Achieving Common Ground under Asymmetric Agency and Social Sentience in Communication for Human-Robot Teaming

(Invited Paper)

Geert-Jan M. Kruijff
DFKI GmbH
Talking Robots@the Language Technology Lab
Saarbrücken, Germany
gj@dfki.de / http://talkingrobots.dfki.de

Abstract — In a human-robot team, actors need to collaborate at different levels to make things work. They work together on performing tasks, to achieve their mission. But, beyond this *task-work*, there is the *team-work* per se. The team needs to coordinate its social dynamics, the assignment of roles and responsibilities, the building up and maintaining of a shared understanding of what the mission is about and what the environment operated in is like. Communication is key, particularly at the level of team-work. This gives rise to particular challenges for human-robot interaction, among others the formation of a common ground. The paper discusses a framework for situated dialogue as part of collaborative activity, and shows how some of the challenges for building and maintaining common ground can be dealt with.

Keywords: *Robot-assisted USAR, human-robot interaction, human-robot teams, situation awareness*

I. Introduction

Robot-assisted search and rescue is often promoted as a driver for robot autonomy, see e.g. [1], [26]. And robots are getting more autonomous, there is no doubt about that. At the same time, practical experience is showing us that robots neither can, nor should, function fully autonomously in a search & rescue mission. Humans are, after all, the ultimate stake-holders. They want to be in control, or at least have the feeling they are in control [22] – even though they cannot (typically) go in and do things themselves. That makes both humans and robots problem-holders in such missions, giving rise to interdependence between humans and robots [12], [25].

Robot-assisted search and rescue is thus simultaneously a driver for human-robot teaming. That's where we stand now – and that's where we are, honestly, standing more or less still. We can take that negatively, and positively, of course. In a team, actors need to address issues at the levels of *task-work* (what they are doing as a team), and *team-work* (how they are doing so as a team) (cf. G. Kaminka). Especially if robots are to perform autonomously, both task-work and team-work require communication. Communication is *key*. Acting autonomously there is all the more reason for robots to keep human actors informed about progress. Without that information, there is the risk that human actors fail to maintain an accurate operational picture shared with the robot [7], [13], [18], [29]. Failing that, human trust in a robot deteriorates, and results in humans reverting to tele-operating robots if and when possible – exactly the situation we are trying to overcome. It is one instance of the human-robot interaction bottleneck [24].

The rabbit hole goes deeper than that, though. It's not just about the robot adequately communicating what is going on. That is task-work. Beyond that, we also have the team-work per se to deal with. And that is where most work on communication, particularly situated dialogue processing in human-robot interaction, stops. The symbol grounding problem is presumably solved [27], we can talk about places, objects, what to do where and how (see e.g. [19], [30]). But, again: that is task-work. Frankly, it is mostly just another form of tele-operating a robot, according the robot only very limited degrees of autonomy notably in decision-making. Arguably, this is where the much deeper issue is with human-robot interaction as a bottleneck to successful robot-assisted search and rescue. We need to deal with the *collaboration*, start from teaming and see situated dialogue as inherently part of a collaborative activity. This paper is not the first to raise this idea as such, see e.g. [14], [17], [28], and our recent work on setting this idea in the context of search and rescue [18].

This paper goes beyond those previous approaches in that we look closer at the issue of common ground from the viewpoint of *asymmetric agency*, and introduce the notion of *social sentience*. Asymmetric agency concerns the observation that humans and robots experience the world fundamentally differently. This changes the view on "meaning" which gets constructed in situated dialogue, from a "truth-value-based" idea to something more akin to a judgement or argument for why an actor (notably a robot) believes it can construct a particular interpretation for another actor (i.e. a human). A way to deal with this is to be more explicit about the conditions under which the robot believes its understanding to be falsifiable, or in need of further verification. We can model this potential (bearing some similarity to the clarification potential of [9]) using the formal notions of assertion [2], [17] including strong negation [15], [31]. Social sentience provides the robot with an explicit sense of its own social role within a team. We extend the previous idea of handling asymmetric agency through a formal notion of assertion, to that of social sentience. The result is a constructive notion of common ground which includes understanding at both the task- and the team-work level.

An overview of the paper is follows. §II starts with a discussion of a fundamental issue in forming common ground between humans and robots, namely that there exists no symmetry between how humans and robots experience and understand reality. The paper captures this issue through the notion of *asymmetric agency*. §III provides a brief discussion of our approach to situated dialogue as part of collaborative activity. §IV continues this discussion, focusing on how we can deal with asymmetric agency at the levels of task-work and team-work, in situated dialogue. We build up a notion of common ground in which use assertions and proofs as principal mechanism for construing coordination in situated and social understanding [4], [5]. The paper ends with conclusions.

II. Asymmetric Agency

Anybody who has ever worked with robots knows this: Robots and humans do not experience reality in the same way. Embodiment, sensors, those are all different. Naturally, we can endow a robot with a set of ontologies, to provide it with human-like concepts for modeling the environment. See for example [19]. This works reasonably well if robots operate in well-structured environments, in which the aspects modeled by these concepts can be more-or-less-easily *individuated* (recognized within observations), and interpreted using (cleanly separable) classifiers. We can reconstruct and constrain the robot's interpretation of its experience to such an extent that it results in an interpretation which, modulo uncertainty and incompleteness, could correspond to how a human would understand "things."

But this is typically a fairly narrow line. We construct within the robot a set of models presumably shared (sufficiently) with human understanding to build up meaning as referring to an objective *truth*. A truth, essentially based in human understand, from which humans and robots can then construct *common ground* as a set of mutually believed true statements. If we consider the types of model theories provided for formal theories of (dynamic) semantics and collaborative dialogue, this is what you find: Given a set of possible worlds, with each world modeling a set of true statements holding for an agent, a shared statement ϕ between agents A and A' is one that holds at a world accessible from worlds at which A and A' privately (individually) believe that ϕ.

This is a notion of *objective truth* which we believe to be untenable for human-robot interaction, as it assumes a symmetry in interpretation and experience between humans and robots. A symmetry which, as we already pointed out above, is arguably absent. Therefore, rather than basing our notion of understanding, or the meanings each actor in a human-robot team constructs, on "truth," we propose to adopt the idea of *judgement* (*op.cit* RobinCooper). A judgement is the assignment of a type to a proposition representing the meaning of an experience. The assignment of this type is based on the construction of an explicit proof, following the perspective of propositions-as-proofs [21]. The proof takes the form of an abductive inference, to reflect that the actor reasons to an *explanation* for why the experience is accorded the given meaning [10], [16], [17], [23]. Abduction in and by itself is a non-monotonic form of reasoning; what we discuss later is how, beyond the possibility to refute the conclusion of the inference as such, we also include explicit statements within the inference itself to indicate what is subject to future verification (or falsification).

This, admittedly, presents a fairly rapid flight into abstract thought, while at the same time leaving several fundamental concepts like "meaning" completely undefined. We correct this in the sections below. What is important to get clear at this point is that we start from an assumption of *asymmetric agency* for modeling human-robot interaction, particularly where it concerns collaborative activity involving a team of humans and robots:

> **Asymmetric agency**: *The notion of* asymmetric agency *characterizes a group of agents, which in connection/relation/comparison to each other experience and understand reality differently (asymmetry in understanding), possibly resulting in different reactions to, and expectations about, acting and interacting in that reality (asymmetry in acting/interacting).*

In the next section, we look in more detail at the characterization of situated dialogue processing for human-robot interaction, particularly the notion of (situated) meaning we adopt. After that, we return to the issue of building up and maintaining common ground.

III. Situated Dialogue as Collaborative Activity

Communication is key in a human-robot team, just like it is in "normal" human teams. We need it to coordinate the task work, what we are doing, whether we are carrying out our actions individually or jointly. Similarly, we need to coordinate how we will carry them, the team work, organizing the team to face the oncoming challenges. And all of that is set against a coordination or alignment of how we understand the environment itself, and how we are even supposed to carry out our tasks there. See also e.g. [13], [14].

That makes situated dialogue an inherent part of collaborative activity – and vice versa, it is difficult to regard situated dialogue simply for what it does linguistically. The meanings it helps construct go beyond the lexical meanings of words. It is about how dialogue can be anchored in an understanding of the world and what is happening there, the intentions of other actors to explain why someone is doing or saying something, and what expectations need to or can be held for what might happen "next." Meaning is situated.

By that we mean the following. Situation Semantics characterizes the truth of a proposition (for an utterance) to be relative to a situation – or rather, a set of connected situations

(a) (b)

Fig. 1. Entering church, pilot observes rubble heap, then UGV crosses it

[8]: The *utterance situation*, the *focus situation*, and possibly one or more *resource situations*. The utterance situation characterizes how the actors involved in the dialogue find themselves in the situation. This includes how they experience the situation being discussed (directly, or mediated e.g. through a user interface), and the social relations between them (e.g. a UGV operator talking to a robot she supervises). The focus situation is the situation under discussion (e.g. a rubble pile right in front of the robot), whereas resource situations cover background knowledge and previous situations which might have a (referential) bearing on talking about the current situation (e.g. the knowledge that crossing the rubble pile presents a threat to the robot, and that it is right in front of the door the robot just passed in order to get into the destroyed building being explored).

Consider Fig. 1. It shows an actual situation we encountered during a disaster response, aiding in structural damage assessment in the aftermath of the earthquakes that hit northern Italy (July 2012). The robot is about to enter the western aisle of the San Francesco d'Assisi church in Mirandola. Right behind the door there is a large pile of rubble, the result of a ceiling having caved in. The following dialogue illustrates the concepts of situated meaning.[1]

1) Situation Fig. 1 (a) H=human, R=robot
   a) H Mission Spc: "There seems to be a rubble pile in front of you."
   b) R: "Acknowledged, I can see some rubble ahead."
   c) R, to H UGV Op: "Shall I attempt to cross it?"
   d) H UGV Op: "Go ahead."
   e) R, to H UGV Op: "Okay, moving forward."
2) Situation Fig. 1 (b)
   a) R, to H UGV Op: "I am trying to cross the rubble pile."
   b) R, to H UGV Op: "However I am not sure whether I can continue."
   c) R, to H UGV Op: "Could you please take over?"

Starting with the situation in Fig. 1 (a), the Mission Spc introduces the rubble pile as the focus situation (1a). The robot acknowledges that it has received the information, and that it can ground it in something which it has observed (1b). The choice of words for the referring expression is deliberately "vague" to reflect the fact that the robot has observed something, which could classify as rubble, without meaning to indicate it has seen exactly what the human means. In (1c) the focus is on the *intended action* event, using the previous focus on the observed situation now as a *resource* situation for the pronominal reference "it." Also, note that the utterance situation changes: The robot shifts the dialogue to the UGV operator, as this person is (currently) in charge of deciding about where the robot is to go. (1d) confirms the intention expressed earlier in (1c). Finally, (1e) uses both the initial situation of (1a) and the (now confirmed) intended action of (1c-d) together: The forward contrasts with the earlier position of the robot, towards the pile to be crossed.

[1]This dialogue is for illustration purposes only. During the actual deployment, the robots were tele-operated, and did not speak.

The continuation of the dialogue in (2) provides further illustrations of *team-work*, beyond simple acknowledgements as in (1b) and (1d). In (2a) the robot shifts the focus to the actual progress it is making. (2b) gives a negative valuation of that progress, shifting the focus from progress to valuation, using (2a) now as resource. (The basis for such a valuation can be that safety thresholds for autonomous 3D structure traversal are close to being violated.) Finally, in reference to the action under execution (2a) and its progress valuation (2b), (2c) shifts the focus even further, purely to the level of the social structure (team-work): The robot suggests the operator to momentarily take over execution.

The dialogue shows how the different aspects of team- and task-work get intertwined in a form of understanding which ties together physical situations, actions, social structure, and communication per se. The system described in [7] provides the basic sources of information we need to reason with: Continual plans and their execution, team structure, environment models up to a level of conceptual-functional understanding. In the next section, we discuss a further development of the notion of *assertion* [2], [17] to aid in reasoning with asymmetric agency, and provide a notion of social sentience to facilitate the explicit reasoning with social dynamics for communication about team-work.

IV. Common Ground, Social Sentience, and Asymmetric Agency

We use abductive inference as the basic mode of reasoning in situated dialogue processing [10], [17], [28]. This type of inference derives an explanation: Δ explains why we believe Γ can happen given our knowledge Σ, i.e. $\Sigma \wedge \Delta \models \Gamma$.

There are two crucial things to observe here. First, we focus explicitly on the *proof* underlying the conclusion, $\Pi[\Sigma \wedge \Delta] \models \Gamma$. The proof steps make explicit *what* information we base the explanation on. Second, the explanation we draw is a *judgment*: We infer that Δ is of a particular *type* t, $\Delta[t]$. As a type it has an internal structure, rather than that it has an objective truth (i.e. a truth value in a model shared by the different actors involved) [6], [21].

Proofs draw from various sources of information to construct their conclusions. In keeping with the characterization of meaning outlined above, we can see that a proof essentially circumscribes a situation in which a certain set of actions is to be, or has been, performed, to achieve an inferable goal. It appeals to information constituting a focus (relative to which a goal is to be achieved), several resources (beliefs about the world, and what other actors might believe [20]; existing plans), as well as a dynamic social structure (e.g. knowledge about actions; roles, their needs and obligations [7]). See also [11], [17] for examples.

As the collaboration progresses, we thus get a sequence of proofs: Proofs explaining how the robot can achieve a particular goal (collaborative action selection and -planning), linked to proofs explaining why a human actor is doing what she is doing (intention recognition). By appealing to situations, these proofs build up a dynamic structure or "universe" over how the robot believes these situations hang together. We can first of all consider this at the level of the dynamics of these situations themselves. Consider σ to be a situation, in the sense of characterizing a focus, a social structure, and (pointers to) reference situations. Furthermore, let α be the non-empty sequence of actions implied by a proof $\Pi[\Sigma \wedge \Delta]$ to help establish the goal $\Delta[t]$. Then, if we understand $\sigma[(\Pi[\Sigma \wedge \Delta])\alpha]$ in the dynamic sense, that is apply the sequence of actions α resulting from Π to (or "in") the situation σ, we should get to a new situation σ' in which the goal $\Delta[t]$ "holds."

More precisely still, the result of the application of α to σ typically is a sequence of situations, of which σ' is only the end-result. And the proof makes explicit, what the information the inclusion of these actions in the inference is based on. Now, given that robots invariably need to act under uncertainty and incomplete knowledge, we need to address this in our inferences. [11], [17] show how uncertainty can be included by constructing a probabilistic version of weighted abduction [10]. They also show how a basic form of incomplete knowledge can be dealt with through the notion of *assertion*, similar to [2]. An assertion is a (logical, probability-weighted) statement about a piece of information which *is* needed to construct the proof, but for which the robot has neither positive nor negative indications. An example is the assertion that there is a door, to gain access into a building, if the goal is to explore the inside of a building. If this assertion turns out to be falsified (i.e. there is no door), we need to reconsider the course of actions to be taken. In continual planning, assertions are therefore used as explicit points in a plan at which re-planning may be needed.

Here, we suggest to extend the notion of assertion, and the (existentially closed) logical language for constructing proofs with the notion of strong negation [31]. Whereas the classical notion of negation basically entails a failure to prove, strong negation states something explicitly as not possible or justified. Strong negation has been considered in several approaches to knowledge representation, to include an explicit notion of (closed) falsifiability – which we can now put "opposite to" the notion of assertion as an explicit notion of (open) verifiability. Strong negation says something cannot be the case on the basis of what is known (or the proof fails), where an assertion states that something is assumed to be the case on the basis of what is essentially *not* known (or, again, the proof fails).

If we now look back at our proofs, as judgements anchored to a complex structure over situations, we thus see that with the inclusion of assertions and strong negation we obtain a framework in which we can represent and reason with the asymmetry inherent to a human-robot team. First of all, attributed and shared beliefs become judgements based in proofs which can be qualified with statements about explicit verifiability and falsifiability. That changes these beliefs from "true statements" into subjective judgements about others, presumed to hold under the continual observations of the other's actions. And if a proof turns out to become invalidated (assertion- or strong negation-wise), this is then immediately traceable to the beliefs these proofs are based on, indicating what needs to be retracted.

We can take this a step further though. There is no reason why we can only reason about beliefs, and how these beliefs lead to actions, already observed or observable. We can *lift* verifiability/falsifiability to the level of intentional reasoning, and reason about what we expect to do or not to do, in the light of what is necessary to do. Fig. 2 illustrates a team structure. It describes roles familiar from [3], [25], identifying for each role the kinds of team- and task-work related information it is responsible for, in relation to another role. For example, the UGV platform (role) reports status to the pilot (playable by either a human or a robot), and the pilot can order the platform to go to a particular location.

With the constructions at hand, we can define an additional level of proofs. This level essentially captures the team work. Each proof is cast as a temporal sequence of actions, with accompanying references to situations, and with explicit verifiable/falsifiable references to the achievability of specific goals by (or through) specific agents. These latter goals in and by themselves can again be translated into proofs, anchoring them in the actual situations. This is crucial: It enables to anchor the team work in the ongoing task work set in a dynamic environment, and it makes it possible to reason about how the team can actually achieve its goals together. The collaboration in (1)-(2) already illustrated this concept. An initial proof planned for autonomously crossing the rubble, under the strong assumption that safety thresholds would not be crossed. This allowed the robot to play the role of Pilot, and the human to be the Pilot-in-Command. Once the safety thresholds were crossed though, this required a revision of the plan. The responsibility was transferred to the human, to pilot the platform across the rubble. This leads to a possibility to deal with what we define here as *social sentience*:

> **Social sentience**: *The notion of* social sentience *implies a capability for an actor to reason explicitly with its role within a social structure, how the assumption of this role requires the assumption of certain responsibilities (goals to be achieved) with respect to other roles – and how the inability to fulfill some or all of these responsibilities may require shifting such responsibilities to other actors, resulting in a shift of roles within the social structure.*

V. Conclusions

The paper outlined an approach to dealing with the problem of asymmetric agency in forming common ground between actors in a human-robot team, particularly as regarded from the viewpoint of the robot. The claim is that this is one of the most fundamental problems to overcome regarding the human-robot interaction bottleneck in robot-assisted USAR. We argued that the (logical-probabilistic) means for doing so, based on the notions of assertion and strong negation, introduce general

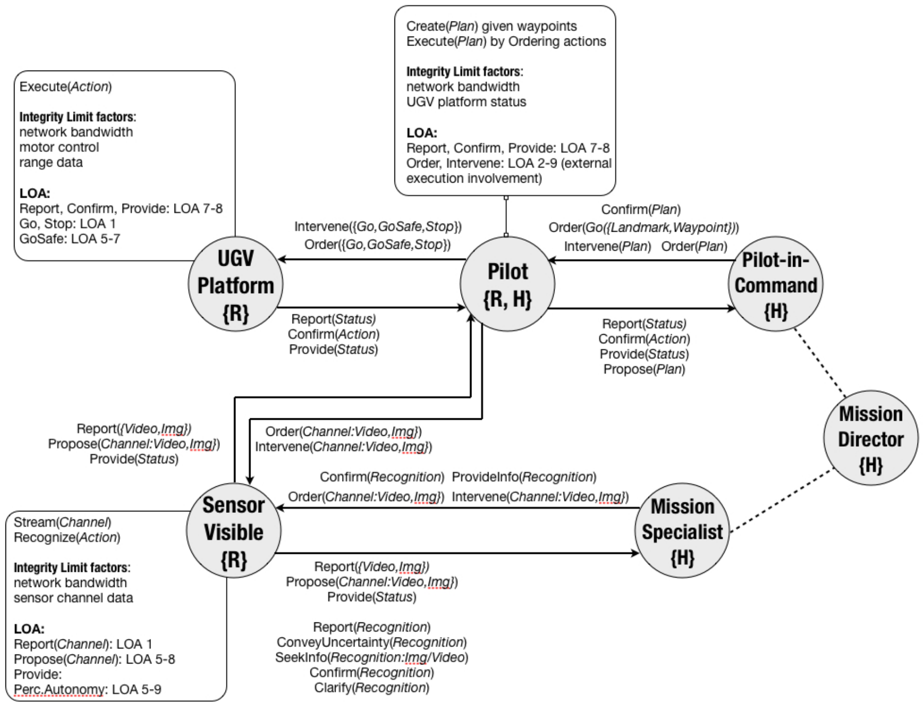

Fig. 2. Team roles with "responsibilities" as upward and downward actions/information needs

concepts of verifiability and falsifiability under incomplete knowledge. We then discussed how these concepts can be lifted to a higher level of reasoning about the dynamics of social structure in a human-robot team, to provide the robot with an explicit notion of social sentience.

At the moment, we are working on a novel continual planner to provide an implementation of the approach described here. The planner is based on multi-agent continual planning with assertions [2]. The planner works of a notion of "state" formulated as a world situation, covering the social situation (or the utterance situation), as well as a focus and set of resource situations. The planner works in conjunction with an execution monitoring strategy, and an "engine" which updates state based on planner-external sensory perception.

Acknowledgments

The work reported in this paper was supported by the EU FP7 Integrated Project NIFTi, "Natural human-robot coordination in dynamic environments" (Grant #247870). For more on NIFTi, see `http://www.nifti.eu`.

References

[1] A. Birk and S. Carpin. Rescue robotics - a crucial milestone on the road to autonomous systems. *Advanced Robotics*, 20(5):595–605, 2006. 1

[2] M. Brenner and B. Nebel. Continual planning and acting in dynamic multiagent environments. *Journal of Autonomous Agents and Multiagent Systems*, 19(3):297–331, 2009. 1, 3, 4

[3] J.L. Burke, R.R. Murphy, M. Coovert, and D. Riddle. Moonlight in Miami: An ethnographic study of human-robot interaction in USAR. *Human Computer Interaction*, 19((1–2)):85–116, 2004. 4

[4] H. Clark. *Using Language*. Cambridge University Press, 1996. 2

[5] H. H. Clark and S. A. Brennan. Grounding in communication. In L.B. Resnick, J.M. Levine, and S.D. Teasley, editors, *Perspectives on socially shared cognition*, pages 127–149. APA Books, Washington VA, 1991. 2

[6] R. Cooper. Records and record types in semantic theory. *Journal of Logic and Computation*, 15(2):99–112, 2008. 3

[7] G.J.M. Kruijff *et al*. Experience in system design for human-robot teaming in urban search & rescue. In *Proceedings of Field and Service Robotics (FSR) 2012*, Matsushima/Sendai, Japan, 2012. 1, 3

[8] J. Ginzburg. Situation semantics and the ontology of natural language. In P. Portner, C. Maierborn, and K. von Heusinger, editors, *The Handbook of Semantics*, pages 830–851. de Gruyter, 2011. 3

[9] J. Ginzburg. Situation semantics: From indexicality to metacommunicative interaction. In P. Portner, C. Maierborn, and K. von Heusinger, editors, *The Handbook of Semantics*, pages 852–872. de Gruyter, 2011. 1

[10] J.R. Hobbs, M.E. Stickel, D. Appelt, and P. Martin. Interpretation as abduction. Technical Report 499, AI Center, SRI International, Menlo Park, CA, USA, Dec 1990. 2, 3, 4

[11] M. Janíček. Abductive reasoning for continual dialogue understanding. In *Proceedings of the ESSLLI Student Session 2011*, Ljubljana, Slovenia, August 2011. 3, 4

[12] M. Johnson, J.M. Bradshaw, P.J. Feltovich, R.R. Hoffman, C. Jonker, B. van Riemsdijk, and M. Sierhuis. Beyond cooperative robotics: The central role of interdependence in coactive design. *IEEE Intelligent Systems*, pages 81–88, May/June 2011. 1

[13] G. Klein, P.J. Feltovich, J.M. Bradshaw, and D.D. Woods. Common ground and coordination in joint activity. In W.B. Rouse and K.R. Boff, editors, *Organizational Simulation*, pages 139–184. John Wiley, New York City, NY, 2004. 1, 2

[14] G. Klein, D.D. Woods, J.M. Bradshaw, R. Hoffman, and P. Feltovich. Ten challenges for making automation a team player in joint human-agent activity. *IEEE Intelligent Systems*, 19(6):91–95, November-December 2004. 1, 2

[15] M. Kracht. On extensions of intermediate logics by strong negation. *Journal of Philosophical Logic*, 27:49 – 73, 1998. 1

[16] G.J.M. Kruijff. Perice's late theory of abduction: A comprehensive approach. *Semiotica*, 153(1/4):431–454, 2005. 2

[17] G.J.M. Kruijff, M. Janíček, and P. Lison. Continual processing of

situated dialogue in human-robot collaborative activities. In *Proceedings of the 19th International Symposium on Robot and Human Interactive Communication (RO-MAN 2010)*. IEEE, 2010. 1, 2, 3, 4

[18] G.J.M. Kruijff, M. Janíček, and H. Zender. Situated communication for joint activity in human-robot teams. *IEEE Intelligent Systems*, 2012. forthcoming. 1

[19] G.J.M. Kruijff, H. Zender, P. Jensfelt, and H.I. Christensen. Situated dialogue and spatial organization: What, where... and why? *International Journal of Advanced Robotic Systems*, 4(2), 2007. 1, 2

[20] P. Lison, C. Ehrler, and G.J.M. Kruijff. Belief modelling for situation awareness in human-robot interaction. In *Proceedings of the 19th International Symposium on Robot and Human Interactive Communication (RO-MAN 2010)*, 2010. 3

[21] P. Martin-Löf. *Intuitionistic Type Theory*. Bibliopolis, Napels, Italy, 1984. 2, 3

[22] D.A. Mindell. *Digital Apollo: Human and Machine in Spaceflight*. MIT Press, 2008. 1

[23] M.P.Shanahan. Perception as abduction: Turning sensor data into meaningful representation. *Cognitive Science*, 29:103–134, 2005. 2

[24] R.R. Murphy. Human-robot interaction in rescue robotics. *IEEE Transactions on Systems, Man and Cybernetics Part C: Applications and Reviews*, 34(2):138–153, 2004. 1

[25] R.R. Murphy and J.L. Burke. The safe human-robot ratio. In M.J. Barnes and F. Jentsch, editors, *Human-Robot Interactions in Future Military Operations*, Human Factors in Defence, pages 31–49. Ashgate, 2010. 1, 4

[26] R.R. Murphy, S. Tadokoro, D. Nardi, A. Jacoff, P. Fiorini, H. Choset, and A.M. Erkmen. Search and rescue robotics. In B. Siciliano and O. Khatib, editors, *Springer Handbook of Robotics*, pages Part F, 1151–1173. Springer Verlag, 2008. 1

[27] L. Steels. The symbol grounding problem has been solved. so what's next? In M. De Vega, G. Glennberg, and G. Graesser, editors, *Symbols, embodiment and meaning*. Academic Press, New Haven, 2008. 1

[28] M. Stone and R.H. Thomason. Coordinating understanding and generation in an abductive approach to interpretation. In *Proceedings of DIABRUCK 2003: 7th workshop on the semantics and pragmatics of dialogue*, 2003. 1, 3

[29] K. Stubbs, P.J. Hinds, and D. Wettergreen. Autonomy and common ground in human-robot interaction: A field study. *IEEE Intelligent Systems*, 22:42–50, March 2007. Special Issue on Interacting with Autonomy. 1

[30] S. Tellex, T. Kollar, S. Dickerson, M.R. Walter, A. Gopal Banerjee, S. Teller, and N. Roy. Understanding natural language commands for robotic navigation and mobile manipulation. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, San Francisco, CA, 2011. 1

[31] H. Wansing. Negation. In L. Goble, editor, *The Blackwell Guide to Philosophical Logic*, pages 415–436. Basil Blackwell Publishers, Cambridge, MA, 2001. 1