

The Role of Intention in Cognitive Robotics

D. Vernon, S. Thill, and T. Ziemke

Abstract We argue that the development of robots that can interact effectively with people requires a special focus on building systems that can perceive and comprehend intentions in other agents. Such a capability is a prerequisite for all pro-social behaviour and in particular underpins the ability to engage in instrumental helping and mutual collaboration. We explore the prospective and intentional nature of action, highlighting the importance of joint action, shared goals, shared intentions, and joint attention in facilitating social interaction between two or more cognitive agents. We discuss the link between reading intentions and theory of mind, noting the role played by internal simulation, especially when inferring higher-level action-focussed intentions. Finally, we highlight that pro-social behaviour in humans is the result of a developmental process and we note the implications of this for the challenge of creating cognitive robots that can read intentions.

1 Introduction

There are many reasons why one would like a robot to exhibit a capacity for cognition. These include the ability to deal with uncertain or poorly specified situations and the ability to act prospectively, anticipating the need for actions and predicting the outcome of those actions [41, 42]. However, perhaps one of the most compelling motivations for research in cognitive robotics is the need for robots to interact naturally and safely with people. People are cognitive agents and consequently when

D. Vernon
School of Informatics, University of Skövde, Sweden, e-mail: david.vernon@his.se

S. Thill
School of Informatics, University of Skövde, Sweden, e-mail: serge.thill@his.se

T. Ziemke
School of Informatics, University of Skövde, Sweden, and Department of Computer and Information Science, Linköping University, Sweden, e-mail: tom.ziemke@his.se

they interact socially with other agents, they exhibit the key attributes of cognition: prospection and adaptive real-time goal-directed behaviour. To interact effectively, a cognitive robot needs to be able to interact on the same basis. When that interaction stretches to helping or collaborating with people, a cognitive robot needs to be able to engage in perspective-taking, i.e. to form a theory of mind [25], to see the world from the other person's perspective.

In this chapter, we first explore the four attributes of cognition involved in interaction with an inanimate world: attention, action, goals, and intentions, focusing in particular on the pivotal role of intention. We then expand the discussion to include interaction with other cognitive agents and discuss how these four attributes are extended in social interaction in the form of joint attention, joint action, shared goals, and shared intentions. In doing this, we explain that shared intention involves more than just the superposition of the intentions of two or more individual agents and highlight the essential role it plays in facilitating safe, engaging, and effective interaction, particularly where two or more agents are helping each other in fulfilling some task. By extension, we argue that cognitive robots as much as people need to have a capacity for reading and sharing intentions when interacting with people if they are to do so effectively [3].

2 The Prospective and Intentional Nature of Action

The movements of cognitive agents are organized: they are defined by goals and guided by prospection. These goal-directed prospectively-controlled movements are called actions [42, 19]. Typically, cognitive agents do not deliberately pre-select the exact movements required to achieve a desired goal. Instead, they select prospectively-guided intention-directed goal-focussed action, with the specific movements being adaptively controlled as the action is executed.

While action and goals are two of the essential characteristics of cognitive interaction, there are two more: intention and attention. The first of these — intention — captures the prospective nature of action and goals. The distinction between intentions and goals is not always clearly made. An intention can be viewed as a plan of action an agent chooses and commits itself to in pursuit of a goal. An intention therefore includes both a goal and the means of achieving it [8, 40]. Thus, an agent may have a goal for some state of affairs to exist and an intention to do something specific in pursuit of that state of affairs. Intentions integrate, in a prospective framework, actions and goals. Finally, there is perception, the essential sensory aspect of cognition. However, in the context of cognitive action, perception is directed. It is focussed on goals and influenced by expectations. In other words, it is attentive. Arguably, one can describe attention in the context of interaction as an intention-guided perception [41].

All of the components of cognitive interaction — action, goals, intention, and attention — have an element of prospection. Our aim in this chapter is to explain what is necessary to transform this characterization to one that is representative of

social interaction between two (or more) cognitive agents. This will involve the notions of *joint action*, *shared goals*, *shared intentions*, and *joint attention*. As we will see, this transformation, and these four notions, go beyond a simple superposition of individual action, goals, intention, and attention from which they derive. Much more is involved in social cognition and the interaction of two or more agents. To set the scene for this, we begin with a brief overview of social cognition and social interaction.

3 Social Cognition and Social Interaction

Social cognition — necessary for effective social interaction with other cognitive agents — embraces a wide range of topics. The abilities required for successful social interaction include reading faces, detecting eye gaze, recognizing emotional expressions, perceiving biological motion, paying joint attention, detecting goal-directed actions, discerning agency, imitation, deception, and empathy, among many others [15].

3.1 *The Basis of Social Cognition*

Social cognition depends on an agent's ability to interpret a variety of sensory data that conveys information about the activities and intentions of other agents. Newborns have an innate sensitivity to biological motion [37] and it has been shown that the ability to process biological motion is a hallmark of social cognition, providing a cognitive agent with a capacity for adaptive social behaviour and nonverbal communication, to the extent that individuals who exhibit a deficit in visual processing of biological motion are also compromised in social perception [33]. The clearest example of this is the ability to read body language, the subtle body movements, gestures, and actions that are an essential aspect of successful interaction between cognitive agents.

For an agent to interact socially with another cognitive agent, it must be (and stay) attuned to the cognitive state of that agent and be sensitive to changes. There is a strong link between the state of an agent's body and its cognitive and affective state, especially during social interaction [1, 22]. There are four aspects to this link:

1. When an agent perceives a social stimulus, this perception produces bodily states in the perceiving agent.
2. The perception of bodily states in other agents frequently evokes a tendency to mimic those states.
3. The agent's own body states trigger affective states in the agent.
4. The efficiency of an agent's physical and cognitive performance is strongly affected by the compatibility between its bodily states and its cognitive states.

Because of the link between bodily states and cognitive and affective states, the posture, movements, and actions of an agent convey a great deal about the agent's cognitive and affective disposition as well as influencing how another agent behaves towards it.

3.2 *Helping and Collaboration*

While social cognition is ultimately about mutual interaction, this interaction can be asymmetric or symmetric: one agent can assist another, or both agents can assist each other. In the following, we will refer to these behaviours as helping (sometimes adding the qualification *instrumental* helping) and collaboration, respectively. Significantly, the development of a capacity for collaborative interaction depends on the prior development of a capacity for instrumental helping and it takes several years for human infants to develop the requisite abilities [28].

During the first year of life the progressive acquisition of motor skills determines the development of the ability to understand the intentions of other agents, from anticipating the goal of simple movements to the understanding of more complex goals [14, 16]. At the same time, the ability to infer what another agent is focussing their attention on and the ability to interpret emotional expressions begins to improve substantially [10, 17].

At around 14 to 18 months of age children begin to exhibit instrumental helping behaviour, i.e. they display spontaneous, unrewarded helping behaviours when another person is unable to achieve his goal. Young children are naturally altruistic and have an innate propensity to help others instrumentally, even when no reward is offered [43]. This is a critical stage in the development of a capacity for collaborative behaviour, a process that progresses past three and four years of age.

Instrumental helping has two components: a cognitive one and an emotional one. The cognitive component is concerned with recognizing what the other agent's goal is: what they are trying to do. The motivational component is what drives the helping agent to act in the first place. This could be the desire to see the second agent achieve the goal or, alternatively, the desire to see the second agent exhibit pleasure at achieving the goal.

The ability to engage in instrumental helping develops with age: 14-month-old infants can help others in situations where the task is relatively simple, e.g. helping with out-of-reach objects, whereas 18-month-old infants engage in instrumental helping in situations where the cognitive task is more complicated [43]. As already mentioned, rewards are not necessary and the availability of rewards does not increase the incidence of helping. Indeed, rewards can sometimes undermine the motivation to help. Infants are willing to help several times and will even continue to help even if the cost of helping is increased.

The second form of helping — collaboration — is more complicated and focuses on mutual helping where two agents work together to achieve a common goal. It requires the two agents to share their intentions, to agree on the goal, share atten-

tion, and engage in joint action. Collaboration requires complex interaction over and above the ability to engage in instrumental helping. It involves the establishment of shared goals and shared intentions and it requires subtle adjustment of actions when the two agents are in physical contact such as when handing items to each other or carrying objects together. Michael Tomasello and Malinda Carpenter argue that *shared intentionality*, i.e. a collection of social-cognitive and social-motivational skills that allow two or more participants engaged in collaborative activity to share psychological states with one another, plays a crucial role in the development of human infants. In particular, it allows them to transform an ability to follow another agent's gaze into an ability to jointly pay attention to something, to transform social manipulation into cooperative communication, group activity into collaboration, and social learning into instructed learning [39].

3.3 The Central Role of Intention in Mutual Interaction

The situation becomes complicated when one progresses from instrumental helping to collaboration. In the latter case, we are dealing with *joint cooperative action*, or *joint action* for short, sometimes referred to as *shared cooperative activity*. Agents that engage in joint action share the same goal, intend to act together, and coordinate their actions to achieve their shared goal through joint attention. That sounds fairly straightforward but as we unwrap each of these issues — joint action, shared intentions, shared goals, and joint attention — interdependencies between them arise. For example, joint action requires a shared intention, a shared goal, and joint attention when executing the joint action; shared intention includes shared goals; and joint attention is effectively perception that is guided by shared intention and is goal-directed (see Fig. 1).

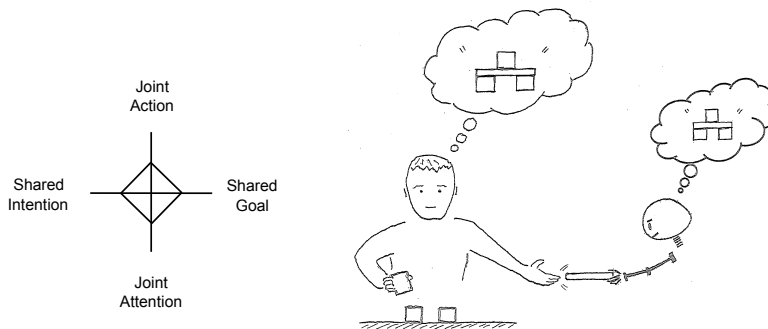


Fig. 1 Collaboration involves joint action, shared intention, a shared goal, and joint attention, each of which are mutually dependent. Here the human and the robot are engaged in a joint action and have a shared intention (and hence a shared goal and commitment to achieving it) and exhibit joint attention where their goal-directed perceptions are guided by their shared intention.

For one agent to be able to help another agent, it must first infer or read the other agent's intentions. This in itself is a complex problem. It can be addressed in two phases: reading low-level intentions associated with movements (e.g. predicting what someone is reaching for) and reading high-level intentions associated with actions (e.g. predicting why someone is reaching for that object or what he or she want to do with it). Elisabeth Pacherie argues that three different levels of intentions can be distinguished: (1) distal intentions, where the goal-directed action to be performed is specified in cognitive terms with reference to the agent's environment; (2) proximal intentions, where the action is specified in terms of bodily action and the associated perceptual consequence; and (3) motor intentions, where the action is specified in terms of the motor commands and the impact on the agent's sensors [30].

To take part in collaborative activities requires an ability to read intentions and infer goals (as was the case in instrumental helping) but it also requires a unique motivation to share psychological states with other agents. By shared intentionality we here mean, following Tomasello *et al.*, "collaborative actions in which participants have a shared goal (shared commitment) and coordinated action roles for pursuing that shared goal" [40]. What is significant is that the goals and intentions of each agent involved in the collaboration must include something of the goals and intentions of the other agent and something of its own goals and intentions. In other words, the intention is a joint intention and the associated actions are joint actions. This differentiates collaboration from instrumental helping and, as we have said, makes it more complicated. Furthermore, each agent understands both roles of the interaction and so can help the other agent if required. Critically, agents not only choose their own action plan, but also represent (or 'mirror') the other agent's action plan in its own motor system to enable coordination in the sense of who is doing what and when.

Assuming collaboration on a shared goal, let us now look more closely at the issues of joint action, shared intention, and joint attention to better understand the role of intention in the interaction between cognitive agents, be they human or robot.

3.4 Joint Action

There are at least six degrees of freedom in joint action [31]. These include the number of participants involved, the nature of the relationship between the participants (e.g. peer-to-peer or hierarchical), whether or not the roles are interchangeable, whether the interaction is physical or virtual, whether or not the participants' association is temporary or more long-lasting, and whether or not the interaction is regulated by organizational or cultural norms. In the following, we will assume physical joint action between two peers that temporarily collaborate on a shared goal.

According to Michael Bratman, joint action, or shared cooperative activity, has three essential characteristics [7]:

1. Mutual responsiveness;
2. Commitment to joint activity;
3. Commitment to mutual support.

Let's assume there are two agents engaged in a shared cooperative activity. Each agent must be mutually responsive to the intentions and actions of the other and each must know that the other is trying to be similarly responsive. Consequently, each agent behaves in a way that is guided partially by the behaviour of the other agent. This is different from instrumental helping where the helping agent is responsive to the intentions of the agent that needs help but not the other way round.

Each agent must also be committed to the activity in which they are engaged. This means that both agents have the same intention but they need not have the same reason for engaging in the activity. This is a subtle point: it means that the outcome of the collaboration is the same for both agents but the reason for adopting the goal of achieving that outcome need not be the same. If a cognitive robot and a disabled person collaborate to do the laundry, the outcome — the goal — may be a wardrobe full of clean clothes but the reason the person has the goal is to have a fresh shirt to wear in the morning whereas the reason the robot has the goal may just be to keep the house clean and uncluttered. If they collaborate to cook a stew, the goal may be nutritious meal, but the person's reason for the goal is to stay healthy whereas the reason the robot adopts the goal may simply be to use up some vegetables that would otherwise have to be thrown out.

Finally, each agent must be committed to supporting the efforts of the other to play their role in the joint activity. This characteristic complements the mutual responsiveness by requiring that each agent will in fact provide any help the other agent requires. It says that each agent treats this collaborative mutual support as a priority activity: even if there are other activities that are competing for the attention of each agent, they will still pay attention to the shared cooperative activity they are both engaged in.

Philip Cohen and Hector Levesque address similar issues in their theory of teamwork [11]. They do so in the context of designing artificial agents that can engage in joint action, setting out the conditions that need to be fulfilled for a group of agents to exhibit joint commitment and joint intention. This builds on their definitions of individual commitment and individual intention which are, roughly speaking, a persistent goal and an action plan in support of achieving that persistent goal.

It is important to note that Bratman's account of joint action has been subject to some criticism in that it appears to require sophisticated shared intentionality and an adult-level theory of mind. Yet, as we have seen, young children develop a capability for joint action. An alternative account that doesn't require sophisticated shared intentionality, but only requires shared goals and an understanding of goal-directed actions has been proposed; see [9, 32, 2]. That said, shared intentions are important for joint action and the intentions of each agent must interlock: each agent must intend that the shared activity be fulfilled in part by the other agent and that their individual activities — both planned actions and actual actions when being executed — mesh together in a mutually-supportive manner.

3.5 Shared Intentions

A shared intention — sometimes called *we-intention*, *collective intention*, or *joint intention* — is not simply a collection of individual intentions, even when those individual intentions are supplemented by beliefs or knowledge that both participating agents share [39]. There is more to it than this.

An agent with an individual intention represents the overall goal and the action plan by which it will achieve that goal and, furthermore, this plan is to be performed by the agent alone. That much is clear. However, agents with a shared intention (and engaged in a joint action) represent the overall shared goal between them but only their own partial sub-plans. Elisabeth Pacherie identifies three levels of shared intentions (shared distal intentions, shared proximal intentions, and coupled motor intentions) [31]; these are extensions of her characterization of individual intentions [30] (Section 3.3).

Shared intentionality appears to be unique to humans and its development seems to depend on a peculiarly-human motivation to share emotions, experience, and activities and a more general motivation to understand others as animate, goal-directed, and intentional [40]. An example of how artificial cognitive systems can exploit these ideas can be found in Peter Ford Dominey's and Felix Warneken's paper "The basis of shared intentions in human and robot cognition." Based on findings in computational neuroscience (e.g. the mirror neuron system) and developmental psychology, it describes how representations of shared intentions allow a robot to cooperate with a human [13].

Each individual agent with a shared intention does not need to know the other agent's partial plan. However, they do need to share the overall goal. When it comes to the realization of a shared intention and the execution of a joint action, the agent must also factor in the real-time coordination of their individual activities. In this case, each agent must also represent its own actions and their predicted consequences *and* the goals, intentions, actions and predicted consequences of the other agent [36]. Furthermore, each agent must represent the effect that their actions have on the other agent, it must have at least a partial representation of how component actions combine to achieve the overall goal, it must be able to predict the effects of their joint actions so that it can monitor progress towards the overall goal and adjust its actions to help the other agent if necessary. The additional requirements imposed by the execution of joint action correspond to Elisabeth Pacherie's shared proximal intentions [31].

It is apparent that, in carrying out a joint intention and executing a joint action, both agents must establish a shared perceptual framework. This is where joint attention (in the sense of perception guided by shared intention) comes in.

3.6 *Joint Attention*

Social interaction, in general, and collaborative behaviours, in particular, depend on the participating agents to establish *joint attention* [21]. Joint attention involves much more than two agents looking at the same thing. As Michael Tomasello and Malinda Carpenter note, joint attention “is not just two people experiencing the same thing at the same time, but rather it is two people experiencing the same thing at the same time and *knowing together that they are doing this*” [39]. The essence of joint attention lies in the relationship between intentionality and attention. This provides the basis for a definition of joint attention as “(1) a coordinated and collaborative coupling between intentional agents where (2) the goal of each agent is to attend to the same aspect of the environment.” [21]. Joint attention, then, requires shared intentionality. Furthermore, the participating agents must be engaged in collaborative intentional action. During this collaboration, each agent must monitor, understand, and direct the attentional behaviour of the other agent, and significantly, both agents must be aware that this is going on.

Joint attention is an on-going mutual activity that is carried on throughout the collaborative process to monitor and direct the attention of the other agent. In a sense, joint attention is, itself, a joint activity.

At least four skills need to be recruited by a cognitive agent to achieve joint attention [21]. First, the agent must be able to detect and track the attentional behaviour of the other agent (we are assuming that there are just two agents involved in joint attention here but of course there could be more). Second, the agent must be able to influence the attentional behaviour of the other agent, possibly by using gestures such as pointing or by use of appropriate words. Third, the agent must be able to engage in social coordination to manage the interaction, using techniques such as taking turns or swapping roles, for example. Finally, the agent must be aware that the other agent has intentions (which, as we noted, could be different provided the goal is the same). That is, the agent must be capable of intentional understanding: it must be able to interpret and predict the behaviour of the other agent in terms of the actions required to reach the shared goal.

4 Reading Intentions and Theory of Mind

The ability to infer intentions is closely linked to what is known as *theory of mind* [23]: the capacity by which one agent is able to take a perspective on someone else’s situation. Theory of mind is defined by Andrew Meltzoff as “the understanding of others as psychological beings having mental states such as beliefs, desires, emotions, and intentions” [23]. To have a theory of mind means to have the ability to infer what someone else is thinking and wants to do. The ability to imitate — the capacity to learn new behaviours by observing the actions of others [4, 24] — forms the basis for the development of a person’s ability to form a theory of mind [25]. It is a key mechanism in cognitive development and it is innate in humans [26, 27].

The link between imitation and theory of mind is the ability of an agent to infer the intentions of another agent. When imitating adults, infants as young as 18 months of age can not only replicate the actions of the adult (and remember: actions are focussed on goals, not just bodily movements) when successfully performing a task but they can also persist in trying to achieve the goal of the action even when the adult is unsuccessful in performing the task. In other words, the infant can read the intention of the adult and infer the unseen goal implied by the unsuccessful attempts. Andrew Meltzoff and Jean Decety summarize the link between imitation and theory of mind (which they also refer to as *mentalizing*) as follows: “Evidently, young toddlers can understand our goals even if we fail to fulfil them. They choose to imitate *what we meant to do*, rather than what we mistakenly did do.” [25], p. 496 (emphasis added). They also remark that “In ontogeny, infant imitation is the seed and the adult theory of mind is the fruit.”

Young children normally differentiate between the behaviour of inanimate and animate objects, attributing mental states to the animate objects. In fact, such is the importance of biological motion to social cognition that if an inanimate object, even two-dimensional shapes such as triangles, exhibit movements that are animate or biological — self-propelled, non-linear paths with sudden changes in velocity — humans cannot resist attributing intentions, emotions, and even personality traits to that inanimate object [18]. In the same way, humans also infer different types of intention depending on whether they are interpreting movements (lower level intentions) or actions (higher level). Whereas movement intention refers to *what* physical state is intended by a certain action, e.g., inferring the end location of a specific observed movement — if the hand moves into the direction of a cup, it is likely that the agent intends to grasp that cup — a higher conceptual level intention refers to *why* that specific action is being executed and the motives underlying the action, e.g., the agent might be thirsty and want a drink. This mirrors the distinction we drew at the beginning between the concrete movements comprising an action and the higher-order conceptual goals of an action.

So, how do humans infer the intentions of others from their actions? Internal simulation is a possible mechanism [5].¹ The key idea is that the ability to infer the intentions of another agent from observations of their actions might actually be based on the same mechanism that predicts the consequences of the agent’s own actions based on its own intentions. Cognitive systems make these predictions by internal simulation using forward models that take either overt or covert motor commands as input and produce as output the likely sensory consequences of carrying out those commands. When a cognitive system observes another agent’s actions, the same mechanism can operate provided that the internal simulation mechanism is able to associate observed movements (and not just self-generated motor commands) and likely, i.e. intended, sensory consequences. This is what the ideo-motor principle suggests [38, 20, 29] and what the mirror-neuron system provides [35, 34]. By exploiting internal simulation, when an agent just sees another agent’s action, not only

¹ For an in-depth discussion of a computational approach to intention recognition, see “Towards computational models of intention detection and intention prediction” by Elisheva Bonchek-Dokow and Gal Kaminka [6].

are the actions activated in it but so too are the consequences of those actions, and hence the intention of the actions can be inferred. With a suitably-sophisticated joint representation and internal simulation mechanism, both low-level movement intentions and high-level action intentions can be accommodated.

Predicting and recognizing intentions in situations where there are groups of agents is particularly challenging because the cognitive system has to do more than track and predict the actions of individual agents, it also has to infer the joint intention of the entire group and this may not simply be “the sum of the intentions of the individual agent” [12]. It is also necessary to recognize the position of each agent in the social structure of the group. Again, this is a difficult challenge because an agent may play more than one role in a group.

5 Conclusions

Our goal in this chapter has been to highlight the pivotal role played by intentions in social interaction and, in particular, to argue that a capacity to infer the intentions of the cognitive agent with which one is interacting is essential if that interaction is to be effective. This is true both in the asymmetric case of instrumental helping where one agent assists another to achieve its goals without implicit instruction to do so and also in the symmetric case where both agents are collaborating. In this latter case, the ability to read intentions, by taking a perspective on the other agent’s view of the interaction and forming a theory of mind for that agent, is doubly important because the other components of successful collaborative interaction — joint action and joint attention — also depends on the other agent’s intentions. The upshot of this is that if we seek to construct cognitive robots that can interact effectively — asymmetrically or symmetrically : helping or collaborating — then it is imperative that these cognitive robots have a capacity for perspective taking and forming a theory of mind. However, it is not simply a question of implanting such a capacity in a cognitive robot. We know from psychology that this capacity is the result of an extended period of cognitive and social development where the emergence of the ability to engage in instrumental helping precedes that of collaborative interaction. Consequently, the challenge in cognitive robotics is to model the developmental process by which these capacities emerge over time as the robot engages with people.

Acknowledgements This work was supported by the European Commission, Project 611391 DREAM: Development of Robot-enhanced Therapy for Children with Autism Spectrum Disorders (www.dream2020.eu). Also supported by the Knowledge Foundation, Stockholm, under SIDUS grant AIR (Action and intention recognition in human interaction with autonomous systems). The authors would like to acknowledge the work of Harold Bekkering, Radboud University Nijmegen, in our characterization of aspects of neonatal development.

References

- [1] Barsalou LW, Niedenthal PM, Barbey A, Ruppert J (2003) Social embodiment. In: Ross B (ed) *The Psychology of Learning and Motivation*, vol 43, Academic Press, San Diego, pp 43–92
- [2] Bekkering H, de Bruijn E, Cuijpers R, Newman-Norlund, van Schie H, Meulenbroek R (2009) Joint action: Neurocognitive mechanisms supporting human interaction. *Topics in Cognitive Science* 1(2):340–352
- [3] Bicho E, Erhagen W, Sousa E, Louro L, Hipolito N, Silva EC, Silva R, Ferreira F, Machado T, Hulstijn M, Maas Y, de Bruijn ERA, Cuijpers RH, Newman-Norlund RD, van Schie HT, Meulenbroek RGJ, Bekkering H (2012) The power of prediction: Robots that read intentions. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vilamoura, Algarve, Portugal, pp 5458–5459
- [4] Billard A (2002) Imitation. In: Arbib MA (ed) *The Handbook of Brain Theory and Neural Networks*, MIT Press, Cambridge, MA, pp 566–569
- [5] Blakemore S, Decety J (2001) From the perception of action to the understanding of intention. *Nature Reviews Neuroscience* 2(1):561–567
- [6] Bonchek-Dokow E, Kaminka GA (2014) Towards computational models of intention detection and intention prediction. *Cognitive Systems Research* 28:44–79
- [7] Bratman ME (1992) Shared cooperative activity. *The Philosophical Review* 101(2):327–341
- [8] Bratman ME (1998) Intention and personal policies. In: Tomberlin JE (ed) *Philosophical Perspectives*, vol 3, Blackwell
- [9] Butterfill S (2012) Joint action and development. *The Philosophical Quarterly* 62(246):23–47
- [10] Butterworth G, Jarrett N (1991) What minds have in common is space: Spatial mechanisms serving joint visual attention in infancy. *British Journal of Developmental Psychology* 9:55–72
- [11] Cohen PR, Levesque HJ (1991) Teamwork. *Nous* 25(4):487–512
- [12] Demiris Y (2007) Prediction of intent in robotics and multi-agent systems. *Cognitive Processing* 8:152–158
- [13] Dominey PF, Warneken F (2011) The basis of shared intentions in human and robot cognition. *New Ideas In Psychology* 29:260–274
- [14] Falck-Ytter T, Gredebäck G, von Hofsten C (2006) Infants predict other people’s action goals. *Nat Neurosci* 9(7):878–879
- [15] Frith U, Blakemore S (2006) Social cognition. In: Kenward M (ed) *Foresight Cognitive Systems Project, Foresight Directorate, Office of Science and Technology*, 1 Victoria Street, London, SW1H 0ET, United Kingdom
- [16] Gredebäck G, Kochukhova O (2010) Goal anticipation during action observation is influenced by synonymous action capabilities, a puzzling developmental study. *Exp Brain Res* 202(2):493–497

- [17] Harris PL (2008) Children's understanding of emotion. In: Lewis M, Haviland-Jones JM, Barrett LF (eds) *Handbook of emotions*, The Guilford Press, New York, pp 320–331
- [18] Heider F, Simmel M (1944) An experimental study of apparent behaviour. *American Journal of Psychology* 57:243–249
- [19] von Hofsten C (2004) An action perspective on motor development. *Trends in Cognitive Sciences* 8:266–272
- [20] Iacoboni M (2009) Imitation, empathy, and mirror neurons. *Annual Review of Psychology* 60:653–670
- [21] Kaplan F, Hafner V (2006) The challenges of joint attention. *Interaction Studies* 7(2):135–169
- [22] Lindblom J (2015) *Embodied Social Cognition*, Cognitive Systems Monographs (COSMOS), vol 26. Springer, Berlin
- [23] Meltzoff AN (1995) Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology* 31:838–850
- [24] Meltzoff AN (2002) The elements of a developmental theory of imitation. In: Meltzoff AN, Prinz W (eds) *The Imitative Mind: Development, Evolution, and Brain Bases*, Cambridge University Press, Cambridge, pp 19–41
- [25] Meltzoff AN, Decety J (2003) What imitation tells us about social cognition: a rapprochement between developmental psychology and cognitive neuroscience. *Philosophical Transactions of the Royal Society of London: Series B* 358:491–500
- [26] Meltzoff AN, Moore MK (1977) Imitation of facial and manual gestures by human neonates. *Science* 198:75–78
- [27] Meltzoff AN, Moore MK (1997) Explaining facial imitation: A theoretical model. *Early Development and Parenting* 6:179–192
- [28] Meyer M, Bekkering H, Paulus M, Hunnius S (2010) Joint action coordination in 2-and-a-half- and 3-year-old children. *Frontiers in Human Neuroscience* 4(220):1–7
- [29] Ondobaka S, Bekkering H (2012) Hierarchy of idea-guided action and perception-guided movement. *Frontiers in Cognition* 3:1–5
- [30] Pacherie E (2008) The phenomenology of action: A conceptual framework. *Cognition* 107:179–217
- [31] Pacherie E (2012) The phenomenology of joint action: Self-agency vs. joint-agency. In: Seemann A (ed) *Joint Attention: New Developments*, MIT Press, Cambridge MA, pp 343–389
- [32] Pacherie E (2013) Intentional joint agency: shared intention lite. *Synthese* 190(10):1817–1839
- [33] Pavlova MA (2012) Biological motion processing as a hallmark of social cognition. *Cerebral Cortex* 22(981–995)
- [34] Rizzolatti G, Craighero L (2004) The mirror neuron system. *Annual Review of Physiology* 27:169–192
- [35] Rizzolatti G, Fadiga L (1998) Grasping objects and grasping action meanings: the dual role of monkey rostroventral premotor cortex (area F5). In: Bock GR,

- Goode JA (eds) *Sensory Guidance of Movement*, Novartis Foundation Symposium 218, John Wiley and Sons, Chichester, pp 81–103
- [36] Sebanz N, Knoblich G (2009) Prediction in joint action: What, when, and where. *Topics in Cognitive Science* 1(353–367)
- [37] Simion F, Regolin L, Bulf H (2008) A predisposition for biological motion in the newborn baby. *Proceeding of the National Academy of Sciences (PNAS)* 105(2):809–813
- [38] Stock A, Stock C (2004) A short history of ideo-motor action. *Psychological research* 68(2–3):176–188
- [39] Tomasello M, Carpenter M (2007) Shared intentionality. *Developmental Science* 10(1):121–125
- [40] Tomasello M, Carpenter M, Call J, Behne T, Moll H (2005) Understanding and sharing intentions: the origins of cultural cognition. *Behavioral and Brain Sciences* 28(5):675–735
- [41] Vernon D (2014) *Artificial Cognitive Systems — A Primer*. MIT Press, Cambridge, MA
- [42] Vernon D, von Hofsten C, Fadiga L (2010) *A Roadmap for Cognitive Development in Humanoid Robots*, *Cognitive Systems Monographs (COSMOS)*, vol 11. Springer, Berlin
- [43] Warneken F, Tomasello M (2009) The roots of human altruism. *British Journal of Psychology* 100(3):455–471