

## Alarms: Heuristics for the control of reasoning attention

**Timothy J. Norman**

Department of Computer Science,  
University College London,  
London, WC1E 6BT, UK.  
T.Norman@cs.ucl.ac.uk

**Derek Long**

Department of Computer Science,  
University College London,  
London, WC1E 6BT, UK.  
D.Long@cs.ucl.ac.uk

### Abstract

Agents in the real world must be capable of autonomous goal creation. One effect of this ability is that the agent may generate a substantial number of goals, but only a small number of these will be relevant at any one time. Therefore, there is a need for some heuristic mechanism to control an agent's reasoning attention. Such a mechanism is presented in this paper: alarms. Alarms serve to focus the attention of the agent on the most salient goals, and thereby avoid unnecessary reasoning. In this way, a resource-bounded agent can employ modern planning methods to effectiveness.

### Introduction

It is widely recognised that if an autonomous agent is required to interact with a real-world domain,<sup>1</sup> a static set of goals is not a sufficiently flexible representation of the agent's purpose (Brooks, 1986; Carbonell, 1982; Long & Fox, 1995; Simon, 1967; Sloman, 1987; Wilensky, 1983). The domain may change at any time such that pursuing a goal may no longer be realistic, required, or even possible. A single goal may need to be satisfied more than once, or periodically, depending on how the domain evolves over time. So, an autonomous agent must have the ability to set itself goals (cf. Luck & d'Inverno (1995)). Such a capability has been investigated to varying degrees and for different purposes in a number of systems including PANDORA (Wilensky, 1983), Pengi (Agre & Chapman, 1987), PRS (Georgeff & Lansky, 1987), and NML1 (Beaudoin & Sloman, 1993; Beaudoin, 1994). In general, goals are set in response to changes in the domain that are relevant to the agent. The work presented here is motivated by two observations about the nature of goals that have important consequences in models of agency.

1. The time at which changes in the domain that may lead to the generation of a goal are detected is not necessarily the time at which the agent should act on that goal. Action may be required some time in the future.
2. Some goals tend to recur periodically, or at particular times of the day/week/etc. (Goals that recur in this way have been referred to as cyclical satisfaction goals (Schank &

<sup>1</sup>The domain of an agent is its internal state and the state of the external environment that the agent perceives. A real-world domain is a domain that can neither be completely nor correctly modelled.

Abelson, 1977) and replenishment goals (Ortony, Clore & Collins, 1988).) These goals persist; they are not abandoned once achieved, their influence is simply mitigated. Then, as time passes, the intensity of the influence increases until the goal recurs. However, the important role that time plays in the generation of these types of goal has not been fully appreciated.

An agent that is capable of autonomous goal creation, potentially can have an unlimited number of goals. However, the number of goals for which action is required *now* will be a small number of these. Furthermore, all real agents are resource-bounded (Simon, 1957), so there is a limit on the number of goals that can be attended to at any one time.<sup>2</sup> An additional complication is that the process of distinguishing between goals that warrant attention and those that do not, itself depletes resources. Therefore, in any resource-bounded agent an heuristic mechanism is required to focus resources on the most salient goals; i.e. to *avoid unnecessary reasoning* (Norman & Long, 1995). Such a mechanism has been characterised as "fast but stupid" by Sloman (1987). An heuristic mechanism for the control of reasoning attention is presented in this paper: *alarms*. This mechanism has been implemented and tested on a simulated warehouse agent, and intuitive examples from this domain are used throughout the paper.

### Goals

In AI planning systems, a goal is seen as a proposition, or a well formed formula in the world model, to be made true. The purpose of the planner is to search for a sequence of operators (a plan) that will transform the current state into one where a given goal or goals hold. Planning research is principally concerned with the creation of good plans to satisfy a conjunction of such goals in an efficient way. However, the planning problem quickly becomes intractable as the number of goals that must be planned for increases.

Planning is also an important cognitive activity, and planning systems are useful in explaining certain types of human decision making. Explanations of human decision making

<sup>2</sup>In fact the load on the agent's reasoning resources is bounded, and this load is related to the number of goals that hold attention. However, some goals are more easy to attain than others, so this is not necessarily a simple relationship.

typically concentrate on the creation and management of multiple goals in every day situations, and less on how these goals are represented or achieved (Beaudoin & Sloman (1993) and Hayes-Roth (1995) are exceptions). Typically, various types of goal are identified. For example, Ortony, Clore & Collins (1988) specify three broad goal types: replenishment goals, active pursuit goals and interest goals (cf. Schank & Abelson (1977) and Carbonell (1982)).

In general, goals are the problems to be solved by an agent through some sort of planning capability. In the work presented in this paper, a goal is viewed as a proposition to be made true through purposeful action, but goals can be generated through different processes. It is different processes through which goals are generated that distinguish the two goal types discussed here. The term D-Goal is used to refer to goals generated through decision, and the term R-Goal is used to refer to goals generated through replenishment. It is important to note that goals from different sources are indistinguishable in content and function after they have been generated. These terms D-Goal and R-Goal refer to the goal *and* the process that generated the goal.

## D-Goals

Ortony, Clore & Collins (1988) describe active pursuit goals (analogous to D-Goals) as states of affairs that the agent wishes to achieve under certain conditions. Necessarily, the agent must have some reason behind the generation of such a goal. In the system presented here, simple recognition mechanisms, similar to “Noticers” (Wilensky, 1983), are used to trigger the agent to consider generating D-Goals. Demons (see figure 1) monitor the external environment and internal state of the agent, reporting to appropriate mechanisms when some set of conditions hold. Some of these demons are dedicated to recognising events that may warrant the generation of a goal. For example, if an agent is informed of a meeting, this event is reported, and a goal may be generated in response.<sup>3</sup> However, the time at which the agent recognises that something must be done is not necessarily the time at which its attention should be directed towards doing it. So, goals may be generated through such a mechanism that are not appropriate to the agent for some time.

D-Goals have a limited life: They are generated, planned for, and once they have either been satisfied or are no longer required, they are deleted. For example, the warehouse agent may receive a request for an order to be satisfied from a potential customer. If the order is accepted, a goal is generated; then if the order is satisfied or if the agent no longer wishes to satisfy the order, the goal is deleted.<sup>4</sup>

---

<sup>3</sup>A decision based on the agent’s beliefs is a prerequisite to the generation of a goal in this way (Castelfranchi, 1995), hence the term D-Goal.

<sup>4</sup>The deletion of a goal for whatever reason may influence other processes, or even cause the agent to consider generating other goals.

## R-Goals

The second type of process through which goals can be generated is replenishment. Replenishment is an autonomic process, and hence does not involve reasoning. Examples: (1) An agent that has the desire to interact with another agent with known behavioural patterns may synchronise its activity with the other agent for its own ends; (2) An agent that is concerned with maintaining the state of a domain variable may periodically affect the variable to keep it within acceptable bounds. Replenishment processes tend to cause the same goals to recur over certain periods of time or at certain characteristic times (e.g. at 5pm every day, or every Thursday). There are two distinct reasons for R-Goals to influence an agent’s behaviour:

1. An agent in a real-world domain is not the sole actor in that domain. Influences that are independent from the agent may produce regularities in the domain; e.g. a lecture timetable. For the agent to affect the domain in useful ways, it may be necessary for it to synchronise its activities with aspects of the environment that are not under its control. So, a student who has the desire to pass a course will be influenced by the recurring goal to have attended particular lectures.
2. It is advantageous for agents to generate goals in this way, rather than through decision. (The process of generating a goal through decision is computationally expensive.) For example, if a warehouse agent, through some autonomic process, causes goals to have restocked a commodity to recur every Thursday, in normal circumstances there is no need to even consider that commodity during the rest of the week. Replenishment processes allow reasoning resources to be redirected to more constructive tasks.

## Alarms

Focusing attention is an important requirement for an agent capable of autonomous goal generation for two reasons: (1) The agent will be influenced by a large number of goals, but only a small number of these goals will warrant attention at any one time; and (2) Resource limitations bound the number of goals that can be considered, or planned for, at any one time.

This section describes alarm management processes; these have the effect of focusing an agent’s reasoning attention on appropriate goals at appropriate times. In the most simple terms, an alarm ( $a_j$ ) is a structure that associates a goal ( $g_j$ ) with a function of intensity over time:  $a_j = (g_j, f_j(t))$ . This function is designed to reflect how appropriate the goal is to the present situation. Typically, as time passes the intensity of this function increases to a point where it exceeds some threshold. Then, the attention of the agent is focussed on that goal. This threshold changes when the situation changes; as the agent becomes more busy, the threshold increases and vice versa. Only when an alarm has triggered in this way will the goal be considered, and possibly activated. Alarms are essentially goal management processes that serve to avoid unnecessary reasoning.

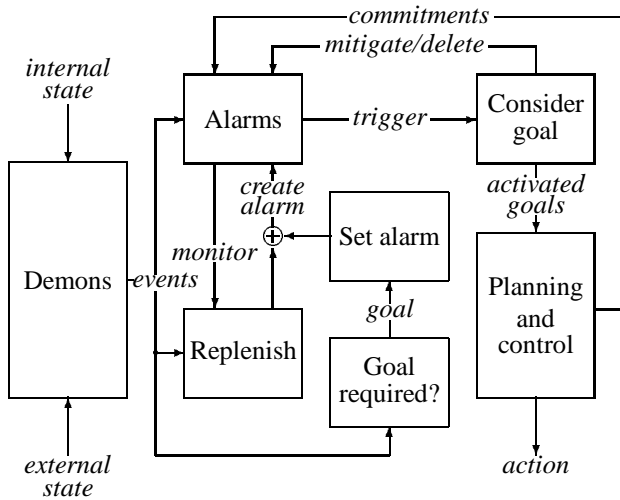


Figure 1: An architecture for goal creation and management.

### Alarm generation

**D-Goals:** When the agent considers that an event (detected by a demon) warrants the creation of a goal, an appropriate alarm is set (see figure 1). For example, the warehouse agent contains a demon that will respond to any message received from a customer. If the customer has placed an order, and this order is acceptable to the agent, then a goal to have met the order is created. The agent predicts when this goal needs to be achieved, how long it will take to achieve, the importance of the goal, etc. With this information, a function of intensity against time is defined, and an alarm is created.

This function of intensity (figure 2) is specified by defining the variables  $t_{dt}$ ,  $\Delta_a t$ ,  $t_{max}$ ,  $i_{dt}$ , and  $i_{max}$ . The deadline ( $t_{dt}$ ) is the time at which the agent wishes the goal to have been satisfied.  $\Delta_a t$  is the period of time that the agent expects will be required to act to satisfy the goal. With this value and the deadline, an estimate can be made of the last point at which the agent should attend to the goal to have it satisfied in time, and hence the time at which the intensity of the alarm should be maximal:  $t_{max}$ . The delay time ( $t_{dt}$ ) is a time, before which it is not appropriate for the agent to act on the goal. For example, there is no point in the warehouse agent preparing an order containing perishable commodities if the customer is expected to arrive after the commodities will have perished. The times  $t_{dt}$  and  $t_{max}$  define a time window where the agent predicts that it is sensible to activate the goal.

Different goals may have the potential to influence the agent to varying degrees; this potential is the maximum intensity of the alarm ( $i_{max}$ ). A modification of the Little Nell problem (McDermott, 1982) is useful in illustrating the effect of varying alarm potentials. The heroine (Nell) is tied to the tracks, a train is approaching, and the hero (Dudley) wishes to save her. The modification is that Dud-

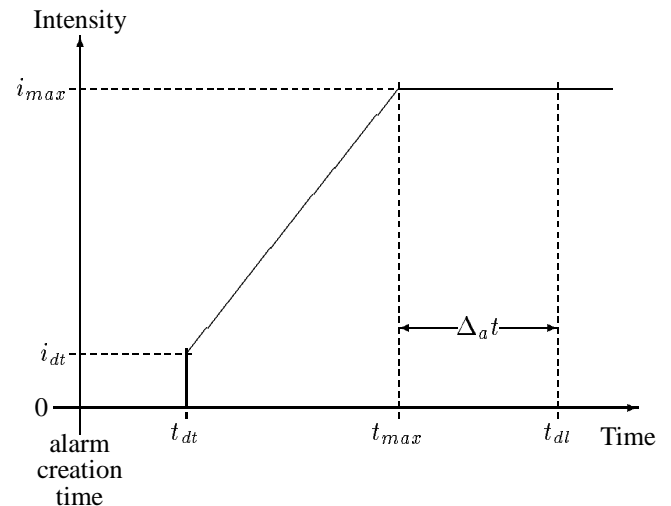


Figure 2: A general alarm function.

ley is not only influenced by the goal to have saved Nell. Dudley also has the goal to have sated his hunger. Furthermore, it is lunch time, so the intensity of the alarm encapsulating the goal to have sated his hunger is near-maximum. Dudley has predicted that Nell will get mashed soon, so the alarm encapsulating the goal to have saved Nell is also near-maximum. It is the fact that saving Nell is far more important to Dudley that distinguishes these two alarms; i.e. the variable  $i_{max}$  for the alarm encapsulating the goal to have saved Nell is greater than for the other alarm. In fact it is possible for a goal to be ignored, even if the alarm's intensity is maximum, if the threshold is sufficiently high (see the extended example).

Finally,  $i_{dt}$  is the initial intensity of the alarm at the delay time. Between  $t_{dt}$  and  $t_{max}$ , the intensity linearly increases from  $i_{dt}$  to  $i_{max}$ . It is possible to define different criteria for an alarm function, but this depends heavily on the accuracy of the predictions, or the information they are based on.

**R-Goals:** The visible effect of an autonomic replenishment process is to generate a stream of alarms, and hence activated goals, at appropriate times. Each goal is treated in the same way as a goal generated through decision by the planning and control mechanisms. The goal will be activated, and acted on in the usual way. The function of the replenishment process is to continually monitor the existing alarms. If there is an R-Goal that is missing (due to it having been deleted), a new alarm encapsulating that goal is automatically created (figure 1).

The function of intensity of this replenished alarm will increase to maximum over some period of time. This period may be fixed and the intensity of the alarm increases as the time since it was last satisfied increases. (This type of mechanism is consistent with the observations of Ortony, Clore & Collins (1988) about the nature of replenishment

goals.) For example, the intensity of the alarm encapsulating the goal to buy milk may increase as the time since the last visit to the shops for this purpose increases, based on some estimate of how long a pint (or quart) of milk lasts. However, the specification of a fixed period of replenishment does not sufficiently express certain types of R-goal behaviour. If hunger is influenced primarily by social habit, the alarm could be replenished to increase in intensity over the periods of time between 08:00, 13:00, and 18:00, causing goals to be activated around these times in the day. This is an example of a timetabled R-Goal. Note, this does not preclude the possibility of a D-Goal to mitigate hunger being generated at any time.

This R-Goal to satiate hunger is an example of an automatic replenishment process that is essentially permanent. However, an agent may have a number of R-Goals that only exist under certain conditions; these are temporary R-Goals. For example, the warehouse agent may accept regular orders from reliable customers under certain conditions. These regular orders become temporary R-Goals, and will only exist if the agent continues to believe that the customer is reliable and that the customer still requires the order. Demon processes are again employed to notify the replenishment process that the conditions for the existence of a temporary R-Goal no longer hold. In general, an agent will synchronise its activities with regularities in the environment, or with other agents, only if the agent considers it advantageous to do so.

### Alarm modification

A real-world domain is intrinsically uncertain; predictions are fallible and the normal replenishment of certain goals can become inappropriate. At any time after an alarm has been set, the domain may change in an unexpected way. Such changes in the domain may directly cause changes in the intensity of an alarm.

**Opportunities:** An opportunity is an action for which all the preconditions of that action hold in the present state, and that the achievement of a goal encapsulated in an existing alarm is a postcondition of that action. By taking an opportunity, it is possible for the agent (if all goes well) to satisfy the goal without the need for further planning. The agent is equipped with a set of opportunity demons, where the conditions for the triggering of that demon are the preconditions of the relevant action. The detection of an opportunity has the effect of giving the appropriate alarm an impulse of intensity.<sup>5</sup> It is possible for such an impulse to increase the intensity of an alarm over the threshold, cause the alarm to trigger, and hence possibly activate the goal. If the goal is activated, the opportunity is communicated to the planning and control mechanisms of the agent.

**Dangers:** A danger to a goal is one of three things: (a) A plan, constructed to satisfy the goal, which has failed; (b) An es-

entially irreversible action<sup>6</sup> that the agent intends to perform which will prevent it from satisfying the goal in time; or (c) All alarms are based on assumptions about how the domain will evolve over time. If the agent detects that the domain has changed such that a salient assumption is no longer valid, this constitutes a danger to the satisfaction of the goal. In the same way as opportunity demons, danger demons notify the agent of a dangerous situation by giving the alarm that encapsulates the goal that is in danger a temporary impulse of intensity. An impulse may cause an alarm to trigger and force the agent to evaluate a possible danger to the goal encapsulated in that alarm.

**Commitments:** In the construction of a plan, the agent commits itself to activity at certain times. Commitments reduce the time available for the agent to act to satisfy other goals. For example, if an agent plans to travel through a desert, whether the agent has recently sated its thirst or not, the goal to have mitigated thirst is appropriate to the situation. Commitments made by the agent have the effect of shifting a function of intensity left along the time axis, and hence the alarm is evaluated at the time  $(t_{now} + \Delta_c t)$ , where  $t_{now}$  is the time now and  $\Delta_c t$  is the effect of the commitments that the agent has made to future action. This alarm may therefore be triggered earlier.

### Alarm triggering

An alarm is triggered, and hence the goal is considered by the agent, if the alarm function evaluates at the time  $(t_{now} + \Delta_c t)$  to a value greater than or equal to some threshold. This threshold is common to all alarms and can have any value above zero. The threshold has the effect of controlling the sensitivity of the agent to alarms, and acts to limit the focus of the agent's attention. If the threshold is low, the agent will be more sensitive to new goals, and vice versa. However, to effectively control the focus of attention, this threshold must change in relation to the changing load on the agent's reasoning resources. As the agent's activity increases, the threshold increases and hence the agent's sensitivity to new goals is reduced. So, if the agent has a large number of urgent tasks, it will tend to leave other alarms that have lower intensity potential to the last minute or ignore them. However, if the agent has fewer tasks demanding attention, alarms that have less intensity potential may trigger (the threshold having been lowered), and goals will generally be considered earlier. The effect of this system of alarms is to focus reasoning attention on a limited number of goals so that the planning and control processes are not swamped with things to do, or even things to think about doing.

Once a goal is activated, it is important for an agent to then reconsider the goal at appropriate times (Bratman, Israel & Pollack, 1988). "Reconsidering a prior intention is an activity that uses up time and other limited resources" (Bratman, 1992), but determining when goals should be reconsidered

<sup>5</sup>This impulse exists only for as long as the opportunity exists.

<sup>6</sup>An essentially irreversible action is one that cannot be reversed in time for the goal to be satisfied in time.

also depletes resources. Therefore, in the same way that an heuristic is required to determine when goals should be considered for achievement, a similar mechanism is required to determine what goals should be reconsidered. For example, as the deadline of the goal approaches, or if a danger to the satisfaction of that goal is detected after it has been activated, it should be reconsidered. When a goal is activated, the alarm management process is not deleted; the effect of activating the goal temporarily mitigates the intensity of the alarm. This mitigation effect decays over time, until the alarm will again be triggered and the goal reconsidered. When a goal is activated, the agent predicts when it should next consider the goal and this determines the rate of decay of the mitigation. In this way, the same alarm heuristic controls both goals to be considered for achievement and goals to be reconsidered.

A goal that is activated by the agent will mitigate the relevant alarm (e.g. the goal to satisfy an order placed by a customer in the warehouse domain). Typically, the goal is activated before the time  $t_{max}$  (figure 2), depending on the threshold. The mitigation effect on the alarm will decay at a rate such that the alarm will again trigger around  $t_{max}$ . At this point the agent will reconsider the goal. If the alarm is not deleted, it is again mitigated until the deadline of the goal ( $t_{dl}$ ). If the goal to satisfy the customer's order is not satisfied and the deadline was correctly predicted, there may be no option other than to delete goal and alarm.<sup>7</sup> This is an example of a goal that is no longer relevant to the agent's alarms once the deadline has passed and it is not satisfied (a goal with a firm deadline). Other goals may still be relevant after the deadline has passed; they may be even more urgent. This type of deadline represents the time at which the agent would like the goal to be satisfied, not the time at which it must be satisfied: a soft rather than a firm deadline.<sup>8</sup> A goal with a soft deadline will continue to be reconsidered periodically until it is satisfied, or the agent decides to delete the alarm for other reasons.

### Managing uncertainty

A planning agent can never know for sure how long it will take to satisfy a goal. If the agent underestimates this time, it may wait too long before acting on the goal and the goal may never be satisfied in time. This is a potential source of error in the alarm heuristic. This period of time can only be known once the goal has been successfully achieved. Therefore, any mechanism for the scheduling of multiple goals may suffer from such errors. This variable ( $\Delta_a t$  (figure 2)) is defined by the time period that the agent will allow for that goal to be satisfied. In fact, such rules-of-thumb are common in decision making. For example, if I intend to travel from one part of London to another, I would simply allow an hour for the journey. However, if I do not know where I must travel to in the process of satisfying my intention, the importance that I

<sup>7</sup>The deletion of a goal may have effects on other parts of the agent; this change in internal state may subsequently cause other alarms to be generated or modified.

<sup>8</sup>The term hard deadline is avoided; this is only relevant for an agent interacting with a real-time process.




Figure 3: Example.

give to satisfying that goal will govern how much time I shall allocate to the goal, and hence how far I could travel.

For the same reason, the time commitments made by the agent in the process of planning cannot necessarily be determined. For example, the duration of the action to travel from point A to point B can only be determined when both these variables are bound. The agent may need to plan to seek further information before this can be done. Again time is allocated to these actions.

### Extended example

Consider an agent that is designed to manage a warehouse. This agent will have various influences on its behaviour including satisfying orders placed by customers, maintaining the levels of stock in the warehouse, ensuring the agent has sufficient energy reserves to act on its intentions, etc. Figure 3 illustrates the behaviour of two alarms in a particular circumstance. These alarms encapsulate the goals to have satisfied an order placed by a customer, and to have restocked the warehouse with a certain commodity. (The time axis has been normalised for clarity.)

At time  $t=0$ , the intensity of the alarm encapsulating the goal to have satisfied the customer exists, but is at zero until  $t=5$  denoting that it is not sensible for the agent to act before this time. This may be due to the order containing commodities that will perish if kept in the unfavourable conditions of the loading bay (e.g. frozen food). After this time (the delay time), the intensity linearly increases until time  $t=20$ . At this point, the agent detects that the customer has arrived earlier than expected. This constitutes a danger to the satisfaction of the goal, and so the intensity immediately increases above the threshold. The agent's attention is directed towards this goal, it determines that it should act, and therefore activates

the goal to have satisfied the customer. The activation of this goal causes the respective alarm to be mitigated. This mitigation decays over time until (at about time  $t=40$ ) the intensity of the alarm has again increased to the threshold, and the attention of the agent is directed towards reconsidering the goal. At this point the agent recognises that at some time between  $t=20$  and  $t=40$ , the goal was satisfied, and so the alarm is deleted.

At time  $t=20$ , the goal to have satisfied the customer's order is activated. It is important that the agent focuses a large amount of its attention towards satisfying this goal; the goal is very urgent and quite important. As the activity of the agent increases, so does the threshold. At time  $t=38$ , the agent has satisfied the goal, its activity decreases, and so does the threshold.

The alarm to restock the warehouse encapsulates a goal that was generated through a replenishment process. The intensity of this alarm increases to maximum (in this case the potential of the alarm is 1) at around time  $t=35$ . However, because the agent is acting on the urgent goal to have satisfied the customer, the sensitivity of the agent to new goals is reduced. The alarm does not trigger the attention of the agent until some time later. At time  $t=50$ , the threshold has reduced sufficiently for the attention of the agent to be directed towards considering restocking the warehouse. The goal is activated, and the alarm mitigated in the usual way.

### Conclusion

An heuristic goal management mechanism has been presented: *alarms*. Alarms serve to focus an agent's attention on the most salient goals at any one time. This concentrates planning and reasoning effort and avoids unnecessary distractions. Modern planning methods can then be employed to effectiveness in the search for solutions to the given problems (the active goals). The focus is continually updated as the situation changes, and limited through manipulation of the threshold. Additionally, goals generated through two distinctive and important types of process have been considered: D-Goals and R-Goals. Goals created through both these processes are effectively managed within the same alarms mechanism.

### Acknowledgements

Maria Fox, Mike Luck and Vijaya Norman gave helpful comments, as did two anonymous reviewers.

### References

- Agre, P. & Chapman, D. (1987). Pengi: An implementation of a theory of activity. In *Proceedings of the Sixth National Conference on Artificial Intelligence* (pp. 268–272).
- Beaudoin, L.P. (1994). *Goal processing in autonomous agents*. Doctoral thesis, School of Computer Science, University of Birmingham.
- Beaudoin, L.P. & Sloman, A. (1993). A study of motive processing and attention. In *Prospects for Artificial Intelligence: Proceedings of AISB-93* (pp. 229–238).
- Bratman, M.E., Israel, D.J. & Pollack, M.A. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4(4):349–355.
- Bratman, M.E. (1992). Planning and the stability of intention. *Minds and Machines*, 2(1):1–16.
- Brooks, R.A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2:14–23.
- Carbonell, J.G. (1982). Where do goals come from? In *Proceedings of the Fourth Annual Conference of the Cognitive Science Society* (pp. 191–194).
- Castelfranchi, C. (1995). Guarantees for autonomy in cognitive agent architecture. In M.J. Wooldridge & N.R. Jennings (Eds.), *Intelligent agents* (pp. 56–70). Volume 890 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag.
- Georgeff, M.P. & Lansky, A.L. (1987). Reactive reasoning and planning. In *Proceedings of the Sixth National Conference on Artificial Intelligence* (pp. 677–681).
- Hayes-Roth, B. (1995). An architecture for adaptive intelligent systems. *Artificial Intelligence*, 72:329–365.
- Long, D.P. & Fox, M. (1995). A hybrid architecture for rational agents. In S. Torrance & C. Thornton (Eds.), *Hybrid models of cognition*. AISB.
- Luck, M.M. & d'Inverno, M. (1995). Agency and autonomy: A formal framework, Tech. Rep. CS-RR-276. University of Warwick, Department of Computer Science.
- McDermott, D. (1982). A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6:101–155.
- Norman, T.J. & Long, D.P. (1995). Goal creation in motivated agents. In M.J. Wooldridge & N.R. Jennings (Eds.), *Intelligent agents* (pp. 277–290). Volume 890 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag.
- Ortony, A., Clore, G.L. & Collins, A. (1988). *The cognitive structure of emotions*. Cambridge University Press.
- Schank, R.C. & Abelson, R.P. (1977). *Scripts, Plans, Goals, and Understanding*. Laurence Erlbaum Associates.
- Simon, H.A. (1957). *Models of man: social and rational*. John Wiley.
- Simon, H.A. (1967). Motivational and emotional controls of cognition. *Psychological Review*, 74:29–39.
- Sloman, A. (1987). Motives, mechanisms, and emotions. *Cognition and Emotion*, 1(3):217–233.
- Wilensky, R. (1983) *Planning and understanding: A computational approach to human reasoning*. Addison-Wesley.