Believing and Acting

A Study of Meta-Knowledge and Meta-Reasoning †

Ernesto J. Morgado
Departamento de Engenharia Mecanica
Instituto Superior Tecnico
Av. Rovisco Pais
1000 Lisboa
Portugal

Stuart C. Shapiro
Department of Computer Science
University at Buffalo
226 Bell Hall
Amherst, NY 14260
U.S.A.

Abstract

In a knowledge representation (KR) system in which assertions and rules are represented in the same way as any other concepts, no special mechanism is needed to represent meta-knowledge, where this is understood to include beliefs about beliefs, rules about beliefs, beliefs about rules, and rules about rules. In a knowledge representation system which has an acting-planning component and which can represent actions and plans, no other mechanism is needed to handle meta-reasoning, where this is understood to include rules about the order of using rules, and reasoning about the process of reasoning. The difference between meta-knowledge and meta-reasoning as formulated above is that the former deals primarily with beliefs while the latter deals with acting. We therefore conclude that, besides the conceptual distinction between the object level and the meta level, a valuable distinction to focus on when building KR systems which can have meta-knowledge and can do meta-reasoning is that between believing and acting.

1. Introduction

We have been involved in the design of a semantic network processing system called SNePS [31]. Our first efforts were put into building the belief and reasoning components of the system.

---

† In *Proc. of EPIA 85* (*"Encontro Portugues de Inteligencia Artificial"*), Oporto, Portugal, September 1985, pp. 138-154.

The belief component was designed to allow a uniform representation of concepts. Each node of the semantic network represents a unique concept, and each concept represented in the network is represented by a unique node. Arcs represent non-conceptual binary relations between nodes.

The reasoning component of the system, SNIP (SNePS Inference Package), relies on a declarative representation of inference rules. The rules are stored in the semantic network like any other concept. SNIP may perform backward inference, forward inference, and bi-directional inference [24]. Every rule may be used in any of these reasoning strategies. When a rule is used, it is activated and remains activated until explicitly de-activated by the user. The set of activated rules are assembled into an active connection graph [25], which is a collection of MULTI processes [26] which carry out the inference. The processes are organized in a producer-consumer model in which, when a consumer wants a result that some producer is already working on, the consumer just adds itself to the list of consumers the producer reports to instead of creating a new producer. The producers also store their results, so that if during some later deduction SNIP needs some of the rules activated by some previous deduction, it uses the results directly instead of rederiving them. The net effect of this is that, until the active connection graph is de-activated, no inference step is done more than once.

The system is being used to build some expert systems, and the results are promising, but some limitations do exist. Two major limitations are: 1) SNIP is currently incapable of representing and using rules which affect the order of using rules -- an aspect of meta-reasoning; 2) SNIP is currently incapable of representing and carrying out general plans, because if a plan requires performing an action more than once, the active connection graph will ensure that it is carried out only once. When investigating ways to overcome these two limitations, we realized that they were really the same, and that that realization provides a general insight into the relations among believing, acting, knowledge, reasoning, meta-knowledge and meta-reasoning.

2. Meta-Knowledge and Meta-Reasoning

In order to accomplish the tasks expected of them, AI KR systems must be capable of a dynamic and flexible integration of strategies of reasoning, and must be capable of reasoning both about states and about activities, in particular about their own beliefs and activities.

Recently, several researchers have suggested the use of *meta-knowledge* and *meta-reasoning* to accomplish the integration of these features in a single system. In a general sense, meta-knowledge is knowledge about other knowledge as opposed to knowledge about "things in the world" [4]. It enables a reasoning system to "know what it knows," and to make multiple use of its knowledge [12]. In addition to using its knowledge directly, the system may have other abilities: knowing what it knows and what it doesn't know (consciousness) [4, 8, 9, 11, 12, 13, 20, 32]; knowing where and how to use knowledge to infer other knowledge (planning reasoning or meta-reasoning) [10, 12, 13, 16, 32, 34]; explaining how and why it used its knowledge (explanation) [10, 12, 13]; and examining its own knowledge, modifying it, abstracting and generalizing it, and acquiring new knowledge (learning) [9, 10, 11, 12].

3. Unifying Knowledge and Meta-Knowledge

Meta-knowledge, like object-knowledge, is composed of assertions (meta-assertions) and rules (meta-rules). Meta-assertions are beliefs about beliefs, and since a rule which is believed to hold is a belief, meta-assertions include beliefs about rules.

Semantic networks [15], and some similar KR systems, are designed according to the philosophy that every concept about which the system has a belief, as well as every belief the system holds, is represented by a node. A node which represents a belief "dominates" (has labeled, directed arcs pointing to) the nodes representing the concepts the belief is about. Such a system can represent meta-assertions as easily as object-assertions. For example, the belief that John loves Jane is an assertion represented by a node. The belief that Henry believes that John loves Jane is a meta-assertion represented by a node which dominates the node representing the proposition that John loves Jane. Similarly, the belief that Henry believes that all men are mortal is a meta-assertion represented by a node which dominates the node which represents the rule that all men are mortal. Other meta-assertions representable in such systems are the beliefs that: John loves Jane is a belief about John; all men are mortal is a rule about men; Bill doesn't know whether John loves Jane; I (the system) don't know about the fishing industry in Venezuela.

Rules tell how to derive beliefs from other beliefs. Since a rule which is believed to hold is a belief, we may have rules about rules as well. There are two types of meta-rules: deduction meta-rules and planning meta-rules.

Deduction meta-rules are rules which use rules to derive beliefs or which derive rules from beliefs. Since a rule is a belief, it is represented by a node. A rule node represents a propositional formula of molecular nodes, using one of a set of possible connectives. A node which represents a rule "dominates" the nodes representing the arguments and quantified variables of the connective. The system can represent deduction meta-rules as easily as deduction object-rules. For example, the rule A -> (B -> C) is a meta-rule which enables the system to derive the rule B -> C if the belief A holds. Similarly, the rule (A -> B) -> C is a meta-rule which enables the system to derive the belief C in case the rule A -> B holds. Both meta-rules are represented by a node which dominates the node representing the rule B -> C, appearing on the consequent and antecedent position of the meta-rule respectively.

In summation, a uniform representation of all kinds of beliefs unifies knowledge and meta-knowledge. No additional mechanism is needed to handle the distinctions "assertions versus meta-assertions" and "deduction rules versus deduction meta-rules", i.e., "knowledge versus meta-knowledge". The second type of meta-rules, planning meta-rules, are rules which encode reasoning strategies. The distinction between deduction rules and planning rules, i.e., between *reasoning* and *meta-reasoning* will be discussed in the next section.

4. Unifying Reasoning and Meta-Reasoning

Believing is a state of knowledge, representing the propositions that the system assumes to be true. Reasoning is the process of inference to form beliefs from other beliefs using deduction rules. Davis proposed the use of meta-rules as a means of encoding strategies for reasoning [10, 12, 13]. They specify which rules should be considered and in which order they should be invoked. For example, the two rules from [10] appearing in Figure 1 are of this type respectively.

These are what we call planning meta-rules. Planning meta-rules are the ones that have to be used differently from all the other rules (deduction object-rules and deduction meta-rules), since they do not express how to derive beliefs, but how to plan the reasoning process. They are inference rules that specify how the deduction rules should be used.

Meta-rule 1

If
(1) you are attempting to determine the best stock to invest in,
(2) the client's tax status is non-profit,
(3) there are rules which mention in their premise the income-tax bracket of the client,
then it is very likely (.9) that each of these rules is not going to be useful.

Meta-rule 2

If
(1) the age of the client is greater than 60,
(2) there are rules which mention in their premise blue-chip risk,
(3) there are rules which mention in their premise speculative risk,
then it is very likely (.8) that the former should be used before the latter.

Figure 1

Davis proposes a layered control structure to handle reasoning. The basic execution cycle in TEIRESIAS consists of selecting the inference strategy to use (backward inference, forward inference, etc.), and applying it to invoke all rules which are relevant to the goal. But before invoking the rules at one level, the system checks for rules at the next higher level which specify which rules should be selected and in what order should they be used.

This process can be seen as a particular case of a more general acting/planning process such as the one proposed by Sacerdoti [30]. Acting is the process of executing a plan. Any complex action has to be planned before being performed. Planning is the process of composing a sequence of actions to be executed to achieve a predetermined goal from a given situation; it is reasoning about how to act to achieve that goal. The basic planning cycle in NOAH consists of looking for a plan to achieve the goal and of an iterative process in which new refinements of the plan are continuously expanded and criticized until a final plan is derived. The expansion phase produces a new more detailed plan. The criticism of the new plan consists of any necessary reordering or elimination of redundant operations to ensure that the local expansions make global sense. After being constructed, a plan of actions may be executed.

Reasoning can be looked at as the sequence of actions performed in applying rules (plans for reasoning) to derive beliefs from other beliefs. Since reasoning is itself an action, and an action has to be planned before being performed, then before reasoning, the system must first plan the reasoning. Since planning is reasoning about acting, and, in this case, the acting is the act of reasoning, then this planning of the act of reasoning is reasoning about how to reason, or meta-reasoning, and Davis' meta-reasoning cycle can be seen as a special case of the general planning cycle.

We conclude, then, that if an acting-planning-reasoning system uses its acting component to carry out its reasoning, its planning component will automatically perform meta-reasoning. The notion of acting becomes the method that unifies reasoning and meta-reasoning.

5. An Alternative Distinction in Reasoning Systems

The distinction that still remains is "believing versus acting." Believing is a state of knowledge, and acting is the process of changing one state into another state.

The difference between meta-knowledge and meta-reasoning as formulated above is that the former deals primarily with beliefs while the latter deals with acting. We therefore conclude that, besides the conceptual distinction between the object level and the meta level, a valuable distinction to focus on when building KR systems which can have meta-knowledge and can do meta-reasoning is that between believing and acting.

A reasoning system must combine a belief model with an acting model to form a single model. It must have a uniform representation for beliefs and actions to reason effectively about the interaction between knowledge and action. In particular, the system should be able to reason about what knowledge it must have to perform an action, what knowledge it may acquire by performing an action, and what knowledge to use to plan an action [1, 27, 28, 29, 33].

6. The Model

We would like to characterize our model in terms of the cognitive activities and types of knowledge it has to deal with. We will use two orthogonal classifications: one considering the duality of cognitive-system/outside-world; the other the distinction between a state and an action, see Figure 2.

The system performs external activities to interact with the outside world as well as internal activities to form new beliefs and to plan its activities. Reasoning is the internal activity which forms beliefs from other beliefs. Planning is the reasoning activity which plans activities. Planning can be used to plan external activities, or to plan the reasoning activity itself (in which case it is called meta-reasoning).

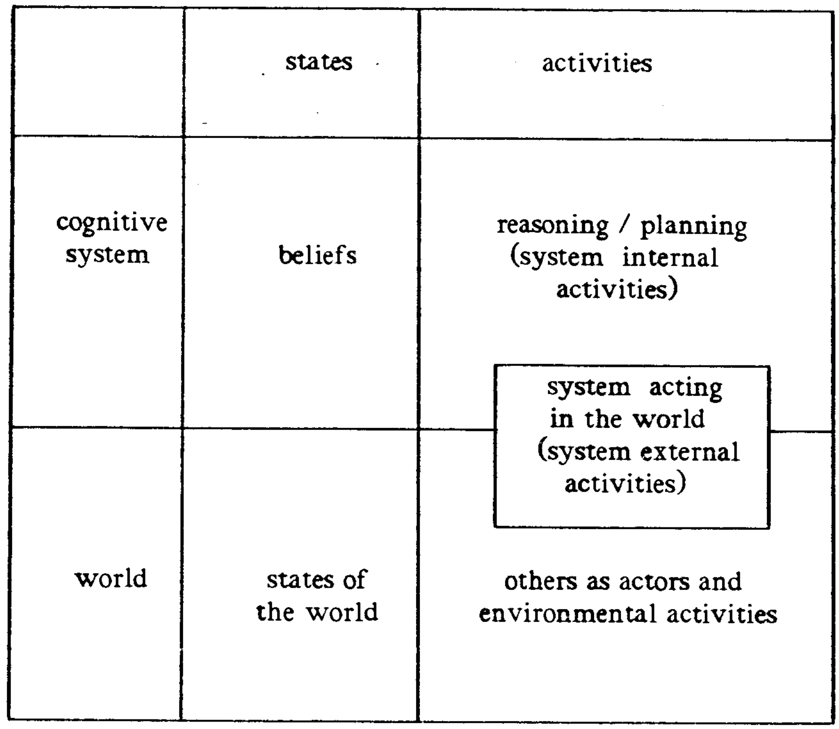

Figure 2

The system has beliefs about external entities as well as beliefs about internal entities (in which case they represent meta-knowledge). External entities are either states of the world or others as actors. Internal entities are either beliefs or system activities.

The system does not need to make any distinctions except those between believing and acting. Thus the system has two components: a belief model and an acting model.

In addition to representing beliefs and actions, the belief model needs to represent sequences of actions in order to represent plans. Plans are like any other concepts; they can be formed, described, and used to form other concepts. Their representation should be consistent with the representation of other concepts. There is however, a peculiarity that makes plans different from other concepts: they can be executed. This suggests that their representation must also be adequate for interpretation by the acting system.

The acting system is the component that carries out the actions according to directions specified in a plan. While executing, it has to create an action history representing the sequence of actions performed, so that the system knows what it is doing and can talk about it (explanation). The action history should thus be represented in the same way as any other concept.

7. Conclusion

We have described a study of meta-knowledge and meta-reasoning which presents an alternative method of addressing these issues in AI KR systems. We claim that a uniform and declarative representation for all types of knowledge, and a unifying model for all types of activities, constitute an alternative design criteria for a system that: knows about states and activities in the world; knows about its own states and activities; and knows how to plan and carry out its own internal and external activities.

An adequate model of believing unifies knowledge and meta-knowledge, and an adequate model of acting unifies reasoning and meta-reasoning. Instead of the distinction between knowledge and meta-knowledge or between reasoning and meta-reasoning, we propose the alternative distinction between believing and acting.

In philosophy there is a substantial literature on the logic of knowledge and belief [18, 19, 22] and on the theory of reason and action [2, 6, 7,]. These topics [1, 3, 23], as well as the topics of meta-knowledge and meta-reasoning [4, 5, 13, 14, 16, 17, 29, 34] and the interaction between knowledge and action [3, 27, 28, 33] have also received considerable attention in AI recently. The contribution of our study is to provide an insight into the relations among these issues in AI KR systems. It demonstrates how a KR system, which integrates an adequate model of believing with an adequate model of acting, can represent meta-knowledge and handle meta-reasoning with no additional mechanisms. Originally motivated by the existence of limitations in the design of SNePS, the study has led to a theory from which KR systems in general may benefit.

References

(1) Amarel, S., On Representations of Problems of Reasoning About Actions, in *Machine Learning and Heuristic Programming*, —, 131-171

(2) Aune, B., *Reason and Action*, Dordrecht, 1977

(3) Appelt, D., A Planner for Reasoning About Knowledge and Action, *Proc. AAAI-80*, Stanford, 1980, 131-133

(4) Bahr, Avron, Meta-knowledge and Cognition, *Proceedings of the Sixth IJCAI* Computer Science Dept., Stanford University, Stanford, CA, 1979, 31-33

(5) Bowen, K., and Kowalski, R., Amalgamating Language and Meta-Language in Logic Programming, Technical Report, School of Computer and Information Science, Syracuse University, New York, 1981

(6) Braud, M., and Walton, D., *Action Theory*, Dordrecht 1976

(7) Castaneda, Hector-Neri, *Thinking and Doing*, Dordrecht, 1975

(8) Collins, A., Fragments of a Theory of Human Plausible Reasoning, *TINLAP-2*, 1978, 194-201

(9) Davis, R., Interactive Transfer of Expertise, *Proc Fifth IJCAI*, Aug 1977a

(10) Davis R., Generalized Procedure Calling and Content-Directed Invocation, *Proc AI/PL Conference*, Aug 1977b

(11) Davis, R., Knowledge Acquisition in Rule-Based Systems - Knowledge about Representations as a Basis for System Construction and Maintenance, in *Pattern-Directed Inference Systems*, Academic Press, 1978

(12) Davis, R. & Buchanan, B.G., Meta-level Knowledge: overview and applications, *Proc Fifth IJCAI*, Cambridge, Mass, 1977.

(13) Davis, R. & Lenat, D., *Knowledge-based Systems in Artificial Intelligence*, McGraw-Hill, 1982, 227-490

(14) Filman, R., Meta-Language and Meta-Reasoning, Computer Research Center, Hewlett-Packard Laboratories, Palo Alto, California 94304

(15) Findler, N., (ed.), *Associative Networks*, Academic Press, Inc., New York, 179-203

(16) Gallaire, H. & Lassere, C., Controlling Knowledge Deduction in a Declarative Approach, *Proc Sixth IJCAI*, 1979

(17) Genesereth, M., An Overview of Meta-Level Architecture, --- , 119-123

(18) Hintikka, J., *Knowledge and Belief*, Cornell University Press, Ithica, New York, 1963

(19) Hintikka, J., Semantics for Propositional Attitudes, in Linski, 1971, 145-167

(20) Kollers, P. A., & Palef, S. R., Knowing Not, in *Memory and Cognition 4-5*, 1976, 553-558

(21) Konolige, K., A Metalanguage Representation of Relational Databases for deductive question-answering system, *Proc. IJCAI-81*, 1981, 496-503

(22) Linski, L., (ed.), *Reference and Modality*, Oxford University Press, London, 1971

(23) Maida, A., and Shapiro, S. Intensional Concepts in Propositional Semantic Networks, in *Cognitive Science 6*, 1982, 291-330

(24) Martins, J., Mckay, D., and Shapiro, S.C., Bi-Directional Inference, Technical Report 174, Dept. of Comp. Sci., SUNY/AB, Buffalo, New York, 1981

(25) McKay, D., and Shapiro, S.C., Using Active Connection Graphs for Reasoning with Recursive Rules, *Proc. IJCAI-81, 1981, 368-374*

(26) McKay, D., and Shapiro, S.C., MULTI - A LISP Based Multiprocessing System, *Proc. LISP Conference 1980*, Stanford University, 1980

(27) Moore, R., Reasoning About Knowledge and Action, *Proc. Fifth IJCAI*, 1977, 223-227

(28) Moore, R., Reasoning About Knowledge and Action, Technical Note 191, AI Center, Comp. Sci. and Tech. Div., SRI International, Menlo Park, California 94025, 1980

(29) Morgado, E., Believing and Acting: An Approach to Meta-Knowledge and Meta-reasoning, a Ph.D. proposal, Dept. of Computer Science, SUNY/AB, 1980

(30) Sacerdoti, E. D., *A structure for Plans and Behavior*, Elsevier North-Holland, New York, 1977

(31) Shapiro, S.C., The SNePS Semantic Network Processing System, in Findler, 1979

(32) Shapiro, S.C., On Representing About, Extended Abstract, Computer Science Dept., SUNY/AB, 1980a

(33) Smith, B., Knowledge Representation Semantics, *Proc. Fifth IJCAI*, 1977, 987-990

(34) Stefik, M., Planning and Meta-Planning, MOLGEN:Part 2, Computer Science Dept., Stanford University, Stanford, California 94305, 1980b