



HAL
open science

Commitments in Human-Robot Interaction

V Fernandez Castro, Aurélie Clodic, Rachid Alami, Elisabeth Pacherie

► **To cite this version:**

V Fernandez Castro, Aurélie Clodic, Rachid Alami, Elisabeth Pacherie. Commitments in Human-Robot Interaction. 2019. ijn_03084105

HAL Id: ijn_03084105

https://hal.science/ijn_03084105

Preprint submitted on 20 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Commitments in Human-Robot Interaction

V. Fernandez Castro^{1,2}, A. Clodic¹, R. Alami¹, E. Pacherie²

¹ LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France

² Institut Jean Nicod, CNRS, ENS, PSL Research University, Paris, France
victor.fernandez-castro@laas.fr

Abstract

An important tradition in philosophy holds that in order to successfully perform a joint action, the participants must be capable of establishing commitments on joint goals and shared plans. This suggests that social robotics should endow robots with similar competences for commitment management in order to achieve the objective of performing joint tasks in human-robot interactions. In this paper, we examine two philosophical approaches to commitments. These approaches, we argue, emphasize different behavioral and cognitive aspects of commitments that give roboticists a way to give meaning to monitoring and pro-active signaling in joint action with human partners. To show that, we present an example of use-case with guiding robots and we sketch a framework that can be used to explore the type of capacities and behaviors that a robot may need to manage commitments.

Introduction

A central challenge in human-robot interaction is to devise robots capable of interacting with humans to perform so-called joint actions, social interactions where two or more participants coordinate their actions in space and time to bring about a change in the environment (Sebanz, Bekkering, and Knoblich 2006). The notion of commitment plays a pivotal role in understanding joint action (Cohen and Levesque 1991; Bratman 1992; Gilbert 2009). In order to perform a joint action, co-agents need to establish and handle commitment regarding the shared goals and plans involved in the joint action. However, the study of commitments in the context of human-robot collaborative task achievement is quite new.

As a first approximation, a commitment is in place when an agent (the author of the commitment) has given an assurance to another agent (the recipient of the commitment) that she will act in a certain way, when, as a result, the recipient has formed an expectation regarding the actions of the author, and when there is mutual knowledge by both parties that this is the case. In order to build our framework of commitments, we discuss two specific philosophical approaches that emphasize different, though complementary, aspects of

commitments: (1) *The functional approach* and (2) *The normative approach*. The first emphasizes the function of commitments, how expectations reduce different forms of uncertainty. The second emphasizes how such a function is carried out; namely, by creating obligations towards one's co-agents and their corresponding entitlements to demand that these obligations be satisfied. In order to demonstrate how these approaches can give hints on the implementation of strategies for HRI, we consider a use-case, in which a robot guide takes a visitor to its destination (e.g. Rackham at the Space City Museum (Clodic et al. 2006)). Although there is already previous research on commitments in social robotics, this work is a first step towards generating a framework that gives theoretical support to this previous research and at the same time generates new avenues for future investigations.

Two Philosophical Aspects on Commitments

Establishing commitments requires the participants involved to generate reliable expectations about the content of the commitment, for instance, making a promise that one will perform a particular objective or giving implicit signals that one act in a particular way. But why do we generate such expectations? Why do we need to establish such commitments?

The Functional Approach

Arguably, what distinguishes a joint action (e.g. walking together) from the mere coordination of behavior (two people randomly walking at the same pace) is that the partners are committed to achieve a goal together (Gilbert 2006). Furthermore, such a participatory commitment comes often with other contralateral commitments (Roth 2004) associated to the general goal or the plan necessary to achieve it (e.g. not speeding up, or waiting for your partner if he's left behind).

The necessity of establishing these different commitments comes from the fact that even assuming that the partners desire the outcome to be the case, different sources of uncertainty could prevent the joint action from being carried out. As Michael and Pacherie (2014) argue, participants can face three sources of uncertainty during joint action:

1. **Motivational uncertainty:** participants may not know whether or not the partner is motivated to engage in the joint action. Thus, they cannot be sure how convergent their respective interests are or whether or not they shared a goal.
2. **Instrumental uncertainty:** participants may not know whether or not they agree about how to proceed; for instance, how the participants must distribute their roles or when and where they must act.
3. **Common ground uncertainty:** it may happen that the instrumental beliefs and motivations are not mutually manifested. Thus, even if the participants share a goal or they agree about how to proceed, they might not know that this is the case.

But how do commitments reduce these sources of uncertainty? According to the *functional approach* supported by Michael and Pacherie (2014), generating commitments reduces uncertainty by stabilizing expectations regarding actions, beliefs and motivations and thus, facilitating the prediction of the other co-actors, which makes the joint action more fluent and efficient. In their view, the key aspects of commitments are related to how the expectations are generated, and thus, how the commitments are established. For instance, how co-agents generate expectations through implicit and explicit communication (Austin 1962; Siposova, Tomasello, and Carpenter 2018) or repetition (Gilbert 2006; Michael, Sebanz, and Knoblich 2016).

The Normative Approach

In a recent paper, Fernandez Castro and Pacherie (2019) have focused on a different aspect of commitments, namely, how they create obligations towards one's co-agents who are entitled to demand that these obligations be satisfied, giving rise to expectations that the agent will act as committed or that, if not, co-agents will demand that she does. This view shares with the functional approach the idea that expectations are the central element of commitments. However, they argue that the expectations generated by commitments are of a special kind: *normative expectations*.

In philosophy of mind, several authors (Greenspan 1978; Paprzycka 1999; Wallace 1998) have emphasized that we can distinguish two types of expectations:

1. **Descriptive expectations:** expectations whose violation or frustration does not necessarily triggers reactive attitudes. These expectations are tied to predictions. For example, you can expect your friend to have a beer because this is what she always does but if she doesn't, this may surprise you but not bother you.
2. **Normative expectations:** expectations whose violation or frustration triggers reactive attitudes: blame ¹, request for justification or sanctions. These types of expectations are connected to the notion of holding someone on a demand. That is, the normal response to its frustration is to

¹The notion of blame must be understood in very general terms and without moral burden in here; that is, as a reaction that serves to sanction negatively a conduct, whether this conduct is a violation of a moral norm or not.

impose a negative reaction to the other agent for not acting as expected. Such negative reactions are more emotionally loaded and aimed at regulating the other's behavior. For instance, you can feel entitled to sanction your friend when he frustrates your expectation that he will cede his seat to an older person on the subway.

On the *normative view*, commitments serve to create for oneself, or the other agent, normative expectations, so we can ensure that every participant behaves as expected. In this view, the key aspect of commitments is how we display different communicative and behavioral strategies to make explicit our obligations towards our partner and their own obligations towards us; for instance, by blaming, making the other aware of what is expected or apologizing when something goes wrong.

At this point, it is worth noticing that the two approaches are not necessarily incompatible. In fact, to the extent that they emphasize different aspects of how agents handle commitments, they can be regarded as complementary. In our view, managing commitments involves both generating expectations that facilitate the reduction of uncertainty and displaying repairing and regulative strategies when such expectations are frustrated.

Related work

Several computational models of commitment management have already been proposed. Joint intention theory (Cohen and Levesque 1990) introduces the notion of commitment in relation to the goal of the collaborative task: "if a team is jointly committed to some goal, then under certain conditions, until the team as a whole is finished, if one of the members comes to believe that the goal is finished but that this is not yet mutually known, she will be left with a persistent goal to make the status of the goal mutually known". Shared Plans theory (Grosz and Kraus 1996) uses individual intentions to establish commitment of collaborators to their joint activity. Interestingly, Shared Plans theory proposes also the notion of agent's commitments to its collaborating partners' abilities, as well as the notion of helpful behavior and contracting action. Other works emphasize grounding as a kind of commitment (Clark 1996; Clark 2006): "Once we have formulated a message, we must do more than just send it off. We need to assure ourselves that it has been understood as we intended to be." Clark also proposes to consider on one side *basic joint activity or joint activity proper* and on the other side the *coordinating joint actions* based on *communicative acts* (Clark 2006). In the same vein, Klein et al. (2005) describes "basic compact" as a kind of contract that need to be shared by the cooperative agents.

Several implementations using commitments have also been proposed ², for instance, Collagen architecture (Rich, Sidner, and Lesh 2001; Sidner et al. 2005) in collaborative conversation. In human-robot interaction, we can find implementations of commitments in different settings such

²In the robotic domain, it is the word "*engagement*" and not "*commitment*" which is often used. For the sake of simplicity we will continue using the term "commitment" in the paper

as in Human-Robot Interaction Operating System (Fong et al. 2005), robot behavior toolkit (Huang and Mutlu 2013), robot tutelage (Hoffman and Breazeal 2007) or linked to supervision, communicative acts and grounding (Clodic et al. 2007), (Clodic 2007).

In the framework of Rackham, a robot guide at the Space City museum³ from 2002 to 2005 (Clodic et al. 2006), we also already used the notion of commitments. In that experiment, we defined a set of commitments to the task : Rackham commitment to the task (obtained when the robot decided to do the task), Rackham's belief concerning the visitor's belief about Rackham commitment to the task (obtained once the robot told the visitor that it will perform the task), Rackham's belief concerning the visitor commitment to the task (obtained once the visitor agrees to do the task). However, commitments were mostly used at task setup and not really handled during task execution.

The model we propose hereafter takes inspiration from this work. It could be seen as complementary to these models, theories and implementation as it helps to focus on the way we could define and manage commitments during the execution of the task and not just on how commitments can be established at the beginning of the task. The Rackham experiment (Clodic et al. 2006) allowed us to note the importance of the quality of interaction. We noticed that: "a continuous interaction all along the mission is fundamental to keep visitor interest. The robot must continuously check the presence of the guided person, show him that it knows where (s)he is (...)". At that time, it was difficult for us to model that since it was not really part of the task by itself and was done in an intuitive way. In one sense, we think that the theoretical framework we propose here will help us to capture and flesh out this intuition.

Implementing Commitment

Framework

We have not yet implemented the proposed framework. Here, we are introducing it using an example, a robot guide such as Rackham at the Space City Museum. In that example, once the destination has been chosen, there should be a joint task of "go to the destination" where Rackham should "execute trajectory" whereas the visitor should "follow the robot". Regarding commitments, we propose that managing commitments is a social skill which requires capacities that are independent of the capacities required to perform the particular actions in place (e.g. *execute trajectory* for Rackham and *follow the robot* for the visitor). In our case, that means that commitment management could be modeled as something aggregated to the actions of the task and linked to expectations.

Given the lessons of the functional and the normative approach, we advance that the function of commitment as a reduction of uncertainty and its normative aspect must be two fundamental elements of the framework. Such elements are present in two different moments of the joint action. First, its function as a reduction tool is exhibited in the repertoire of

actions dedicated to insuring the mutual understanding that the commitment is in place before and during the interaction. Second, its normative aspect becomes evident when the commitment is frustrated and certain reactions, responses, and repairs became necessary. In what follows, we present how our framework would depict these two moments.

Generating Commitments

Considering a commitment as a reduction tool implies that ascribing or representing a partner as undertaking a commitment must be tied to a repertoire of actions aiming at insuring the establishment of a mutual understanding of the expectations involved, e.g. social signals and cues or regulative actions. Such actions can vary depending on whether or not the actor is the agent expecting the other to perform the action (assigning a commitment) or the agent who should act as expected (undertaking the commitment). In our example, Rackham can expect the human to follow him and undertake the commitment of adapting the speed to human pace (which can also be modeled as the human expectation regarding the robot). For a sake of simplicity, from the robot point of view, such difference can be represented in terms of assignations of expectations:

- Expectation (*Human follows*)
- Expectation (*Robot adapts speed*)

Given the function of commitments, the robot must be able to actively provide different cues that facilitate human understanding of the expectations in place, but also, he must be able to monitor that the human is behaving as he should (See Figure 1). In the guiding example, when Rackham expects the human to follow (Expectation (*Human follows*)), he must possess the following capacities:

- **Exhibiting/Signaling the expectation** Rackham must provide reliable cues regarding what he expects from the visitor in order to make him aware of it. One way of establishing such a signal is through verbal communication, for instance, saying "I accompany you to your destination, please follow me".
- **Monitoring:** Rackham must be able to recognize the visitor's action and be sure he is behaving as expected, that is, monitoring that the visitor is following.

Certainly, the type of reliable signals in place varies with the assignation of expectations. In this case, when the expectation involves the robot itself, he must be sure that the human knows what he is going to do, so the human does not face any type of uncertainty regarding robots behavior. Furthermore, the dynamic nature of joint actions requires the participants to reinforce the pertinent expectations. In our example, thus, assignation of expectations to the human (*Robot adapts speed*) must be tied to the following capacities:

- **Anticipating Expectation:** Rackham must make clear that he is capable of adapting to the speed to the human, for instance, using verbal expressions like indirect speech acts (Are you in a rush? We can go faster if you want).

³<http://www.laas.fr/robots/rackham>

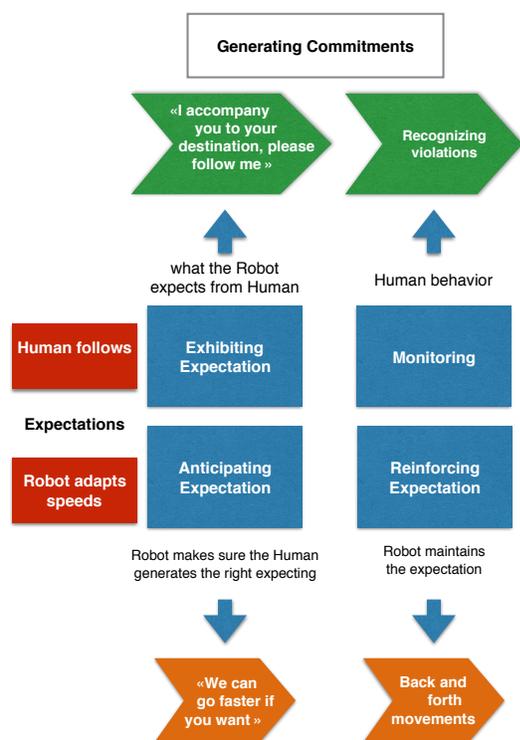


Figure 1: Framework for generating commitments for a guiding task: 1) When the robot assigns a commitment to the human (the human must follow (*Human follows*)), the following capacities are required: **Exhibiting expectations** to make the other aware of her obligations; **Monitoring expectations**: observe and analyse the other agent behavior and detect if she violates expectations. 2) When the robot assigns a commitment to itself (he must adapt speed (*Robot adapts speed*)), the following capacities are required: **Anticipating expectations** to make the human aware of the robots obligations ; **Reinforcing expectations**: to make the human aware that the robot that he keeps committed to the task.

- **Reinforcing Expectation**: Rackham must maintain the human expectation that the robot will adapt to his speed. For instance, exhibiting some back and forth movement that makes the human aware that Rackham is still committed to the task.

Repairing

Of course, expectations can be frustrated or fulfilled, and this is where the normative aspects of commitments are made explicit. As the normative approach emphasizes, commitments enable participants to hold others on demand and regulate and influence their actions when a particular action is frustrated. This gives to social robotics an important element to improve social interaction by using commitment management as a mechanism for reparation. Given our representation of expectations (Expectation (*Human follows*)/Expectation (*Robot adapts speed*)), we can assume two plausible violations in the context of our example:

- The violation of Expectation (*Human follows*): the visitor

does not follow.

- The violation of Expectation (*Robot adapts speed*): the robot is going too slow (or too fast) or the human bump into the robot

Now, considering the first type of violation, the robot must be able to execute two different actions: reaction and repair. The reaction is necessary to make the human understand that the expectation has been frustrated, and thus, that he is responsible for his commitment. Furthermore, such a reaction must be a facilitator for negotiating a re-engagement or repair. In our example, we can implement these two actions as follows:

- **Reaction**: when the expectation is frustrated (e.g. the visitor does not follow), Rackham must be able to *react* to the violation in a way that makes the visitor aware that he should be following. For instance, Rackham could stop, look back and adopt a posture of wait-and-see attitude.
- **Repair**: Rackham must be equipped with *repairing* strategies that facilitate the re-engagement in the task, for instance, approaching the visitor and asking whether everything is all right. Such a reparation would facilitate the conclusion of the general task, but also, it would make the human aware of his obligations to the robot, reinforcing the motivation for interaction.

Concerning the violation of the robot expectation, the necessary abilities of the robot should differ in several respects. In addition to being able to react and repair in a different way to when it is the human who is at fault, the robot must be able to recognize plausible signals that it is not behaving as expected⁴. In our example, these capacities can be regarded as follows:

- **Recognition of reactions**: when the expectation is frustrated (e.g. Rackham goes too fast or too slow), the robot must be able to notice when the visitor does not follow anymore or, on the contrary, bumps into the robot because it is too slow. Moreover, Rackham must be able to recognize some natural reactions to violation of expectations, like verbal petitions such as “Hey!” or “Wait for me!”.
- **Reaction**: the recognition that the robot has violated an expectation must trigger an excusing reaction to make the visitor aware that it knows that it has violated the commitment; for instance, apologizing to the visitor.
- **Repair**: the reaction must be followed by a repair strategy of the commitment to ensure that the joint action takes place; for instance, slowing down or speeding up according to the human’s needs.

Future Work

In the previous sections, we have proposed that designing social robots with the capacity for managing commitments can facilitate human-robot interaction. Our central

⁴Certainly, we could provide more autonomy to the capacity of the robot to handle commitments by given the capacity of evaluating the situation and recognize that he may not be fulfilling the expectation, and thus, reacting and repairing before the human expresses his discontent.

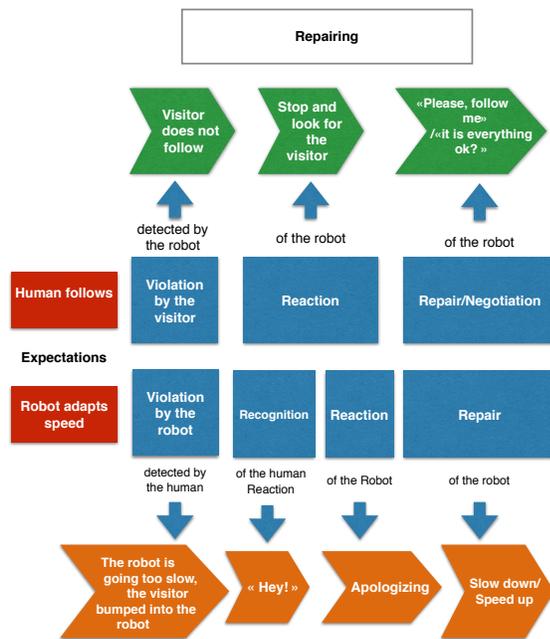


Figure 2: Framework for repairing when a case of violation is detected: 1) When the robot detects a violation of the human, he must **React** to violation of expectations to make the other aware that their obligations are not satisfied; **Repair** the joint action. 2) In the case of the robot violation, he must **Recognize reactions** in order to detect the violation; **React** to make the human aware that he knows that he made a mistake; and **Repair** the joint action.

contention is that such a capability, understood as necessary for, but independent of, joint action, plays a central role in the reduction of different types of uncertainties and as a way of enabling participants to react to frustrations of expectations that may undermine the task. Drawing on philosophical approaches to commitments, we have proposed that a commitment management mechanism must be composed of different sub-capabilities for signaling, recognizing signals, monitoring actions, reacting to violations, and repairing different expectations on both the human and robot side. In this section, we would like to propose two lines of development that, we believe, could enrich our framework and open future research avenues.

First, an important way to in which our framework could be extended is by considering the capacity for selecting different signaling strategies depending on the partner. Conforming to expectations or commitments strongly depends on a variety of motivational sources including emotions (Fehr and Camerer 2007; Sugden 2000; Colombo 2014), pro-social dispositions (Godman 2013), need to belong (Fernandez Castro and Pacherie 2019) or reputation management (Michael, Sebanz, and Knoblich 2016). These types of motivations are important when deciding on and selecting the appropriate signal for a particular expectation. For instance, a robot should be capable of deciding on the most appropriate way to ask someone to do something (e.g. using a flat or jovial tone of voice) depending on the assumed mo-

tivation in the context (e.g. reputation in a context of work vs pro-social disposition in a playful context). Moreover, ideally, taking into account these motivational components would not only help the robot to select the appropriate signal for a particular expectation but also, to identify and adapt to the particular type of signal used by the human. Being capable of distinguishing someone's motivation and readiness to interact (e.g. distinguish a good teammate from a bad teammate) would provide the robot with better means of adapting and reacting to possible violations during the joint action.

Second, another future direction of research for our framework would be to explore more robust and flexible ways of dealing with the reactions to violations of expectations. Following the functional approach to the notion of commitment, we can advance that an important number of frustrations of expectations would depend on one of the different types of uncertainties (motivation, instrumental or common ground). One plausible improvement of our framework would be dedicated to designing robots capable for determining which source of uncertainty causes the violation and being able to select different repair strategies and reactions depending on the source. To give an example, if Robot could determine the reason why the visitor does not follow it, e.g. motivation: the visitor found a friend; instrumental: the visitor found an obstacle, it could adapt its answer depending on this factor, for instance, waiting for the human, or approaching her and giving assistance.

Conclusion

Managing commitments is a central aspect of joint action, and thus, it should be taken into consideration in the design of social robots. In order to do that, we have explored how different philosophical approaches to commitment can provide useful suggestions to roboticists for designing more efficient robots capable of engaging in joint actions. We have presented two complementary proposals that emphasize different behavioral and cognitive aspects of commitments that we must concentrate on for the study of joint action for human-robot interaction.

To design capacities for managing commitments, we must give to the robot abilities to signal its expectations on the one hand and to monitor humans expectations, reactions and signals on the other. We must also give the human different means of understanding the robot's expectations, reactions and signals on the one hand and of conveying his expectations to the robot on the other. These monitoring and signaling abilities should not be considered part of the action in itself but are needed for the interaction. In this sense perhaps, we should design a kind of code to help the management of these aspects on both the robot's and on human side (e.g. a kind of common interface that is understandable by every robot).

Finally, we have advanced two different research avenues that may enrich our framework. First, one may explore how to design more flexible and adaptable signals depending on the assumed source of motivation behind the joint action. Moreover, such a detection of motivation may help the robots to adapt in case of violation. Second, we have speculated that, when there is a violation of expectation, the source

of the violation could come from the different uncertainties explored by the functional approach. As a result, a plausible line of investigation may research how to design the ability to determine the source of violation and adapt depending on the source.

Acknowledgments

This work was supported by the ANR JointAction4HRI project (jointaction4hri.laas.fr), grant ANR-16-CE33-0017 of the French Agence Nationale de la Recherche.

References

- [Austin 1962] Austin, J. 1962. *How to do things with words*. Oxford: Clarendon Press.
- [Bratman 1992] Bratman, M. E. 1992. Shared Cooperative Activity. *Philosophical Review* 101(2):327–341.
- [Clark 1996] Clark, H. H. 1996. *Using Language*. Cambridge: Cambridge University Press.
- [Clark 2006] Clark, H. H. 2006. Social actions, social commitments. In Enfield, N., and Levinson, S., eds., *Roots of human sociality: Culture, cognition, and interaction*. New York: Berg. 126–150.
- [Clodic et al. 2006] Clodic, A.; Fleury, S.; Alami, R.; Chatila, R.; Bailly, G.; Brethes, L.; Cottret, M.; Danes, P.; Dollat, X.; Elisei, F.; and et al. 2006. Rackham: An interactive robot-guide. In *The 15th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2006, Hatfield, Herthfordshire, UK, 6-8 September, 2006*, 502509. IEEE.
- [Clodic et al. 2007] Clodic, A.; Ransan, M.; Alami, R.; and Montreuil, V. 2007. A management of mutual belief for human-robot interaction. In *2007 IEEE International Conference on Systems, Man and Cybernetics*, 15511556. IEEE.
- [Clodic 2007] Clodic, A. 2007. *Supervision pour un robot interactif: action et interaction pour un robot autonome en environnement humain*. Theses, Université Paul Sabatier - Toulouse III.
- [Cohen and Levesque 1990] Cohen, P. R., and Levesque, H. J. 1990. Intention is choice with commitment. *Artificial Intelligence* 42(2):213 – 261.
- [Cohen and Levesque 1991] Cohen, P. R., and Levesque, H. J. 1991. Teamwork. *Nous* 25(4):487–512.
- [Colombo 2014] Colombo, M. 2014. Caring, the emotions, and social norm compliance. *Journal of Neuroscience, Psychology, and Economics*.
- [Fehr and Camerer 2007] Fehr, E., and Camerer, C. F. 2007. Social neuroeconomics: the neural circuitry of social preferences. *Trends in Cognitive Sciences* 11(10):419 – 427.
- [Fernandez Castro and Pacherie 2019] Fernandez Castro, V., and Pacherie, E. 2019. The credibility problem in human robot interaction. In *27th Conference of the European Society for Philosophy and Psychology*.
- [Fong et al. 2005] Fong, T. W.; Nourbakhsh, I.; Ambrose, R.; Simmons, R.; Schultz, A.; and Scholtz, J. 2005. The peer-to-peer human-robot interaction project. In *AIAA Space 2005*.
- [Gilbert 2006] Gilbert, M. 2006. Rationality in collective action. *Philosophy of the Social Sciences* 36(1):317.
- [Gilbert 2009] Gilbert, M. 2009. Shared intention and personal intentions. *Philosophical Studies* 144(1):167–187.
- [Godman 2013] Godman, M. 2013. Why we do things together: The social motivation for joint action. *Philosophical Psychology* 26(4):588–603.
- [Greenspan 1978] Greenspan, P. S. 1978. Behavior control and freedom of action. *The Philosophical Review* 87(2):225.
- [Grosz and Kraus 1996] Grosz, B. J., and Kraus, S. 1996. Collaborative plans for complex group action. *Artificial Intelligence* 86(2):269 – 357.
- [Hoffman and Breazeal 2007] Hoffman, G., and Breazeal, C. 2007. Effects of anticipatory action on human-robot teamwork: Efficiency, fluency, and perception of team. In *2007 2nd ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 1–8.
- [Huang and Mutlu 2013] Huang, C.-M., and Mutlu, B. 2013. The repertoire of robot behavior: Enabling robots to achieve interaction goals through social behavior. *J. Hum.-Robot Interact.* 2(2):80–102.
- [Klein et al. 2005] Klein, G.; Feltovich, P. J.; Bradshaw, J. M.; and Woods, D. D. 2005. *Common Ground and Coordination in Joint Activity*. John Wiley and Sons. chapter 6, 139–184.
- [Michael and Pacherie 2014] Michael, J., and Pacherie, E. 2014. On Commitments and Other Uncertainty Reduction Tools in Joint Action. *Journal of Social Ontology* 1(1):89–120.
- [Michael, Sebanz, and Knoblich 2016] Michael, J.; Sebanz, N.; and Knoblich, G. 2016. The Sense of Commitment: A Minimal Approach. *Frontiers in Psychology* 6.
- [Paprzycka 1999] Paprzycka, K. 1999. Normative expectations, intentions, and beliefs. *The Southern Journal of Philosophy* 37(4):629652.
- [Rich, Sidner, and Lesh 2001] Rich, C.; Sidner, C. L.; and Lesh, N. 2001. Collagen: Applying collaborative discourse theory to human-computer interaction. *AI Magazine* 22(4):15.
- [Roth 2004] Roth, A. S. 2004. Shared agency and contralateral commitments. *The Philosophical Review* 113(3):359–410.
- [Sebanz, Bekkering, and Knoblich 2006] Sebanz, N.; Bekkering, H.; and Knoblich, G. 2006. Joint action: bodies and minds moving together. *Trends in Cognitive Sciences* 10(2):70–76.
- [Sidner et al. 2005] Sidner, C. L.; Lee, C.; Kidd, C. D.; Lesh, N.; and Rich, C. 2005. Explorations in engagement for humans and robots. *Artificial Intelligence* 166(1):140 – 164.
- [Siposova, Tomasello, and Carpenter 2018] Siposova, B.; Tomasello, M.; and Carpenter, M. 2018. Communicative eye contact signals a commitment to cooperate for young children. *Cognition* 179:192–201.
- [Sugden 2000] Sugden, R. 2000. The motivating power of expectations. In Nida-Rümelin, J., and Spohn, W., eds.,

Rationality, Rules, and Structure, 103–129. Dordrecht:
Springer Netherlands.

[Wallace 1998] Wallace, R. J. 1998. *Responsibility and the
moral sentiments*. Harvard Univ. Press, 2. print edition.