How Could We Model Cohesiveness in Team Social Fabric in Human-Robot Teams Performing Under Stress?

Geert-Jan M. Kruijff
German Research Center for Artificial Intelligence
Saarbrücken, Germany
gj@dfki.de

Abstract

The paper discusses how a human-robot team can remain "cohesive" while performing under stress. By cohesive the paper understands the ability of the team to operate effectively, with individual members being interdependent-yet-autonomous in carrying out tasks. For a human-robot team, we argue that this requires robots to (1) have an adequate sense of that interdependency in terms of the social dynamics within the team, and to (2) maintain transparency towards the human team members in terms of what it is doing, why, and to what extent it can achieve its (possibly jointly agreed upon) goals. The paper reports of recent field experience showing that failure in transparency results in reduced acceptability of robot autonomous behavior by the human team members. This reduction in acceptability can have two negative impacts on cohesiveness: Humans and robots fail to maintain common ground, and as a result they fail to maintain trust.

Introduction

We envision humans and robots to team up, work together. And not just that. We intend them to do so under circumstances that are physically, mentally, stressful. Military missions. Search and rescue missions. These are missions that typically demand the utmost from those involved.

Teams survive these ordeals because, or rather – when, they stick together. This is a matter of leadership, this is a matter of bonding between the members of the team; See for example (Siebold 2000; Kolditz 2007). The stronger the social fabric and moral fibre of the team, the more resilient it is against adverse circumstances. That is much to ask. But it is something we can train people, teams, to achieve.

How could we ever achieve that in human-robot teams?

Why would we ever even need to achieve that?

The question of social fabric in human-robot teams can only be asked if we can consider robots to be, at least to some degree, autonomous actors. If they are not, we should strictly speaking not consider them as team-members, but simply as tools. This notwithstanding the human tendency to anthropomorphize robots (Duffy 2003); See (Singer 2009) for interesting anecdotes from soldiers bonding with their *packbot* robots during deployments in the Iraq theatre. Cohesiveness is about the strength between actors, between whom there exists an interdependency.

This naturally does raise the question whether robots *ought to be* autonomous, whatever the degree, and whatever precisely we might understand by that notion; See e.g. (Parasuraman, Barnes, and Cosenzo 2007). The literature shows this to be a matter of what the robot is supposed to be capable of doing autonomously. For robots with lethal capability, opinions point to restricting autonomy, always keeping the human (soldier) in full control (Moshkina and Arkin 2007). On the other hand, for robots to be deployed in search & rescue missions, increased autonomy is often seen as highly desirable (Birk and Carpin 2006). The contention here is that, without some degree of autonomy in the robot, there is no human-robot team to speak of.

But back to interdependency. In a human-robot team, there is an interdependence between the different actors (Johnson et al. 2011). Shared control is a *relation* between a human and a robot; See also (Bradshaw et al. 2004). Autonomy is not just a matter of the behavior of an individual robot. It is how this behavior is interleaved with human behavior, human task assignments and delegations. Humans and robots *together* build up a common ground, providing complementary information, playing by a common set of rules and agreements for a joint activity (Klein et al. 2004b; Klein et al. 2004a; Bradshaw, Feltovich, and Johnson 2011).

One way to model a perspective on interdependency is to use policies which (directly) govern action- and interaction behavior of a robot; See e.g. (Johnson et al. 2006). Here we would like to complement that (fairly low-level) perspective with a (higher-level) characterization of the social structure of a human-robot team. This social structure is based on an analysis of possible roles within a human-robot team, following up on the communicative analysis of (Burke et al. 2004; Murphy and Burke 2010). A role for an actor essentially specifies the functions it is responsible for (action, interaction) and with which other roles these functions are connected (relational view); under what conditions these functions can be performed (integrity limits); and what information is required, and provided, by the role. We would like to argue that the advantages of the role-based perspective are the following:

- Human-robot teaming can be modeled from an event-driven, systemic viewpoint on the entire socio-technical

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

system; See (Stanton, Baber, and Harris 2008).

- Context-driven demands on functioning within the team can be explicitly tied to social structure: roles between actors, and context-specific information demands, yielding a *distributed* notion of situation awareness which is necessary for modeling geographically distributed teams; See (Salmon et al. 2009).
- An explicit model of social structure and distributed situation awareness, coupled to explicitly modeled and actively measurable integrity limits, aid in providing means for active trust management in the team (particularly from the human to the robot); See (Fitzhugh, Hoffman, and Miller 2011).

How Could We Model?

This paper *does not* provide a model. There is no fully functional model of human-robot teaming yet, particularly not of human-robot teams performing under stress. The point we are trying to make is that we need to construct an explicit account of the social dynamics within a human-robot team. A robot needs to be able to reason with social structure. It needs to understand why, when, and how it is to work with others. Otherwise it cannot adapt its own behavior to best fit the circumstances – particularly if it just does not have enough (certain) information to act upon in a particular situation.

For this to work we need to ground these social dynamics in a notion of distributed situation awareness. Human-robot teams are often distributed across a larger area. Typically robots are deployed in a hostile area, whereas human operators are at a remote command post, or work in line-of-sight of the robot but not necessarily directly next to it. This in and by itself already means there is a variation in the perspectives the various actors develop on "the" situation, an issue further compounded by the way roles might imply different (possibly mediated) views on reality (Murphy and Burke 2010). Furthermore, not everybody needs to know the same, needs to have the same situation awareness. This very much depends on their assignments, the roles they play (Murphy and Burke 2010). Situation awareness is distributed between the actors, depending on who they are connected "socially" through the roles they play; a notion quite different from the individualistic or everybody-is-sharing notions of situation awareness, See (Salmon et al. 2009).

Finally, tracking the dynamics of role assignments and - shifts within a human-robot team provides a context within which we can consider active trust management (Fitzhugh, Hoffman, and Miller 2011). Trust between a human and a complex technological artifact such as a semi-autonomous robot can not be considered a static notion. Trust of the human in the robot, i.e. with the robot being the object of trust, is very much dependent on context. Fitzhugh et al mention multiple dimensions: Trust in the technology per se (trust in reliability of the platform, sensors, to network infrastructure); trust in decision-making (trust in the robot to deal accurately with incomplete and uncertain information); trust in the sociotechnical work system (trust in the robot and the overall system to assist the human); and, trust in the overall capability of the human-robot to be able "to get the job done." An important issue is then, how to provide the quantitative means for a robot to measure its performance, to communicate that to the user to guide trust management, and to *actively* communicate and resolve issues in trust. This is where the interdependency comes in again: If a robot is uncertain, it can (and should) indicate this, and solicit the human in helping to resolve it; See e.g. (Kruijff, Brenner, and Hawes 2008).

To facilitate this, we have been working on a model of team structure. We close this paper by briefly outlining its main structure. At the core of the model is the notion of a role. A role defines several dimensions which are intended to cover the requirements for raise for collaborative role models for robots in C2 activities, following (Klein et al. 2004b) and (Murphy and Burke 2010).

- **Scope of Action**: A role needs to define what actions it performs in information- and decision-management processes, (relative to one or more C2 activities). This outlines the scope of the contributions this role can make to the overall effort.
- **Bandwith of Autonomy**: A role needs to define the lower- and upper limits on its autonomy in acting in information- and decision-management processes (LOA, (Parasuraman, Sheridan, and Wickens 2000)).
- **Integrity Limits**: Complementary to the LOA bandwidth for specific actions a role needs to define a notion of "integrity limits." These need to describe the limits to which contributions can be made, and provide for contingency management e.g. through strategies for direct reporting to identified (active) roles in the team ("who should know when I fail").

In highly interdependent settings like a team, the behavior of the individual actors needs to be transparent to others (*mutual predictability*). It needs to be clear why someone is performing a particular action, so that outcomes and follow-up behavior can be predicted. This is necessary for interdependent action and coordination to be efficiently plannable and executable. Challenges here include predictability itself, the ability of an agent to make pertinent aspects of his actions and intentions clear to others, and the ability to observe and interpret such signals from others.

- **Direct Reports on Acting: To**: A role needs to define strategies for reporting its actions and its reasons for performing them, to one or more roles in a team. This is connected to the Bandwidth of Autonomy, and the issue of keeping the human in the loop.
- **Direct Reports on Acting: From** A role needs to define conditions that reflect interdependence of its own role on others, in terms of information- and decision-making state.

Mutual directability involves both the capability to assess and modify the actions of other agents, and to be responsive to the influence of others on one's own actions. This involves the ability to be directable, to negotiate goals, and to support a continual approach to collaboration. The latter reflects the need to allow for the actions and plans to be collaboratively adjusted, as the situation demands. Finally, actors must be able to participate in managing attention, particularly to be able to align perspectives on situations (Murphy and Burke 2010; Zender, Kruijff, and Kruijff-Korbayová 2009; Zender et al. 2010).

Again, from the viewpoint of collaborative role models and keeping the human in the loop, this requires roles to combine an upper limit to autonomy in decision processing, with reporting (and negotiating) on decision status, as indicated above. Furthermore, a role needs to make dynamic and role-based authority explicit.

- **Authority** A role needs to make explicit authority relations: from what roles it accepts directability, what goals it can negotiate with other roles (and when), and for what actions (states) it has authority to decide (which may include actions associated with other roles). Authority may be inherent to the role itself, or be derived from the adoption of a role in the current team context (e.g. being designated team leader over other agents).
- **Attention** A role needs to define strategies for directing the attention of others in specific roles, on the basis of signals, activities and changes in its own state. Any attention-direction action needs to be mediated through the viewpoints associated with the roles in question.

Finally, effective coordination between team members requires them to establish and maintain a form of common ground, distributed situation awareness capturing (activated) knowledge, beliefs, and intentions.

- **Viewpoints** A role needs to define its own viewpoint(s) (Murphy and Burke 2010) and the modalities through which these viewpoints are mediated.

We formulate a role as a tuple (*PossibleActors, Functions,Relations,Viewpoints*). *PossibleActors* is a set of actors (types) that can take on this role. *Functions* is a set of function primitives. The function primitives identify what they operate on (information INFO, actions ACT) , and (where relevant) with respect to what (ENVIRONMENT), or whom (the role's primary actor SELF, or other actors ACTOR). ACTORs are identified by roles. ACTs have assigned ACTORs including SELF. Table 1 lists the primitives we consider. The list is an extension of the RASAR-CCS scheme presented in (Burke 2003). For a function, the role also specifies a bandwidth of autonomy indicating what level(s) of autonomy are required for performing this function, and which integrity limits condition the possible execution of the function. *Relations* is a set of directed links between the role, and other roles. Each link indicates how an actor "playing" this role can construct connections to other roles active in the team. Finally, *viewpoints* specify from what (mediated or non-mediated) perspectives an actor in this role perceives the situation; See (Murphy and Burke 2010).

We are currently working on a logical formalization of the role-based model of social dynamics. The model is grounded in logical-probabilistic models of multi-agent situation awareness, and is coupled to an explicit account of trust (Kruijff and Janíček 2011); See also (Herzig et al. 2010). One of the current challenges is to provide a notion of situatedness which allows us to accurately capture trust as an active concept, with the possibility to vary the attribution, reliance, and management of trust within a team with the degree of stress under which the human actors are operating.

Recent experience

Over 2011, we have been deploying systems for robot-assisted Urban Search & Rescue in various realistic settings.

	Function	Explanation
1.	SEEKINFO	Ask for INFO from an ACTOR
2.	REPORT	Share INFO about SELF, ENVIRONMENT, or other ACTOR
3.	CLARIFY	Make previous INFO more precise
4.	CONFIRM	Affirm previous INFO, or (selected) ACT
5.	CONVEYUNC	Express doubt, disorientation, or loss of confidence in INFO
6.	PROVIDEINFO	Share INFO other than REPORT, either in response to a SEEKINFO request from another ACTOR, or to provide unsolicited information
7.	PLAN	Project future, spatially situated GOALs, or ACTs to GOALs
8.	SELECT	Select ACT
9.	EXECUTE	Execute ACT
10.	ORDER	Authority: Order another ACTOR to ACT, or allow another ACTOR to order SELF
11.	INTERVENE	Authority: Allow another ACTOR to intervene in ACT
12.	PROPOSE	Propose ACT(s) to ACTOR

Table 1: Function primitives for robot collaborative role models; (1–7 from Burke et al/RASAR-CCS).

In an initial deployment at the training center of the Fire Department of Dortmund (FDDO; January 2011), first responders were given the task to explore a tunnel accident-like setting, using a fully tele-operated robot(ActivMedia P3-AT, with 2D laser and omni-directional camera). The setting was built up in a large garage, and involved several crashed cars, debris, and a motor bike. Victims were distributed over the entire setting, inside and outside cars, and near the motorbike. The exploration was conducted under time pressure.

We observed relatively constant, average cognitive load for each subject (subjective rating), although subjects did display heightened stress signals as soon as the robot would be operating under smokey conditions. Subjects were able to build up a relatively accurate assessment of the situation, despite difficulties in observing for example victims in cars. (For more detailed discussions, see (Larochelle et al. 2011; Mioch, Smets, and Neerincx 2012).)

Adding more robots into the setting to provide for more flexibility in operating in the environment would make it necessary though for more humans to become involved; See also the discussion in (Murphy and Burke 2010). To study human-robot teaming in more detail, we first organized NJEx 2011, a joint exercises event.[1] During this event, teams consisting of several humans, a microcopter (UAV) and a rover (UGV) would explore several complex environments, including a multi-story residential building "on fire." Team members included both first responders and scientists. Both human and robot team members were geographically dispersed. Human team members in the roles of Mission Commander, UGV Operator, and UGV/UAV Mission Spe-

[1]Video: `http://youtu.be/MTtEEsoEhnw`

(a) (b) (c)

Figure 1: Human-robot teams at a control post (a; FDDO Dortmund); team interface setup, at the table at the wall in the back (b; SFO Montelibretti); UGV operator GUI showing different 2D and 3D map information, as well as panoramic and zoomable camera views (c)

cialist were located in a remote control post (Fig. 1(a)), whereas the UAV Operator remained in-field within line-of-sight of the UAV to fly it.

For NJEx 2011 we particularly focused on human-human interaction within the human-robot team, as the robots were fully tele-operated. To this end, we recorded all the radio traffic between human members which were distributed over the area. Interesting observations here were that communication particularly concerned the communication of situation awareness ("we see a victim under the shelves in the room at the end of corridor, right"), and goals ("we are going to look under the staircase, at the end of the corridor, left"). All of these messages were explicitly situated (or located) in the environment. The Mission Commander mostly communicated situation awareness, to maintain common ground within the team, whereas the UGV operator would indicate what next actions the UGV would perform (and where). Planning exploration tasks was typically done within the control post. Coordination with the in-field team members was done through the Mission Commander. The UAV Operator would be tasked to fly the UAV to a particular point, with an explicitly communicated purpose – typically, what kinds of observations the control post would like to make. Video feeds from the UAV were inspected by the UAV Mission Specialist in the command post, with the Mission Commander providing feedback to the UAV Operator.

For safety reasons each team included two Safety Directors, one for the UGV and one for the UAV. The Safety Directors had the best awareness of the situations around the robots, as they were right there. The protocol was such that they were not allowed to describe the environment, just in some cases they could give hints about how to get a robot out of a tight spot. In practice, as no situation awareness was allowed to come from a Safety Director (protocol), the team members in the remote control post often decided to rely on their own insights.

Altogether we collected over 12 hours of audio data (radio traffic) during NJEx. This data has been segmented, and annotated for speaker roles, and perceived stress. Analysis so far has yielded that the Mission Commander and the UGV Operator generate the most radio traffic, with one or the other taking on a leading role. (In the most effective teams, this was always the Mission Commander.) Furthermore, variations in stress levels can be detected acoustically in this data in real time, currently with a correct classification rate of 74% (Charfuelan, p.c.). Stress variation was particularly observable for team members with high radio traffic, i.e. the Mission Commander, UGV Operator, and the UGV Safety Commander.

We took these insights into account for the end user evaluations we performed mid-December 2011, at the SFO training area of the Vigili del Fuoco, in Montelibretti (Italy). At these evaluations, experiment subjects took the role of UGV Operator, in a team consisting of a Mission Commander doubling as UAV Mission Specialist, an UAV Operator, an UGV, and a UAV. The Mission Commander and the UGV Operator were in a remote command post, whereas the robots and the UAV Operator were in-field. The task of the team was to explore a complex tunnel accident involving multiple cars, debris, and a truck, set up in a real-life tunnel structure.[2] Operating in the tunnel happened under strong variations in lighting conditions, and(artificial) smoke.

The team members in the control post had access to variety of information views, in a multi-screen multi-modal graphical user interface. Fig. 1(b) shows the overall setup. Views included robot-specific interfaces, for example Fig. 1(c) illustrates the view for the UGV Operator, and (qualitative) views for team-level situation awareness. Communication between the command post and the in-field UAV Operator was via hand-held radio, through the Mission Commander. The UGV Operator communicated with the in-field UGV using the multi-modal GUI (touchscreen) and possibly spoken dialogue. The UGV was capable of autonomous navigation, and could use spoken dialogue to inform the UGV Operator about observations, and to provide basic feedback on actions (action-possibility, action-onset).

Altogether the exploration task was run with 7 subjects over the course of a week. Each subject took about 4 hours to instruct, run the task with, and debrief. The exploration task itself took about 45 minutes. For each subject several objective biometric data (heart rate, facial expression, task- and communication observation protocols) were recorded, as well as subjective ratings of cognitive load (intrusive, every 2 minutes). For 4 subjects we recorded audio of their interactions with the robot, and with other team members. In addition, logs were created for the entire user interface, and the UGV. These logs enable a complete replay of each

[2] See e.g. `http://youtu.be/1hWEIV0XL_4`

run of the UGV.

All this data (approximately 200GB) is currently still under analysis. At the same time, we have been able to make some initial observations. Robots can assist humans in complex missions like Urban Search & Rescue (Murphy et al. 2008). To make this possible, we arguably need more autonomy in the robot (Birk and Carpin 2006) – whereby we must understand "autonomy" as a complex notion, spanning a wide variety of capabilities, ranging from perception to navigation. However. Disaster areas are harsh places, for humans and robots alike. In practice we *always* experience what Woods et al (Woods et al. 2004) termed "***(Robin) Murphy's Law**: any deployment of robotic systems will fall short of the target level of autonomy, creating or exacerbating a shortfall in mechanisms for coordination with human problem holders*."

Things break down. Inevitably. Adaptive autonomy or shared control might be one way out of this (Parasuraman, Barnes, and Cosenzo 2007; Miller and Parasuraman 2007), but the problem really goes much deeper than that. All autonomy is for naught if the humans in the team do not *accept* a robot's autonomous capabilities and intelligence. Recent experience in Fukushima (S. Tadokoro, p.c.) and in our own end user evaluations at SFO underline this. A robot's abilities, behavior, and possible achievements need to be transparent to a human operator: Whether the robot is doing something, what it is doing and why, whether it thinks it has achieved a goal (or not). If an operator is unclear about what to expect from the robot, he or she is unlikely to delegate control to the robot. Instead, no matter what the robot is able to do on its own, the operator reverts to tele-operation.

That's not quite what anybody wants. But currently, this is where we seem to stand. There is an issue of (lacking) transparency in experience, behavior and intentions (Clark 1996; Thomaz and Breazeal 2008). As this directly impacts user expectations, a lack of transparency can seriously affect the interaction (Lohse 2011; Komatsu and Yamada 2011); cf. also (Dautenhahn 2007). All of this results in a lack of acceptability, thus serving as a possible explanation for *why* human-robot interaction appears to be a bottleneck in USAR (Murphy 2004). This problem gets exacerbated in the context of USAR. Even though the subjects at the SFO experiments indicated average cognitive workload, other (more objective) observations indicated that there was frustration, stress. Partly this was due to technical failures, but in a not insignificant way this was also due to things simply being unclear to the users. As soon as that would happen, they would regain control of the UGV, and operate it under much more stringent levels of autonomy: For example, users would move from waypoint navigation back to small movement commands, or even revert to basic tele-operation.

Thus, if we want humans and robots to operate as a team, transparency is key to achieving effective communication, coordination, and collaboration.

Acknowledgement

This paper describes research done under the EU-FP7 ICT 247870 NIFTi project. For more about NIFTi, please visit `http://www.nifti.eu`. The author would like to thank Nanja Smets, Tina Mioch, Mark Neerincx, Benoit Larochelle, Hendrik Zender, Jeff Bradshaw, and Satoshi Tadokoro for fruitful discussions. The FDDO pilot experiments in January 2011 were led by Tina Mioch, and the SFO end user evaluations in December 2011 by Nanja Smets (both from the TNO project partner in NIFTi).

References

[Birk and Carpin 2006] Birk, A., and Carpin, S. 2006. Rescue robotics - a crucial milestone on the road to autonomous systems. *Advanced Robotics* 20(5):595–605.

[Bradshaw et al. 2004] Bradshaw, J.; Feltovich, P.; Jung, H.; Kulkarni, S.; Taysom, W.; and Uszok, A. 2004. Dimensions of adjustable autonomy and mixed-initiative interaction. In Nickles, M.; Rovatsos, M.; and Weiss, G., eds., *Agents and Computational Autonomy: Potential, Risks, and Solutions*, number 2969 in Lecture Notes in Computer Science. Berlin, Germany: Springer Verlag. 17–39.

[Bradshaw, Feltovich, and Johnson 2011] Bradshaw, J. M.; Feltovich, P.; and Johnson, M. 2011. Human-agent interaction. In Boy, G., ed., *Handbook of Human-Machine Interaction*. Ashgate. 283–302.

[Burke et al. 2004] Burke, J.; Murphy, R.; Coovert, M.; and Riddle, D. 2004. Moonlight in Miami: An ethnographic study of human-robot interaction in USAR. *Human Computer Interaction* 19((1–2)):85–116.

[Burke 2003] Burke, J. 2003. Moonlight in Miami: A field study of human-robot interaction in the context of an urban search and rescue disaster training response training exercise. Master's thesis, Department of Psychology, College of Arts and Sciences, University of South Florida.

[Clark 1996] Clark, H. 1996. *Using Language*. Cambridge University Press.

[Dautenhahn 2007] Dautenhahn, K. 2007. Socially intelligent robots: dimensions of human-robot interaction. *Philosophical Transactions of The Royal Society B* 362:679–704.

[Duffy 2003] Duffy, B. 2003. Anthropomorphism and the social robot. *Robotics and Autonomous Systems* 42:177190.

[Fitzhugh, Hoffman, and Miller 2011] Fitzhugh, E.; Hoffman, R.; and Miller, J. 2011. Active trust management. In Stanton, N., ed., *Trust In Military Teams*, Human Factors in Defence. Ashgate. 197–218.

[Herzig et al. 2010] Herzig, A.; Lorini, E.; Hbner, J.; and Vercouter, L. 2010. A logic of trust and reputation. *Logic Journal of the IGPL* 18(1):214–244.

[Johnson et al. 2006] Johnson, M.; Bradshaw, J.; Feltovich, P.; Jeffers, R.; and Uszok, A. 2006. A semantically-rich policy-based approach to robot control. In *Proceedings of the International Conference on Informatics in Control, Automation, and Robotics (ICINCO 2006)*.

[Johnson et al. 2011] Johnson, M.; Bradshaw, J.; Feltovich, P.; Hoffman, R.; Jonker, C.; van Riemsdijk, B.; and Sierhuis, M. 2011. Beyond cooperative robotics: The central role of interdependence in coactive design. *IEEE Intelligent Systems* 81–88.

[Klein et al. 2004a] Klein, G.; Feltovich, P.; Bradshaw, J.; and Woods, D. 2004a. Common ground and coordination in joint activity. In Rouse, W., and Boff, K., eds., *Organizational Simulation*. New York City, NY: John Wiley. 139–184.

[Klein et al. 2004b] Klein, G.; Woods, D.; Bradshaw, J.; Hoffman, R.; and Feltovich, P. 2004b. Ten challenges for making automation a team player in joint human-agent activity. *IEEE Intelligent Systems* 19(6):91–95.

[Kolditz 2007] Kolditz, T. 2007. Leading as if your life depended on it. In Crandall, M. D., ed., *Leadership Lessons from West Point*. Jossey-Bass / John Wiley. 160–187.

[Komatsu and Yamada 2011] Komatsu, T., and Yamada, S. 2011. Adaptation gap hypothesis: How differences between users expected and perceived agent functions affect their subjective impression. *Journal of Systemics, Cybernetics and Informatics* 9(1):67 – 74.

[Kruijff and Janíček 2011] Kruijff, G., and Janíček, M. 2011. Using doctrines for human-robot collaboration to guide ethical behavior. In *Proceedings of the AAAI 2011 Fall Symposium Robot-human team-work in dynamic adverse environments*. AAAI.

[Kruijff, Brenner, and Hawes 2008] Kruijff, G.; Brenner, M.; and Hawes, N. 2008. Continual planning for cross-modal situated clarification in human-robot interaction. In *Proceedings of the 17th International Symposium on Robot and Human Interactive Communication (RO-MAN 2008)*.

[Larochelle et al. 2011] Larochelle, B.; Kruijff, G.; Smets, N.; Mioch, T.; and Groenewegen, P. 2011. Establishing human situation awareness using a multi-modal operator control unit in an urban search & rescue human-robot team. In *Proceedings of the 20th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE.

[Lohse 2011] Lohse, M. 2011. The role of expectations and situations in human-robot interaction. In Dautenhahn, K., and Saunders, J., eds., *New Frontiers in Human-Robot Interaction*. John Benjamins. 35–56.

[Miller and Parasuraman 2007] Miller, C., and Parasuraman, R. 2007. Designing for flexible interaction between humans and automation: Delegation interfaces for supervisory control. *Human Factors* 49:57–75.

[Mioch, Smets, and Neerincx 2012] Mioch, T.; Smets, N.; and Neerincx, M. 2012. Predicting performance and situation awareness of robot operators in complex situations by unit task tests. In *Proceedings of the Fifth International Conference on Advances in Computer-Human Interactions 2012*.

[Moshkina and Arkin 2007] Moshkina, L., and Arkin, R. 2007. Lethality and autonomous systems: Survey design and results. Technical Report GIT-GVU-07-16, Georgia Tech. Prepared for the US Army Research Office.

[Murphy and Burke 2010] Murphy, R., and Burke, J. 2010. The safe human-robot ratio. In Barnes, M., and Jentsch, F., eds., *Human-Robot Interactions in Future Military Operations*, Human Factors in Defence. Ashgate. 31–49.

[Murphy et al. 2008] Murphy, R.; Tadokoro, S.; Nardi, D.; Jacoff, A.; Fiorini, P.; Choset, H.; and Erkmen, A. 2008. Search and rescue robotics. In Siciliano, B., and Khatib, O., eds., *Springer Handbook of Robotics*. Springer Verlag. Part F, 1151–1173.

[Murphy 2004] Murphy, R. 2004. Human-robot interaction in rescue robotics. *IEEE Transactions on Systems, Man and Cybernetics Part C: Applications and Reviews* 34(2):138–153.

[Parasuraman, Barnes, and Cosenzo 2007] Parasuraman, R.; Barnes, M.; and Cosenzo, K. 2007. Adaptive automation for human-robot teaming in future command and control systems. *International Journal of Command and Control* 1(2):43–68.

[Parasuraman, Sheridan, and Wickens 2000] Parasuraman, R.; Sheridan, T. B.; and Wickens, C. D. 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.

[Salmon et al. 2009] Salmon, P.; Stanton, N.; Walker, G.; and Jenkins, D. 2009. *Distributed Situation Awareness: Theory, Measurement, and Application to Teamwork*. Human Factors in Defence. Ashgate.

[Siebold 2000] Siebold, G. 2000. Military group cohesion. In Britt, T.; Castro, C.; and Adler, A., eds., *Military Life: The psychology of serving in peace and combat*. Praeger Security International. 185–201.

[Singer 2009] Singer, P. 2009. *Wired for War: The Robotics Revolution and Conflict in the 21st Century*. Penguin Books.

[Stanton, Baber, and Harris 2008] Stanton, N.; Baber, C.; and Harris, D. 2008. *Modelling Command and Control: Event Analysis of Systemic Teamwork*. Human Factors in Defence. Ashgate.

[Thomaz and Breazeal 2008] Thomaz, A. L., and Breazeal, C. 2008. Experiments in socially guided exploration: Lessons learned in building robots that learn with and without human teachers. *Connection Science* 20(2&3):91–110. Special Issue on Social Learning in Embodied Agents.

[Woods et al. 2004] Woods, D.; Tittle, J.; Feil, M.; and Roesler, A. 2004. Envisioning human-robot coordination for future operations. *IEEE Transactions on Systems, Man and Cybernetics Part C: Applications and Reviews* 34(2):210–218.

[Zender et al. 2010] Zender, H.; Koppermann, C.; Greeve, F.; and Kruijff, G.-J. M. 2010. Anchor-progression in spatially situated discourse: a production experiment. In *Proceedings of the Sixth International Natural Language Generation Conference (INLG 2010)*, 209–213. Trim, Co. Meath, Ireland: Association for Computational Linguistics.

[Zender, Kruijff, and Kruijff-Korbayová 2009] Zender, H.; Kruijff, G.-J. M.; and Kruijff-Korbayová, I. 2009. Situated resolution and generation of spatial referring expressions for robotic assistants. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09)*, 1604–1609.