Information processing, motivation and decision making

Luis Miguel Botelho[1] and Helder Coelho[2]

Abstract

In (Botelho and Coelho, 1995) the authors presented a model of memory for autonomous artificial agents (SALT: Schema-Associative Long Term memory). The main feature of the SALT model is that it allows agents to exhibit context-dependent cognition. This is an important feature since it enables us to gain a better understanding of the reasons why someone may produce different decisions about a given problem, in different situations. In our research we focus on decisions regarding personnel selection. Namely, we are interested in addressing situations in which a manager with a particular task to be done has to decide which of his or her subordinates will be assigned to its execution. In this paper we present the COMINT model, a COgnition and Motivation INTegration model of decision making. The COMINT model extends the original SALT model to explain the influence of motivation in the information processing mechanism of the decision maker. We show that the quality of the decision making process depends on both the motivation to search information in the long term memory of the decision maker, and the motivation to ignore information. We conclude that the COMINT model is able of generating the four different decision making styles described in (Forgas, 1994), using a single information processing mechanism. Along the paper, some aspects of the implementation of the model are discussed and an example is presented.

1 - Introduction

It is widely accepted that human decision makers cope with the complexity of an ever changing environment, trading off accurate decisions by fast decisions. However, the extent to which a decision process is more or less accurate depends on the actual situation. A growing body of theory and evidence suggests there is a significant relation between the effort allocated to the decision process and the motivation of the decision maker. More generally, motivation serves adaptive behavior since it directs the individual to act in conformity to his or her basic instinct of survival as well as to other acquired needs, values and goals. Therefore if we want to understand the way people make decisions, or if we want to build successful artificial autonomous agents capable of deciding, acting and surviving in complex and challenging environments we'd better understand the role of motivation in the decision making process. This paper aims at giving a contribution to this problem.

A great deal of the Artificial Intelligence literature on autonomous artificial agents concerns the representation and properties of several motivational states (e.g., goals, intentions and commitments) and relate them to the agent's overt behavior (Cohen and Levesque, 1990), (Rao and Georgeff, 1991), (Singh, 1991), (Konolige and Pollack, 1993). However, the mentioned literature doesn't address the relationship between motivation and information processing. This constitutes a serious drawback since agents built or defined within this approach have no easy way of saying "That's enough of processing!". If they have a large body of cognitive structures relevant to some decision problem, they are committed to use them all - they can't escape.

The relationship between motivation and information processing has been one of the major concerns in the literature on cognitive social psychology (Srull and Wyer, 1986), (Pyszczynski and Greenberg, 1987), (Forgas, 1994), (Baumeister and Newman, 1994). It is even becoming a central issue in neurology (Damásio, 1994). Nevertheless, this relationship has not been explained with the computational level of detail required for Artificial Intelligence purposes. Hence, there is still a lot to be done along the integration of motivational factors within cognitive models of information processing. In this paper we present a model that extends the SALT model to achieve that integration. The extended model is called COMINT (COgnition and Motivation INTegration model of decision making). Traditionally, motivation is viewed as a conscious and volitional process. COMINT represents an hypothesis regarding possible automatic aspects of motivation.

Joe Forgas (1992; 1994) presented a model (AIM - Affect Infusion Model) with four different styles of decision making corresponding to four different information processing strategies: direct access to stored information, heuristic information processing, motivated information processing, and substantive information processing. However, it has long been suggested that different styles of decision making can be explained in terms of a single information processing

[1]Luis Miguel Botelho, Department of Computer Science of ISCTE, Av. das Forças Armadas, Edifício ISCTE, 1600 Lisbon, Portugal, Tel (351) 1-7935000, Ext 209, Fax: (351) 1-7964710, Email: luis@iscte.pt

[2]Helder Coelho, Group of Artificial Intelligence of the Faculty of Sciences of the University of Lisbon, Bloco C5, Piso 1, Campo Grande, 1700 Lisbon, Portugal, Tel (351) 1-7577831, Fax (351) 1-7577831, Email: helder.coelho@di.fc.ul.pt

mechanism, if we consider differences in the extent of and biases affecting the process (Pyszczynski and Greenberg, 1987), (Baumeister and Newman, 1994). The COMINT model is able of generating the four decision making styles proposed by Forgas. This is done, accounting the role motivation plays on a single information processing mechanism. On one hand, the extent to which information is searched in the agent's long term memory depends on its motivation to search. On the other hand, the agent may be motivated to ignore information inconsistent with its current goals. Both motivation to search and motivation to ignore information depend on other motivational variables stored in long term memory: epistemic needs *(need for accuracy, need for structure, need for a desired outcome),* values, goals and mood (Kruglanski and Klar, 1985), (Feather, 1990) and (Forgas, 1994). In our approach, the motivation to search information is seen as the amount of resources available to the information search process; and the motivation to ignore is seen as a device that filters generated answers to given problems, in given situations.

If we assume an agent is not consciously aware of the automatic operation of the information processing mechanism (procedural knowledge), then whenever it ignores information, it remains unaware of that fact. Therefore, contrary to (Pyszczynski and Greenberg, 1987) and (Baumeister and Newman, 1994), we argue that there is no need for the operation of an affective mechanism of evaluation in order to keep the agent unaware of having ignored information.

Section 2 briefly describes the original SALT model; section 3 describes the influence of the motivation on information processing; section 4 shows how the four decision making styles postulated by Forgas may be explained in terms of the interaction between the agent's motivation and its information processing; section 5 describes the operation of the model in a realistic example; finally section 6 presents some conclusions and directions for future investigation.

2 - SALT: Schema-Associative Long Term Memory

In this section we describe a model of the organization of information in long term memory and the corresponding access methods. As explained in (Botelho and Coelho, 1995), the main advantage of the SALT model results from the fact that it combines the idea of spreading activation with the schema-based model of memory. Therefore it explains all the priming phenomena responsible for the context dependent cognition people exhibit, while permitting the simultaneous activation of a whole body of knowledge necessary to deal with each situation, and enabling a symbolic level of reasoning, usually not provided in connectionism-like models.

According to the SALT model, long term memory is an associative network that may be represented by a directed labeled graph. However, this network is different from other associative and semantic networks (Collins and Loftus, 1975), (Shapiro, 1979), (Anderson and Pirolli, 1984) in three aspects: the contents of the nodes, the meaning and labels of the arcs, and the kind of inference made in the network. Subsections 2.1, 2.2 and 2.3 concern each of these aspects.

2.1 - Contents of the nodes and search in long term memory

A node contains a set of procedures and propositions[3]. So, a node is equivalent to a scheme in many schema-based representation systems. Besides the set of propositions and procedures, a node is characterized by an activation value that represents its accessibility in the network -- nodes more activated are nodes more accessible in long term memory. Like bins in the bin storage model of Wyer and Srull (1986), a node also contains a header describing the concepts involved by the set of propositions and procedures represented in the node. When memory is searched for a node suited to handle a particular situation, the features of the situation are matched against the header of the sampled nodes. The node selected is the first one (i.e., the most activated) that matches the situation. In its present form the model postulates only one search procedure that samples nodes in descending order of activation - nodes more activated are sampled first.

2.2 - Meaning and labels of arcs

A directed arc from node N_1 to node N_2 represents an asymmetric association between the two nodes: if one thinks about concepts in node N_1, then it is likely that concepts represented in node N_2 may come to one's mind, but the converse is not necessarily true. The labels in the arcs represent the strength of the correspondent association. The strengths of associations may vary between but excluding 0 (no association at all) and 1 (strongest association possible). The sum of the strengths of all associations of a given node must be less than or equal to 1.

[3]Like most models of human memory (specially, [Bower81]), it is assumed that some of these propositions may represent affective states. In current implementations of the model, each proposition is expressed in a dialect of the first order logic (horn clauses). This option is not a personal stance of the authors. Actually, the knowledge contained in a node could be represented in any other representation system.

2.3 - Inference possible in the network

For the sake of clarity, in the discussion that follows we assume that a node contains only a set of propositions expressed in FOL (first order logic) -- we do not consider the possibility of procedures[4]. There are only two kinds of inference possible in our associative network: activation and inheritance. Having only these two inference mechanisms doesn't constitute a limitation of the SALT model when compared to other representation systems using FOL. In fact our definition of inheritance exceeds the whole power of FOL inference capabilities.

activation

Activation corresponds to the same concept introduced in theories of spreading activation (Collins and Loftus, 1975), (Anderson and Pirolli, 1984). Each node is characterized by an activation that represents its accessibility in the network. Each time a node is selected to handle a situation it receives a fixed amount of activation per time period. Whenever a node's activation increases the activation's increment spreads to the network (almost) instantaneously, (Anderson and Pirolli, 1984)[5], through the arcs getting out of the activated node. The proportion of the activation's increment that spreads through a particular arc is determined by the product of that increment by the strength of the association: the stronger the association, the greater the activation that spreads through it. Just like in (Anderson and Pirolli, 1984), our model postulates that activation decays exponentially with time.

inheritance

If node N_1 is associated to node N_2, then inheritance is the property that the concepts in N_2 become available to reasoning performed over the concepts in N_1. Stated more formally, suppose node N_1 contains the set of propositions Δ_1, and nodes N_2, ..., N_n represent theories T_2, ..., T_n. Suppose also that there are n–1 directed arcs from node N_1 to nodes N_2, ..., N_n. Then, the theory represented by node N_1 is the closure of $\Delta_1 \cup T_2 \cup ... \cup T_n$ under logical implication.

3 - Influence of Motivation in Reasoning and Search

Section 2 described the way information is organized in long term memory. In section 3, we present the COMINT model of decision making. COMINT extends the SALT model to integrate cognition and motivation in the information processing mechanism of the agent. In subsection 3.1 we describe the general mechanism of information processing postulated by the SALT model. In subsection 3.2 we consider the influence of motivation in the information processing mechanism. Subsection 3.3 deals with the cognitive representation of motivation. Finally, subsection 3.4 focuses on some dynamic aspects of the motivation, and describes the process involved in determining the motivational variables that interfere with the main information processing mechanism.

3.1 - Information processing mechanism

The information processing mechanism described in this section assumes two kinds of memory: long term memory where information is permanently stored, and working memory where selected pieces of information are kept and manipulated while arising problems are faced. According to the SALT model, whenever a problem is put to the agent it searches a node in its long term memory whose contents are relevant to solve the problem -- the initial node. The initial node is activated and the activation spreads to the network through the arcs emanating from it. If the contents of the initial node are not enough to handle the problem, the contents of one or more nodes to which it is associated have to be inherited. Each time a new node is inherited, it is activated and the activation spreads to the network changing the accessibility of nodes. The contents of both the initial node and all inherited nodes are then copied to working memory. Then the agent tries to solve the current problem using the knowledge structures stored in working memory. If the problem cannot be solved, the information processing mechanism backtracks and another set of inherited nodes is tried. If the problem cannot be solved with any possible set of nodes inherited starting with the initial node, the program searches another initial node and the procedure is repeated until a solution is found or no more initial nodes can be selected. After a problem has been handled (successfully or unsuccessfully) another problem may be put to the agent. Actually, this information processing mechanism may be interrupted if a new high priority problem is put to the agent, but for the sake of simplicity, we will not consider this possibility here. Fig. 3.1 depicts the general information processing mechanism and shows the points where motivation may influence its functioning.

---

[4] This simplification enables us to avoid extending the concept of logical implication in order to acomodate all data transformations performed by the execution of such procedures. However, once this extension is made, all the given definitions continue to be valid.

[5] In previous forms of the ACT theory (Anderson 1976) the spreading of the activation was not thought to be instantaneous.

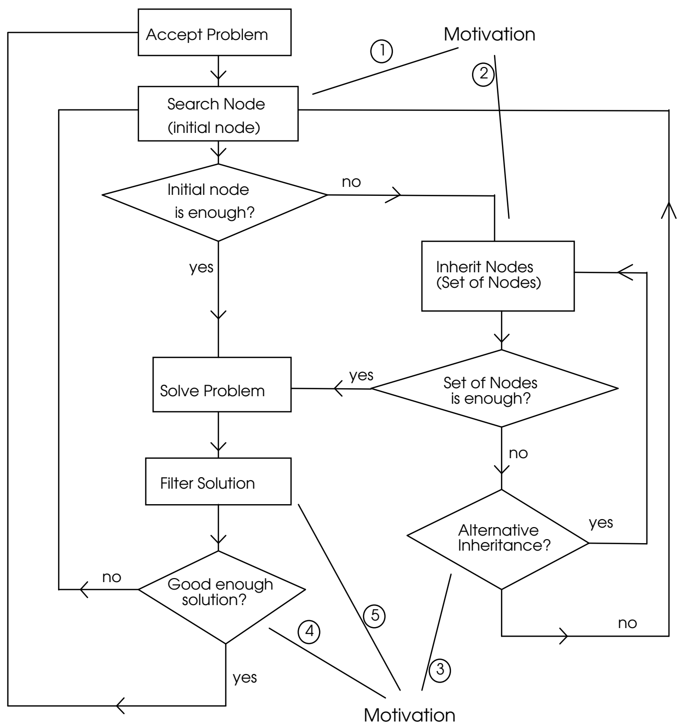

Figure 3.1 - Information processing mechanism

3.2 - Motivation and the information processing mechanism

Fig. 3.1 shows five entry points (encircled numbers) where motivation may interfere with the information processing mechanism. In the present subsection we consider each of these points. First of all it should be noticed the difference in the effects of motivation between entry point 5, on one hand, and entry points 1 to 4, on the other hand. While in entry point 5, some or all the solutions found to the current problem may be ignored (filtered) if the agent is motivated to ignore them; in entry points 1 to 4 the extent of the search process to solve the current problem is determined.

Entry point 1 concerns the amount of search steps the agent is motivated to perform when the initial node to solve the current problem is first sought. In case the motivation to search is very weak, the agent may not be able to find the first node that matches the current problem. In that case, even if the agent possesses enough knowledge to handle the problem, it will give up and fail. This can be understood if the motivation to search is viewed as the amount of resources available to the agent for the search process. If the resources are not enough the search process fails.

In our model, the resource unit is the amount of motivation needed to sample one single node from long term memory (i.e., a search step). A search step involves accessing the (next) most activated node in memory and matching the features of the current problem to the features contained in the header of the node. The search process consists of a sequence of search steps that stops when the agent runs out of motivation to search. If a match between a node and the current problem succeeds, the node is selected, activated and copied to working memory, the activation spreads to the network, and the contents of working memory are used to solve the current problem.

Since COMINT is a model intended to generate computational behavior, we have to represent the motivation to search in a formal way. In that representation we have to express the idea that an agent has a certain amount of search steps available for searching information relevant to handle a given problem. So we need to attach quantities of search steps to problems. A quantity of search steps is represented by an integer. As explained in (Botelho and Coelho, 1995), a problem is formally expressed through a FOL formula. Then, the motivation to search must be represented by a formula that associates another formula to an integer. The obvious way of doing that is to use a modal formula like the following:

mSearch(Problem, SearchSteps)

in which mSearch is a modal operator meaning "motivation to search", Problem is a first order formula stating the problem to be solved, and SearchSteps is an integer specifying the quantity of search steps allowed in the search process.

If the contents of the initial node selected are not enough to handle the current problem, the contents of other nodes have to be inherited[6]. The inheritance operation involves another search process for which more motivation is needed. This is depicted by entry point 2 in Fig. 3.1. Entry point 3 represents the motivation needed to try alternative sets of nodes to be inherited starting with the initial node. Entry point 4 is related to the motivation needed to try another initial node and initiate the process again. If the motivation to search is strong enough, the agent will eventually come up with alternative solutions to the problem.

Entry point 5 involves another kind of motivation -- the motivation to ignore some solutions to the current problem. The motivation to ignore a given solution is represented by a modal formula like the following

mIgnore(Answer)

in which mIgnore is a modal operator meaning "motivation to ignore" and Answer is a first order formula representing a particular solution to the current problem. If mIgnore(Answer) is the case, then the agent ignores the solution specified by Answer. In this way, motivation to ignore may be regarded as a device that filters generated answers to given problems, in given situations.

3.3 - Cognitive representation of motivation

The previous subsection described how the motivation to search and the motivation to ignore interfere with the information processing mechanism. Here we present the hypothesis that the motivation both to search and to ignore may be related to other well-known motivational variables. Those variables include epistemic needs (need for structure, need for accuracy, need for a desired outcome, (Kruglanski and Klar, 1985)), values (Feather, 1990), goals (Srull and Wyer, 1986), (Wyer and Srull, 1986), and mood (Forgas, 1992; 1994). In accordance with (Feather, 1990) and (Bower, 1981), those variables (epistemic needs, values, goals and mood) are represented by information structures stored in the agent's long term memory. Therefore, as a consequence of the spreading activation, the ongoing information processing may change the accessibility of the information structures that represent those motivational variables.

Although we don't specify the concrete contents of the relation between both the motivation to search and the motivation to ignore, on one hand, and certain motivational variables (epistemic needs, values, goals and mood), on the other hand, they might take a form similar to implications such as (3.1) through (3.3).

(3.1) If the need for a desired outcome is very high and the agent is in a moderate negative mood, then its motivation to search is very strong;

(3.2) If the agent has a goal(G), and the need for a desired outcome, G, is very high, then it is motivated to ignore undesired information;

(3.3) If the need for accuracy is very high and the agent is in a moderate negative mood, then its motivation to search is very strong and it is not motivated to ignore information.

We have written a Prolog program that implements our cognitive model. For the sake of simplicity we have considered only five strengths of motivation to search: VERY WEAK, WEAK, MEDIUM, STRONG and VERY STRONG. If the motivation to search is VERY WEAK, the information processing mechanism may perform only a small fixed number of search steps that may not be enough to find the initial node. In this motivational condition, the information processing mechanism won't be able of performing any inheritance. If the motivation to search is WEAK, the information processing mechanism is always able of finding the initial node to handle the current problem, but it won't be able of performing any inheritance as well. If the motivation to search is MEDIUM, the information processing mechanism is able of finding both the initial node and one set of nodes from which to inherit so that a single solution is found. If the motivation to search is STRONG, the information processing mechanism is able of finding alternative sets of nodes from which to inherit (starting with the initial node). Finally, if the motivation to search is VERY STRONG, the information processing mechanism is able of finding all alternative solutions to the current problem. After finding all alternative sets of nodes from which to inherit, starting with the initial node, the system backtracks, searches another initial node and repeats the process again and again until no more initial nodes can be found. In this implementation of the model, the relationships (3.1) through (3.3) are naturally expressed by formulas (3.4) through (3.6) respectively.

---

[6] We could consider another possibility: instead of trying to inherit the contents of other nodes, the information processing mechanism could try to find another initial node. Choosing between these two possibilities, while of theoretical and practical interest, is irrelevant to the present discussion.

(3.4) mSearch(P, veryStrong) :- needFor(outcome, P, veryHigh), mood(bad, moderate).
(3.5) mIgnore(P) :- goal(G), needFor(outcome, G, veryHigh), conflict(P, G).
(3.6) mSearch(P, veryStrong) :- needFor(accuracy, P, veryHigh), mood(bad, moderate).

Without loosing generality, we assume these relationships are represented at the level of the information processing mechanism, as if they were hardwired in the cognitive architecture of the agent (whether innate or acquired and compiled along the agent's lifetime)[7]. However, the motivational variables themselves (epistemic needs, values, goals and mood) are stored in the agent's long term memory. For instance, the agent may have nodes in long term memory like those presented in Fig. 3.2.

<u>Node</u>: goal	<u>Node</u>: need
goal(goodSelfEvaluation). goodSelfEvaluation :- evaluation(self, X), X >= 3.	needFor(accuracy, evaluation(self, X), veryLow).

Figure 3.2 - Two nodes in long term memory representing motivational variables

Since the representation of motivational variables in long term memory must have a special meaning to our computer programs, they must be expressed by formulas with special constants such as **mood**, **goal** and **needFor**.

3.4 - A secondary search process

Until this point we have explained how and when do motivation to search and motivation to ignore interfere with the information processing of an agent. We have further said that the motivation to search and the motivation to ignore depend on motivational variables stored in the agent's long term memory. Therefore, those motivational variables must be searched in long term memory, in order to compute the motivation to search and to ignore information. At this point we have to make an important distinction. When a problem is put to an agent, its information processing mechanism executes a search process aimed at finding long term memory nodes relevant to the solution of the problem. This is called the main search process. The search process executed for finding the motivational variables on which the motivation to search and the motivation to ignore depend, is called the secondary search process. To make it more clear, we might say the main search process is shaped by the content of information structures looked up by the secondary search process[8]. The secondary search process is the central topic of this subsection.

Since motivation to search and motivation to ignore are dependent on motivational variables stored in the agent's long term memory, changes in the accessibility of the motivational variables may entail changes in both the motivation to search and motivation to ignore. With this in mind, the fundamental questions addressed in this subsection are: When to determine the strength of the motivation to search? When to determine if the agent is motivated to ignore a certain piece of information? What should be the nature of the secondary search process needed to find the motivational variables necessary to determine the motivation to search and the motivation to ignore?

The motivation to search and the motivation to ignore, at a particular point in time, both depend on the motivational variables more accessible (i.e., more activated) in long term memory at that point in time. According to the SALT model, the activation of long term memory changes only when a node is selected to handle a given problem, but not before that instant. As a consequence, the motivational variables more accessible in memory may change, only after a node has been selected. Therefore the motivation to search must be determined when the search begins and, after that, whenever a node is selected and the activation spreads to the network (possibly changing what motivational variables are more accessible). To determine the motivation to ignore a given piece of information, that piece of information must be known to the information processing mechanism. Therefore, the motivation to ignore a certain answer to a given problem must be determined after the answer has been produced.

[7]Of course these relationships could be explicitly represented by information structures stored in the agent's long term memory, but this option, while acceptable and probably more consonant with the present approach, would result in useless complexity in the programs we have written to test the model.

[8]Of course, if we ask the agent what is its current goal (a motivational variable), its information processing mechanism executes a main search process to find an answer. Once more, this search process is shaped by information sought by a secondary search process.

There are several differences between the main and the secondary search processes. The COMINT model assumes a number of hypothesis regarding the nature of the secondary search process involved in finding the motivational variables on which the motivation to search and the motivation to ignore depend.

First hypothesis. The search for motivational variables performs a maximum of N search steps. N is a fixed predetermined value.

Second hypothesis. The search stops as soon as relevant motivational variables are found. That is, the information processing mechanism uses the most accessible set of relevant motivational variables represented in the agent's long term memory.

Third hypothesis. If the search for motivational variables stops before they can be found, then it is assumed that the motivation to search has a fixed predetermined strength[9] and that the agent is not motivated to ignore the current answer.

Fourth hypothesis. When a node of long term memory containing the representation of a relevant motivational variable is selected, it is copied to an internal memory buffer (not to working memory), and its activation doesn't change. Therefore the activation of the network doesn't change as a consequence of the secondary search process.

Hypotheses one, two and three express the intuitive idea that a meta-process controlling a main process should be as fast as possible. Hypothesis number four appears for two reasons. First, we feel that a meta-process controlling a main process should not consume resources used by the main process. Therefore, since working memory has limited capacity (Simon, 1990), we should avoid using it to store data used by the control process. Second, since a fundamental operation slowing down the performance of our programs is the spreading of activation, we decided not to include it in the control process.

4 - Four Decision Making Styles

Joe Forgas (1992; 1994) presented a model of decision making (AIM - Affect Infusion Model) according to which, there are four different decision styles corresponding to four different information processing strategies: direct access to information stored in memory, heuristic information processing, motivated information processing, and substantive information processing. However, both Tom Pyszczynski and Jeff Greenberg (1987), and Roy Baumeister and Leonard Newman (1994) argue that differences in decision styles may be explained in terms of differences in the extent of information processing. In this section we show that the four decision styles proposed by Forgas arise naturally in our model as a result of different motivation strengths to search and of the motivation to ignore information, using a single information processing mechanism.

First, the direct access to information in memory results whenever the motivation to search is relatively weak. When the agent is not much motivated to search information, the search process stops as soon as the first relevant node is found in memory. However, in order for a solution to be produced, the first relevant node found in memory must contain information enough to produce a solution for the current problem. Although the COMINT model doesn't specify how the nodes in memory are created and modified, we feel that if the target problem is a familiar one, then it is likely that the agent possesses a long term node containing enough information to handle it. This suggestion is consistent with the hypothesis of the AIM model according to which, the direct access style occurs only when the target problem is a familiar one.

Second, the heuristic processing style should occur when motivation to search is weak, specially when the agent is in a good mood. If we assume that a node representing good mood is strongly associated to a node containing heuristic knowledge to solve the current problem, then if the agent is in a good mood, it is likely that the more accessible node relevant to the current problem is the node containing heuristic knowledge. Therefore, as the motivation to search is weak, the search will stop when the node containing heuristic knowledge (i.e., the most accessible relevant node) is found in long term memory. This means we think of heuristic decision making as a matter of content of accessible knowledge structures, not as a matter of differences in the information processing.

Third, our cognitive model is also able to generate a motivated decision style, when the motivation to search is relatively strong and the agent is motivated to ignore information inconsistent with a desired outcome. Under those circumstances, when the reasoning process finds an answer to the current problem, three things may happen.

(4.1) The answer is not consistent with the desired outcome

(4.2) The answer is consistent with the desired outcome

(4.3) The answer is the desired outcome

In situation (4.1) the answer is ignored. However, if the agent is strongly motivated to search, the information processing proceeds until a new answer is found. In situation (4.2), the answer is not ignored. Like in (4.1) it is possible that the motivation to search is still enough for the information processing to proceed until another answer is found. Finally, in

---

[9] In case of failing to find a set of motivational variables, the current implementation of the model assumes that the strength of the motivation to search is MEDIUM.

situation (4.3), the answer is not ignored. Besides, we may assume the agent's motivation to search becomes very weak (it has already found a desired outcome) and the search process stops. However, COMINT doesn't specify the mechanism governing this dissipation of the motivation to search. Nevertheless, the current implementation of the model assumes that, in cases of high need for a desired outcome, the motivation to search depends on the answers already produced. When the system has already produced a desired answer, the motivation to search disappears.

Joe Forgas defined the motivated decision making style as a process that avoids undesirable information and actively searches desirable information, such as what happens when someone is trying to repair a bad mood or to maintain a good mood. This is exactly the kind of behavior our cognitive model generates when the motivation to search is strong and the agent is motivated to ignore information inconsistent with a desired outcome.

Finally, if the motivation to search is very strong and the agent is not motivated to ignore information, then the search process will produce all alternative answers to the current problem - that's a substantive processing style.

5 - A Case Study: Personnel Selection

In the previous sections we have outlined the COMINT model and presented some aspects of its implementation. In this section we illustrate the functioning of the model, through a realistic example built upon recent findings reported in the personnel selection literature. In this example, both the motivation to search and the motivation to ignore information interfere with the information processing of the decision maker, yielding a motivated decision making style. We have chosen an example of this decision style because it involves all the relevant features of the COMINT model.

According to Jackson and Schuler (1995), it is very likely that a manager who wants to gain power in his or her company, has a strong motive to control valued human resources. On the other hand, since there is an increasing belief that trust is very important to reduce uncertainty within organizations, it is also likely that a manager has a strong motive to present a self image of a trustworthy person, to their subordinates. According to (McAllister, 1995) a person tends to trust a second person, if the second person's decisions regarding the first take into account his or her interests. Therefore it is plausible that a manager who wants to present a self image of trustworthiness takes into account the interests of his or her subordinates when making personnel decisions regarding them.

The scenario that serves as a case study for the COMINT model is based on the ideas just presented. In this scenario, a given manager, Mr. Abraracourcix, has two motivational structures stored in his long term memory: one represents a motivation to control valued human resources, the other represents a motivation to present a self image of a trustworthy person. In the line of (Botelho and Coelho, 1995), we assume that some contexts enhance the accessibility of one of those motivational structures in Mr. Abraracourcix's long term memory, while other contexts enhance the accessibility of the other motivational structure. Mr. Abraracourcix has a specific task to be done: he wants a report on the environmental impact of a new power plant, and he has to decide which of his two subordinates will make the study - Mr. Simpson or Miss Yoko Tsuno. Mr. Simpson is not a valued employee for the company, however he would like to execute the study. Miss Yoko Tsuno is quite the opposite. She is a very efficient employee and she has a high instruction level. Several other departments would like her to work for them. Unfortunately she is not interested in making that environmental study.

It is natural to expect that, in contexts in which the motivation to control valued human resources becomes more accessible in Mr. Abraracourcix's long term memory, he would assign the environmental study to Miss Yoko Tsuno because she is a valued employee; while in contexts that enhance the accessibility of the self-presentation motive, Mr. Abraracourcix would assign the study to Mr. Simpson, because he is very interested in it. That's exactly the behavior generated by the COMINT model. Fig. 5.1 depicts two interactions with a COMINT agent playing the role of the manager in the imagined scenario just described.

Q1: How would you gain power in your company?
A1: Controlling valued human resources.
Q2: Who are you going to assign to the execution of the environmental study?
A2: Miss Yoko Tsuno.

Q3: Tell me a way to reduce uncertainty in your department?
A3: Presenting a self image of a trustworthy person.
Q4: Who are you going to assign to the execution of the environmental study?
A4: Mr. Simpson.

Figure 5.1 - English transcript of an interaction with a COMINT agent playing the role of a decision maker

The first question is intended to create a context that enhances the accessibility of the motive to control valued human resources in the agent's memory (see section 2 and (Botelho and Coelho, 1995) for details). In this context the agent assigns a valued employee to the task at hand, even though the employee is not interested in performing that task. This decision outcome enables the manager to control valued human resources and hence to gain some more power in its company. Question Q3 is intended to create a context that enhances the accessibility of the self-presentation motive. This time, the agent assigns the task to an employee interested in executing it, even though the best thing to do, in the interest of the company, would be not to assign this task to this particular employee. This decision outcome enables the manager to present a self image of trustworthiness, decreasing the organizational uncertainty. In the remaining of this section we describe the way the COMINT model generates the behavior presented in Fig. 5.1.

We have created a model of a decision agent (the manager). The agent is in a moderate negative mood, has a model of a task to be executed, the models of its two subordinates (Mr. Simpson and Miss Yoko Tsuno), two motivational structures, and two decision rules for assigning people to the execution of tasks, Fig. 5.2. There are also other cognitive structures just to ease the task of creating contexts that enhance the accessibility of each motivational structure.

<u>Self-presentation motive:</u>

```
needFor(outcome, assign(T, F), veryHigh).

goal(trustworthy(self)).

trustworthy(Y) :-
    assign(T, X),
    not interested_in(X, T),
    !,
    fail.
trustworthy(Y).
```

<u>Affective state:</u>

```
affectiveState(midNeg).
```

<u>Decision rule 1:</u>

```
assign(T, X) :-
    task(T),
    valued(X).
```

<u>Decision rule 2:</u>

```
assign(T, X) :-
    task(T),
    interested_in(X, T).
```

Figure 5.2 - Partial contents of the agent's long term memory

According to (Forgas, 1992; 1994) and (Kruglanski and Klar, 1985), a moderate negative mood and a strong need for a desired outcome entail a strong motivation to search relevant information, and a motivation to ignore information inconsistent with the agent's motives. This relationship is encoded in the agent's information processing mechanism (section 3.3).

Let us consider what happens after the agent has answered question Q3. In this context, the most accessible motivational structure in the agent's long term memory is the self-presentation motive. At this point the agent is faced with a decision problem expressed in question Q4 - it has to assign one of its employees to the execution of a particular task. Now, we assume the agent uses decision rule number one (fig 5.2). This rule, together with the information the agent has about its employees (not presented in fig 5.2), entail the selection of Miss Yoko Tsuno (valued employee). However, this solution is not consistent with the agent's more accessible motivation structure (fig 5.2), because Miss Yoko Tsuno is not interested in the task. Since the agent is motivated to ignore information not consistent with its current goal, this first answer is ignored. However, the agent is strongly motivated to search relevant information for the current problem, hence it continues to try. This time, it finds decision rule number two (fig 5.2). This new rule entails the selection of Mr. Simpson (interested in executing the task). This outcome is consistent with the agent's current goal, therefore it is not ignored and it is presented to the user. If the agent is still motivated to search, other answers will be tried and the process continues the same way as before.

One of the most important features of the described decision process is that none of the decision rules used to select employees to be assigned to tasks (fig 5.2) makes any reference whatsoever to the current motives of the agent. This means the COMINT model enables goals and other motives to control apparently unrelated cognitive tasks, without the agent possessing methods or plans to do so. This feature has two important consequences. First, it provides the agents with a way to avoid having to learn all possible methods by which all possible motives may control all possible cognitive tasks. Second, as the agent doesn't have any declarative cognitive structure relating its current motivation and its current cognitive task, it remains unaware of having manipulated its decision process in order to produce a self serving outcome (instead of the best outcome for the company), and so it does not feel guilty.

Let us now suppose that instead of a high need for a desired outcome, the agent has a high need for accuracy. In this conditions, according to expression (3.3) it will continue to exhibit a strong motivation to search, but no motivation to ignore answers inconsistent with its goals. This time, the first answer found (assign the task to Miss Yoko Tsuno) wouldn't have been ignored. So, the agent would have produced all different alternatives (from which it would have to choose one). This would have been the substantive processing style.

Finally, we suppose the agent doesn't have a high need for accuracy nor a high need for a desired outcome, and consequently it is not strongly motivated to search nor motivated to ignore information. In this way, it wouldn't have ignored the first answer produced (Miss Yoko Tsuno) and it wouldn't have tried to find another answer. This would have been what Forgas would have called the direct access to information in memory. But, since the decision rule used to produce the answer ("if valued(x) then assign(x)") may be considered an heuristic rule, this may also be considered the heuristic decision making style. Indeed, we view the heuristic decision style as a matter of the content of the knowledge structure used, not as a matter of a different processing mechanism.

6 - Conclusions and Future Investigation

We have presented COMINT, a unifying cognitive model of decision making. First, both beliefs and motivational variables (epistemic needs, values, goals, affect) have a similar representation in long term memory. Second, it presents a single information processing mechanism.

The model presented is not only a descriptive model, but also a generative model because it produces behavior. The model is able of generating the four decision making styles postulated by Joe Forgas, making use of a single unified information processing mechanism, in accordance with (Pyszczynski and Greenberg, 1987) and (Baumeister and Newman, 1994). More than that, since the motivation to search controls the number of search steps performed by the information processing mechanism, the model is actually capable of generating a continuum of decision making styles ranging from the direct access to information stored in memory, to the substantive information processing style. The model also accounts for dynamic changes in motivation due to the operation of the information processing mechanism, itself. Therefore, the decision style may also change during the course of its operation. Finally, according to COMINT, a manager may produce self serving decisions without having explicit cognitive structures representing plans to achieve his or her motives. This means the manager can't be consciously aware of having decided in ways that suited his or her motives instead of the interests of the company. Besides, if we assume no one is consciously aware of the automatic operation of one's information processing mechanism (procedural knowledge), then whenever one ignores information, one remains unaware of that fact. Therefore, contrary to (Pyszczynski and Greenberg, 1987) and (Baumeister and Newman, 1994), we argue that there is no need for the operation of an affective mechanism of evaluation in order to keep the agent unaware of having ignored information. So we feel, in this case, it would be unfair to blame a manager for his or her sub optimal decisions. Instead, top management should spend more effort understanding this natural phenomena and trying to create conditions to minimize their negative consequences.

It should be noticed that COMINT model makes it conceivable that an agent may produce different decisions about a given problem, in different situations, not due to different accessibilities of the cognitive structures used directly to face the decision problem[10], but due to different accessibilities of certain motivational structures. This feature is quite important since it represents a mechanism through which an agent may possess a contradictory set of motives and beliefs, and still be able of deciding without performing any belief revision. It is important to understand that this kind of behavior is only possible due to the properties of the SALT model of long term memory.

In consonance with the tradition, we suspect that besides the automatic effects of motivation on the information processing mechanism, there is another kind of effects controlled by the conscious volition of the individual. We haven't dealt with this kind of effects, yet. The next step in our research project will be the extension of the model to account for this conscious and volitional effects.

Another direction for the future will be to investigate the exact nature of the relationship between the motivation to search and the motivation to ignore, on one hand, and the motivational variables such as epistemic needs, values, goals and affect, on the other hand.

Finally it is necessary to answer questions related to the adaptive qualities of human beings and try to transpose them to our naturistic[11] agents. Such questions include the dynamic creation and modification of long term memory nodes, as well as the dynamics of the strengths of the associations between nodes in long term memory.

Acknowledgments

The authors are grateful to Pedro Ramos for permanent discussions and comments on early drafts of this work. He also suggested the names of the characters used in the example.

---

[10] In this example, the cognitive structures representing decision rules number one and number two, have the same accessibility.

[11] Naturistic agents: artificial agents whose behavior is the closest possible to their natural counterparts.

References

Anderson, J.R. (1976) Language, memory and thought, Earlbaum, Hillsdale, N.J.

Anderson, J.R. and Pirolli P.L. (1984) Spread of activation, *Journal of Experimental Psychology: Learning, Memory and Cognition, 10, 791-798*

Baumeister, R.F. and Newman, L.S. (1994) Self-regulation of cognitive inference and decision processes, *Personality and Social Psychology Bulletin, 20, 3-19*

Beach, L.R. (1990) Image theory: decision making in personal and organizational contexts, John Wiley and Sons, New York

Botelho, L.M. and Coelho, H. (1995) A Schema-Associative Model of Memory, *Proc. of the Fourth Golden West International Conference on Intelligent Systems, pp.81-85*

Bower, G.H. (1981) Mood and memory, *American Psychologist, 36, 129-148*

Cohen, P.R. and Levesque, H. (1990) Intention is choice with commitment, *Artificial Intelligence, 42, 213-261*

Collins, A.M. and Loftus, E.F. (1975) A spreading-activation theory of semantic processing, *Psychological Review, 82, 407-428*

Damásio, A.R. (1994) Descartes' Error: Emotion, Reason and Human Brain, Picador, London

Feather, N.T. (1990) Bridging the gap between values and actions: recent application of the expectancy-value model. In Higgins, T.E. and Sorrentino, R.M. (eds) *Handbook of Motivation and Cognition: Foundations of Social Behavior, Vol. 2, pp151-192*, Guilford Press, New York

Forgas, J.P. (1992) Affect in social judgments and decisions: a multiprocess model, *Advances in Experimental Social Psychology, 25, 227-275*

Forgas, J.P. (1994) The role of emotion in social judgments: an introductory review and an affect infusion model (AIM), *European Journal of Social Psychology, 24, 1-24*

Jackson, S.E. and Schuler, R.S. (1995) Understanding human resource management in the context of organizations and their environments, *Annual Review of Psychology, 46, 237-264*

Janis, I., and Mann, L. (1977) Decision making. A psychological analysis of conflict, choice and commitment, The Free Press, a division of Macmillan Publishing Co. Inc., New York

Konolige, K. and Pollack, M.E. (1993) A representationalist theory of intention, *IJCAI'93, pp390-395*

Kruglanski, A.W. and Klar, Y. (1985) Knowing what to do: on the epistemology of actions. In Kuhl, J. and Beckmann, J. (eds) *Action Control: from cognition to behavior, pp101-128,* Springer, Heidelberg

McAllister, D.J. (1995) Affect- and cognition-based trust as foundations for interpersonal cooperation in organizations, *Academy of Management Journal, 38, 24-59*

Pyszczynski, T. and Greenberg, J. (1987) Toward an integration of cognitive and motivational perspectives on social inference: a biased hypothesis-testing model, *Advances in Experimental Social Psychology, 20, 297-340*

Rao, A.S. and Georgeff, M.P. (1991) Modeling rational agents within a BDI-architecture, *KR' 91, Proc. of the 2nd International Conference on Principles of Knowledge Representation and Reasoning, pp.473-484*

Shapiro, S.C. (1979) The SNePS semantic network processing system. In Findler, N. (ed) *Associative Networks: The Representation and Use of Knowledge by Computers,* Academic Press, N.Y.

Simon, H.A. (1990) Invariants of human behavior, *Annual Review of Psychology, 41, 1-19*

Singh, M.P. (1991) On the commitments and precommitments of limited agents, *IJCAI'91 Workshop on Theoretical and Practical Design of Rational Agents*

Srull, T.K. and Wyer, R.S. (1986) The role of chronic and temporary goals in social information processing. In Sorrentino, R. and Higgins, T. (eds) *Handbook of Motivation and Cognition: Foundations of Social Behavior, Vol. 1, pp503-549,* Guilford Press, New York

Wyer, R.S. and Srull, T.K. (1986) Human cognition in its social context, *Psychological Review, 93, 322-359*