

IN 63-TM

91522

P-45

**Computer Modeling of Human  
Decision Making**

**WILLIAM B. GEVARTER**

**AI RESEARCH BRANCH, MAIL STOP 269-2  
NASA AMES RESEARCH CENTER  
MOFFETT FIELD, CA 94035  
(415) 604-6527**

(NASA-TM-107868) COMPUTER MODELING OF HUMAN  
DECISION MAKING (NASA) 45 p

N92-26630

Unclas  
63/63 0091522

**NASA Ames Research Center**

**Artificial Intelligence Research Branch**

**Technical Report FIA-91-32**

**December, 1991**



# COMPUTER MODELING OF HUMAN DECISION MAKING

William B. Gevarter  
Artificial Intelligence Research  
NASA Ames Research Center  
Moffett Field, CA 94035

Report

Oct. 31, 1991

# COMPUTER MODELING OF HUMAN DECISION MAKING

William B. Gevarter

## Abstract

### I. Introduction

### II. Motivations in Decision Making

### III. Unimplemented Models that Ignore Motivations

PASS

Recognition-Primed Decision Making

Decision Making with Planning

GEMS (Generic Error Modeling System)

Other Models of Human Decision Making

### IV. Computer Models that Ignore Motivations

ACT\*

Soar

### V. Computer Models that Include Motivations

Hot Cognition

MoCog1

DAYDREAMER

Goals in DAYDREAMER

Emotions

Basic Information Flow

The DAYDREAMER Planner

## **VI. Computer Simulations of Humans Interacting with a Simulated World**

**MIDAS**

## **VII. Discussion**

**References**

**Figures**



# COMPUTER MODELING OF HUMAN DECISION MAKING

William B. Gevarter

## Abstract

This report reviews models of human decision making. Models which treat just the cognitive aspects of human behavior are included as well as models which include motivation. Both models which have associated computer programs, and those that do not, are considered. As flow diagrams, that assist in constructing computer simulation of such models, were not generally available, such diagrams have been constructed and are presented in this report. The result provides a rich source of information, which can aid in construction of more realistic future simulations of human decision making.

## I. Introduction

This report briefly reviews the current state of computer simulation of human decision making. Such models can aid in the design of intelligent machines that interact with humans. They are also useful in the design of equipment, automation and missions compatible with humans. In addition, such models can aid in the design of training programs, and in understanding and preventing human errors. In these ways, improved modeling of human behavior can be of substantial benefit to NASA in its lead role in manned space activities and its role in research for manned aircraft operations.

At present, there are few computer models of human decision making. Thus, we begin by constructing flow diagrams of some human decision paradigms. Most of these paradigms have not yet been converted to simulations, but these flow diagrams are suggestive of the architecture of computer programs that can be created to implement them. We then examine some of the few computer simulations that are now available.

## II. Motivations in Decision Making

It is generally acknowledged (cf. Izard, 1984) that emotions are involved in all human decision making. However, the nature and mechanisms of the associated motivations are still incompletely understood. Thus, in most paradigms of human decision making, and in cognitive science in particular, motivations are ignored and just the information processing aspects are considered. Therefore, we will start with models that do not explicitly include motivations and then move on to those that do include them.

### III. Unimplemented Models that Ignore Motivations

#### PASS

Maher (1991) offers us a model that he believes is descriptive of air crew cockpit behavior. PASS is derived from the analytical core of the process (Problem identification, Acquisition of information, Strategy survey, and Solution selection). A flow diagram we have constructed reflecting PASS is given in Figure 1. The central core of PASS consists of a straightforward sequential process of identifying a problem, gathering information, and moving on to a strategy survey and solution selection. After selecting a solution, the participants decide if the problem is solved by it, if not they reenter the loop. They may choose to re-define the problem, seek more information or choose the next alternative.

Maher indicates that the decision problem is often complicated by what he refers to as pilot tendencies to shortcut or short circuit the information gathering process. These tendencies are:

#### Representativeness

When faced with a problem experienced and successfully solved in the past, humans tend to bypass the information gathering process in PASS, and apply the past solution. (Note that this is also the paradigm for straightforward "procedure following" in the accomplishment of a task).

### Availability

This is a tendency to recall recent or vivid, compelling solutions and use them, bypassing information gathering.

### Anchoring

This is a tendency to fasten on to a piece of information on the first pass through a problem and use it to support an early conclusion, strongly resisting any change based on having to consider an alternate piece of information.

### Overconfidence

This is the tendency for humans to be overly sure of their decisions and to resist change, even when faced with compelling evidence to do so.

When activated, these tendencies produce alternative paths through the decision process. These paths are also shown in Figure 1.

Maher's paper also include observations on personality, recall, attention, and learning, which are pertinent to air crew decision making and error avoidance. In particular, Maher suggests that the tendencies depicted in Figure 1 are more likely to be actuated when attention is reduced. He indicates that attention is a function of stress, cognitive load, degree of information saturation, and fatigue. Suitable algorithms using these relationships could be employed in a simulation to account for tendency actuation.

### Recognition-Primed Decision Making

Klein (1989) interviewed skilled decision makers and discovered that they usually do not have to consciously consider alternatives. They found that based on experience, the experts can usually judge the prototypicality of a situation and that this assessment includes a recognition of a typical way to react. In terms of the PASS model, this means that the representiveness and availability tendencies are active. Figure 2 is a flow diagram of recognition-primed decision making. Following the flow in Figure 2, the expert automatically recognizes the task as a familiar one, and automatically retrieves the associated goal appropriate to deal with the situation. From the possible actions associated with that goal, a salient one is

automatically retrieved, and blending it with the available resources, the expert acts on it, modifying the situation, and continues to loop through this process until the situation is resolved.

In more complex cases, diagramed in Figure 3, experts do perform some conscious evaluation of their reaction, serially evaluating each recalled possible action as it become salient. The expert uses mental imagery to simulate the likely result of implementing an action, either modifying the action, seeking more information or evaluating the next salient action until one is found that appears satisficing. After implementing the action deemed as satisficing, the expert then monitors the result looking for cues that will confirm or disconfirm the expert's expectancies, continuing to loop through the process as required.

### **GEMS (Generic Error Modeling System)**

Reason (1990) has devised a flow diagram of human decision making in his quest to understand human errors. Figure 4 is my version of a GEMS flow diagram. Reason focuses on three levels of human decision making -- Skill-based, Rule-based, and Knowledge-based; and on two modes of cognitive control -- attentional control associated with working memory, and schematic control derived from long-term memory. Reason (p. 50) indicates that the (conscious) attentional control mode ". . . is limited, sequential, slow, effortful and difficult to sustain for more than brief periods."

In referring to the schematic control mode, Reason (p.51) observes that, "The cognitive system is extremely good at modelling and internalizing the useful regularities of the past and then reapplying them whenever their 'calling conditions' are supplied by intentional activity or by the environment. . . . This schematic control mode can process familiar information rapidly, in parallel and without conscious effort."

The schemas in long term memory, in addition to being activated by attentional activity via plans (descriptions of intended activities), are activated by contextual cueing, recency, frequency of prior use, features shared with other schemas, and emotional factors. When the calling conditions are underspecified, the cognitive system tends to default to contextually appropriate, high frequency responses.

## Decision Making with Planning

Bratman (1987, 1989) has proposed a model of human decision making that incorporates planning. Because humans have limited mental capacity, coordination capacity, time, and other resources, humans usually do not make and enforce complete plans. Rather, humans usually settle for plans that are partial and fill them in as required as time goes by. Action derives from intention to act, which in turn is predicated on what the agent desires and what the agent believes. These intentions are reflected in plans for actions in the future. Plans are hierarchical, with more specific plans being imbedded in very general plans such as those reflecting life goals.

Plans are subject to constraints. Plans tend to be consistent with beliefs and with potential means for carrying them out. Future-directed intentions and partial plans have a characteristic stability. Prior intentions, not up for reconsideration, constrain future intentions. Partial plans associated with prior intentions constrain future options for consideration, enabling humans not to be overwhelmed with information processing.

Personal policies are like intentions, but do not give rise to future plans. Instead, they guide what to do in specific situations. In a manner similar to intentions, personal policies provide a filter on admissible options -- including other policies, relatively limited options, and extended plans of action (Bratman, 1989).

Bratman did not include a computer program reflecting his concepts, but it can be observed that his theory of human decision making fits nicely into Artificial Intelligence approaches to planning with constraints.

## Other Models of Human Decision Making

Klein and Zsombok (1991) discuss other models of skilled decision making. One interesting approach used by humans in trying to decide between alternative hypotheses to explain prior events is to use their domain knowledge to decide what is consistent with their prior experience. This is a form of explanation-based decision making.

Another approach (Coward, 1990), not yet developed into a working computer program, takes the neural net operation of the brain as a basis for developing a general theory of human behavior.

#### IV. Computer Models that Ignore Motivations

##### ACT\*

Act\* (Anderson, 1983; Newell, 1990), illustrated in Figure 5, is a theory of the human cognitive architecture. It focuses on the memory and processing structures that form the base of human task performances. It has a declarative memory in the form of a semantic net. It has three basic data types (for sequences, objects, and images) which are linked together by associations. Long term procedural memory is in the form of productions. The "If" part of the productions are to be matched with information in declarative memory. The "Then" part of the productions (when fired) produces new nodes or associations in declarative memory.

The nodes of the long-term declarative memory each have a degree of activation. The working memory is that part of long-term memory that is highly activated. Using algorithms, the spreading (and the rate of spreading) of the activation is automatically determined.

Act\* has a theory of learning. Declarative learning is implemented by highly activating the new nodes and associations created by productions, but giving them only a probabilistic chance of becoming permanent. Cognitive skill acquisition is implemented by forming new productions by chaining together productions that fire in the accomplishment of a task. The strengths that determine the likelihood that productions will fire (when their actuating conditions are satisfied) increase or decrease depending on their frequency of use.

ACT\* is activation governed. ACT\* is specified by (1) a set of rules about its structure and the way actions occur, (2) a set of equations for the flow of activation, and (3) a set of equations governing the fluctuations of production strengths.

Though ACT\* was designed to simulate human performance, most of its simulation is governed by the equations chosen to represent its dynamics and its input and output functions, rather than resulting directly from its architecture.

## Soar

Soar (Laird et al., 1987; Newell, 1990) is the best known and most advanced of the computer systems designed to be capable of general intelligence. The architecture of Soar is predicated to be capable of working on the entire range of cognitive tasks, from highly routine to extremely difficult open-ended problems. As such, it is intended to be able to employ the full range of problem solving methods and representations required for these tasks and to be able to learn about these tasks as part of its problem solving efforts. Soar has been successfully demonstrated on the weak ((essentially heuristic search) methods as well as on a number of well known knowledge-intensive expert systems, such as R1. Soar has not thus far been explored with affects and emotions.

Soar takes the approach that all tasks are formulated as heuristic search. Thus, every task of attaining a goal is formulated as finding a desired state in a problem space (a space with a set of operators that can be applied to a current state to yield a new state). If due to inconsistent or incomplete immediate knowledge, an impasse is reached in selecting an operator in the current problem space, a new goal is generated to resolve the problem. Soar then continues at some initial state in a new problem space associated with this subgoal. This property, to set up a subgoal to resolve any problematic decision, is referred to as "universal subgoaling."

To use Soar for a problem, an associated knowledge base in the form of production rules must first be developed. This knowledge base consists of domain-specific task-implementation knowledge of the structure of the problem, the problem spaces that are needed, and operators that can transform one state to the next. Though, Soar has default search control knowledge (look-ahead), efficient search requires that domain-specific search-control knowledge (heuristics) also be included in the knowledge base.

In Soar, task implementation generates (or retrieves) new problem spaces, states and operators; and then search control selects among the alternatives generated. Other functions needed to form a

complete system -- goal generation, goal selection, goal termination, memory management and learning ("chunking") -- are performed by the architecture of the system.

Soar employs a frame-like structure called a context. This consists of four slots: the current goal, the current problem space, the current state, and the operator to be applied to change to the next state. These slots become bound to objects that fill them as the problem solving proceeds.

Figure 6 presents the architectural structure of Soar. Working memory consists of a context stack, a set of objects linked to the context stack, and preferences that help in selecting among the applicable operators.

Soar's problem solving approach can be depicted as a sequence of decision cycles. At any cycle, processing begins with an elaboration phase which consists of a single firing of all productions that are applicable based on the contents of working memory. This elaboration phase adds new objects, augmentations of existing objects, and preferences. The decision procedure is executed when the elaboration phase reaches quiescence. It determines which slot in the context stack should have its contents replaced and with which object. This procedure proceeds by processing the context stack from the oldest to the newest. Making a change to a higher order slot makes processing of lower slots irrelevant. Figure 7 is a simplified flow diagram of Soar's overall problem solving actions. The actual system is slightly more complex to promote greater processing efficiency. Thus, each time a decision is made to change a context element, the entire goal stack is reexamined to see if as a result of the change, progress can be made at a higher level. If so, all the current work at lower levels is discarded and the system processing proceeds from this higher level.

Soar uses a single learning mechanism, referred to as "chunking." This replaces the process of finding a solution to a subgoal with a single production (a chunk). This chunk is then generalized so that given the same type of relevant conditions, the chunk can be applied to directly obtain the subgoal solution.

Chunking is based on a dependency analysis of traces of the productions fired while finding the subgoal solution. Generalization is achieved by replacing the identifiers in the working memory

elements with variables, and removing conditions from the chunk that were not used to fire the relevant productions. The production conditions are then ordered to make the conditions match faster.

Though Soar appears capable of being developed to cover virtually the full range of cognitive activity, it currently lacks motivations and has only a limited resemblance to the human problem solving architecture and to normal human decision making.

## V. Computer Models that Include Motivations

There are very few computer models that include motivations and/or emotions. Virtually all that do include only a very specific aspect of human behavior. Thus, we have only found that DAYDREAMER (Mueller, 1990) includes the full range of internal goals, motivations, emotions, control of memory, planning, and internal and external behavior in response to the internal and external environment. A few of the significant motivation-oriented computer models are described below.

### Hot Cognition

Thagard and Kunda (1987, p. 753) observe that "People make motivated inferences when their conclusions are biased by their general motives or goals." Thagard and Kunda have written a Common Lisp computer program of motivated inference referred to as "Hot Cognition." This computational model is designed to account for a variety of phenomena that have been investigated empirically:

1. Motivated changes of self-conceptions.  
How people see themselves may be influenced by how they would like to see themselves.
2. Motivated changes of theories about the world.  
People tend to generate those theories about the causal determinants of events that are most likely to support their goals.
3. Motivated changes of inferential rules.  
Individuals threatened by some evidence are less likely to believe it.

4. Motivated changes of goals.

When people realize that they are less likely to obtain their goals, they diminish the importance of the goals to the self, thereby maintaining positive self-evaluation (the "fox and the grapes" phenomenon).

The model includes the following components:

1. A representation of the self.

This includes the individual's fundamental motives (such as staying healthy) and their priority and activation, and attributes of the self (such as smokes) together with their importance and activation.

2. A mechanism for evaluating the relevance of a potential conclusion to the motives of the self.

3. Mechanisms for motivated search.

4. Mechanisms for adjusting the parameters of the inference rules to bias them to ensure that inferences favorable to the self are more likely to be made.

The interaction of the elements in this model, in inferring whether to accept the stated conclusion in response to a proposition (e.g., "Smokers die young."), is shown in Figure 8.

### MoCog1

MoCog1 (Gevarter, 1991) is the first of a NASA series of computer programs that include motivations in attempting to develop successively increasingly sophisticated models for simulating human cognitive behavior. MoCog1 (for Motivated Cognition 1) differs from the models previously considered by not only focussing on motivations, but by seeking to correlate human decision making with what is known about information processing in the brain.

Most human decision-making is of an experience-based, relatively straight-forward, largely automatic response to internal goals and drives, utilizing cues and opportunities perceived from the current environment. MoCog1 limits itself to this type of decision making, which can be recognized as being similar to Klein's (1989) Recognition-Primed Decision Making, reviewed earlier. The major

differences are 1) that motivations have been included in MoCog1, and 2) that an attempt has been made to utilize, in the development of the model, knowledge about information flow in the human brain.

The central ideas underlying the MoCog1 model are:

1. In the brain, stored along with each experience are the emotions that were present at the initiation of the experience and those that resulted from the experience. The affect pattern thus associated with the pre-conditions and post-conditions of the experience are accessible during future interactions.
2. Affect vectors are formed by combining existing internal affect states with those elicited by information processing of the various levels of the situational response (from affects involving the lowest level of implicit motivations, to affects at the highest level resulting from cognitive involvement with the "self").
3. In the MoCog1 version of decision-primed decision making, situation assessment (and accompanying affects), and access to applicable procedures are accomplished by automatic associative recall of stored experience.
4. The current affect state and the expected affect states, resulting from the automatic assessment of the current event, act as inputs to the brain's control mechanism, which generates needs and goals to move the anticipated resultant affect state to a more desirable condition. These needs and the current context elicit applicable stored procedures. The predicted results and affiliated affect patterns (associated with the various applicable procedures) are then fed to the brain's decision making mechanism. This mechanism then automatically seeks to select procedures that would produce the most desirable overall satisfaction of the generated needs, considering the weights or priorities given each affect and their current degree of activation.

Figure 9 presents the resultant flow diagram.

The resultant computer program was employed successfully in simulating Dweck and Leggett's (1988) findings that relate how an individual's implicit theories orient them toward particular general goals, with resultant cognitions, affects and behavior in response to their environment.

## **DAYDREAMER**

DAYDREAMER (Mueller, 1990) is a computer program that simulates the emotional control of an individual's train of thought, as exemplified in human daydreaming. Figure 10 is a simplified flow diagram of the DAYDREAMER program.

### **Goals in DAYDREAMER**

The DAYDREAMER individual has Personal Goals common to all humans. These are composed of

Cyclic Goals (which require repeated satisfaction):

- Food
- Sex
- Love-Giving
- Love-Receiving
- Companionship
- Entertainment
- Money
- Possessions.

Achievement Goals (corresponding to physical or mental states to be achieved or maintained):

- Self-Esteem
- Social Esteem
- Lovers
- Friends (and Positive Relations)
- Employment.

In addition to Personal Goals, there are Domain Goals associated with activity currently salient. For DAYDREAMER, these are taken to be:

Emotional Daydreaming Goals (to generate daydream scenarios to recover emotionally from a personal goal failure):

- Rationalization
- Roving (shifting attention away from the failure)
- Revenge.

Learning Daydreaming Goals [either associated with learning from past (or possible future) goal failures or with the the generation of possible future scenarios for achieving an active personal goal]:

- Reversal

Recovery  
 Repercussions  
 Rehearsal.

### Emotions

Associated with each goal success or failure are specific emotions, and more general emotions such as: pleasure and relief, and displeasure and regret. Relief is a positive emotion resulting from an imagined alternative past scenario in which an imagined action was taken that resulted in a goal failure. Regret results from an imagined past action (that could have been taken) that would have resulted in a goal success.

Positive DAYDREAMER emotions (resulting from succeeded goals) are:

pleasure			
hope			
gratitude			
amusement	from	ENTERTAINMENT	goal
satiation	"	FOOD	"
pride	"	SELF-ESTEEM	"
poise	"	SOCIAL ESTEEM	"

Negative DAYDREAMER emotions (resulting from failed goals) are:

displeasure			
worry			
anger			
shame	from	SELF-ESTEEM	goal
embarrassment	"	SOCIAL ESTEEM	"
humiliation	"	SOCIAL ESTEEM	"
rejection	"	POSITIVE RELATION	"
heartbreak	"	LOVERS	"

### Basic Information Flow

Referring to Figure 10, we can now follow the basic information flow in DAYDREAMER. All goals have their priorities. When the satisfaction of a personal goal falls below a critical level, a concern is actuated to achieve the goal. Associated with each concern is an emotion whose activation level depends on the goal priority. Several concerns may be active at the same time. It is predicated in

DAYDREAMER that the individual automatically selects the concern to concentrate on that has the highest emotional activation level. This selection is referred to as "Emotion Driven Control." DAYDREAMER then develops a train of thought, or daydream, to achieve that goal. The process by which DAYDREAMER does this is referred to as planning. In response to the daydream's success or failure in achieving its goal an associated response emotion results. If the resultant emotion is negative, DAYDREAMER goals are actuated to seek a more positive emotional state. The priority of these new goals is related to the emotions associated with the failed goal.

Figure 11 is an expanded version of Figure 10. For the selected concern, the planning is done one step at a time. At the end of that planning step, the system is checked to see if the concern has succeeded or failed. If the concluded concern has resulted in a positive emotion, the associated plan which has been generated is generalized and put into episode storage for future use. If another active concern, B (not currently being worked on), has fortuitously succeeded while working on the selected concern, A, a positive response emotion is added to A in proportion to the Emotion associated with B.

If, fortuitously while taking a step to achieve A, a subgoal of B succeeds or a plan is found to achieve B, a positive response emotion is added to B proportional to the emotional level of B.

If during the planning step, several planning rules unify with the current subgoal, each rule is applied in its own separate world model (planning context), giving rise to alternative states of a hypothetical world.

At the end of the planning step, the system checks the modified emotion levels and switches to a new concern if its emotion level warrants. Having completed the current planning step, DAYDREAMER then performs the next planning step. Theoretically, the system can run forever.

### The DAYDREAMER Planner

DAYDREAMER uses the process of planning to generate imaginary sequences of events (daydreams) to try to satisfy its selected concerns. Figure 12 is a simplified diagram of the DAYDREAMER Planner. The associated planner functions of Analogical Rule

Application, Reminding, Serendipity, and Mutation are diagrammed in Figures 13, 14, 15, and 16, respectively.

DAYDREAMER has been demonstrated with "constructed" daydreaming examples, both daydreaming trains of thought and trains of thought in response to constructed external inputs.

## VI. Computer Simulations of Humans Interacting with a Simulated World

### MIDAS

MIDAS (A<sup>3</sup>I Executive Summary, 1990) is a simulation of an aircraft pilot carrying out an assigned mission. The purpose of the simulation is to 1) develop criteria for cockpit design, and 2) to determine the associated pilot training requirements. Both cognitive and spatial/temporal aspects of the pilot are simulated. This includes physical motion responses, achieved by anthropomorphic modeling, and perceptual responses associated with focussing attention and sensory processing delays.

As insufficient human operator performance data is available, initially normative, rather than descriptive, models have been chosen. Figure 17 is a very simplified MIDAS flow diagram. The following quote from The A<sup>3</sup>I Executive Summary (1990, p. 13) allows us to follow through the flow diagram.

During a simulation, the model attempts to execute assigned mission activities subject to specified constraints, state variables, and other simulation object requirements. This model accomplishes this action by:

- 1) updating the simulated operator's goal list to delete terminated or inappropriate goals,
- 2) examining equipment and world state variables to determine if event-response activities are required,
- 3) tracing the decomposition of mission goals to their lowest level, finding matching equipment operation patterns or activities which will satisfy them,

- 4) sorting these matched goal-activity patterns by priority,
- 5) interacting with the scheduling and task loading model components as appropriate, and
- 6) executing these activities subject to physical resource (hand, eye, etc.) requirements, Visual, Auditory, Cognitive and Motor load values, as well as temporal/logical constraints.

MIDAS uses a constraint-based opportunistic model of operator scheduling behavior. The modular design provides the opportunity to try different scheduling strategies. The operator task loading model classifies individual tasks in terms of their demands on the Visual, Auditory, Cognitive and Motor processing dimensions, based on the attributes of the mission tasks, world state, operator and crew station equipment.

The actual simulation has been chosen to be object-oriented. Each module of the simulation is modeled as a separate object. These modules communicate by message passing. Much of the communication in such a system is in the query/response (Q/R) form. Figure 18 provides a more accurate detailed diagram of the MIDAS simulation. Corker et al. (1990) provide a more complete description of this type of simulation, (including simulation of short and long term memory, and the process of forgetting).

## VII. Discussion

Only a few of the models reviewed include the modeling of human motivations, rather than just cognition. Hot Cognition, MoCog1, and DAYDREAMER particularly strive to include motivation in their simulation of humans. Hot Cognition focuses on motivated inferences; MoCog1 on the effects of emotions on human decision making; and DAYDREAMER on developing human streams of thought in response to emotions generated by goal needs, failures, or successes. Each of these three simulations employs a somewhat different theory of the role of emotions. The response to emotions in MoCog1 is to select actions that produce the most overall favorable response considering all the emotions. The actions generated in DAYDREAMER function to satisfy a single selected concern and its associated emotion. However, the DAYDREAMER focus on only satisfying a single concern at a time, appears appropriate because,

during conscious thought about immediate concerns, humans tend to think about one thing at a time. Nevertheless, Mueller (1990, p. 38) observes that "A complete system would have to have both meta-planning for long term planning [to attempt to satisfy multiple concerns simultaneously], and emotion-driven planning for short-term decisions about what to do next and what to daydream about next." In MoCog1, the decision to attempt to satisfy the total affect vector is done subconsciously, which is consistent with the parallel processing capabilities of the subconscious. The conditions under which these two types of responses are appropriate needs to be further explored. Figure 19 provides a generalization of the DAYDREAMER'S approach to the effects of personal and domain goals on emotions and resultant behavior. Figure 20 provides a rough example of a view of human decision making when the emotional control aspect of Figure 19 is combined with the GEMS Human Decision Model given in Figure 4. This greatly simplified model highlights the relationships between motivations, memory, external inputs, subconscious processing, and the limited role of attention -- moving us closer to a more realistic view of human decision making.

MIDAS is an effort to simulate humans considering both cognitive and physical aspects. If we add a motivation module to such a model, we can begin to achieve a more holistic model of humans. Models of this type could be very valuable for simulating human operator behavior, including errors, such as have contributed to aircraft mishaps and operator control problems of nuclear reactors. Such models, coupled with the use of flight simulators at NASA Ames, could be very valuable toward developing, testing, and refining theories of human operator behavior. Though additional syntheses of the findings of cognitive science and other disciplines are needed, the models of human decision making discussed in this report provide a start in developing such theories.

## Computer Model References

Anderson, John R., *The Architecture of Cognition*, Harvard University Press, Cambridge, 1983.

Bratman, Michael E., *Intention, Plans and Practical Reason*, Harvard U. Press, Cambridge, 1987.

Bratman, Michael E., "Intention and Personal Policies," *Philosophical Perspectives 3, Philosophy of Mind and Action Theory*, 1989.

Corker, Kevin M.; Cramer, Michael L.; and Henry, Eugene H., *Methodology for Evaluation of Automation Impacts on Command and Control C2 Systems*, Report No. 7242, BBN Systems and Technologies, Cambridge, July 1990.

Coward, L. A., *Pattern Thinking*, Praeger, New York, 1990.

Dweck, C. S. and Leggett, E. L., "A Social-Cognitive Approach to Motivation and Personality," *Psychological Review*, Vol. 95, No. 2, 1988, pp. 256-273.

Gevarter, William B., "MoCog1: A Computer Simulation of Recognition-Primed Human Decision Making," NASA TM, in publication.

Izard, C. E., "Emotion-Cognition Relationships and Human Development," *Emotions, Cognition and Behavior*, C. E. Izard, J. Kagan and R. B. Zajonc Eds., Cambridge, MA: Cambridge University Press, 1984, pp. 17-37.

Klein, Gary, "Recognition-Primed Decisions," *Advances in Man-Machine Systems Research*, Vol. 5, W. R. Rouse (Ed.), JAI Press, Greenwich, Conn., 1989, pp. 47-92.

Klein, Gary A., and Zsombok, Caroline E., "Models of Skilled Decision Making," *Proceeding of the 35th Annual Human Factors Society*, September 1991, San Francisco.

Laird, John E.; Newell, Allen; and Rosenbloom, Paul S., "Soar: An Architecture for General Intelligence," *Artificial Intelligence*, Vol. 33, No. 1, 1987, pp. 1-64.

Maher, John W., "Air Crew Judgement, Decision Making, Situational Awareness: A View from the Cockpit," In publication.

Mueller, Erik T, *Daydreaming in Humans and Machines*, Ablex, Norwood, NJ, 1990.

Newell, Allen, *Unified Theories of Cognition*, Harvard University Press, Cambridge, 1990.

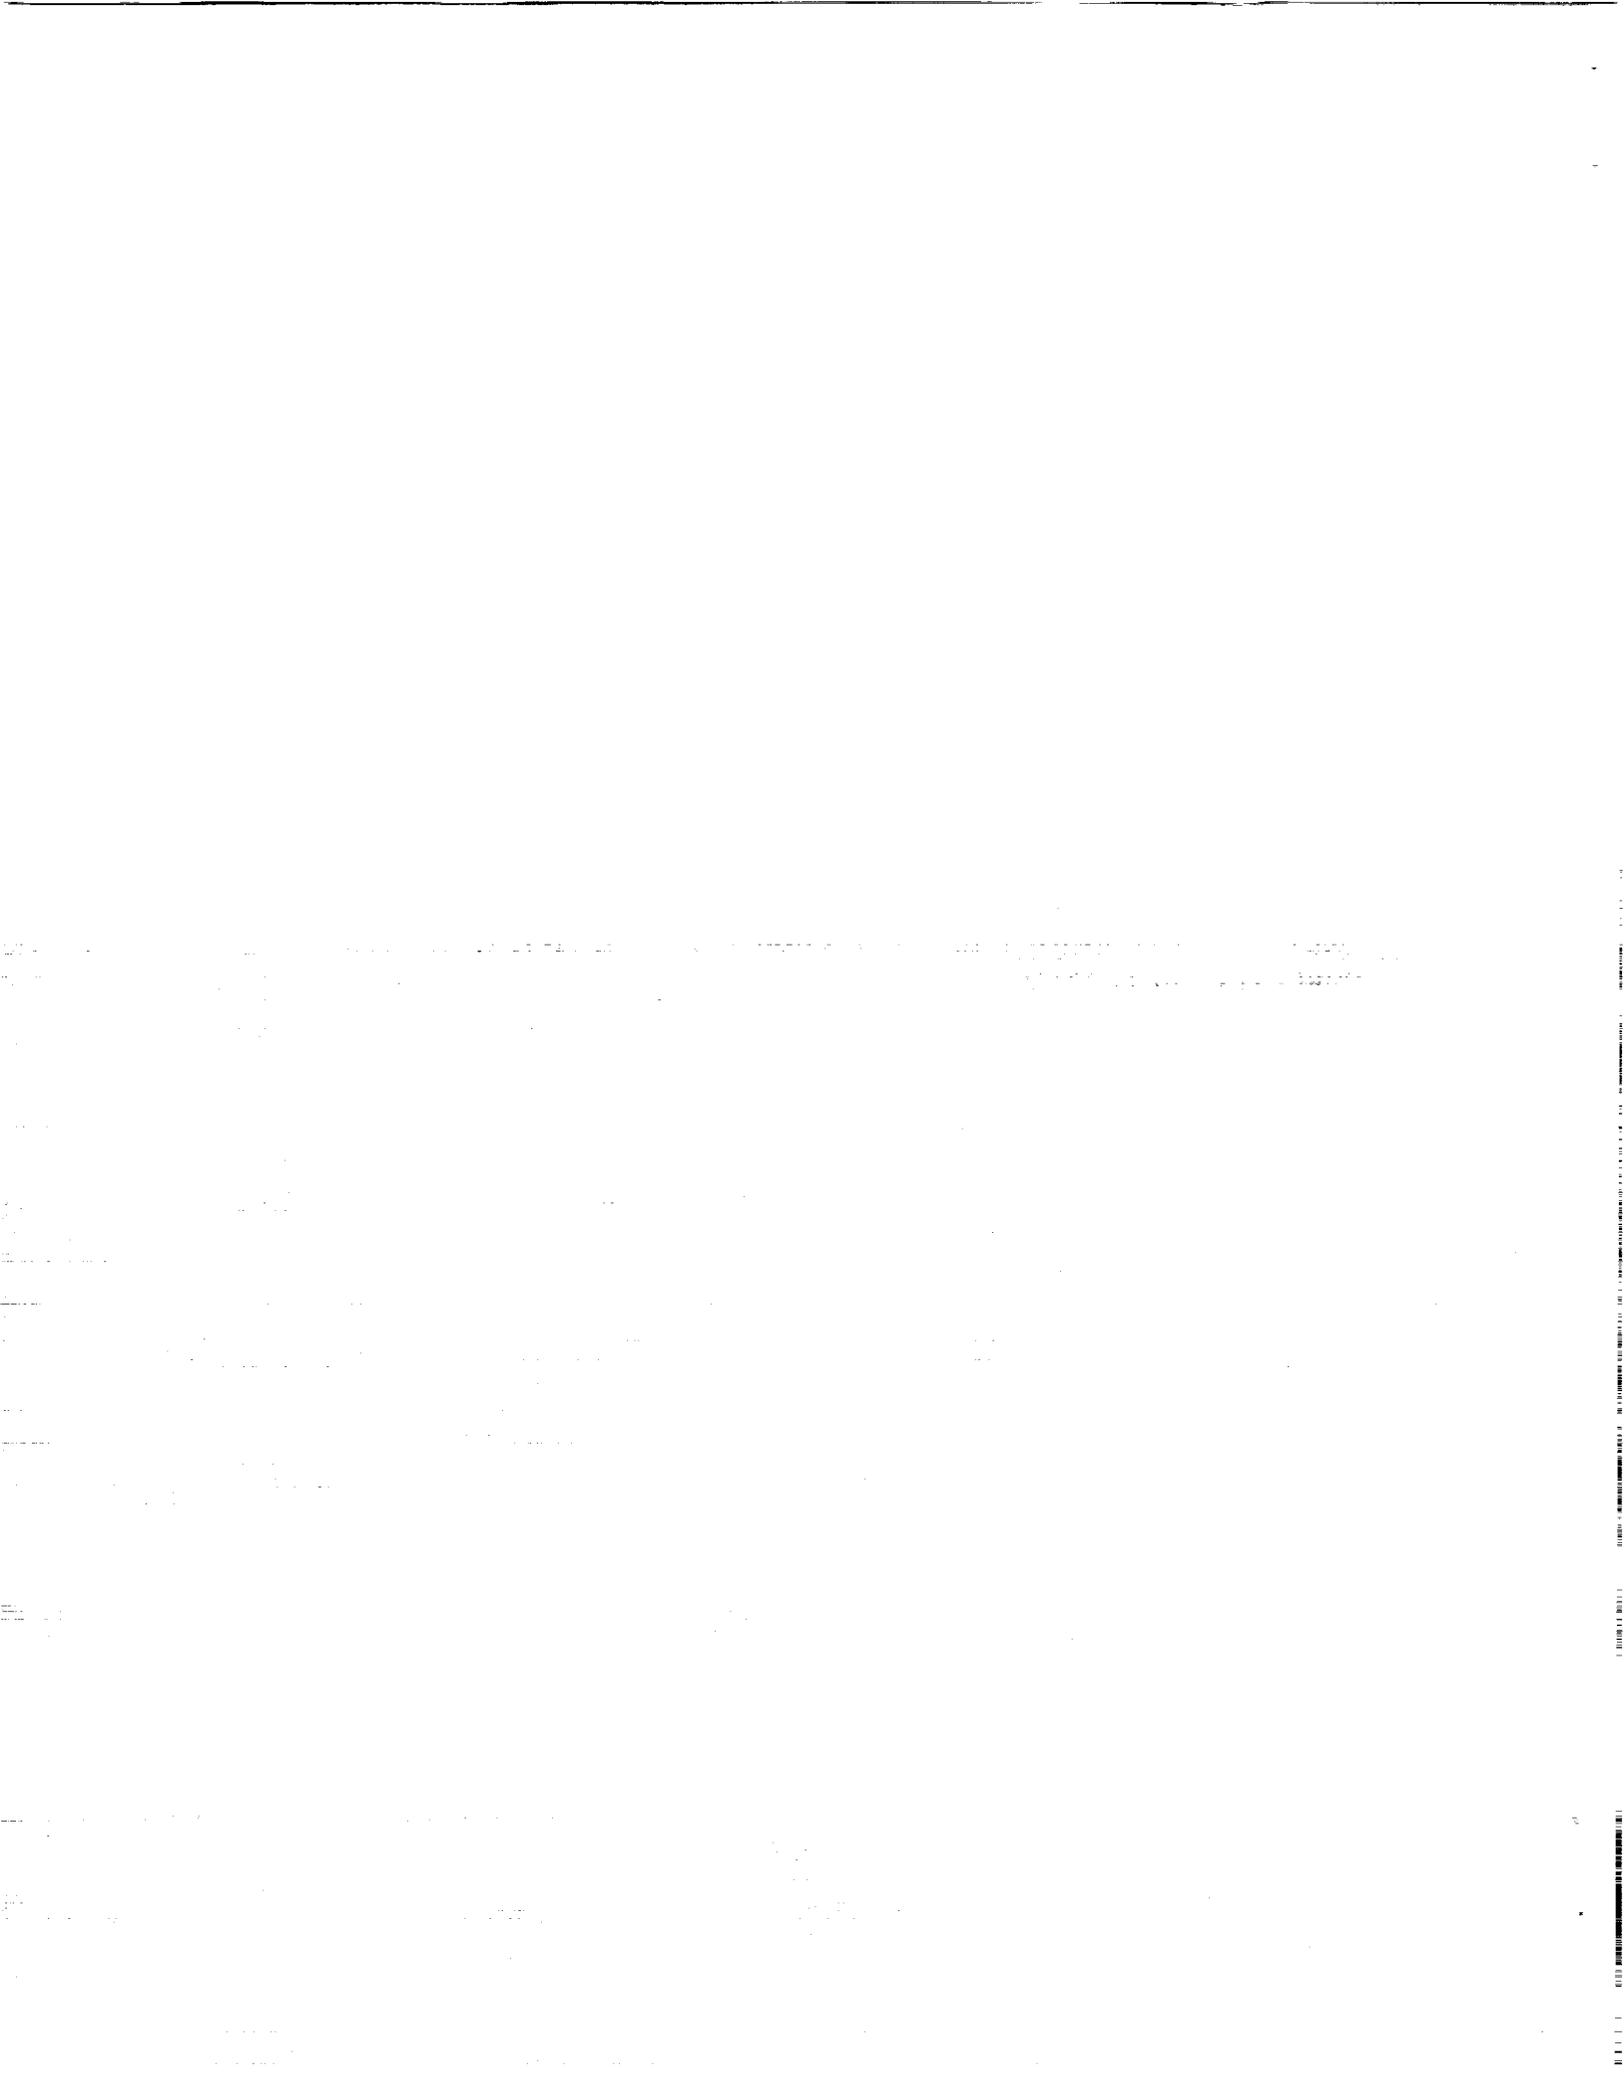
Paritsis, Nicholas C., "Man as a Hierarchical and Purposeful Intelligent System," *Systems Research*, Vol. 4, No. 3, 1987, pp. 169-176.

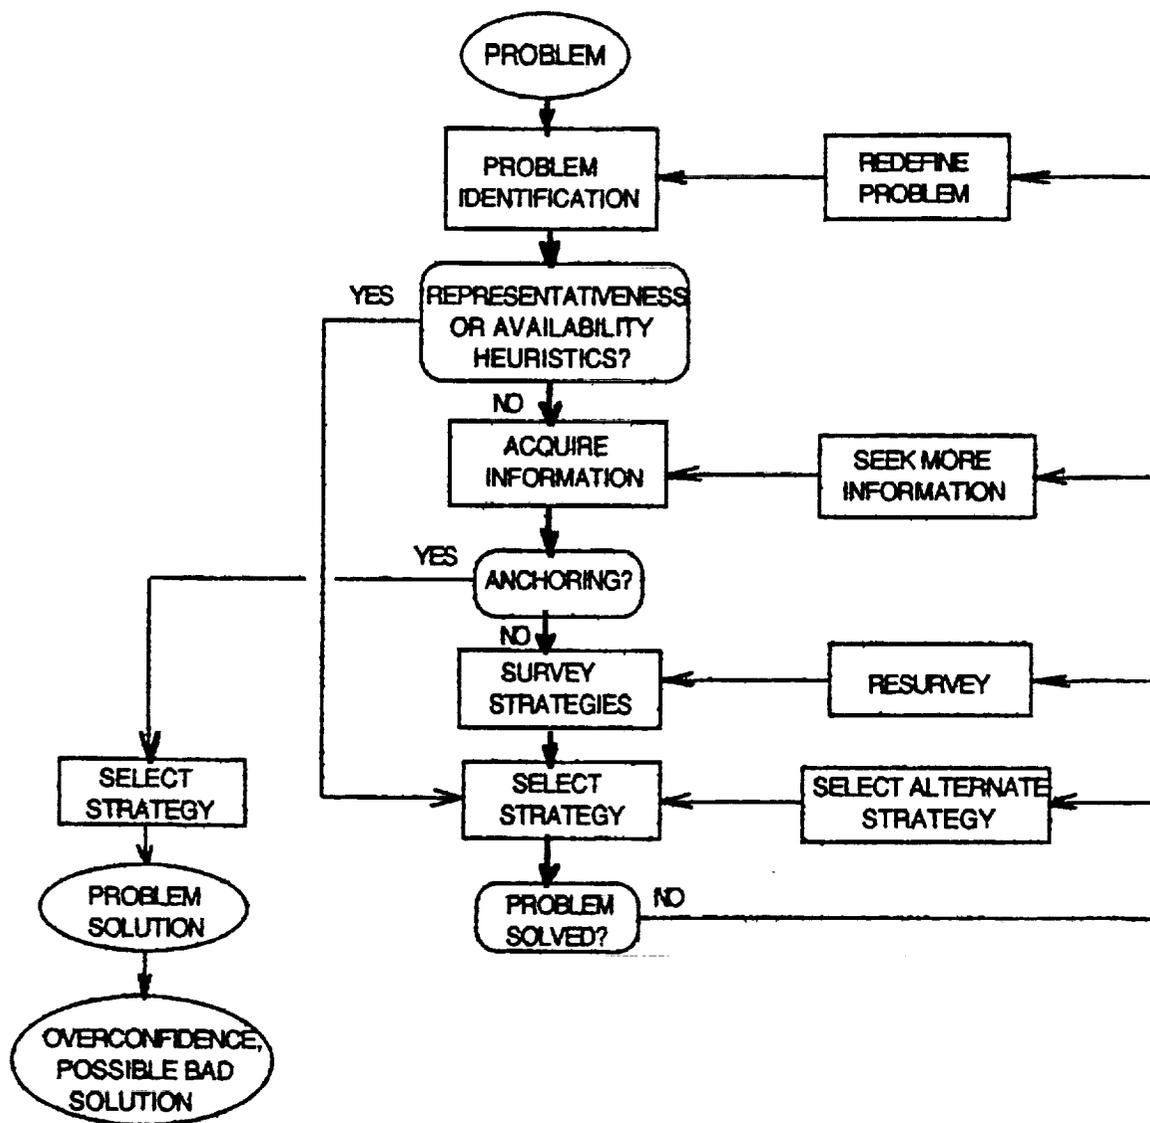
Reason, James, *Human Error*, Cambridge, MA: Cambridge University Press, 1990.

Restak, Richard M., *The Mind*, Bantam Books, Toronto, 1988.

Thagard, P.; and Kunda, Z., "Hot Cognition: Mechanisms for Motivated Inference," *Proceedings of the 9th Annual Conference of Cognitive Science*, 1987, pp. 752-763.

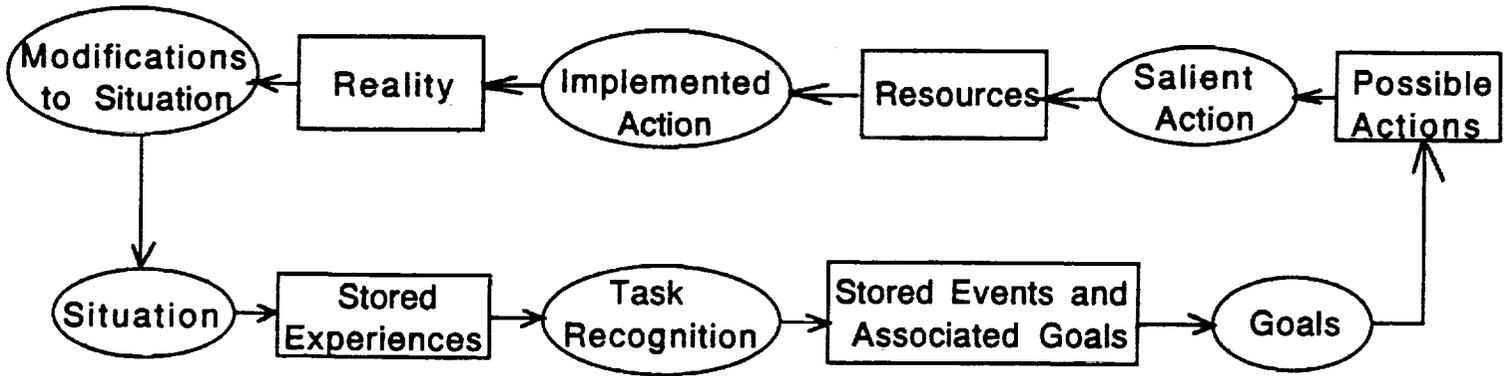
*Army-NASA Aircrew/Aircraft Integration Program: A<sup>3</sup>I*, Executive Summary, Sept. 1, 1990.





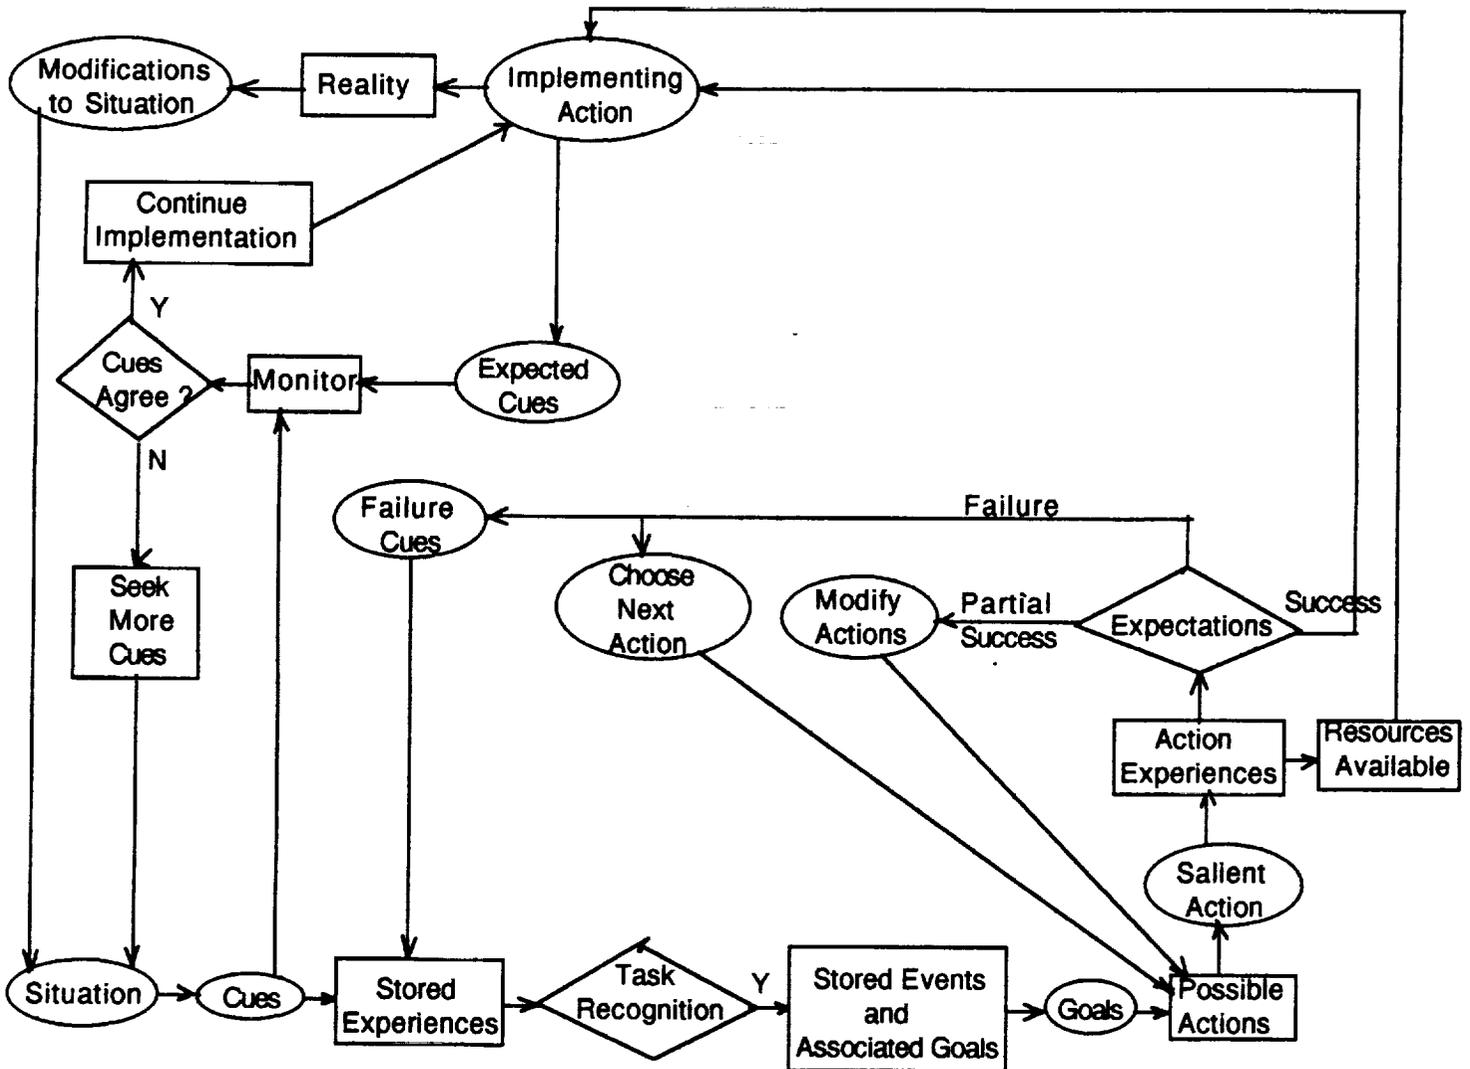
# PASS ARCHITECTURE OF AIRCREW DECISION MAKING

Figure 1



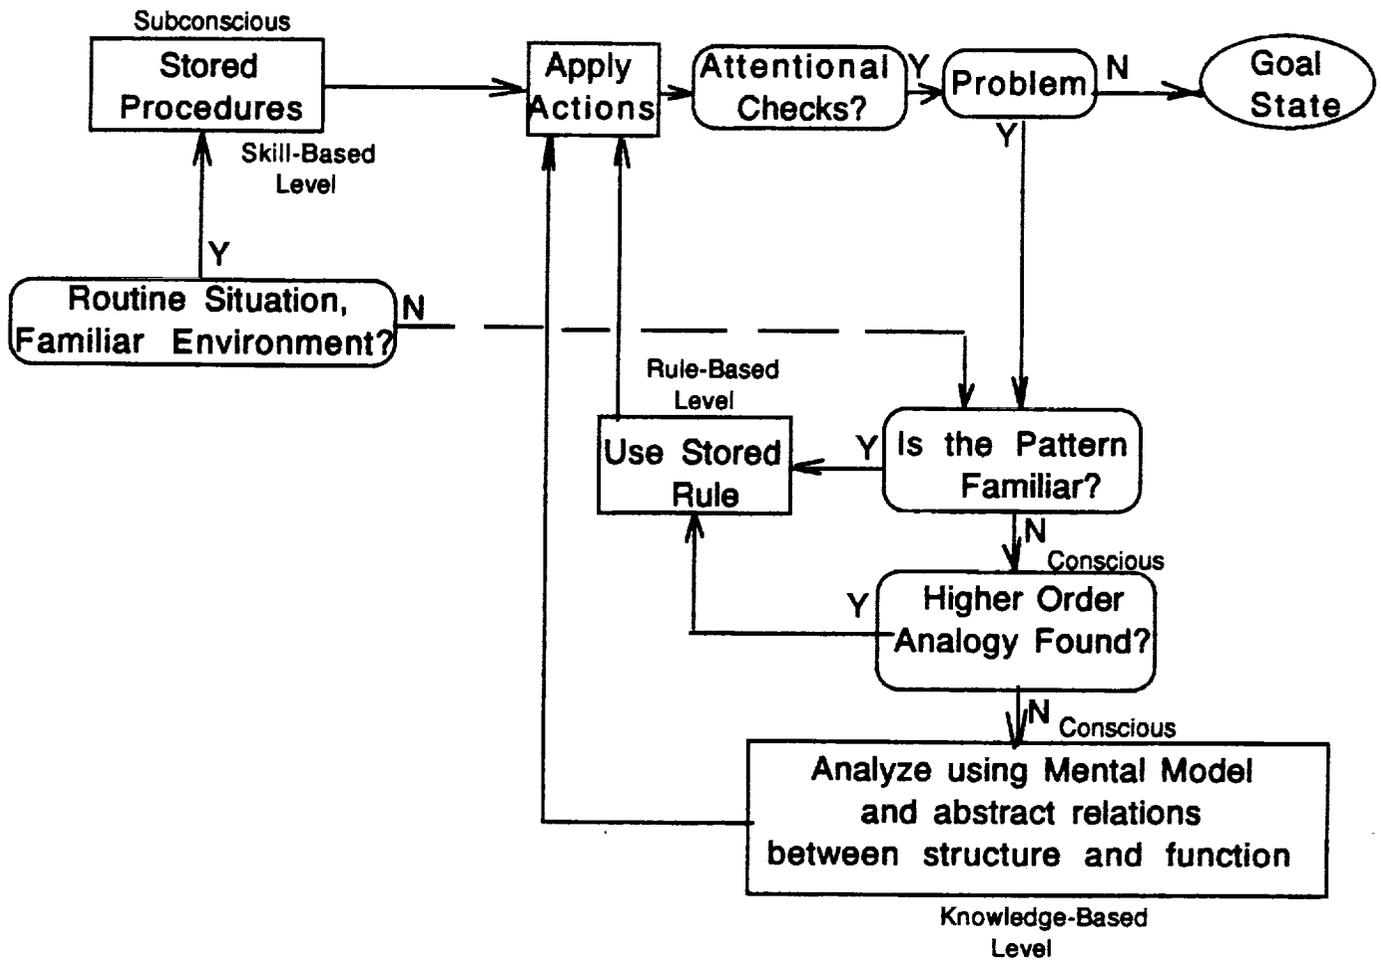
## Recognition-Primed Decision Making

Figure 2



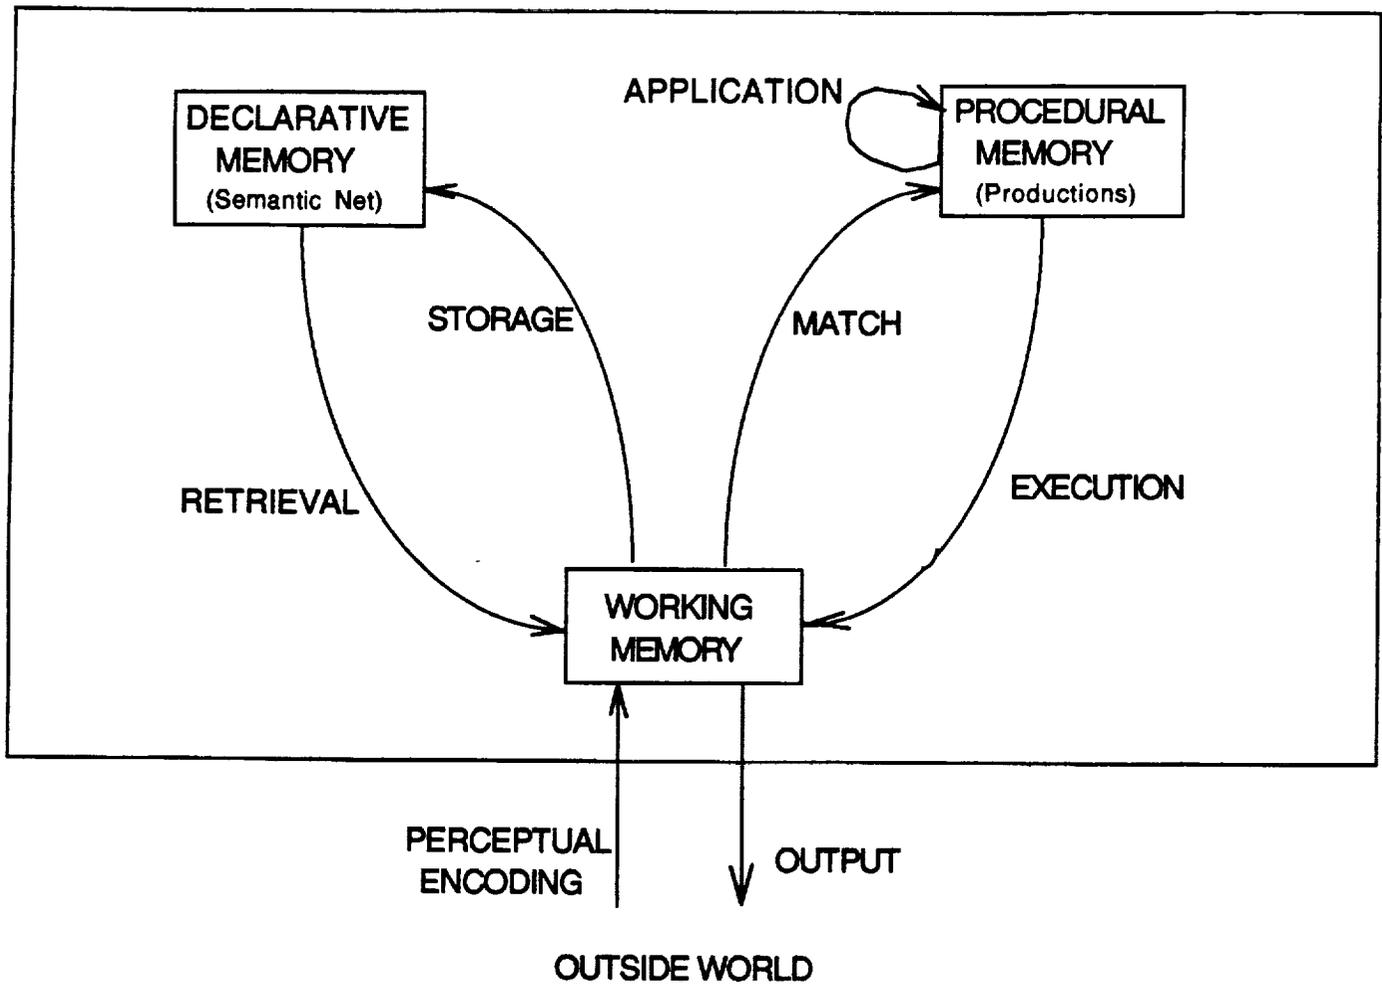
## Recognition-Primed Decision Making with Deepening

Figure 3



# GEMS

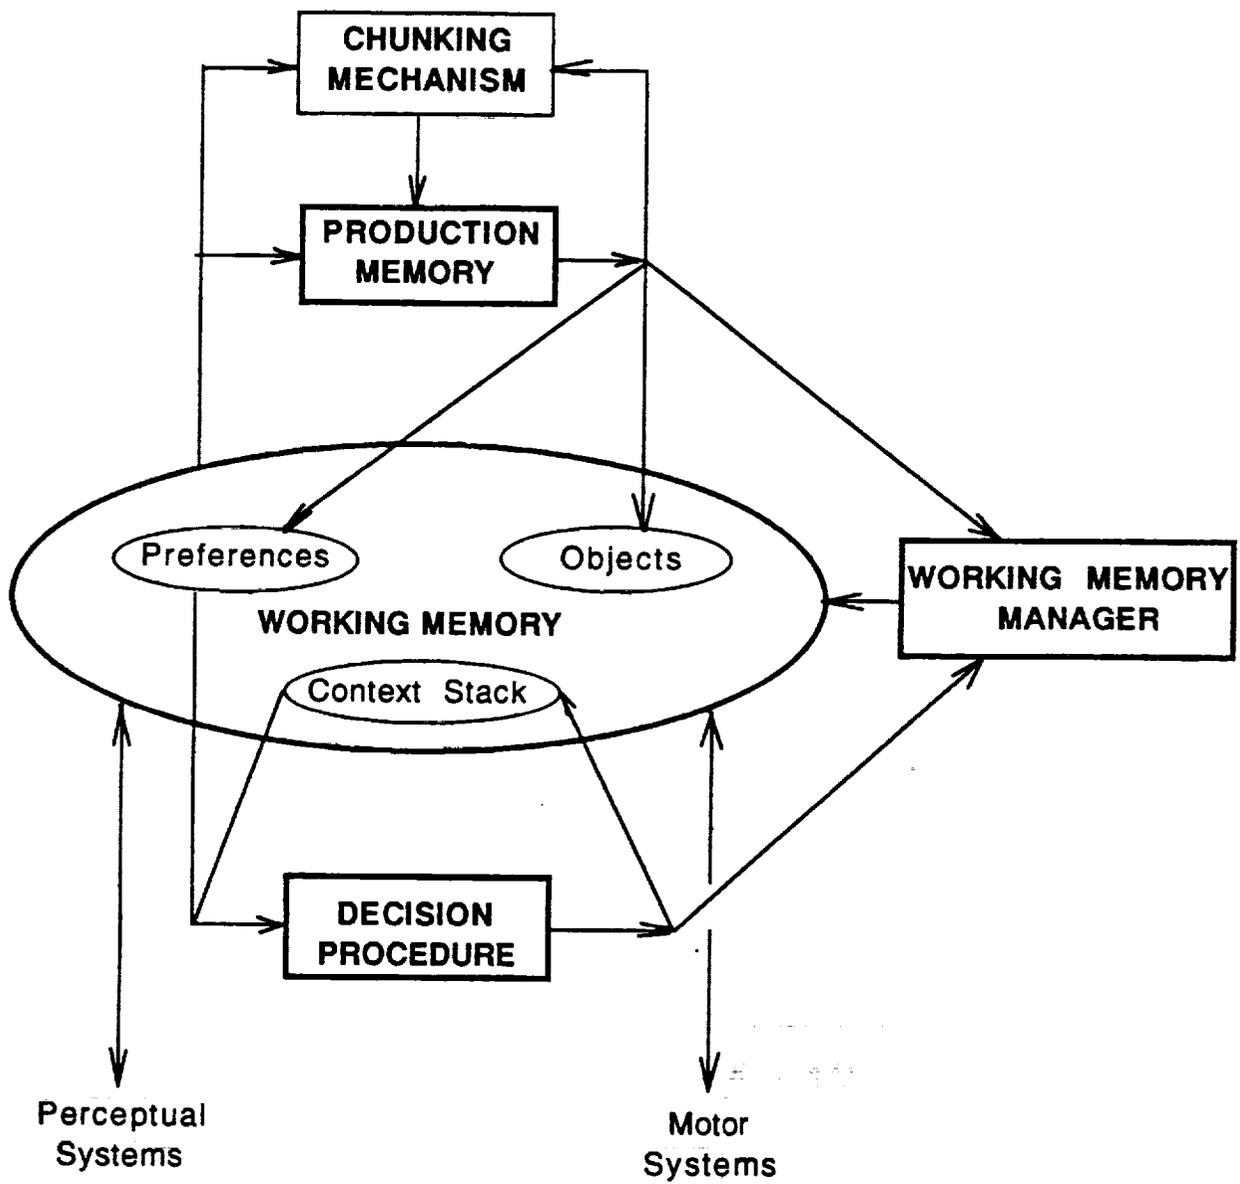
Figure 4



# THE ACT\* COGNITIVE ARCHITECTURE

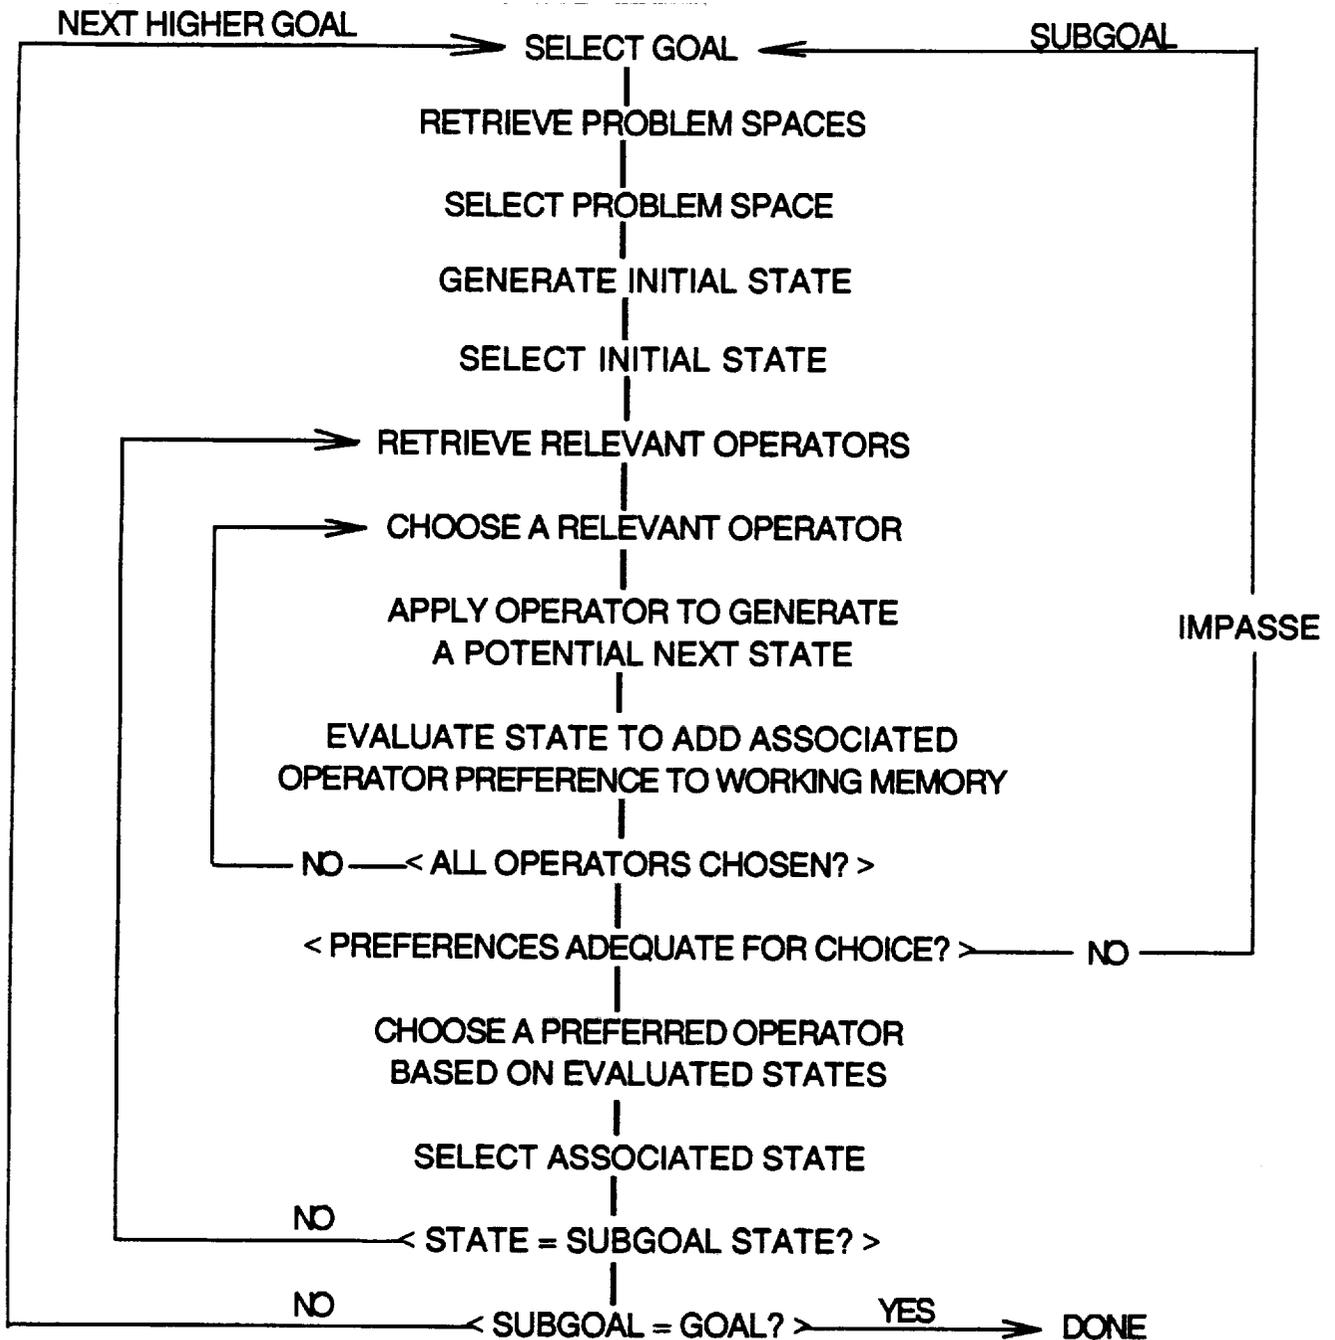
(After Anderson, 1983)

Figure 5



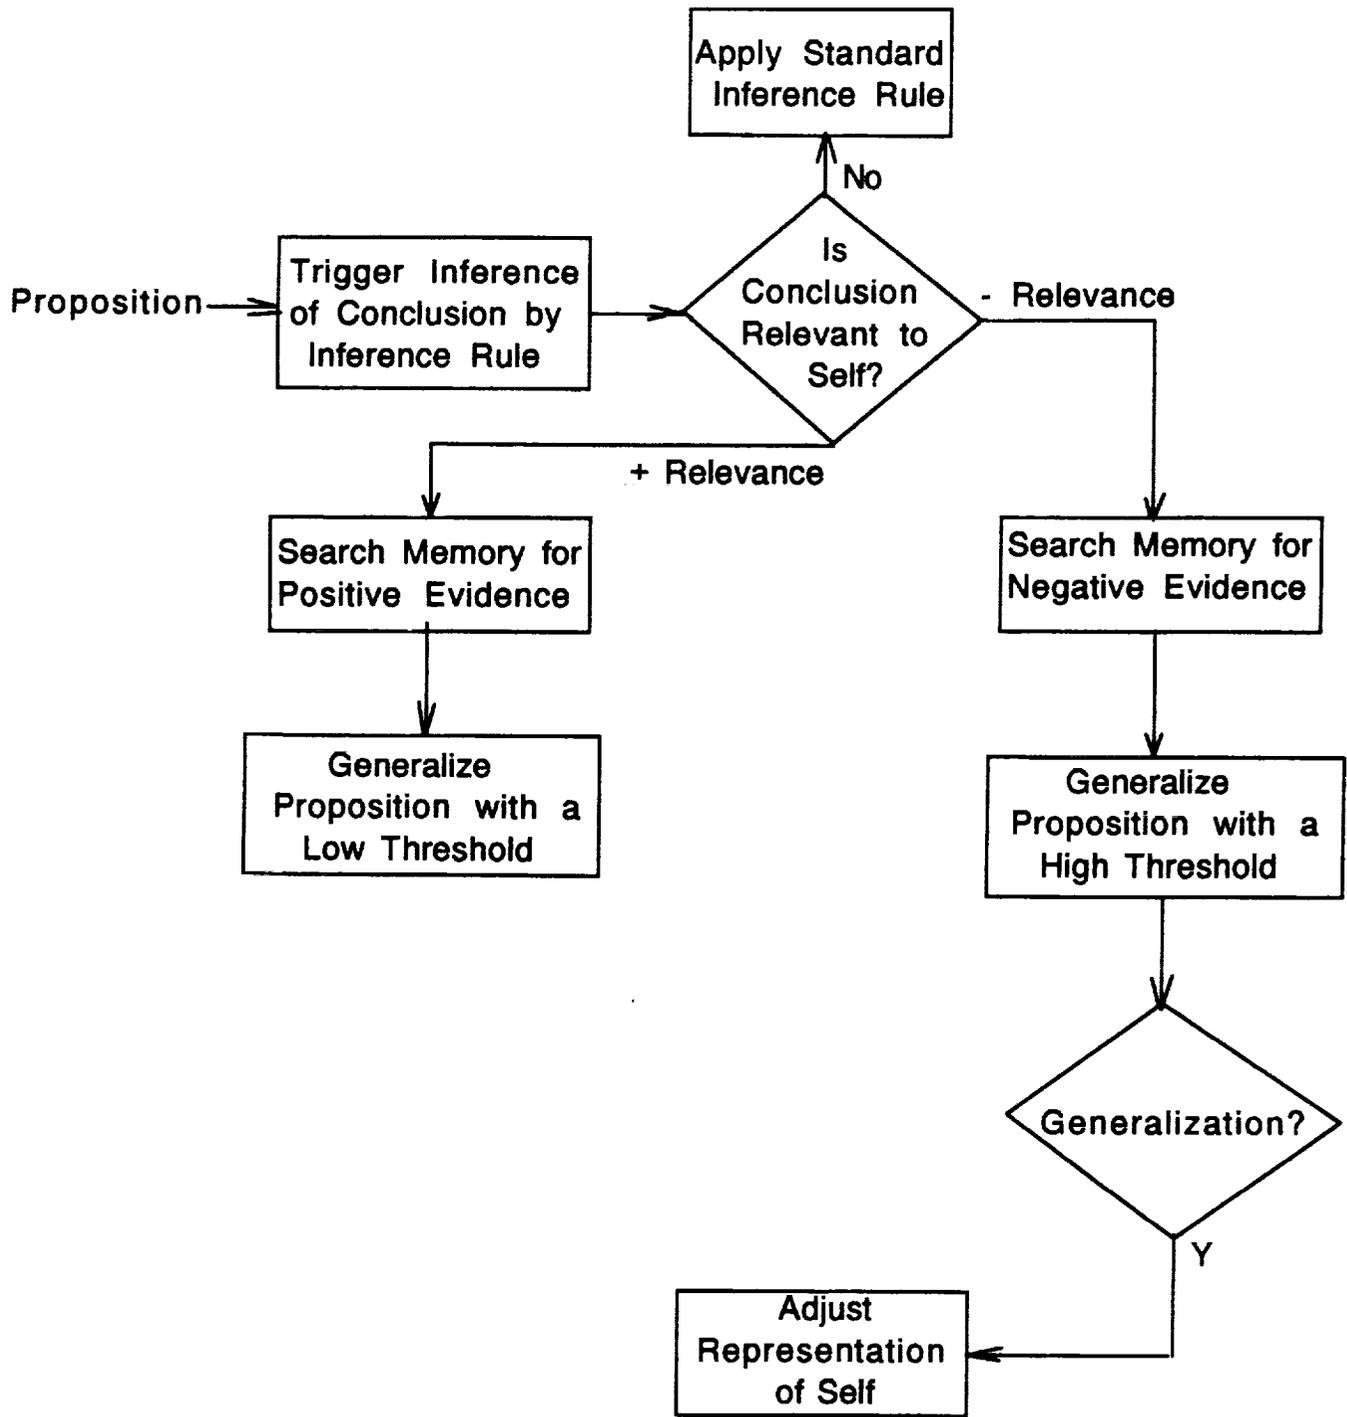
# SOAR ARCHITECTURE

FIGURE 5



## FLOW DIAGRAM OF SOAR DECISION PROCEDURE

Figure 7



# GENERAL MODEL OF MOTIVATED INFERENCE

FIGURE 8

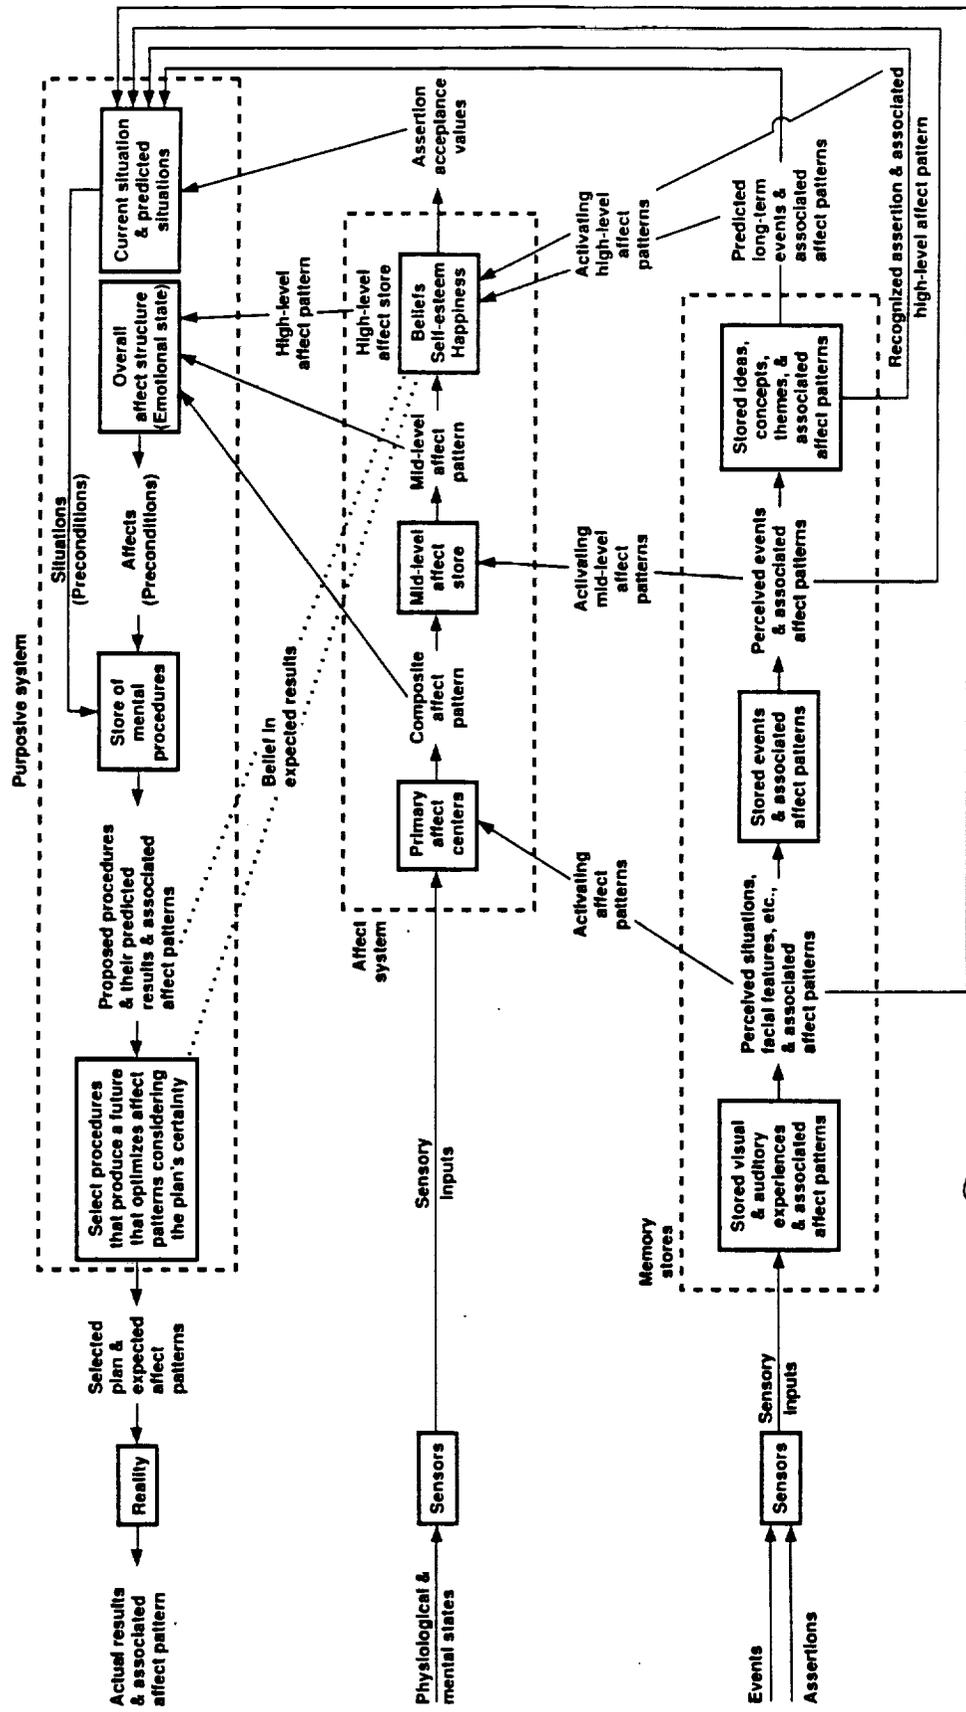
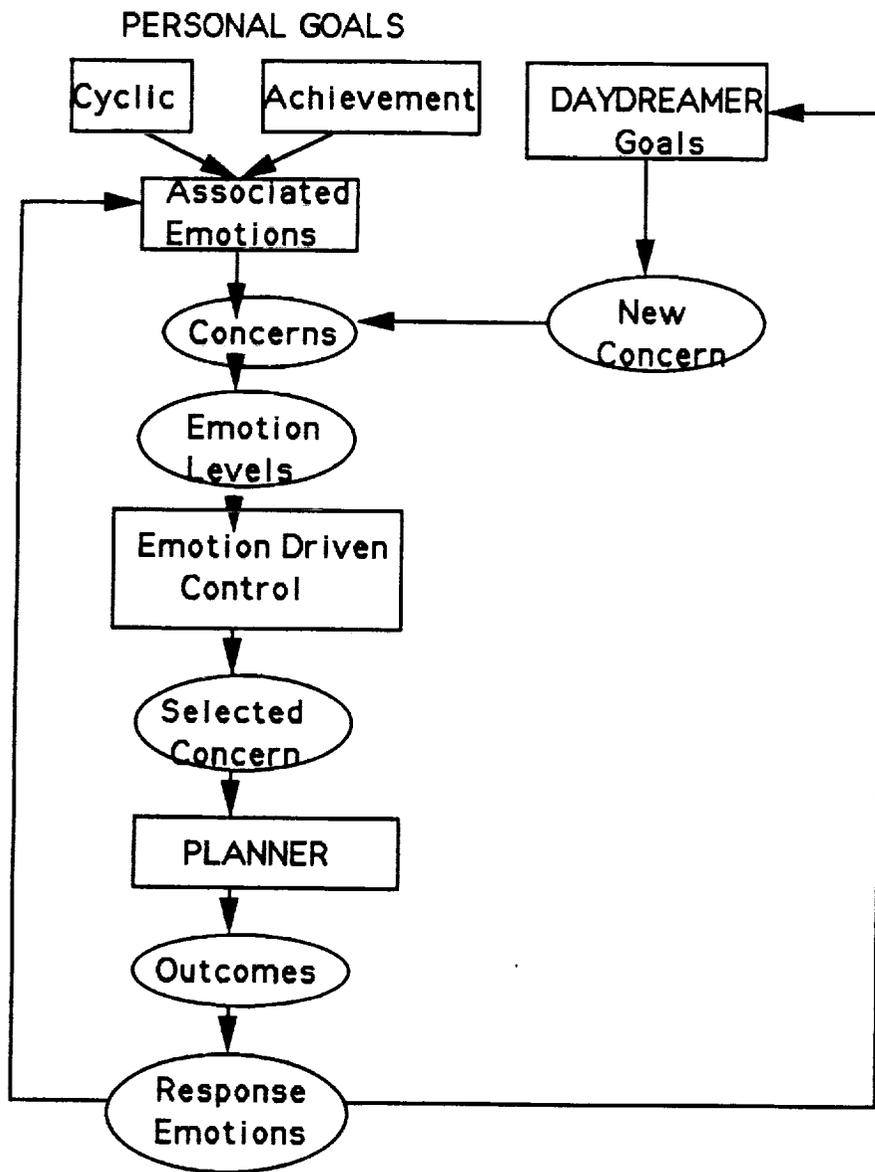
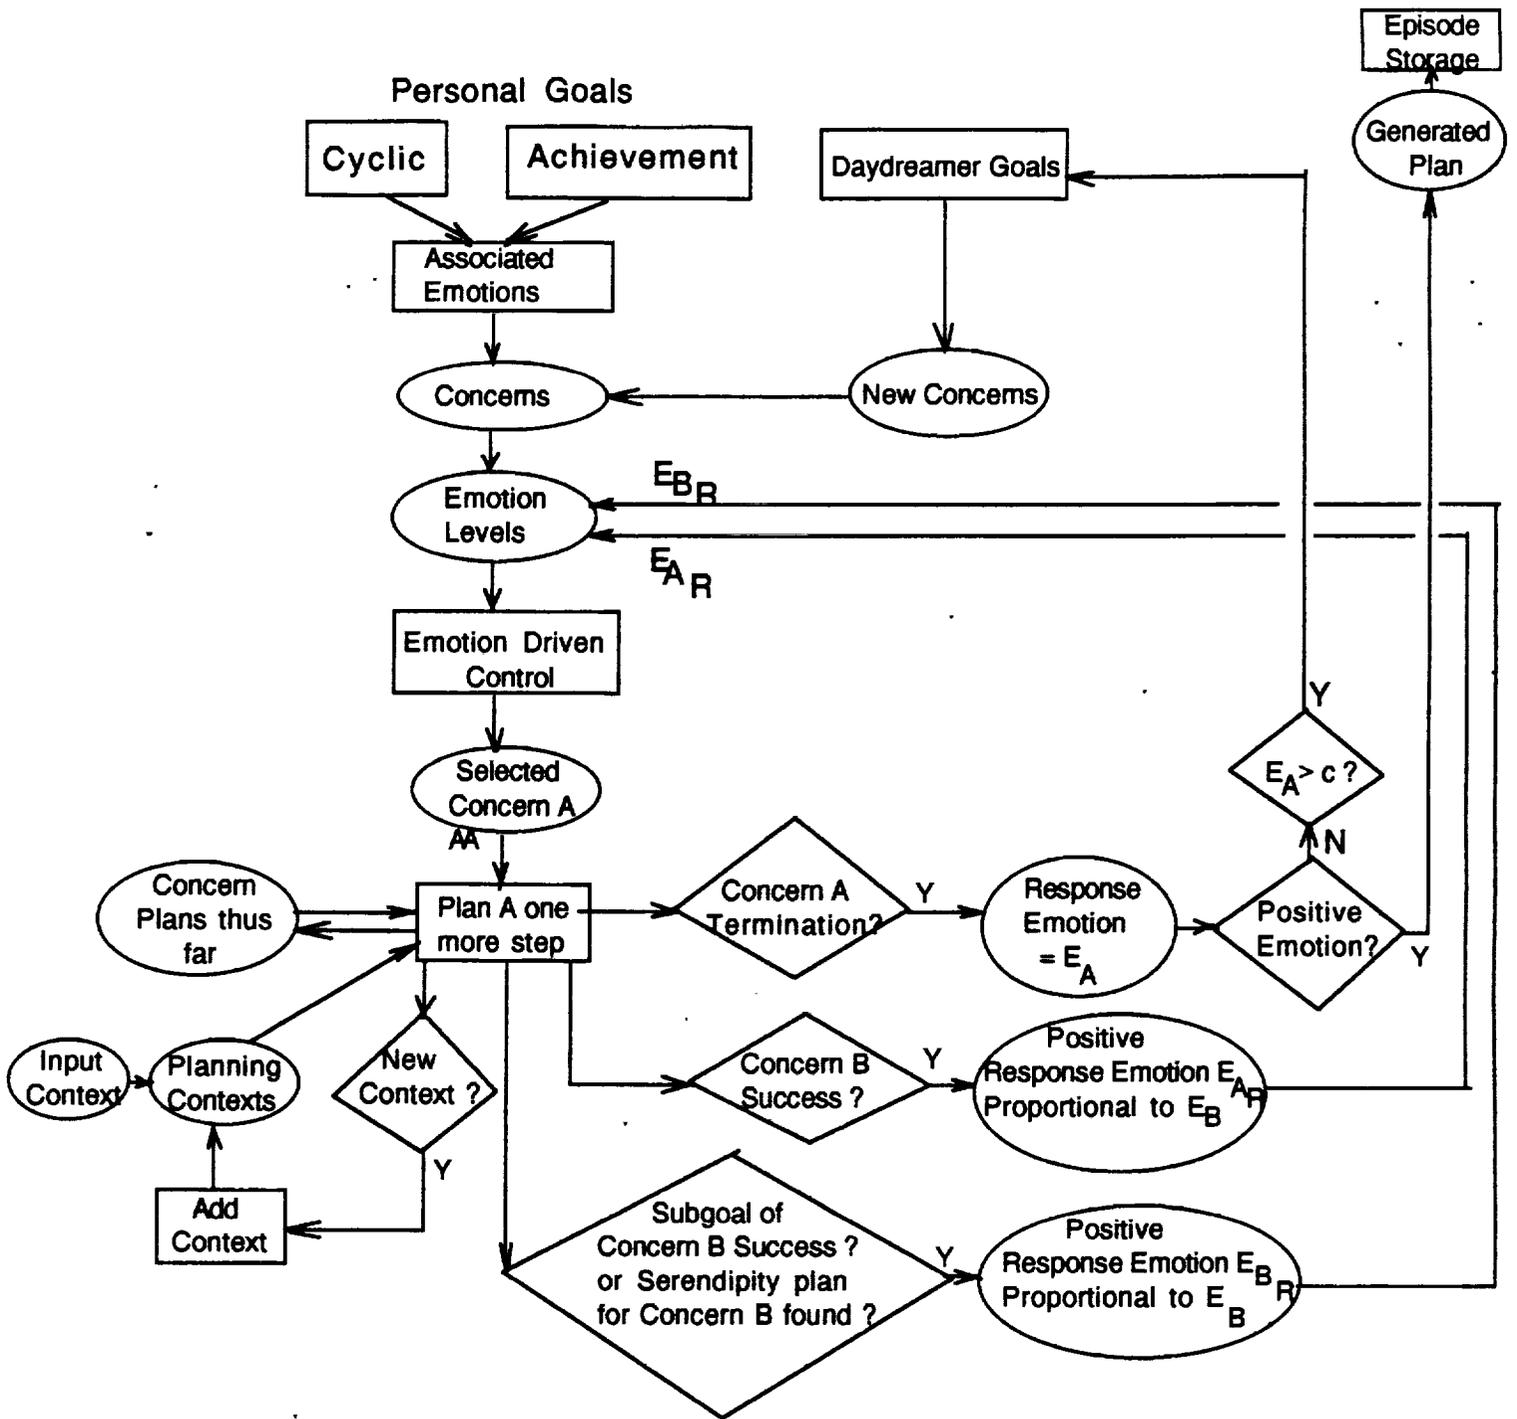


Figure 9. Flow diagram for recognition-primed human decision making.



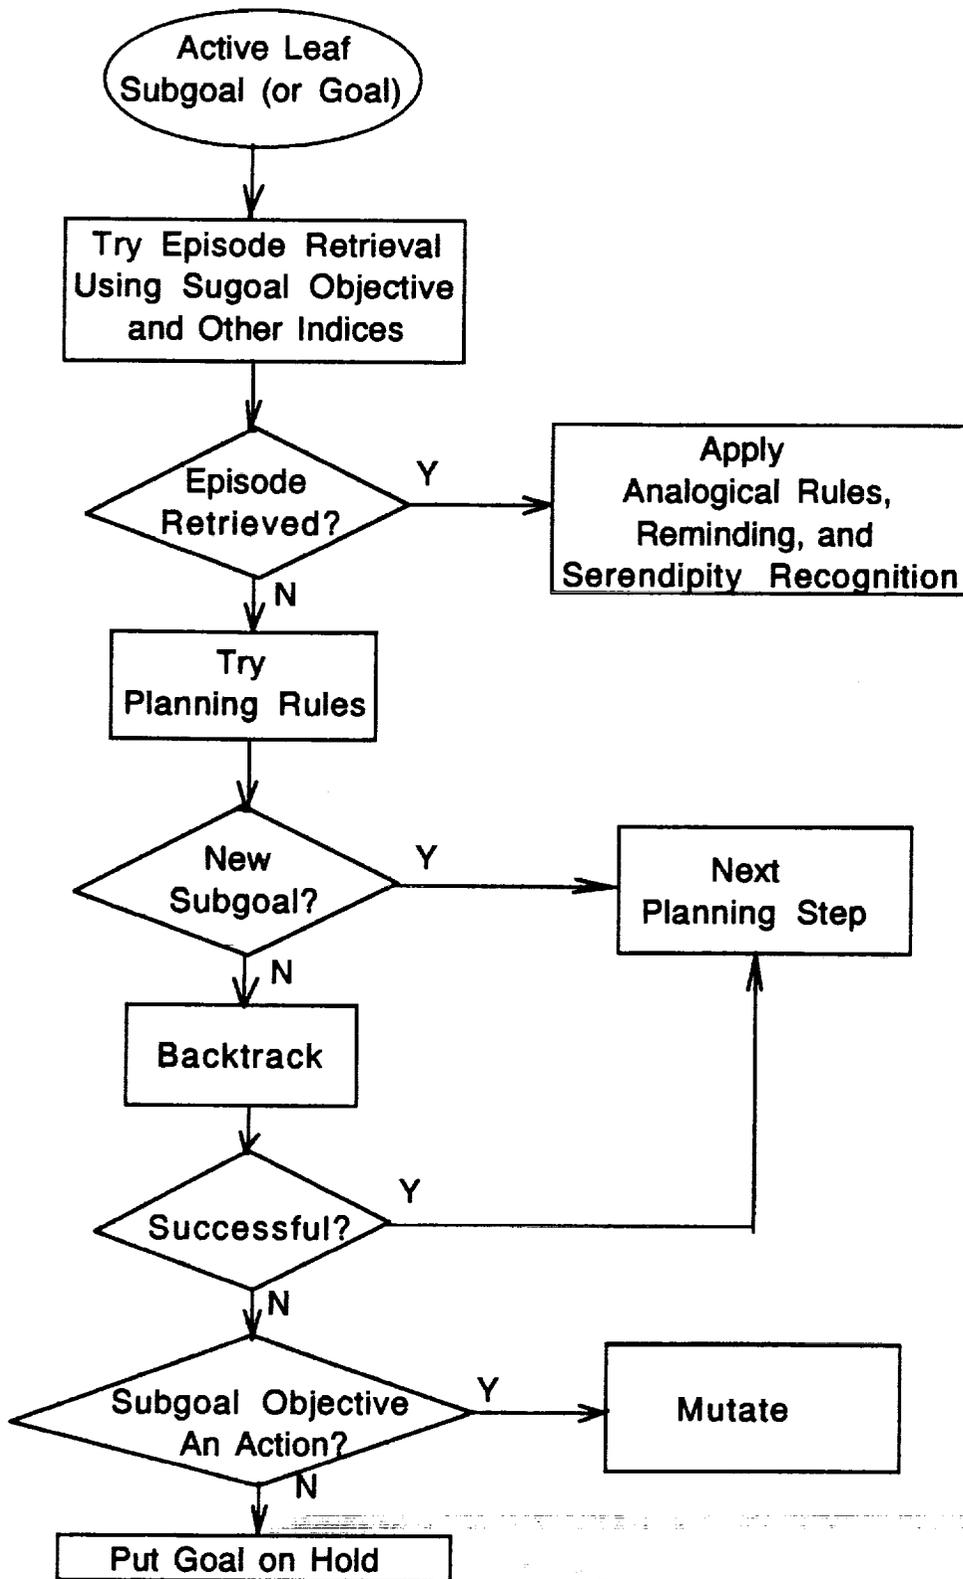
# SIMPLIFIED DAYDREAMER DIAGRAM

Figure 10



## DAYDREAMER FLOW DIAGRAM

Figure 11



# SIMPLIFIED DAYDREAMER PLANNER

FIGURE 12

# ANALOGICAL RULE APPLICATION

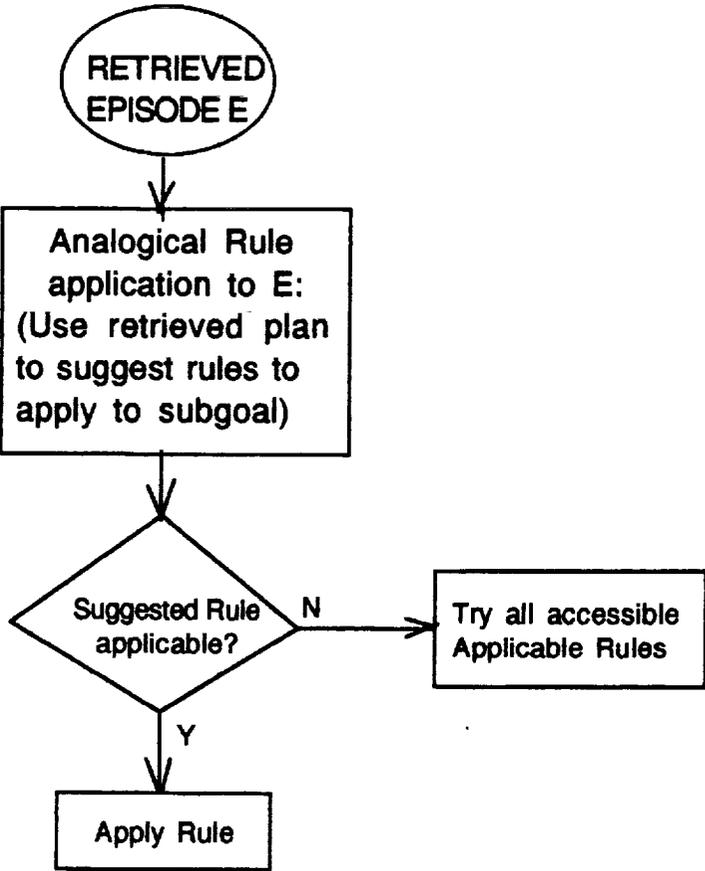


Figure 13

# REMINDING

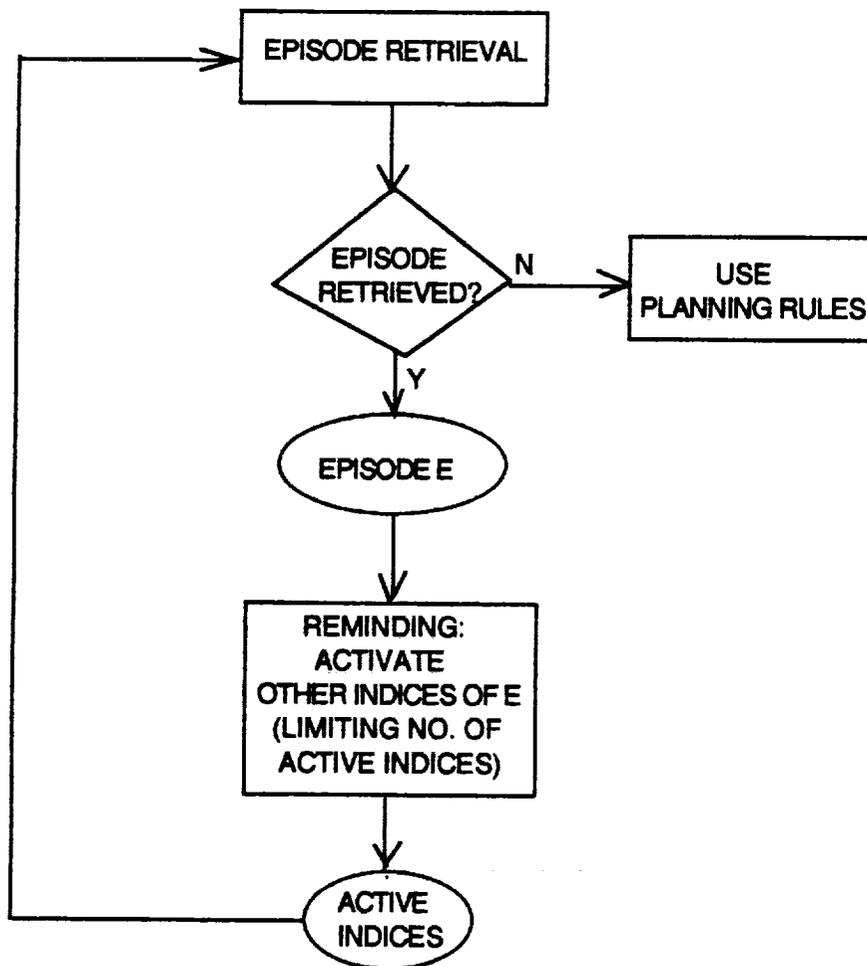


Figure 14

# Serendipity Recognition

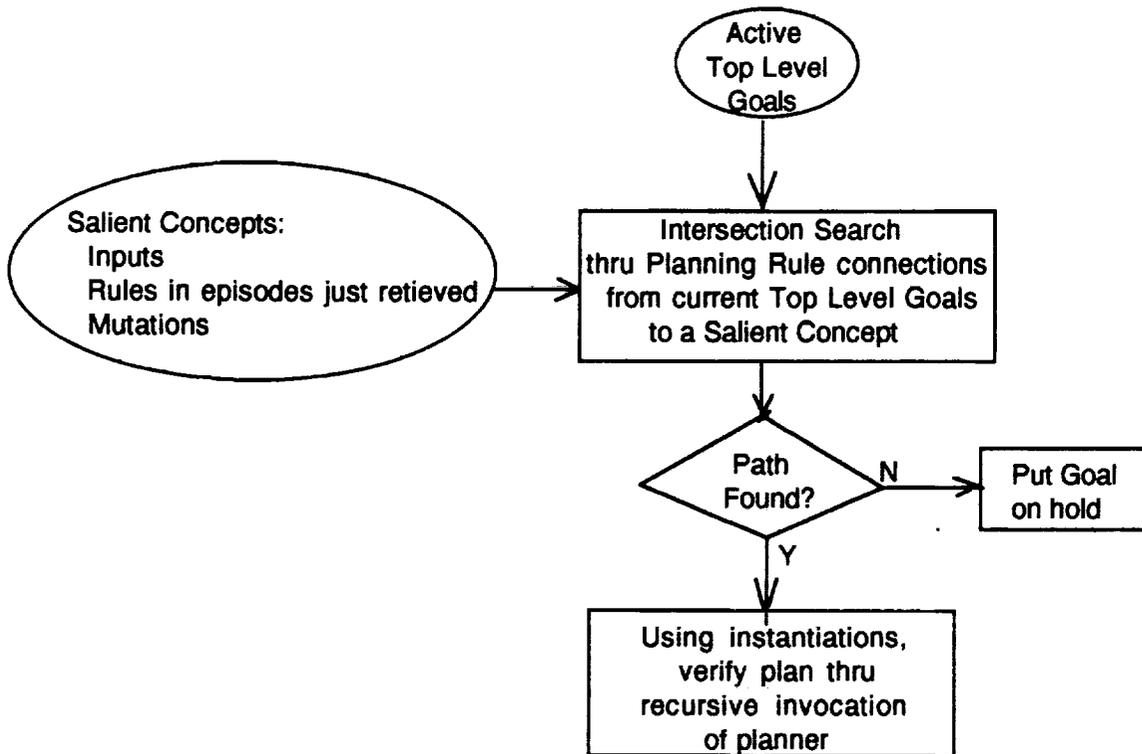
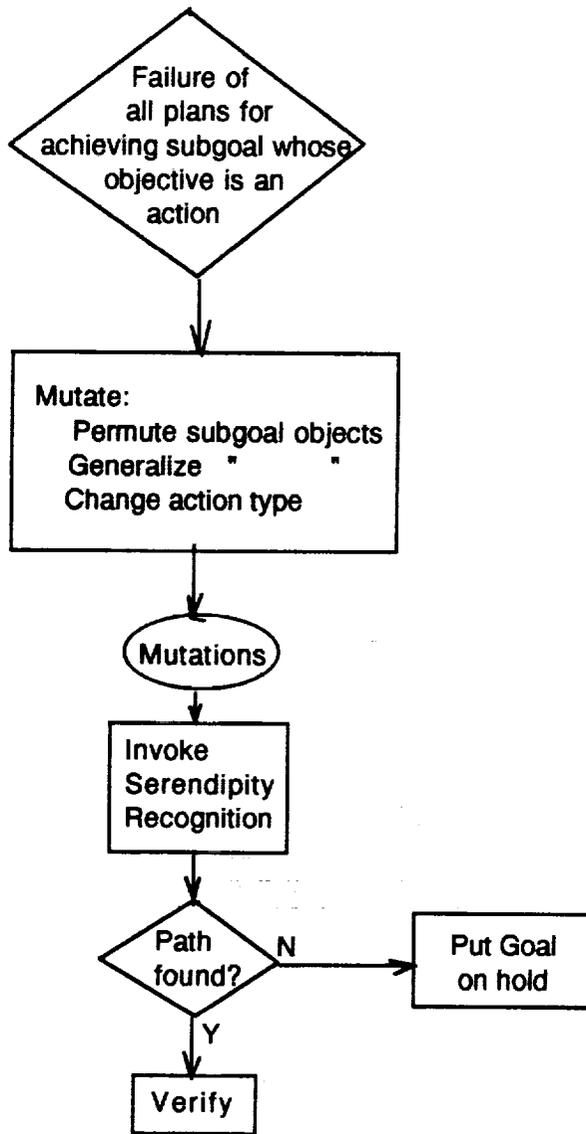
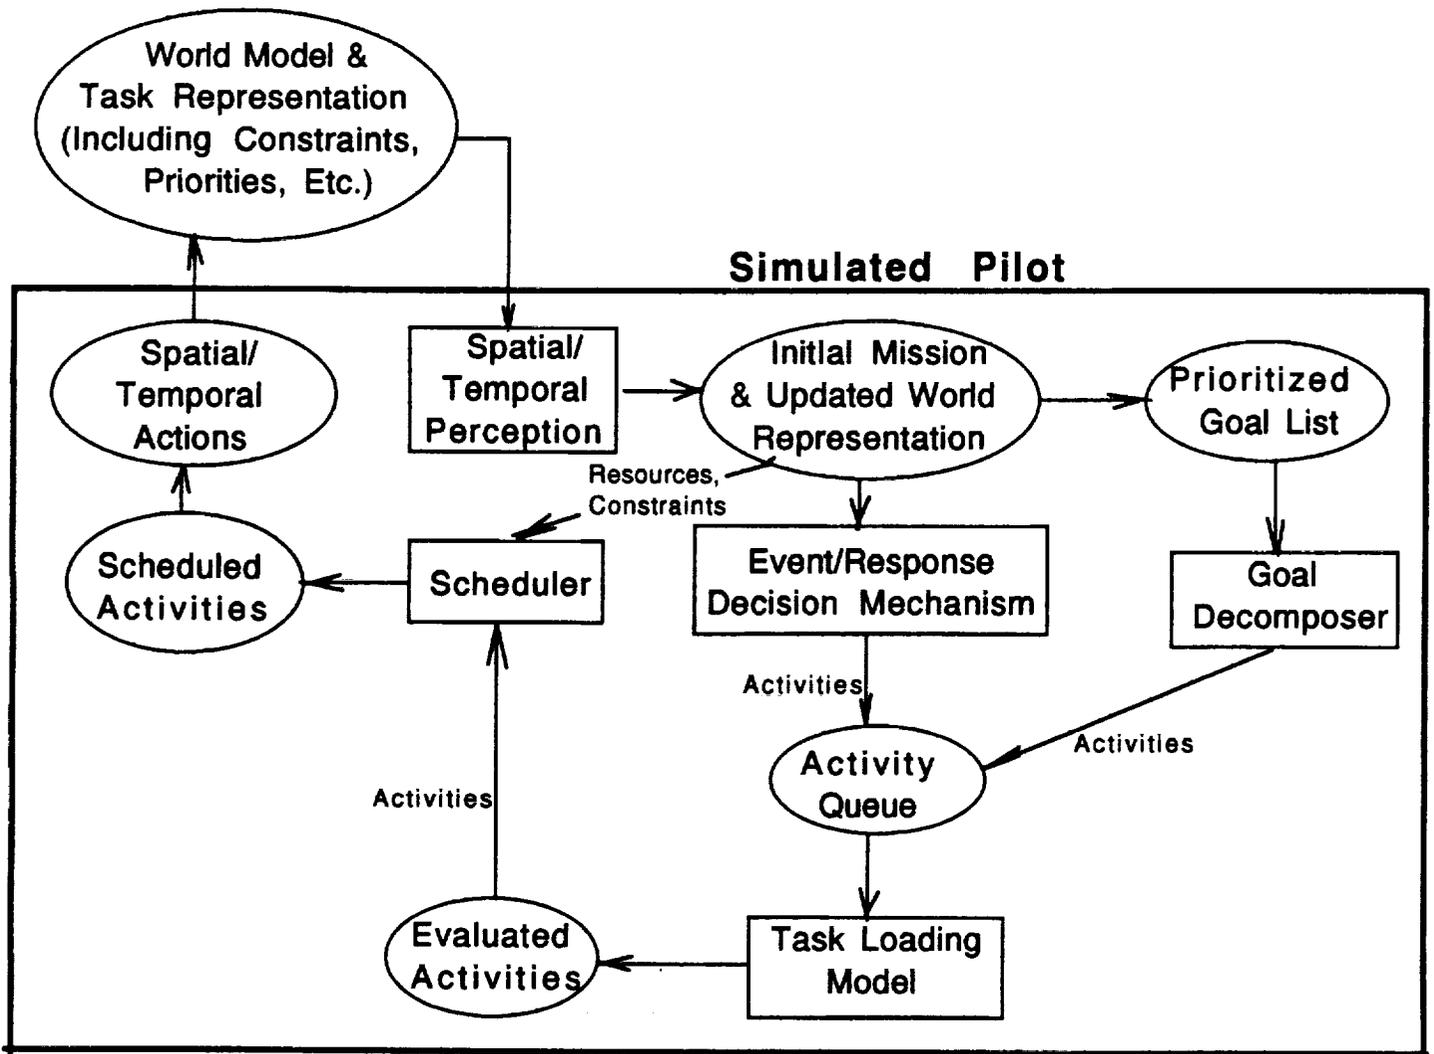


Figure 15



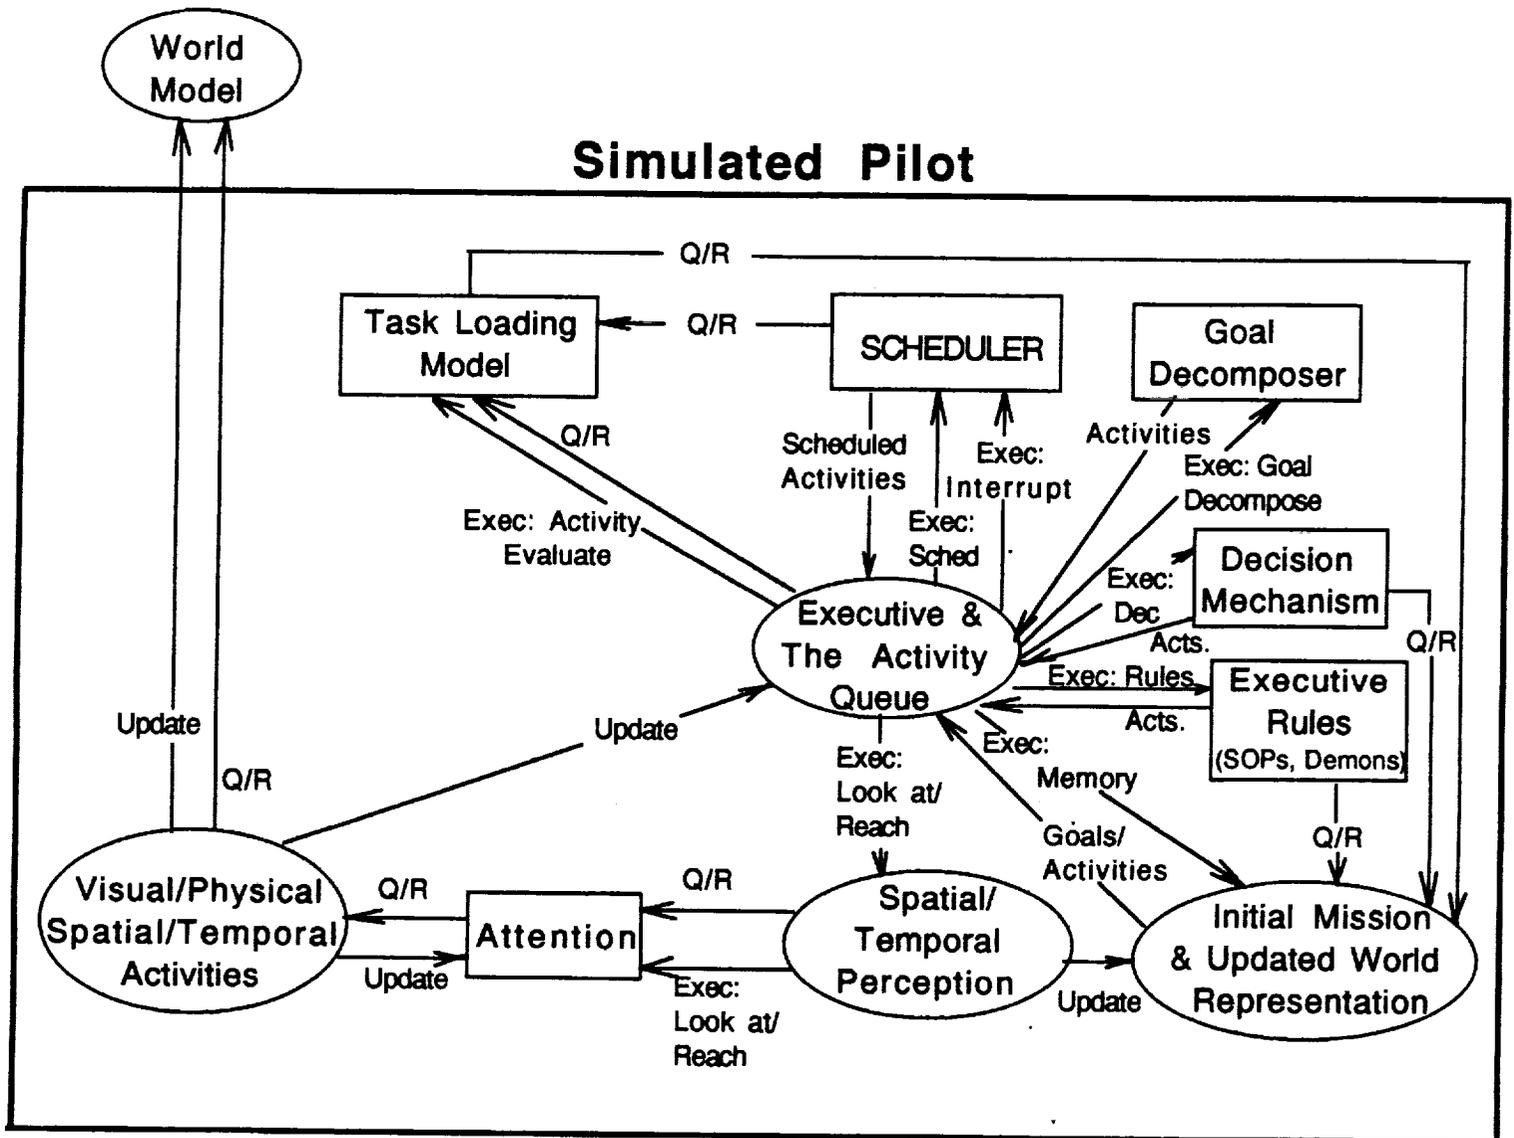
## MUTATION

Figure 16



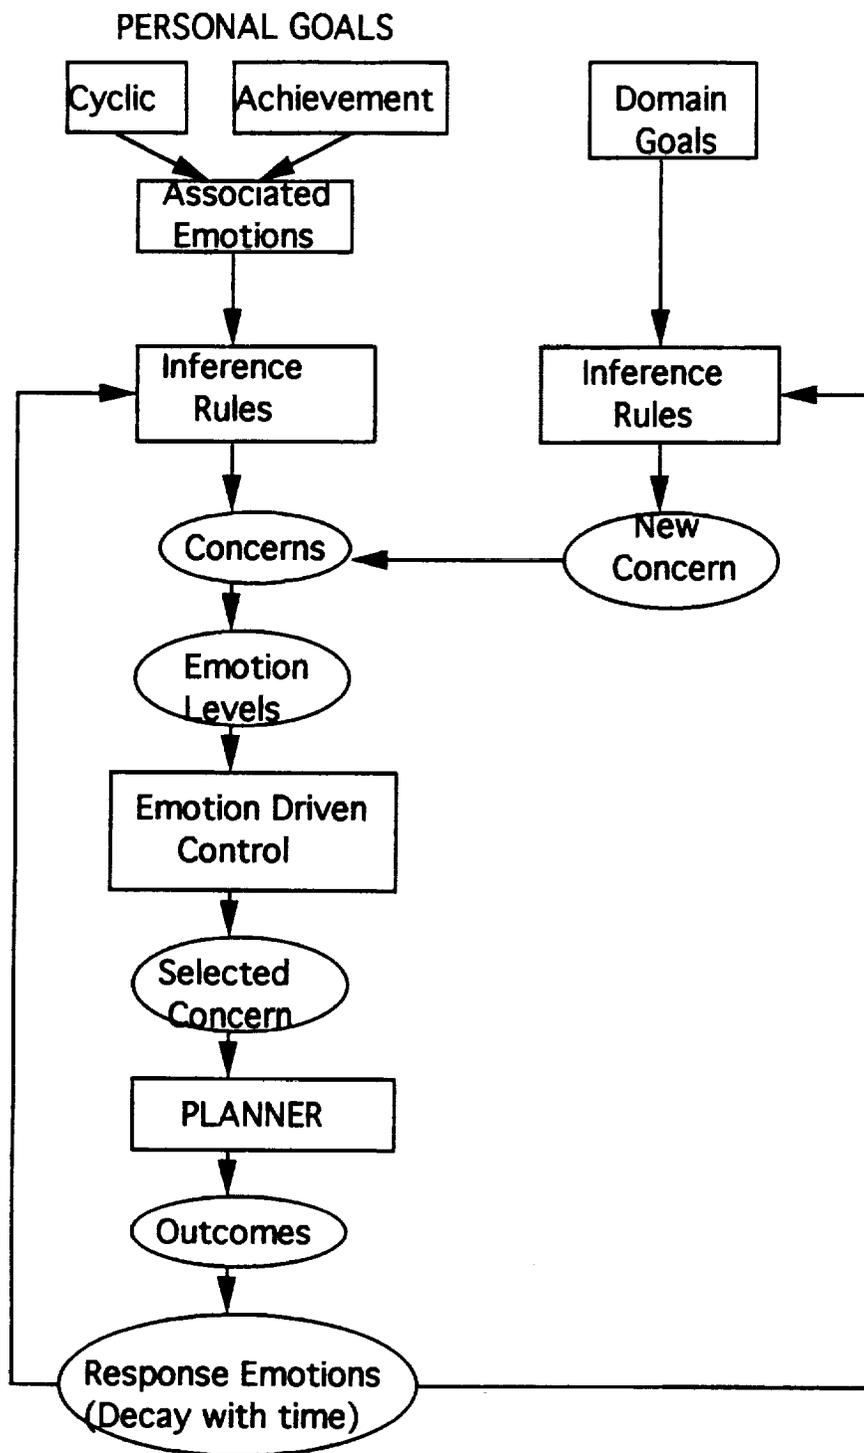
# Simplified MIDAS Flow Diagram

Figure 17



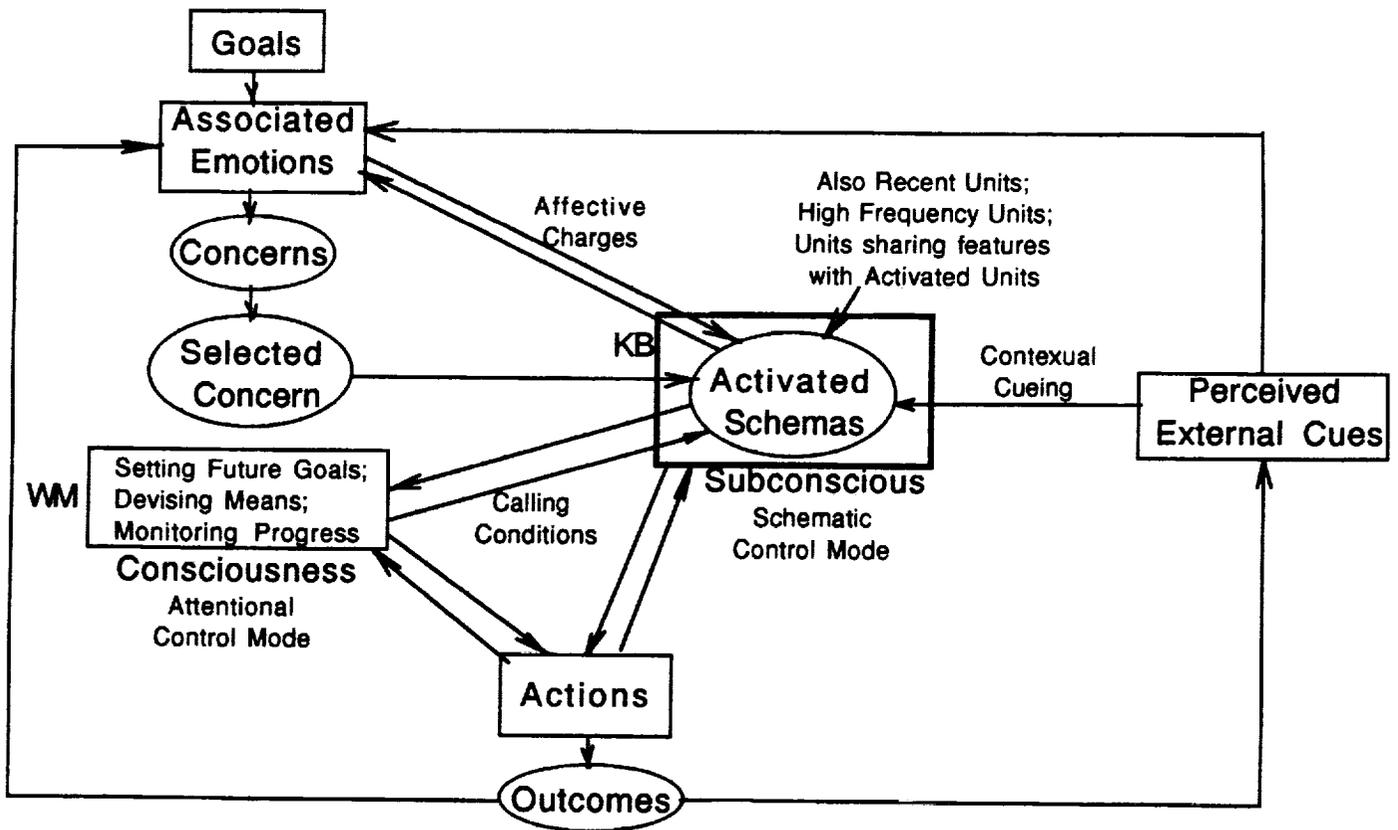
# MIDAS

Figure 18



## EFFECTS OF GOALS ON EMOTIONS AND RESULTANT BEHAVIOR

Figure 19



# HUMAN DECISION MAKING

Figure 20