SOFT CONTROL OF COGNITIVE PROCESSES

Michael R. Fehling

USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90291

Submitted to:
Fourth Annual Conference
of the Cognitive Science Society
Ann Arbor, Michigan August 4-6, 1982

Topic:
Problem Solving and Knowledge Structures
Sponsor:
Gary M. Olson

1. INTRODUCTION

A critical feature of any problem solving system is its control structure. This, of course, refers to a mechanism (and its associated knowledge) used allocate processing resources among the various components of the system as they are needed to carry out some task within some problem domain. It is clear that the control structure of a problem solver is a major determiner of that system's ability to efficiently and effectively carry out any task.

As important as the notion of control is, it is surprising that so little work has been devoted to it by either computer scientists interested in developing expert systems or psychologists interested in modeling human cognition. It has been the feeling in Artificial Intelligence that, if there were enough knowledge available in the construction of an expert system, the problem of selecting an appropriate control structure would be a minor one (Feigenbaum, 1977). And, as we shall discuss below, although some recent psychological models have addressed issues that are closely related to the control problem, little or no research has directly addressed the general question of control of cognitive processes.

In this paper we report some work we are doing on the control problem. The ultimate goal of this research is to design and implement an expert system that controls other expert systems. That is, we are developing a problem solving system that is specialized to select and maintain a control regime for components of another "embedded" expert system. Our Expert System Controller (ESC) is able to reason about control. It uses both general knowledge and domain specific knowledge of the embedded system to create and maintain control plans for scheduling the use of the embedded system's component processes.

It is our belief that the issue of reasoning about control is one that must be addressed by anyone interested in developing more powerful problem solving systems, whether those systems are intended as expert systems or as models of human cognition. Moreover, it is a central premise of our research that such systems require soft control. By this term we mean the following:

> The ability to apply problem solving techniques to the problem of control itself (i.e., to reason about control)
>
> The ability to select from alternative control plans the one that is most appropriate in a particular task environment
>
> The ability to apply general (albeit less powerful) knowledge when specific domain knowledge is unavailable
>
> The ability to opportunistically deviate from a selected control strategy as a response to new information.

Soft control yields a flexibility of interaction among the various components of domain and control knowledge that allow for opportunistically allocating resources to activities most likely to make efficient progress in completing the task at hand.

2. META-COGNITION and CONTROL

First, let us discuss control in terms of human cognition by consider the vast amount of psychological research on the use of strategies to guide processing. A brief examination of research on this topic shows that, in any given task context, some particular processing strategy may be proposed as the organizer and controller for a more basic set of cognitive skills. Strategy guided models have been offered as a description of many types of cognitive skills. Examples include models for text processing (e.g., Clark, 1978), logical inference (Revlin & Leirer, 1978), memory retrieval (Brown, 1978), perception (Kolers, 1972), and so forth.

Perhaps a generalization and expansion of the idea of processing strategy is Flavell's (1976) concept of meta-cognition. Meta-cognition refers to cognitive processing involving knowledge about other cognitive processes or the results of other cognitive processes. One place where this concept has been used extensively is the research on the topic of learning strategies (e.g., O'Neil, 1978; O'Neil & Speilberger, 1979).

This research demonstrates the ubiquity of task specific strategies. Each strategy appears to be a kind of specialized "control plan" that organizes the cognitive processes underlying performance in a particular task domain. This, in turn, suggests that there exists some general mechanism to produce these specialized control plans and to monitor their use. Although little work has been done to determine the characteristics of the meta-cognitive mechanism, we note the following important features.

> First the diversity of strategies that arise in different contexts indicates that these meta-cognitive structures are typically highly "tuned" to the specific problem domain Thus, both creation and selection for use of such control plans is a function of specific domain knowledge.
>
> Second, the use of strategies is opportunistic in that use of one strategy may be interrupted or even abandoned in favor of another known strategy as a response to some special circumstance that is noticed during task performance.
>
> Third, control can revert to more general knowledge and problem solving techniques when situation specific knowledge is insufficient.

These observations together indicate that metacognition is probably best modeled as what we refered to above as a mechanism for soft control.

Now let us consider the need for soft control in the context of expert systems in Artificial Intelligence research. We wish to show that there is a need for soft control in expert system just as that required for models of human cognition.

Recall Feigenbaum's argument, mentioned earlier, that the control problem for expert systems is secondary to the problem of representing sufficient knowledge about the problem domain. The knowledge in an expert system embodies primarily expert "rules of thumb" and descriptions for when such knowledge is applicable. Any such rule of thumb is typically a large chunk of domain specific knowledge that has compiled into it the control that would have been necessary to take the several smaller steps that are equivalent to it. Reasoning with such large chunks produces shorter inference chains which, therefore, greatly reduce the magnitude of the control problem for managing these inferences. In this sense most expert systems simply finesse the control problem by relying upon a very powerful set of domain specific principles that embody both domain knowledge and control assumptions for use of that knowledge.

Unfortunately, the exclusive use of expert rules can have severe limitations. The powerful domain principles of the expert system are usually only plausible rules of inference which do not embody logically necessary relationships. Hence, expert systems of this sort can fail precipitously at the limit of their knowledge, that is, when the system encounters new situations for which the special rules do not apply. When such rules fail the system is unable to retreat to weaker but more general methods of inference to determine such things as why the rule failed in this case, how to modify it to fit, or at least, how to start from smaller and less efficient but more universal principles to derive a response to the new situation. That is, control is too rigid to allow the system's performance to gracefully degrade as the limits of its expert knowledge are reached.

A general solution, which we have adopted, is to provide the expert system with an ability to revert to the more basic form of problem solving when the expert rules do not apply. However, control methods for using the expert rules will probably be useless for the more complex inferences required when using general principles. So the expert system must be able to select (or construct) a new control plan that is appropriate for the type of knowledge being used at each point in the task. Moreover, the system must detect when and how to make a graceful shift from one mode of inference to another. In general, the system must be able to develop or select from a stock of control plans that allow the system to use a variety of types of knowledge during task performance.

Therefore, to build a more flexible expert system or a more general cognitive model, one must design a system that has the ability to reason about control. Furthermore, the system must be able to select an appropriate control regime for a specific task context. It must have the ability to apply expert "rules of thumb", or, when such rules are not available, it must be able to engage in novel reasoning using finer grained and less specific logical rules. And it must be able to decide when to do which. That is, either a cognitive model or an expert system needs a means to provide soft control.

3. THE EXPERT SYSTEM CONTROLLER (ESC)

Next, we briefly describe some features of an expert system we are developing that realize the concept of soft control. As stated earlier, the primary application of this system is as an expert system to control other expert systems. However, in creating a problem solving architecture in which both domain and control plan reasoning are supported, we are developing a type of model that may also be valuable as a framework for developing cognitive models in which meta-cognitive processes, as well as the processes and knowledge they control, can be explicitly described.

In order that our Expert System Controller (ESC) have the capability to provide soft control, it must have the following features.

> An architecture which supports problem solving about selection, modification, and use of control plans as well as problems within a substantive problem domain.
>
> A representation language for expressing control relations (e.g., sequencing, tests, parallelism, etc.)
>
> The ability to opportunistically modify or

abandon a control strategy in response to new information.

ESC is an extension of the Hearsay-III problem solving system (Erman et al., 1981). Hearsay-III is a "blackboard model" in which knowledge is represented by a collection of individual processing components called "knowledge sources" (KSs). KSs embody the knowledge associated with a particular part of a problem solving task and are activated by the occurence of patterns on a "communications blackboard". KSs can interact during problem solving by leaving new "triggering" patterns on the blackboard that activate other KSs. Since more than one KS can be activated at a time, a "scheduler" is provided that makes decisions as to the firing order of the activated KSs. (For the reader unfamiliar with the architecture of blackboard models, see Rummelhart, 1977, pp. 103-116.)

Hearsay-III provides blackboard structures for both domain and scheduling purposes and provides for the implementation and use of knowledge sources for scheduling as well as domain knowledge sources. Thus, reasoning about scheduling can be accomplished by methods that are consistent with those used for problem domain reasoning. In order to extend the problem solving capabilities of this model to the full domain of control concepts we are adding an explicit control representation and a mechanism to react to that representation. The control notions that can be represented include

Programmatic control relations (e.g., sequential, parallel, or conditional)

Non-programmatic control relations (e.g., co-operative subprocesses [all of which combine to contribute to some goal] versus competitive subprocesses [each of which provides an alternative way to achieve a subgoal])

Descriptors of problem structures, goals, and knowledge sources.

Descriptors of hierarchical plans as well as descriptors of conditions under which control "jumps" out of such a plan in a non-hierarchical way.

Methods based on some work we have been doing using the Dempster-Shafer calculus (Shafer, 1976), to express preference relations among plans and activities (cf. also, Barnett, 1981).

Besides developing an explicit representation for reasoning about control, we are augmenting the architecture of Hearsay-III to fully support the control domain as a problem-solving activity. This extension provides abilities such as the following

Interpretation of control plans in the representation alluded to above.

Filling out of partially specified control plans using domain independent control knowledge to affect the elaboration

Optimization of execution within plans by using any applicable general knowledge.

Construction of control plans using preference relations supplied by scheduling knowledge sources when more specific control constraints are not available.

Communication of plan progress to scheduling knowledge sources, thus allowing the scheduling knowledge sources to modify and improve plans opportunistically.

4. CONCLUSIONS

The explicit control representation and other modifications we are making to the basic Hearsay-III architecture provide a means to achieve soft control in a problem solving system. We believe this model will be useful as a framework for building expert systems that have greater flexibility and power than those currently available.

This framework should also interest cognitive scientists whose concern is models of human cognitive since it provides a framework for a model in which meta-cognitive processes are treated uniformly with all other cognitive processes. Newell (1980) has pointed out that, if we are to get rid of the homunculus that always controls the processes of cognitive models, we must incorporate into those models a representation of the way that strategies arise from general and domain specific knowledge as a response to task conditions. Perhaps the issues that we have raised in developing our notion of soft control will help to evict this homunculus.

Note: This research was supported by Defense Advanced Research Projects Agency contract MDA 903-81-C-0335. Views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official opinion or policy of DARPA, the U.S. Government, or any other person or agency connected with them.

References

Barnett, J. A. Computational Methods for a Mathematical Theory of Evidence. In Proc. 7th Joint Conf. on Artificial Intelligence. IJCAI, Vancouver, B. C., 1981.

Brown, A. L. Knowing When, Where, and How to Remember: a Problem of Meta-Cognition. In R. Glaser (Ed.), Advances in Instructional Psycholgoy, Hillsdale, N.J.: Lawrence Erlbaum Associates, 1978.

Clark, H. H. Linguistic Processes in Deductive Reasoning. Psychological Review, 1978, 72(4), 387-404.

Erman, L. D., London, P. E., & Fickas, S. F. The Design and Example Use of Hearsay-III. In Proc. 7th Joint Conf. on Artificial Intelligence. IJCAI, Vancouver, B. C., 1981.

Feigenbaum, E. A. The Art of Artificial Intelligence. In Proc. 5th Int. Joint Conf on Artificial Intelligence, pages 1014-1029. IJCAI, Cambridge, MA, 1977.

Flavell, J. H. Meta-cognitive Aspects of Problem Solving. In L. B. Resnick (Ed.), The Nature of Intelligence, Hillsdale, N. J.: Lawrence Erlbaum Associates, 1976.

Kolers, P. A. Some Problems of Classification. In J. F. Kavanaugh & I. G. Mattingly (Eds.), Language by Ear and by Eye: The Relationships between speech and Hearing, Cambridge: M.I.T. Press, 1972.

Newell, A. Reasoning, Problem Solving, and Decision Processes: The Problem Space as a Fundamental Category. In R. S. Nickerson (Ed.), Attention and Performance, Hillsdale, N. J.: Lawrence Erlbaum Associates, 1980.

O'Neil, H. F. (Ed.). Learning Strategies. New York: Academic Press 1978.

O'Neil, H. F. & Speilberger, C. D. (Eds.). Cognitive and Affective Learning Strategies. New York: Academic Press 1979.

Revlin, R. & Leirer, V. O. The Effect of Personal Biases on Syllogistic Reasoning: Rational Decisions from Personalized Representations. In R. Revlin & E. Mayer (Eds.), Human Reasoning, Washington, D.C.: Winston, 1978.

Rummelhart, D. E. Human Information Processing. New York: Wiley 1977.

Shafer, G. A Mathematical Theory of Evidence.: Princeton University Press 1976.