

Running Head: Common Ground

Common Ground and Coordination in Joint Activity

Gary Klein
Klein Associates Inc.

Paul J. Feltovich
Jeffrey M. Bradshaw
Institute for Human and Machine Cognition

David D. Woods
The Ohio State University
Cognitive Systems Engineering Laboratory

June 7, 2004

ABSTRACT

Generalizing the concepts of *joint activity* developed by Clark (1996), we describe key aspects of team coordination. Joint activity depends on *interpredictability* of the participants' attitudes and actions. Such interpredictability is based on *common ground*—pertinent knowledge, beliefs and assumptions that are shared among the involved parties. Joint activity assumes a *basic compact*, which is an agreement (often tacit) to facilitate coordination and prevent its breakdown. One aspect of the Basic Compact is the commitment to some degree of aligning multiple goals. A second aspect is that all parties are expected to bear their portion of the responsibility to establish and sustain common ground and to repair it as needed.

We apply our understanding of these features of joint activity to account for issues in the design of automation. Research in software and robotic agents seeks to understand and satisfy requirements for the basic aspects of joint activity. Given the widespread demand for increasing the effectiveness of team play for complex systems that work closely and collaboratively with people, observed shortfalls in these current research efforts are ripe for further exploration and study.

1. INTRODUCTION

Attending Anesthesiologist: {re-entering the operating room mid-case} “Nice and tachycardic.”

Senior Resident Anesthesiologist: “Yeah, well, better than nice and bradycardic.”

To most of us, this exchange is a mystery. We don’t have a clue about what the two anesthesiologists were saying to each other. To the anesthesiologists, this highly coded conversation carried a great deal of meaning. They relied on their common ground to clearly communicate even when using very abbreviated comments, and their comments helped them maintain their common ground.

The two anesthesiologists were working together on a neurosurgery case to clip a cerebral aneurysm. The patient had been anesthetized, but the surgeons had not yet exposed the aneurysm. The anesthesiologists had little to do during a relatively long maintenance phase, so the attending anesthesiologist used this slow period to check on another case.

Then the anomaly occurred—a fall in heart rate that is termed a bradycardia. The resident anesthesiologist quickly detected the disturbance and acted to correct it. He also paged the attending anesthesiologist.

The resident’s intervention resulted in a temporary overswing—the patient’s heart rate became too rapid, rather than too slow. In this surgery, an increased heart rate is better than a decreased heart rate.

In their discussion above, the attending anesthesiologist, who had looked at the monitors as he entered the operating room, described his understanding of the change in the patient’s status and invited some explanation. The resident, in three words, corrected the attending anesthesiologist’s understanding, enabling them to work together to diagnose the reason for the bradycardia.

Performance depends on coordination, as cognitive work is distributed among different team members. If we want to improve teamwork, we need to better understand the nature of coordination and its requirements. That is the purpose of this chapter, to build on Clark’s (1996) in-depth analysis of the coordination needed to carry on a conversation, in order to explore issues of coordination more generally.

Clark (1996) has described how the simple transaction of two people talking to each other places a surprising number of demands on coordination. Sometimes we have conversations with people we have just met; sometimes with people we have known all our lives. We have to follow what the other person is saying while preparing to insert our own thoughts and reactions. We have to manage the turn-taking in order to start speaking at an appropriate time. These handoffs in a conversation are difficult to coordinate in order to sustain a smooth flow of comments that are at least nominally linked to each other—unlike long, awkward silences or tangles where everyone tries to talk at once. Conversations are in large part undirected—the route taken depends on the topics raised and the reactions made, rather than the following of a rigorous script or the guidance of an external leader. In many ways, a conversation is a microcosm of coordination.

Clark's observations are valuable for understanding joint activity in general. Though previous accounts of team coordination (e.g., Klein, 2001; Malone & Crowston, 1994; Zalesny, Salas, & Prince, 1995) have identified features of effective coordination, Clark's description of joint activity during conversations seems to provide a much stronger basis for understanding team coordination. This chapter extends Clark's work on conversations to a whole range of other coordinated activities.

This much being said about the usefulness of Clark's work in clarifying the nature of joint activity, there is a caution that should be raised. As Spiro, Feltovich, Coulson, and Anderson (1989) have shown, any single model or any single analogy can be limiting and misleading as an account of a complex phenomenon. Thus, the dynamics of a conversation will capture some aspects of coordination, such as taking turns and developing collaborative products, but miss other aspects such as following a plan or taking direction from a leader. The analogy of a conversation misses situations marked by the need for precise synchronization. It misses the aspect of planned coordination involving a dedicated team, such as a relay team, that relies on arduous training and preparation to carry out tight sequences of actions.

In preparing this chapter, we studied the nature of joint activity in three (non-conversational) domains that require coordination. Relay races demand very tight coordination and careful planning (Klein, 2001). Driving in traffic requires coordination among strangers who will likely never meet, and who must be counted on to follow official as well as informal rules (Feltovich, Bradshaw, Jeffers, Suri, & Uszok, in press). Coaching high school football teams involves understanding and adhering to the directives of leaders (Flach & Dominguez, 2003). Throughout this chapter we use examples from these domains to elaborate on or contrast with the activity of having a conversation.

We considered several definitions of team coordination. Olson, Malone, and Smith (2001) offer the following definition: "Coordination is managing dependencies between activities." This definition includes only one of the processes involved in teamwork and organizational dynamics. One of its omissions has to do with the process of resolving issues of conflicting and interacting goals.

The definition of coordination offered by Zalesny et al. (1995) emphasizes temporal dependencies oriented around a common goal. They conceptualize coordination as "the complementary temporal sequencing (or synchronicity) of behaviors among team members in the accomplishment of their goal." (p. 102) █

Klein (2001) has stated that "Coordination is the attempt by multiple entities to act in concert in order to achieve a common goal by carrying out a script they all understand." This definition anchors the concept of coordination to a script. (p. 70) This emphasis is understandable, because coordination requires some partially shared scaffold around which to structure and synchronize activity. However, in examples such as driving on a highway or engaging in a conversation the shared scaffold is minimal. There need be no overriding common goal for a conversation, and there is almost surely none for highway drivers (except self-protection). There is no overriding script in either case, other than routines and conventions (e.g., of speech) used by a society. The participants are not working together to carry out a plan. The participants are not a formal team. They are coordinating in order to have productive or safe transactions.

The Zalesny et al. (1995) and Klein (2001) definitions address the issue of goals. However, the concept of working toward a common goal fails to capture the interdependencies of how multiple goals interact and conflict (Clancey, 2004; Cook & Woods, 1994). It misses cases where the parties do not have any common goal except for the goal of working cooperatively in order to achieve their individual goals. Coordinating parties have to relax their own dominant goals and develop a joint agreement to work together—to adjust and communicate with each other as they all benefit from participating in the joint activity despite the costs of coordination.

None of these definitions does a good job of covering all of the examples we considered. We will provide our own definition at the end of this chapter, after we have had a chance to explore the nature of joint activity further.

The concepts of joint activity and team coordination are obviously very closely related. Entering into a joint activity requires the participants to coordinate because at least some of their actions affect the goal-directed activities of others.

The next three sections examine joint activity in some detail: the *criteria* for joint activity, the *requirements* for carrying out joint activities, and the *choreography* of joint activity (see Figure 1). The criteria for joint activity are that the parties intend to work together, and that their work is interdependent, rather than performed in parallel without need for interaction. If these criteria are to be satisfied, the parties to the joint activity have to fulfill requirements such as making their actions predictable to each other, sustaining common ground (the concept of common ground will be elaborated later in the chapter), and letting themselves be directed by the actions of the other parties. The form for achieving these requirements—the choreography of the joint activity—is a series of phases that are guided by various signals and coordination devices, in an attempt to reduce the coordination costs of the interaction.

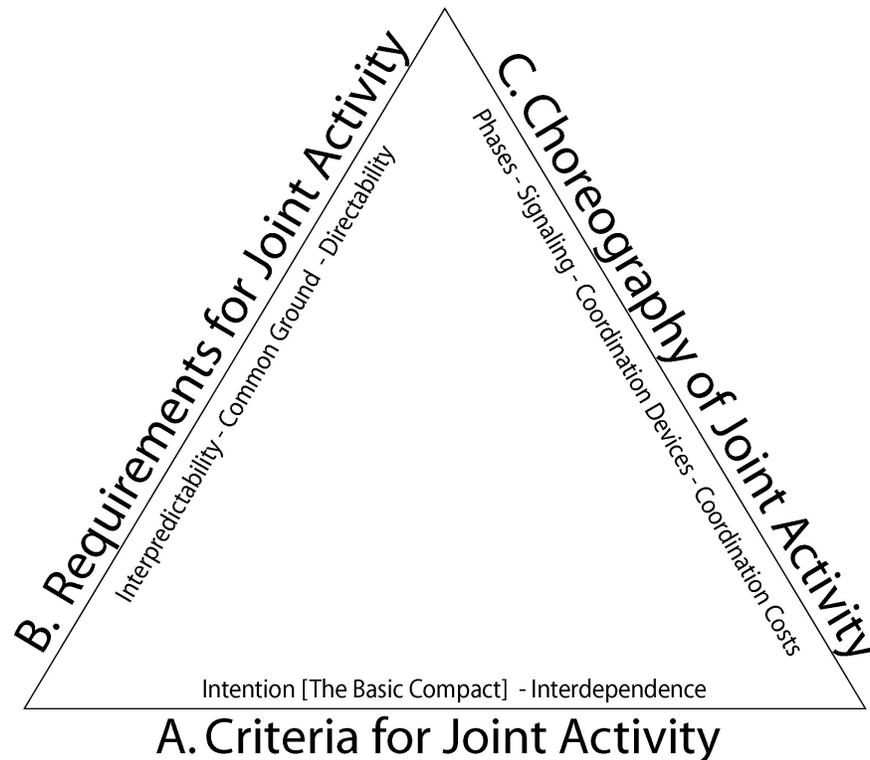


Figure 1. Description of Joint Activity.

2. CRITERIA FOR JOINT ACTIVITY

A joint activity is an extended set of behaviors that are carried out by an ensemble of people who are coordinating with each other (Clark, 1996 p. 3). The length of time of a joint activity can vary greatly. For example, the exchange between a customer and a cashier in a convenience store during a purchase may be on the order of a minute or two. Performance of a duet may take an hour or so. Two researchers writing a paper in collaboration may take months, or even years.

We have identified two primary criteria for a joint activity: the parties have to *intend* to work together, and their work has to be *interdependent*.

2.1 Intention to Generate a Multi-Party Product

“It’s not cooperation if either you do it all or I do it all.” (Woods, 2002). To be truly joint activity, the activity should be aimed at producing something that is a genuine joint project, different from what any one person could do working alone.

People engage in joint activity for many reasons: because of necessity (neither party, alone, has the required skills or resources), enrichment (while each party could accomplish the task, they believe that adding complementary points of view will create a richer product), coercion (the boss assigns a group to carry out an assignment), efficiency (the parties working together can do the job faster or with fewer resources), resilience (the different perspectives and knowledge broaden the exploration of possibilities and cross check to detect and recover from errors) or even collegiality (the team members enjoy working together).

We propose that joint activity requires a “Basic Compact” that constitutes a level of commitment for all parties to support the process of coordination. The Basic Compact is an agreement (usually tacit) to participate in the joint activity and to carry out the required coordination responsibilities. Members of a relay team enter into a Basic Compact by virtue of their being on the team; people who are angrily arguing with each other are committed to a Basic Compact as long as they want the argument to continue.

One aspect of the Basic Compact is the commitment to some degree of goal alignment—typically this entails one or more participants relaxing some shorter-term local goals in order to permit more global and long-term goals to be addressed. These longer-term goals might be shared goals (e.g., a relay team) or individual goals (e.g., drivers wanting to ensure safe journeys). A second aspect of the Basic Compact is a commitment to try to detect and correct any loss of common ground that might disrupt the joint activity.

We do not view the Basic Compact as a once-and-for-all prerequisite to be satisfied, but rather as a continuously reinforced or renewed agreement. Part of achieving coordination is investing in those things that promote the compact as well as being sensitive to and counteracting those factors that could degrade it.

The Basic Compact can vary in strength. Physically or culturally distributed teams such as international peacekeeping missions may have difficulty maintaining the Basic Compact. When that happens, the amount of coordination that is possible is reduced. Some teams will deliberately simplify their plans because they realize they are not capable of complex interactions.

All parties have to be reasonably confident that they and the others will carry out their responsibilities in the Basic Compact. In addition to repairing common ground, these responsibilities include such elements as acknowledging the receipt of signals, transmitting some construal of the meaning of the signal back to the sender, and indicating preparation for consequent acts.

The Basic Compact is also a commitment to ensure a reasonable level of interpredictability. Moreover, the Basic Compact requires that if one party intends to drop out of the joint activity, he or she must inform the other parties.

The Basic Compact is illustrated by an example from Cohen and Levesque (1991)—arranging a convoy of two cars—as a prototype for establishing common ground in a distributed team. A critical requirement

for making the convoy work is that all team members understand their obligations to each other and sign onto the Basic Compact.¹

Example 1. A Convoy as a Distributed Team.

There are two agents in this example: Alice and Bob. Bob wants to go home after a party, but doesn't know the way. He does know that Alice knows the way and that her destination is not far from his home. So all Bob has to do is to follow Alice. But that may not be enough. Bob cannot follow Alice if she drives too quickly through traffic, runs red lights, and so forth. Therefore, it is best if Alice knows Bob is going to be following her until he gets within familiar range of his home. If they use an intermediary to organize their convoy, they may be uncertain about what the intermediary arranged. So Bob and Alice need a direct affirmation that they are both signing on to drive in a convoy.

Bob and Alice may fail to affirm their commitment to drive in a convoy. For instance, if an intermediary mentions to Alice that Bob will be following her, but Bob isn't sure if Alice has agreed, then the door is open for all kinds of confusions. For example, if Bob sees Alice going down lots of side streets, he may infer that Alice doesn't really know her way and abandon the convoy, leaving Alice to worry about where she lost Bob, and circling back to try to find him. And if Bob makes a momentary wrong turn, Alice might conclude Bob now knows his way so she can abandon the convoy and leave him to find his own way.

While she is in the convoy, Alice has to drive differently, adapting her own goals (e.g., not letting her attention wander too much, slowing her speed, refusing to maneuver around knots in the traffic, refusing to take a shorter but more complicated route). The compact signals her willingness to adapt what would be her ways to achieve her goals in order to participate in the joint activity and the goals that it supports.

If one party abandons his or her intention, he or she must make this known to the rest of the team. All sorts of events can occur to make an agent believe a commitment is no longer reasonable. But the Basic Compact requires parties to communicate to each other if they plan to abandon the joint activity. If Bob and Alice had talked to each other before the party ended and agreed to ride in a convoy, then their agreement—their Basic Compact—would require each of them to signal to the other before heading off alone. The convoy example shows how difficult it is to arrange a Basic Compact by proxy.

2.2 Interdependence of the Actions of the Parties

In joint activity what party "A" does must depend in some significant way on what party "B" does and vice versa (Clark 1996, p. 18). Two musicians playing the same piece in different rooms are not involved in the joint activity of performing a duet. If two researchers decide to have one author write the first half of the paper and the other the last, with the final product merely assembled together at the end, then this is parallel—not joint—activity. A group of Army units that set out on related missions which fit within some sort of synchronization schedule but have no further interaction, are not engaged in a joint activity

¹ We have simplified the example described by Cohen and Levesque (1991).

because their efforts have no mutual influence. Joint activity emphasizes how the activities of the parties *interweave* and *interact*.

3. REQUIREMENTS FOR EFFECTIVE COORDINATION IN JOINT ACTIVITY

In reviewing different forms of coordination, we have identified three primary requirements that cut across domains: the team members have to be *interpredictable*, they have to have sufficient *common ground*, and they have to be able to *redirect* each other.

3.1 Interpredictability

Coordination depends on the ability to predict the actions of other parties with a reasonable degree of accuracy. It follows that each party also has a responsibility to make his or her actions sufficiently predictable to enable effective coordination. Predictability includes accurate estimates of many features of the situation—for example, the time needed by all of the participants to complete their actions, the skill needed, and the difficulty of the action.

Shared scripts aid interpredictability because they allow participants in joint activities to form expectations about how and when others will behave, thus enabling absence of activity to be seen as meaningful (e.g., inasmuch as the other is not doing X, then the element Y must not be present because the doctrine or script would lead someone who sees Y to do X).

Interpredictability is greatly increased when the partners can take on the perspective of the others. The heuristic of assuming that others will see the events the same way we see them, and that they will react the way we would, is reasonable in many settings. But it is not as powerful as being able to decenter and imagine the events from the other person's point of view.

3.2 Common Ground

Perhaps the most important basis for interpredictability is common ground (Clark & Brennan, 1991), which refers to the *pertinent* mutual knowledge, mutual beliefs and mutual assumptions that support interdependent actions in some joint activity. Common ground permits people to use abbreviated forms of communication and still be reasonably confident that potentially ambiguous messages and signals will be understood. Short of being able to rely on common ground to interpret such communications, every vague or ambiguous referent would have to be unpacked, at great cost to the parties in the transaction. For example, in a relay race, as runner A approaches runner B and announces “stick,” this one word has a clear meaning to runner B and requires no breath-wasting elaboration. Runner A doesn't have to tell runner B, “I am getting close to you and am just now reaching out the baton, so why don't you extend your right hand back to let me place the baton in your palm.”

The Basic Compact includes an expectation that the parties will repair faulty knowledge, beliefs, and assumptions when these are detected. Common ground is not a state of having the same knowledge, data, and goals. Rather, common ground refers to a *process* of communicating, testing, updating, tailoring, and repairing mutual understandings (cf. Brennan, 1998). Moreover, the *degree of quality* of common ground demanded by the parties can vary due to the particulars of people, circumstances, and their current objectives. Two parties in friendly chit-chat may be satisfied by simple head-nods or sustained attention to confirm that they are understanding each other well enough; two people working out a contractual agreement may demand and work toward—through more extensive mutual testing, tailoring, and repair—evidence for more precise mutual understanding (see Cahn & Brennan, 1999, p. 2 “grounding criteria”). In general, common ground is what makes joint activity and coordination work (although each joint action, in turn, serves to change common ground).

Common ground can be characterized in terms of three basic categories (Clark, 1996): *initial common ground*, *public events so far*, and the *current state of the activity*.

Initial common ground includes all the pertinent knowledge and prior history the parties bring to the joint activity. It involves not only their shared general knowledge of the world, but also all the conventions they know that are associated with their particular joint task. For an example of such conventions, think of the general procedure to be followed by a surgical team in conducting a particular kind of surgical procedure in a particular hospital setting (different hospitals can have slightly different procedures, as can particular surgical teams). It also includes what parties know about each other prior to engagement—for example, the others' background and training, habits, and ways of working.

Public events so far includes knowledge of the event history—the activity the participants have engaged in together up to the present point in the joint activity. For instance, it is important to know that options may have been lost irretrievably due to time or resource dependence (e.g., in surgery, tissue may have been cut or excised and cannot be uncut or put back). Once everyone is informed about the pertinent aspects of previous history, the participants' ongoing work together provides them with further information about each other. Also, as the course of the joint activity unfolds, it establishes precedents regarding how certain situations have been handled, how the parties have negotiated the meaning of things and events, who has been established as leader or follower, and the like. Precedents established during the process of carrying out the joint activity are important coordination devices. As we bring new people into an evolving situation we have to work out a range of issues about how they will come up to speed in the absence of their participation in the public events so far.

The *current state of the activity* also provides cues to enable prediction of subsequent actions and the formulation of appropriate forms of coordination. The physical “scene” provided by the current state serves as a kind of accumulated record of past activity, and often makes salient what is most critical in the scene for further operation. It must be remembered, however, that parties can also be fooled by such situational framing into falsely thinking they are viewing the same thing and interpreting it the same way (Koschmann, LeBaron, Goodwin, & Feltovich, 2001). In addition to the state of the literal work environment at a particular point in the activity, “props” of various kinds are often available to portray the current state—for example, the pieces placed on the chessboard. Sarter, Woods, and Billings (1997) and Christoffersen and Woods (2002) have discussed the importance of “observability” for effective coordination, and we see the various forms of observability as bases for solidifying common ground.

In reviewing various forms of team coordination, we found that some of the most important types of mutual knowledge, beliefs, and assumptions are about:

- the roles and functions of each participant;
- the routines that the team is capable of executing;
- the skills and competencies of each participant;
- the goals of the participants, including their commitment to the success of the team activity; and
- the “stance” of each participant (e.g., his or her perception of time pressure, level of fatigue, and competing priorities).

Although common ground depends greatly on these kinds of shared knowledge, we obviously do not mean to imply that all parties in joint activity need to know and think identically. Parties in joint activity represent different points-of-view by way of their particular roles or duties in the activity, or because they

represent distinct backgrounds, training or skills. It is likely that understanding and aligning these different perspectives can at times require added effort, for instance, in the need for more “repair” in the process of communication. But such diversity of perspective may actually improve performance by requiring the participants to negotiate and reconcile different pertinent understandings (e.g., Feltovich et al., 1989; Spiro, Feltovich, Coulson, & Feltovich, 1996).

We can differentiate several activities in which teams often engage in order to support common ground. These are:

- 1) Structuring the *preparations* in order to establish initial calibration of content, and to establish routines for use during execution.
- 2) *Sustaining* common ground by inserting various clarifications and reminders, whether just to be sure of something or to give team members a chance to challenge assumptions.
- 3) *Updating* others about changes that occurred outside their view or when they were otherwise engaged.
- 4) *Monitoring* the other team members to gauge whether common ground is being seriously compromised and is breaking down.
- 5) Detecting *anomalies* signaling a potential loss of common ground.
- 6) *Repairing* the loss of common ground.

Common ground can vary in quality over time and in different situations but can never be perfect. That is one of the functions of the sustaining and monitoring functions—to catch the discrepancies before they become serious. Common ground enables a team to coordinate in order to achieve tasks, but it is not a work output in itself.

3.3 Directability

Christoffersen and Woods (2002) have identified “directability” as an important aspect of coordination, because it is so central to the resilience of a team. We also see it as central to the interdependence of actions. If the way that one partner performs a task has *no* effect on the other partner, then they are stove-piped, perhaps working in the same space but not coordinating with each other. Sometimes, one partner runs into difficulty, and the other partner notices this and adapts. At other

In relay races, despite prior planning, the runners must also monitor progress during the race itself and adjust accordingly. Runner B should study the progress of runner A, changes in the situation (e.g., a competitor moving up rapidly), and the “stance” of

In relay races, baton passes are tricky and are often the key to success or failure. The exchange must happen within a limited stretch of the track. The current runner and the new (handoff) runner try to make the exchange with both of them going at as near full speed as possible. Hence the waiting runner starts to accelerate before the baton pass so that he or she can be in full motion at the time of the handoff. Ragged hand-offs waste precious fractions of a second. But the precise timing of the handoff cannot be directed by the coach of the team. A coach’s attempts to tell the runners exactly when to make the exchange would add confusion.

times, one partner will find some way to signal to the other to alter course, and to modify the routine they are carrying out. Directability refers to deliberate attempts to modify the actions of the other partners as conditions and priorities change.

Taking direction and giving direction are another whole aspect of the interplay in coordination when there are levels of management and control (Shattuck & Woods, 2000).

4. THE CHOREOGRAPHY OF JOINT ACTIVITY

Each small phase of coordination is a joint action, and the overall composite of these is the joint activity (Clark, 1996, pp. 3, 30-35, 125). For example, a couple on a dance floor is engaged in a joint activity (the entire dance) and in joint actions (each sequence within the dance). Two waltz dancers must recognize in an ongoing manner the changes of body pressure and posture of the other that indicate the direction of the next movement and they must respond appropriately. Failure to recognize the appropriate cues and, just as importantly, failure to respond appropriately in the necessary time frame (in this case nearly instantaneously) results in at best bad dancing, and at worst something one might not characterize as “dancing” at all.

The choreography of a joint activity centers on the *phases* of the activity. The choreography is also influenced by the opportunities the parties have to *signal* to each other and to use *coordination devices*. *Coordination costs* refer to the burden on joint action participants that is due to choreographing their efforts.

4.1 Three-Part Phases

According to Clark (1996), what actually gets coordinated is a phase. A phase is a joint action with an entry, a body of action, and an exit (although there may also be other embedded joint actions within any joint action). Coordination is accomplished one phase at a time in a joint activity. Therefore, the most basic representation of the choreography of joint action is the entry, actions and exits for each phase, with coordination encompassing each of these. Often, the phases do not have any official demarcations—the parties themselves determine which part of the phase they are in, based on the work they have accomplished together.

The exiting of a phase can be difficult to coordinate—each person performing an action needs evidence that he or she has succeeded in performing it (e.g., when you press an elevator button, you want to see a light go on). Similarly, joint closure is needed for coordinated actions. For example, during a conversation the passive team members have a role in signaling to generate the belief in the one talking that the action was understood. This signaling can be done through head nods, grimaces, paraphrases, and so forth.

Thus, communication proceeds on two tracks: the official business can be thought of as the task work (Salas & Cannon-Bowers, 2001), while the conformations, corrections and completions can be thought of as the teamwork or the communication overhead. This overhead constitutes the coordination costs of the interaction. Note that the distinction can sometimes blur—for example, what is the task work versus the teamwork in dancing a waltz?

Both conversations and relay races can be described in terms of joint action phases. In a relay race, one phase is the handoff of the baton. That is when the teammates have to coordinate with each other. In a conversation, one phase is the handoff of turn-taking, as one person yields the floor to the other. That is when the participants engage in joint action to effect a smooth transition.

The overall structure of a joint activity (e.g., “writing a collaborative research article”) is one of embedded sets of a series of phases throughout many possible layers of joint activity. The “writing of the paper” is itself a joint activity, which will have within it sub-joint activities, e.g., “creating the outline,” which will themselves have embedded joint activities, e.g., “collecting the major resources to be used,” and so on, all the way down. Synchronizing entry and exit points of the many embedded phases involved in complex joint activity is a major challenge.

The phases of a joint activity can be more or less “scripted,” that is, more or less prescribed in how the activity is to be conducted. Precise scripting may be mandated or may emerge because of law, policy, norms of practice, and the like (Clark, 1996, p. 30). An example of tight scripting is the handling of evidence in a crime lab. A somewhat less scripted example is the carrying out of a particular kind of surgical procedure by a surgical team. Even less scripted is the work of a design team working to produce a new company product or logo.

The entire joint activity can be more or less scripted—as can be each of the phases. In addition to regulatory coordination mechanisms that are officially mandated, various macro guides can serve to coordinate the course of entire joint activities. Examples are plans for an activity worked out in advance by the participants, and the prior extensive outline worked out by authors involved in writing a joint manuscript. Motivations for standardizing procedures include the desire to prevent interactions where some earlier move impedes some necessary later move, and also better anticipation of necessary coordination points (e.g., Shalin, Geddes, Bertram, Szczepkowski, & DuBois, 1997). However, no matter the extensive nature or degree of foresight incorporated into any prior “plan,” there will always be a need for adjustment, further specification, or even drastic revision (Clancey, 1997).

We observe that conversations are not as scripted or planned as relay races. The direction that conversations take is continually being reshaped, depending on what has been said earlier, what has just been said, and the motivations of the parties for what they want to accomplish through the conversation (e.g., to get to the bottom of a deep issue, or, at a social gathering, to get rid of an unwanted conversational partner).

4.2 Signaling

The choreography of joint activity depends on the way the participants signal to each other about transitions within and between phases. The participants may also signal such things as their intentions, the difficulties they are having, and their desires to redirect the way they are performing the task. The importance of signaling can be seen in the way drivers coordinate with each other on the highway.

Feltovich et al. (in press) have identified the task of driving in traffic as a challenge of coordination. One of the distinctive properties of driving in traffic is that it depends on coordination among people who have never met each other and in no sense form a team. Yet if their actions are not skillfully coordinated, the result can be chaotic and even fatal. We rely on the coordination of strangers as we travel at high speeds on crowded expressways, and when we make split-second decisions about when to make a left turn against oncoming traffic. When we look away from the road if no immediate hazard is apparent, we trust the driver in front of us not to unexpectedly come to a halt when we aren’t watching.

Driving in traffic captures a new item of coordination: the drivers cannot rely on personally established precedent for coordination. Hence, there is a great deal of local control. Except in extraordinary circumstances, nobody is directing traffic, telling the drivers what to do. Yet, interpredictability is crucial if hazards of the road are to be avoided.

Driving in traffic depends heavily on signaling (see also Norman, 1992). The individual drivers are constantly providing cues for coordination, the most salient being signals to indicate intention to turn or to change lanes, or (through the brake lights) the act of slowing down. If drivers see a potential problem ahead, they may pump their brakes in advance of slowing down in order to alert the car behind them and reduce the chance of being rear-ended.

Driving in traffic depends on a Basic Compact that affirms that everyone will follow the rules. However, there may be defensible motives for drivers to escape from the Basic Compact, as when they are responding to an emergency by rushing someone to the nearest hospital. At such times, drivers may turn on their emergency blinkers to signal to other drivers that their actions are no longer as predictable and that they temporarily consider themselves not bound by certain aspects of the Basic Compact. Automobiles are all equipped with emergency blinkers to permit drivers to signal this kind of state during unusual events.

Signaling only works if the other participants notice the signals. Clark (1996) explains that coordinated conversation depends on the way the participants manage attention—their own attention and the attention of others. Generally speaking, we can only give full conscious attention to one thing at a time. We can redirect attention very quickly, sometimes too quickly. Our efforts to attend to a particular strand of a situation is easily subverted by distractions. Thus, part of the coordination effort in a conversation is to coordinate attention, helping the participants gain, hold, and relinquish attention. The coordination of attention is done via such things as gazes, pointing, and instructions—and by making repairs when coordination goes wrong. Appropriate direction of attention is such a challenge that people often expect interference, disruptions, and interruptions. Then they try to work around these when they occur.

Thus, part of the choreography of joint action is to direct the attention of other participants to signals, or to use signals to direct their attention to relevant cues in the situation (Moore & Dunham, 1995; Woods, 1995). If one partner believes a coordination phase has been completed and wants to exit that phase, it may be necessary to direct the other partner's attention to the cue that marks the phase completion.

Signaling carries a responsibility to judge the interruptability of the other partners, as well as to understand what signaling modality would be the most appropriate for the situation. Careless signaling can degrade performance. Dismukes, Young, and Sumwalt (1998) have described how interruptions have contributed to a number of aviation accidents as when, for example, air traffic controllers request information from pilots who are going through a flight checklist, leading to the omission of critical steps. The very use of a checklist is a coordination device intended to protect pilots from factors such as interruptions because of how the checklist can structure activities in ways that are visible to others.

4.3 Coordination Devices

The choreography of joint activity is shaped by the use of coordination devices (Clark, 1996, pp. 64-66). These devices include highly diverse mechanisms of signaling. Many of these signals, coupled with common ground, serve to increase interpredictability among the parties. Examples of typical coordination devices are agreement, convention, precedent, and situational salience.

Agreement: Coordinating parties can explicitly communicate their intentions and work out elements of coordination. This category includes, in addition to language, diverse other forms of signaling that have shared meaning for the participants, including signs, gestures, and displays.

Convention: Often prescriptions of various types and degrees of authority apply to how parties interact. These can range from rules and regulations to less formal codes of

appropriate conduct. These less formal codes include norms of practice in a particular professional community as well as established practices in a workplace. Coordination by convention depends on structures outside of a particular episode of joint activity.

Precedent: Coordination by precedent is like coordination by convention, except that it applies to norms and expectations developed within the ongoing experience of the joint activity. As the process unfolds, decisions are made about the naming and interpretation of things, standards of acceptable behavior and quality (e.g., what is established by this particular surgical team, during the course of a surgical procedure, as the standard for adequate cauterization of a vessel), who on the team tends to take the lead, and so forth. As these arise and develop during the course of the activity, they tend to be adopted as devices (or norms) of coordination for the remainder of the activity.

Salience: Salience has to do with how the ongoing work arranges the workspace so that next move becomes apparent within the many moves that could conceivably be chosen. During surgery, for example, exposure of a certain element of anatomy in the course of pursuing a particular surgical goal can make it clear to all parties involved what to do next. Coordination by salience is produced by the very conduct of the joint activity itself. It requires little overt communication and is likely to be the predominant mode of coordination among long-standing, highly practiced teams.

Let us look at an example, one in which the authors, having recently communicated only via e-mail regarding their paper, decide that this has been insufficient for resolving some issues and that they need to talk on the phone (see also Brennan, 1998 for a complementary account of a set of e-mail exchanges). They try to arrange this phone conversation by e-mail:

F: Hi, K, our e-mails have left a lot of questions and confusion in my mind. I really think we need to talk if we are going to get some of this settled.

K: Yes, F, I've been feeling the same way. I'm free all tomorrow afternoon. How's that for you?

F: Any time between two and three is okay for me, how about two o'clock?

K: Let's book two o'clock! Want me to call you or do you want to call?

F: I'll do it. Catch you at two tomorrow. By the way, what is your cell number in case I don't reach you in the office?

There are several coordination devices involved in this exchange. There is a convention that one responds to the overtures of friends and active collaborating colleagues with a relatively prompt return communication. There is precedent at work: such exchanges on the paper have been held from the location of the authors' offices in the past, perhaps because most of the pertinent materials for these conversations reside there. Although the offices are the default origination places by precedent—F feels the need for a back-up if the usual does not hold true, essentially saying, "If you are not in your office when I call, give me your cell phone number to call so we can re-coordinate." Some aspects are communicated simply by their salience in the field of joint activity of the authors. For instance, the "paper" itself is never mentioned. The authors may have communicated by e-mail about all sorts of things: friends and family, an academic meeting they will be attending together in a month, a workshop they are designing on some topic. Yet, somehow, all the prior interaction between the parties has conspired to make the "paper" the salient focus of this exchange in time and space.

4.4 The Joint Action Ladder

Clark uses the concept of a joint action ladder (1996, p. 152) to describe the successive inferences that we make during a conversational interaction (Figure 2).

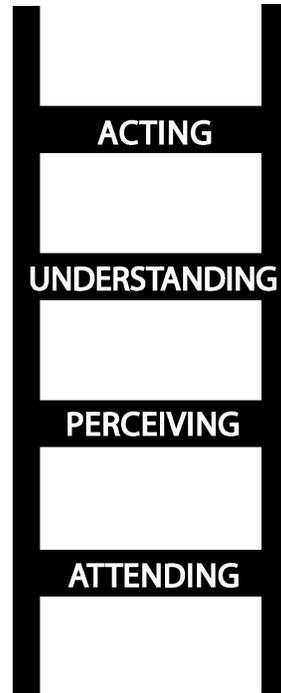


Figure 2. Joint Action Ladder.

When someone sends a message to a team member, as in the earlier example of the e-mail that F sent to K, four things have to happen for the exchange to be successful in advancing the overarching joint activity—in this case writing the collaborative academic paper. First, when F sends the note to K, K must *attend* to it. This step seems trivial, but it is not. A virus may have temporarily disabled K’s e-mail system, leading F to erroneously assume that K has received the message. Second, while K may attend to the note (i.e., he sees that it has arrived), he must also *perceive* its contents. He may fail to do so. For example, he may be in a rush, have hundreds of new email items, open F’s message to see if F is writing about the paper or about some other topic, and never read the detail. Third, while the recipient may attend to and perceive the note, he must also *understand* it. Fourth, the recipient must *act* on the message, by doing a number of additional things, among them acknowledging that he has “climbed the action ladder” or indicating how far he has been able to climb it (e.g., “I got your message with the attachment but for some reason I can’t open the attachment”). In the best case, he states what he will do next as a result of receiving the note, which indicates that he has attended to it, perceived it, and understood it.

Certain aspects of the ladder may not need to be addressed directly because of the principle of “downward evidence” (Clark, 1996, p. 148): lower rungs on the ladder are subsumed by higher ones. For example, the initial response by K in our example interchange (“Yes, F, I’ve been feeling the same way. I’m free all tomorrow afternoon. How’s that for you?”) lets F know that K has attended to and perceived F’s message, and also suggests that K understands the note and intends to act on it to continue the joint work.

Although the interchange in our example of joint activity involved language, this need not always be the case. Anything that carries meaning for the participants—sign, signal, gesture, posture, actions, and so

forth—can be used to coordinate joint activity. Indeed, even vertebrate animals participate in simple forms of joint activity, using various forms of signaling and display (Feltovich et al., in press; Smith, 1977, 1995).

The choreography of joint activity does not just cover instances in which the transmission of the joint action signal is more or less direct, face-to-face. For example, the signal can be delayed, mediated by a machine, or mediated by another person. Also, the “cycle time” of any joint action can vary greatly, from the almost instantaneous interactions of two waltz partners to the days or weeks of two authors creating a research paper together. There may be a defined leader or conductor of the joint activity (who decides, for instance, where to move next—commonly the role of the “attending” surgeon during a surgical procedure) or there might not be.

4.5 Coordination Costs in Choreography

The effort expended to choreograph joint activity is one type of coordination cost. Fundamental to coordination is the willingness to do additional work and to narrowly downplay one’s immediate goals in order to contribute to the joint activity. Consider the example of a relay race. Each handoff takes its toll in terms of time lost. One typical arrangement for a relay race is to use four runners to cover 400 meters. It would be inefficient to have 24 runners stationed around a 400-meter track. The coordination costs would outweigh advantages of the freshness and energy of each new runner.

Schaeffer (1997) and Klinger and Klein (1999) have identified types of coordination costs that are incurred by joint activity: synchronization overhead (time wasted in waiting for one entity to complete its work before the next one can begin); communication overhead (effort to manage the handoff); redirection overhead (wasted time and energy in going in the wrong direction after a new direction is called out but before all entities can be told to change course); and diagnosis overhead (the additional burden of diagnosing a performance problem when multiple moving parts are involved).

Consider also how unwieldy and frustrating a conversation with many participants can be. The coordination costs of choreographing the phases, the time wasted in trying to sort out how to enter a comment into the flow of discussion—each of these increase as the number of participants becomes greater. Sometimes the wait time is so long that when people do get to speak they forget what they wanted to say.

As common ground is lost, coordination costs can rise. Effort spent in improving common ground will ultimately allow more efficient communication and less extensive signaling. On the other hand, failure to invest in adequate signaling and coordination devices, in an attempt to cut coordination costs, increases the likelihood of a breakdown in the joint activity. The next section describes some typical ways that such breakdowns occur.

5. THE FUNDAMENTAL COMMON GROUND BREAKDOWN

Many different factors can lead to breakdowns in joint activity. In this section, we discuss one type of breakdown that seems to arise repeatedly in the instances we examined. We call this the Fundamental Common Ground Breakdown. We also discuss a form of this breakdown that is particularly troublesome—when one party defects from the joint activity and leaves the other parties believing that the Basic Compact is still in force.

5.1 The Logic of the Fundamental Common Ground Breakdown

We assert that, no matter how much care is taken, breakdowns in common ground are inevitable: no amount of procedure or documentation can totally prevent them. During a transaction, while common ground is being established in some ways (e.g., through the record and activities of a conversation or through physical markers), it is simultaneously being lost in others (e.g., as each party has differential access to information and differential interpretation of that information). That is why Weick, Sutcliffe, and Obstfeld (1999) have asserted that high reliability organizations and work processes are marked by a continual mindfulness—a continual searching for early indicators of problems, including indications of a loss of common ground.

Why do teams lose common ground? There are a number of typical reasons for the loss of common ground. These have been identified by Klein, Armstrong, Woods, Gokulachandra, and Klein (2000). In the incidents Klein et al. studied, they found that common ground is continually eroding and requiring repair. As a situation changes, people are likely to make different interpretations about what is happening, and what the others know. So the baseline state is one in which people are detecting and repairing problems with common ground—and not attempting to document all assumptions and duplicate the contents of each person's mind.

Our research has shown that teams tend to lose common ground for the following reasons:

- the team members may lack experience in working together;
- they may have access to different data;
- they may not have a clear rationale for the directives presented by the leader;
- they may be ignorant of differences in stance (e.g., some may have higher workload and competing priorities);
- they may experience an unexpected loss of communications or lack the skill at repairing this disruption;
- they may fail to monitor confirmation of messages and get confused over who knows what.

This last reason, confusion over who knows what, is sufficiently frequent that it deserves more careful examination. We refer to this confusion as the Fundamental Common Ground Breakdown.

The script for the Fundamental Common Ground Breakdown is illustrated by the following example:

Example 2. The Viennese Waltz.

Gary gave a presentation on decision making to a conference in Scotland. One of the attendees requested that Gary give the same presentation in Vienna a few months later. Gary had a schedule conflict but arranged for his close colleague, Dave, to give the presentation in his stead. Dave was willing to make the trip to Vienna, but did not have the time to put together a presentation. Gary assured him that this was no problem—all he had to do was give the same presentation that Gary used in Scotland.

Gary asked his production associate, Veronica, to give Dave the Scotland presentation. She did so.

The day before Dave was going to leave for Vienna, Gary stopped in to talk, and noted that Dave was feeling pressured from the workload because he still had to pull together materials. This surprised Gary.

He verified that Dave had gotten the presentation from Veronica. Dave said that there were still some big gaps to fill.

Gary asked to see the presentation, and Dave opened it up on his computer and they went through it, slide by slide. Gary then asked where the exercise materials were, and Dave said he didn't know anything about exercise materials—that was what he had to pull together.

Gary went back to Veronica, and they discovered that Veronica had only sent Dave a file containing the presentation. She didn't know she was also supposed to send a text file containing the handouts with the exercise.

In this incident, Gary assumed that Dave had gotten everything he needed because Gary had requested that from Veronica, Veronica had complied, and Dave had received the file Veronica had sent. The disruption occurred because the word “presentation” meant one thing to Gary (i.e., everything he used during his talk in Scotland) and another to Veronica (i.e., the file containing the presentation—a text file is not a presentation). As a result, Dave almost went off to Vienna without the critical materials needed to run a decision making exercise.

We present this example to illustrate the fundamental common ground breakdown:

- Party A believes that Party B possesses some knowledge.
- Party B doesn't have this knowledge, and doesn't know he is supposed to have it.
- Therefore, he or she doesn't request it.
- This lack of a request confirms to Party A that Party B has the knowledge.

As a result, they fail to catch the mismatch in their beliefs about what Party B knows and is supposed to know.

Further, Party A interprets Party B's subsequent statements and comments with the assumption that Party B possesses the critical knowledge, thereby constructing a cascading set of incorrect inferences about Party B's beliefs and actions.

Too often, people discover a serious loss of common ground when they experience a coordination surprise. When something happens that doesn't make sense in terms of their beliefs, that event may trigger a deeper inquiry. Sometimes, if the parties are sensitive to each other, they will catch the error early enough to repair it, as in the Viennese Waltz example. Too often the discovery is made following an accident or performance failure (For a discussion of coordination surprises, see Patterson, Woods, Sarter, & Watts-Perotti, 1998; Sarter, Woods, & Billings, 1997.).

The Fundamental Common Ground Breakdown is depicted in Figure 3. Person A assumes that person B knows item “X,” whereas person B doesn't have this information. As subsequent events occur, person A interprets B's beliefs and actions in a way that diverges from person B's actual beliefs and rationale for actions. The discrepancy grows greater until they encounter a coordination surprise.

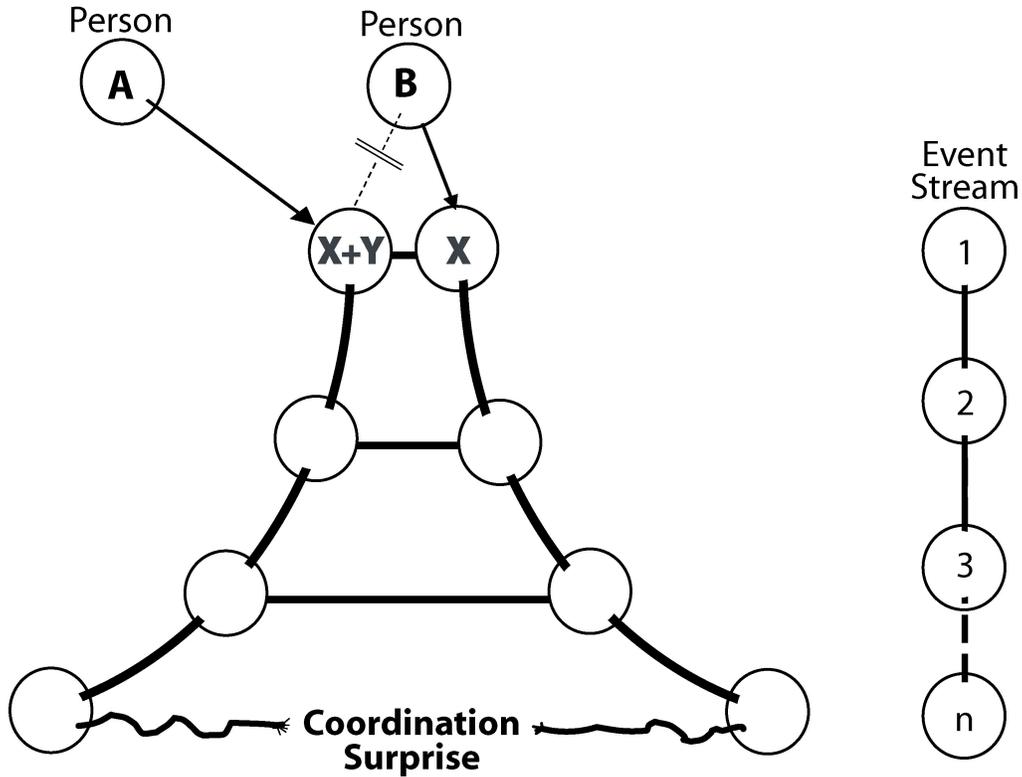


Figure 3. Fundamental Common Ground Breakdown.

This account is different from a case where the parties just don't have a shared understanding of the situation. The breakdown is more than the lack of a shared understanding. The breakdown arises because person A is making incorrect assumptions about what person B knows, and also because of the ripple effects as person A has to explain away the initial anomalies and prediction failures by coming up with more and more elaborate accounts which just end up adding to the distortions.

Careful inspection will show that the Fundamental Common Ground Breakdown is actually a corruption in the execution of the normal Joint Action Ladder sequence. The breakdown occurs when people wrongly assume that the ladder has been scaled. The joint action ladder has four rungs: attending to a signal, perceiving its contents, understanding it, and acting on it. While downward inferences are usually benign (if people act on a message, they have probably read it), upward inferences are very risky (e.g., if a person sends a message then he or she cannot safely assume the recipient has read it).

In the Viennese Waltz, Gary assumed that Dave had gotten and appreciated all the materials because Gary asked Veronica to send them, and Dave said he received them. Unfortunately, Veronica failed to understand what Gary had wanted to request (because of Gary's carelessness in making the request), and therefore failed to act in the way Gary expected. Dave's confirmation that he had received materials from Veronica reinforced Gary's mistaken belief that the joint action ladder had been scaled successfully, and that Veronica had understood his intentions.

In the Fundamental Common Ground Breakdown, activities lower in the ladder are used to hypothesize that the ladder was scaled. We do not, and cannot, take the effort to ensure that every message has been read and properly understood. We often have to make these assumptions, just to get anything done. Then we have to be on the lookout for evidence that the assumptions were unwarranted, although when we hit these coordination surprises we usually think "Is he stupid, or crazy?"

However there are practices that can be engaged in proactively in order to make it easier to detect a breakdown before a possibly harmful surprise. One of these practices is to elaborate—to the extent possible—all the steps of the “joint action ladder” during coordination events. The recipient of a message can report back to the sender some indication of the recipient’s understanding of the message (in addition to the fact that it has been attended to and perceived). In this case Veronica may have said to Gary’s original request: “I got your message about helping Dave, and as a result I have sent him the presentation file named “Scotland,” and dated...” This may have prompted Gary to ask about the auxiliary materials for the talk. Even so, we are dubious about the general effectiveness of a reporting back procedure. Veronica would likely only report back in this manner if she was unsure about whether to send the text file—and in that case, she would simply have asked Gary whether she should do so. The breakdown in The Viennese Waltz could best have been avoided if Gary had explained his overarching intent to Veronica—that he wanted Dave to carry out the same program in Vienna that Gary had accomplished in Scotland.

Here are some additional examples of the Fundamental Common Ground Breakdown, taken from Klein et al. (2000):

Example 3. The Fuel Leak.

In a flight simulation to study the decision making of commercial pilots, researchers introduced a malfunction into an otherwise standard scenario: a leak from the fuel tank in one of the wings. This malfunction was patterned after an actual event. Nevertheless, it is an unlikely malfunction, and the aircrew doubted the problem after they spotted it. They thought it was likely that the fuel gauge was malfunctioning, and they spent some time trying to determine if the leak was really happening.

The captain realized that if there were a real leak, the affected wing would be getting lighter, changing the handling characteristics of the craft. He turned to the first officer (who was flying the plane) and said, “Are you having any trouble?” The first officer answered, “No.” That seemed to argue for a gauge problem, and it took another ten minutes to establish that the fuel was really being lost.

Afterwards, during the debriefing, the crew was shown the videotape of their attempts to diagnose the problem. Upon seeing the first officer say “No,” the captain and flight engineer seemed to find the culprit. But the first officer asked for the tape to be rewound. He argued that the question he was asked was problematic. Yes, the airplane was handling poorly. He was so busy trying to keep it straight and level that he lost track of their discussion, while assuming that the captain and flight engineer knew about the leak. So, when asked if he was having trouble, he answered honestly that he wasn’t. He was able to keep the plane leveled out. He should have been asked if the plane was imbalanced.

Here is another example of the Fundamental Common Ground Breakdown, involving air traffic controllers. In some ways, this example is like a relay race, with the aircraft serving as the baton passed between air traffic controllers. Unhappily, this baton was dropped.

Example 4. Handing off an Airplane.

In aviation, airline dispatchers and flight crews work with air traffic controllers (ATC) to adapt flight plans in the face of constantly changing conditions. Weather can rapidly and unpredictably turn

unpleasant, creating obstacles in the planned flight path of an airplane, and air traffic densities can suddenly ebb or flow, changing the assumptions of the plan. Hence, plans must have the ability to be modified to accommodate these changes.

An example of a potentially significant breakdown in communication occurred during a flight from Dallas/Ft. Worth to Miami (Smith et al., 1998). The airline dispatcher in charge of the flight filed an initial flight plan. However, as part of his further responsibility to the flight plan, the dispatcher was also required to warn the pilot of any hazardous weather enroute. In this case, after the flight was on its way, the dispatcher noticed a line of thunderstorms that he felt could potentially endanger the safety of the flight and, with the Captain's concurrence, issued a reroute instruction to the aircraft. During this process, the Captain was informed about the situation and the conditions of weather that prompted the reroute. The reroute was coordinated with the ATC and approved.

The dispatcher in charge of the flight, who originally noticed the bad weather, was also responsible for about 30 other planes. Once he had filed the reroute, briefed the Captain, and obtained permission from the first ATC center, he felt that the issue had been resolved and devoted his attention to other airplanes and tasks in his charge.

Once the flight was underway, however, the Receiving Center, (i.e., the ATC that was taking on responsibility for the flight as the aircraft transitioned from one ATC sector to another) rejected the reroute and put the airplane back on its originally filed flight plan. The Captain assented, assuming that the receiving ATC also knew about the weather front.

However, following the original flight plan trapped the aircraft south of the line of thunderstorms. The weather system forced the aircraft to circle, waiting for a break. As the aircraft began to run low on fuel, the crew decided to fly through the weather front to land in Miami (the only usable airfield in the area at the time).

In this case there was a breakdown in common ground between most of the dyads involved. The Captain and the second ATC had trouble modifying the flight plan due to a disrupting event. The Captain knew about the disrupting event, assumed that the ATC did as well, and thought that the ATC had complete, detailed weather information about the relevant portion of the airspace. The Captain assumed that the joint action ladder had been scaled in his interaction with the second ATC, but that ATC did not understand the situation in the way the Captain expected. As a consequence, neither the Captain nor the ATC took the appropriate action.

The second ATC rejected the modified plan and did not take the weather conditions into account. The ATC had a different set of tasks, a different focus of attention (north-south routes on the east coast of Florida), and did not have complete data about the weather situation in the region the pilot was concerned about. Neither party realized that they were on different wavelengths and neither made any inquiries to verify that their assumptions about the other's knowledge, goals, and activities matched their own.

The third party to the plan modification, the dispatcher, had reached closure on this issue and did not realize that it had been reopened. Thus, the dispatcher was unable to participate by elaborating the intent behind the original modification.

Upon modifying the plan following the disruptive event, each of the three parties communicated with others and assumed that all were operating on a common assessment of the situation, that all had examined the issue at hand, and that all were working to the same subset of goals. Yet each of the parties

misunderstood the others' perspectives, miscommunicated, and miscoordinated, leaving the flight crew to deal with the obstacle of the weather system—an obstacle that the replanning process had anticipated and tried to avoid.

5.2 Defections from the Basic Compact

When the partners in a joint activity agree to work together, they usually are not making a legally binding contract of infinite duration. At any point they may find that they have to abandon the activity, and often this is reasonable. Breakdowns occur when the abandonment is not signaled, with one set of partners continuing to believe that the Basic Compact is still in force while the party who has abandoned the joint activity believes it has become void.

Sometimes the partners will explicitly discuss the fact that they are abandoning the joint activity. But in most cases, the agreement itself was tacit and the partners depend on more subtle types of signaling to convey that they are no longer going to participate in the joint activity. When this signaling fails, the result is a form of Fundamental Common Ground Breakdown.

In the relay race, whether or not participants are remaining in the Basic Compact is visible and fairly unambiguous. The fact that a runner has abandoned the Basic Compact would most often be fairly obvious when, for example, he does not accept the baton at handoff.

Remaining in compact during a conversation is also highly visible in the process of accepting turns, relating understandings, detecting the need for and engaging in repair, displaying a posture of interest, and the like. When these sorts of things are not happening, it may be inferred that one or more of the parties is not wholeheartedly engaged.

On the highway, evidence of breakdown in the compact is often cued by deviations from norms or rules related to the safety of all, such as tailgating, cutting in too fast after passing, or not providing appropriate signals. In this sense, breakdown of the compact on the highway is perhaps easier to verify than its maintenance—the latter being consistent with “nothing unusual.”

In football games, there are so many hidden coordination activities (e.g., messages and signals communicated across telephones, or spoken or displayed among players in the throes of the game) embedded in complex activity that knowing the state of the basic compact across the many parties can be quite difficult. The tighter the coupling—the degree to which one's action is dependent on the other—the more likely that lapses in the Basic Compact will show negative results and be noticed. For example, contrast the degree of interdependence between a quarterback and center to that between the two offensive tackles.

We have finished our discussion of the primary factors that affect joint activity. These factors include both support for coordination and the reasons for a loss of common ground and resultant breakdowns in coordination. In the next section we explore the implications of these factors for the design of automation.

6. MAKING AUTOMATION A TEAM PLAYER

The concept of automation—which began with the straightforward objective of replacing whenever feasible any task currently performed by a human with a machine that could do the same task better, faster, or cheaper—became one of the first issues to attract the notice of early human factors researchers. These researchers attempted to systematically characterize the general strengths and weaknesses of humans and machines (Fitts, 1951). The resulting discipline of *function allocation* aimed to provide a

rational means of determining which system-level functions should be carried out by humans and which by machines.

Over time it became plain to researchers that things were not that simple. For example, many functions in complex systems are shared by humans and machines; hence the need to consider synergies and conflicts among the various performers of joint actions (Hoffman, Klein, & Laughery, 2002). Also, the suitability of a particular human or machine to take on a particular task may vary by time and over different situations; hence the need for methods of function allocation that are dynamic and adaptive (Hancock & Scallen, 1998). Moreover, it has become clear that function allocation is not a simple process of transferring responsibilities from one component to another (Boy, 1988). Automated assistance of whatever kind does not simply enhance our ability to perform the task: it changes the nature of the task itself (Christoffersen & Woods, 2002; Feltovich, Hoffman, Woods, & Roesler, 2004; Norman, 1992). Those who have had a five-year-old child help them by doing the dishes know this to be true—from the point of view of an adult, such help does not necessarily diminish the effort involved, it merely effects a transformation of the work from the physical action of washing the dishes to the cognitive task of monitoring the progress (and regress) of the child.

The ultimate desire of researchers and developers is to make automation a team player (Christoffersen & Woods, 2002; Malin et al., 1991). A great deal of the current work to determine how to build automated systems with sophisticated team player qualities is taking place within the software and robotic agent research communities, albeit in many forms and with somewhat divergent perspectives (Allen, 1999; Allen et al., 2000; 2001a, b; 2002; Bradshaw et al., 1997; 2003; 2004a, b, c; in preparation; Christoffersen & Woods, 2002; Clancey, 2004; Cohen & Levesque, 1991; Grosz, 1996; Jennings, 1995; Tambe et al., 1999). In contrast to early research that focused almost exclusively on how to make agents more autonomous, much of current agent research seeks to understand and satisfy requirements for the basic aspects of joint activity, either within multi-agent systems or as part of human-agent teamwork.

Clancey (2004) argues that the kind of coordinated joint activity we see in multi-agent systems or joint human-agent activities is of such a shallow nature—being based on artificially-construed goals rather than the rich realm of activities—that it is inappropriate to apply the terms “collaboration” or “teamwork” to them. He prefers to restrict his characterization of agents to the use of the term “assistant” and argues that true collaboration requires the kind of consciousness that allows them to have a personal *project*, not just a job, task, or problem. Few researchers would disagree with his point that human-machine interaction is currently of a very different nature than joint activity among humans. Although we are sympathetic to such arguments, we have decided to use these terms in this chapter because the concepts of teamwork and collaboration may be helpful in guiding automation development so that it becomes less likely to result in coordination surprises.

Researchers in human-agent teamwork have used the term in two broad ways: 1) as a conceptual analogy for heuristically directing research (e.g., to build systems that facilitate fluent, coordinated interaction between the human and agent elements of the system as “team players”) and 2) as the subject matter for research (e.g., to understand the nature of teamwork in people). The first activity focuses on practical engineering of useful systems through application of human-centered design principles, empirical studies of the use of these systems, and often a limited commitment to studying teamwork among people. The second activity is explicitly framed as a scientific study, and may have two angles: 1) providing information relevant to the design of successful human-agent systems, and 2) independent of application, understanding the nature of cognition, communication, and cooperation in people and animals. The latter activity is seen by these researchers as essential for achieving the ultimate goals of artificial intelligence. Because the authors of this chapter are drawn from all of these traditions, our perspective attempts to reflect sensitivity to both: neither undervaluing the independent study of social and cognitive aspects of human teamwork, nor slavishly imitating superfluous aspects of natural systems in the development of

artificial ones—like an engineer who insists that successful airplane designs must necessarily feature flapping wings because all birds have them (Ford & Hayes, 1998).

Given the widespread demand for increasing the effectiveness of team play for complex systems that work closely and collaboratively with people, a better understanding of the state of current agent research and observed shortfalls is important. In the remainder of this section we build on our previous examination of the nature of joint activity to describe six common ground requirements that are required to some degree for automation to be a team player. Automation in general, and agent implementations in particular, should be designed and evaluated with a view to their relative sophistication in satisfying each of these requirements.

6.1 The Basic Compact

To be a team player, an agent must fulfill the requirements of a Basic Compact to engage in common grounding activities. The Basic Compact is an agreement to work together in a coordinated fashion and to communicate events and changes in status that the other parties need to know in order to coordinate. Not only does the agent need to be able to enter into such a compact, it must also understand and accept the joint goals of the enterprise, understand and accept its roles in the collaboration, be capable of signaling if it is unable or unwilling to fully participate in the activity, and be capable of understanding other team members' signals of their status and changes in participation.

In the limited realm of what software agents can communicate and reason about among themselves, there has been some limited success in the development of theories and implementations of multi-agent cooperation. Teamwork has become the most widely accepted metaphor for describing the nature of such cooperation between software agents. The key concept usually involves some notion of shared knowledge, goals, and intentions that function as the glue that binds the software agents' activities together (Cohen & Levesque, 1991; Jennings, 1995; Tambe et al., 1999). By virtue of a largely reusable explicit formal model of shared “intentions,” multiple software agents attempt to manage general responsibilities and commitments to each other in a coherent fashion that facilitates recovery when unanticipated problems arise. For example, a common occurrence in joint action is when a software agent fails and can no longer perform in its role. General-purpose teamwork models typically entail that each team member be notified under appropriate conditions of the failure, thus reducing the requirement for special-purpose exception handling mechanisms for each possible failure mode. In this way researchers have been using computational analogues of human joint activity and the Basic Compact to coordinate multiple software modules acting in parallel.

Addressing human-agent teamwork presents a new set of challenges and opportunities for agent researchers. No form of automation today or on the horizon is capable of entering fully into the rich forms of Basic Compact that are used among people. Thus, agents cannot be full-fledged members of human-agent teams in the same sense that other people are—a basic coordinative asymmetry between people and automata. By stretching, we can imagine in the future that some agents will be able to enter into a Basic Compact with diminished capability (Bradshaw, et al., 2004b, c; in preparation). They may eventually be fellow team members with humans in the way a young child can be—subject to the consequences of brittle and literal-minded interpretation of language and events, inability to appreciate or even attend effectively to key aspects of the interaction, poor anticipation, and insensitivity to nuance.

Consider the activity of a traveler who calls up navigation directions from a web-based service. These services have developed extensive maps to enable travelers to select a route. They can recommend a route and even let a traveler pick the fastest route, the shortest route, or a route that avoids highways. That seems so helpful. But a person needing directions may have a different goal—the least confusing route. The computerized aids don't know what that means. “Least confusing” depends on common ground. To provide a “least confusing” route to someone, it is necessary to appreciate what that person already knows

of the area and what landmarks the person is likely to be able to recognize. It is necessary to engage with that person, to appreciate that person's perspective, to draw on common ground. Without this kind of capability software agents will not be able to enter fully into a Basic Compact with humans in these sorts of situations. In such situations the human's ability to appreciate nuance and understand the shortcomings of the machine agent will have to compensate for those aspects of joint activity that are missing in the interaction of machine agents (Clancey, 2004; Woods, Tittle, Feil, & Roesler, 2004).

6.2 Interpredictability

To be a team player, an agent has to be reasonably predictable, and has to have a reasonable ability to predict the actions of others. Similarly, one aspect of the Basic Compact is a commitment to attempt to be predictable to others (i.e., to act neither capriciously nor unobservably), and to be sensitive to the signals that others are sending (which often contribute to *their* predictability).

Thus, agents have three challenges for supporting interpredictability: 1) acting predictably and being directable, 2) signaling their status and intentions, and 3) interpreting signals that indicate the status and intentions of other team members. As with any team coordination arrangement, the signaling and interpretation aspects of interpredictability are reciprocal parts of the effort to increase predictability.

6.2.1 Acting Predictably and Being Directable

The intelligence and autonomy of machine agents directly works against the confidence that people have in their predictability. Although people will rapidly confide tasks to simple deterministic mechanisms whose design is artfully made transparent, they naturally are reluctant to trust complex software agents to the same degree (Bradshaw et al., 2004c). On the one hand, their autonomy and intelligence grants agents the flexibility and additional capability needed to handle situations that require more “wobble room” than traditional software. On the other hand, their blindness to the limits of their competence, their non-transparent complexity, and their inadequate directability can be a formula for disaster (Billings, 1997).

In response to these concerns, agent researchers have increasingly focused on developing means for controlling aspects of agent autonomy in a fashion that can both be dynamically specified and humanly understood—directability. Policies are a means to dynamically regulate the behavior of a system without changing code or requiring the cooperation of the components being governed:

- Through policy, people can precisely express bounds on autonomous behavior in a way that is consistent with their appraisal of an agent's competence in a given context.
- Because policy enforcement is handled externally to the agent, malicious and buggy agents can no more exempt themselves from the constraints of policy than benevolent and well-specified ones can.
- The ability to change policies dynamically means that poorly performing agents can be immediately brought into compliance with corrective measures.

Policy-based approaches to regulation of machine agents have explicitly bound “wobble room,” making those aspects of agent behavior about which people have special concern more predictable (Bradshaw et al., 2004a, b, c; Kagal, Finin, & Joshi, 2003; Myers & Morley, 2003).²

6.2.2. Signaling Status and Intentions

Agents in general must be able to signal their own status, including, for example, their current goals, stance, state of knowledge and upcoming actions to coordinate with others. Wiener (1989) has commented on the fact that the highest levels of automation on the flight deck of commercial jet transport aircraft (Flight Management Systems or FMS) often leave commercial pilots baffled, wondering what the automation is currently doing, why it is doing that, and what it is going to do next. These confusions demonstrate a failure to establish adequate interpredictability. In essence, the machine agent needs to make its own targets, changes, and upcoming actions externally accessible as people manage and re-direct these systems. This challenge runs counter to the advice that is sometimes given to automation developers to create systems that are barely noticed. We are asserting that people need to have a model of the machine as an agent participating in the joint activity. People can often effectively use their own thought processes as a basis for inferring the way their teammates are thinking. But this self-referential heuristic is not usually available in working with agents. In this regard, there is a growing concern that agents in human-agent teams should do a better job of communicating their current states, capacities, and intentions (e.g., Feltovich et al., in press; Norman, 1992).³

Consider this example from a study of pilot interaction with cockpit automation (Sarter & Woods, 2000).

Example 5. Having an altitude problem.

A pilot prepares his descent into his destination airport and receives an initial ATC clearance for an instrument landing approach to runway 24 L together with a number of altitude constraints for various waypoints of the arrival. The pilot programs these constraints into the flight automation.

Shortly after the entire clearance has been programmed, an amended clearance is issued to now make an instrument landing approach to runway 24 R. (The change in runway was made because of an equipment failure). When the pilot changes the runway in the instructions to the automation, the automation signals it understands the runway change and begins to act based on this new target.

Question: Does the automation continue to use the altitude constraints *which still need to be respected*? The pilot assumes that the automation will remember the previously entered altitude constraint, but the system is designed to delete the previous entry when a new entry is made.

² Feltovich, Bradshaw, Jeffers, Suri., and Uszok (in press) argue that policies have important analogues in animal societies and human cultures that can be exploited in the design of artificial systems.

³ To be sure, this hoped-for gain in adaptivity would mean some loss in predictability. Moreover, second-order issues of limited competence would no doubt now emerge at the level of the component doing the adjusting. Policies governing social aspects of agent interaction can be used to help assure that humans are kept properly appraised of agent state and intentions (Bradshaw, Bradshaw et al., 2004a, b, c).

If the pilot was working with a human team member, it would be reasonable to expect that person to continue to apply the constraints—to “remember” the previous intent and know that it is still relevant to the situation. Both team members know that the altitude constraints apply to the current situation; and that the change in runways is not relevant to the altitude constraints. However, in this instance the automation doesn’t behave this way—the software drops all of the altitude constraints entered for the originally planned approach. In addition, the only signal that these constraints have been dropped is through the disappearance of two indications that would be present if the constraints were still in the program.

The pilots may not understand or anticipate that the automation does not remember the constraints following a change but rather reverts to a baseline condition. This lack of interpredictability (and the lack of observability in the displays of automation activity) creates the conditions for an automation surprise—the automation will violate the altitude constraints as it flies the descent unless the human pilot notices and intervenes (in the simulation study 4 of 18 experienced pilots never understood or noticed the automation’s behavior and 12 of 14 who noticed at some point were unable to recover before the constraints were violated). This illustrates how a Fundamental Common Ground Breakdowns occur in human-automata interaction unless special measures are taken in design.

6.2.3 Interpreting Status and Intention Signals

Agents must be able to appreciate the signals given by human teammates. The ideal agent would grasp the significance of such things as pauses, rapid pacing, and public representations that help to mark the coordination activity. Olson and Sarter (2001) have described the way Flight Management Systems may obscure the downstream implications of current decisions. Pilots may believe they are making a choice at point A and not realize that their choice will have unexpected consequences later on, at point B (see also Feltovich, Hoffman et al., 2004). These kinds of automation surprises are the antithesis of team coordination. An ideal team member would instead identify the kinds of outcomes that a pilot might not anticipate, and alert the pilot to these implications.

Few agents are intended to read the signals of their operator teammates with any degree of substantial understanding, let alone nuance. As a result, the devices are unable to recognize the stance of the operator, much less appreciate the operator’s knowledge, mental models, or goals.

With respect to the design of automation that monitors human state and changes based on the machine's assessment of that state, Billings (1997) and Woods (2002) have elucidated many fundamental concerns and have voiced their general skepticism about the value of this line of research. More generally, they are concerned about the basic asymmetry in coordinative competencies between people and machines. Given that asymmetry, the design of human-agent teams will always be difficult.

A few researchers are exploring ways to stretch the performance of agents in order to overcome this asymmetry, such as exploiting and integrating available channels of communication from the agent to the human, and conversely sensing and inferring cognitive state through a range of physiological measures of the human in real time so they can be used to tune agent behavior and thus enhance joint-human machine performance (Bradshaw et al., 2004a; Forsythe & Xavier, in press; Kass, Doyle, Raj, Andrasik, & Higgins, 2003; Raj, Bradshaw, Carff, Johnson, & Kulkarni, 2004). Similarly, a few research efforts are taking seriously the agent’s need to interpret the physical environment (e.g., Rybski & Veloso, in press). Efforts such as these can help us appreciate the difficulty of this problem. For example, by making the machine agents more adaptable, we also make them less predictable.

We see this tradeoff with approaches to “adjustable autonomy” that enable policies to be adjusted as circumstances change, without requiring a human in the loop, essentially amounting to an automated way to “wobble the bounds of the wobble room” (e.g., Bradshaw, et al., 2004b, c; Falcone & Castelfranchi, in

press; Maheswaran, Tambe, Varakantham, & Myers, 2003; Scerri, Pynadath, & Tambe, 2002). As Klein (2004) has pointed out, the more a system takes the initiative in adapting to the existing working style of its operator, the more reluctant operators may be to adapt their own behavior because of the confusions these adaptations might create. This is another example of how the asymmetry between people and agents creates a barrier to forming a full-fledged Basic Compact.

6.3 Goal Negotiation

To be a team player, an entity has to be able to enter into goal negotiation, particularly when the situation changes and the team has to adapt. Agents need to convey their current and potential goals so that operators can participate in the negotiations. Agents need to be readily re-programmable, to allow themselves to be re-directed, and to improve the resilience of the people who work with them to adapt to unexpected events.

Unlike the aviation automation described by Sarter and Woods (2000) and Olson and Sarter (2001), agents have to clearly announce their current intent and permit the operator to easily anticipate the consequences of making a change, the way any human team member would convey the implications of departing from a game plan. If agents are unable to readily represent, reason about, or modify their goals, they will interfere with common ground and coordination.

Traditional planning technologies for software and robotic agents typically take an *autonomy-centered* approach, with representations, mechanisms, and algorithms that have been designed to ingest a set of goals and output as if they can provide a complete plan that handles all situations. This approach is not compatible with what we know about optimal coordination in human-agent interaction. A *collaborative autonomy* approach, on the other hand, takes as a premise that people are working in conjunction with autonomous systems, and hence adopts the stance that the processes of understanding, problem solving, and task execution are necessarily incremental, subject to negotiation, and forever tentative (Bradshaw, et al., 2003; 2004d). Thus, a successful approach to collaborative autonomy will require that every element of the autonomous system be designed to facilitate the kind of give-and-take that quintessentially characterizes natural and effective teamwork among groups of people.

Allen's research on a Collaboration Management Agent (CMA) is a good initial example. It is designed to support human-agent, human-human, and agent-agent interaction and collaboration within mixed human-robotic teams (Allen et al., 2000; Allen et al., 2001; Allen & Ferguson, 2002). The CMA interacts with individual agents in order to 1) maintain an overall picture of the current situation and status of the overall plan, as completely as possible based on available reports, 2) detect possible failures that become more likely as the plan execution evolves and to invoke replanning; 3) evaluate the viability of proposed changes to plans by agents, 4) manage replanning when situations exceed the capabilities of individual agents, including recruiting more capable agents to perform the replanning, 5) manage the re-tasking of agents when changes are made, 6) adjust its communications to the capabilities of the agents (e.g., graphical interfaces work well for a human but wouldn't help most other agents). Because the agents will be in different states based on how much of their original plan they have executed, the CMA must support further negotiation and re-planning among team members while the plan is being executed. These sorts of capabilities should provide a foundation for more ambitious forms of goal negotiation in future agent research.

6.4 Coordination Phases

To be a team player, agents have to partner with the operator in carrying out all of the coordination phases. In many cases, the human-agent coordination involves the handoff of information, from the operator to the system or from the system to the operator.

Clark described coordination as embodying phases that can be thought of as event patterns. This pattern consists of an entry, a body of action, and an exit. The exit portion of the phase is particularly interesting in considering automation issues; each member of the team has to provide evidence to the other members that a particular phase has been completed. Agent-based systems are sometimes lax when it comes to providing feedback to operators, especially about completing or abandoning tasks.

Common ground is created or lost during handoffs between team members (Patterson, Roth, Woods, Chow, & Gomes, 2004). Schreckenghost and her colleagues (Schreckenghost, Martin, Bonasso et al., 2003; Schreckenghost, Martin, & Thronesbery, 2003) have attempted to address some of these problems in their work on agent support for teams of operators. Their vision of future human-agent interaction is that of loosely coordinated groups of humans and agents. As capabilities and opportunities for autonomous operation grow in the future, agents will perform their tasks for increasingly long periods of time with only intermittent supervision. Most of the time routine operation is managed by the agents while the human crews perform other tasks. Occasionally, however, when unexpected problems or novel opportunities arise, people must assist the agents. Because of the loose nature of these groups, such communication and collaboration must proceed asynchronously and in a mixed-initiative manner. Humans must quickly come up to speed on situations with which they may have had little involvement for hours or days (Patterson & Woods, 2001). Then they must cooperate effectively and naturally with the agents. Schreckenghost's group has developed interesting capabilities for managing notification and situation awareness for the crewmembers in these situations. Also relevant is work on generic teamwork phases as it is being developed in an ongoing manner in Brahm's and KAoS (Sierhuis et al., 2003; Bradshaw et al., 2004d).

6.5 Attention Management

In coordinated activity, team members help each other direct their attention to signals, activities and changes that are important (Moore & Dunham, 1995). Similarly, machine agents should refrain from working autonomously but silently, because this places the burden on the operators to discover changes. The other extreme doesn't help either—the case where a partner generates a large volume of low-level messages with little signal but a great deal of distracting noise.

The Basic Compact in part shows how responsible team members expend effort to appreciate what the other needs to notice, within the context of the task and the current situation. Similarly, open shared workspaces provide the means for one agent to notice where another's attention is directed, the activity ongoing, and their stance toward that activity (Carroll, Neale, Isenhour, Rosson, & McCrickard, 2003; Patterson & Woods, 2001; Woods, 1995).

To see the issues, take one example of a coordination breakdown between crews and flight deck automation—bumpy transfer of control (Sarter & Woods, 2000). Trouble begins and slowly builds, for example asymmetric lift worsens as wing icing develops or trouble in an engine slowly reduces performance. Automation can compensate for the trouble, but silently (Norman, 1990). Whether the automation is acting at all, a little, or working more and more to compensate is not visible (an example of a private workspace). Crews can remain unaware of the developing trouble until the automation nears the limits of its authority or capability to compensate. The crew may take over too late or be unprepared to handle the disturbance once they take over, resulting in a bumpy transfer of control and significant control excursions. This general problem has been a part of several incident and accident scenarios.

In contrast, in a well-coordinated human team, the active partner would comment on the unusual difficulty or increasing effort needed to keep the relevant parameters on target. Or, in an open environment, supervisors could notice the extra work or effort exerted by their partners and ask about the difficulty, investigate the problem, or intervene to achieve overall safety goals.

Notice how difficult it is to get the machine to communicate as fluently as a well-coordinated human team working in an open visible environment. The automation should signal (or its activities should be open and visible to other agents so they can see):

- when it is having trouble handling the situation (e.g., turbulence);
- when it is taking extreme action or moving towards the extreme end of its range of authority.

These are quite interesting relational judgments about another agent's activities including: How do we tell when an agent is having trouble in performing a function, but not yet failing to perform? How and when does one effectively reveal or communicate that they are moving towards a limit of capability?

Adding threshold-crossing alarms is the usual answer to these questions in the design of machine agents. But this is a design dilemma as the thresholds are inevitably set too early (resulting in an agent that speaks up too often, too soon) or too late (resulting in an agent that is too silent, speaking up too little). For example, designers have added auditory warning that sounds whenever the automation is active but practitioners usually remove them because the signal is a nuisance, distraction, or false alarm. Sarter and her colleagues (Ho, M., Waters, & Sarter, in press; Sklar & Sarter, 1999) have developed a successful way to overcome these difficulties by designing tactile feedback to signal automation activities. The tactile cues are non-disruptive to other ongoing activities, yet allow the human partner to stay peripherally aware of automation changes and to focus only on those that are unusual given the context or the practitioner's assessment of the situation, thus reducing coordination surprises.

An important aspect of cooperative communications is gauging the interruptibility of other practitioners. In a study of an emergency call-in center, Dugdale, Pavard, and Soubie (2000) found that directing another's attention depended on being able to see what the other agent is doing in order for one agent to be able to judge when another was interruptible. In other words, interruptibility is a joint function of the new message and the ongoing activity. This requires one agent being able to see the activity of the other in enough detail to characterize the state of the other's activities—what line of reasoning are they on? Are they having trouble? Does their activity conform to your expectations about what the other should be doing at this stage of the task? Are they interruptible?

For example, Patterson, Watts-Perotti, and Woods (1999) observed the role of voice loops in mission control at Johnson Space Center. They noticed that controllers gauge the interruptibility of practitioners outside their immediate team before communicating with them. When a controller needed to communicate with another controller working on a different subsystem, the controller would first listen in on that controller's voice loop. By listening to that loop, she could estimate the controller's current workload to judge how interruptible the controller would be in terms of the criticality of the issues that she or he is addressing. Using this strategy reduced the number of unnecessary interruptions and allowed controllers to judge the priority of their item against the ongoing work context. This reduced situations where a controller was forced to direct her attention away from current tasks in order to receive information about a lower priority item.

The Department of Defense has tried to develop a means of attention management in command posts described as a "common operating picture." The concept was to provide a single map display that showed everything on the battlefield—location of friendly and enemy forces, aviation assets, topography, and so forth. The intent to provide common ground by presenting a single shared picture misses the active nature of building common ground.

To illustrate how a single picture is not the same as common ground, consider an exercise in which a Brigade Commander acted like an aide, in order to ensure that a staff member had seen a key piece of information on the display. During the exercise, a critical event occurred and was entered into the large screen display. The commander heard about it on his radio, and noted the change in the display. But he was not sure that one of his staff members had seen the change to the large screen display. Therefore, the commander made a radio call to the staff member to point out the event, as marked on the display. The commander made the call because he felt it was so important to manage the attention of his subordinate, and because the technology did not let him see if the staff member had noticed the event.

The common operating picture eliminated the feedback portion of the information hand-off. As a result, the Brigade Commander could not infer whether the staff member had scaled the joint action ladder, and he had to intervene personally to get confirmation. The Brigade Commander appreciated how the technology created a potential for a Fundamental Common Ground Breakdown and sought to prevent a coordination surprise.

A common picture is not necessarily a barrier to common ground. It can serve as a useful platform that enables distributed team members to calibrate the way they understand events. However, it does not automatically produce common ground. Even worse, the use of a common picture can create the conditions for a Fundamental Common Ground Breakdown, as illustrated by the example involving the Brigade Commander, because of the way it can interfere with attention management.

6.6 Controlling the Costs of Coordinated Activity

The Basic Compact commits people to coordinating with each other, and to incurring the costs of providing signals, improving predictability, monitoring the other's status, and so forth. All of these take time and energy. These coordination costs can easily get out of hand, and therefore the partners in a coordination transaction have to do what they reasonably can to keep coordination costs down. This is a tacit expectation—to try to achieve economy of effort.

Achieving coordination requires continuing investment and hence the power of the Basic Compact—a willingness to invest energy and accommodate to others, rather than just performing alone in one's narrow scope and sub-goals. The six challenges are ongoing investments. Coordination doesn't come for free; and coordination, once achieved, does not allow one to stop investing. Otherwise the coordination breaks down.

Agents can pose difficulties in serving as team members, and they can also impose penalties when team members are trying to coordinate with each other. Here, interaction with agents can increase the coordination costs of an information handoff. To find out if someone received an e-mail we might request a receipt, but in so doing we incur and inflict coordination costs of five or ten seconds per transaction from each person. To gauge the effect of an anesthetic on a patient, a physician may have to watch a computer monitor for thirty seconds after the initial physiological reaction becomes measurable. To trick the system into showing critical measures, the physician has to endure coordination costs for what used to be a simple information seeking action.

Keeping coordination costs down is partly a matter of good human-computer interface design. But it takes more than skillful design: the agents must be able to actively seek to conform to the needs of the operators, rather than requiring operators to adapt to them. Information handoff, which is a basic exchange in coordination phases involving humans and agents, depends on common ground and interpredictability. As we have seen, agents have to become more understandable and predictable, and more sensitive to the needs and knowledge of people.

The work of Horvitz (1999; Horvitz, Jacobs, & Hovel, 1999) provides one example (among the many that could be provided) of explicitly taking the cost of coordination into account in planning interaction with the user. For example, he has used measures of expected utility to evaluate the tradeoffs involved in potentially interrupting the ongoing activities of humans.

6.7 Concluding Remarks About the Challenges

The six challenges we have presented can be viewed in different lights. They can be seen as a blueprint for inspectors who want to evaluate automation. These challenges can be viewed as requirements for successful operation, so that the design of software and robotic agents and other kinds of automation are less likely to result in coordination breakdowns. The six challenges can also be viewed as a cautionary tale, ways that the technology can disrupt rather than support coordination. They can be taken as reminders of the importance of shared experience and expertise with regard to common ground. Simply relying on explicit procedures, such as common operating pictures, is not likely to be sufficient. The challenges can also be used to design team and organizational simulations that capture coordination breakdowns and other features of joint activity. Finally, the six challenges can be viewed as the underlying basis for human-agent systems.

7. SUMMARY

Based on our review of alternative analogies and previous results, we identified a set of central concepts for coordination in joint activity. The foundation for coordinated activity is the Basic Compact or intent to work together to align goals and to invest effort to sustain common interests. The Basic Compact reflects reciprocal commitment that is renewed and sustained. From this foundation, the criteria for engaging in a joint activity include the interdependence of the work performed by the participants, and the expectation that the joint activity will be resilient and adaptive to unexpected events. If these criteria are to be met, then the participants must meet requirements for making their actions predictable to each other, for sustaining common ground, and for being open to direction and redirection from each other as the activity unfolds. The choreography for carrying out these requirements involves coordinating a series of phases, and it is accomplished through employing various forms of signaling and the use of coordination devices, all of which incur coordination costs. Participants have to signal each other, and they also have to direct each other's attention to ensure that the signals are received, and to ensure that public events are noticed.

One of the key aspects of joint action is the process of sustaining common ground to enable coordination. Common ground is not a binary or constant feature—it is both continuous in its degree and constantly changing over time (cf. Feltovich et al., 1989; 2004, regarding the “reductive bias”). This includes the role of shared experience and expertise, and poses a limitation to those who think coordination can be manufactured through procedures and explicit guidelines.

Key aspects of common ground include: 1) The types of knowledge, beliefs and assumptions that are important for joint activity, including knowledge of roles and functions, standard routines, and so forth; 2) Mechanisms for carrying out the grounding process: to prepare, monitor and sustain, catch and repair breakdowns; 3) The Basic Compact committing the parties in a joint activity to continually inspect and adjust common ground.

Common ground is likely to become degraded during team interactions, unless effort is put into calibrating the perspectives of the team members. We described the Fundamental Common Ground Breakdown as a paradigm for how team members can lose their calibration without knowing it, and continue to interact until they run into a coordination surprise.

Common ground is reflected in the amount of work needed in order to manage the communications for a joint activity. As common ground builds mutual knowledge, beliefs and assumptions, participant's communications become coded and abbreviated, and economical. As common ground improves, the effort needed to clarify and explain should decrease. That is why effective teams are so careful to establish as much common ground as they can in advance of critical periods of activity. Furthermore, effective teams have also learned where to expect coordination breakdowns, and in preparation they elaborate common ground in these areas: e.g., clarifying the roles and functions of the participants, the goals and goal tradeoffs, the skills and competencies of the participants, and the preexisting differences in mental models.

Our examination of joint activity and coordination did not particularly focus on the use of technology to facilitate team coordination—the field of Computer Supported Collaborative Work (CSCW). The application of Clark's work on joint activity and common ground to CSCW is being pursued by others (e.g., Carroll et al., 2003). Since the costs of coordinative communication can be considerable, particularly for a large organization, another topic for additional investigation is how these workload costs are managed relative to the benefits of establishing and sustaining common ground.

This chapter should provide the reader with a deeper appreciation of the nuances of coordination, and some ideas about how to evaluate coordination mechanisms in an organization and in a technology. The concepts that we have discussed in this chapter should be useful for researchers who need to observe teams in action. The chapter may also be helpful for modelers trying to describe and simulate teamwork, by identifying important variables, especially those pertinent to coordination. Finally, our account of joint activity and coordination should be informative to developers of automation, by serving as guidelines they can use for improving coordination in mixed human-agent work configurations.

ACKNOWLEDGEMENTS

We would like to thank Joan Feltovich for reading drafts and making suggestions that improved the manuscript. John Flach generously provided us with the results of his observational study of high school football coaches. Herb Bell, Donald Cox, and Dave Klinger provided us with very valuable and constructive critiques of earlier drafts. This research was supported by ARLADA contracts to Klein Associates, Ohio State, and to IHMC. We prepared this chapter through participation in the Advanced Decision Architectures Collaborative Technology Alliance, sponsored by the U.S. Army Research Laboratory under cooperative agreement DAAD19-01-2-0009.

References

- Allen, J., Byron, D. K., Dzikovska, M., Ferguson, G., Galescu, L., & Stent, A. (2000). An architecture for a generic dialogue shell. *Journal of Natural Language Engineering*, 6(3), 1-16.
- Allen, J. F. (1999). Mixed-initiative interaction. *IEEE Intelligent Systems*, 14(5), 14-16.
- Allen, J. F., Byron, D. K., Dzikovska, M., Ferguson, G., Galescu, L., & Stent, A. (2001). Towards conversational human-computer interaction. *AI Magazine*, 22(4), 27-35.
- Allen, J. F., & Ferguson, G. (2002). Human-machine collaborative planning. In *Proceedings of the NASA Planning and Scheduling Workshop*. Houston, TX.
- Billings, C. E. (1997). *Aviation automation: The search for a human-centered approach*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Boy, G. A. (1988). *Cognitive function analysis*. Norwood, NJ: Ablex.
- Bradshaw, J. M. (Ed.). (1997). *Software Agents*. Cambridge, MA: The AAAI Press/The MIT Press.
- Bradshaw, J. M., Acquisti, A., Allen, J., Breedy, M., Bunch, L., Chambers, N., Galescu, L., Goodrich, M., Jeffers, R., Johnson, M., Jung, H., Lott, J., et al. (2004d). Teamwork-centered autonomy for extended human-agent interaction in space applications. In *Proceedings of the AAAI Spring Symposium* (pp. 136-140). Stanford, CA: The AAAI Press.
- Bradshaw, J. M., Beautement, P., Breedy, M., Bunch, L., Drakunov, S. V., Feltovich, P. J., Hoffman, R. R., Jeffers, R., Johnson, M., Kulkarni, S., Lott, J., Raj, A., Suri, N., & Uszok, A. (2004a). Making agents acceptable to people. In N. Zhong & J. Liu (Eds.), *Intelligent Technologies for Information Analysis: Advances in Agents, Data Mining, and Statistical Learning*. Berlin: Springer Verlag.
- Bradshaw, J. M., Boy, G., Durfee, E., Gruninger, M., Hexmoor, H., Suri, N., Tambe, M., Uschold, M., & Vitek, J. (Ed.). (in preparation). *Software Agents for the Warfighter. ITAC Consortium Report*.
- Bradshaw, J. M., Feltovich, P. J., Jung, H., Kulkarni, S., Taysom, W., & Uszok, A. (2004b). Dimensions of adjustable autonomy and mixed-initiative interaction. In M. Klusch, G. Weiss & M. Rovatsos (Eds.), *Computational Autonomy*. Berlin, Germany: Springer-Verlag.
- Bradshaw, J. M., Jung, H., Kulkarni, S., Allen, J. Bunch, L., Chambers, N., Feltovich, P., Galescu, L., Jeffers, R., Johnson, M., Taysom, W. & Uszok, A. (2004c). Toward trustworthy adjustable autonomy and mixed-initiative interaction in KAoS. Proceedings of the AAMAS 2004 Trust workshop, New York City, NY, July.
- Bradshaw, J. M., Sierhaus, M., Acquisti, A., Feltovich, P., Hoffman, R., Jeffers, R., Prescott, D., Suri, N., Uszok, A., & Van Hoof, R. (2003). Adjustable autonomy and human-agent teamwork in practice: An interim report from space applications. In H. Hexmoor, C. Castelfranchi & R. Falcone (Eds.), *Agent autonomy* (pp. 243-280). Boston, MA: Kluwer Academic Press.
- Brennan, S. E. (1998). The grounding problem in conversations with and through computers. In S. R. Fussell & R. J. Kreuz (Eds.), *Social and cognitive psychological approaches to interpersonal communication* (pp. 210-225). Mahwah, NJ: Lawrence Erlbaum Associates.

- Cahn, J. E., & Brennan, S. E. (1999). A psychological model of grounding and repair in dialog. In *Proceedings of AAAI Fall Symposium on Psychological Models in Collaborative Systems* (pp. 25-33). North Falmouth, MA: American Assoc. for Artificial Intelligence.
- Carroll, J. M., Neale, D. C., Isenhour, P. L., Rosson, M. B., & McCrickard, D. S. (2003). Notification and awareness: synchronizing task-oriented collaborative activity. *International Journal of Human-Computer Studies*, 58, 605-632.
- Christoffersen, K., & Woods, D. D. (2002). How to make automated systems team players. *Advances in Human Performance and Cognitive Engineering Research*, 2, 1-12.
- Clancey, W. B. (2004). Roles for agent assistants in field science: Understanding personal projects and collaboration. *IEEE Transactions on Systems, Man, and Cybernetics--Part C: Applications and Reviews*, 32(2), 125-137.
- Clancey, W. J. (1997). The conceptual nature of knowledge, situations, and activities. In P. J. Feltovich, K. M. Ford & R. R. Hoffman (Eds.), *Expertise in context: Human and machine* (pp. 248-291). Menlo Park, CA: AAAI/MIT Press.
- Clark, H. (1996). *Using language*. Cambridge: Cambridge University Press.
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine & S. D. Teasley (Eds.), *Perspectives on socially shared cognition*. Washington: D.C.: American Psychological Association.
- Cohen, P. R., & Levesque, H. J. (1991). Teamwork. *Nous*, 25, 487-512.
- Cook, R. I., & Woods, D. D. (1994). Operating at the Sharp End. In M. S. Bogner (Ed.), *Human Error in Medicine*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Dismukes, K., Young, G., & Sumwalt, R. (1998). Cockpit Interruptions and Distractions: Effective Management Requires a Careful Balancing Act. *ASRS Directline*.
- Dugdale, J., Pavard, B., & Soubie, J. L. (2000). A Pragmatic Development of a Computer Simulation of an Emergency Call Centre. In *Proceedings of COOP 2000, Fourth International Conference on the Design of Cooperative Systems*. Cannes, France.
- Falcone, R., & Castelfranchi, C. (in press). Adjustable Social Autonomy. In J. Pitt (Ed.), *The Open Agent Society*. New York: John Wiley & Sons.
- Feltovich, P. J., Bradshaw, J. M., Jeffers, R., Suri, N., & Uszok, A. (in press). Social order and adaptability in animal and human cultures as analogues for agent communities: Toward a policy-based approach. In A. Omacini, P. Petta & J. Pitt (Eds.), *Engineering societies in the agents world IV (Lecture Notes in Computer Science Series)*. Heidelberg, Germany: Springer-Verlag.
- Feltovich, P. J., Hoffman, R. R., Woods, D., & Roesler, A. (2004). Keeping it too simple: How the reductive tendency affects cognitive engineering. *IEEE Intelligent Systems*, 19(3), 90-94.
- Feltovich, P. J., Spiro, R. J., & Coulson, R. L. (1989). The nature of conceptual understanding in biomedicine: The deep structure of complex ideas and the development of misconceptions. In D. Evans & V. Patel (Eds.), *Cognitive science in medicine: Biomedical modeling*. Cambridge, MA: MIT Press.

Feltovich, P. J., Spiro, R. J., Coulson, R. L., & Feltovich, J. (1996). Collaboration within and among minds: Mastering complexity, individually and in groups. In T. Koschmann (Ed.), *CSCL: Theory and practice of an emerging paradigm* (pp. 25-44). Mahwah, NJ: Lawrence Erlbaum Associates.

Fitts, P. M. (Ed.). (1951). *Human engineering for an effective air navigation and traffic control system*. Washington, D.C.: National Research Council.

Flach, J., & Dominguez, C. (2003). Supporting the adaptive human expert: A critical element in the design of meaning processing systems. In L. Hettinger & M. Hass (Eds.), *Virtual and adaptive environments*.

Ford, K. M., & Hayes, P. J. (1998). On computational wings: Rethinking the goals of Artificial Intelligence. *Scientific American*. *Special issue on "Exploring Intelligence"*, 9(4), 78-83.

Forsythe, C., & Xavier, P. (in press). Cognitive models to cognitive systems. In C. Forsythe, M. L. Bernold, & T. E. Goldsmith (Eds.), *Cognitive Systems: Cognitive Models in System Design*. Hillsdale, N.J.: Lawrence Erlbaum.

Grosz, B. (1996). Collaborative Systems. *AI Magazine*, 2(17), 67-85.

Hancock, P. A., & Scallen, S. F. (1998). Allocating functions in human-machine systems. In R. Hoffman, M. F. Sherrick & J. S. Warm (Eds.), *Viewing Psychology as a Whole* (pp. 509-540). Washington, D.C.: American Psychological Association.

Ho, C.-Y., M., N., Waters, M., & Sarter, N. B. (in press). Not now: supporting attention management by indicating the modality and urgency of pending task. *Human Factors*.

Hoffman, R. R., Klein, G., & Laughery, K. R. (2002). The state of cognitive systems engineering. *IEEE Intelligent Systems*, 17(1), 73-75.

Horvitz, E. (1999). Principles of mixed-initiative user interfaces. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '99) held in Pittsburgh, PA*. Pittsburgh, PA: ACM Press.

Horvitz, E., Jacobs, A., & Hovel, D. (1999). Attention-sensitive alerting. In *Proceedings of the Conference on Uncertainty and Artificial Intelligence (UAI '99)* (pp. 305-313). Stockholm, Sweden.

Jennings, N. R. (1995). Controlling cooperative problem-solving in industrial multiagent systems using joint intentions. *Artificial Intelligence*, 75, 195-240.

Kagal, L., Finin, T., & Joshi, A. (2003). A policy language for pervasive systems. In *Proceedings of the Fourth IEEE International Workshop on Policies for Distributed Systems and Networks, June*. Lake Como, Italy.

Kass, S. J., Doyle, M., Raj, A. K., Andrasik, F., & Higgins, J. (2003). Intelligent adaptive automation for safer work environments. In J. C. Wallace & G. Chen (Co-Chairs) (Eds.), *Occupational health and safety: Encompassing personality, emotion, teams, and automation. Symposium conducted at the Society for Industrial and Organizational Psychology 18th Annual Conference*. Orlando, FL, April.

- Klein, G. (2001). Features of team coordination. In M. McNeese, M. R. Endsley & E. Salas (Eds.), *New trends in cooperative activities* (pp. 68-95). Santa Monica, CA: HFES.
- Klein, G. (2004). *The power of intuition*. New York: A Currency Book/Doubleday.
- Klein, G., Armstrong, A., Woods, D. D., Gokulachandra, M., & Klein, H. A. (2000). *Cognitive wavelength: The role of common ground in distributed replanning* (Final Technical No. Report No. AFRL-HE-WP-TR-2001-0029). Wright-Patterson AFB, OH: United States Air Force Research Laboratory.
- Klinger, D. W., & Klein, G. (1999). Emergency response organizations: An accident waiting to happen. *Ergonomics In Design*, 7(3), 20-25.
- Koschmann, T. D., LeBaron, C., Goodwin, C., & Feltovich, P. J. (2001). Dissecting common ground: Examining an instance of reference repair. In *Proceedings of the 23rd Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Maheswaran, R. T., Tambe, M., Varakantham, P., & Myers, K. (2003). Adjustable Autonomy challenges in Personal Assistant Agents: A Position Paper. In *Proceedings of Computational Autonomy--Potential, Risks, Solutions (Autonomy 2003)*. Melbourne, Australia.
- Malin, J. T., Schreckenghost, D. L., Woods, D. D., Potter, S. S., Johannesen, L., Holloway, M., & Forbus, K. D. (1991). *Making intelligent systems team players: Case studies and design issues* (NASA Technical Memorandum 104738). Houston, TX: NASA Johnson Space Center.
- Malone, T. W., & Crowston, K. (1994). The interdisciplinary study of coordination. *ACM Computing Surveys*, 26, 87-119.
- Moore, C., & Dunham, P. (Eds.). (1995). *Joint attention: Its origins and role in development*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Myers, K., & Morley, D. (2003). Directing agents. In H. Hexmoor, C. Castelfranchi & R. Falcone (Eds.), *Agent Autonomy* (pp. 143-162). Dordrecht, The Netherlands: Kluwer.
- Norman, D. A. (1990). The "problem" with automation: Inappropriate feedback and interaction, not "over-automation." *Philosophical transactions of the Royal Society of London*, 327, 585-593.
- Norman, D. A. (1992). Turn signals are the facial expressions of automobiles. In *Turn Signals Are the Facial Expressions of Automobiles*. (pp. 117-134). Reading, MA: Addison-Wesley.
- Olson, G. M., Malone, T. W., & Smith, J. B. (2001). *Coordination theory and collaboration technology*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Olson, W. A., & Sarter, N. B. (2001). Management by consent in human-machine systems: When and why it breaks down. *Human Factors*, 43(2), 255-266.
- Patterson, E. S., Roth, E. M., Woods, D. D., Chow, R., & Gomes, J. O. (2004). Handoff strategies in settings with high consequences for failure: Lessons for health care operations. *International Journal for Quality in Health Care*, 16(2), 125-132.

- Patterson, E. S., Watts-Perotti, J. C., & Woods, D. D. (1999). Voice loops as coordination aids in Space Shuttle Mission Control. *Computer Supported Cooperative Work*, 8, 353-371.
- Patterson, E. S., & Woods, D. D. (2001). Shift changes, updates, and the on-call model in space shuttle mission control. *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, 10(3), 317-346.
- Patterson, E. S., Woods, D. D., Sarter, N. B., & Watts-Perotti, J. C. (1998). Patterns in cooperative cognition. In *COOP '98, Third International Conference on the Design of Cooperative Systems*. Cannes, France.
- Raj, A. K., Bradshaw, J. M., Carff, R. W., Johnson, M., & Kulkarni, S. (2004). An agent based approach for Aug Cog integration and interaction. In *Proceedings of Augmented Cognition-Improving Warfighter Information Intake Under Stress, Scientific Investigators Meeting*. Orlando, FL 6-8 Jan 04.
- Rybski, P. E., & Veloso, M. M. (in press). Inferring human interactions from sparse visual data. In *Proceedings of the Autonomous Agents and Multi-Agent Systems Conference (AAMAS 2004)*. New York: ACM Press.
- Salas, E., & Cannon-Bowers, J. A. (2001). The science of training: A decade of progress. *Annual Review of Psychology*, 52, 471-499.
- Sarter, N., & Woods, D. D. (2000). Team Play with a Powerful and Independent Agent: A Full Mission Simulation. *Human Factors*, 42, 390-402.
- Sarter, N. B., Woods, D. D., & Billings, C. (1997). Automation surprises. In G. Salvendy (Ed.), *Handbook of human factors/ergonomics* (Second ed.). New York: John Wiley & Sons, Inc.
- Scerri, P., Pynadath, D., & Tambe, M. (2002). Towards adjustable autonomy for the real world. *Journal of AI Research (JAIR)*, 17, 171-228.
- Schreckenghost, D., Martin, C., Bonasso, P., Kortenkamp, D., Milam, T., & Thronesbery, C. (2003). Supporting group interaction among humans and autonomous agents. Submitted for publication.
- Schreckenghost, D., Martin, C., & Thronesbery, C. (2003). Specifying organizational policies and individual preferences for human-software interaction. Submitted for publication.
- Schaeffer, J. (1997). *One jump ahead: Challenging human supremacy in checkers*. New York: Springer-Verlag.
- Shalin, V. L., Geddes, N. D., Bertram, D., Szczepkowski, M. A., & DuBois, D. (1997). Expertise in dynamic, physical task domains. In P. J. Feltovich, K. M. Ford & R. R. Hoffman (Eds.), *Expertise in context: Human and machine* (pp. 195-217). Menlo park, CA: AAAI/MIT Press.
- Shattuck, L. G., & Woods, D. D. (2000). Communication of intent in military command and control systems. In C. McCann & R. Pigeau (Eds.), *The human in command: Exploring the modern military experience* (pp. 279-291). New York: Plenum Publishers.
- Sierhuis, M., Bradshaw, J. M., Acquisti, A., Van Hoof, R., Jeffers, R., & Uszok, A. (2003). Human-agent teamwork and adjustable autonomy in practice. In *Proceedings of the Seventh International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS), 19-23 May*. Nara, Japan.

- Sklar, A. E., & Sarter, N. B. (1999). "Good Vibrations": The Use of Tactile Feedback In Support of Mode Awareness on Advanced Technology Aircraft. *Human Factors*, 41(4), 543-552.
- Smith, P., Woods, D., McCoy, E., Billings, C., Sarter, N., R., D., & Dekker, S. (1998). Using forecasts of future incidents to evaluate future ATM system designs. *Air Traffic Control Quarterly*, 6(1), 71-85.
- Smith, W. J. (1977). *The Behavior of Communicating*. Cambridge, MA: Harvard University Press.
- Smith, W. J. (1995). The biological bases of social attunement. *Journal of Contemporary Legal Issues*, 6.
- Spiro, R. J., Feltovich, P. J., Coulson, R. L., & Anderson, D. K. (1989). Multiple analogies for complex concepts: Antidotes for analogy-induced misconception in advanced knowledge acquisition. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 498-531). Cambridge, England: Cambridge University Press.
- Tambe, M., Shen, W., Mataric, M., Pynadath, D. V., Goldberg, D., Modi, P. J., Qiu, Z., & Salemi, B. (1999). Teamwork in cyberspace: Using TEAMCORE to make agents team-ready. In *Proceedings of the AAAI Spring Symposium on Agents in Cyberspace*. Menlo Park, CA: The AAAI Press.
- Weick, K. E., Sutcliffe, K. M., & Obstfeld, D. (1999). Organizing for high reliability: Processes of collective mindfulness. *Research in Organizational Behavior*, 21, 13-81.
- Wiener, E. L. (1989). *Human factors of advanced technology ("glass cockpit") transport aircraft* (No. NASA Report 177528). Moffett Field, CA: Ames Research Center.
- Woods, D. D. (1995). The alarm problem and directed attention in dynamic fault management. *Ergonomics*, 38(11), 2371-2393.
- Woods, D. D. (2002). *Steering the reverberations of technology change on fields of practice: Laws that govern cognitive work*, Proceedings of the 24th Annual Meeting of the Cognitive Science Society [Plenary Address] url: <http://cse1.eng.ohio-state.edu/laws>
- Woods, D. D., Tittle, J., Feil, M., & Roesler, A. (2004). Envisioning human-robot coordination in future operations. *IEEE SMC Part C*, 34(2), 210-218.
- Zalesny, M. D., Salas, E., & Prince, C. (1995). Conceptual and measurement issues in coordination: Implications for team behavior and performance. In M. D. Zalesny, E. Salas & C. Prince (Eds.), *Research in personnel and human resources management* (Vol. 13, pp. 81-115). Greenwich, CT: JAI Press Inc.

BIOSKETCHES

Gary Klein

Klein Associates Inc.
 1750 Commerce Center Blvd. North
 Fairborn, OH 45324
 PH: 937-873-8166
 FX: 937-873-8258
 gary@decisionmaking.com

Gary Klein, Ph.D., is Chief Scientist of Klein Associates, Inc., a company he founded in 1978 to develop stronger models and methods for studying and improving the way people make decisions in natural settings. In 1985, Dr. Klein and his colleagues developed the Recognition-Primed Decision (RPD) model to explain how people can make effective decisions under time pressure and uncertainty. The research has been extended to training individuals and teams, and to applying a decision-centered design approach to increase the impact of information technology. In addition, Dr. Klein and his colleagues have developed a set of cognitive task analysis methods to study decision making and other cognitive functions in field settings.

Dr. Klein received his Ph.D. in experimental psychology from the University of Pittsburgh in 1969. He was an Assistant Professor of Psychology at Oakland University (1970-1974) and worked as a research psychologist for the U.S. Air Force (1974-1978). He is the author of "Sources of Power: How People Make Decisions" (1998, MIT Press) and "The Power of Intuition" (2004, Doubleday Currency).

Paul J. Feltovich

Institute for Human and Machine Cognition
 40 S. Alcaniz Street
 Pensacola, FL 32502
 850-202-4470; 850-232-4345
 pfeltovich@ihmc.us

Paul J. Feltovich, Ph.D., is a Research Scientist at the Institute for Human and Machine Cognition, Pensacola, FL. He has conducted research and published on topics such as expert-novice differences in cognitive skills, conceptual understanding for complex knowledge, and novel means of instruction in complex and ill-structured knowledge domains. More recently (with Jeffrey Bradshaw) he has been studying regulatory and coordination systems for mixed human-agent teams.

Dr. Feltovich received a Ph.D. in educational psychology from the University of Minnesota in 1981. He was also a post-doctoral fellow in cognitive psychology at the Learning, Research, and Development Center, University of Pittsburgh, from 1978 to 1982. Before joining IHMC in 2001, he served as Professor in the Department of Medical Education, and Director of the Cognitive Science Division at Southern Illinois University School of Medicine, Springfield, IL. He is co-editor (with Ken Ford and Robert Hoffman) of *Expertise in Context: Human and Machine* (AAAI/MIT) and (with Ken Forbus) *Smart Machines in Education* (AAAI/MIT).

Jeffrey M. Bradshaw

Institute for Human and Machine Cognition
 40 S. Alcaniz Street
 Pensacola, FL 32502

PH: 850-202-4470; 850-232-4345
jbradshaw@ihmc.us

Jeff Bradshaw, Ph.D., is a Senior Research Scientist at the Institute for Human and Machine Cognition where he leads the research group developing the KAoS policy and domain services framework. Formerly, he has led research groups at The Boeing Company and the Fred Hutchinson Cancer Research Center. In research sponsored by DARPA, NASA, ONR, and ARL, he is investigating principles of human-robotic teamwork, human-agent interaction, and trust and security for semantic web and semantic grid services.

Dr. Bradshaw received his Ph.D. in cognitive science from the University of Washington in 1996. He has been a Fulbright Senior Scholar at the European Institute for Cognitive Sciences and Engineering (EURISCO) in Toulouse, France, is a member and former chair of the NASA Ames RIACS Science Council, and is former chair of ACM SIGART. Jeff serves on the editorial board of the Journal of Autonomous Agents and Multi-Agent Systems, the International Journal of Human-Computer Studies, the Web Semantics Journal, and the Web Intelligence Journal. Among other publications, he edited the books Knowledge Acquisition as a Modeling Activity (with Ken Ford, Wiley, 1993), Software Agents (AAAI Press/MIT Press, 1997) and the forthcoming Handbook of Agent Technology.

David D. Woods

The Ohio State University
Cognitive Systems Engineering Laboratory
210 Baker Systems
1971 Neil Avenue
Columbus, OH 43210
PH: 614-292-1700
FX: 614-292-7852
woods.2@osu.edu

David Woods, Ph.D. is Professor in the Institute for Ergonomics at the Ohio State University. He has developed and advanced the foundations and practice of Cognitive Systems Engineering since its origins in the aftermath of the Three Mile Island accident in nuclear power. He has studied team work between people and automation, including automation surprises, through studies in anesthesiology and aviation. He has studied cooperative work systems in space mission operations and he has designed new concepts for aiding cognitive work such as visual momentum and applied them in aviation, space operations, nuclear power, and critical care medicine. Multimedia overviews of his research are available at url: <http://csel.eng.ohio-state.edu/woods/>.

Safety in complex systems is a constant theme in his work; see his monographs: "Behind Human Error" and "A Tale of Two Stories: Contrasting Views of Patient Safety." He was one of the founding board members of the National Patient Safety Foundation, Associate Director of the MidWest Center for Inquiry on Patient Safety of the Veterans Health Administration, and advisor to the Columbia Accident Investigation Board.

Dr. Woods received his Ph.D. from Purdue University in 1979 and has been President of the Human Factors and Ergonomic Society. He is a Fellow of that society and a Fellow of the American Psychological Society and the American Psychological Association. He has shared the Ely Award for best paper in the journal Human Factors (1994), a Laurels Award from Aviation Week and Space Technology (1995) for research on the human factors of highly automated cockpits, the Jack Kraft Innovators Award from the Human Factors and Ergonomics Society (2002), and five patents for computerized decision aids.