In *Prospects for Artificial Intelligence*, Proc. AISB'93, (1993)
Eds. A. Sloman, D. Hogg, G. Humphreys, D. Partridge & A. Ramsay,
pp. 229–238, IOS Press, Amsterdam

A Study of Motive Processing and Attention

Luc P. Beaudoin and Aaron Sloman
Cognitive Science Research Centre, University of Birmingham
L.P.Beaudoin@cs.bham.ac.uk, A.Sloman@cs.bham.uk.ac

Abstract. We outline a design based theory of motive processing and attention, including: multiple motivators operating asynchronously, with limited knowledge, processing abilities and time to respond. Attentional mechanisms address these limits using processes differing in complexity and resource requirements, in order to select which motivators to attend to, how to attend to them, how to achieve those adopted for action and when to do so. A prototype model is under development. Mechanisms include: motivator generators, attention filters, a dispatcher that allocates attention, and a manager. Mechanisms like these might explain the partial loss of control of attention characteristic of many emotional states.

1. Introduction

A mother whose child has been run over reported that, for a long time afterwards, she found her thoughts constantly and unwillingly drawn back to the event, to what she might have done to prevent the accident, to what the child might have been doing had it not occurred. She was so tormented by this that only the desire not to harm her remaining children prevented her taking her own life. A common feature of infatuation is that thoughts, desires, questions and worries concerning the loved one keep intruding. An Olympic winner will find it hard not to return in thought to the moment of success and the associated struggles and triumph. Emotional states like grief, infatuation and elation seem to have in common a partial loss of control of one's own thought processes: in other words, attention is repeatedly drawn back to thoughts, desires, regrets, what might have happened, what may yet happen, and the like. This is part of what it means to be "moved", or even to be "beside oneself". Why is this such a pervasive feature of human experience? We shall not quarrel about the definition of "emotion" since the word is full of ambiguity and vagueness. All we are concerned with is that there are certain features found in some mental states that many people would describe as "emotional".

Our conjecture is that design requirements for intelligent agents in a world like ours lead to mechanisms that have these tendencies as "emergent" properties. Our "Attention and Affect" project aims to develop this conjecture into a high level theory of the architecture of a human mind, integrating many of the phenomena studied by AI and Psychology. As a first step we are exploring design options within a simulated domain, the "robot nursemaid domain". This paper reports some preliminary work.

2. Requirements

Any explanation of how autonomous agents (such as human beings) work must account for a number of features concerned with attention and motivation, as follows. The agent will have multiple independent sources of motivation operating asynchronously, triggered by external and internal events (e.g. hunger, thirst, seeing a friend in trouble). Its mental processing is parallel at least at a coarse grained level: for instance controlling a physical action while monitoring the environment, being sensitive to new possibly unrelated problems, and planning next steps (compare [18]). There are often time constraints (which it may discover dynamically) on the goals or desires it attempts to fulfil. Its information about the world (and itself) will generally be incomplete and may contain inaccuracies.

In addition an agent's capacity to deal with problems is limited both mentally and physically [11]. For instance, it cannot be in two different places, look in two quite different directions, simultaneously plan courses of action for two unrelated and complicated problems. Mental resource limits require various *selections* to be made: which goals to try to achieve, which ones to think about now, which data to look at, which procedures to apply to the data, and so on. In other words attention is directed to meet a *subset* of current needs.

3. Example: a Real Nursemaid

The problems of controlling attention arise in many human activities. Consider a nursemaid in charge of a collection of babies: She needs to be able to detect and react to a great variety of problems, for instance, babies becoming hungry, sick or injured, or their causing damage to themselves or others. While taking care of a problem she is able to react to other, more pressing, problems. For instance, if while feeding one baby and thinking about how to handle a peculiar rash on another she notices that a third baby is choking on a string, she may then interrupt the feeding as well as her train of thought about the rash, in order to decide how to care for the choking baby. Having dealt with it, she remembers her previous feeding task and returns to it if no more pressing problem has arisen. Depending on the nursemaid's constitution, her skill and familiarity with such situations, and the speed with which she needs to react, she may go through states of alarm, dismay, anxiety, relief or joy.

4. Towards an Architecture

We shall sketch a theory of the processing of motives which postulates mechanisms that both meet the general high level design requirements, and also account for these phenomena. The theory emphasises the need for the control of attention, where attention includes the selection of information to be processed and the selection of procedures to apply to it. The mechanisms create representations of problems or goals, known as "motivators", and include processes which operate on them. Our explanation will extend ideas in our previous papers on this topic. We shall not here repeat the explanation of how emotional states emerge when there is partial loss of control of attention (see [19, 22].) We are not directly concerned in this paper with the processes underlying pleasure, pain, and creation of attitudes.

We have designed a simplified "nursemaid" domain involving an agent faced with the requirements listed. We propose a partial design for this agent below, and report on a partial implementation in Poplog [3]. This modelling work has already indicated deficiencies in earlier versions of our theory and led to a number of refinements. The models could also have a useful tutorial function.

5. Related Work

Three areas of research are particularly relevant to our theory. Firstly, there is a long history of work in Psychology dealing with affect and attention. For instance, McDougall underscored the importance of purpose in the regulation of human behaviour [13]. Gordon Allport emphasised that human beings have multiple top level ("functionally autonomous") motives [1]. More recently, Julius Kuhl [12] has developed an extensive theory of the regulation of motivation. However, this psychological research is not "design based" [21], and hence lacks some of the rigour and explanatory power that we seek. The work of Norman and Shallice on attention and behaviour is potentially relevant, but it is still too sketchy. It doesn't elucidate in any detail the role of motivation in attention and action.

Secondly, recently, there has been interest in Cognitive Science on agents meeting the requirements mentioned above [7, 14, 15]. Herbert Simon's seminal paper on emotion and motivation [17], which links them to attention, has greatly influenced a number of cognitive theorists of affect, including ourselves. Simon proposes that mental processing can be asynchronously interrupted when humans detect motivationally relevant states or events. Oatley and Johnson-Laird propose that the mind comprises a hierarchy of processors modulated by a top level processor. Asynchronous detection of changes in the likelihood of success of a plan leads to the global broadcasting of a signal which interrupts processing and directs it in predetermined ways to the relevant plan's goal. This occurrence corresponds to an emotional disturbance. A problem with this theory is that its actual scope is larger than its intended scope. The interrupts, for instance, would seem to apply equally well to examples where emotions are *not* generated. That is, global interruption and redirection of effort can occur without any emotional state. Beaudoin and Sloman [4] discuss this argument and apply it to Frijda's work as well.

Thirdly, AI research on activity has until recently mainly focused on systems with a single top level goal. The systems that do have multiple goals tend to have very specific tasks. Georgeff and Lansky's Procedural Reasoning System [9], however, takes seriously the constraints on time and interruptability, and the importance of an agent's production and management of its own goals. However, their provisions for controlling attention are limited compared to what we propose. Moreover the links with motivation and emotion are not explicitly drawn. The same comment applies to Wilensky's theory [23], which emphasises the agent's production and management of multiple goals.

6. Representations and Mechanisms

There is a need for a design based theory that accounts for the control of motivation and attention irrespective of whether emotion-like phenomena are generated. As a result of analysing the general requirements, taking note of a range of common phenomena in human beings, and exploring the needs of a simulated nursemaid, we propose that a minimal set of mechanisms would include the following, which we believe go beyond the sophistication of robot architectures proposed hitherto.

6.1. Motivators

A motivator is a representation of a possible state of affairs towards which the agent has a motivational attitude. This means that the representation has the dispositional power to produce action, though the disposition may be suppressed or over-ridden by other factors. The concept of a motivator is similar to that of a *goal* as frequently used in AI (*e.g.,* [5]), except that its structure is richer, and our concept covers more cases, including ideals, principles, etc. However in this paper we consider only goal-like motivators.

A motivator structure has the following ten fields: (1) a proposition **P** denoting a possible state of affairs, which may be true or false, e.g. "John has eaten"; (2) a motivational attitude to **P**, which is one of "make true", "make false", "keep true", "make true faster", or the like; (3) a value representing what is believed about **P**, which may be "true", "false", "nearly true", "unknown", or some kind of "probability", and which together with field (2) disposes the agent towards certain sorts of actions (so that if the attitude is "make true", and the belief-value is "false", the agent is disposed to try to make **P** true, e.g. if "John has eaten" is believed false, and the attitude is "make true", then the agent may be disposed to get John close to food): (4) an "importance value", such as "neutral", "low", "medium", "high", "unknown", or possibly a numerical value, or a conditional description ("high if so and so agrees"), where the value is based on analysis of the benefits of fulfilling, or the costs of not fulfilling, the motive, or assessment according to ideals or principles, such as ethical or aesthetic principles (though our current simulation ignores principles): (5) an "urgency descriptor", which in simple cases merely measures time till it's too late, or in more complex cases relates time of action to costs and benefits (e.g. without action a motivating situation may get worse over time, making action more important); (6) an "insistence value" heuristically representing importance and urgency, and determining potential to attract attention; (7) a plan or set of plans for achieving the motive (if available); (8) a commitment status, such as "adopted", "rejected", "undecided" (possibly with attached conditions, and possibly with a numerical degree), reflecting decisions that have been taken about the motivator, e.g. whether it is an intention or not; (9) "management information" in the form of a set of condition action pairs associated with a motivator, determining, for example, when action should begin or be resumed (where the actions may be mental or physical); (10) a dynamic state such as "being considered", "consideration deferred", "current", "plan suspended", "plan aborted". A more complete model would add a field for intensity, a measure of the strength of the disposition to act on the motive, and possibly also associated pleasure and displeasure.

Many of the problems of designing an autonomous agent arise out of the fact that a lot of motivators can exist simultaneously in different states of processing, and new ones can be generated at any time, potentially disturbing current processing. When there are several motivators requiring attention that are important, urgent, adopted, and unsatisfied (but possibly satisfiable), the situation will be described as having a high level of "hastiness".

6.2. The processing of Motivators

A number of processes apply to motivators. Before they can be considered, motivators first need to be generated. In highly trained individuals some motivators produce action directly via "cognitive reflexes", but in general a new motive has to be attended to, e.g. to determine its commitment status (8), or form a plan. But considering a motivator requires diversion of attentional resources, and such a diversion can in some cases be dangerous or counter-productive. (For example, being diverted while making a tricky right turn in traffic, or while being given important instructions, can have fatal consequences.)

Therefore, we assume a process which associates with each motivator an "insistence value", which is a quickly computed heuristic measure of importance and urgency, and which determines how likely it is that the motivator will attract attention, e.g. by exceeding a threshold in an attention filter. A motivator that penetrates a filter is said to be surfacing. Note that the role of insistence measures and filtering is to prevent unsuitable diversion of resources. This attention protection mechanism is heuristic and may therefore sometimes fail. That is, in some cases it may prevent pertinent motivators from being considered; in others it may allow relatively unimportant motivators to disrupt attention.

If a motivator manages to penetrate the filter, a further requirement is to decide whether, when, and how it may be acted upon, and whether it ought even to be considered, and if so whether it should be considered now or later. These are the tasks of the management process. At any one time, the management process can be in the midst of adoption-assessment, scheduling, expansion (planning), or meta-management. *Adoption-assessment* aims at deciding whether to adopt or reject a motivator. It is influenced mainly by the motivator's importance and the cost of satisfying it. *Scheduling* aims at deciding when to execute a plan for a motivator. It is primarily influenced by a motivator's urgency. Motivator *expansion* aims at deciding how to satisfy a motivator. It may involve constructing new plans (*e.g.,* [5]) and/or retrieving preformed plans or procedures (*e.g.,* [8]). Motivator expansion is often referred to as 'planning', though planning often involves other processes (*e.g.,* scheduling). Expansion needs to be sensitive to beliefs concerning the instrumentality [10] of possible actions, *i.e.,* the extent to which they increase or decrease the likelihood of the motivator being satisfied. Plans may be partial, with details left to be filled in at execution time. *Meta-management* is described below.

Although we described the different functions of the management process separately, they are often inextricably linked. For instance, *how* a goal may be expanded will depend on *when* it can be acted upon, as well as on how important it is. And when a certain motivator is pursued may impact on the chances of success. Hence each kind of objective of the management process may involve the others. Moreover, the scheduling, expansion, adoption-assessment and meta-management functions may be triggered in any order, contradicting Sloman's [20] earlier suggestion that the state transitions in the post-filtering processing of motivators follow a rigid order. For example, it is possible to start expanding a motivator before even adopting it, in order to assess the costs of adoption.

If and when a motivator penetrates the filter the management process is interrupted and given the objective of producing one of the four kinds of decision regarding the surfacing motivator (*i.e.,* expansion, scheduling, adoption-assessment, or meta-management). This raises the question: how is the system to decide what will be the objective of the management process, given that it will be considering a certain motivator? We assume that there is a mechanism, which we call *dispatching*, which takes this decision very rapidly. Dispatching is a mechanism for controlling attention. It is instructive to compare dispatching to deliberation scheduling [6], which quickly determines "what to think about when" (*p.* 50). The function of dispatching differs from deliberation scheduling in that (1) rather than deciding *what* goal to think about, it determines *how* to think about it, and (2) it takes a decision about a motivator which is already guaranteed to be considered next; *i.e.,* it doesn't have to decide *when* to think about something. Variants of dispatching are explored in the task domain described below.

In contrast to dispatching, the management process may take arbitrary amounts of time and resources to produce a decision. However, urgency and hastiness (defined above) may affect its decision time. The management process is a major bottle-neck in the processing of motivators, because it requires the most sophisticated decisions to be taken. In order judiciously to utilise the management process, we assume that it itself may (in some circumstances) decide whether and when a motivator ought to be managed, and (possibly) how it ought to be managed. This function of the management is referred to as *meta-management*.

For example, while trying to decide how to deal with a baby's rash a nursemaid might start to wonder what to do about another baby who has been misbehaving. She might then give herself a meta-management objective, to decide whether she ought to be thinking about the sick baby or the misfit. Meta-management is often useful when there are current demands on the management process and a motivator which is tangential to the management process surfaces. It is also useful under those conditions where the filter threshold should be high (for instance when distraction can have dangerous side effects). Of course not every process can be preceded by meta-management processes, since then nothing could ever start. So normally some 'routine' management strategy will be followed automatically.

The meta-management function is similar to that of filtering, *i.e.*, to protect attention. Filtering requires a simple and rapid mechanism, whereas meta-management can use whatever inference engine is available to the management process and therefore it can take longer to terminate.

7. The Nursemaid Domain

In order to explore our theory we designed a domain in which an autonomous agent is given tasks the requirements of which are similar to those listed in the introduction. It was essential to simplify the domain in order to make the problems tractable without first solving all the problems of AI, including 3-D vision, motor control, and naive physics. We therefore chose a domain that presents the problems we wished to address while avoiding problems best left to others. The domain involves a robot nursemaid which must look after a collection of robot babies roaming around a two dimensional nursery. The babies can get into various difficulties, like falling into fatal ditches, running out of battery charge, becoming sick, ill, or violent. The nursemaid can detect these problems using a simulated camera which can observe a part of the nursery at a time. And she can respond to them by using her simulated claws to transport babies, move them away from ditches, bring them to a battery recharge point, bring them to an infirmary, isolate them if they are violent, or dismiss them if they are dead or have reached a certain age. The actual 'physics' and 'psychology' of the domain can be extended indefinitely as required for testing later more complex versions of our theory.

In a typical scenario, the nursemaid detects that a baby, say Ba, has a low charge. As the nursemaid is starting to expand a solution to this problem, she notices that another baby, Bb, is dangerously close to a ditch. She decides to take care of the problem concerning Bb first, since it is more urgent. As she is expanding a solution to this urgent problem, she perceives that Bc is ill. This generates a new motivator to cure Bc. But this motivator is not noticed because of an attention filter, and Bc later dies of its illness because it wasn't attended to early enough. We shall revisit this scenario in the light of a description of the nursemaid's architecture. The domain should not be interpreted as having any political, social, or economic significance. It merely happens to embody many of the problems that interest us. Moreover, it can be extended gradually by adding new complications, requiring increasing architectural sophistication.

8. Design of the Initial Nursemaid

The initial design of the nursemaid comprises a number of databases, including a world model and a motivator database. There are simultaneously active modules, including a perception module, motivator generators, an attention filter, a dispatcher, a motivator manager, a motivator monitor, and a plan executor. There are two sorts of effectors: claw controllers and gaze controllers. A perceptual module collects information from the nursery and incorporates it into the nursemaid's world model, which is distinct from the nursery and may contain dated, incorrect, information.

Motivators are represented as objects with the ten components listed previously, though some of the permitted values are simplified for the first prototype. For example the commitment status (8) has only two options, "true" or "false", urgency isn't explicitly represented and importance (4) is represented numerically. The effect of having condition-action pairs is achieved by maintaining a schedule of motivators to execute or consider, as described below. Insistence (6) is represented as a number between 0 and 1. The dynamic state of a motivator (10) denotes whether it is being managed and/or executed. Other information about the motivator is contained in the motivator database (see below).

Motivator generators respond to the world model and information about extant motivators and generate new motivators when the appropriate conditions are met, including allocating a numeric insistence value. For example, if a baby is dangerously close to a ditch, then a motivator is generated with an insistence value which is proportional to the baby's proximity to the ditch. It is possible for more than one motivator to be generated or activated at a time. We had to decide how to handle the case in which some information calls for a motivator generator to produce a motivator whose proposition and attitude are the same as an extant motivator. This raises the issues of individuation and recognition of motivators. We arbitrarily decided to individuate motivators by their propositions and attitudes, and to make them unique, so that in the case in question the motivator generators would *activate* the existing motivator rather than generate a copy.

There is a variable threshold *filter* and a filter threshold *updater*. Recall that the role of the filter is to prevent insufficiently insistent motivators from disrupting the management process. Motivators whose insistence are less than the filter's threshold are ignored, the others are passed on to the dispatcher thereby diverting attention from other activities. The filter updater varies the filter threshold as the hastiness of the situation varies.

The *dispatcher* interrupts the management process when a motivator surfaces, and directs it to the new motivator, while specifying what the management's objective should be (either scheduling, adoption-assessment, expansion or meta-management). Different dispatching rules are being studied; here is an arbitrary example. If the hastiness is high, or the motivator has been postponed until some conditions which aren't currently true, then favour a meta-management objective for this motivator. This will allow the management to postpone consideration of the motivator, which it is likely to want to do if the motivator has already been scheduled for later execution. Or the dispatcher could randomly choose a type of management objective which hasn't already been achieved for this motivator.

The *management process* decides whether, when, and how to act on a motivator, by using either an adoption-assessment, a scheduling, or an expansion routine, respectively. Motivators are adopted unless they conflict with others in such a manner that only one of the conflicting motivators may be executed at all. Scheduling decisions are recorded in the motivator database (described below). These decisions may be to execute a motivator before or after another, to execute a motivator at a certain time, or to execute a motivator when some other condition is met. For instance, the management may decide to recharge a baby when there is sufficient room in the infirmary. Expansion is currently simulated by retrieving pre-stored plans and requiring that this retrieval be extended in simulated time. *Meta-management* provides a means for the system to control its management of motivators. It may decide to postpone consideration of a motivator or to manage it now. And it may decide on the kind of management process which is most appropriate.

The *motivator database* contains information about extant motivators. A schedule contains triples, referred to as schedule items. The first element of a *schedule item* is its activation-condition; this is a proposition, which evaluates to true or false. The second is its action type, which is an expression recording what kind of action should take place, that is either "physical action" or a management objective. The third item is a list of motivators. When the activation-condition of a schedule item is found to be satisfied, its motivators are executed or considered, depending on the action-type. The schedule is updated by the management process, and accessed by the management and motivator monitor (described below). There is a list of *conflicting motivators* known to be incompatible (typically because two motivators have a high urgency). And there is a list of importance *priorities*, which contains partial orders of motivators. The latter two lists are also set by the management process which can, in the process of adoption-assessment, detect and resolve conflicts.

When the management process is interrupted, a record of its state is kept in the *management record*. This contains information about the motivators with which the process was concerned, as well as the objective of the process (scheduling, adoption-assessment, expansion, or meta-management), and whatever intermediate conclusions may have been produced. The next time the management is directed to process the same motivators with the same objective, the record may be used to resume the previous processing.

A *motivator monitor* examines the information about motivators and can activate some of them, generate new motivators, and trigger the plan executor. It examines the schedule to determine if there are motivators whose plans should currently be executed. If it discovers more than one, it generates a 'management motivator' noting the conflict. (So the monitor is also a motivator generator.) If it discovers exactly one motivator to be executed, it triggers the *plan executor*, which executes the plan associated with the motivator, or creates a motivator to create a plan if necessary. The motivator monitor may also activate motivators whose conditions of execution haven't been set, or which haven't been considered for a while. This activation will cause them to go through the usual process of filtering and possibly dispatching and management.

9. A Scenario in the Nursemaid Domain Revisited

In this section we expound the nursemaid's processing in the scenario described above. The nursemaid receives information about Ba which she incorporates into her model. A motivator generator responds to the belief that the baby's charge is of a certain level by generating a motivator the proposition of which is "low_charge(Ba)," and the attitude of which is "con". The insistence of this motivator is greater than the current filter threshold, so this motivator surfaces. The dispatcher directs the management process to decide whether to adopt the motivator. The management decides to adopt it, and proceeds to schedule it for immediate execution. But execution requires that the motivator be expanded. An expansion process is therefore dispatched. However, as this expansion process is running, a motivator regarding Bb's proximity to the ditch surfaces. The dispatcher interrupts the management process and instructs it to schedule the "Rescued(Bb)" motivator. The hastiness of the situation is recorded and this drives the filter threshold up. While all this is happening, perceptual information about another baby, Bc, being injured is inserted into the nursemaid's world model, which generates a motivator to cure Bc. However, this motivator's insistence isn't sufficiently high for it to surface. The management decides to execute the motivator concerning Bb now, and postpones the one concerning Ba.

10. Limitations and Future Work

Several limitations still need to be addressed, only some of which can be mentioned here. The theory needs to be more specific about how different kinds of objectives of the management process are to be attained. Existing AI work on plan formation could be incorporated in *goal expansion*. There has also been a lot of work on scheduling, though not so much on scheduling in autonomous systems with their own objectives. The processes of adoption-assessment and meta-management should be more principled. We also need to state more clearly how high-resource processes differ from low-resource processes and how resource limits arise.

Representation and processing of motivators are still too limited. When scheduling the nursemaid considers only the problems that are actually present. A better version would make decisions that are sensitive to the probability that other problems might occur. Moreover, although in real life some motivators have a hierarchical structure (cf. [16]), the existing nursemaid's motivators do not. (This is partly due to the fact that the prototype nursemaid matches motivators against opaque, pre-stored plans.) Ideally, motivators that are necessary and/or sufficient for satisfying other motivators should be treated differently.

The nursery domain is a deliberately simplified initial test-bed for some general ideas about intelligent agents. As our research progresses, the domain should be made more demanding with regard to the requirements of interest. For instance, at present all of the conflicts in the domain are due to resource limits of one sort of another rather than uniqueness constraints. (A uniqueness constraint states that only one object can be in a certain type of relationship.) There are many more possible developments. For instance, we could allow for more kinds of sensors; more different kinds of babies, with degrees of likeableness, the nursemaid getting attached to some babies. We could allow more nursemaids, and add ideals and ethical principles to the nursemaid. Many forms of learning could be investigated to reduce the arbitrariness of decision procedures and thresholds. Genetic algorithms might be used to evolve different forms of the architecture in different environments.

We conjecture that many emotional states, such as grief or elation, involve processes in which certain high-insistence motivators or thoughts tend to divert attention from currently adopted motivators of higher current 'importance'. We expect to be able to study emerging emotional phenomena in a more advanced model producing 'insistent' motivators that tend to resurface even if the nursemaid has previously rejected, postponed, or scheduled them. We hope that this will help us to understand phenomena like the cases described in the introduction, where people seem partially to lose control of their own thought processes and attention. It should also provide the foundation for a design based explanation of anxiety disorders [2], such as obsessive compulsion disorder.

We do not, however, claim that there is a unique architecture even for human-like systems, so further work should relate design variations to evolutionary pressures, variations in human development, and personality differences.

References

[1] Allport, G. (1961), *Pattern and growth in personality,* New York: Holt, Rinehart and Winston.

[2] *American Psychiatric Association (1987), Diagnostic and statistical manual of mental disorders,* 3rd Ed Rev. Washington, D.C.: APA.

[3] Anderson, J., ed. (1989), *Pop-11 Comes of Age,* Ellis Horwood: Chichester.

[4] Beaudoin, L. and A. Sloman (1991), A proposal for a study of motive processing (Cognitive Science Research Paper No. 91-6), School of Computer Science, University of Birmingham.

[5] Cohen, P.R. and E.A. Feigenbaum, ed. (1982), *The handbook of artificial intelligence* Vol. III. William Kaufmann, Inc.: Los Altos, CA.

[6] Dean, T. and M. Boddy (1988), An analysis of time-dependent planning, *AAAI-88* AAAI.

[7] Frijda, N.H. and J. Swagerman (1987), Can computers feel? Theory and design of an emotional system, *Cognition and Emotion,* 1: 235-257.

[8] Georgeff, M.P. and A.L. Lansky (1986), Procedural Knowledge, *Proceedings of the IEEE,* 74(10), 1383-1398.

[9] Georgeff, M.P. and A.L. Lansky (1987), Reactive Planning, *AAAI-87,* Seattle, WA:

[10] Heckhausen, H. and J. Kuhl (1985), From wishes to action: The dead ends and short cuts on the long way to action, *Goal directed behavior: The concept of action in psychology,* M. Frese and J. Sabini, eds. Lawrence Earlbaum Associates: London.

[11] Horvitz, E.J. (1987), Reasoning about beliefs and actions under computational resource constraints, *Third AAAI Workshop on uncertainty in artificial intelligence,* Seattle, WA: Elsevier Science Publishers.

[12] Kuhl, J. (1992), A theory of self-regulation: Action versus state orientation, self-discrimination, and some applications, *Applied Psychology: An international review,* 41(2), 97-129.

[13] McDougall, W. (1923), *Outline of psychology,* New York: Scribner.

[14] Oatley, K. (1992), *Best laid schemes: The psychology of emotions,* Studies in emotion and social interaction, ed. P. Ekman and K.R. Scherer. Cambridge: CUP

[15] Oatley, K. and P.N. Johnson-Laird (1987), Towards a cognitive theory of emotions, *Cognition and Emotion,* 1: 29-50.

[16] Ortony, A., G.L. Clore, and A. Collins (1988), *The cognitive structure of emotions,* Cambridge: CUP

[17] Simon, H.A. (1967), Motivational and emotional controls of cognition, *Psychological Review,* 74: 29-39.

[18] Sloman, A. (1985), Real time multiple-motive expert systems, *Expert systems 85,* M. Merry, ed. Cambridge: CUP, 1-13.

[19] Sloman, A. (1987), Motives, mechanisms and emotions, *Cognition and Emotion,* 1: 217-234.

[20] Sloman, A. (1991) AI, Neural Networks, Neurobiology, Architectures and Design Space, part of Special Feature on ‘Hybrid Models’, *AISB Quarterly,* No. 78 Autumn/Winter 1991, 10-13

[21] Sloman, A. (1993), Prospects for AI as the General Science of Intelligence *This volume*

[22] Sloman, A. (1992), Prolegomena to a theory of communication and affect, *Communication from an Artificial Intelligence Perspective: Theoretical and Applied Issues,* A. Ortony, J. Slack, O. Stock, eds. Springer: Heidelberg, 229-260.

[23] Wilensky, R. (1983) *Planning and understanding: A computational approach to human reasoning,* Reading, MA: Addison-Wesley.