.

This distinction between the ascription of function by convention and by disposition is not only important for understanding the differences between social and engineering systems. In fact both types of functional ascription are required in order to model the means-end relations in e.g. a power plant. As an example consider a coolant system. The function of the water circulation is here to provide cooling. The ascription of this function to the system explains why the system is there (its teleology) and is valid because the system is able to provide the function (it has the appropriate dispositions). However, if we want to explain why the water is recirculated it should be described as part of a context of value exchange where its function or role is to be a valuable object. This functional ascription is based on social conventions.

The Means-End Relations

In addition to the problems of the five levels of means and ends in the AH there are also problems with the meaning of the relations connecting them. Thus, it is not obvious that the relations between two adjacent levels in the AH are of the same nature. For example; is the relation between purposes and abstract functions of the same nature as the relations between physical functions and generic functions? Furthermore, these relations cannot be the same as the relation between physical form and physical function. Finally, it is not explained how an item on a given level is related to another item on the same level of abstraction. For example; how are two physical functions related and are their relations of the same type as the relations between two abstract functions? All these questions can only be answered by a refinement of the semantics of the means-end relation and by understanding the nature of the levels themselves.

The vague semantics of the means-end relations can in some cases lead to descriptions that are cyclic and therefore contradicting the notion of a hierarchy which is an acyclic structure. Consider for example the circulation pump in a cooling loop. The pumping (physical function) is a means of circulating the water (generalized function) that again is means of transporting energy from a heat source to a sink (abstract function). However, the pump can only circulate the water if it is not boiling and the water temperature depends on the heat transfer rate. This means that the energy transport is a means of ensuring proper conditions for the pumping. The mapping of this chain of means-end arguments into the AH requires a link between the level of abstract function (the means) and the level of physical function (the end) that seems to contradict the directions of causality (and intentionality) in the AH. The problem with this example can be resolved by clarifying the semantics of the means-end relation through a distinction between sufficient (causation) and necessary (enabling) conditionings between the means and the end. The AH does not make that distinction and therefore cannot make sense of the example. Even more problematic is the modeling of a coolant system with natural circulation where the movement of the water is caused by temperature gradients between the heat source and the sink. In this case it is questionable whether it is useful to talk about levels at all!

The Concepts of Wholes and Parts

Another cluster of problems concerns the interpretation of the levels of wholes and parts. The scope of these problems will not be explored here but only illustrated by referring to the work of Nagel (1961). Nagel analyzed different meanings of these concepts and came up with the following eight possible interpretations. A whole can be 1) something with a spatial extension and anything is then called a part of such a whole that is spatially included in it, 2) a temporal period, whose parts are temporal intervals in it, 3) any class, set or aggregate of elements, and part may then designate either any proper subclass of the initial set or any element in the set, 4) a property of an object or process, and part to some analogous property that stands to the first in certain specified relations; 5) a pattern of relations between certain specified kinds of objects or events, 6) a process, one of its parts being another process that is some discriminated phase of the more inclusive one, 7) any concrete object, and part to any of its properties, 8) any system whose spatial parts stand to each other in various relations of dynamic dependence.

It is quite clear that this variety of possible interpretations of the part-whole relation is a potential source of confusion when the actual meaning to be used in the AH is not well specified. Parts and wholes should obviously have different meanings on the five levels of means-end abstraction. On the levels of physical form and function the relation between wholes and parts is spatial inclusion (i.e. the first option in Nagel's analysis). However, on the levels of generalized and abstract functions the temporal meaning (i.e. the second option) seems more appropriate. Further confusion is introduced because one of the levels of part-whole abstraction in the AH is labeled the level of "function units" (Rasmussen, 1986, p. 119). This labeling confuses concepts of function belonging to the means-end distinction with concepts belonging to the part-whole decomposition.

DISCUSSION AND SYNTHESIS

The analysis above has shown that the AH suffer from both methodological and conceptual problems. The conceptual problems have been demonstrated through detailed examples from the power plant domain. In the following we will discusss the means-end concepts and specific problems of the AH in the broader contexts of decision making and modeling methodology. The following two questions will be addressed: 1) What are the advantages and uses of means-end concepts in modeling complex systems? and 2) How are levels of means and ends identified? The first question is adressed by Rasmussen but not completely answered and the second has not been raised. The discussions reveal that the AH in addition to the specific problems mentined bove also suffer from more basic problems both when seen in an application perspective and as a modeling framework.

Advantages and uses of Means -End Concepts

Means-ends concepts play an essential role in theories of action and practical reasoning (VonWright, 1963) and has potential application in many domains. Rasmussen(1984) argue that the means-end representation instantiated in the AH provides a systematic framework for identifying and evaluating alternative courses of action and that the AH in this way can reduce the complexity of decision making in supervision and control of anomalous plant situations. The main feature of means-end representations (an thereby also the AH) is that they define opportunities for decision making and support an organized sequential decision process.

Rasmussen combines the AH with a model of decision making. This decision model provides in principle a decomposition of the supervisory control problem into three basic decision problems in supervisory control each coping with a separate aspect of uncertainty. The three problems are state identification, goal selection and action planning. In state identification the problem is to select between alternative interpretations of the system state caused by multiple, insufficient or uncertain observations. In goal selection the problem is to choose or make compromises between multiple contradicting goals. The problem in action planning is to select among alternative courses of action. The three problems cope with three distinct resources (and preferences) of the decisionmaker – the set of possible observations and perceptions of the system state, the range of possible goals and the possibilities for action. The AH is proposed as a framework for solving the overall supervisory control problem, but it is not clear whether the AH can support all three subordinate decision problems. The AH seems to be most fit for the action planning problem since the entities described on the level of physical function and functional purpose represent means and ends of intervention. However, other representations seems to be required for state identification where the problem is to manage multiple and possibly conflicting means of observation and state interpretation. Rasmussen emphasizes the interdependence of the three decision problems in supervisory control and suggest that the causal and intentional reasoning processes supported by the AH provide a framework for their integration. However, even though the decision problems are interdependent they may require application of three separate means-end representations because the problem involves the management of different categories of means and ends.

Identification of Levels of Means and Ends

It is clear from the discussion above that the quality of decision making is directly dependent on the problem framing provided by the levels of means and ends i.e. the number and semantics of the levels and the means-ends relations. Improperly defined levels will lead to inefficient or directly wrong decisions. If the levels of description are too abstract decision alternatives may be overlooked. If the levels are too detailed they may include irrelevant alternatives and may lead to uncertain decisions. It is therefore essential that the methodology used to build the means-end representation has explicit principles, rules or guidelines for identification of levels and their semantics. Actually changes of representation of an artifact through redefinition of levels of descriptions may sometimes be a necessary step in a problem solving process. The problem framing provided by the levels of description may be a hindrance rather than a help in situations that call for the creation of new opportunities for decision making.

The AH framework does not present the number of levels and their semantics as subjects for decision in the modelling proces. The AH with the five levels of means-end abstraction seems to pretend that there is no problem of levels at all. A user will therefore tend to regard the number of levels of means-end abstraction in the AH as paradigmatic and may find it difficult to fit the levels to a given problem. The levels should rather be seen as resulting from a process of framing that may have different outcomes in different modeling situations. In other words, the number and contents of the levels and their relations cannot be defined in general but are features that characterize a particular modeling domain or even a specific problem. The problem is to understand how levels are constructed i.e. to know the criteria used in the selection of perspectives and particular aspects of a system or a situation and how these perspectives and aspects are combined into a representation that can be used to solve a given decision problem. Only through such an understanding can means-end concepts be used to construct models that makes sense.

ACKNOWLEDGEMENT

The present work has been supported by the Danish Foundation for Basic Research through a contract with the CHMI Center for Human-Machine Interaction.

REFERENCES

Bisantz, A. M. & Vicente, K. J. (1994). Making the Abstraction Hierarchy Concrete. *Int. J. Human-Computer Studies*, 40, 83-117.

Fillmore, C. (1968). The Case for Case. In: E. Bach and R. T. Harms (Eds), *Universals in Linguistic Theory*. New York: Holt, Rinehart and Winston.

Lind, M. (1993). *Functional Architectures for Systems Management and Control*. In Lind, M et. al., Interactive Planning for Integrated Supervision and Control in Complex Plant. Final report from CEC JRC project: 4937-92-08-ED ISP DK.

Lind, M. (1994). Modelling Goals and Functions of Complex Industrial Plant. *Journal of Applied Artificial Intelligence*, 8, 259-283.

Lind, M. (1996). Status and Challenges of Intelligent Plant Control. *Annual Review of Control*, 20, 22-31.

Lind, M. (1999). *Plant Modeling for Human Supervisory Control*. Submitted for publication.

Nagel, E. (1961). *The Structure of Science: Problems in the Logic of Scientific Explanation*. New York: Hartcourt, Brace & World, Inc.

Rasmussen, J. (1986). *Information Processing and Human-Machine Interaction*. New-York: North Holland.

Rasmussen, J., Pejtersen, A. M. & Goodstein, L. P. (1994). *Cognitive Systems Engineering*. New York: John Wiley.

Rescher, N. (1966). Aspects of Action. In N. Rescher (Ed), *The Logic of Decision and Action*. Pittsburgh: University of Pittsburgh Press.

Rescher, N. (1996). *Process Metaphysics – An Introduction to Process Philosophy*. New-York: State University of New York Press.

VonWright, G. H. (1963). *Norm and Action*. London: Routledge & Kegan Paul.