The role of emotions in multiagent teamwork

Ranjit Nair
Computer Science Department
University of Southern California
941 W 37th Pl
Los Angeles CA 90089
nair@usc.edu

Milind Tambe
Computer Science Department and Information Sciences Institute
University of Southern California
941 W 37th Pl
Los Angeles CA 90089
tambe@usc.edu

Stacy Marsella
Information Sciences Institute
University of Southern California
4676 Admiralty Way, #1001
Marina del Rey CA 90292
marsella@isi.edu

Running Title: The role of emotions in multiagent teamwork

Contact: Ranjit Nair
Computer Science Department
University of Southern California
941 W. 37th Place, #300
Los Angeles, CA 90089-0781
Phone:(213) 740-6747 Fax: (213) 740-7285
Email: nair@usc.edu

Abstract

Emotions play a significant role in human teamwork. However, despite the significant progress in multiagent teamwork, as well as progress in computational models of emotions, there have been very few investigations of the role of emotions in multiagent teamwork. This chapter attempts a first step towards addressing this shortcoming. It provides a short survey of the state of the art in multiagent teamwork and in computational models of emotions. It considers three cases of teamwork, in particular, teams of simulated humans, agent-human teams and pure agent teams, and examine the effects of introducing emotions in each. Finally, it also provides preliminary experimental results illustrating the impact of emotions on multiagent teamwork.

1 Introduction

The advantages of teamwork among humans have been widely endorsed by experts in sports (Jennings, 1990) and business organizations (Katzenbach and Smith, 1994). The following quote from Andrew Carnegie, one of America's most successful businessmen, highlights the crucial role of teamwork in any organization.

> "Teamwork is the ability to work together toward a common vision. The ability to direct individual accomplishments toward organizational objectives. It is the fuel that allows common people to attain uncommon results."– Andrew Carnegie

When team members align their personal goals with the goals of the team, they can achieve more than any of them individually.

Moving away from human organizations to organizations of artificial intelligence entities called agents, we find that the same advantages of teamwork are greatly felt. An agent is defined as "a computer system that is *situated* in some *environment*, and is capable of *autonomous action* in this environment in order to meet its design objectives." (Wooldridge, 2000). This computer system could either be a software agent that exists in a virtual environment or a hardware entity like a robot that operates in a real environment and the design objectives of the system can be thought of as the goals of the agent. A system comprising of multiple agents working collaboratively or competitively in an environment is referred to as a *multiagent system*, a subfield of *distributed artificial intelligence*. In this chapter, we will focus only collaborative multiagent systems where agents can benefit greatly by working as a team.

In today's multiagent applications, like simulated or robotic soccer (Kitano et al., 1997), urban search and rescue simulations (Kitano et al., 1999), battlefield simulations (Tambe, 1997) or personal assistant agents (Scerri et al., 2002), agents have to work together in order to jointly complete some task. For instance, ambulance and fire-brigade agents need to work together to jointly save as many civilians in an urban search and rescue simulation (Kitano et al., 1999) or personal assistant agents representing different humans need to work together to schedule a meeting between these humans (Scerri et al., 2002). This involves choosing personal goals that are aligned with the overall team goal. To that end, several teamwork theories and models (Cohen and Levesque, 1991; Grosz and Kraus, 1996; Tambe, 1997; Jennings, 1995) have been proposed that help in the coordination of teams, deciding when and what they should communicate (Pynadath and Tambe, 2002), how they should form and re-form these teams (Hunsberger and Grosz, 2000; Nair et al., 2003), etc. Through the use of these models of teamwork, large-scale multiagent teams have been deployed successfully in a variety of complex domains (Kitano et al., 1999; Kitano et al., 1997; Tambe, 1997; Scerri et al., 2002).

Despite the practical success of multiagent teamwork, the role of emotions in such teamwork remains to be investigated. In particular, in practical human teams, a lot of emphasis is placed on the emotional state of the members of the team and on methods of making sure that members of the team understand each others' emotions and help keep each other motivated about the team's goal (Katzenbach and Smith, 1994; Jennings, 1990). Work in human and animal behavior (Lazarus, 1991; Darwin, 1998; Oatley, 1992; Goleman, 1995) suggests several roles for emotions and emotional expression in teamwork. First, emotions act like a value system, allowing each individual to perceive its situation and then arrive at a decision rapidly. This can be very beneficial in situations where the individual needs to think and act quickly. Second, the emotional expressions of an individual can act as a cue to others communicating to them something about the situation that they are in and also about the likely behavior of the individual displaying emotion. For instance, if we detect fear in someone's behavior, we are alerted that something dangerous might be present. Also, a person displaying emotion like fear may behave in an irrational way. Being receptive to the emotional cues of this fearful team member, allows us to collaborate with that person or compensate for that person's behavior.

In spite of these advantages to human teams, the role of emotions has not been studied adequately in multiagent research and in particular in multiagent teams. In this chapter, we will conjecture about how multiagent teams can stand to gain through the introduction of emotions. Section 2 describes the state of the art in multiagent teamwork and in agent emotions. Section 3 describes how multiagent teamwork and emotions can be intermixed and the benefits of such a synthesis of these two threads of research. In particular, we will consider three types of teams: teams of simulated humans, mixed agent-human teams and pure agent teams. In Section 4, we will demonstrate empirically, the effect of introducing emotions in a team of helicopters involved in a mission rehearsal.

2 State of the art in multiagent teamwork and agent emotions: A quick survey

This section provides an overview of the state of the art in multiagent teamwork (Section 2.1) and of current research in agent emotions (Section 2.2). Section 3 then attempts to discuss how these two threads of research could interact.

2.1 State of the art in Multiagent Teamwork

There is an emerging consensus among researchers in multiagent systems that teamwork can enable flexible coordination among multiple heterogeneous entities and allow them to achieve their shared goals (Cohen and Levesque, 1991; Tambe, 1997; Grosz and Kraus, 1996; Jennings, 1995). Furthermore, previous work (Cohen and Levesque, 1991; Tambe, 1997; Grosz and Kraus, 1996; Jennings, 1995) has also illustrated that effective teamwork can be achieved through team coordination algorithms(sometimes called teamwork models) that are independent of the domain in which the agents are situated. Given that each agent is empowered with teamwork capabilities via teamwork models, it is feasible to write a high-level Team-Oriented Program (TOP) (Tambe, 1997; Tidhar, 1993), and the teamwork models then automatically generate the required coordination. In particular, team-oriented programs omit details of coordination, thus enabling developers to provide high-level specifications to the team to perform team actions rather than invest effort in writing code for low-level detailed coordination. The teamwork models that govern coordination are based on a belief-desire-intention (BDI) architecture, where beliefs are information about the world that an agent believes to be true, desires are world states that the agent would like to see happen and intentions are effects that the agent has committed to achieve. The beliefs of an agent need not be true, for instance, an agent may receive incorrect observations owing to faulty sensors and the different agents in the team could have different beliefs owing to differences in how and what they observe. Also, different agents in the team could have desires that are in conflict with each other. The goal of the team coordination algorithms is to achieve *mutual belief* among the team members and to form joint intentions to allow the agents to work towards the same goal.

We illustrate the use of *Team-oriented Programming* through several example domains: Mission rehearsal (Tambe, 1997), RoboCupRescue (Nair et al., 2003) and Electric Elves (Scerri et al., 2002). A description of these domains is also helpful for our discussions of emotions in Section 3. For expository purposes, we have intentionally simplified the mission rehearsal domain: A helicopter team is executing a mission of transporting valuable cargo from point X to point Y through enemy terrain (see Figure 1). There are three paths from X to Y of different lengths and different risk due to enemy fire. One or more scouting sub-teams must be sent out on different routes (some routes may not have a scouting team), and the larger the size of a scouting sub-team the safer it is. When scouts clear up any one path from X to Y, the transports can then move more safely along that path. However, the scouts may fail along a path, and may need to be replaced by a transport at the cost of not transporting cargo. Of course, we wish for the most amount of cargo to be transported in the quickest possible manner within the mission deadline.

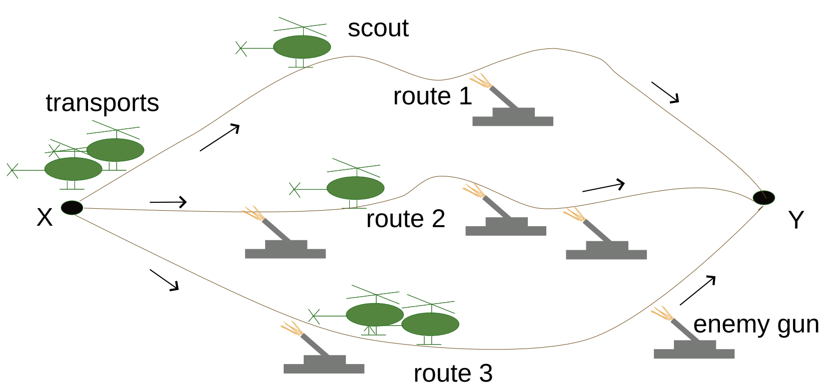

Figure 1: Helicopter Mission Rehearsal Domain. Helicopters take on the role of either scouting a route or transporting cargo along a scouted route. Helicopters may be shot down by enemy guns on unscouted routes. The goal is to determine what roles each of the helicopters should take so as to get as much cargo as possible from point X to Y within the mission deadline.

The TOPs for domains such as these consist of three key aspects of a team: (i) a team organization hierarchy consisting of roles; (ii) a team (reactive) plan hierarchy; and (iii) an assignment of roles to execute plans. Thus, the developer need not specify low-level coordination details. Instead the TOP interpreter (the underlying coordination infrastructure) automatically enables agents to decide when and with whom to communicate and how to reallocate roles upon failure. In the TOP for this example, we first specify the team organization hierarchy (see Figure 2(a)). *Task Force* is the highest level team in this organization and consists of two roles *Scouting* and *Transport*, where the *Scouting* sub-team has roles for each of the three scouting sub-sub-teams. Next we specify a hierarchy of reactive team plans (see Figure 2(b)). Reactive team plans explicitly express joint activities of the relevant team and consist of: (i) initiation conditions under which the plan is to be proposed; (ii) termination conditions under which the plan is to be ended; and (iii) team-level actions to be executed as part of the plan. In Figure 2(b), the highest level plan **Execute Mission** has three sub-plans: **DoScouting** to make one path from X to Y safe for the transports, **DoTransport** to move the transports along a scouted path, and **RemainingScouts** for the scouts which haven't reached the destination yet to get there.

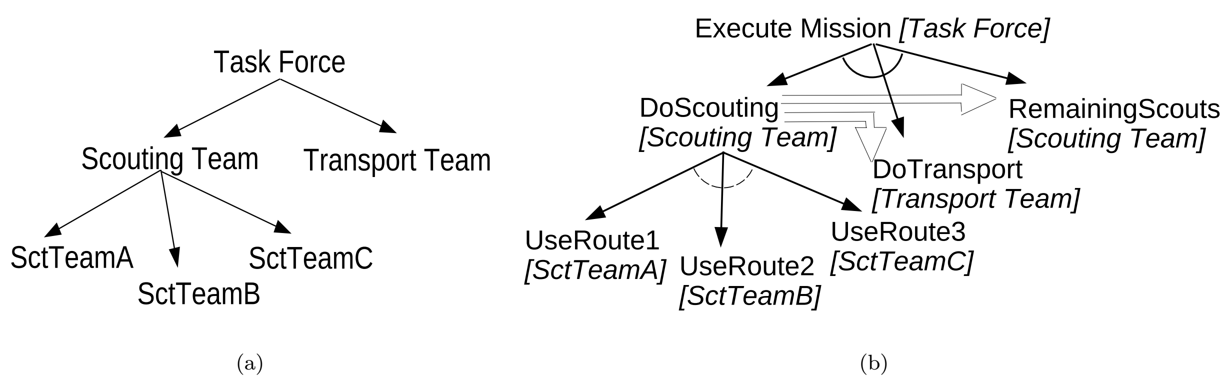

Figure 2: Team Oriented Program for the helicopter domain. a: Organization hierarchy; b: Plan hierarchy

Figure 2(b) also shows coordination relationships: An AND relationship (depicted with a solid arc) indicates sub-plans that need to be all completed successfully for the parent plan to succeed, while an OR relationship (depicted with a dotted arc) indicates that success of any one of the sub-plans will result in the parent sub-plan succeeding. Thus, **DoScouting**, **DoTransport** and **RemainingScouts** must all three be successful while at least one of **UseRoute1**, **UseRoute2** or **UseRoute3** need be performed successfully. There is also a temporal dependence relationship among the sub-plans (depicted with a double arrow), which implies that succeeding sub-plans cannot begin until the preceding sub-plan has been successfully completed. Thus, the sub-teams assigned to perform **DoTransport** or **RemainingScouts** cannot do so until the **DoScouting** plan has succeeded. However, **DoTransport** and **RemainingScouts** execute in parallel. Finally, we assign roles to plans — Figure 2(b) shows the assignment in brackets adjacent to the plans. For instance, *Task Force* team is assigned to jointly perform **Execute Mission**.

New techniques for combining BDI approaches like team-oriented programming with decision-theoretic approaches based on partially observable Markov decision processes (POMDPs) have recently emerged (Schut et al., 2001; Nair et al., 2003). The advantage of the BDI approach is that it allows the specification of large plans for complex domains. Unfortunately, such complex domains generally contain uncertainty. An Markov Decision Process (MDP) (Howard, 1960) is a formal representation of a domain with a single agent where there is uncertainty because the agent's actions have probabilistic outcomes. However, MDPs make an unrealistic assumption that the agent can sense the world state precisely. A partially observable Markov decision processes (POMDP) is a generalization of an MDP, where the single agent may not observe the entire world state, but only some observations drawn from some probability distribution. However, both MDPs and POMDPs are for single agents. Distributed POMDP models (Bernstein et al., 2000; Boutilier, 1996; Pynadath and Tambe, 2002; Nair et al., 2003) are generalization of POMDPs to the case where there a multiple agents, each with a possibly different partial view of the world state. Both POMDPs and decentralized POMDPs are computationally expensive to use to find the optimal plan for very large domains. However, they are very useful for analyzing exisiting teams plans and coordination algorithms. For instance, Schut *et al.* compare various strategies for intention reconsideration (deciding when to deliberate about its intentions) by modeling a BDI system using a POMDP, but their work is confined to a single agent.

Nair *et al.* (Nair et al., 2003) use Role-based Markov Team decision problem (RMTDP), a distributed POMDP model, for analysis of TOPs, where the results of RMTDP analysis are fed back into the BDI-based TOPs (see Figure 3). The RMTDP for a team of n agents is defined as a tuple $\langle S, A, P, \Omega, O, R, \mathcal{RL} \rangle$. It consists of a finite set of states S. $P(s, < a_1, \ldots, a_n >, s')$ gives the probability of transitioning from state s to state s' given that the agents perform the actions $< a_1, \ldots, a_n > \in A$ jointly. Each agent i receives an

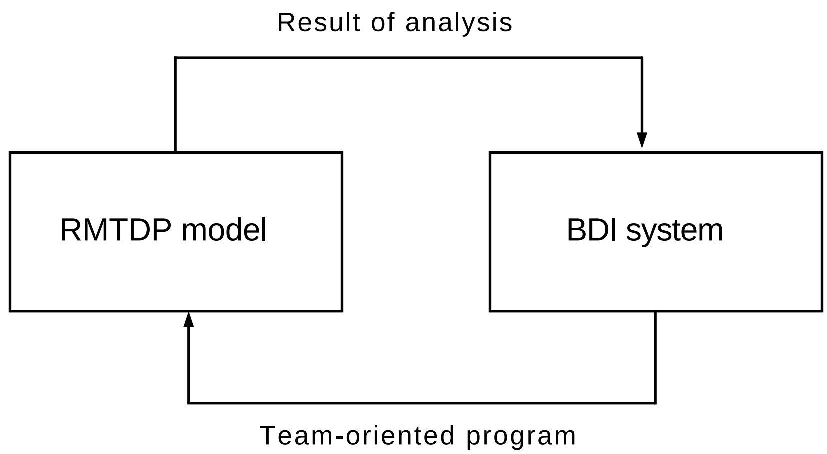

Figure 3: Integration of Belief-Desire-Intention(BDI)-based Team-Orient Programming approach with Role-based Multiagent Team Decision Problem (RMTDP), a distributed POMDP model. Analysis of the team-oriented program using the RMTDP model is fed back to improve the team-oriented program

observation $\omega_i \in \Omega_i$ based on the function $O(s, < a_1, \ldots, a_n >, \omega_1, \ldots, \omega_n)$, which gives the probability that the agents receive the observations, $\omega_1, \ldots, \omega_n$ given that the world state is s and they perform $< a_1, \ldots, a_n >$ jointly. $\mathcal{RL} = \{r_1, \ldots, r_s\}$ is a set of all roles that the agents can undertake . Each instance of role r_j may be assigned some agent i to fulfill it. Each agent's actions are distinguishable into role-taking and role-execution actions. The agents receive a single joint reward $R(s, a_1, \ldots, a_n)$. This model is used for evaluating various role allocations and reallocation strategies for TOPs. In Section 4, we show how the emotional state of the agents in the team can be modeled using RMTDP and empirically show how the emotional state of the agents can affect how roles should be allocated.

Another example domain, where TOPs have been applied is RoboCupRescue (Nair et al., 2003; Kitano et al., 1999). This is a large-scale simulation of a city that has just undergone an earthquake. Here, teams of ambulance and fire-brigade agents have to be formed and dispatched to various buildings to put out fires and rescue civilians trapped within them. TOPs allow us to flexibly coordinate the activities of these agents, monitoring the situations and reforming the teams if necessary.

The third example domain we discuss is the Electric Elves (E-Elves) project at USC/ISI, which deploys an agent organization in support of the daily activities of a human organization (Scerri et al., 2002; Chalupsky et al., 2001). We believe this application to be fairly typical of future generation applications involving teams of agents and humans. The operation of a human organization requires the performance of many everyday tasks to ensure coherence in organizational activities, e.g., monitoring the status of activities, gathering information and keeping everyone informed of changes in activities. Teams of software agents (proxy agents) can aid organizations in accomplishing these tasks, facilitating coherent functioning and rapid, response to crises. The notion of an agent proxy is similar to the idea of "robot as avatar" in Chapter ? (Brezeal and Brooks). While the goal of both is to reduce the coginitive load on the humans, the key difference is that agent proxies exist only in software and need to interact with several other agent proxies in addition to the human it represents. Each agent proxy is called Friday (after Robinson Crusoe's man-servant Friday) and acts on behalf of its user in the agent team. Currently, Friday can perform several tasks for its user. If a user is delayed to a meeting, Friday can reschedule the meeting, informing other Fridays, who in turn inform their users. If there is a research presentation slot open, Friday may respond to the invitation to present on behalf of its user. Friday can also order its user's meals and track the user's location, posting it on a Web page. Friday communicates with users using wireless devices, such as personal digital assistants. Each Friday's team behavior is based on a teamwork model, called *S*hell for *TEAM*work (STEAM) (Tambe, 1997). STEAM encodes and enforces the constraints between roles that are required for the success of the joint activity, e.g., meeting attendees should arrive at a meeting simultaneously. When a role within the team needs to be filled, STEAM requires that a team member is assigned responsibility for that role. To find the best suited person, the team auctions off the role, allowing it to consider a combination of factors and assign the best suited user.

2.2 Computational Models for Emotion

The interest in general computational models of emotion and emotional behavior has been steadily growing in the agent and artificial intelligence research communities. Although the creation of general computational models is of potential interest in understanding human behavior, much of the interest in the agent community has been fueled by the application areas for such models. For example, there has been a growing body of work in the design of virtual humans, software artifacts that act like people but exist in virtual worlds, interacting with immersed humans and other virtual humans. Virtual human technology is being applied to training applications (Rickel et al., 2002), health interventions (Marsella et al., 2000), marketing (André et al., 2000) and entertainment (Cavazza et al., 2002). Emotion models are a critical component of this technology, providing virtual humans that are better facsimiles of humans as well as providing a more engaging experience. Emotion models have also been proposed as a critical component of more effective human computer interaction that factors in the emotional state of the user (Lisetti and Schiano, 2000; Picard, 1997).

Much of the work on computational models of emotion has been strongly influenced by Cognitive Appraisal theories of emotion (Ortony et al., 1988; Frijda, 1987; Lazarus, 1991), although some computational research (Velásquez, 1998) has also been influenced by theories that posit non-cognitive sources of emotion (Izard, 1993). Appraisal theories argue that emotion stems from a person's assessment of their relationship to the environment in terms of a set of appraisal variables or dimensions covering such factors as whether an event facilitates or inhibits the individual's goals, how desirable the impacted goals are, who deserves blame or credit, etc. Among the various cognitive appraisal theories, there is broad agreement on the set of appraisal variables and these have provided a detailed framework for building computational models of the causes of emotion. Of course, emotions also impact behavior in myriad ways. In particular, appraisal has been related to action readiness (or tendencies) and facial expressions (Smith and Scott, 1997).

The work of Richard Lazarus also tightly integrates appraisal with human coping behavior (Lazarus, 1991), the process of managing one's emotions, either by acting externally on the world (problem-focused coping), or by acting internally to change one's beliefs or attention (emotion-focused coping). Specifically, coping manages emotions by attempting to alter the person's appraisal through a combination of altering the factors in the environment that are leading the emotion, altering beliefs about those factors or altering their attention to the factors. For example, a person can focus effort and attention on countering a threat to a goal, decide that the goal is not so important or try to avoid thinking about the threat. Which of these responses will be an effective coping response depends on the seriousness/likelihood of the threat, the relevance of the goal and a person's potential to deal with the threat, among other factors. A person may or may not make an effective response and thus emotional stress may lead to adaptive or maladaptive coping behavior. The interaction of appraisal and coping unfolds over time, modeling the temporal character of emotion noted by several emotion researchers (Lazarus, 1991; Scherer, 1984): an agent may 'feel' distress for an event (appraisal), which motivates the shifting of blame (coping) to another person, which leads to anger at the now blameworthy other person (re-appraisal).

One of the appeals of cognitive appraisal models as the basis of a computational model is the ease with which appraisal can be tied to the belief, desire and intention (BDI) framework often used in agent systems. And in fact, computer science researchers have realized a range of appraisal-based approaches to modeling how emotions arise. For example, Elliott's Affective Reasoner (Elliott, 1992) is a computational realization of Ortony, Clore and Collins' appraisal model (Ortony et al., 1988). Elliott's model characterizes events in terms of specific appraisal variables and has the capability to appraise the same event from multiple perspectives (from the agent's own perspective, and the supposed perspective of other agents). Clearly, a useful capability from a social interaction and teamwork perspective. However, the model required domain specific rules to appraise events.

More recent approaches have increasingly moved towards incorporating emotions into general artificial intelligence architectures that fit well with the planning frameworks typically used in multi-agent teams. For example, Neil Reilly's work (Reilly, 1996) uses a reactive planning framework that associates desirability with probability of goal attainment but uses domain-specific rules to derive the likelihood of threats or goal attainment. El Nasr et al. (El Nasr et al., 2000) uses Markov decision processes (MDP) to provide a very general framework for characterizing the desirability of actions and events. A key advance of this method is that can represent indirect consequences of actions by examining their impact on future reward (as encoded in the MDP). The Will architecture (Moffat and Frijda, 1995) ties appraisal variables to an explicit model of plans (which capture the causal relationships between actions and effects).

One aspect of emotional behavior that is missing from most computational models is a detailed account of the myriad ways that humans cope with emotion. Rather the emphasis of the models has been on simple action selection and facial expression. The work of Marsella and Gratch (Marsella and Gratch, 2002) addresses this limitation by providing a domain independent model of coping that attempts to capture the full range of human coping behavior including not only action selection but also more sophisticated problem-focused and emotion-focused strategies. Among these strategies are planful problem solving, positive reintrepretation (finding positive meaning in an otherwise negative event such as a loved one's illness), acceptance (that a future negative event is inevitable), shifting blame and denial/wishful thinking (Marsella and Gratch, 2003). In their model, coping is essentially the inverse of appraisal, changing one or more of the appraisal factors that contributed to the emotion. Both appraisal and coping are integrated within a general BDI planning framework that employs a uniform causal interpretation of the world from the agent's perspective. The causal interpretation incorporates both past events (an episodic memory) and the agent's plans concerning future events. Thus, appraisal of past events can lead to coping responses that operate on the beliefs about that event; an agent may for example begin to believe a bad accident was someone else's fault. Similarly, intentions about future events can impact present emotions; forming the intent to redress a wrong will make an agent feel better even before the intent is acted upon. Coping is modeled as a set of basic operations that manipulate appraisal factors and can be combined to create a range of coping strategies. For example, an agent may plan a response to a threat while engaging in wishful thinking about the likelihood of success or potential negative consequences. This mirrors how coping processes are understood to operate in human behavior whereby, for example, people may employ a mix of problem-focused coping and emotion-focused coping to deal with stress. Further, this work tightly integrates appraisal and coping in an effort to model the unfolding temporal dynamics of appraisal and coping.

3 How could emotions affect multiagent teamwork

In this section, we will discuss the implication of introducing emotions into different kinds of multiagent teams. In particular, we consider three types of teams, viz. teams of simulated humans, mixed agent-human teams and purely agent teams. The need for including emotions in multiagent teams is different depending on the nature of the team. For instance, in the case of teams of simulated humans, emotions need to be modeled for an accurate simulation of a human teams. The implications of introducing emotions will vary depending on the constituents of the team. In the case of mixed-human agent teams, the introduction of emotions may actually improve the team performance. In this section, we discuss the role that emotions will play in each type of team and the issues involved in adding emotions to the team. We conclude that at least in the first two cases, i.e., in teams of simulated humans and in mixed agent-human teams, computational models of emotions based on appraisal theories (see Section 2.2) can play a critical role.

3.1 Teams of simulated humans

The central role of emotion in human teamwork becomes apparent when one works through the wide ranging impacts it has on decision-making, goal prioritization, perception, belief changes, action selection, etc. Hence in teams where each agent is a facsimile of a human being, one would have to introduce emotions in order to represent human behavior faithfully. For example, domains such as the mission rehearsal domain (see Section 2.1) are focused on simulation-based training, to provide the right set of training environment for human participants, by requiring each agent in the mission rehearsal simulation to simulate human behavior. In order to analyze or predict the behavior of humans in adverse scenarios it is important to study the influence of emotions like fear that such scenarios bring about in the humans. For example, in a demonstration of the helicopter agents that didn't model emotions, it was found that even after all its team-mates were shot down the sole remaining helicopter continued executing its mission completely unaffectedly much to the consternation of the military experts. In particular, human team's fear for their own self-survival might motivate them to abandon the teams goals in the face of the high number of fatalities. Further an individual fear would tend to spread across members of the team, influencing the decision-making of the surviving members of the team.

Introducing emotions like fear could result in the team's performance worsening but as this example clearly highlights, in order for an accurate portrayal of human organizational behavior, it is important to include such emotions. In particular, within such teams and organizations, emotions play a big role in choosing between a human's private and team goals. In the helicopter scenario, each helicopter should have a private goal of self-preservation and a public team goal to accomplish the mission. As the scenario changes, the emotional state of the agent should change as well. Appraisal-based models can express the impact of such survival fear by providing survival as one of the agent's goals. Threats to that survival would lead to fear, with fear increasing with the expectation that the threat was more certain. At some juncture, this fear for individual survival would override the agent's desire to achieve the team's mission goals. Further, the contagion-like process (Hatfield et al., 1994) of one agent's fear affecting another could be modeled as a process whereby one agent's fear would affect another agent's appraisal of the seriousness of the threat.

Thus, going back to our helicopter example, depending on the current emotional state, the helicopter agent would choose between saving itself by retreating or continuing with the mission. Even after the agent chooses which goal to perform, emotions play a role in action selection. Thus, emotions could act as a value system that the agent uses in assessing a situation.

At the inter-agent-level, emotions act could act as cues that can allow individuals to synchronize their goals/plans, to synchronize perceptual biases, to compensate for other's emotional state, or to establish a shared mental model. Again, we can look at natural systems for inspiration. For example, the expression of emotion on a mother's face has been shown to bias how a baby acts in ambiguous situations (Campos and Sternberg, 1981). This communicative role for emotions in social situations is also well recognized in psychological theories of emotion (Oatley, 1992). In the helicopter domain, the pilots could perceive fear in the voices of the other pilots and hence conclude that the situation is dangerous even though the danger may not be visible to them. In addition, humans can, in essence, appraise events from others' perspectives and thus know how they will feel. In the absence of more detailed information that can form the basis of more accurate threat assessment, these emotional signals can be very useful in threat detection.

Based on these discussions, we conclude that in teams where agents simulate humans, we may need to introduce emotions because humans use emotions for goal prioritization and action selection. We also need to build into agents the ability to perceive other agents' emotions and use these emotional cues to conclude something about the state and decision-making of the other agents. In Section 4, we demonstrate empirically, how emotions affect the decision-making of a team of helicopter agents.

3.2 Mixed Agent-human teams: Virtual Organizations

In mixed agent-human teams, it is no longer important for agents to simulate humans and human emotions. However, it would be very beneficial if the agents model the emotions of the human team-mates in order to get a better understanding of his/her needs and expected behavior (Lisetti and Schiano, 2000; Picard, 1997). For example, in Electric Elves (see Section 2.1), it would be useful if the agents could perceive the emotional state or mood of humans in order to know how to behave with them. For instance, if the "elf" could sense that the human was upset it could decide not to disturb him/her with trivial questions. If the "elf", knew something about the emotional state of the human, it could anticipate the human's actions and be prepared ahead of time with information that the human would find useful.

Although agents need not have emotions themselves, they could display emotions, in order to achieve a better rapport with the humans, achieve the cooperation of the humans and make the human feel less like he/she was interacting with an inanimate entity. Thus, in agent-human teams it maybe useful for the agents to not only model the humans' emotions but also display emotion itself.

3.3 Pure Agent Teams

The argument for including emotions in purely agent or robotic selfless teams is more challenging to make. If we just focus on the role of emotion as a signal to other members of a team, a key question is whether emotion, as a signal, provides some capability not subsumed by existing models of communication in agent teams. In human teamwork, emotional signals like facial expressions inform other members of the team something about the state of the world as well as the state and intentions of their teammates. This communication differs in many respects from how agent teams communicate. In particular, the content of this communication is not simply specific statements about the world, as it often is in agent teams, but rather the individual's attitudes and emotional reactions to the world. Further, emotional signals can be intentional, unintentional or a mixture of the two as is the case when people try to suppress their emotional expressions (Ekman, 2001). In contrast, pure agent teams communicate via explicit, intended communication acts or by the intended actions they take in the world. Further emotional signals are communicated across a variety of channels, verbally and nonverbally. These channels vary in their capacity, the specificity of the information effectively communicated and the cognitive overhead in using them. A person can smile at a cute baby without much thought but may need more resources to verbally express happiness. Agent teams typically have two channels, their communication channel and their actions. These differences suggest potential benefits for using emotions in purely agent teams. For instance, there might be an advantage to having agent teams communicate attitudinal or emotional information as well as an advantage to exposing this information to teammates automatically, across low cost channels. Consider that the agents are built such that apart from being able to communicate and act deliberately after a accurate and possibly computationally intensive assessment of the state, they can also can emit some low cost "emotional" signal based on an approximate state assessment. For example, a robot could have hardwired circuitry that triggers l.e.d.s (light emitting diodes), that represent emotional cues like fear to indicate a state where the robot is in danger, worry to indicate low likelihood of success, helplessness to indicate that it needs to help. These emotional cues can be computed and transmitted quickly and could result in the team being able to coordinate itself without having to wait for the accurate state estimation to be performed. If, for example, agents could use these emotional cues to determine action selection of the other agents in the team, it could result in greater synchronization and consequently, better teamwork.

4 Experimental illustration

In this section, as an illustration of the effect of emotions on multiagent teamwork, we will demonstrate how allocation of roles in a team get affected by emotions like fear. Our approach is to first build a Role-based Markov Team Decision Problem (RMTDP) (Nair et al., 2003) for the team of agents (See Section 2.1). Here we model the agents such that their emotional states are included in the model.

We will now concretely demonstrate how emotions can affect decision-making in a team of helicopters. To this end, recall the RMTDP analysis of team-oriented programs mentioned in Section 2.1. The emotional state of the agent could skew how the agent sees the world. This could result in the agent applying different transition, observation or reward functions. In this discussion, we will focus on how fear may affect the reward function used in RMTDP. For instance, in a fearful state, agents may consider the risk of failures to be much higher than in a non-fearful state. In the helicopter domain, such agents might penalize heavily those states where a helicopter crashes. We will now demonstrate how such a change in the emotional state of the agents would affect the best role allocation.

We consider a team of 6 helicopters and vary the number of agents in the team who were scared of losing a helicopter to enemy fire. Such an agent would place a big penalty on those states where one or more helicopter crashed. Figures 4(a) and 4(b) show the number of scouts allocated to each route (X-axis) as we vary the number of fearful agents in the team (Y-axis) from none to all 6 for 2 different penalties for helicopter crashes. In Figure 4(a), when all the agents were fearless, the number of scouts sent out was 3, all on route 2, however when fearful agents were introduced the number of scouts sent out changed to 4, also on route 2, because the team was now prepared to lose out on the chance of getting higher reward if they could ensure that each scout that was sent out would be safer. In Figure 4(b), we reduced the amount of penalty the agents ascribed to a helicopter crash. Here, when fearful agents were introduced the number of scouts remained unchanged but the scouts now used route 1, a safer albeit longer route, instead of route 2 which was more dangerous but allowed the mission to be completed quicker. Thus, with the introduction of “fear”, we found that the team’s decision making behavior changed such that they either deployed more scouts or assigned the scouts to a safer route.

Although, the emotion “fear” was modeled simply as a penalty for states where a helicopter crashes, the purpose of the experiment was simply to show that emotional response affects what the team perceives is its best allocation. In order to evaluate teams where emotions are represented more realistically , we would need the following:

- a more realistic model of how an agent’s emotional state would change based on new percepts. This

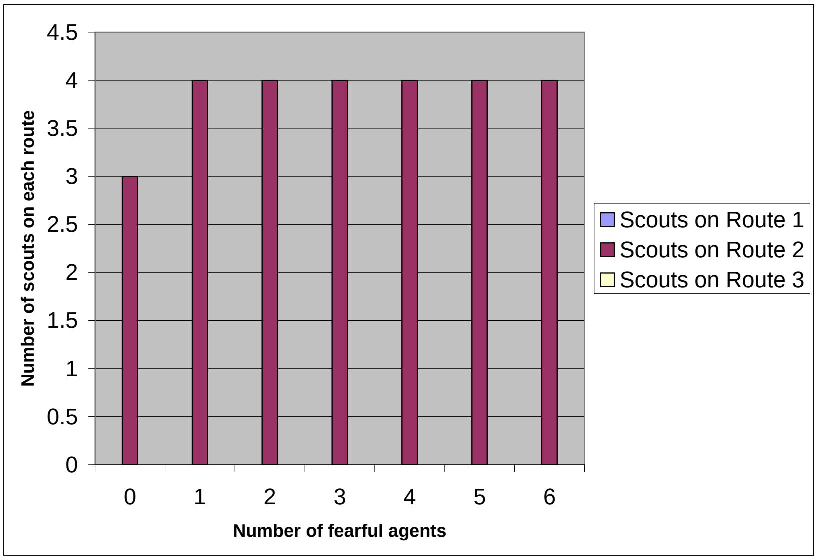

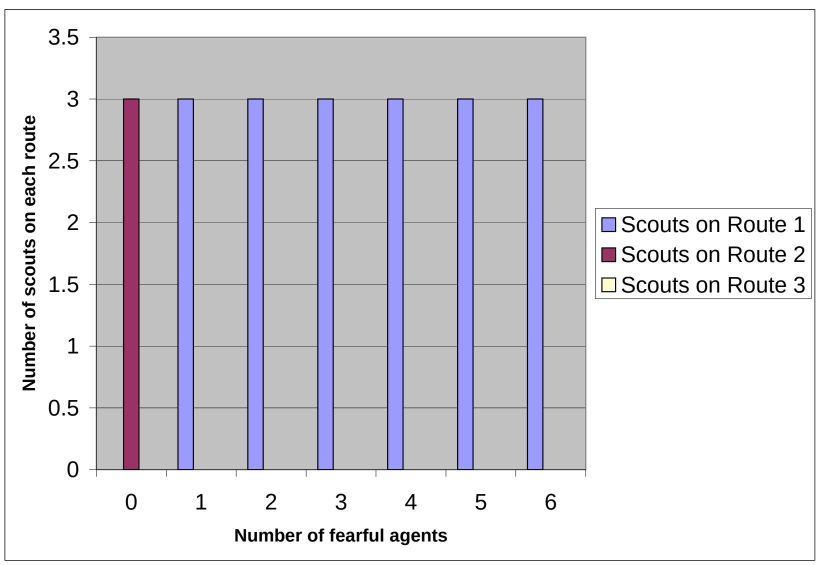

Figure 4: Role allocations in fearful teams with different reward functions. a) Graph on the top shows role allocations for reward function 1. Increasing the number of fearful agents results in more scouts being sent together to increase the safety of the scouting team. b) Graph at the bottom show role allocations for reward function 2. Increasing the number of fearful agents results in moving scouts from a shorter but more risky route to a safer but safer route

model of how the emotional state transitions can be incorporated as part of the transition function in the RMTDP model in order to evaluate the team's performance in the presence of emotion.

- a more realistic model of how humans (that the agents are simulating), would respond based on their emotional state. This would form part of the team-oriented program where the individual agent's action selection is specified.

Both the model of how emotional state changes as well as the model of human behavior in the presence of emotion should ideally be informed by human behavior in such task domains.

5 Conclusion

This chapter represents the first step in introducing emotions in multiagent teamwork. We examined the role of emotions into three different kinds of teams. Firstly, in teams of simulated humans, introducing emotions in the agents results in more believable agent behavior and consequently better simulations. Secondly, in virtual organizations, where agents could simulate emotions to be more believable and engaging to the human and can also anticipate the human's needs by modeling the human. Finally, in pure agent teams, the introduction of emotions could help bring in the same advantages that emotions bring to human teams.

Teams of simulated agents and mixed-human teams can greatly benefit with compuational models of emotion. In particular, to evaluate and improve such teams, we would need the following:

- a model of how an agent's emotional state would change based on new percepts.

- a model of how humans would respond based on their emotional state.

6 Acknowledgement

This research was supported by grant #0208580 from the National Science Foundation.

REFERENCES

André, E., Rist, T., Mulken, S. V., and Klesen, M., 2000, The automated design of believable dialogues for animated presentation teams, In Cassell, J., Sullivan, J., Prevost, S., and Churchill, E., editors, Embodied Conversational Agents, pages 220–255, Cambridge, MA: MIT Press.

Bernstein, D. S., Zilberstein, S., and Immerman, N., 2000, The complexity of decentralized control of MDPs, In UAI.

Boutilier, C., 1996, Planning, learning & coordination in multiagent decision processes, In TARK.

Campos, J. and Sternberg, C., 1981, Perception, appriasal and emotion: The onset of social referencing, In Lamb, M. and Sherrod, L., editors, Infant Social Cognition, pages 273–314, Erlbaum.

Cavazza, M., Charles, F., and Mead, S. J., 2002, Interacting with virtual characters in interactive storytelling, In Proceedings of First International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS-02), pages 318–325.

Chalupsky, H., Gil, Y., Knoblock, C., Lerman, K., Oh, J., Pynadath, D., Russ, T., and Tambe, M., 2001, Electric Elves: Applying agent technology to support human organizations, In Proceedings of International Conference on Innovative Applications of AI (IAAI-01), pages 51–58.

Cohen, P. R. and Levesque, H. J., 1991, Teamwork, Nous, 25(4):487–512.

Darwin, C., 1998, The expression of emotions in man and animals, New York: Oxford University Press, 3rd edition.

Ekman, P., 2001, Telling Lies: Clues to Deceit in the Marketplace, Politics and Marriage, Norton.

El Nasr, M. S., Yen, J., and Ioerger, T., 2000, Flame: Fuzzy logic adaptive model of emotions, Journal of Autonomous Agents and Multiagent Systems, 3(3):219–257.

Elliott, C., 1992, The affective reasoner: A process model of emotions in a multi-agent system, PhD dissertation 32, Northwestern, IL, Northwestern University Institute for the Learning Sciences.

Frijda, N., 1987, Emotion, cognitive structure, and action tendency, Cognition and Emotion, 1:115–143.

Goleman, D., 1995, Emotional intelligence, New York: Oxford University Press.

Grosz, B. and Kraus, S., 1996, Collaborative plans for complex group action, Artificial Intelligence, 86(2):269–357.

Hatfield, E., Cacioppo, J., and Rapson, R., 1994, Emotional Contagion, Cambridge University Press.

Howard, R. A., 1960, Dynamic Programming and Markov Processes, Cambridge, MA: The MIT Press.

Hunsberger, L. and Grosz, B., 2000, A combinatorial auction for collaborative planning, In Proceedings of Fourth International Conference on Multiagent Systems (ICMAS-2000), pages 151–158.

Izard, C., 1993, Four systems for emotion activation: Cognitive and noncognitive processes, Psychological Review, 100(1):68–90.

Jennings, J., 1990, Teamwork: United in victory, Silver Burdett Press, Englewood Cliffs, NJ.

Jennings, N., 1995, Controlling cooperative problem solving in industrial multi-agent systems using joint intentions, Artificial Intelligence, 75:195–240.

Katzenbach, J. and Smith, D. K., 1994, The Wisdom of Teams, Harper Business.

Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., and Osawa, E., 1997, RoboCup: The robot world cup initiative, In Proceedings of First International Conference on Autonomous Agents (Agents'97), pages 340–347.

Kitano, H., Tadokoro, S., and Noda, I., 1999, RoboCup-Rescue: Search and rescue for large scale disasters as a domain for multiagent research, In Proceedings of IEEE Conference on Systems, Men, and Cybernetics (SMC-99).

Lazarus, R., 1991, Emotion and Adaptation, NY: Oxford University Press.

Lisetti, C. L. and Schiano, D., 2000, Facial expression recognition: Where human-computer interaction, artificial intelligence, and cognitive science intersect, Pragmatics and Cognition, 8(1):185–235.

Marsella, S. and Gratch, J., 2002, A step toward irrationality: Usign emotion to change belief, In Proceedings of First International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS-02).

Marsella, S. and Gratch, J., 2003, Modeling coping behavior in virtual humans: Don't worry, be happy, In Proceedings of Second International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS-03).

Marsella, S., Johnson, W. L., and LaBore, C., 2000, Interactive pedagogical drama, In Proceedings of Fourth International Conference on Autonomous Agents (ICMAS-2000).

Moffat, D. and Frijda, N., 1995, Where there's a will there's an agent, In Proceedings of Workshop on Agent Theories, Architectures and Languages (ATAL-95).

Nair, R., Tambe, M., and Marsella, S., 2003, Role allocation and re-allocation in multiagent teams: Towards a practical analysis, In Proceedings of Second International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS-03).

Oatley, K., 1992, Best laid schemes: The psychology of emotions, Cambridge University Press.

Ortony, A., Clore, G., and Collins, A., 1988, The Cognitive Structure of Emotions, Cambridge University Press.

Picard, R. W., 1997, Affective Computing, Cambridge, MA: MIT Press.

Pynadath, D. V. and Tambe, M., 2002, The communicative multiagent team decision problem: Analyzing teamwork theories and models, Journal of Artificial Intelligence Research, 16:389–423.

Reilly, N., 1996, Believable Social and Emotional Agents, PhD dissertation CMU-CS-96-138, Carnegie Mellon University, Pittsburgh, PA.

Rickel, J., Marsella, S., Gratch, J., Hill, R., Traum, D., and Swartout, W., 2002, Toward a new generation of virtual humans for interactive experiences, IEEE Intelligent Systems, July-August:32–38.

Scerri, P., Pynadath, D., V., and Tambe, M., 2002, Towards adjustable autonomy for the real-world, Journal of Artificial Intelligence Research, 17:171–228.

Scherer, K., 1984, On the nature and function of emotion: A component process approach, In Scherer, K. R. and Ekman, P., editors, Approaches to emotion, pages 293–317, Erlbaum, Lawrence Associates, Incorporated.

Schut, M. C., Wooldridge, M., and Parsons, S., 2001, Reasoning about intentions in uncertain domains, In Proceedings of European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty.

Smith, C. A. and Scott, H. S., 1997, A componential approach to the meaning of facial expressions, In Russell, J. A. and Fernández-Dols, J. M., editors, The Psychology of Facial Expression, pages 229–254, Paris: Cambridge University Press.

Tambe, M., 1997, Towards flexible teamwork, Journal of Artificial Intelligence Research, 7:83–124.

Tidhar, G., 1993, Team-oriented programming: Social structures, Technical Report 47, Australian A.I. Institute.

Velásquez, J., 1998, When robots weep: emotional memories and decision-making, In Proceedings of Fifteenth National Conference on Artificial Intelligence (AAAI-98).

Wooldridge, M., 2000, Intelligent agents, In Weiss, G., editor, Multiagent Systems: A Modern Approach to Distributed AI, pages 27–78, Cambridge, MA: MIT Press.