A Process-Oriented Definition of Teamwork with Consequences for Team Cognition and Performance

Christopher Martin

Land Operations Division,
Defence Science and Technology Organisation,
Department of Defence, Australian Government.[1]
christopher.martin@dsto.defence.gov.au

Abstract

In attempting to predict and improve team performance, researchers have sought to define the essential elements of teamwork. Although component-oriented frameworks of team performance dimensions appear to be in a state of flux, many include some concept of *shared mental models*. Based on theoretical requirements for the organisation and functionality of mental models, we argue for a process-oriented approach that could underpin many of the components and coordinating mechanisms proposed by other frameworks, providing a lens through which one may re-examine team phenomena and re-assess how to develop more effective teams. We consider tasks and teams as organisations of processes and examine in detail their process components and their interactive characteristics and constraints. Key to this approach are several elements: (1) hierarchical decomposition of tasks into *alternative sets* composed of sequences of specific subtasks and concrete activities (Miller & Parasuraman, 2007); (2) delineation of the function and content of mental models as *situation knowledge webs* (Bickhard, 2005); (3) extension of such mental models from the individual into the team domain (Bosse et al., 2005); (4) a concept of *apperceptive value* in relation to the dynamic operation and reorganisation of team-on-task processes; and (5) a definition of teamwork as being constituted in the selection of *mutually beneficial methods* for task performance. This approach offers a novel perspective on the assessment and improvement of team performance. It also invites a theoretical critique of some popular research methodologies and associated findings.

Introduction

There are longstanding conceptual and definitional issues that have clouded applied research into teaming such that there is still no widely recognised theoretical framework sufficient to construct an operational model of effective teamwork and determine how teamwork could be measured and improved. Nor is it clear how a teamwork facilitator might choose one framework over another. This is particularly so in the context

---

[1] This material is UNCLASSIFIED and approved for public release.

of *action teams* (Marks et al., 2002) as exemplified by many military organisations. Action teams are characterised by the following:

- Performance in less familiar and more challenging environments
- Membership and/or tasking may be relatively temporary
- Expertise, information and tasks are distributed
- Members have specialised task-related skill sets
- Members have honed teamwork abilities to coordinate their activities
- Team performance and effectiveness depends on:
  - task contributions from each member
  - successful interaction among the team members
  - rapid, complex, and coordinated task behaviour
  - ability to dynamically adapt to the shifting demands of the situation

Cannon-Bowers et al. (1995) defined teamwork as a three-dimensional complex of cognitions, skills and attitudes. But for Marks et al. (2001), teamwork was a set of cognitive, verbal, and behavioral activities (or mediating processes) operating on team properties and directed toward organising team members' interactions with tasks, tools, machines, and systems. Then Salas et al. (2005) proposed that teamwork was a 'big five' set of task-independent competencies with three coordinating mechanisms. But a year later, Driskell & Salas (2006) were describing teamwork in terms of eight broad dimensions of activity or function and six categories of contextual moderating factors. Meanwhile, others were drawing from this heritage to propose categorisations in a similar vein. For example, Baker et al. (2004) proposed a complex of attitudes, experiences and implicit theories as one dimension, with a second dimension composed of three broad categories of cognitive, personal and interpersonal skills underpinned by a fourth category of communication skills. But Shanahan et al. (2007) reworked (with additions) the eight dimensions and many subskills of Driskell & Salas (ibid.) into three dimensions of broad process categories; namely, behavioural, cognitive and motivational.

So researchers and practitioners alike have been asking for decades, "What is teamwork?" The quest for an answer has produced a profusion and confusion of concepts and terminology. Teamwork's defining dimensions and categories have been in flux with elements being added, repositioned, redefined or discarded. The influence of each proposed factor is highly variable and context-dependent, encouraging the expansion of taxonomies and concepts in attempts to determine and classify contexts in which these factors come into play to varying degrees. Approaches to resolving conflicting findings have continued to focus on classification schemes for teams, tasks and/or contexts. Hypotheses have taken the general form of "factor *X* (say, closed-loop communication) is important for teams of type *Y* working on tasks of type *Z*, but not for teams of type *A* or in phase *B* or on tasks of type *C* or in contexts of type *D*." For example:

- Salas et al. (2005) claim that teams-on-task go through an early phase in which closed loop communication and leadership are relatively important to team function while in a later phase, mutual performance monitoring and back-up behaviours are more important.

- Espinosa et al. (2004) and MacMillan et al. (2004) make a distinction between explicit coordination versus implicit coordination with respect to newly formed versus established teams.
- Drawing on prior research, Lee (2007) differentiates task-related versus team-related shared mental models, and pre-process versus in-process coordination, to suggest that team-related modelling might be more important during pre-process coordination, and task-related modelling more important during in-process coordination.

But it is not clear that we are closer now to an operational model for improving teamwork. There is a tension between trying to reduce the field to essential elements versus expanding the field to cater for vagaries in observed results in real-world contexts. We do not wish to debate the value or validity of the many classification schemes proposed in the literature. Rather, we wish to propose an alternative approach which may offer a more comprehensive resolution of the field and some conflicting findings, leading to a model of teamwork that can be both empirically supported and practically relevant.

The first step is to note that many frameworks have included some notion of a *sharing* of individual team members’ *mental models* containing knowledge of team- and task-related properties and processes and having considerable influence over team performance through increasing the accuracy of expectations of team needs and anticipation of the actions of team members (Cannon-Bowers et al., 1990). But current research is not conclusive with burgeoning questions concerning with the content, structure, development and influence of shared mental models on team processes and individual/team performance. Much debate has arisen as attempts to empirically test and apply models and frameworks have produced mixed, even conflicting, results. For example:

- Espinosa & Carley (2001) found some task coordination was influenced by supposed team-related shared mental models. Langan-Fox et al. (2004) found a lack of both team- and task-related sharing of mental models was associated with poor coordination.
- Lim & Klein (2006) found both task- and team-related sharing of mental models improved team performance, with team-related being more important. Levesque et al. (2001) and Mathieu et al. (2000) found that team-related sharing of mental models, but not task-related, was closely related to team performance. But Mathieu et al. (2005) contradicted their previous findings, finding that it was task-related, not team-related, sharing of mental models that related to team performance. And Lee (2007) found neither task- nor team-related sharing of mental models significantly affected team performance.
- Espinosa & Carley (2001) found increases over time in sharing of task-related mental models while Levesque et al. (2001) found decreases instead. Lee (2007) found individual performance overall decreased over time as the sharing of team-related mental model ‘structure’ increased, although the detailed findings were interpreted in favour of starting a project with a very high degree of team-related model ‘structure’ exchange, then switching focus to increasing task-related exchange as the project progresses.

In the field of adaptive automation, Parasuraman et al. (2000) had once considered tasks as decomposable into four abstract functions based on an assumed four-stage information-processing model of task performance, but later saw greater value in decomposing tasks hierarchically into alternative sets composed of sequences of specific subtasks and concrete activities (Miller & Parasuraman, 2007). Similarly, rather than advocate for explanatory frameworks based on general phases or categorisations, we consider tasks and teams as organisations of processes and examine in detail their process components and their interactive characteristics and constraints. Central to our approach is a delineation of the function and content of mental models and how such mental models, when extended from the individual into the team domain, both enable and constrain functional team organisations. This allows a different approach to the analysis of team performance. Rather than seeking to explain team performance in terms of categorised factors, we seek to show how team performance emerges from the interaction of multilayered organisations of team and task processes with mental models at the crux.

Components of Task and Team Organisations

Miller & Parasuraman (2007) define a *task* as “an encapsulated set of behaviours that constitute a known method of accomplishing a domain goal” (p. 66). A complete task model is formed by functionally decomposing a task into related subtasks, each subtask associating a subgoal with one or more alternate methods of realisation. These methods are themselves subtasks forming a hierarchy descending through partial plans at higher levels to ‘primitive actions’ at the lowest level of resolution required for the purpose of the task model. The notion that a goal may have alternative methods of realisation (Brazier et al., 2000; Miller & Parasuraman, 2007; Seibt, 2009) will be crucial to discussion below.

Salas et al. (1992) define a team as “two or more people who interact dynamically, interdependently and adaptively toward a common and valued goal, who have been assigned specific roles or functions to perform, and who have a limited life-span of membership” (p. 4). Terrell (2006) is more succinct: “The very definition of teamwork refers to the activities of groups whose members work towards a common goal” (p. 123; see also Andersen, 2001; Langan-Fox et al., 2004; Yen et al., 2001). A comprehensive framework for modelling organisations (including teams) and analysing their operational dynamics is that of Sharpanskykh (2008; see also Jonker et al., 2007; Popova & Sharpanskykh, 2007; Stroeve et al., 2007). In this framework, four aspects of an organisation are defined and interrelated: performance-oriented, process-oriented, organisation-oriented and agent-oriented. This detailed framework is highly consistent with the team definition of Salas et al. (ibid.).

Because the term *agent* is sometimes taken to refer to an artificially intelligent software stand-in or assistant for a human operator, we will instead use the term *actor* to refer to one who performs some purposeful action. An actor could be a human, animal or synthetic entity. The seminal definition of teams by Salas et al. (ibid.) suggests that only humans can qualify as team members. We accord with more recent definitions (eg. Christoffersen & Woods, 2002; Cohen et al., 1997) in accepting the possibility that “information technology can function as a respected decision maker within the team” (Terrell, 2006, p. 115). This position will be supported in later discussion.

Shared Mental Models, Team Mental Models, and Team Cognition

According to Rouse & Morris (1986), mental models serve to generate: (1) descriptions of system purpose and form; (2) explanations of system functioning and observed system states; and (3) predictions of future system states. But what does it mean for two or more members of a team to 'share' such a thing? If we view the team itself as a higher-order system, what form could the team's mental model take and where could it reside?

The literature can be divided into two groups distinguished by the dependence or independence of an individual's mental model. In one group, each individual's model is independent; no external system is required for the model to fully function. But in the other group, the functionality of each individual's model depends to some extent on an external system. So the former group conceives of sharing in terms of similarity or aggregation. A shared mental model could be construed as:

- the union of knowledge and belief held in common by at least two members;
- the intersection of knowledge and belief held in common by all members;
- the sum of that held by all members (regardless of commonality);
- the average of that held by all members.

The former two treat mental models as sets, the latter two treat mental models as vectors. The above list is not exhaustive and there is ongoing debate over the relative merits of various points in this space (eg. Espinosa & Carley, 2001; Mohammed & Dumville, 2001; Resick, 2004; Terrell, 2006). However, for the purposes of this paper we will not be exploring it any further.

The other group, adopting a *distributed cognition* approach (Hutchins, 1991), conceives of *team cognition* as a higher-order construct that is generated by, and may affect, the cognition of its members. This group subdivides further on the notion of what is shared. For one subdivision, team cognition is about processes that let members use the environment, artifacts and other members as cognitive resources. Members thus 'share out' the burden of cognition; part of *my* mental model can reside in *your* mind to be accessed when required (eg. Tollefsen, 2006). But for the other subdivision, team cognition is a multi-level organisation of processes that generate a virtual model which does not itself reside in any member but enables and constrains activity such that each member has a 'share' in the larger process (eg. Bosse et al., 2005). This is very much a hierarchy- and process-oriented view.

For completeness, it must be said that some authors distinguish between: (1) *shared mental model* as "a team-level well-constructed cognitive structure... created by team members' intellectual efforts combining into a cognitive unit"; and (2) *team cognition* as "the broadest concept that focuses on the entire team's cognitive processes involving shared mental models, shared knowledge, and shared understanding; the whole process of team thinking" (Lee, 2007, p. 12). This creates a clear distinction between lower-order structure (shared mental model) and higher-order process (team cognition). Although such a distinction may appeal, it also requires commitment to a similarity-based conception of shared mental model: the 'shared' part of the model must be replicated in each team member's mind. But how much can be shared in this way? And how much should be shared?

A 'strong' concept in this vein would define the shared mental model as that knowledge which is common to all members (Mohammed & Dumville, 2001). But this is disputed (Espinosa & Carley, 2001) and highly unlikely to be the case for action teams. Tasks typically assigned to action teams will be constituted in specialised subtasks encapsulated by roles which are in turn assigned to members with specialised capabilities (Jonker et al., 2007; Popova & Sharpanskykh, 2007), so each member may know only that part of the model with subtasks encapsulated by their role(s). Action teams are problematic for those who take a position that a shared mental model is a body of knowledge common to team members such that the more is shared, the more effective the team.

It is more pragmatic to concede that there may be no knowledge of task or team that is common to all members, and pertinent information about the current state of a team-on-task is not always known by all members of a team. Yet even so, 'weak' models have been shown not to handicap teams and may even be advantageous in some contexts. Terrell (2006) found that in 911 Emergency Response teams, it was "not necessary for all members of the dispatch team ... to fully understand the entirety of an emergency situation in order to accurately respond to an emergency" (p. 118). "It seems to be that the team mental model is only shared by some members of the team, i.e. the members who are ultimately responsible for consolidating the information, dispatching the resources, and relaying the consolidated information to the first responders who are also members of the meta-level emergency response team" (p. 115). Note also that in this case, model sharing occurred *outside* the team.

Such findings invite questions about the content of shared mental models, and about what kind and degree of overlap between a model's content held by different team members is beneficial for team performance and in what contexts; in short, questions of representational content. Consideration of representational content is commonally restricted to internal mental states of an individual cognitive agent. This would be the case when considering the content of an individual team member's privately held mental model. But what of the content of a mental model which, being shared among multiple teammates, can have no monolithic structure?

Key to resolving these questions is to accept that a team may have properties that are not reducible to independent properties of individual team members. If those team-level properties include mental states, these constitute a *collective mind* (Wilson, 2004) or *shared extended mind* (Bosse et al., 2005). The notion of shared extended mind allows *collective* representational content for shared extended (team-level) mental states by extending notions of representational content in two ways: (1) for external, rather than internal, mental states; and (2) for groups of individuals rather than single individuals (Bosse et al., 2006a). Collective representational content for shared external mental state properties can be formulated as **hierarchical relations between states of team members and the environment over time**. To describe and relate mental models at individual and team levels, explicit reference is made to time. Dynamic properties can be formulated that relate a state at one point in time to a state at another point in time. More broadly, the occurrence of a mental state property can be related to sets of past and future interaction traces. In this process-hierarchy view, **these relations are the constituents of shared mental models**. They need not be explicit nor fully contained in the mind of any team member. And it is just such relations that are required to fulfil the functions of mental models as originally proposed by Rouse & Morris (1986).

Constitution and Function of Individual Mental Models

Mental models must contain relatively stable knowledge and beliefs concerning team and task organisation and processes. They must also capture the momentary transit of any number of active threads (Stojanov & Kulakov, 2003) in support of managing and directing an unfolding team-on-task situation. As stated previously, we argue that mental models must therefore contain a hierarchy of relations between state properties over space and time. This section expands the argument and begins to explore the implications for teamwork.

In accordance with an interactivist approach, we propose that task and team mental models are actually examples of a more general *situation knowledge web* construct of potential interactions (Bickhard, 1980, 2005, 2008; Bickhard & Richie, 1983; Bickhard & Terveen, 1995). The fundamental unit is a pairing of action with potential outcome. Through an action-outcome pair, a situation may be classified in simplest terms as affording either a successful interaction (benefit realised) or failure (benefit not realised).[2] Some actions performed in some contexts lead to beneficial outcomes for the actor; other actions, or the same actions in other contexts, lead to detrimental outcomes. For some actions, there may be multiple possible outcomes of which only some are beneficial. Whichever outcome is realised may depend upon factors in the current situation arrived at through a (potentially very long) sequence of overlapping outcome-action-outcome triads. For our discussion, it does not matter where the influencing factors lie, in the actor or the environment, or both; what matters is that a beneficial outcome will be realised in only some, not all, situations. Through the ability of an actor to perform an action and perceive whether the anticipated outcome is realised or not, action-outcome pairs thus serve to differentiate situational classes.[3] Situation knowledge webs are in practice highly recursively interwoven so that the success (or failure) of one action-outcome pair—that is, one interaction—may indicate many other possibilities for further interaction:

> In still more complex organisms, there may be vast webs of indications of interactive potentialities, with some of them indicating the potentialities of still others, should those first interactions be engaged in and proceed as anticipated. These webs constitute the organism's

---

[2] It should be noted that we are using 'benefit' in a sufficiently general sense so as to include the *avoidance* of some detrimental happening. For brevity in the following discussion, unless otherwise specified, the unqualified term 'outcome' should be understood to mean 'beneficial outcome'.

[3] The problem of correctly identifying causal relationships between actions and outcomes—"Did I do that?"—is of course crucial but outside the focus of this paper because humans engaged in the kind of teamwork under discussion tend to have fewer difficulties determining causality than do synthetic agents. The web of causal relationships is worked out by the human designers of a teamwork support system and embedded in the system's operating procedures. More will be said below.

> knowledge of its current environment, organized in terms of how some interactive possibilities could be reached via various intermediary interactions. This web must be updated and continuously maintained. Parts and aspects of it will change with various interactions of the organism, and other changes will occur whether or not the organism engages in particular interactions. The process of maintaining the web of indications of interactive possibilities is that of *apperception* (Bickhard, 2005, pp. 207–208, emphasis added).

If there is no possibility of any detrimental outcome (including loss of non-disposable resources) then an actor has nothing to lose by trying an action. However, it does the actor no good to determine that it was unsafe to perform an action by performing that action! A risky situation for an actor can be characterised by three properties: (i) an action could have either a beneficial or detrimental outcome; (ii) the actor intends to realise the beneficial outcome; and (iii) the beneficial outcome cannot (to the actor's knowledge at that moment) be realised through alternative, safer actions. It becomes vital for an actor to somehow determine that the current situation is a member of the beneficial-outcome class, as opposed to the detrimental-outcome alternative, before performing the action.[4]

Bifurcation of the situation knowledge web is the outcome of a process by which an actor uses an action-outcome pair to differentiate a situational class. Bifurcations of the situation knowledge web are behind sentences that begin, “Had I known then what I know now…” Consider the simplest possible case as shown in Figure 1 below. At time (t=-2) the actor projected the performance of a particular action at time (t=-1). Proceeding from that action were two parallel arcs to two mutually exclusive possible outcomes[5] at time (t= 0). But which outcome would be realised? The actor's situation at time (t=-2) was undifferentiated with respect to the two outcomes in that the actor could not predict which outcome would be realised at time (t= 0). At time (t=-1) the actor performed the action. By time (t=0) there was no longer any question. One of the outcomes, either A or ¬A, was realised and the actor was able to bifurcate its situation knowledge web to determine which class of situation held at time (t=-2). That is, at time (t=0) the actor had the hindsight to know which situation it was in at time (t=-2): either a situation leading to particular outcome A, given a particular action; or a situation leading to some outcome other than A (*i.e.* ¬A) given that same action. Figure 1 below shows the stages of the bifurcation process from the outcome realisation back to the identification of the situational class that was.

---

[4] In practice, there could be any number of actions that the actor may perform. However, whether there be one action or many, the principle remains that a wise actor would determine the class of situation before acting.

[5] More precisely: mutually exclusive sets of possible outcomes functionally differentiated with respect to a relevant performance indicator.

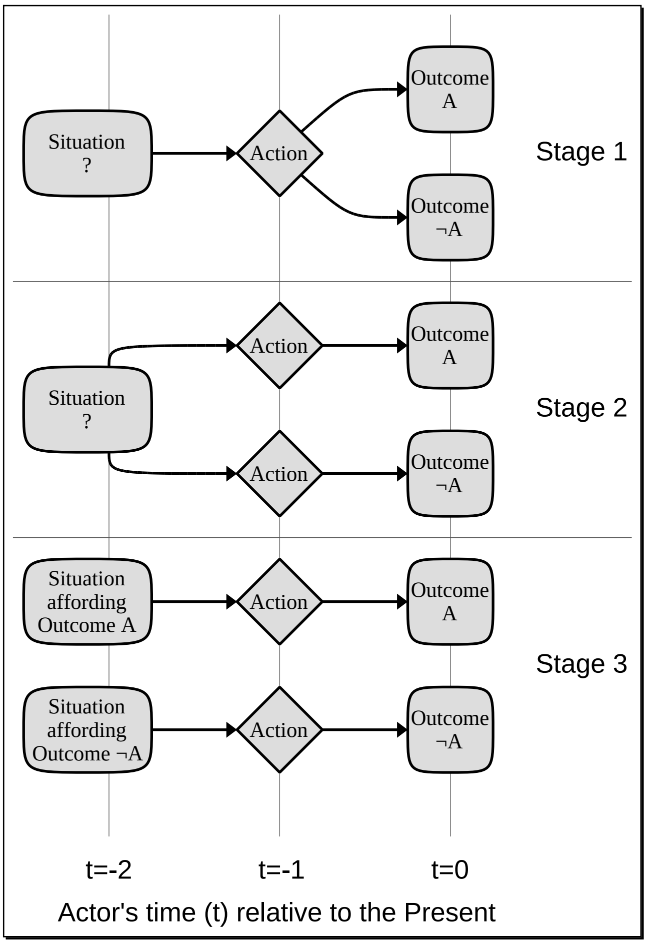

Figure 1: The process of situation knowledge web bifurcation by which situations are differentiated according to the outcomes of actions.

Suppose that Outcome A in Figure 1 above is a beneficial outcome while Outcome ¬A is a detrimental alternative. The actor then needed to determine at time (t=-2) if it was safe to perform the Action at time (t=-1). This could have been achieved if the actor could have identified some other action—with multiple possible outcomes but none detrimental—to perform between times (t=-2) and (t=-1) such that the outcome realised would bifurcate the situation knowledge web. A stage in this process of *active sensing* is shown in Figure 2 below.

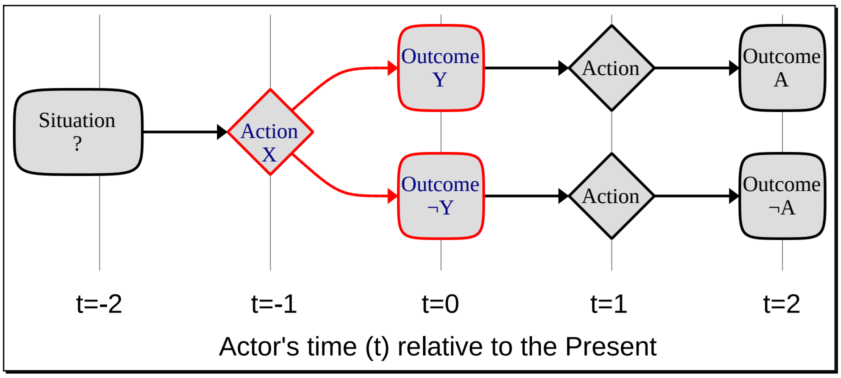

Figure 2: 'Active sensing' whereby an action is performed not for the outcome *per se* but for what the outcome indicates about other interactions.

The actor in Figure 2 chose to perform Action X not because either Outcome Y or ¬Y is a beneficial end in itself but because of what the realised outcome indicates for further action-outcomes. This is different from either passively waiting for whatever sensory input arrives or applying some filter to all sensory input received. Rather, the actor engages in *probing* activity that is intended to elicit or invite a particular class of response from the environment whereby the absence of stimulation may be just as telling as its presence. In principle, bifurcation could be continued indefinitely (cognitive resources permitting) into the actor's history of interactions; but in practice this does not happen. For the actor to predict the realisation of either Outcome A or ¬A, it is sufficient merely to stop bifurcation at Action X. Bifurcation is cognitively expensive and therefore the resolution of present and past situations is performed only as much as is necessary for the resolution of future interaction possibilities.

Situation knowledge web architectures and implementations have received extensive coverage elsewhere (eg. Kulakov et al., 2002; Stojanov & Kulakov, 2003, 2011). The depictions that follow in this paper have been greatly simplified in order to highlight certain properties as clearly as possible for the purposes of this discussion. Furthermore, for the purposes of this paper, the situation knowledge web is a psychological construct that serves functions for anticipation, action-selection and error-correction. We are making no claim to its existence in any particular neurophysiological structure. The depiction of situation knowledge webs herein should not be taken as a guide to system implementation. (But see Mareschal et al., 2007, for proposals concerning action-based knowledge representations which are neurological, partial, distributed and context-dependent.) Readers who are familiar with traditional rule-based reasoning systems should note that to encode a situation knowledge web would require more complex rule expressions than is traditional. The essential difference is illustrated in the following table. In the simplest case, a traditional condition-action rule will be bracketed by a preceding action and an expected post-action condition:

Traditional	Situation Knowledge Web
If condition C holds, Then perform action A.	If action X was performed Resulting in condition C holding, Then perform action A Expecting condition Y to hold.

Table 1: Traditional rule-based systems perform actions conditional on immediate system states. But to take full advantage of situation knowledge webs for reasoning, systems must be able to relate actions-with-outcomes in hierarchical and temporal dimensions.

Extending Mental Models into Teams

The situation knowledge web provides the substrate required to create mental models that can extend into the team domain as a shared extended mind.

The following discussion refers to a team-on-task specification, depicted in Figure 3 below, which has two actors (i.e. team members) performing three subtasks each having two alternate performance methods (i.e. the actions). The intervals between time points are unspecified; an interval may be zero.

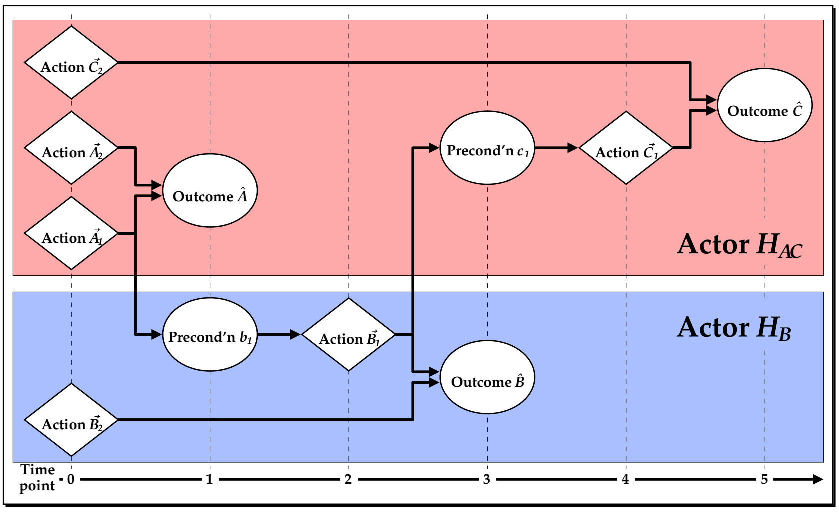

Figure 3: Alternate methods of satisfying goals which also set beneficial conditions for other goals’ methods.

Let $\langle \hat{A}, \{\langle \vec{A}_1, \varnothing \rangle, \langle \vec{A}_2, \varnothing \rangle\} \rangle$ represent a task specification indicating that outcome $\hat{A}$ can be realised by performing either action $\vec{A}_1$ or action $\vec{A}_2$, neither having any precondition for success.

Let $\langle \hat{B}, \{\langle \vec{B}_1, b_1 \rangle, \langle \vec{B}_2, \varnothing \rangle\} \rangle$ represent a task specification indicating that outcome $\hat{B}$ can be realised by performing either action $\vec{B}_1$ when precondition b_1 is satisfied, or action $\vec{B}_2$ which has no precondition for success.

Let $\langle \hat{C}, \{\langle \vec{C}_1, c_1 \rangle, \langle \vec{C}_2, \varnothing \rangle\} \rangle$ represent a task specification indicating that outcome $\hat{C}$ can be realised by performing either action $\vec{C}_1$ when precondition c_1 is satisfied, or action $\vec{C}_2$ which has no precondition for success.

Outcomes $\hat{A}$, $\hat{B}$ and $\hat{C}$ are subtask goals. (Associated performance indicators and higher-order goals are not essential to the argument and so are not depicted.) Outcomes $\hat{A}$ and $\hat{C}$ are assigned to one actor H_{AC}, and outcome $\hat{B}$ is assigned to another actor H_B.

Teamwork as the selection of *mutually beneficial methods*

Given the above team-on-task specification, suppose that the performance of action $\vec{A}_1$ also creates a situation that satisfies precondition b_1; likewise $\vec{B}_1$ and c_1: $(\vec{A}_1 \vDash b_1)$ and $(\vec{B}_1 \vDash c_1)$. Then in choosing action $\vec{A}_1$ to realise outcome $\hat{A}$, actor H_{AC} also makes it possible for actor H_B to choose action $\vec{B}_1$ to realise outcome $\hat{B}$. And in choosing action $\vec{B}_1$ to realise outcome $\hat{B}$, actor H_B also makes it possible for actor H_{AC} to choose action $\vec{C}_1$ to realise outcome $\hat{C}$.

Under these conditions, we say that $\langle \vec{A}_1, \vec{B}_1, \vec{C}_1 \rangle$ forms a superprocess composed of mutually beneficial subprocesses $\langle \vec{A}_1, \vec{C}_1 \rangle$ and $\langle \vec{B}_1 \rangle$. And should this superprocess be performed, we expect that actors H_{AC} and H_B will respond positively to the question, "Did you work as a team?" We propose that it is the performance of supertasks through the selection of mutually beneficial subtask methods that is the essence of *teamwork*. Construction of such methods occurs at higher levels of teamwork.

How a shared mental model enables teamwork

We are now in a position to clearly indicate the minimum shared content required of mental models to support teamwork as we have defined it. That content is represented in Figure 1 by elements b_1 and c_1. Although it appears in the figure that b_1 and c_1 are each contained within a single individual model, each element serves as a common factor relating two outcome-action-outcome triads and, through those triads, the two actors' individual models. This is how the models are 'shared' (in a way that is more than by common knowledge).

Note that the elements occupy different relational positions in each model: element b_1 is an outcome on the one hand and a precondition on the other for actors H_{AC} and H_B respectively; likewise c_1 but reversed. Taking b_1 for example:

- $(\vec{A}_1 \vDash b_1)$, meaning b_1 is an outcome of method $\vec{A}_1$ for actor H_{AC} ; and
- $\langle \vec{B}_1, b_1 \rangle$, meaning b_1 is a precondition for the success of method $\vec{B}_1$ for actor H_B.

These two expressions taken together capture the beneficial relationship between the two actors' subtasks.

Even if actor H_{AC} is unaware of the content of b_1, **the relationship itself** provides actor H_{AC} with a rational basis for selection of action $\vec{A}_1$ over alternatives. In our view, instances of such selections exemplify what Salas et al. (2005) describe as team orientation enabled by shared mental models and mutual trust.

This is a very simple example. When several actors are involved, it may be that no two actors engage in activities that directly benefit each other but there is a complex transitive chain or web such that each actor benefits, or benefits from, some activity of some other actor.

Teamwork, teams and *being-in-a-team*

Having defined what we claim is the essential character of a *teamwork* process, we can now define a *team* as a group of two or more actors which, at least some of the time, engage in what we have defined as *teamwork* processes. In other words, if there is a group who at least sometimes engages in teamwork, then one can justifiably call that group a team. The classical requirements for teams and teamwork outlined above—for example, that there be a common goal—are subsumed in our definition of teamwork by the hierarchical task organisation and its relation to roles, goals and performance indicators. But rather than follow the same route as classical definitions which begin with a definition of teams then proceed to define teamwork as what teams do, we have begun with a definition of teamwork then proceeded to define a team as the thing that does teamwork. This reversal of the relationship between the terms will produce insights on team performance in the discussion that follows. But let us highlight that our conception of teamwork makes a significant contribution to the literature by strengthening and clarifying the classical requirement for interdependence.

It may be that an actor need only believe that another actor chooses to act in a way that benefits the former for that actor to experience *being-in-a-team*. The belief may not be true. The other actor may have perceived no choice in acting as they did. Or there may have been other methods, but none as personally appealing as the method chosen. Perhaps what matters is not that there were indeed better alternative methods, but that the receiver of the benefit believes that the other chose to act with the provision of that benefit in mind.

Some Implications for Teamwork Studies

In the sections that follow, a few illustrative aspects of teamwork will be considered in light of the approach presented above. Some areas of focus for future work will be sketched.

Fluent multi-actor coordination through *situation/action recognition*

When two or more actors are coordinating the execution of their respective subtasks, each actor must recognise the appropriate moment at which to perform an appropriate action. This section considers the consequences of this recognition requirement. From the preceding exposition of the situation knowledge web concept, it becomes possible to distinguish between several classes of coordination mechanism and, for each class, the prerequisite knowledge each actor must have of the others in relation to the task.

The following series of figures depict several cases of a scenario in which two actors are coordinating their actions to realise a common goal. (The graphical notation is identical to that of Figure 3 above but for the labelling of elements which has been simplified.) For illustrative purposes, the scenario is a very simple one in which a common goal can be realised with the successful performance of just two actions in sequence. Each actor is capable of performing one of the required actions. Extended versions of such scenarios are very common, from factory assembly lines to racket-sport doubles games. Although simple, each case illustrates a distinct combination of coordination mechanism and prerequisite knowledge. In particular, it will be shown what prerequisite knowledge each actor requires of the other in order to enable coordination in each case.

All the following figures conform to a format showing a sequence of events comprised of actions and outcomes where each action is performed once by one actor. Time is shown progressing as a series of steps, but this should not be taken to express a requirement that time be discrete nor that the durations between time steps be equal. In all cases, the task progresses through a series of situations from an initial situation A to a final situation A+B+C that satisfies the task goal. The transition from situation A to situation A+B occurs through the competent performance of action I by actor 1. The transition from situation A+B to situation A+B+C occurs through the competent performance of action J by actor 2.

The first case depicted in Figure 2 shows the second actor performing action J after it recognises that outcome B has occurred. For the second actor, so long as outcome B is recognisable and guarantees the successful transition to C through J, then no knowledge of the first actor’s role is required. Nor is any knowledge of any predecing situations required. Before perceiving B, all the second actor need know is that B will indicate that J should be performed to realise C.

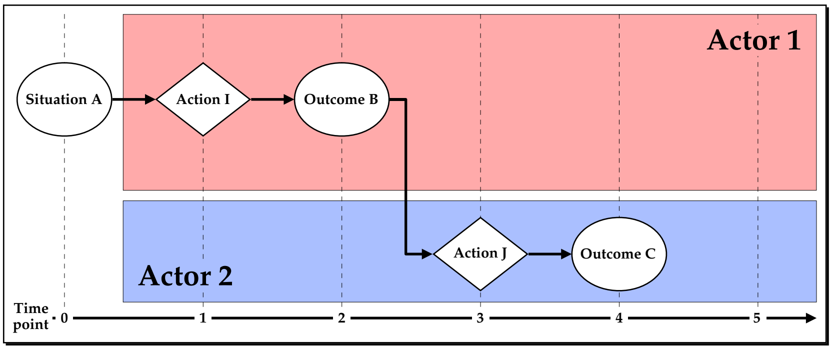

Figure 4: Action selection following decontextualised situation recognition.

The next case depicted in Figure 3 shows the second actor performing action J after it recognises that outcome B has occurred **as a result of action I**. In other words, the actor knows that it is the performance of

I realising B that indicates the possibility of performing J to realise C. In this case, to make the decision to perform J, the second actor must be able to recognise not only outcome B but also the first actor performing action I. So if the second actor cannot recognise the performance of I by another actor, then the ability to recognise B alone will not induce the actor to perform J. Such contextualisation is demanding and offers no advantage in this simplistic task scenario, but would be essential in a more complicated task scenario in which B could also occur through other means in situations not related (or even counter) to achieving the goal via J. As the need for such contextualisation increases, the cost to the second actor is a task knowledge web increasingly expanded and complicated with actions and outcomes that are not in that actor's repertoire.

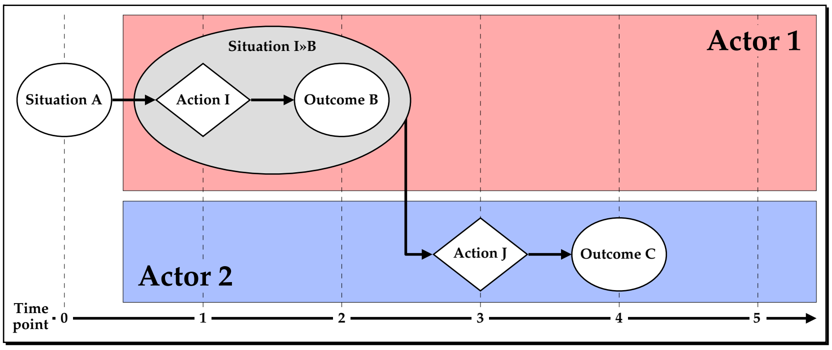

Figure 5: Action selection following contextualised situation recognition.

The preceding cases have required explicit situation recognition. However, there is another possibility: implicit situation recognition through action recognition. This capability is key to improving team performance through what Hoffman & Breazeal (2008) call *fluency*. As in the preceding case, it requires an actor to be capable of recognising (which presupposes perceiving) the performance of actions by other actors. However, it does not necessarily require the ability to recognise the outcomes of those actions (or situations in general). Recognising an action may be significantly less costly than recognising a situation and may also allow for decreased latency through anticipation of the outcome (Dousson et al., 1993; Hoffman & Breazeal, ibid.; Knoblich & Sebanz, 2008). Figure 4 depicts the case in which the second actor will start to perform action J when it recognises that the first actor is performing I. This is all the prerequisite knowledge that the second actor needs in this case. This knowledge presupposes that the first actor will perform action I when it recognises a situation that will afford outcome B, thereby realising B and allowing action J to realise outcome C.

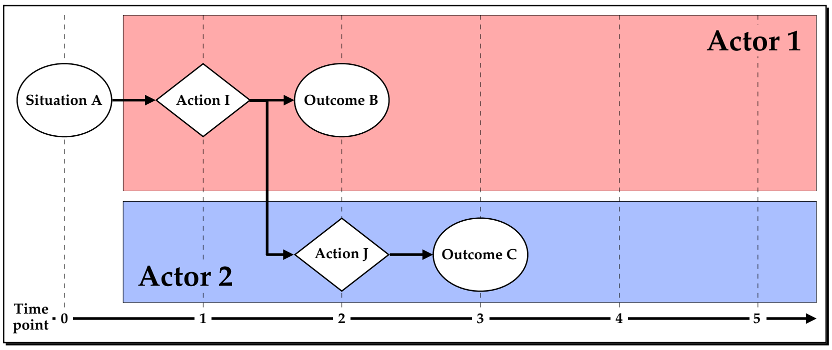

Figure 6: Action selection following decontextualised action recognition.

As with the transition from the first to second cases (Figure 2 to Figure 3) contextualisation is possible. That is, the task knowledge that is presupposed can be made explicit if the first actor may perform action I in other situations or with other possible outcomes. Figure 5 depicts such a case. In such cases, and in relation to cases exemplified by Figure 4, the ability to recognise situations will need to be reintroduced and latencies may increase again (when outcomes must be awaited). The difference between these cases and those exemplified by Figure 3 is subtle; it lies in whether the actor is attending preferentially to either other actors' behavioural cues or situational cues.

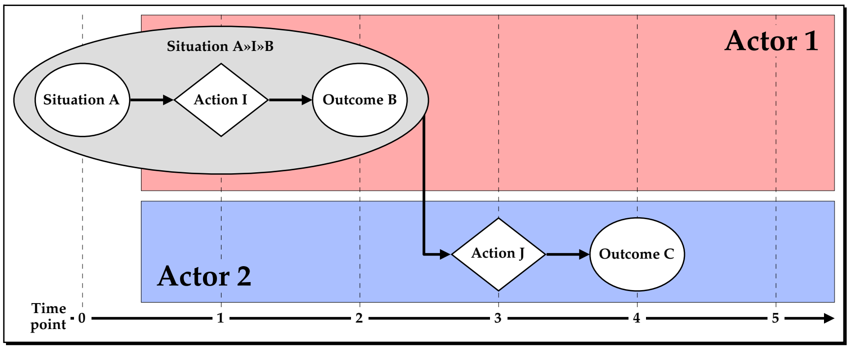

Figure 7: Action selection following contextualised action recognition.

Action recognition allows a special case depicted in Figure 6. It allows the first actor to relieve the second actor of some of the recognition load and guesswork. The first actor must already be capable of recognising a situation that affords performing action I to realise outcome B. Should that situation arise and action I then be performed with outcome B realised, rather than leave it to the second actor to recognise this, the actors could have established a prior convention that the first actor will perform an action ∅ as a signal to the second actor. There need be no significant outcome of ∅. All that is required of ∅ is the following criteria:

- Its meaning in relation to the task is already agreed, or there is a process by which it can be agreed. That is, it constitutes a *situation convention* (Bickhard, 2008).
- For the first actor, it is economical to perform.

- For the second actor, it is more recognisable than the equivalent combination of A, I and B.

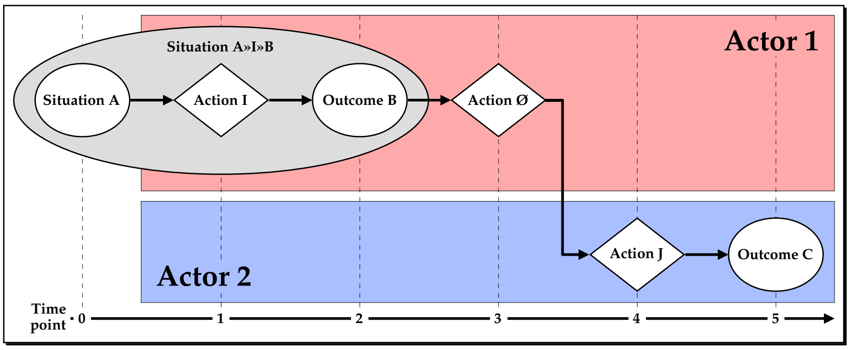

Figure 8: Using action recognition to minimise aggregate situation recognition costs.

Formal expressions of situation/action recognition

This section includes formal expressions in simplified Temporal Trace Logic (TTL) (Sharpanskykh, 2008) corresponding to each case in the preceding section.[6] The ability to express teamwork relations in this way has great potential for experimentation. One can imagine populating a database of logged events, then formulating queries to detect examples of teamwork for analysis.

With reference to Figure 2, the following TTL expressions capture the actors' task knowledge webs:

$$\begin{aligned} \text{Actor 1:} \quad & \text{state}\left(t,\ \text{input}\left(a_1\right)\right) \models \text{recognises}\left(A\right) \\ \Rightarrow\ & \text{state}\left(t+1,\ \text{output}\left(a_1\right)\right) \models \text{perform}\left(I\right) \\ \text{Actor 2:} \quad & \text{state}\left(t,\ \text{input}\left(a_2\right)\right) \models \text{recognises}\left(B\right) \\ \Rightarrow\ & \text{state}\left(t+1,\ \text{output}\left(a_2\right)\right) \models \text{perform}\left(J\right) \end{aligned} \tag{1}$$

Note that because this case presupposes that action I will produce outcome B and action J will produce outcome C, so these action-outcome dyads are not included in the actors' task knowledge webs. In other words, so long as the environment ensures that these action-outcome dyads hold true, actors need not hold internal representations of them.

With reference to Figure 3, the following TTL expressions capture the actors' task knowledge webs:

$$\text{Actor 1:} \quad \text{state}\left(t,\ \text{input}\left(a_1\right)\right) \models \text{recognises}\left(A\right) \tag{2}$$

[6] For a comprehensive introduction to Temporal Trace Logic (TTL), see Sharpanskykh (2008); also Bosse et al. (2006b). For the sake of brevity, we omit trace parameters and arguments. Also omitted are expressions that indicate actor's recognition of the goal outcome C, performance indicators, and operations performed by the actors' environment. These omissions have been made in order to highlight the key differences between the cases. A complete description of any non-trivial system would require many more expressions than it would be appropriate to show here.

$$\Rightarrow \text{state}\left(t + 1, \text{output}\left(a_1\right)\right) \models \text{perform}\left(I\right)$$

$$\text{Actor 2:} \quad \text{state}\left(t, \text{input}\left(a_2\right)\right) \models \text{observes}\left(I\right)$$
$$\wedge\, \text{state}\left(t + 1, \text{input}\left(a_2\right)\right) \models \text{recognises}\left(B\right)$$
$$\Rightarrow \text{state}\left(t + 2, \text{output}\left(a_2\right)\right) \models \text{perform}\left(J\right)$$

Note that in this case, actor 1's task knowledge web remains the same and includes nothing of actor 2's role or repertoire. However, actor 2's web now explicitly includes an observational reference to an action (I) in actor 1's repertoire.[7] It is the cost of evaluating the *observes* predicate that establishes the crucial difference between this case and the simpler one.

The following expressions capture the key features of the webs in Figure 4:

$$\text{Actor 1:} \quad \text{state}\left(t, \text{input}\left(a_1\right)\right) \models \text{recognises}\left(A\right) \tag{3}$$
$$\Rightarrow \text{state}\left(t + 1, \text{output}\left(a_1\right)\right) \models \text{perform}\left(I\right)$$
$$\text{Actor 2:} \quad \text{state}\left(t, \text{input}\left(a_2\right)\right) \models \text{observes}\left(I\right)$$
$$\Rightarrow \text{state}\left(t + 1, \text{output}\left(a_2\right)\right) \models \text{perform}\left(J\right)$$

The essential difference between these webs and those in Figure 2 may appear to be merely the substitution for actor 2 of *recognises* by *observes*. However, remember that with this substitution comes an explicit reference to other actors' behaviour. Thus the decision to employ situation recognition or action recognition hinges on the cost difference between a socially-blind perception of the state of the environment versus a socially-aware perception of the behaviour of another actor. (There is also the consideration due to reduced latency, although that is not clear in the TTL.)

The TTL for Figure 5 indicates that once actors need to represent extensive webs of other actors, there is no clear distinction between situation recognition and action recognition; both are needed:

$$\text{Actor 1:} \quad \text{state}\left(t, \text{input}\left(a_1\right)\right) \models \text{recognises}\left(A\right) \tag{4}$$
$$\Rightarrow \text{state}\left(t + 1, \text{output}\left(a_1\right)\right) \models \text{perform}\left(I\right)$$
$$\text{Actor 2:} \quad \text{state}\left(t, \text{input}\left(a_2\right)\right) \models \text{recognises}\left(A\right)$$
$$\wedge\, \text{state}\left(t + 1, \text{input}\left(a_2\right)\right) \models \text{observes}\left(I\right)$$
$$\wedge\, \text{state}\left(t + 2, \text{input}\left(a_2\right)\right) \models \text{recognises}\left(B\right)$$
$$\Rightarrow \text{state}\left(t + 3, \text{output}\left(a_2\right)\right) \models \text{perform}\left(J\right)$$

However, the TTL for Figure 6 indicates the significant efficiencies possible at the team level with the introduction of situation conventions (Bickhard, 2008) that can indicate and stand-in for extensive webs:

$$\text{Actor 1:} \quad \text{state}\left(t, \text{input}\left(a_1\right)\right) \models \text{recognises}\left(A\right) \tag{5}$$

[7] In a more fully specified system description, the *observes* and *recognises* predicates would include parameters specifying the actor in question.

$$\Rightarrow \text{state}\left(t+1, \text{output}\left(a_1\right)\right) \models \text{perform}\left(I\right)$$

$$\text{state}\left(t, \text{input}\left(a_1\right)\right) \models \text{recognises}\left(B\right)$$
$$\Rightarrow \text{state}\left(t+1, \text{output}\left(a_1\right)\right) \models \text{perform}\left(\varnothing\right)$$

$$\text{Actor 2:} \quad \text{state}\left(t, \text{input}\left(a_2\right)\right) \models \text{observes}\left(\varnothing\right)$$

$$\Rightarrow \text{state}\left(t+1, \text{output}\left(a_2\right)\right) \models \text{perform}\left(J\right)$$

Receiving and computing the meaning of a brief, transient, fleeting communication—a signal—to indicate that a particular situation has arisen may be less costly than computing the same through situation recognition. This is especially so for distributed teams denied opportunities to exploit action recognition. Because situation knowledge must be continually updated, especially as plans go awry, efficient update mechanisms are highly desirable in rapidly changing complex environments. **Communication may therefore be highly functional as a stand-in for situation and action recognition.** In this sense, communication is less about information transfer and more about solving coordination problems (Bickhard, ibid.) and reducing cognitive load. In asynchronous teams, coordination occurs in part through communications embodied in the production and exchange of artifacts (eg. letter drops, emails).

However, the potential for aggregate savings may come at some personal cost to actor 1 with additions to the action repertoire and knowledge web. In particular, for the first time in the TTL expressions here, actor 1 must recognise outcome B, something which was previously only of express concern to actor 2 (if anyone). Furthermore, the additional action ∅ for actor 1 is unrelated to the task itself and only required for the sake of coordination with actor 2; that is, team-imposed overhead. To make this clear, consider how the TTL could be simplified to remove ∅ entirely from Equation 5 if the task was being performed by a single actor:

$$\text{Actor:} \quad \text{state}\left(t, \text{input}\left(a\right)\right) \models \text{recognises}\left(A\right) \tag{6}$$
$$\Rightarrow \text{state}\left(t+1, \text{output}\left(a\right)\right) \models \text{perform}\left(I\right)$$
$$\text{state}\left(t, \text{input}\left(a\right)\right) \models \text{recognises}\left(B\right)$$
$$\Rightarrow \text{state}\left(t+1, \text{output}\left(a\right)\right) \models \text{perform}\left(J\right)$$

Implicit goals, classes of error and team member trust

Note that in every case depicted in Figures 1 through 5, Actor 1 holds no representation of the goal outcome C and thus no representation of the goal situation A+B+C. It may seem clear, therefore, that in lacking a common goal the two actors do not form a team. Certainly the first four cases could represent scenarios in which Actor 2 simply recognises in the conditions created by Actor 1 an opportunity for personal advancement and seizes it. For example, Actor 2 could be a predator and Actor 1 the prey who makes itself vulnerable to attack.

But in the last case, the signalling behaviour ∅ of Actor 1 to Actor 2 changes the relationship between the two actors such that a parasitic interpretation is no longer possible. Yet Actor 1 still holds no representation of outcome C. Without holding such a representation, how can the two actors be considered to be working towards a common goal? The answer lies in attributing a presupposed common goal to Actor 1's signalling behaviour $\varnothing$ (there being no other reason for it). Instead of holding a representation of the goal, Actor 1 need only trust that Actor 2 (i) **does** hold a representation of a goal that is beneficial to both actors, and (ii) **will** act **competently** to realise that goal. In effect, in this last case there is an implied goal for Actor 1: whatever the goal is for Actor 2. Of course, such relationships are not sufficient to define teams. The point here is to demonstrate that explicitly held common goals are not necessary to produce team-like relationships. Given sufficient trust relationships, common goals can be presupposed. Furthermore, the trust relationships need not be explicitly represented but may also be presupposed in the interactions of the team members.

An expectation that a situation affords performing a particular method to achieve a particular outcome (or that the outcome will hold after the method is performed) can always turn out to be wrong in any given instance. Even if an actor does not know what to do about the fact that a method produced an unexpected outcome, it must at least be able to recognise that fact as necessary (though not sufficient) for adaptivity. Roda (2007) distinguishes between four classes of error which we have expanded to six in the context of teamwork. The six classes of error can be subdivided into two groups of three. The first group of three pertains to matters of situation recognition and method selection. The second group of three pertains to issues of method performance:

(1) The situational class may have been mistakenly identified.

(2) The method selected may not have been appropriate to the subtleties of the actual situation. That is, the identification of the situation is insufficiently refined; bifurcation is warranted.

(3) The method selected may in fact not be capable of achieving the outcome in any case.

(4) The method, albeit selected and activated, was not actually performed.

(5) The method was performed but with inadequate resources, skill or attention.

(6) The method appeared to be performed well but was in fact intended to fail (i.e., treachery).

Topics for future work include: processes for differentiating classes of error; detection not only of error but also the *source* of error; the mapping of classes of error to processes of error mitigation in teams; and the effect of error class on perceived degree of severity in breaches of trust.

"Weak" models do not necessarily handicap teams

The entire hierarchy of a task model need not be known, completely specified from top to bottom, by all the actors in a team (Cannon-Bowers et al., 1995; Entin & Serfaty, 1999; Miller & Parasuraman, 2007; Terrell, 2006) for that team to function well. Because parts of the hierarchy are encapsulated by roles which are in turn assigned to actors (Jonker et al., 2007; Popova & Sharpanskykh, 2007), by default each actor will know only that part of the model with subtasks encapsulated by their role(s). Yet it appears that such “weak” models do not necessarily handicap teams and may even be advantageous in some contexts providing they are “accurate” (Terrell, ibid.). Having said that, it invites questions about what kind and degree of overlap between the model’s parts held by different team members is beneficial for team performance and in what contexts. For example, automated decision support systems appear most useful when they have embedded as comprehensive a specification of the model as it is possible to construct (eg. Brazier et al., 2000; Miller & Parasuraman, ibid.; Terrell, ibid.). However, Terrell (ibid.) found that it was “not necessary for all members of the dispatch team at either [911 Emergency Response] centre to fully understand the entirety of an emergency situation in order to accurately respond to an emergency” (p. 118). This issue will be explored further below.

Team training and development

In reframing “skills” in terms of interdependent subprocesses, we can conceptually distinguish between those which do not require preconditions to be established by other subprocesses, and those which do. Those subprocesses (or parts thereof) which are performed by an individual actor, and which do not impinge on any other actor’s subprocesses, may be learned and practiced by that individual in isolation. Furthermore, subprocesses whose preconditions can be established by other means may also be practiced outside a team setting. But in cases where one subprocess either impinges another or establishes preconditions for another, these must be developed and practiced in a team setting.

But there is another important aspect of team functioning that must not be overlooked: *situation recognition* and *action recognition*. These relate not to the establishment of preconditions *per se*, but to the **recognition of the preconditions being established**. This recognition ability cannot be developed in isolation. Situational indicators of success will present themselves in contextualised variations, so teams will need exposure to a variety of contexts in order to learn to recognise the situational indicators in each context. Team members will also need to learn the idiosyncracies of their teammates’ performances. Teams that will rely heavily on action recognition will require more time practicing together, observing each other’s behaviour, than teams that will rely heavily on situation recognition. However, given the potential unboundedness of contextual variation in situations, those teams that learn to exploit action recognition may gain an advantage in dealing with a wider variety of contexts through being able to spread the situation recognition load across all team members. If one team member recognises a particular situational affordance, and through some action indicates that affordance to teammates, then those teammates are spared the cost of recognising the affordance for themselves.

Role emergence, team member awareness, and constraint of task evolution

If a novel task organisation must be formed to implement a solution to a novel problem, roles will not be well formed initially. Partially formed roles may exist if the team members have interacted before on other kinds of task. But a well formed role organisation will emerge with the task organisation out of regularities in the way that team members formulate, decompose and take on the performance of subtasks. The emergence of roles is exemplified by transitions from "Fred did x" to "x is something that Fred does" and "because Fred does x, Fred also does y." Note that the detection of such regularities requires some episodic memory and temporal reasoning capability. Note also that beliefs about team members' capabilities, such as "Fred is *good* at doing x," can also emerge in this way given that the requirement for performance indicators allowing such assessments may be satisfied by the same mechanisms allowing discrimination of beneficial and detrimental outcomes.

In the case that there is an established partial role organisation which does not fit what may be an optimal task organisation for solving the new problem, the risk is high that the established role organisation will act as a strong constraint on the evolution of the task, preventing that optimal solution from being found. Such is the risk when well-established military organisations are engaging in novel theatres and operations.

Emergence of team leadership functions

In general, the performance of some action will incur a cost on the performing actor; that is, each action has an associated cost function. Actions can be compared on the basis of their cost functions. All other things being equal, a rational actor will prefer a less costly action over one that is more costly but otherwise functionally equivalent. In this context, we can describe one aspect of *team leadership* as a function which seeks to minimize aggregate costs across actors and subtasks.

Referring to the general teamwork situation described above, suppose that $\text{cost}(\vec{A}_2) < \text{cost}(\vec{A}_1) < \text{cost}(\vec{B}_1) < \text{cost}(\vec{B}_2)$. In the absence of all other considerations, actor H_{AC} will prefer action $\vec{A}_2$ over action $\vec{A}_1$, while actor H_B will prefer action $\vec{B}_1$ over action $\vec{B}_2$. Recall that action $\vec{B}_1$ requires precondition b_1 for success and that b_1 is established by action $\vec{A}_1$. Should actor H_{AC} (quite rationally from a personal perspective) select the less costly action $\vec{A}_2$, the preconditions for action $\vec{B}_1$ will not be satisfied and actor H_B will be forced to select the more costly action $\vec{B}_2$. Now suppose that $\text{cost}(\vec{A}_1) + \text{cost}(\vec{B}_1) < \text{cost}(\vec{A}_2) + \text{cost}(\vec{B}_2)$. In this case, it is less costly from a systemic perspective to force actor H_{AC} to select action $\vec{A}_1$ (and bear the extra personal cost) so that actor H_B can select action $\vec{B}_1$ thus minimising the aggregate cost.

So this aspect of team leadership can be cast as a higher-order control function, setting environmental conditions so as to force detrimental subprocess selection in favour of beneficial superprocess outcomes. Note that this function can be performed by the actors immediately involved; an external agency may not be necessary. All that is required is for actor H_{AC} to have some knowledge of the relations highlighted above and repeated here:

- that $(\vec{A}_1 \vDash b_1)$; and

- that $\langle \vec{B}_1, b_1 \rangle$.

It can be seen in the figure that b_1 and $\vec{B}_1$ are external to the domain of H_{AC}, being surplus to the subtask organisation assigned to H_{AC}. In other words, H_{AC} doesn't need to know about these elements to fulfil the subtasks specifically assigned to it. Thus should H_{AC} hold some representation of these additional elements, it would exemplify what others have designated *team knowledge* as distinct from *task knowledge*, hence the tendency in the literature to designate "leadership" as a "team skill." But it is clear from the discussion herein that this function of leadership, at least, cannot be independent of task knowledge (cf. Salas et al., 2005).

However, it must be stressed that the appropriateness of such a selection presupposes all the other conditions prescribed for this example. In other contexts, selecting action $\vec{A}_1$ over the less costly $\vec{A}_2$ on the basis of these two facts alone **may not** be appropriate. For the purposes of discussion, we declared that $\text{cost}(\vec{A}_2) < \text{cost}(\vec{A}_1) < \text{cost}(\vec{B}_1) < \text{cost}(\vec{B}_2)$ and that $\text{cost}(\vec{A}_1) + \text{cost}(\vec{B}_1) < \text{cost}(\vec{A}_2) + \text{cost}(\vec{B}_2)$. But if it were not strictly so, then selection of $\vec{A}_1$ would be detrimental to both actor and system. For instance, what if action $\vec{B}_1$ was highly costly and would therefore never be selected by actor H_B in any case? It is therefore quite possible for one actor to select an action in the expectation of setting favourable conditions for other actors—i.e. being helpful—when in fact the sacrifice is at best in vain. In summary, very little team knowledge is required to enable higher-order control of method selections to be distributed among the team members themselves, but the benefits or detriments realised on that basis will be highly contextualised. This is not problematic when the context is favourable; that is, when the presupposed conditions do in fact hold. But detecting and addressing invalid presuppositions that lead to team dysfunction requires a higher class of team functioning to be addressed in future work.

Sharing of mental models between humans and non-human systems

When functionally grounded in actions and outcomes-leading-to-further-actions, the mental models of two co-actors need not be structurally nor mechanistically alike in order to permit successful interaction (Smart et al., 2009). Thus the finding by Levesque et al. (2001) that mental models in software development project teams diverged over time is not surprising from our perspective. Indeed, as roles and associated subtasks become more defined and differentiated, mental model divergence is to be expected. Some researchers have made observations leading them to accord non-human systems the ability to engage in some degree of teamwork:

> Information technology can function as a respected decision maker within the team... The function of the [Computer Aided Dispatch] CAD system in dispatch teams, for instance, is more than a vehicle for the storage and retrieval of information. Rather, it processes information and provides problem solutions depending upon the circumstance; this occurs even in cases where the circumstance changes and a modification of an initial decision should be made. Of course it must be recognized that, similar to other technologies, the

> output of the decision came originally from human knowledge. However, the direct output of a decision is from the CAD itself and is frequently consulted by the dispatch team for its ability to provide efficient solutions to a given situation. Recent research, in fact, suggests that human decision making can be supplemented by the **"recognition" of certain situations** by information technology which, in turn, **uses this information to suggest problem solutions based upon past situations** (Terrell, 2006, p. 115, emphasis added).

By not needing to commit to a structural isomorphism for the sharing of mental models, our process-oriented approach supports the proposition that non-human systems can participate in teamwork and proposes an architecture for shared mental models under which it could happen. If indeed a system can perform some functions required of a shared mental model for effective teamwork—situation recognition and the consequent selection of methods likely to benefit other team members—then it does not matter how idiosyncratic that system's internal organisation might be; the system can still perform some of the functions of a team member. Bifurcation is an important apperceptive process for developing situation knowledge webs. However, it is increasingly expensive the further back it goes. Pre-bifurcation is an investment, allowing an organism to save time and resources when a snap decision needs to be made. Fully bifurcated webs of interactions encoded and embedded within a computerised teamwork support system, especially one of a diagnostic or decision-support nature, enable that system to attain to some degree the affective status of team member.

Communication and coordination in distributed teams

Members of dispersed and/or asynchronous teams may not be able to immediately perceive each other at all, let alone well enough to perceive their actions or the outcomes thereof. Consider two team members geographically dispersed at night such that one member must signal the other with a flare. It is tempting to think of this as an example of Action ∅ (seen in Figure 8 and discussed above). However, this is not so. The specific action is the pulling of the flare gun's trigger. There is in fact an intended and significant outcome: the flare itself lighting up a region of the night sky for a brief period. To clarify, think of how this process may fail: the gun may not be loaded; the gun may be loaded but misfire; the flare may be fired prematurely at a low elevation; the flare may shoot straight up but then fail to burn... In this case, what we might call a signal isn't really a simple signalling action like ∅ which *per se* has no significant environmental effect. Rather, it is the encapsulation of an extended and potentially error-prone process with environmental inputs and outputs. A general difficulty for distributed teams is that simple signalling *a la* ∅ is not usually possible. Instead, signalling (and coordination in general) often relies on actions that generate, transport and process artifacts. This brings an additional cost burden which is only justified if it would be impossible or exceedingly costly for the signalled actor to recognise (which presupposes perceiving) the action or situation which provides the signal's content. It also highlights the importance of being able to deal with error through on-the-fly generation and selection of novel methods and intra-team conventions. This is an area for future work.

Decreased communications, presumably through the development of more 'efficient' shared mental models, have been associated with improved team performance (eg. Langan-Fox et al., 2004). Based on prevailing theory stressing the importance of communication as an element of teamwork and correlations with team awareness measures, the *anticipation ratio* has been proposed as a teamwork effectiveness metric and predictor of mission success (MacMillan et al., 2004). In theory, a 'pushed' message contains information that the recipient has not requested. A 'pulled' message is sent in response to a request for information. So 'pulled' information is classified as such because it is preceded by a 'pushed' request for information; 'pushed' information is that which is not preceded by a request.[8] Considering all information exchanges between members of a team, a (push ÷ pull) ratio greater than 1 is taken to indicate more anticipation of other team members' needs—implying higher team awareness, more efficient communications and lower team workload, all presumably through more accurate and comprehensive shared mental models (Langan-Fox et al., ibid.)—than a (push ÷ pull) ratio less than 1.

A question naturally arises among operations researchers: what is a 'good' anticipation ratio? Could there be an optimal anticipation ratio calculable for each characteristic scenario and activity phase? If so, and if a training mission failed, practitioners might be able to attribute that failure to poor teamwork—specifically, through inaccurate or rudimentary shared mental models leading to poor team awareness, inefficient communications or high workload—if the anticipation ratio was found to be non-optimal. Unfortunately for this appealingly simple teamwork metric, some military simulation exercises conducted within my organisation have not sustained the expected correlations between anticipation ratio and mission success. One explanation offered is that inefficient shared mental models in mixed teams of software agents and human subjects cause software agents to produce a large volume of communication far beyond the capability of human subjects to process (Cao, 2005; Miller & Volz, 2001). But this then throws the value of the ratio itself into question. Firstly, it is not clear that large communication volumes have been constituted in requests for information. And secondly, where is the value in a large volume of unsolicited 'information'?

The approach outlined herein invites a closer examination of the anticipation ratio concept. The key question to be asked is: Did the receipt of a message provide some benefit to the recipient's apperception (Bickhard, 1980, 1998, 2005); that is, the ongoing appreciation of the possibilities for action and potentially beneficial method selections within the currently evolving situation? In short, did the message have *apperceptive value*? Without knowledge of the effect of pushed or pulled messages on the conduct of the mission, push/pull ratios are a blunt instrument with low predictive or analytic value. What matters more than the sheer volume of messages pushed or pulled is the proportion of messages with apperceptive value. We suggest that the results of future studies into team communication and performance, particularly those based on a push/pull communications paradigm, would be clarified by distinguishing between communications with high apperceptive value and those with little or no such value. Pushed communication is to be expected when one team member recognises a situation which they believe a team-mate cannot (easily) recognise but would have apperceptive value to them. Pulled communication is to be expected when a team member needs to

---

[8] In practice, it has not always been clear whether the content of a message is information pushed or pulled. The theoretical definitions have been found to be too coarse and simplistic for human subtlety in the use of language. Some refinement and interpretation has been required for application to military operations.

resolve a partially recognised situation (eg. through the costly process of bifurcation as discussed above). In our view, it would be counterproductive to try to eliminate pulled communications by maximising pushed communications because there is no apperceptive value in pushing so-called ‘information’ which does not, in fact, resolve a pertinent but only partially recognised situation.

We argue that communication is about more than “providing the right information at the right time to the right person” (to use a common military expression). It is more appropriate to see communication as serving primarily a complex team coordination function and hence it would be a mistake to always strive to decrease communications in the belief that it will necessarily lead to improved team performance. We propose that previously observed communications overload phenomena could be subdivided into two distinct cases: (1) when communications served no apperceptive function for the receiver; and (2) when the rate of apperceptive communications exceeded the rate at which the receiver could perform the corresponding transformations (eg. bifurcation) to their situation knowledge web.[9]

Comparisons to classical *Team Orientation*

It is instructive to compare *teamwork* as defined herein with the canonical ‘big five’ of Salas et al. (2005), particularly *team orientation*. It may be argued that the teamwork definition herein is merely a re-statement of the team orientation concept which is just one component of many competing teamwork frameworks. A team member can exemplify team orientation by (1) readily sharing information with other team members; or (2) considering, as input to a problem-solving process, alternative solutions proposed by other team members.

Our approach is not diametrically opposed to classical frameworks; the difference is more a matter of direction, perspective and focus. It proposes that information will be shared when it has (or is considered likely to have) apperceptive value. Effective information sharing will increase as a consequence of an increase in team members’ abilities to correctly estimate the apperceptive value of information. This approach can address team members’ alternative solutions as outputs as well as inputs. It proposes that team members will also consider alternative solutions for what they might offer as a means to set beneficial conditions for others to solve their problems. This approach could thus promote a transformation of focus from communication bandwidth to the development of apperceptive capability across the team.

---

[9] An example of the latter case, familiar to designers of computer systems, is when a peripheral device interrupts a computer’s CPU for a second time while the CPU is busy servicing the first interrupt from that device or another. The second interrupt may be ignored in which case data will be lost. But if the service handler is re-entrant, the CPU may attempt to service the second interrupt before returning to finish the first. And if a third interrupt arrives in the meantime (then a fourth, and so on) then processing will effectively halt until the interrupt stack overflows and the system crashes.

Conclusion

Based on a particular notion of the logical organisation and function of mental models, we have offered an approach to defining and studying teamwork that could underpin many components and coordinating mechanisms prevalent in the literature. It proposes that teamwork is characterised by team members' ongoing selection of mutually beneficial methods for subtask performance as enabled by apperceptive processes. Through such an approach, we may re-examine team phenomena and re-assess how we might develop new operational models and measures of teamwork with an eye to developing more effective teams.

References

Andersen, A. (2001). Implementation of engineering product design using international student teamwork. *European Journal of Engineering Education*, 26(2), 179–186.

Baker, D. P., Horvath, L., Campion, M., Offermann, L., & Salas, E. (2004). *Teamwork*. Technical report, Adult Literacy and Lifeskills Survey. URL: `http://www.ets.org/Media/Tests/ETS_Literacy/ALLS_TEAMWORK.pdf`.

Bickhard, M. H. & Richie, D. M. (1983). *On the nature of representation: A case study of James Gibson's theory of perception*. New York: Praeger.

Bickhard, M. H. & Terveen, L. (1995). *Foundational issues in artificial intelligence and cognitive science: Impasse and solution*. Amsterdam: Elsevier Scientific.

Bickhard, M. H. (1980). *Cognition, Convention and Communication*. New York: Praeger.

Bickhard, M. H. (1998). Levels of representationality. *Journal of Experimental & Theoretical Artificial Intelligence*, 10(2), 179–215.

Bickhard, M. H. (2005). Consciousness and reflective consciousness. *Philosophical Psychology*, 18(2), 205–218.

Bickhard, M. H. (2008). Social ontology as convention. *Topoi*, 27(1), 139–149.

Bosse, T., Jonker, C. M., Schut, M. C., & Treur, J. (2005). Simulation and analysis of a shared extended mind. *SIMULATION*, 81(10), 719–732.

Bosse, T., Jonker, C. M., Schut, M. C., & Treur, J. (2006a). Collective representational content for shared extended mind. *Cognitive Systems Research*, 7, 151–174.

Bosse, T., Jonker, C. M., van der Meij, L., Sharpanskykh, A., & Treur, J. (2006b). *A Temporal Trace Language for the Formal Analysis of Dynamic Properties*. Technical Report 234, Vrije Universiteit Amsterdam. URL: `http://www.few.vu.nl/~treur/TTL.pdf`.

Brazier, F. M. T., Jonker, C. M., Treur, J., & Wijngaards, N. J. E. (2000). On the use of shared task models in knowledge acquisition, strategic user interaction and clarification agents. *International journal of human-computer studies*, 52(1), 77–110.

Cannon-Bowers, J. A., Salas, E., & Converse, S. A. (1990). Cognitive psychology and team training: Training shared mental models and complex systems. *Human Factors Society Bulletin*, 33, 1–4.

Cannon-Bowers, J. A., Tannenbaum, S. I., Salas, E., & Volpe, C. E. (1995). Defining competencies and establishing team training requirements. In R. A. Guzzo & E. Salas (Eds.), *Team Effectiveness and Decision Making in Organizations* (pp. 333–380). San Francisco: Jossey-Bass.

Cao, S. (2005). *Role-Based and Agent-Oriented Teamwork Modeling*. PhD thesis, Texas A&M University.

Christoffersen, K. & Woods, D. D. (2002). How to make automated systems team players. In E. Salas (Ed.), *Advances in Human Performance and Cognitive Engineering Research*, volume 2 (pp. 1–12). Emerald Group Publishing Limited.

Cohen, P., Levesque, H., & Smith, I. (1997). On team formation. In G. Holmstrom-Hintikka & R. Tuomela (Eds.), *Contemporary Action Theory*, volume 2 (pp. 87–114). Dordrecht: Kluwer Academic.

Dousson, C., Gaborit, P., & Ghallab, M. (1993). Situation recognition: Representation and algorithms. In *Automated Reasoning* (pp. 166–172). Chambéry, France.

Driskell, J. E. & Salas, E. (2006). Groupware, group dynamics, and team performance. In C. Bowers, E. Salas, & F. Jentsch (Eds.), *Creating High-Tech Teams: Practical Guidance on Work Performance and Technology* (pp. 11–34). Washington, DC: American Psychological Association.

Entin, E. E. & Serfaty, D. (1999). Adaptive team coordination. *Human Factors*, 41, 312–325.

Espinosa, J. A. & Carley, K. M. (2001). Measuring team mental models. In *Proceedings of the Academy of Management Conference, Organizational Communication and Information Systems Division*. Washington, D.C.

Espinosa, J. A., Lerch, F. J., & Kraut, R. E. (2004). Explicit versus implicit coordination mechanisms and task dependencies: One size does not fit all. In Salas & Fiore (2004).

Hoffman, G. & Breazeal, C. (2008). Anticipatory perceptual simulation for human-robot joint practice: Theory and application study. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence* (pp. 1357–1362).

Hutchins, E. (1991). The social organisation of distributed cognition. In L. Resnick (Ed.), *Perspectives on socially shared cognition* (pp. 238–287). Washington, DC: American Psychological Association.

Jonker, C. M., Sharpanskykh, A., Treur, J., & Yolum, P. (2007). A framework for formal modeling and analysis of organizations. *Applied Intelligence*, 27(1), 49–66.

Knoblich, G. & Sebanz, N. (2008). Evolving intentions for social interaction: from entrainment to joint action. *Philosophical transactions of the Royal Society of London, Series B, Biological sciences*, 363(1499), 2021–2031.

Kulakov, A., Stojanov, G., & Davcev, D. (2002). A model of an expectancy-driven and analogy-making actor. In *Proceedings of the 2nd International Conference on Development and Learning (ICDL'02)*: IEEE Computer Society.

Langan-Fox, J., Anglim, J., & Wilson, J. R. (2004). Mental models, team mental models, and performance: Process, development, and future directions. *Human Factors and Ergonomics in Manufacturing*, 14(4), 331–352.

Lee, M. Y. (2007). *Understanding changes in team-related and task-related mental models and their effects on team and individual performance*. PhD thesis, The Florida State University.

Levesque, L. L., Wilson, J. M., & Wholey, D. R. (2001). Cognitive divergence and shared mental models in software development project teams. *Journal of Organizational Behavior*, 22, 135–144.

Lim, B. C. & Klein, K. J. (2006). Team mental models and team performance: A field study of the effects of team mental model similarity and accuracy. *Journal of Organizational Behavior*, 27, 403–418.

MacMillan, J., Entin, E. E., & Serfaty, D. (2004). Communication overhead: The hidden cost of team cognition. In Salas & Fiore (2004), (pp. 61–82).

Mareschal, D., Johnson, M., Sirois, S., Spratling, M., Thomas, M., & Westermann, G. (2007). *Neuroconstructivism*, volume I. Oxford University Press.

Marks, M. A., Mathieu, J. E., & Zaccaro, S. J. (2001). A conceptual framework and taxonomy of team processes. *Academy of Management Review*, 26, 356–376.

Marks, M. A., Sabella, M. J., Burke, C. S., & Zaccaro, S. J. (2002). The impact of cross-training on team effectiveness. *Journal of Applied Psychology*, 87(1), 3–13.

Mathieu, J. E., Heffner, T. S., Goodwin, G. F., Cannon-Bowers, J. A., & Salas, E. (2005). Scaling the quality of teammates' mental models: Equifinality and normative comparisons. *Journal of Organizational Behavior*, 26, 37–56.

Mathieu, J. E., Heffner, T. S., Goodwin, G. F., Salas, E., & Cannon-Bowers, J. A. (2000). The influence of shared mental models on team process and performance. *Journal of Applied Psychology*, 85(2), 273–283.

Miller, C. A. & Parasuraman, R. (2007). Designing for flexible interaction between humans and automation: Delegation interfaces for supervisory control. *Human Factors*, 49(1), 57–75.

Miller, M. S. & Volz, R. A. (2001). Training for teamwork. In *5th World Multi-Conference on Systemics, Cybernetics and Informatics*. Orlando, FL.

Mohammed, S. & Dumville, B. (2001). Team mental models in a team knowledge framework: Expanding theory and measurement across disciplinary boundaries. *Journal of Organizational Behavior*, 22, 89–106.

Parasuraman, R., Sheridan, T., & Wickens, C. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man and Cybernetics – Part A: Systems and Humans*, 30, 286–297.

Popova, V. & Sharpanskykh, A. (2007). A formal framework for modeling and analysis of organizations. In J. Ralyté, S. Brinkkemper, & B. Henderson-Sellers (Eds.), *Situational Method Engineering: Fundamentals and Experiences*, volume 244 of *IFIP International Federation for Information Processing* (pp. 343–358). Boston: Springer. (Proceedings of the IFIP WG 8.1 Working Conference, 12–14 September 2007, Geneva, Switzerland.)

Resick, C. J. (2004). *An investigation of the antecedents and consequences of shared mental models in teams*. PhD thesis, Wayne State University, Michigan.

Roda, C. (2007). Supporting attention with dynamic user models. In *Proceedings of the Interactivist Summer Institute 2007*, Paris, 26–29 May 2007.

Rouse, W. B. & Morris, N. M. (1986). On looking into the black box: Prospects and limits in the search for mental models. *Psychological Bulletin*, 100(3), 349–363.

Salas, E., Dickinson, T., Converse, S., & Tannenbaum, S. (1992). Towards an understanding of team performance and training. In R. W. Swezey & E. Salas (Eds.), *Teams: Their Training and Performance*. Norwood, New Jersey: Ablex Publishing Corporation.

Salas, E. & Fiore, S. M., Eds. (2004). *Team Cognition: Understanding the Factors that Drive Process and Performance*. Washington, DC: American Psychological Association.

Salas, E., Sims, D., & Burke, C. (2005). Is there a `big five' in teamwork? *Small Group Research*, 36(5), 555–599.

Seibt, J. (2009). Forms of emergent interaction in general process theory. *Synthese*, 166(3), 479–512.

Shanahan, C., Best, C., Finch, M., & Sutton, C. (2007). *Measurement of the Behavioural, Cognitive, and Motivational Factors Underlying Team Performance*. Technical Report DSTO-RR-0328, Air Operations Division, Defence Science and Technology Organisation, Fishermans Bend, Victoria.

Sharpanskykh, A. (2008). *On computer-aided methods for modeling and analysis of organizations*. PhD thesis, Vrije Universiteit Amsterdam, The Netherlands.

Smart, P. R., Mott, D., Sycara, K., Braines, D., Strub, M., & Shadbolt, N. R. (2009). Shared understanding within military coalitions: A definition and review of research challenges. In *Proceedings of KSCO-09* Chilworth Manor, Southampton, UK: Knowledge Systems for Coalition Operations.

Stojanov, G. & Kulakov, A. (2003). Interactivist approach to representation in epigenetic agents. In C. G. Prince, L. Berthouze, H. Kozima, D. Bullock, G. Stojanov, & C. Balkenius (Eds.), *Proceedings of the Third International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems 101* (pp. 123–130). Boston, MA, USA.

Stojanov, G. & Kulakov, A. (2011). Modelling attention within a complete cognitive architecture. In C. Roda

(Ed.), *Human Attention in Digital Environments* (pp. 210–241). Cambridge, UK: Cambridge University Press.

Stroeve, S. H., Sharpanskykh, A., & Blom, H. A. P. (2007). Systemic safety modelling and analysis of organizational processes in air traffic. In *Proceedings of the EUROCONTROL Annual Safety R&D Seminar*, Rome, 24–26 Oct 2007.

Terrell, I. S. (2006). *Understanding 911 Dispatch Teams across context: Implications for theory, information technology, and practice*. PhD thesis, The Pennsylvania State University.

Tollefsen, D. P. (2006). From extended mind to collective mind. *Cognitive Systems Research*, 7, 140–150.

Wilson, R. (2004). *Boundaries of the mind*. Cambridge: Cambridge University Press.

Yen, J., Yin, J., Ioerger, T. R., Miller, M. S., Xu, D., & Volz, R. A. (2001). Cast: Collaborative agents for simulating teamwork. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'01)*, Seattle.