Why You Should Buy an Emotional Planner

Jonathan Gratch

University of Southern California Information Sciences Institute 4676 Admiralty Way, Marina del Rey, CA 90405

gratch@isi.edu

ABSTRACT

Computation models of emotion have begun to address the problem of how agents arrive at a given emotional state, and how that state might alter their reactions to the environment. Existing work has focused on reactive models of behavior and does not, as of yet, provide much insight on how emotion might relate to the construction and execution of complex plans. This article focuses on this later question. I present a model of how agents appraise the emotion significance of events that illustrates a complementary relationship between classical planning methods and models of emotion processing. By building on classical planning methods, the model clarifies prior accounts of emotional appraisal and extends these accounts to handle the generation and execution of complex multi-agent plans.

Keywords

Emotions, cognitive appraisal, planning

1 INTRODUCTION

Though emotions are central to our daily experience, they are peripheral at best in most computational models of intelligent behavior. This tradition is being increasingly challenged on two fronts. First, there is increasing interest in modeling "believable agents" (Bates, 1994) and realistic human organizations (NRC, 1998) which demand some model of how emotions relate to behavior. Second, there is growing evidence that emotions play a distinct and crucial role in human decision making (Damasio, 1994; LeDoux, 1996). Many have thus concluded that emotional processing is indispensable for any model of intelligence, be it human or artificial. If we accept this leap, then what role, if any, do traditional models of decision-making play in modeling "true" intelligence. Are they "inappropriate for building believable agents" (Reilly, 1996)?

Many who have studied emotion have proposed an intimate relationship between emotions and plans (Simon, 1967; Oatley and Johnson-Laird, 1987; Sloman, 1987). Nonetheless, existing computational approaches to emotion have focused exclusively on reactive reasoning methods and have not addressed the relationship of emotions to more proactive reasoning. Such models do not address a number of questions relating emotional processing to the construction and execution of complex multi-agent plans. For example, it is unclear how simply thinking about one's goals could change one's emotional state. Nor is it clear how to manage the complex expectations associated with plans and how expectations interact to influence emotions.

In this article, I illustrate the role classical planning methods can plan in addressing such issues (Chapman, 1987; McAllester and Rosenblitt, 1991). To understand the emotional impact of future expectations, one needs the ability to plan and manage expectations. Planning methods provide simple yet powerful techniques for both forming plans and assessing the impact of new information on future expectations. Additionally, models of emotion can potentially address a number of difficulties facing classical planning methods. Many emotion researchers have proposed that emotions guide planning in the face of conflicting goals, limited resources, and imperfect information (Frijda and Swagerman, 1987; Oatley and Johnson-Laird, 1987), three problems that bedevil existing planning techniques. Emotional models hint at possible solutions and certainly shed a different perspective on these problems. (Interestingly, there is a close relationship between the computations needed for emotional processing and the computations planners use to guide their reasoning.)

This paper makes two contributions. The first is to illustrate the relevance of planning methods to computational accounts of these phenomena. Specifically, I consider the problem of *cognitive appraisal* (Ortony et al., 1988; Frijda 1986), by which an agent evaluates the emotional significance of events with respect to its goals and expectations. Besides building plans, planning algorithms possess many properties that facilitate cognitive appraisal. Building atop these properties eases the task of implementing an appraisal model, and results in one that is more concise and easily understood. I illustrate this by implementing an appraisal model based on Elliott's construal theory (Elliott, 1992) which already supports a limited form of expectation management.

The second contribution is a model of cognitive appraisal called *plan-based appraisal*. Though based on construal theory, it makes some significant departures. Elliott centers appraisals on events and uses domain-specific rules to derive the relationship between events and goals. In contrast, I center appraisals on goals and use a planner's domain-independent functions to determine the significance of events to goals. This approach is closer to Oatley's (1992) notion of emotions being related to changes in the probability of goal satisfaction, and generalizes Elliott's model in

several ways. It extends the theory to handle complex plans and, rather than using domain-specific rules to compute various properties of events, these computations are recast into domain-independent rules that key off of the syntactic structure of the plans in working memory. Thus, whereas construal theory's domain-specific rules are outside the theory, the rules that I propose have precise meanings and can be judged as theoretical claims about the appraisal process.

2 PLANNING

In the model that follows, the emotional significance of events is appraised through a syntactical analysis of an agent's current plans. I build on classical planning models, which, although not without limitations, provide a wellunderstood and coherent account of how complex plans are produced and modified. This is not to say that planning is all there is to intelligent behavior. (A complete agent architecture must support both reactive and deliberative reasoning strategies and I discuss how these strategies interact later in this section.) I do, however, argue that plan-based appraisal is sufficient to account for all emotional appraisals that relate to an agents goals (i.e., cognitive appraisals). How these appraisals influence behavior, however, may be expressed through both reactive and plan mediated action selection (see below). In this sense, I clearly separate where appraisals come from (plans) from how they are used.

Following the classical definition, a plan is some sequence of tasks constructed in the service of a set of goals. For the planning community, a goal has a very specific meaning that is derived from a logical representation of the world. The world is represented as a series of states, each state corresponding to a characterization of the world at some point in time and is described by a conjunction of logical predicates. The standard "blocks world" example describes a stack of blocks as:

On(A, B) & On(B, Table) & Clear(A) & Blue(A) & Red(B) A sequence of state descriptions describes how the world changes over time. A goal describes a desired partial state description. For example, a goal could be to have block A on block B, without specifying the color of the blocks.

It's not enough to achieve a goal. The goal should persist for some time. (When we want a good job, we implicitly rule out the possible getting fired the day we show up for work.) Following this, a *goal* is defined to be a partial state description that must be satisfied with respect to some maintenance constraints (e.g. satisfied over some time interval).

To construct a plan, a planner is provided (1) a set of goals, (2) the current state, and (3) a description of a set of tasks it may use to attain its goals. Planners typically use a variant of the STRIPS formalism proposed by Fikes and Nilsson (1971) to represent tasks where a task is described by its preconditions and effects. Preconditions are conjunctions of predicates, which must simultaneously hold for the

task to be initiated (i.e., a set of goals that must be maintained simultaneously).

2.1 Establishment

A key idea in planning is the notion of establishment. The idea is that for a goal to be satisfied (1) something must make it true and (2) it must persist until needed. Establishment is the process of selecting some specific effect of some task that could be used to achieve the goal. For example, if our goal is On(A,Table) and On(A,Table) is an effect of some task, then we could decide to establish our goal by this effect. The act of establishment can be equated with the act of forming an intention (in the sense of Bratman, 1987), in that the planner intends to use some specific task to achieve its goal. The specific effect of the task is referred to as the establisher of the goal.

Choosing an establisher is only part of satisfying a goal; the effect must also persist until it is needed. When a planner establishes a goal it elaborates the goal's maintenance constraints by asserting that the goal must be maintained from when the effect is achieved to when the goal is needed. Specifically, the goal must be maintained during the interval between the establisher's task and the goal's task.

Planning is a process of detecting and resolving threats to maintenance constraints in the evolving plan. A planning algorithm organizes threats into several classes and provides specific procedures for addressing each of class:

Open precondition threat: there is an unestablished goal (precondition). This is remedied by *simple-establishment*, where the planner chooses an effect of some task in the current plan as an establisher, or *step-addition*, where a new task is added to the plan that has an effect which establishes the precondition. Both remedies introduce a maintenance constraint to protect the establishment of the precondition.

Protection threat: there is an effect of a task that could violate an existing maintenance constraint. This may be remedied by: (1) reordering tasks (known as promotion or demotion) so the problematic effect no longer threatens the maintenance constraint; (2) removing the task with the problematic effect with the hope that goals achieved by the task could be achieved in some other manner; or (3) retracting the establishment that introduced the maintenance constraint

Figure 1 illustrates a plan in the working memory of an agent named Jo. Jo's top-level goal is to have money. Her plan is to first drive to work and then work to obtain the money. The Drive task has two preconditions that are established to the current state. The Work task has one precondition that is established by the drive task. Finally, the top-level goal, which is represented as the precondition of a "top-goal" task, is established by the Work task.

2.2 Dynamic and Social Planning

Classical planning techniques were developed under some restrictive assumptions that would seem to rule out their use



Figure 1: An example plan

in modeling agents interacting in a social and dynamic world. Thus it is worth pointing out that planning algorithms can be more useful than a strict theoretical treatment of their capabilities would suggest. For example, planning algorithms can easily represent and manipulate plans containing inconsistent goals (though this is not typically done). Some have also considered how to integrate plan generation, plan execution, and plan repair into a single planning approach (Ambrose-Ingerson and Steel, 1987). These approaches allow an agent to execute tasks in a plan, even if plans are incomplete or contain inconsistencies. They also extend the representation to model not just what is planned, but what behavior actually occurs in the world.

Planners have also been extended to act in multi-agent environments where they may reason about the plans and intentions of other agents. I specifically work with the CFOR planning system that is designed to support dynamic multi-agent planning (Gratch, 1998; Gratch, 1999). The CFOR planner allows the activities of several agents to be maintained in the same plan memory. The planner can make distinctions such as: these tasks correspond to my plan to work on the truck, whereas those tasks correspond to your plan to drive to the rodeo. The planner uses the same threat detection processes to recognize threats between plans of different agents. In this case, it can reason that you taking the truck to the rodeo violates my maintenance constraint of keeping the truck in the garage while I work on it. The CFOR planner responds to inter-threat plans differently depending on the social relationship between the agents involved, though that capability is not discussed in this article (see Gratch, 1999 for details on social planning).

3 AGENT ARCHITECTURE

Before moving on to a discussion of plan-based appraisals, it is helpful to illustrate the role the planning algorithm plays in the overall behavior of an agent and how it interacts with other behavioral components. Figure 2 illustrates the overall



Figure 2: Overall agent architecture

architecture of an agent. We have explored this basic architecture (without the appraisal component) in the context of developing teams of automated pilot agents inhabiting a synthetic world (Hill, et. al, 1997). An agent interacts with its environment through its sensors and effectors. Raw sensory input is processed by a situation assessor module into a set of declarative facts representing the current state of the world. This current world description serves as input both to the planning module and to an action selection mechanism. This actor may chose to reflexively act on the environment, or it may enable mental actions associated with planning. Any plans developed by the planner also serve as input to the actor, and influence its reactions. Thus, for example, the actor may "reflexively" choose to take the next valid action in its current plans. Currently, the actor and situation assessor are implemented as rule-based systems and do not embody any theoretically interesting structure. The entire agent architecture is implemented in Soar (Newell, 1990), though this point does not have any bearing on the model of appraisal.

The appraisal component derives the agents current emotional state using as input the agent's current plans and current world description. In theory, the emotional state can be used as input to the planner, modulating the way plans are developed, and as input to the actor, modulating the selection of external actions. Neither of these capabilities are implemented at this time and I discuss some possible approaches at the end of the article. The architecture is neutral as to whether there may be other, non-cognitive inputs to the agent's current emotional state. For example, one might allow a somatic appraisal mechanism that acts in parallel with plan-based appraisals in determining an overall emotional state.

4 CONSTRUAL THEORY

Cognitive appraisal involves evaluating the significance of events with respect to goals and expectations. But evaluating the significance of events is also a core function underlying planning. What differs between cognitive appraisal and planning is the intended meaning of "significance". What I show is that the representation of expectations (in the form of plans) and a planner's inference mechanisms (such as threat detection) provide a powerful and well-understood infrastructure, some of which may be leveraged when building a cognitive appraisal mechanisms. As an illustration, I show how to re-implement Elliott's (1992) construal theory, which, in turn, is a computational account of Ortony et al.'s (1988) theory of cognitive appraisal. As the reimplementation significantly enriches construal theory's ability to reason about complex plans, I refer to the resulting model as plan-based appraisal.

In construal theory, events are matched against knowledge structures called *construal frames*. These frames evaluate events against an agent's goals, social standards (norms of behavior), and preferences (the appealingness of domain objects). Construal frames make two determinations. First they assess if the event is of relevance to the agent. If so, they extract several high-level features of the event that are later used to assess the emotional response. Collectively, these features are referred to as *an emotion eliciting condition relation*. The individual conditions of the relation are:

Self: The agent that the event is judged with respect to.

Desire-self: Is the event considered desirable to "self"?

Status: Represents the status of an expectation. Expectations can be *unconfirmed*, *confirmed*, or *disconfirmed*.

Evaluation: Does the event uphold or violate social standards. Values can be *praiseworthy* or *blameworthy*.

Responsible agent: The agent that upheld or violated a social standard, if any.

Desire-other: The desirability of the event to other agents. Values can be *desirable* or *undesirable*.

Other: Agent used in assessing "desire-other"

Pleased: Self's reaction to the "desire-other" status of another agent

Appealingness: Does the event contain an attractive or repulsive object a judged by self's preferences?

Emotion eliciting condition relations serve as the input for an emotion generating process. Elliott provides a set of domain-independent rules based on Ortony et al.'s work for mapping emotion eliciting condition relations into several possible emotion classes.

5 PLAN-BASED APPRAISAL

A classical planning view of an agent's mental state nicely extends two aspects of construal theory. First, construal theory only considered agents with reactive reasoning, and thus has difficulties representing future expectations. Construal theory does manage some expectations through largely domain-specific mechanisms. Planning techniques, in contrast, provide a much cleaner, more powerful, and domain-independent solution to managing expectations.

Second, although emotion eliciting condition relations provide a useful structure to the appraisal process, the theory does not provide much guidance in how to derive the value of these conditions. For example, construal frames contain domain-specific rules to compute the desirability of events. How to make this assessment is beyond the scope of the theory. The plan-based re-implementation, in contrast, computes these conditions using a few domain-independent rules that key off of the syntactic structure of the plans in working memory. Thus, whereas Elliott's construal frames are outside the theory, the rules that I propose have precise meanings and can be judged as theoretical claims about the appraisal process. For example, the desirability of an event is based on the extent to which it introduces threats (in the technical sense defined above) to an agent's goals.

5.1 Events vs. Goals

Plan-based appraisal departs somewhat from the spirit of construal theory, specifically in the handling of events. Construal theory centers appraisal on events and uses the construal frames to derive the relationship between events and goals. In contrast, plan-based appraisal centers appraisals on goals and uses the planner's domain-independent threat detection processes to determine the significance of events to goals. I associate something analogous to a construal frame with each goal in working memory, which is closer to Oatley's (1992) notion of emotions being related to changes in the probability of goal satisfaction. Thus, the construal frame for each goal is essentially a goal monitor.

A consequence of this approach is that an event that benefits or harms an agent may not be immediately recognized as such. Whereas some events create an immediate and direct relationship with goals, others may only become apparent in the course of subsequent planning. For example, an agent that hears a radio report of a jack-knifed trailer on The 5 freeway will not recognize the significance of this information until, after a bit of planning, it realizes that The 5 is the only freeway it can take to Disneyland.

5.2 Emotion Eliciting Conditions

To implement the appraisal scheme, it suffices to provide definitions for each of the emotion eliciting conditions in terms of syntactic features of the current plans in working memory. Elliott defines nine emotion-eliciting conditions. Currently, I offer a partial implementation to give a flavor of how plan representations support the appraisal process. Whereas Elliott's theory distinguishes between goals, standards, and preferences, this re-implementation focuses on goals. I describe a single domain-independent standard, which does not cover the scope of standards that Elliott envisions, and ignore preferences. These simplifications are mainly for convenience, though a more complete treatment of standards and preferences could force a departure from a strictly goal-centered appraisal process. I associate an appraisal frame with each goal in the plan structure. Emotion eliciting conditions describe properties of each goal:

Self: This is the agent that "owns" the goal. All goals are associated with some task. The owner is the agent who executes the associated task. In that an agent can represent plans of other agents, it can also appraise their goals.

Desire-self: A goal can be in a desirable or undesirable disposition. It is considered desirable if (1) it is established and (2) there are no known threats to the goal's maintenance constraints. If these conditions hold, the agent has every reason to believe the goal will be achieved. A goal is in an undesirable disposition if there is no establisher, or the establisher is threatened. In this case, the agent has reason to believe that the goal will not be achieved. In general, this appraisal is simply an expectation and may change as planning proceeds or events occur in the world. Until the goal has actually been achieved, there will always be opportunity for the situation to change.

Status: As the "desire-self" appraisal is only an expectation (until the goal has been achieved), the "status" condition summarizes the status of the expectation: is it subsequently confirmed or disconfirmed. Plans allow more flexibility than a simple binary distinction between confirmation and disconfirmation, however to remain consistent with construal theory I draw a line that seems consistent with Elliott's intent. I describe the undesirable (goal threatened) and desirable (goal unthreatened) cases separately. Say we have a goal where effect E of task1 is used to establish goal G of task2. This would be threatened if some other task3 has an effect that undoes E, AND task3 might occur between task1 and task2. This is a possible or unconfirmed threat. It becomes a confirmed threat if task3 necessarily occurs between task1 and task2. For example, if we have already executed task1 and task3, in that order, and not yet executed task 2, then the maintenance constraint is permanently violated - we can't change history. (Although we still retain the option of establishing the goal in some other manner.)

To say a goal is confirmed unthreatened should be the dual of this: the maintenance constraint is guaranteed not to be violated. However, this is unrealistic for dynamic social environments. Instead we adopt a weaker definition. A goal is confirmed to be unthreatened if (1) it is currently true (we have successfully executed the task that establishes the predicate) and (2) we know of no threats to the maintenance constraints associated with this goal.

Evaluation: This specifies if the event contains a praiseworthy or blameworthy act. This involves reasoning about standards of behavior and this is where I have made my greatest simplification. Currently I model a single standard: "thou shalt not introduce threats into someone else's plans." As it stands, even this single standard is pretty simplistic as it avoids the issue of intent, but it is enough to get things started. In terms of the planning model described

above, this standard is violated if a maintenance constraint of one agent is threatened by a task associated with another agent. It should be relatively easy to model other standards. For example, it could be considered praiseworthy if one agent proposes a task that achieves another agent's goals.

Responsible agent: This specifies who was responsible for the praiseworthy or blameworthy act. This only has a value if one agent threatens another agent's goals.

Desire-other: This specifies any other agent involved in the event. Construal theory allows sophisticated domainspecific reasoning about how events might effect the concerns of other agents. In the implementation I have a rudimentary model of this. Agents can have friends and enemies. It is desirable if something good happens to a friend or something bad happens to an enemy. For each appraised event relevant to "self", the appraisal mechanism searches through a list of relationships and asserts that the event would be desirable or undesirable to those other agents. One subtlety here is that an event shouldn't lead to an appraisal if the agent is unaware of it. For example, my wife might tell me she is taking the care to work. I know that I didn't fill up the tank last time I drove the car. Thus, I can reason that she will be unhappy, but not until she finds out. The planner supports this subtlety, but I have not yet incorporated it into the appraisal mechanism.

Other: This indicates what other agent is involved in the appraisal. Currently it lists whatever agents were involved in computing the "evaluation" and "desire-other" features.

Pleased: This specifies one agent's reaction the state of another agent's goals. The appraisal uses the friendship relationship between agents in conjunction with the desire-self condition of goals. Thus, if the planner represents the goal of a friend, it would be please if the goal is achieved and displeased if it is thwarted.

The values of emotion eliciting conditions are combined into an emotion eliciting condition relation. As in construal theory, my re-implementation uses a set of domain-independent rules to map these relations into emotion classes. I implement a subset of Elliott's twenty-four classes since I don't currently model all nine emotion eliciting conditions. Each instantiation of one of these rules generates an emotion instance. These instances are associated with the agents that "feel" them. They are also associated with specific plans. Thus, an agent could "feel" good but be troubled with the progress of a specific plan.

Figure 2 extends the example begun in Figure 1. Now working memory has been expanded with the knowledge of Bo's plans as well as Jo's. Bo wishes to take the car to the shop before Jo can drive it to work. The planner's threat detection processes recognize a potential goal violation: Bo driving the car to the shop violates Jo's maintenance constraint that the car remain at home till she drives it. The appraisal mechanisms perform several inferences based on this new information. Whereas Jo was previously hopeful



Figure 3: An example of inter-plan threats

that this subgoal would be achieved, she is now fearful that it will be violated. Furthermore, Jo should resent Bo as he has violated a social standard by threatening her (glossing over whether Bo intends this).

6 DISCUSSION

This article claims two contributions. The first is a model of plan-based appraisals that generalizes construal theory in a general way to deal with complex multi-agent plans. This allows an agent to exhibit emotional reactions to future contingencies as well as immediate events and to have these reactions change as the agent reasons out the consequences of events. The broader claim is that classical planning algorithms and representations are well-suited for developing models of cognitive appraisal and result in models that are concise and easily discussed. Planning models provide crisp definitions for terms like goal and expectation. Even if one doesn't agree with the specific definitions, they form a starting point for discussion and many variations have already been explored within the basic planing framework.

6.1 Intensity

Many issues arise when we consider cognitive appraisal over complex plans. Given length constraints I can only elaborate a few of them here. One is the issue that some appraisals should lead to a more intense response than others.

It seems a key aspect in determining intensity, as argued by Oatley, is the probability of goal attainment. Unfortunately, probability of goal attainment is difficult (if not impossible) to assess given complex plan and goal structures, though numerous heuristics have been proposed in the planning and constraint satisfaction communities. Interestingly some of the computational accounts of assessing emotional intensity look quite similar to these heuristics as they involve syntactic analysis of goal and domain-theory structure (Sloman, 1987; Beaudoin, 1995)

This suggests that there is a possibility of fruitful interchange between work in search control heuristics for planners and work on emotional intensity. Search control heuristics try to assess what are the key problems to work on next, which is a role that some like Damasio and Oatley have argued that emotions play. More on this point later.

6.2 Reasoning about Belief

The plan-based view forces us to confront the issue of who knows what as it relates to appraisal. Research into multiagent planning has long recognized the need to distinguish between what I know vs. what I know that you know when developing plans in a social world (Grosz and Kraus, 1996). Computational models of appraisal have largely avoided this issue by focusing on reactive reasoning mechanisms. In contrast, by allowing us to represent the plans of other agents, the plan-base model forces us to distinguish between activities which are known or unknown by other agents. The example in Figure 2 illustrates this. Imagine that this example represents Bo's view of the world: he knows that Jo plans to drive the car to work and that he plans to drive the same car to the shop. For Bo to assess Jo's emotional state he must know if Jo knows his plans. If she does, she is angry. If she doesn't she will be angry but currently isn't. In fact, construal theory cannot even make this distinction. The planner supports such a distinction by providing the means to keep track of which agents know which activities.

6.3 Hot vs. cold emotions

One distinction that is oft made in the emotion literature is between hot and cold emotions. Cold emotions are characterized as dispassionate cognitive process. A cold theory of emotion (as Ortony et. al. is often characterized) allows a computer to reason about the emotional responses of people, but the computer in no way "has" emotions: they don't compel the computer to act; they provide no color to its perception of the world. This distinction quickly becomes philosophical, but it seems that two aspects of "hot" emotions are that they influence behavior and that they are not under complete volitional control of the agent. Given this characterization, the model I propose is not inconsistent with a hot theory of emotion. Although appraisals arise from a deliberate process, the appraisals themselves are non-volitional. The planner cannot choose to view a threat to its goals as a beneficial occurrence, though it could indirectly influence its appraisals by what it chose to deliberate about. This depends on how appraisals are used. If they are wired into routines that determine behavior (e.g. by changing the heuristics that govern modifications to the agent's plans), then we could argue for a "hot" theory. If they used simply to augment the current state (e.g., I perceive that I am angry), we would have a "cold" emotional model.

6.4 How does all this help us plan?

My ultimate goal is to use an understanding of emotional appraisal to guide planning in useful ways. I do not, as of yet, have anything concrete to say on this point, but a few high-level approaches suggest themselves. One obvious connection is to use appraisals as search control. It is important to emphasize here a key distinction with prior models of emotion-influenced action selection. Whereas prior reactive reasoning models focused on modulating immediate action selection, the plan-based view allows appraisals to influence plan construction as well as immediate action selection. For example, a planner could be guided to focus its planning effort on goals that elicit the strongest appraisals. One can alter the balance between plan generation and plan execution by being more or less eager to execute steps in a plan before completely reasoning through their consequence.

Classical planning methods perform a lot of expensive bookkeeping in the service of generating valid plans. This can be a problem both from the standpoint of efficiency and the standpoint of believability. People are notoriously bad (at least lazy) about reasoning through the consequences of planned activities. One idea I am exploring is using a "lazy" constraint propagation method that only partially reasons through the consequences of plan commitments. The idea is to use something like spreading activation based on emotional appraisals to focus bookkeeping on "emotionally salient" aspects of the plan network. Such a method could generate plans more quickly but the plans would be more likely to contain problems or inconsistencies. The hope is that these "bugs" will be non-critical or easily repaired.

Another possible use of emotional appraisals relates to my work on social planning (Gratch, 1998; Gratch 1999). The CFOR planner implements a notion of planning stances that alter the way the planner behaves based on its social relationship to other agents. For example, Figure 3 illustrated a interaction between the plans of two agents, Bo and Jo, that wish to use the same vehicle to go to different locations. When Bo learns of Jo's plans, he has several options to resolve the threat. He could run to the car and get there before Jo. This corresponds to a "rude" stance where the planner tries to resolve threats in its own plans but disregards any threats that it introduces into the plans of others. Alternatively, it could meekly give up the use of the car, a deferential stance. Currently, agents based on the CFOR planner deliberately decide which planning stance to adopt given the current plans and world state. Obviously, one's emotional state is a believable input to such a decision.

The planning community will only take interest in theories of emotion if the information processing suggested by emotion research is functionally different from that already investigated by planning researchers. One apparent difference lies in social aspects of behavior. Emotion researchers have richer theories to describe how an agent's behavior is perceived by others. These distinctions could help make planners "socially appropriate" through the appropriate use of search control. There are other ways that this social knowledge could guide planning. For example, appraisals could influence an agent's perception of the current state or its perception of the intentions of other agents. Currently I implement a simplistic scheme for reasoning about the behavior of other agents. For Bo to reason about Jo's plans, Jo must tell Bo her plans and goals. A more realistic interaction would allow agents to infer the plans and goals of other using plan recognition techniques. As plan recognition is inherently ambiguous, we could use appraisals to resolve ambiguities. Thus, based on our feelings we can decide whether Bill Clinton's goal was to lie or mislead.

7 ACKNOWLEDGMENTS

This work benefited from numerous discussions with Paul Rosenbloom, Jeff Rickel, Stacy Marsella, and Clark Elliott. We gratefully acknowledge the support of the Army Research Institute under contract TAPC-ARI-BR.

8 **REFERENCES**

Ambros-Ingerson, J. A. and Steel, S. 1988. "Integrating Planning, Execution and Monitoring," *AAAI98*.

Bates, J., 1992. Virtual Reality, art, and entertainment. *Presence: The Journal of Teleoperators and Virtual Environmentt.* V1(1) MIT Press.

Beaudoin, L. 1995. Goal Processing in Autonomous Agents. Ph.D Thesis (CSRP-95-2), University of Birmingham.

Bratman, M. E. 1987. *Intentions, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA.

Chapman, D. 1987. Planning for conjunctive goals. Artificial Intelligence, 32, pp. 333-377.

Damasio, A. 1994. *Descartes' Error*. Avon Books, New York, NY.

Elliott C. D. The Affective Reasoner: A Process Model of Emotions in a Multi-agent System. Ph.D Thesis (TR#32), Northwestern University.

Fikes, R.E. and Nilsson, N. J. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence*, 2(3-4)

Frijda, N. 1986. The emotions. Cambridge Univ. Press.

Frijda, N. and Swagerman, J. 1987. Can computers feel? Theory and design of an emotional system. *Cognition and Emotion*, 1, 235-257.

Gratch, J., 1998. Metaplanning for multiple agents. AIPS98 Workshop on Plan Execution, Pittsburgh, PA

Gratch, J., 1999. How to Make Your Planner Rude and Other Issues in Multi-agent Planning. *Information Sci*ences Institute Research Report ISI/RR-99-464. Grosz, B., and Kraus, S. 1996. Collaborative Plans for Complex Group Action. *Artificial Intelligence*, 86(2).

Hill, R., Chen, J., Gratch, G., Rosenbloom, P., and Tambe, M. 1997. "Intelligent Agents for the Synthetic Battlefield," in AAAI-97/IAAI-97, pp. 1006-1012.

McAllester, D. and Rosenblitt, D. 1991. Systematic Nonlinear Planning. AAA191.

National Research Council (NRC), 1998. *Modeling Human and Organizational Behavior*. National Academy Press, Washington D.C.

Newell, A. 1990. Unified Theories of Cognition. Harvard Press.

Oatley, K. and Johnson-Laird, P.N. 1987. Towards a Cognitive Theory of Emotions. *Cognition and Emotion*, 1 (1) pp. 29-50

Oatley, K. 1992. Best Laid Schemes: the psychology of emotions. Cambridge University Press.

Ortony A., Clore, G. L., and Collins, A. *The Cognitive Structure of Emotions*. Cambridge University Press.

Reilly, W.S.N., 1996. *Believable Social and Emotional Agents*. Ph.D Thesis CMU-CS-96-138. Carnegie Melon University.

Simon, H. A. 1967. Motivational and Emotional Controls of Cognition. *Psychological Review*, 74, pp. 29-39.

Sloman, A. 1987. Motives, mechanisms and emotions. *Cognition and Emotion*, 1, pp 217-234.