ELSEVIER

Neurocomputing 44–46 (2002) 735–742

NEUROCOMPUTING

www.elsevier.com/locate/neucom

# Problem-solving behavior in a system model of the primate neocortex

Alan H. Bond

*California Institute of Technology, Mailstop 136-93, Pasadena, CA 91125, USA*

**Abstract**

We show how our previously described system model of the primate neocortex can be extended to allow the modeling of problem-solving behaviors. Specifically, we model different cognitive strategies that have been observed for human subjects solving the Tower of Hanoi problem. These strategies can be given a naturally distributed form on the primate neocortex. Further, the goal stacking used in some strategies can be achieved using an episodic memory module corresponding to the hippocampus. We can give explicit falsifiable predictions for the time sequence of activations of different brain areas for each strategy. © 2002 Published by Elsevier Science B.V.

*Keywords:* Neocortex; Modular architecture; Perception–action hierarchy; Tower of Hanoi; Problem solving; Episodic memory

## 1. Our system model of the primate neocortex

Our model [4–6] consists of a set of processing modules, each representing a cortical area. The overall architecture is a perception–action hierarchy. Data stored in each module is represented by logical expressions we call descriptions, processing within each module is represented by sets of rules which are executed in parallel and which construct new descriptions, and communication among modules consists of the transmission of descriptions. Modules are executed in parallel on a discrete time scale, corresponding to 20 ms. During one cycle, all rules are executed once and all inter-module transmission of descriptions occurs. Fig. 1 depicts our model, as a set of cortical modules and as a perception–action hierarchy system diagram. The action of the

*E-mail address:* bond@cs.caltech.edu (A.H. Bond).


PII: S0925-2312(02)00466-6

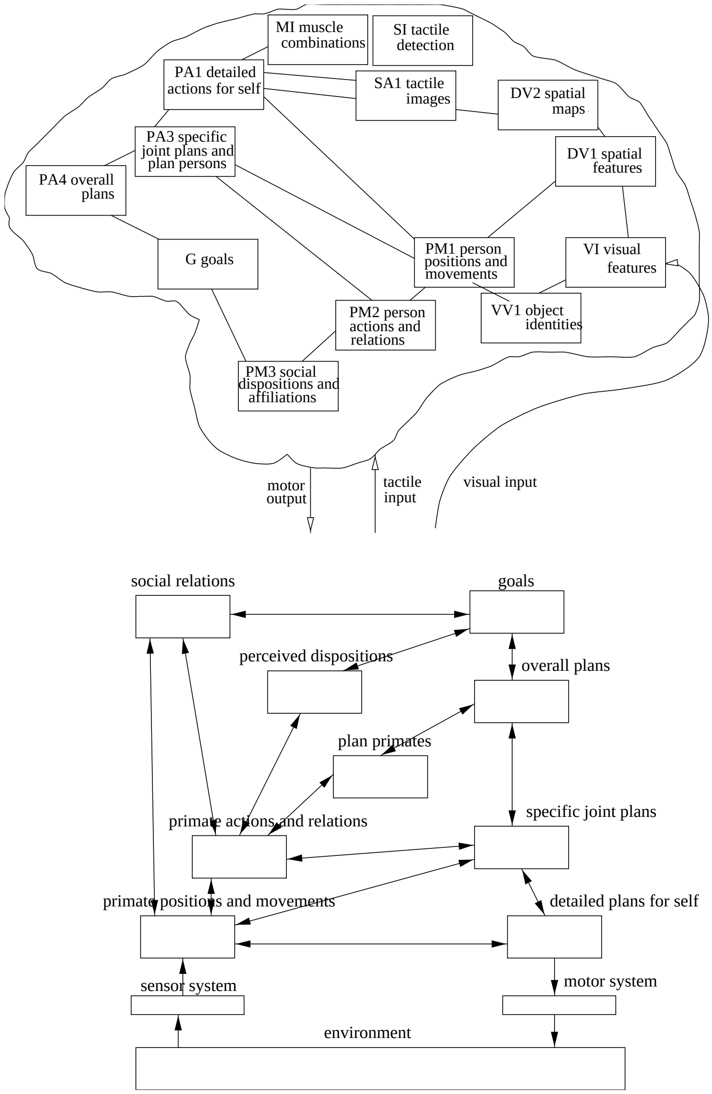


Fig. 1. Our system model shown in correspondence with the neocortex, and as a perception–action hierarchy.

system is to continuously create goals, prioritize goals, and elaborate the highest priority goals into plans, then detailed actions by propagating descriptions down the action hierarchy, resulting in a stream of motor commands. (At the same time, perception of the environment occurs in a flow of descriptions up the perception hierarchy. Perceived descriptions condition plan elaboration, and action descriptions condition perception.) This simple elaboration of stored plans was sufficient to allow is to demonstrate simple socially interactive behaviors using a computer realization of our model.

## 2. Extending our model to allow solution of the Tower of Hanoi problem

### 2.1. Tower of Hanoi strategies

The Tower of Hanoi problem is the most studied, and strategies used by human subjects have been captured as production rule systems [9,1]. We will consider the two most frequently observed strategies—the perceptual strategy and the goal recursion strategy. In the general case, reported by Anzai and Simon [3], naive subjects start with an initial strategy and learn a sequence of strategies which improve their performance. Our two strategies were observed by Anzai and Simon as part of this learning sequence. Starting from Simon's formulation [8], we were able to represent these two strategies in our model, as follows:

### 2.2. Working goals

Since goals are created dynamically by the planning activity, we needed to extend our plan module to allow *working goals* as a description type. This mechanism was much better than trying to use the main goal module. We can limit the number of working goals. This would correspond to using a fixed size store, corresponding to working memory. The module can thus create working goals and use the current working goals as input to rules. Working goals would be held in dorsal prefrontal areas, either as part of or close to the plan module. Main motivating topgoals are held in the main goal module corresponding to anterior cingulate.

### 2.3. Perceptual tests and mental imagery

The perceptual tests on the external state, i.e. the state of the Tower of Hanoi apparatus, were naturally placed in a separate perception module. This corresponds to Kosslyn's [7] image store. The main perceptual test needed is to determine whether a proposed move is legal. This involves (a) making a change to a stored perceived representation corresponding to making the proposed move, and (b) making a spatial comparison in this image store to determine whether the disk has been placed on a smaller or a larger one. With these two extensions, we were able to develop a representation of the perceptual strategy, depicted in Fig. 2.

## 3. Episodic memory and its use in goal stacking

In order to represent the goal recursion strategy, we need to deal with goal stacking, which is represented by push and pop operations in existing production rule representations.

Since we did not believe that a stack with push and pop operations within a module is biologically plausible, we found an equivalent approach using an *episodic memory* module.

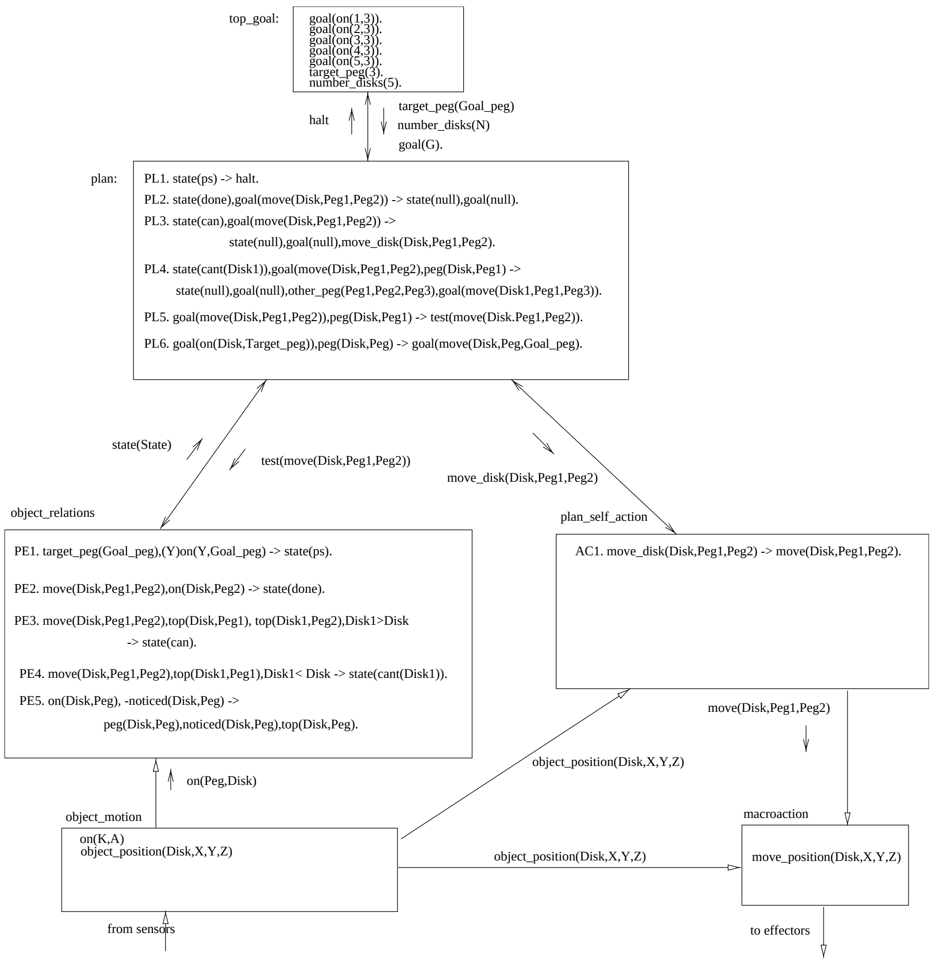


Fig. 2. Representation of the perceptual strategy on our brain model.

This module creates associations among whatever inputs it receives at any given time, and it sends these associations as descriptions to be stored in contributing modules. In general, it will create episodic representations from events occurring in extended temporal intervals; however, in the current case we only needed simple association.

In the Tower of Hanoi case, the episode was simply taken to be an association between the current working goal and the previous, parent, working goal. We assume that these two working goals are always stored in working memory and are available to the plan module. The parent forms a context for the working goal. The episode description is formed in the episodic memory module and transmitted to the plan module where it is stored. The creation of episodic representations can proceed in parallel with the problem solving process, and it can occur automatically or be requested by the plan module. Rules in the plan module can retrieve episodic descriptions using

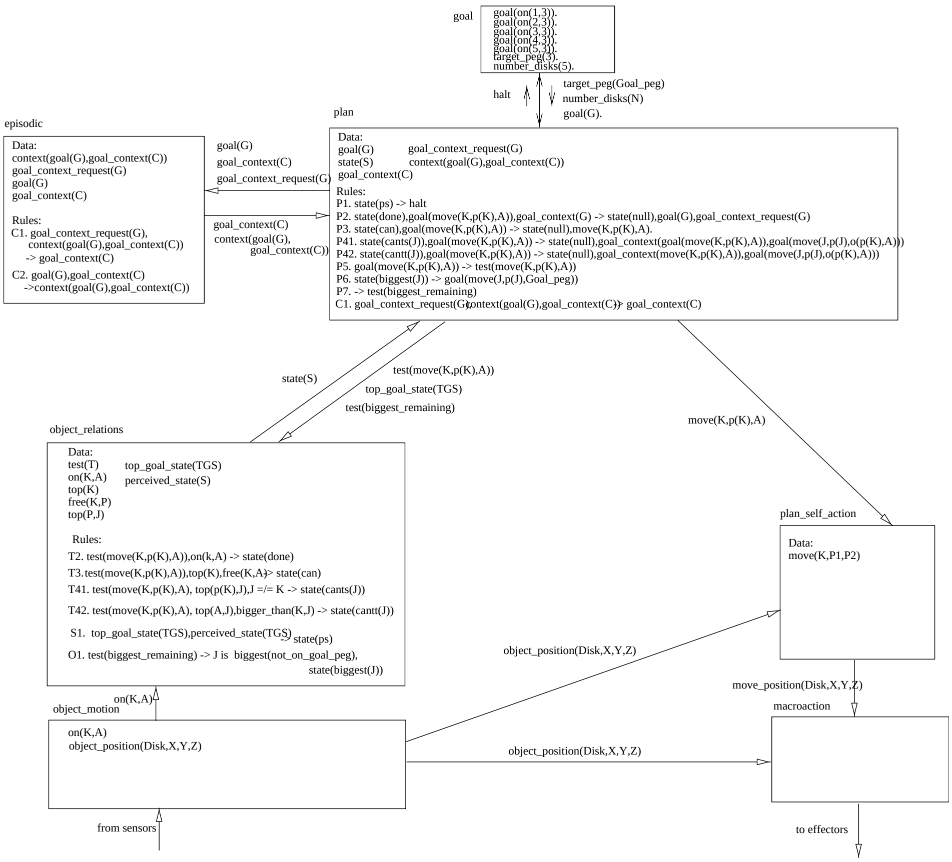


Fig. 3. Representation of the goal recursion strategy on our brain model.

the current parent working goal, and can replace the current goal with the current parent, and the current parent with its retrieved parent. Thus the working goal context can be popped. This representation is more general than a stack, since any stored episode could be retrieved, including working goals from episodes further in the past. Such effects have, in fact, been reported by Van Lehn et al. [10] for human subjects.

With this additional extension, we were able to develop a representation of the goal recursion strategy, depicted in Fig. 3. Descriptions of episodes are of the form context(goal(G),goal_context(C)). goal(G) being the current working goal and goal_context(C) the current parent working goal. The figure shows a slightly more general version, where episodes are stored both in the episodic memory module and the plan module. This allows episodes that have not yet been transferred to the cortex to be used.

We are currently working on extending our model to allow the learning a sequence of strategies as observed by Anzai and Simon. This may result in a different representation of these strategies, and different performance.

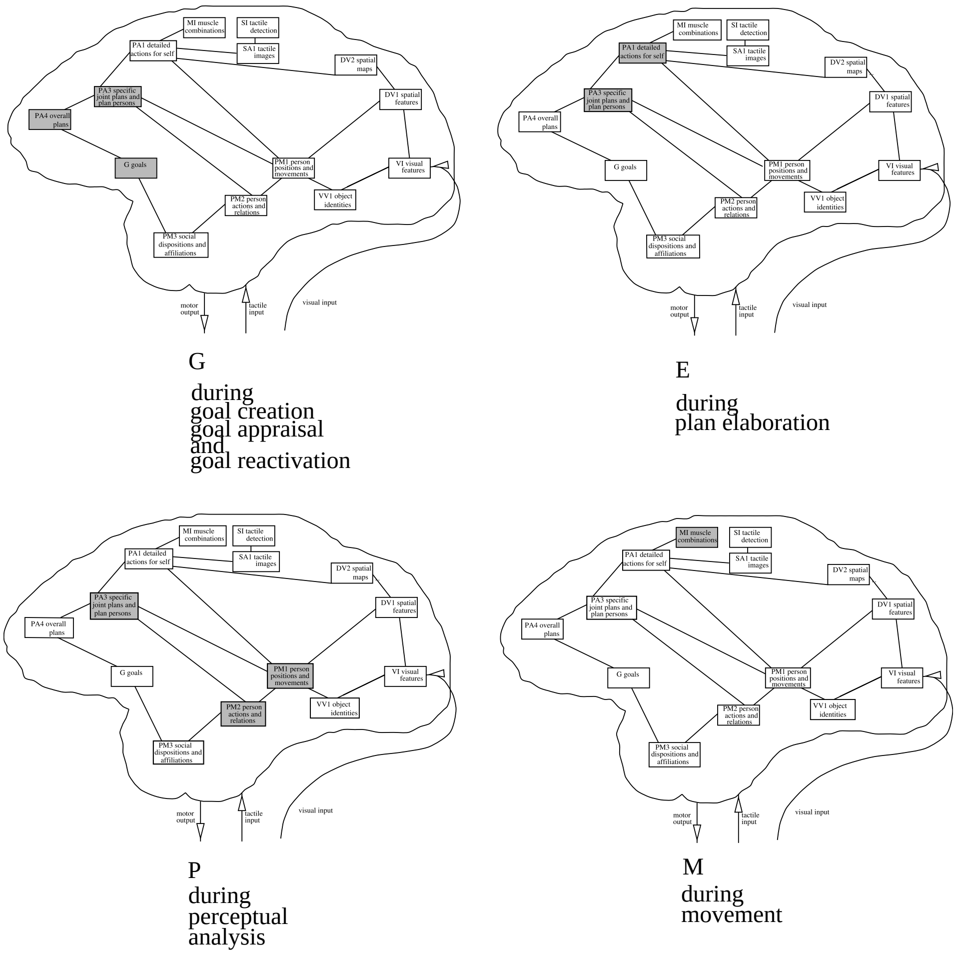


Fig. 4. Predictions of brain area activation during Tower of Hanoi solving.

## 4. Falsifiable predictions of brain area activation

For the two strategies, we can now generate detailed predictions of brain area activation sequences that should be observed during the solution of the Tower of Hanoi problem. Using our computer realization, we can generate detailed predictions of activation levels for each time step. Since there are many adjustable parameters and detailed assumptions in the model, it is difficult to find clearly falsifiable predictions. However, we can also make a simplified and more practical form of prediction by classifying brain states into four types, shown in Fig. 4.

Let us call these types of states G, E, P and M, respectively. Then, for example, the predicted temporal sequences of brain state types for 3 disks are:

For the perceptual strategy:

$$\mathrm{G0, G, E, P, G, E, P, G, E, P, E, M, P, G, E, P, G, E, P, E, M, P, G, E, P, G, E, P, E, M, P,}$$

$$\mathrm{G, E, P, E, M, P, G, E, P, G, E, P, E, M, P, G, E, P, E, M, P, G, E, P, E, M, P, G0.}$$

and for the goal recursion strategy:

$$\mathrm{G0, G, E, P, G+, E, P, G+, E, P, E, M, P, G^*, E, P, E, M, P, G^*, E, P, G+, E, P, E, M, P,}$$

$$\mathrm{G^*, E, P, E, M, P, G, E, P, G+, E, P, E, M, P, G^*, E, P, E, M, E, G, E, P, E, M, P, G0.}$$

We can generate similar sequences for different numbers of disks and different strategies. The physical moves of disks occur during $M$ steps. The timing is usually about 3.5 s per physical move, but the physical move steps probably take longer than the average cognitive step. If a physical move takes 1.5 s, this would leave about 300 ms per cognitive step.

The perceptual strategy used is an expert strategy where the largest disk is always selected. We assume perfect performance; when wrong moves are made, we need a theory of how mistakes are made, and then predictions can be generated. In the goal recursion strategy, we assume the subject is using perceptual tests for proposed moves, and is not working totally from memory. G indicates the creation of a goal, G+ a goal creation and storing an existing goal (push), and G* the retrieval of a goal (pop). Anderson et al. [2] have shown that pushing a goal takes about 2 s, although we have taken creation of a goal to not necessarily involve pushing. For us, pushing only occurs when a new goal is created and an existing goal has to be stored. G0 is activity relating to the top goal.

It should be noted that there is some redundancy in the model, so that, if a mismatch to experiment is found, it would be possible to make some changes to the model to bring it into better correspondence with the data. For example, the assignment of modules to particular brain areas is tentative and may need to be changed. However, there is a limit to the changes that can be made, and mismatches with data could falsify the model in its present form.

## Acknowledgements

This work has been partially supported by the National Science Foundation, Information Technology and Organizations Program managed by Dr. Les Gasser. The author would like to thank Professor Pietro Perona for his support, and Professor Steven Mayo for providing invaluable computer resources.

## References

[1] J.R. Anderson, Rules of the Mind, Lawrence Erlbaum Associates, Hillsdale, NJ, 1993.

[2] J.R. Anderson, N. Kushmerick, C. Lebiere, The Tower of Hanoi and Goal structures, in: J.R. Anderson (Ed.), Rules of the Mind, Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1993, pp. 121–142.

[3] Y. Anzai, H.A. Simon, The theory of learning by doing, Psychol. Rev. 86 (1979) 124–140.
[4] A.H. Bond, A computational architecture for social agents, Proceedings of Intelligent Systems: A Semiotic Perspective, An International Multidisciplinary Conference, National Institute of Standards and Technology, Gaithersburg, Maryland, USA, October 20–23, 1996.
[5] A.H. Bond, A system model of the primate neocortex, Neurocomputing 26–27 (1999) 617–623.
[6] A.H. Bond, Describing behavioral states using a system model of the primate brain, Am. J. Primatol. 49 (1999) 315–388.
[7] S. Kosslyn, Image and Brain, MIT Press, Cambridge, MA, 1994.
[8] H.A. Simon, The functional equivalence of problem solving skills, Cognitive Psychol. 7 (1975) 268–288.
[9] K. VanLehn, Rule acquisition events in the discovery of problem-solving strategies, Cognitive Sci. 15 (1991) 1–47.
[10] K. VanLehn, W. Ball, B. Kowalski, Non-LIFO execution of cognitive procedures, Cognitive Sci. 13 (1989) 415–465.

**Alan H. Bond** was born in England and received a Ph.D. degree in theoretical physics in 1966 from Imperial College of Science and Technology, University of London. During the period 1969–1984, he was on the faculty of the Computer Science Department at Queen Mary College, London University, where he founded and directed the Artificial Intelligence and Robotics Laboratory. Since 1996, he has been a Senior Scientist and Lecturer at California Institute of Technology. His main research interest concerns the system modeling of the primate brain.