Literature Review of Teamwork Models

Katia Sycara   Gita Sukthankar

CMU-RI-TR-06-50

November 2006

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Abstract

Both human collaboration and software agent collaboration have been thoroughly studied, but there is relatively little research on hybrid human-agent teamwork. Some research has identified the roles that agents could play in hybrid teams: supporting individual team members, being a teammate, or supporting the team as a whole [99]. Some other work [57] has investigated trust concepts as the fundamental building block for effective human-agent teamwork or posited the types of shared knowledge that promote mutual understanding between cooperating humans and agents [9, 68]. However, many of the facets of human agent teamwork models, such as communication protocols for forming mutual intelligibility, performing team monitoring to assess progress, forming joint goals, addressing task interdependencies in hybrid teamwork are still unexplored. In this report, we address the following questions:

1. what factors affect human team task performance and cognition?
2. how can agent coordination mechanisms be adapted for human-agent teams?
3. with current technologies, what roles can agents successfully fill in hybrid human-agent teams?
4. what are the barriers to human-agent interaction?

Contents

1 Introduction

In this report, we give an overview of the literature on teamwork: human-only, agent-only, and human-agent teamwork models. Cohen et al. [23] defined agent teamwork as *"a set of agents having a shared objective and a shared mental state"*, whereas Salas et al. [85] characterizes human teams as *"a distinguishable set of two or more people who interact dynamically, interdependently, and adaptively towards a common and valued goal/objective/mission"*. Researchers desire to make agents an integral part of teams [20]; however, this desire has not yet been fully realized because current software agents lack the dynamism and adaptiveness in Salas's description of human teams. The next section gives an overview of human teamwork models and team cognition.

2 Human Teamwork

2.1 Representative Theories

Human team processes have been studied by psychologists since the 1950s. Paris et al. [71] group the representative theories influencing our understanding of human teamwork into the following eight categories:

1. social psychological approaches: how team members' relate and interact with each other
2. sociotechnical approaches: work-related implications of team members' relationships and interactions
3. ecological approaches: how organizational or working environments affect teamwork
4. human resource approaches: how teams utilize the members' capabilities and talents
5. technological approaches: relating to technological progress
6. lifecycle approach: how team performance changes during the lifecycle of existence
7. task-oriented approach: team roles, functions, and tasking
8. integrative approach: a fusion of multiple different approaches

Cannon-Bowers et al. [17] divide human teamwork into three dimensions: cognitions, skills, and attitudes. The cognition or knowledge category includes information about the task such as as team mission, objectives, norms, problem models, and resources. Teamwork skills include behaviors such as adaptability, performance monitoring, leadership, communication patterns, and interpersonal coordination. Attitudes measure the participants' feelings about the team: team cohesion, mutual trust, and importance of teamwork.

2.2 Team Cognition

Research in human team performance suggests that experienced teams develop a shared understanding or *shared mental model* utilized to coordinate behaviors by anticipating and predicting each others needs and adapting to task demands [39]. Further, for such teams, both tacit and explicit coordination strategies are important in facilitating teamwork processes. Explicit coordination occurs through externalized verbal and non-verbal communications, whereas tacit coordination is thought to occur through the meta-cognitive activities of team members who have shared mental models of what should be done, when, and by whom [31, 37, 52]. A teams shared mental models thus allow the team members to coordinate their behavior and better communicate depending upon situational demands. Team training researchers have most clearly articulated theories involving shared cognition in general, and definitions of shared mental models in specific. Initial theorizing on training shared mental models suggested that, for teams to successfully coordinate their actions, they must possess commonly held knowledge structures that allow them to predict team behavior based upon shared performance expectations [16]. Generally, this includes knowledge of team objectives and goals but more specifically, it encompasses knowledge of teammates roles and responsibilities along with the team tasks and procedures and the timing/sequencing of the task.

Two important elements of successful communication between humans include the ability for each of the communicators to generally understand what the other person is thinking, and to determine what his/her intentions (or goals) are [30]. For non-living entities, Dennett proposed that humans have three options when interpreting an object's actions: (a) a physical stance, (b) a design stance, or (c) an intentional stance. A physical stance is the application of the laws of nature in predicting what an object will do. A design stance involves one's attempt to make predictions about an object based on their beliefs about the designer's intentions. Finally, an intentional stance is derived from a person's perceptions about the beliefs or desires that they suspect drive the object in question. This last stance, intentional stance, is what people use to read each others minds and predict behaviors.

Another important key to team performance is congruence of team cognition.

Common cognition among team members is associated with higher team effectiveness and is an important element of training human military teams [96, 97]. Commonality of cognition can be measured by rating team member schema similarity (TMSS) [78]. A schema is defined as a set of structured knowledge that humans use to understand incoming stimuli. There are two components of team member schema similarity: team member schema congruence and team member schema accuracy. Congruence is the degree of matching between team members' schema; accuracy is a comparison between team members' schema and the "true score" (quantified by an external observer). Metacognition, "what group members know about the way groups process information" [50], is another important element governing human team performance. If all team members have similar beliefs about the operation and functioning of the group, team performance is improved.

The RETSINA agents (see Section 3) use the ATOM model of teamwork proposed by Smith-Jentsch et al. [96]. The ATOM model postulates that, besides their individual competence in domain specific tasks, team members in high performance teams must have domain independent team expertise, that is comprised of four different categories: information exchange, communication, supporting behavior, team initiative/leadership. The performance of teams, especially in tightly coupled tasks, is believed to be highly dependent on these interpersonal skills.

2.3 Dimensions

Team cognition at the macro level involves many characteristics that affect the collaborative process, the cognitive skills required, and ultimately the quality of the outcome. Understanding the impact of these dimensions is critical to modeling dynamic teamwork and human-agent collaborative processes. Below, we describe the most important dimensions as suggested by Warner et al. [118, 117]:

2.3.1 Collaboration system characteristics

1. Synchronous versus asynchronous collaboration: Is the collaborative process conducted in a same-time manner or are participants collaborating at different times?
2. Proximity of collaborators: Are the participants located proximally or are individuals geographically distributed?

2.3.2 Team characteristics

1. Command structure: Are the participants organized in a hierarchical or flat structure?
2. Homogeneity of knowledge: Do all participants possess the same knowledge or is there information asymmetry?
3. Team size: How many individuals are required to collaborate on a team?

2.3.3 Task dimensions

1. Collaborative output: Is the goal of the team to deliberate and process information or to determine a course of action (COA)?
2. Time stress: Is the team subject to time pressure?
3. Task complexity: How large and complex is the task?
4. Task familiarity: Is the task a one-time or a recurring event?
5. Nature of constituent subtasks: e.g. whether subtasks involve planning, decision making, cognitive conflict, creative and intellective subtasks etc.

3 Agent Teamwork

3.1 Theories

Theoretical work on agent teamwork [110, 46] characterizes team behavior as having the following features: First, the agents need to share the goals they want to achieve, share an overall plan that they follow together and to some degree share knowledge of the environment (situation awareness) in which they are operating. Second, the agents need to share the intention to execute the plan to reach the common goal. Third, team members must be aware of their capabilities and how they can fulfill roles required by the team high level plan. Fourth, team members should be able to monitor their own progress towards the team goal and monitor team mates activities and team joint intentions [22]. Using these basic teamwork ideas, many systems have been successfully implemented, including teams supporting human collaboration [18], teams for disaster response [66], and for manufacturing [108].

3.2 Frameworks

In addition to identifying suitable roles for agents to play in human teams, to implement a software system, we must also select coordination and communication mechanisms that the agents can use. For some domains, simple pre-arranged coordination schemes like the locker-room agreement [100] in which the teams execute pre-selected plans after observing an environmental trigger are adequate. Although this coordination model has been successful in the Robocup domain, the locker-room agreement breaks down when there is ambiguity about what has been observed; what happens when one agent believes that the trigger has occurred but another agent missed seeing it? The TEAMCORE framework [110, 111] was designed to address this problem; the agents explicitly reason about goal commitment, information sharing, and selective communication. This framework incorporates prior work by Cohen and Levesque [21] on logical reasoning about agent intention and goal abandonment. Having agents capable of reasoning about fellow agents' intentions makes the coordination process more reliable, since the agents are able to reason about sensor and coordination failures. By giving all team members proxies imbued with this reasoning capability, it is possible to include agents, robots, and humans in a single team [87].

Similarly other formalisms such as SharedPlans [46] have been successfully used by collaborative interface agents used to reason about human intentions. The RETSINA software framework [107] uses reasoning mechanisms based on SharedPlans to:

1. identify relevant recipients of critical information and forward information to them
2. track task interdependencies among different team members
3. recognize and report conflicts and constraint violations
4. propose solutions to resolve conflicts
5. monitor team performance

To be an effective team member, besides doing its own task well, an agent must be able to receive tasks and goals from other team members, be able to communicate the results of its own problem solving activities to appropriate participants, monitor team activity, and delegate tasks to other team members. A prerequisite for an agent to perform effective task delegation is (a) to know which tasks and actions it can perform itself (b) which of its goals entail actions that can be performed by others and (c) who can perform a given task. The RETSINA agent architecture [103] includes a communication module that allows agents to send messages, declarative representation of agent goals and planning mechanisms for fulfilling these goals. Therefore, an agent is aware of the objectives it can plan for and the tasks it can perform. In addition, the planning mechanism allows an agent to reason about actions that it cannot perform itself but which should be delegated to other agents.

Adjustable autonomy, the agent's ability to dynamically vary its own autonomy according to the situation, is an important facet of developing agent systems that interact with humans [88, 89]. Agents with adjustable autonomy reason about transfer-of-control decisions, assuming control when the human is unwilling or unable to do a task. In many domains, the human teammates possess greater task expertise than the software agents but less time; with adjustable autonomy the human's time is reserved for the most important decisions while the agent team members deal with the less essential tasks. Scerri et al. [88] demonstrated the use of Markov decision processes to calculate an optimal multiple transfer-of-control policy for calendar scheduling user interface agents. Having agents with adjustable autonomy, is beneficial to agent teams. For example, a robot may ask a software agent for help in disambiguating its position; a software agent may relinquish control to a human to get advice concerning the choice between two decision making alternatives; a human can relinquish control to a robot in searching for victims. However there are many interesting research challenges that remain: how control can be relinquished in ways that do not cause difficulties to the team, how to maintain team commitments, how to support large-scale interactions with many agents.

To coordinate distributed autonomous agents into effective teams in dynamic environments, we embrace the following principles: (1) the environment is open and the team constitution could vary dynamically through addition, substitution or deletion of teammates; (2) team members are heterogeneous (having different or partially overlapping capabilities); (3) team members share domain independent teamwork models; (4) individual and team replanning is necessary while supporting team goals and commitments; (5) every team member can initiate team goals. In forming a team to execute a mission team members with different capabilities will be required. The location and availability of potential teammates is not known at any given point in time. Moreover, during a mission, teams may have to be reconfigured, due to loss, total or partial, of team member capabilities (e.g. a robot loses one of its sensors) or necessity of adding team members to the team. Automated support for addition, substitution or deletion of team members requires extensions to current teamwork models: (1) development of robust schemes for agents to find others with required capabilities, i.e. agent discovery; (2) development of robust algorithms for briefing new team members so as to make them aware of the mission and current plans of the team; and (3) individual role adjustment and (re)negotiation of already existing plans due to the presence of the new (substitutable) team mate. We collectively call these three parts capability-based coordination.

After a team has been formed (or reconfigured) through capability based coordination, team members must monitor teammates activities in order to pursue and maintain coherent team activity, team goals and commitments. Once team goals have been formulated, team members perform domain specific planning, information gathering and execution according to their expertise while maintaining team goals. An agent's planning must take into consideration temporal constraints and deadlines, as well as resource constraints and decision tradeoffs.

3.3 Plan Execution for Agent Teams

Theoretical work on team behavior [22, 45] stresses that the agents need to share the goals they want to achieve, a plan that they follow together and knowledge of the environment in which they are operating. In addition they need to share the commitment to execute the plan to reach the common goal. Furthermore, team members must be aware of their capabilities and how they can fulfill roles required by the team high level plan, should be able to monitor their own progress towards the team goal and monitor team mates activities and team joint intentions. The theoretical work and the operationalization of these representational and inferential abilities constitute a generic model of teamwork. System implementation of team coordination [110] has resulted in the creation of an agent wrapper [74] that implements a generic model of teamwork. An instance of such a wrapper can be associated with any agent in the multiagent system and used to coordinate the agents activity with team activity. Teamwork wrappers can be used to wrap non-social agents to enable them to become team oriented.

Besides capability-based coordination, discussed above, for open agent societies, and domain-specific multiagent planning information gathering and execution that is discussed in the next section, several enhancements to the agent teamwork models that have been reported thus far are necessary to adapt these models to human-agent teams.

3.3.1 Goal Creation

Any agent can generate a team goal, thus becoming a team initiator. To become a team initiator requires the ability to perceive and assess events as meriting initiation of team formation activities. The team initiator must be able to generate a skeletal team plan, determine roles for himself (by matching his capability to plan requirements) and find additional teammates through capability-based coordination . The team initiator is also responsible for checking that the collective capabilities of the newly formed team cover all requirements of the team goal.

3.3.2 Proactive (Reactive) Mutual Assistance and Altruism

Current models of teamwork are agnostic with respect to agent attitudes but implicitly assume a general cooperative attitude on the part of individuals that make them willing to engage in teamwork. Experimental work on human high performance teams [73], with numerous human military subjects in time-stressed and dangerous scenarios, has demonstrated that attitudes of team help are important factors in achieving high performance. Currently, there is no theoretical framework that specifies how such agent attitudes can be expressed; whether it is possible to incorporate such attitudes into current teamwork theories (e.g. joint intentions or SharedPlans); what additional activities such attitudes entail during teamwork. In prior work, we have explored such agent attitudes in the context of software agent assistance provided to human teams [43]. Agents unobtrusively eavesdropped on human conversations to determine proactively what kind of assistance humans could use (e.g. finding flight schedules in the context of a non combatant evacuation scenario), and engaged in capability-based coordination to find agents that could supply the needed information.

3.3.3 Monitoring Individual and Team Activity

One consequence of agents attitude of proactive assistance is a clearly increased need for team monitoring. In prior work, monitoring teammates activities was done so as to maintain joint intentions during plan execution. Therefore, monitoring was done [110] to ascertain (a) team role non-performance, (e.g. a team member no longer performs a team role) or (b) whether some new event necessitates team goal dissolution (e.g. realizing that the team goal has already been achieved). When developing schemes for tracking intentions of heterogeneous human-agent teams and dealing with the issue of appropriate levels of information granularity for such communications, additional monitoring must be done as a consequence of the proactive assistance agent attitude. Agents should volunteer to get information that is perceived to be useful to a teammate or the team as a whole. Moreover, agents should send warnings, if it perceives a teammate to be in danger (e.g. Agent A warns Robot B of impending ceiling collapse in Bs vicinity). Additional monitoring mechanisms, such as timeouts and temporal- and location-based checkpoints that are established during briefing, may also be useful. Timeouts are used to infer failure of an agent to continue performing its role (e.g. if it has not responded to some warning message for example). In large dynamic environments, monitoring individual and team activity via plan recognition is not possible since the agents activities are not directly perceivable by others most of the time. Moreover plan recognition is computationally expensive. Hence, we assume that communication (via Agent Communication Languages such as FIPA or KQML) is the main means by which agents monitor one anothers activities. The content of the communication can be expressed in a declarative language (e.g. XML, or DAML). Agents communicate significant events, e.g. a new victim may have been heard; events pertaining to their own ability to continue performing their role, e.g. I lost my vision sensor; requests for help, e.g. can someone who is nearby come help me lift this rubble? These types of communication potentially generate new team subgoals (e.g. establish a team subgoal to from subteam to get the newly heard victim to safety) and possibly the formation of subteams. The new team subgoal can be initiated by the sender or the receiver of the message.

4 Human-Agent Teamwork

The Webster dictionary defines teamwork as *"work done by several associates with each doing a part but all subordinating personal prominence to the efficiency of the whole."* How well do agent-human teams fit this definition of teamwork? By the strict dictionary definition, interacting with your web browser might qualify as teamwork since the web browser is working while subordinating personal prominence. For the purpose of this survey, we limit our discussion of agents to pieces of software that 1) are autonomous, capable of functioning independently for a significant length of time 2) proactively act in anticipation of future events 3) are capable of self-reflection about their and their teammates' abilities. When a group of actors coordinates via teamwork, they can flexibly and robustly achieve joint goals in a distributed, dynamic and potentially hostile environment.

4.1 Agent Roles within Human Teams

Researchers desire to make agents an integral part of teams [20]; however, this desire has not yet been fully realized. Researchers must identify how to best incorporate agents into human teams and what roles they should assume. Sycara and Lewis [106] identified three primary roles played by agents interacting with humans:

agents supporting individual team members in completion of their own tasks These agents often function as personal assistant agents and are assigned to specific team members [18]. Task-specific agents utilized by multiple team members (e.g. [19]) also belong in this category.

agents supporting the team as a whole Rather than focusing on task-completion activities, these agents directly facilitate teamwork by aiding communication, coordination among human agents, and focus of attention. The experimental results summarized in [106] indicate that this might be the the most effective aiding strategy for agents in hybrid teams.

agents assuming the role of an equal team member These agents are expected to function as "virtual humans" within the organization, capable of the same reasoning and tasks as their human teammates [114]. This is the hardest role for a software agent to assume, since it is difficult to create a software agent that is as effective as a human at both task performance and teamwork skills.

4.1.1 Research Challenges

Creating shared understanding between human and agent teammates is the biggest challenge facing developers of mixed-initiative human/agent organizations. The limiting factor in most human-agent interactions is the user's ability and willingness to spend time communicating with the agent in a manner that both humans and agents understand, rather than the agent's computational power and bandwidth [106]. Horvitz [53] formulates the problem of mixed-initiative user interaction as a process of managing uncertainties: (1) managing uncertainties that agents may have about user's goals and focus of attention (2) uncertainty that users' have about agent plans and status. Regardless of the agents' role, creating agent understanding of user intent and making agents' results intelligible to a human are problems that must be addressed by any mixed-initiative system, whether the agents reduce uncertainty through communication, inference, or a mixture of the two. Also, protecting users from unauthorized agent interactions is always a concern in any application of agent technology.

Additionally, there are additional research challenges, specific to the role assumed by the agent. Agents that support individual human team members face the following challenges: 1) modeling user preferences; 2) determining optimal transfer-of-control policies [87]; 3) considering the status of user's attention in timing services [53]. Agents aiding teams [60, 59, 57, 56], face a additional set of problems: 1) identifying information that needs to be passed to other team members before being asked; 2) automatically prioritizing tasks for the human team members; 3) maintaining shared task information in a way that is useful for the human users. Agents assuming the role of equal team members [114, 34, 33] must additionally be able to: 1) competently execute their role in the team 2) critique team errors; 3) independently suggest alternate courses of action.

Human-agent teams have been used in a variety of applications from: command and control scenarios [13, 120], disaster rescue simulations [90], team training in virtual environments [114], and personal information management [18]. Also recently there has been increasing interest in the problem of creating effective human-robot interfaces [67, 49, 98, 4, 102, 116]. Since these applications have widely different requirements, the generality of the models and results between domains is questionable. The following distinctions are instructive:

1. how many humans and agents are there in the team? Are the agents supporting a team of humans or is it a team of agents supporting one user?
2. how much interdependency is there between agents and humans? Can the humans perform the task without the agents?
3. are the agents capable of unsolicited activity or do they merely respond to commands of the user?

4.1.2 Agents Supporting Team Members

In this class of applications, the software agents aid a single human in completing his/her tasks and do not directly interact with other human team members. The two organizational structures most commonly found in these types of human/agent teams are: 1) each human is supported by a single agent proxy. Agent proxies interact with other agents to accomplish the human's tasks; 2) each human is supported by a team of agents that work to accomplish the single human's directives. Often there are no other humans involved in the task, and the only "teamwork" involved is between the software agents. Examples of these type of agent systems include: agents assisting humans in allocating disaster rescue resources [90] and multi-robot control systems in which teams of robot perform tasks under the guidance of a human operator [98].

4.1.3 Agents Acting as Team Members

Instead of merely assisting humans team members, the software agents can assume equal roles in the team, sometimes replacing missing human team members. It can be challenging to develop software agents of comparable competency with human performers unless the task is relatively simple. Agents often fulfill this role in training simulation applications, acting as team members or tutors for the human trainees. Rickel and Johnson [81] developed a training simulator to teach human boat crews to correctly respond to nautical problems, using STEVE, a SOAR based agent with a graphical embodiment. The Mission Rehearsal Environment [114] is a command training simulation that contains multiple "virtual humans" who serve as subordinates to the human commander trainee. The human must negotiate with the agents to get them to agree to the correct course of action. It is less common in human-robot applications to have robots acting as team members, rather than supporters; however limited examples of human-robot teamwork are starting to emerge in the Segway Robocup division where each soccer team is composed of a Segway-riding human paired with a robotically-controlled Segway [3].

4.1.4 Agents Supporting Human Teams

In this class of applications, the agents facilitate teamwork between humans involved in a group task by aiding communication, coordination, and focus of attention. In certain applications, this has shown to be more effective than having the agents directly aid in task completion. For the TANDEM target-identification control and command task, Sycara and Lewis [106] examined different ways of deploying agents to support multi-person teams. Different agent-aiding strategies were experimentally evaluated within the context of a group target-identification task : 1) supporting the individual by maintaining a common visual space; 2) supporting communication among team members by automatically passing information to the relevant team member; 3) supporting task prioritization and coordination by maintaining shared checklist. The two team-aiding strategies (supporting communication and task prioritization) improved team performance significantly more than supporting team members with their individual tasks. Aiding teamwork also requires less domain-knowledge than task-aiding which makes the agents potentially reusable across domains.

4.2 Human-Agent Interaction

We suggest that three important facets of human-agent interaction are: (a) mutual predictability of teammates [106] (b) team knowledge (shared understanding), and (c) ability to redirect [20] and adapt to one another. Mutual predictability means that parties must make their actions sufficiently predictable to the teammates to make coordination effective and also try to form some estimate of many features of the team activity (e.g., how long it will take a teammate to perform an action). Team knowledge refers to the pertinent mutual knowledge, beliefs and assumptions that support the interdependent actions and the construction or following of plans to fulfill the team goals. Team knowledge (shared understanding/shared knowledge) could exist before the team is formed (through previous experiences) or must be formed, maintained and repaired after the team has started its collaboration. Directability and mutual adaptation is a key component of teamwork because it expresses the interdependencies of team activity. If the way a player performs an activity has no effect on another, then the two players work in parallel but do not coordinate.

4.2.1 Team Knowledge

Underpinning team knowledge, mutual predictability, directability and adaptation is clear and effective knowledge transfer. Within the team context, knowledge transfer usually occurs through communication between team members across different channels (e.g., verbal, text, nonverbal). Clear and effective transfer of knowledge between is essential in human teams [99]. Team knowledge is facilitated through effective communication of individual understanding. Knowledge representation is affected by interaction among team members as they are engaged in stages of the OODA loop: observing, orienting, deciding, and acting.

A number of theories developed in the communication discipline for human-only interaction may be adapted for human-agent coordination and communication. Humans form norms about behaviors they expect from both humans and agents. Expectancy Violations Theory (EVT) [11] explains reactions to behaviors that violate norms and expectations for usual interaction. Another view, Interaction Adaptation Theory (IAT) [12], articulates how an actors goals, expectations, and requirements for an interaction and interaction partner generate either reciprocal or compensatory behavioral patterns in response to the partners communication. Additionally, Adaptive Structuration Theory [44], helps to address the joint evolution of technology and social systems. It emphasizes the social and intrapersonal affects that the use of technology has on groups, as opposed to the technology itself. These interactions often affect how the technology is used. Interactions create or adapt structures to maintain and manage information. This theory has specifically has been used to examine new forms of interaction and communication technologies.

Team knowledge is critical to understanding team performance because it explains how members of effective teams interact with one another [15]. Team knowledge is also termed shared understanding, collective cognition, shared cognition, team mental models, shared knowledge, transactive memory, shared mental models, and so forth [55, 69]. Team knowledge does not refer to a unitary concept; it refers to different types of knowledge that need to be shared in effective teams. Teams build knowledge about specific tasks (both declarative and procedural task-specific knowledge), items related to tasks (e.g. expectations of how teams operate), characteristics of teammates (e.g. strengths, preference, weaknesses, tendencies or each individual), and attitudes and beliefs of teammates [16]. Knowledge of the strength and weaknesses of teammates and their attitudes and beliefs is generalizable across a variety of tasks. Knowledge that is task-related can be used across similar tasks.

Increased knowledge and understanding in any of these categories should lead to increased task performance. Team knowledge has been hypothesized to explain variance in team development, team performance, strategic problem definition, strategic decision making, and organization performance [55]. Expert teams have even been shown to operate without communication when the team exhibits high team knowledge [16]. Research has shown that when team members share knowledge, team knowledge enables them to interpret cues in a similar manner, make compatible decisions, and take appropriate action [55, 65].

Team knowledge and shared understanding need to be formed between humans and agent despite the presence of multiple representations. As Cooke et al. [24] point out members of human teams have both shared and unshared knowledge of the teams task and state. Agents are likely to have varying levels of intelligibility to other members of their team because their tasks and conditions they are responding to will be known to different extents by other team members. One of the ways to address this is through customization of the agent communication for each team member based on the agents estimation of what the human teammate knows. Another way is to always give human team members the maximum amount of information the agent considers relevant but without customization. Team knowledge implies that the agents will have a clear idea of important features of the team composition. For agents to perform as full fledged team members, they must have a clear model of the teams goals, its members, their capabilities, and their roles in procedurally defined tasks.

4.2.2 Mutual Predictability

Mutual predictability, as well as team knowledge, entails knowledge transfer between team members. It enables teammates to communicate and coordinate in a meaningful manner. Humans represent most knowledge implicitly in their brains. This knowledge needs to be represented in some explicit manner for other teammates to understand it. Some work explores representing knowledge within human teams using the concept of collective intelligence [95]. This work builds on several Information Processing System (IPS) models and architectures of individual cognition to identify key components and functions within collaborative groups.

The introduction of agents into teams creates impediments to mutual predictability. The greatest impediment to agents assisting human users lies in communicating their intent and making results intelligible to them [62]. To this end, representation schemes that are declarative and intelligible both by humans and agents are most useful. Research on knowledge representation within agents is primarily based on logic-based formalisms [8]. High level messaging languages such as KQML contain message types intelligible both to agents and humans (e.g. inform, tell, ask) and have been used in systems such as RETSINA for successful human agent collaboration. Different forms of communication (e.g. text, pictures, menus) might vary in effectiveness as carriers of knowledge transfer between teammates, both human and agent.

Situation Theory [6] suggests that the meaning transferred between teammates should be a function of shared knowledge, shared context, and the communication itself. While simple co-reference problems have been studied for reconnaissance and map reading it is likely that team problems involving the sharing of hidden knowledge, convergence to a team mental model, and team pattern recognition [117] may also be impeded by differences in external representation.

4.2.3 Directability and Mutual Adaptation

Directability and mutual adaptation enable teams to be flexible and agile in different contexts and task requirements. Directability refers to assigning roles and responsibilities to different team members. Mutual adaptation defines how team members alter their roles to fulfill the requirements of the team. Researchers acknowledge that the most effective agents will need to change their level of initiative, or exhibit adjustable autonomy, in response to the situation to be most effective [53, 89]. For agents to appropriately exercise initiative, they must have a clear model of the teams goals, member roles and team procedures. We suggest that Expectancy Violation Theory is a basis for determining how agent members of a team should adjust their autonomy. Agents can execute transfer-of-control strategies [88] which specify a sequence of changes of autonomy level in the case a human is occupied and cannot make timely decisions, but such strategies must be designed to fit with human expectations, rather than violating them. An important research question is whether teams perform better when agents have a constant, but potentially sub-optimal level of autonomy, or when agents constantly adjust to the teams context.

4.2.4 Communication

Underpinning team knowledge, mutual predictability, and directability and mutual adaptation is clear and effective communication. These constructs are extremely difficult to formalize and measure. Adding agents to human teams introduces additional complexity. Cooke et al. [27] suggest that shared understanding and team cognition might be measured by extracting the level and structure of knowledge from each team member and then measuring the team process behaviors. The level and structure of knowledge from each group member may be obtained through existing, validated psychometric tests and questionnaires.

Measuring group process behaviors is much more difficult. Group process behaviors are most evident in communication between group members. Thus substantial knowledge and experience in human communication, computer-aided collaboration, and human-agent interaction are required to analyze these behaviors. Since communication between team members underlies team knowledge, mutual predictability, and shared directability and mutual adaptation, developing and validating measures of group process behaviors is crucial.

4.3 Military Applications

Although it is instructive to examine different applications of human-agent teamwork, we focus our discussion on applications of human-agent teamwork to military planning. Joint military planning is a fruitful domain for software agent technology because: 1) it is often done with teams of geographically-separated individuals who are communicating over voice or electronic channels that can be monitored by software agents; 2) it can be time-stressed, especially in a crisis situation; 3) even minor improvements in team performance can be important since the cost of inadequate mission planning is high.

Cognitive activities abound in the network-centric battlefield. Teams plan, make decisions, design, debrief, share knowledge, and execute actions as an integrated unit. These are team-level cognitive activities that are the focus of work on team cognition [86]. The processes and structures associated with team cognition impact team performance and thus are central targets for assessment and remediation. One aspect of team cognition that is prominent in networked environments is team coordination. We view team coordination as an essential component of networked performance that can benefit from training.

The network-centric battlefield demands intense coordination among network effectors (humans and automation) that are part of a larger interconnected social organization. In this context we define coordination as the timely and adaptive distribution of information among network effectors. We think of team coordination as analogous to cognitive processing at the individual level. Coordination is challenging in network-centric environments because entities are often geographically dispersed and may be unfamiliar with other entities as well as the specific task or mission. This situation leads to what has been called "team opacity" [38]. and has been frequently associated with differences in process behaviors, poorer shared understanding, and lean communication, relative to co-located teams [26]. In fact, teams often adapt to these situations through spontaneous self-organization of their coordination structure [25].

It is important to note that we do not consider coordination in information theoretic terms [92] in which information is encoded, decoded and passively moved from effector to effector with some degree of uncertainty based on channel capacity. Rather, coordination involves active communication or mediation among effectors in a social network [41]. Consequently, our coordination metrics do not measure amount of information passed or uncertainty, but instead extend social network theory or coordination theory by quantifying the effectiveness of coordination patterns.

Team coordination in network-centric battlefield settings is predictive of the performance of the team, and to some degree, the social system in which the team is embedded. However, team coordination is not identical to team performance. Sometimes poor coordination can result in fortuitously positive outcomes and even the best coordination can sometimes fail to prevent a negative outcome. Coordination is, however, a precursor of team performance, and in our view, a critical precursor for the network-centric battlefield, in that effector competencies, as well as effectors themselves, are dispersed across the battlefield.

Based on our experimental data coordination improves with team experience and training, but decays over long retention intervals [25]. The development of coordination skill is a large part of the development of collective competence of the social group. Coordination, therefore, is a team skill that can be trained. It is also a skill that can be quantified and modeled. The measurement and modeling of the development of coordination in networked command and control is challenging due to the nonlinearities associated with interactions in complex distributed systems [26]. For instance, coupled effectors have capabilities for contributing second-hand information to the information available in the local environments of other, reciprocating effectors. This positive feedback mechanism entails nonlinear changes in overall system state as information is adaptively dissipated through the system.

Fan et al. [34, 33] have evaluated the use of cognitive agents with a collaborative Recognition-Primed Decision model (RPD) for supporting human teams performing military decision tasks. In one task [34], the human teams had to maximize the amount of supplies delivered to the troops by successfully protecting an airport. The RPD agents were able to participate in human teams and increase the amount of supplies delivered when the humans were under time-stress. In a second task, C2 teams had to react to incoming threats menacing a metropolis from crowds, insurgents, and improvised explosive devices. The task is difficult because it is 1) real-time and 2) involves context-switching. The human-agent C2 teams performed much better than the human-only teams at the same level of task complexity; moreover the human-agent team performances are significantly more stable than the human-only performances. These evaluations of human-agent teams for military tasks are encouraging because they demonstrate that agents can produce a measurable difference at task performance in this domain.

Using the RETSINA software architecture we were able to develop a set of agents to aid a human-team during a simulation of Noncombatant Evacuation Operation planning [106]. In the scenario, three human decision-makers (a U.S. ambassador, Joint Forces Commander, and USTRANSCOM) must coordinate to evacuate civilians from a developing protest situation in Kabul Afghanistan.

The agent team consisted of three types of agents:

task agents capable of specialized tasks Examples include the Weather Agent, a weather information provider, the Visual Recognition agent, which interprets visual surveillance information, and Route Planning agents that can plot cost minimizing routes according to a set of constraints.

middle agents: agents that facilitate the discovery and matching of other agents

interface agents: agents that interact with human team members. Examples include the Voice Agent, an agent that receives voice input from human team members, and Messenger Agents that infer the humans' goals from conversation transcripts.

The NEO system demonstration includes agents that fulfill all possible agent roles.

1. Aiding an individual human team member in information gathering or planning tasks
2. Acting as team members themselves. In this capacity, RETSINA agents: (a) provide proactive and reactive information and planning support, (b) perform information gathering and planning tasks to promote the team goals, (c) perform task decomposition and task allocation to other members of the agent team so as to efficiently contribute to the team goals.
3. Aiding the human team as a whole. In this capacity RETSINA agents (a) provide situation assessment, (b) monitoring and alerting team members to important situation changes, (c) communicating their results in unambiguous, multimodal and non-intrusive ways, (d) discover (through middle agents) suitable additional team members and information sources that can aid the team.

Unlike most examples of human teams where the team members are statically known *a priori*, RETSINA does not make any such closed world assumptions but allows dropping, adding, and discovering new teammates dynamically. This functionality reflects the requirements of real situations (especially military situations where teammates may become incapacitated and others must be found to take up their roles).

5 Conclusion

Agent assistance will be particularly critical to military teams as their operations become more agile and situation specific. As unfamiliar forces are brought together for one time missions the infosphere they establish between their networked information systems will become a primary mechanism for coordination. In this uncertain environment supporting teamwork becomes crucial. Our results suggest that software agents are well suited for this task. Because the domain independence of teamwork agents would allow them to be rapidly deployed across a broad range of tasks and settings teamwork appears to be a particularly high payoff area for future agents research.

6 Acknowledgements

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation heron.

References

[1] J. Allen. Mixed initiative interaction. *IEEE Intelligent Systems*, 14(5), 1999.

[2] J. Allen and G. Ferguson. Human-machine collaborative planning. In *NASA Planning and Scheduling Workshop*, 2002.

[3] B. Argall, Y. Gu, and B. Browning. The first Segway soccer experience: Towards peer-to-peer human robot teams. In *Proceedings of the First Annual Conference on Human-Robot Interactions*, 2006.

[4] J. Atherton, B. Harding, and M. Goodrich. Coordinating a multi-agent team using a multiple perspective interface paradigm. In *Proceedings of the AAAI 2006 Spring Symposium: To Boldly Go Where No Human-Robot Team Has Gone Before*, 2006.

[5] S. Banks and C. Lizza. Pilot's associate: A cooperative, knowledge-based system application. *IEEE Intelligent Systems and their Applications*, 6(3), 1991.

[6] J. Barwise and J. Perry. *Situations and Attitudes*. Bradford Books, 1983.

[7] M. Beyerlein, D. Johnson, and S. Beyerlein, editors. *Virtual teams: Advances in interdisciplinary studies of work teams*. Elsevier Science/JAI Press, 2001.

[8] R. Brachman and H. Levesque. *Knowledge representation and reasoning*. Margan Kaufmann, 1984.

[9] J. Bradshaw et al. Teamwork centered autonomy for extended human-agent interaction in space applications. In *Proceedings of the AAAI 2004 Spring Symposium*, 2004.

[10] J. Bradshaw, M. Sierhuis, A. Acquisti, Y. Gawdiak, D. Prescott, R. Jeffers, N. Suri, and R. van Hoof. What we can learn about human-agent teamwork from practice. In *Workshop on Teamwork and Coalition Formation, Autonomous Agents and Multi-agent Systems*, 2002.

[11] J. Burgoon. A communication model of personal space violation: Explication and an initial test. *Human Communication Research*, 4, 1978.

[12] J. Burgoon, L. Stern, and L. Dillman. *Interpersonal Adaptation: Dyadic Interaction Patterns*. Cambridge University Press, 1995.

[13] M. Burstein and D. Diller. A framework for dynamic information flow in mixed-initiative human/agent organizations. *Applied Intelligence*, 20(3), 2004.

[14] J. Cannon-Bowers and E. Salas. Individual and team decision making under stress: theoretical underpinnings. In J. Cannon-Bowers and E. Salas, editors, *Making decisions under stress: Implications for individual and team training*. American Psychological Association, 1998.

[15] J. Cannon-Bowers and E. Salas. Reflections on shared cognition. *Journal of Organizational Behavior*, 22, 2001.

[16] J. Cannon-Bowers, E. Salas, and S. Converse. Shared mental models in expert team decision making. In J. Castellan, editor, *Individual and group decision making*. Lawrence Erlbaum Associates, 1993.

[17] J. Cannon-Bowers, S. Tannenbaum, E. Salas, and C Volpe. Defining team competencies: implications for training requirements and strategies. In R. Guzzo and E. Salas, editors, *Team Effectiveness and Decision Making in Organizations*. Jossey Bass, 1995.

[18] H. Chalupsky et al. Electric Elves: Applying agent technology to support human organizations. In *Proceedings of the Innovative Applications of Artificial Intelligence Conference*, 2001.

[19] L. Chen and K. Sycara. Webmate: a personal agent for browsing and searching. In *Proceedings of the Second International Conference on Autonomous Agents*, 1998.

[20] K. Christoffersen and D. Woods. How to make automated systems team players. *Advances in Human Performance and Cognitive Engineering Research*, 2, 2004.

[21] P. Cohen and H. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42, 1990.

[22] P. Cohen and H. Levesque. Teamwork. *Nous*, 25(4), 1991.

[23] P. Cohen, H. Levesque, and I. Smith. On team formation. In G. Holmstrom-Hintikka and R. Tuomela, editors, *Contemporary Action Theory*. Kluwer Academic, 1997.

[24] N. Cooke et al. Measuring team knowledge. *Human Factors*, 42, 2000.

[25] N. Cooke and J. Gorman. Assessment of team cognition. In P. Karwowski, editor, *International Encyclopedia of Ergonomics and Human Factors*. Taylor and Francis Ltd., 2007.

[26] N. Cooke, J. Gorman, M. Pedersen, and B. Bell. Distributed mission environments: Effects of geographic distribution on team cognition, process, and performance.

In S. Fiore and E. Salas, editors, *Towards a science of distributed learning and training*. American Psychological Association, 2007.

[27] N. Cooke, E. Salas, P. Kiekel, and B. Bell. Advances in measuring team cognition. In E. Salas and S. Fiore, editors, *Team Cognition: Understanding the Factors that Drive Process and Performance*. American Psychological Association, 2004.

[28] K. Decker, A. Pannu, K. Sycara, and M. Williamson. Designing behaviors for information agents. In *Proceedings of the First International Conference on Autonomous Agents*, 1997.

[29] K. Decker, A. Pannu, K. Sycara, and M. Williamson. Middle-agents for the internet. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, 1997.

[30] D. Dennett. *Kinds of Minds: Towards an Understanding of Consciousness*. Basic Books, 1997.

[31] E. Entin and D. Serfaty. Adaptive team coordination. *Human Factors*, 41, 1999.

[32] J. Espinosa, F. Lerch, and R. Kraut. Explicit vs. implicit coordination mechanisms and task dependencies: One size does not fit all. In E. Salas and S. Fiore, editors, *Team Cognition: Understanding the Factors that Drive Process and Performance*. American Psychological Association, 2004.

[33] X. Fan, B. Sun, S. Sun, M. McNeese, and J. Yen. RPD-Enabled agents teaming with humans for multi-context decision making. In *Proceedings of International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2006.

[34] X. Fan, S. Sun, M. McNeese, and J. Yen. Extending the recognition-primed decision model to support human-agent collaboration. In *Proceedings of International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2005.

[35] X. Fan and J. Yen. Modeling and simulating human teamwork behaviors using intelligent agents. *J. of Physics of Life Reviews*, 1(3), 2004.

[36] X. Fan, J. Yen, and R. Volz. A theoretical framework on proactive information exchange in agent teamwork. *Artificial Intelligence*, 169:23–97, 2005.

[37] S. Fiore, E. Salas, and J. Cannon-Bowers. Group dynamics and shared mental model development. In M. London, editor, *How people evaluate others in organizations: Person perception and interpersonal judgment in industrial/organizational psychology*. Lawrence Erlbaum Associates, 2001.

[38] S. Fiore, E. Salas, H. Cuevas, and C. Bowers. Distributed coordination space: towards a theory of distributed team performance. *Theoretical Issues in Ergonomic Science*, 4, 2003.

[39] S. Fiore and J. Schooler. Process mapping and shared cognition: Teamwork and the development of shared problem models. In E. Salas and S. Fiore, editors, *Team Cognition: Understanding the Factors that Drive Process and Performance*. American Psychological Association, 2004.

[40] H. Foushee. Dyads and triads at 35,000 feet: Factors affecting group process and aircrew performance. *American Psychologist*, 39, 1984.

[41] N. Friedkin. *A structural theory of social influence*. Cambridge University Press, 1998.

[42] J. Giampapa, M. Paolucci, and K. Sycara. Agent interoperation across multiagent boundaries. In *Proceedings of the Fourth International Conference on Autonomous Agents*, 2000.

[43] J. Giampapa and K. Sycara. Conversational case-based planning for agent team coordination. In *Proceedings of the Fourth International Conference on Case-Based Reasoning*, 2001.

[44] A. Giddens. *The Constitution of Society: Outline of the Theory of Structuration*. University of California, Berkeley, 1984.

[45] B. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86, 1996.

[46] B. Grosz and S. Kraus. The evolution of SharedPlans. In A. Rao and M. Wooldridge, editors, *Foundations and Theories of Rational Agency*. Kluwer Academic, 2003.

[47] Y. Gu and M. Veloso. Effective team-driven multi-model motion tracking. In *Proceedings of the First Annual Conference on Human-Robot Interactions*, 2006.

[48] C. Gutwin and S. Greenberg. The importance of awareness for team cognition in distributed collaboration. In E. Salas and S. Fiore, editors, *Team Cognition: Understanding the Factors that Drive Process and Performance*. American Psychological Association, 2004.

[49] T. Harris, S. Banerjee, and A. Rudnicky. Heterogeneous multi-robot dialogues for search tasks. In *Proceedings of the AAAI Spring Symposium Intelligence*, 2005.

[50] V. Hinsz. Metacognition and mental models in groups: An illustration with metamemory of group recognition memory. In E. Salas and S. Fiore, editors, *Team Cognition: Understanding the Factors that Drive Process and Performance*. American Psychological Association, 2004.

[51] R. Hoeft. *Investigating the mechanisms that drive implicit coordination in teams*. PhD thesis, University of Central Florida, Orlando, 2006.

[52] R. Hoeft, J. Kochan, and F. Jentsch. Automated team members in the cockpit: Myth or reality. In A. Schulz and L. Parker, editors, *Series: Advances in Human Performance and Cognitive Engineering Research*. Elsevier Science, 2006.

[53] E. Horvitz. Principles of mixed-initiative user interfaces. In *Proceedings of SIGCHI*, 1999.

[54] F. Jentsch, S. Fiore, S. Rehfeld, and E. Salas. Military teams and tele-operation of unmanned ground vehicles: An initial literature review. Technical Report TPL-03-2004, University of Central Florida, Orlando, 2004.

[55] R. Klimoski and S. Mohammed. Team mental model: construct or metaphor? *Journal of Management*, 20, 1994.

[56] T. Lenox. *Supporting teamwork using software agents in human-agent teams*. PhD thesis, Westminster College, 2000.

[57] T. Lenox, S. Hahn, M. Lewis, T. Payne, and K. Sycara. Agent-based aiding for individual and team planning tasks. In *Proceedings of IEA 2000/HFES 2000 Congress*, 2000.

[58] T. Lenox, S. Hahn, M. Lewis, T. Payne, and K. Sycara. Task characteristics and intelligent aiding. In *Proceedings of International Conference on Systems, Man, and Cybernetics*, 2000.

[59] T. Lenox, M. Lewis, E. Roth, R. Shern, L. Roberts, T. Rafalski, and J. Jacobson. Support of teamwork in human-agent teams. In *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics*, 1998.

[60] T. Lenox, L. Roberts, and M. Lewis. Human-agent interaction in a target identification task. In *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics*, 1997.

[61] H. Levesque and P. Cohen. On acting together. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*, 1990.

[62] M. Lewis. Designing for human-agent interaction. *AI Magazine*, 19(2), 1998.

[63] J. MacMillan, E. Entin, and D. Serfaty. Communication overhead: The hidden cost of team cognition. In E. Salas and S. Fiore, editors, *Team Cognition: Understanding the Factors that Drive Process and Performance*. American Psychological Association, 2004.

[64] M. McNeese, E. Salas, and M. Endsley, editors. *New trends in collaborative activities: understanding system dynamics in complex environments*. Human Factors and Ergonomics Society, 2001.

[65] S. Mohammad and B. Dumville. Team mental models in a team knowledge framework: expanding theory and measurement across disciplinary boundaries. *Journal of Organizational Behavior*, 22, 2001.

[66] R. Nair, M. Tambe, and S. Marsella. Role allocation and reallocation in multiagent teams: Towards a practical analysis. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*, 2003.

[67] C. Nielsen, M. Goodrich, and J. Crandall. Experiments in human-robot teams. In A. Schulz and L. Parker, editors, *Multirobot Teams: From Swarms to Intelligent Automata*. Kluwer Academic, 2003.

[68] S. Okamoto, P. Scerri, and K. Sycara. Towards an understanding of the impact of personal assistants in human organizations. In *Proceedings of the Fifth International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2006)*, 2006.

[69] J. Orasanu. Shared mental models and crew performance. In *Proceedings of the 34th Annual Meeting of the Human Factors Society*, 1990.

[70] J. Orasanu and E. Salas. Team decision making in complex environments. In G. Klein, J. Orasanu, R. Calederwood, and C. Zsanbok, editors, *Decision making in action: Models and methods*. Ablex, 1993.

[71] C. Paris, E. Salas, and J. Cannon-Bowers. Teamwork in multi-person systems: a review and analysis. *Ergonomics*, 43(8), 2000.

[72] T. Payne, K. Sycara, and M. Lewis. Varying the user interaction within multi-agent systems. In *Proceedings of the Fourth International Conference on Autonomous Agents*, 2000.

[73] C. Prince, D. Baker, L. Shreshtha, and E. Salas. Situation awareness in team performance. *Human Factors*, 37, 1995.

[74] D. Pynadath and M. Tambe. An automated teamwork infrastructure for heterogeneous software agents and humans. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS): Special Issue on Infrastructure and Requirements for Building Research Grade Multi-Agent Systems*, 2002.

[75] D. Pynadath, M. Tambe, N. Chauvat, and L. Cavedone. Towards team-oriented programming. In *Agent Theories, Architectures, and Languages*. Springer, 1999.

[76] S. Rehfeld, F. Jentsch, M. Curtis, and T. Fincannon. Collaborative teamwork with unmanned ground vehicles in military missions. In *Proceedings of 11th Annual Human-Computer Interaction International Conference*, 2005.

[77] J. Rentsch. Quantifying congruence in cognition: Social relations modeling and team member schema similarity. In E. Salas and S. Fiore, editors, *Team Cognition: Understanding the Factors that Drive Process and Performance*. American Psychological Association, 2004.

[78] J. Rentsch and R. Hall. Members of great teams think alike: a model of team effectiveness and schema similarity among team members. *Advances in Interdisciplinary Studies of Work Teams*, 1, 1994.

[79] C. Rich and C. Sidner. COLLAGEN: When agents collaborate with people. In *Proceedings of the International Conference on Autonomous Agents*, 1997.

[80] C. Rich, C. Sidner, and N. Lesh. COLLAGEN: Applying collaborative discourse theory to human-computer interaction. *AI Magazine*, 22(4), 2001.

[81] J. Rickel and W. Lewis Johnson. Extending virtual humans to support team training in virtual reality. In G. Lakemeyer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann Publishers, 2002.

[82] W. Rouse, J. Cannon-Bowers, and E. Salas. The role of mental models in team performance in complex systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(6), 1992.

[83] E. Salas, C. Bowers, and E. Edens. *Improving teamwork in organizations: Applications of resource management training*. Lawrence Erlbaum Associates, 2001.

[84] E. Salas, J. Cannon-Bowers, and J. Johnston. How can you turn a team of experts into an expert team? Emerging training strategies. In C. Zsambok and G. Klein, editors, *Naturalistic decision making*. LEA, 1997.

[85] E. Salas, T. Dickinson, S. Converse, and S. Tannenbaum. Towards an understanding of team performance and training. In R. Swezey and E. Salas, editors, *Teams: Their Training and Peformance*. Norwood, 1992.

[86] E. Salas and S. Fiore, editors. *Team Cognition: Understanding the Factors that Drive Process and Performance*. American Psychological Association, 2004.

[87] P. Scerri, D. Pynadath, L. Johnson, P. Rosenbloom, N. Schurr, and M. Tambe. A prototype infrastructure for distributed robot-agent-person teams. In *Proceedings of International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2003.

[88] P. Scerri, D. Pynadath, and M. Tambe. Adjustable autonomy in real-world multiagent environments. In *Proceedings of the International Conference on Autonomous Agents*, 2001.

[89] P. Scerri, D. Pynadath, and M. Tambe. Towards adjustable autonomy for the real world. *Journal of Artificial Intelligence Research*, 17, 2002.

[90] N. Schurr, J. Marecki, P. Scerri, P. Lewis, and M. Tambe. The DEFACTO system: Training tool for incident commanders. In *Proceedings of the Seventeeth Innovative Applications of Artificial Intelligence Conference*, 2005.

[91] J. Scolaro and T. Santarelli. Cognitive modeling teamwork, taskwork, and instructional behavior in synthetic teammates. In *Proceedings of the Eleventh Conference on Computer-Generated Forces and Behavior Representation*, 2002.

[92] C. Shannon and W. Weaver. *The mathematical theory of communication*. University of Illinois Press, 1949.

[93] O. Shehory and K. Sycara. The RETSINA communicator. In *Proceedings of the Fourth International Conference on Autonomous Agents*, 2000.

[94] M. Sierhuis, J. Bradshaw, A. Acquisti, R. van Hoof, R. Jeffers, and A. Uszok. Human-agent teamwork and adjustable autonomy in practice. In *Proceedings of 7th International Symposium on Artificial Intelligence (I-SAIRAS)*, 2003.

[95] J. Smith. *Collective intelligence in computer-based collaboration*. Lawrence Erlbaum Associates, Inc., 1995.

[96] K. Smith-Jentsch, J. Johnson, and S. Payne. Measuring team-related expertise in complex environments. In J. Cannon-Bowers and E. Salas, editors, *Decision making under stress: Implications for individual and team training*. American Psychological Association, 1998.

[97] K. Smith-Jentsch, J. Johnston, and S. Payne. Team dimensional training. In J. Cannon-Bowers and E. Salas, editors, *Decision making under stress: Implications for individual and team training*. American Psychological Association, 1998.

[98] P. Squire, G. Trafton, and R. Parasuraman. Human control of multiple unmanned vehicles: Effects of interface type on execution and task switching times. In *Proceedings of the First Annual Conference on Human-Robot Interactions,* 2006.

[99] G. Stasser, D. Stewart, and G. Wittenbaum. Expert roles and information exchange during discussion: the importance of knowing who knows what. *Journal of Experimental Social Psychology*, 31(1), 1995.

[100] P. Stone and M. Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence*, 110, 1999.

[101] R. Stout, J. Cannon-Bowers, E. Salas, and D. Milanovich. Planning, shared mental models, and coordinated performance: an empirical link is established. *Human Factors*, 41, 1999.

[102] K. Stubbs, P. Hinds, and D. Wettergreen. Challenges to grounding in human-robot interaction: Sources of errors and miscommunications in remote exploration robotics. In *Proceedings of the First Annual Conference on Human-Robot Interactions,* 2006.

[103] K. Sycara, K. Decker, A. Pannu, M. Williamson, and D. Zeng. Distributed intelligent agents. *IEEE Expert Intelligent Systems and Their Applications*, 2(6), 1996.

[104] K. Sycara, M. Klusch, S. Widoff, and J. Lu. Dynamic service matchmaking among agents in open information environments. *SIGMOD Record (ACM Special Interests Group on Management of Data),* 28, 1999.

[105] K. Sycara and M. Lewis. From data to actionable knowledge and decision. In *Proceedings of the Fifth International Conference on Information Fusion*, 2002.

[106] K. Sycara and M. Lewis. Integrating agents into human teams. In E. Salas and S. Fiore, editors, *Team Cognition: Understanding the Factors that Drive Process and Performance*. American Psychological Association, 2004.

[107] K. Sycara, M. Paolucci, J. Giampapa, and M. van Velsen. The RETSINA multiagent infrastructure. *Autonomous Agents and Multi-agent Systems*, 7(1), 2003.

[108] K. Sycara, N. Sadeh-Koniecpol, S. Roth, and M. Fox. An investigation into distributed constraint-directed factory scheduling. In *Proceedings of the Sixth IEEE Conference on AI Applications*, 1990.

[109] K. Sycara and D. Zeng. Coordination of multiple intelligent software agents. *International Journal of Intelligent and Cooperative Information Systems*, 5, 1996.

[110] M. Tambe. Towards flexible teamwork. *Journal of AI Research*, 7, 1997.

[111] M. Tambe et al. Teamwork in cyberspace: Using TEAMCORE to make agents team-ready. In *Proceedings of AAAI Spring Symposium on Agents in Cyberspace*, 1999.

[112] M. Tambe and W. Zhang. Towards flexible teamwork in persistent teams. *Autonomous Agents and Multi-agent Systems*, 3(2), 2000.

[113] A. Tate. Mixed initiative interaction in O-Plan. In *Proceedings of AAAI Spring Symposium on Computational Models for Mixed Initiative Interaction*, 1997.

[114] D. Traum, J. Rickel, J. Gratch, and S. Marsella. Negotiation over tasks in hybrid human-agent teams for simulation-based training. In *Proceedings of International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2003.

[115] A. Wagner, Y. Endo, P. Ulam, and R. Arkin. Multi-robot user interface modeling. In *Proceedings of Eighth International Symposium on Distributed Autonomous Robotic Systems (DARS)*, 2006.

[116] J. Wang, M. Lewis, and P Scerri. Cooperating robots for search and rescue. In *Proceedings of International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2006.

[117] N. Warner, M. Letsky, and M. Cowen. Cognitive model of team collaboration: macrocognitive focus. In *Human Factors and Ergonomics Society (HFES) 49th Annual Meeting*, 2005.

[118] N. Warner and E. Wroblewski. The cognitive processes used in team collaboration during asynchronous, distributed decision making. In *Command and Control Research and Technology Symposium*, 2004.

[119] H. Wong and K. Sycara. A taxonomy of middle-agents for the Internet. In *Proceedings of the Fourth International Conference on Multi-Agent Systems*, 2000.

[120] D. Xu, R. Volz, M. Miller, and J. Plymale. Human-agent teamwork for distributed team training. In *Proceedings of the 15th IEEE International Conference on Tools with Artificial Intelligence*, 2003.

[121] J. Yen, X. Fan, S. Sun, T. Hanratty, and J. Dumer. Agents with shared mental models for enhancing team decision-making. *Decision Support Systems*, 2004.

[122] J. Yen, X. Fan, S. Sun, M McNeese, and D. Hall. Supporting anti-terrorist analyst teams using agents with shared RPD process. In *IEEE International Conference on Computational Intelligence for Homeland Security and Personal Safety*, 2004.

[123] J. Yen, J. Yin, T. Ioerger, M. Miller, D. Xu, and R. Volz. CAST: Collaborative agents for simulating teamwork. In *Proceedings of International Joint Conference on Artificial Intelligence*, 2001.

[124] W. Zachary and J. Mentec. Modeling and simulating cooperation and teamwork. *Military, Government, and Aerospace Simulation*, 32, 2000.

[125] Y. Zhang, L. He, K. Biggers, J. Yen, and T. Ioerger. Simulating teamwork and information flow in tactical operations centers using multi-agent systems. In *Proceedings of the Tenth Conference on Computer-Generated Forces and Behavior Representation*, 2001.