

# A Multi-Category Theory of Intention

Henny Admoni (henny@cs.yale.edu) and Brian Scassellati (scaz@cs.yale.edu)

Department of Computer Science, 51 Prospect Street  
New Haven, CT 06511 USA

## Abstract

People excel at attributing intentionality to other agents, whether in simple scenarios such as shapes moving in two dimensions or complex scenarios such as people interacting. We note that intentionality attributions seem to fall into two categories: low-level intentionality in which an observer has a theory of mind about an agent, and high-level intentionality in which an observer believes the agent has a theory of mind about something else. We introduce the terms *L-intentionality* and *H-intentionality* to refer to these attributions, respectively, and describe this division by using examples from previous research. Social robots provide a particularly good platform for evaluating the presence of different types of intentionality, and we discuss how robots can help distinguish the relationship between H- and L-intentionality, based on a number of possible models that we enumerate. We conclude by highlighting some interesting questions about intentionality in general and the interplay between H- and L-intentionality in particular.

**Keywords:** intention; animacy; computer model; human-robot interaction; robotics

## Introduction

Much research in psychology has focused on people's ability—and eagerness—to attribute intentions and animacy to simple shapes based on motion. From Michotte's (1963) and Heider and Simmel's (1944) experiments with animacy and intention to recent work decomposing intentional actions such as chasing (Gao, Newman, & Scholl, 2009), psychologists have found that intention attributions to moving shapes appear to be immediate and irresistible. Animacy is often observed in a display of simple shapes when the motion in the display cannot be explained as ordinary inanimate motion, for instance when speed and direction change without direct contact with other objects (Tremoulet & Feldman, 2000).

At the same time, evidence shows that people attribute intentions based on high-level behavioral evaluations, as well. For instance, 18-month-old toddlers can recognize and imitate intentional actions performed by adults, even if those actions are unsuccessful (Meltzoff, 1995). By pre-school age, children begin to represent others' beliefs, even when those beliefs are mistaken, in order to correctly predict a person's intentional action (Wellman, Cross, & Watson, 2001). As adults, neurological evidence indicates that a certain region of the brain is sensitive to whether peoples' motions are consistent or inconsistent with their purported intentions (Pelphrey, Morris, & McCarthy, 2004).

While abundant evidence demonstrates peoples' attributions of intentionality, the types of attributions they make seem to differ. Cues that prompt intention attributions come in two categories: low-level, perceptual cues, such as motion, and high-level cues that must be reasoned about, such as facial expression. To distinguish between intentions cued

in these different ways, we introduce two novel terms, referring to intention attributions made from low-level cues as *L-intentionality* and to attributions made from high-level cues as *H-intentionality*. To date, little work has explored such categorical differences of intentionality. In the Types of Intentionality section, we use examples from previously published research to define our hypothesis that L-intentionality and H-intentionality are separate kinds of intention attributions.

Robotics has provided a valuable experimental platform to test perceptions of intentionality. Because robots are extremely flexible (in terms of appearance, motions, sounds, and so on), researchers can manipulate specific variables of a human-robot interaction to test specific features of intentionality attributions. In the Social Robots as Experimental Platforms section, we describe past work with robots and other computational models of intentionality, and we discuss the benefits social robotics can offer intentionality research.

The next section, Models of Intentionality, enumerates possible models for the relationship between H-intentionality and L-intentionality based on the hypothesis that these are distinct observations. We describe what each model implies about real-world intentionality attributions, and we note how each model can be tested to confirm or deny our hypothesis.

We conclude this paper by discussing some likely starting points for research on the different categories of intention, and describing some interesting questions about intentionality that have yet to be addressed.

## Types of Intentionality

In this paper, we define an intentional action as a goal-directed action that is performed deliberately. Intentionality is the capacity to express or perform intentional actions. A theory of mind for other agents enables us to attribute intentionality to those agents (Leslie, 1987; Baron-Cohen, 1995), an ability that develops early in life (Meltzoff, 1995). Note that for our purposes, animacy and intentionality are strongly correlated, in that it is impossible to attribute animacy without the presence of intentional, goal-directed behavior (Tremoulet & Feldman, 2006).

In this section, we distinguish L-intentionality and H-intentionality as distinct but related categories. We can define each category by how an observer perceives and recognizes intentionality. To put the categorical difference simply, L-intentionality in an agent involves an observer having a theory of mind for that agent; H-intentionality involves an observer believing that the agent has a theory of mind for something else (Figure 1).

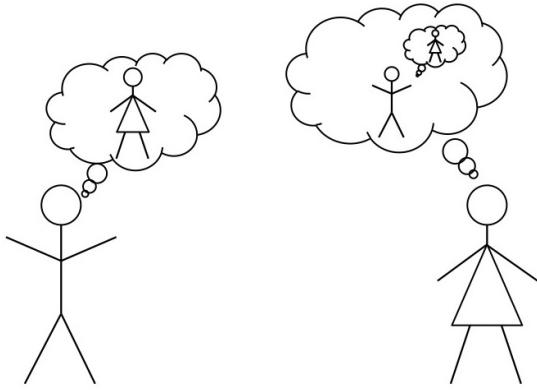


Figure 1: Low-level intentionality (left) is attributed when an agent has a theory of mind for another agent. High-level intentionality (right) is attributed when an agent believes another agent has a theory of mind for something else.

### L-Intentionality

To illustrate the different types of intentionality, picture a preschool boy named Billy. Billy is good at putting away his toys: he readily brings them to the toy chest whenever it is time to clean up. This kind of goal-directed behavior—carrying toys directly to the toy chest and depositing them— involves a series of coordinated actions and knowledge, but it can occur independently of theories of mind about others in the environment. Observing this, we might attribute L-intentionality to Billy. (Assume for the sake of the example that we cannot infer that Billy is H-intentional based solely on what we know about human beings.)

Actions that cue L-intentionality reveal goal-directed, deliberate behavior. L-intentionality is often elicited from low-level perceptions, such as those arising from visual displays of moving shapes. The perception of intentionality and animacy from simple moving shapes has been well-established in psychology literature (see (Scholl & Tremoulet, 2000) for a review); even when people do not know the type of intentional action they are looking for, they show high validity with the ground truth and high reliability with each other when evaluating the motion of animated shapes (Pantelis et al., 2011). Most often, these shapes exhibit basic actions such as chasing, fighting, or foraging.

For example, in Gao et al. (2009), an animated “wolf” chases after an animated “sheep” by moving toward the sheep within some degree of direct heat-seeking behavior. When the degree of chase is sufficiently small, participants identify chasing consistently. When the wolf deviates more strongly from direct heat seeking, however, the perception of intentionality disappears. In these L-intentionality experiments, goal-directed motion leads to an attribution of intentionality.

### H-Intentionality

Now imagine that we observe Billy hiding his shoes inside the toy chest. He watches his mother place his shoes in a cov-

ered cubby. While his mother looks away, Billy retrieves the shoes, walks to the chest, looks to make sure his mother is not watching and puts his shoes in the closed chest. The action of watching his mother put the shoes in one place, and then covertly moving them to another, suggests that Billy has a theory of mind for his mother, understanding that she has her own (mistaken) beliefs; Billy is displaying H-intentionality.

While L-intentionality is based on a theory of mind for the agent in question, H-intentionality is based on that agent’s theory of mind for others. Therefore, H-intentionality is seen in more complex visual scenes than the simple moving shapes of L-intentionality; it is often cued by a combination of stimuli such as facial expression, vocal prosody, and physical motion. H-intentionality does not “pop out” in the same way as L-intentionality because attributions of H-intentionality require additional cognitive processing to account for the agent’s beliefs about its environment.

## Social Robots as Experimental Platforms

Experiments with human-robot interaction are a particularly rich source of intentionality attributions. In one experiment, a robot received greater attributions of animacy and intelligence when it cheated at a game of rock-paper-scissors than when it played the game correctly (Short, Hart, Vu, & Scassellati, 2010). In each round of the game, a human participant and the robot both selected an item (rock, paper or scissors) in secret, then simultaneously displayed their selection through hand signals. Each of the items loses to one other item and wins over one other item, so that one’s performance in the game depends on the opponent’s selection as well as one’s own. In conditions where the robot verbally cheated—declaring “I win!” when it had lost—participants tended to report that the robot was broken and less intelligent; in conditions where the robot physically cheated—changing its losing hand signal after viewing the participant’s selection—participants were significantly more likely to use active verbs when describing the robot. In other words, a small change in behavior (switching hand signals) led to a dramatic increase in intentionality attributions.

Social robots are a valuable platform for experiments involving intentionality. Robots are available in a huge variety of appearances (from anthropomorphic to simple shapes) and motion abilities (fully mobile to stationary; with a broad range of physical capabilities). Being programmable, robot behavior can be carefully manipulated to alter individual features (such as moving an arm at a particular speed) and to ignore subtle (and potentially subconscious) social cues from others, a feat that a human experimenter might find difficult. As machines, robots can perform exactly the same action again and again, but as social tools, robots appear socially neutral: while most participants in experiments from our laboratory have seen or heard about robots, most have no experience with actual robot capabilities, allowing robots to act a blank slate on which social characteristics (such as intentionality) can be drawn at will.

Animations or videos of people provide many of the same benefits as robots, but they lack an embodiment that may affect interactions. Research has shown that people follow commands more consistently from physically present robots than from virtual robots, even when the virtual robot looks and acts exactly like the embodied robot (Bainbridge, Hart, Kim, & Scassellati, 2011). This reflects the common wisdom of sales, which asserts that you should visit someone in person to close a deal. For intentionality research, in which subtle features may make a difference in intention attributions, having an embodied agent observed in real time allows for highly realistic experimental setup while maintaining strict control over experimental variables.

H-intentionality has been of particular interest to social robotics researchers, who are motivated to design human-robot interactions that are natural and communicative. One part of natural interaction is identifying and displaying a theory of mind for others through non-verbal intentional behavior, whether with gaze (Mutlu, Yamaoka, Kanda, Ishiguro, & Hagita, 2009), hand gestures (Nehaniv, Dautenhahn, Kubacki, Haegele, & Parlitz, 2005), or facial expression (Breazeal & Scassellati, 1999). Therefore, understanding intentionality is important for the design of robots that will interact with people, such as service or assistive robots.

## Models of Intentionality

In this section, we enumerate the possible relationships between H-intentionality and L-intentionality (Figure 2); in the following Discussion section, we explain which models we believe are most viable. To better describe the models, we will return to the previous example with Billy and his toys.

Researchers have attempted to computationally recognize intentions through Bayesian models (Baker, Tenenbaum, & Saxe, 2006; Schrempf, Albrecht, & Hanbeck, 2007), hidden Markov models (Aarno & Kragic, 2006), and algorithmic methods (Feldman & Tremoulet, 2008). As observed actions and possible intentions become more complex, specifying a reasonably-sized state space for intention-action mapping becomes increasingly challenging, which limits the power of current computational models of intention recognition. The models in this paper are intended to be abstract, high-level views of how intention attributions can be conceptually organized, not algorithmic specifications for functional programs.

The descriptions of the models in this section are based on *features of intentionality*—observations that can be empirically measured. For instance, goal-directed movement toward another agent in an approximately heat-seeking manner (as in Gao et al. (2009)) is one feature of chasing, an L-intentional action. Behavior based on anticipation of others’ responses, as when Billy looked to see whether his mother was watching him hide his shoes, is a feature of H-intentional actions. Many features for each type of intention have yet to be identified, and are part of the novelty of this area of research.

Of all of the models for the relationship between H- and L-intentionality, the simplest option is that there is no cor-

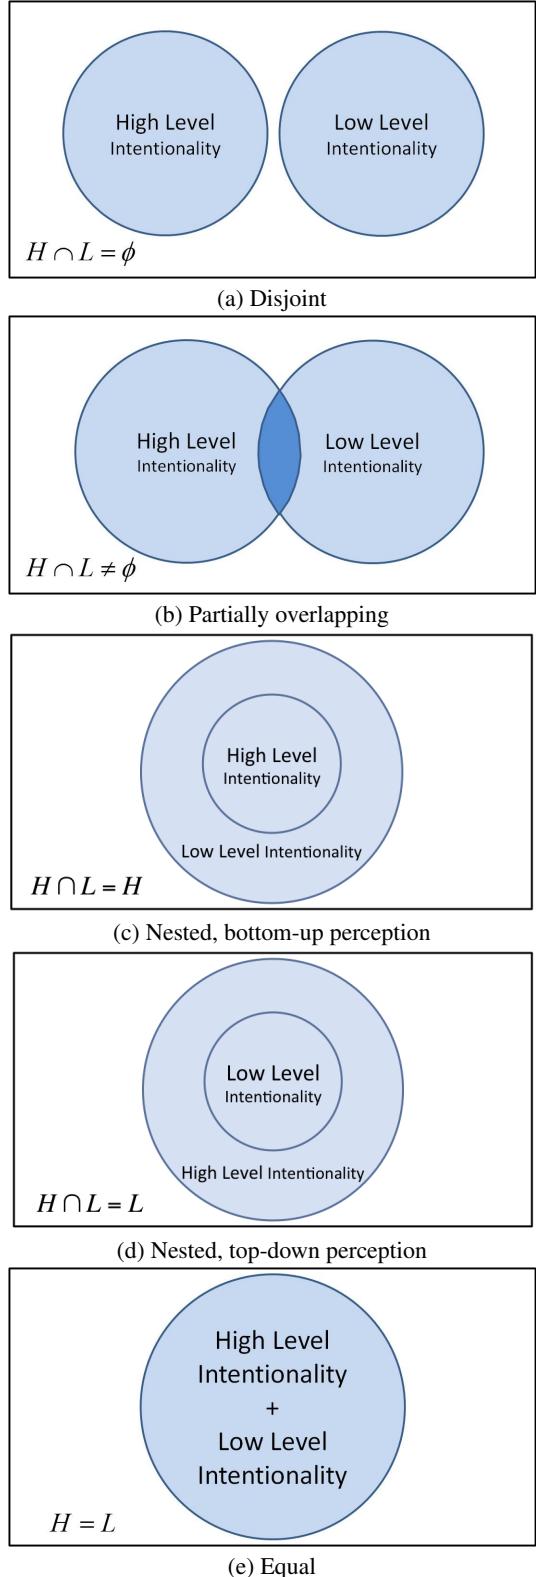


Figure 2: Potential models of the interaction between H- and L-intentionality. Each labeled circle represents the set of features that cue that type of intentionality. Set notation below each image mathematically describes the relationship between  $H$ , the set of features that cue H-intentionality, and  $L$ , the set of features that cue L-intentionality.

relation between the two (Figure 2a). In this model, the two types of perceptions share no features. Seeing Billy put his toys in the chest (and subsequently attributing L-intentionality to Billy) would not affect later judgements about H-intentionality, and vice versa. In the disjoint case, recognizing a feature of one type of intentionality would immediately identify the type of intentionality present, because that feature could not correspond to the other type of intentionality. To falsify this model, researchers must show that both H- and L-intentionality are cued by some feature.

A second possibility is that the two types of intentions share some features, but do not overlap completely (Figure 2b). In this case, seeing shared features would cue both types of intentionality attributions, though not every feature would be a shared feature. Identification of shared features would not confirm which type of intention is being perceived, but identification of disjoint features would allow experimenters to pinpoint the type of intention present. To prove that this model is correct, researchers would have to identify one feature that is unique to each type of intentionality, and one feature that is shared by both types.

Another possibility is that the types of intention are nested, so that one is wholly contained in the other. In the first form of nesting, which we'll call bottom-up perception, H-intentionality features are a strict subset of L-intentionality features (Figure 2c). In this case, L-intentionality is cued whenever H-intentionality is perceived, though the former can be present without the latter. In our example with Billy, merely perceiving H-intentionality—from watching Billy hide his shoes—would be enough to perceive L-intentionality, even without watching Billy put away his toys. This model can be falsified by identifying a scenario in which H-intentionality but not L-intentionality is perceived.

The complementary model posits that L-intentionality features are a strict subset of H-intentionality features (Figure 2d). We call this top-down perception, because the more complex H-intentionality can be cued without L-intentionality. In this model, H-intentionality is automatic whenever L-intentionality is perceived. This model can be falsified by identifying a feature of L-intentionality that does not cue H-intentionality.

A final possibility is that the feature sets of H-intentionality and L-intentionality are equal; that is, there is no difference between the features involved in cueing each type of intentionality. In this model, H-intentionality arises whenever L-intentionality is cued, and vice versa. To falsify this model, researchers would need to identify at least one feature from one type of intentionality that does not cue the other. Failing to falsify this model would challenge the hypothesis that H-intentionality and L-intentionality are separate processes.

## The Importance of Intentional Duality

Researchers have long established that people attribute intentionality to other people and to simple moving shapes under some conditions. The existence of H- and L-intentionality as

distinct types of intention attributions, if proven, might indicate a categorical distinction that runs more deeply in people's cognitive systems. H-intentionality and L-intentionality may not only be perceived differently; they may be understood and processed in different ways as well. After all, if the perceptual pathways for the two types of intentionality are different, perhaps the cognitive pathways that process them are also different. Perhaps there are distinct brain regions or neural pathways that process L-intentional and H-intentional stimuli. Perhaps, even, there is a different developmental time course for each type of intentionality, and infants can perceive one type of intentionality before the other. This difference might even extend to the perception of one agent as more complete or more animate than another, based on the type of cues that were used to establish its intentionality. All of these possibilities are consequences of our two-intention hypothesis, and will need to be further explored.

Understanding intentionality is also important for designing human-robot interactions. For robots that must interact naturally with people, being able to both recognize and perform intentional behavior is essential. To date, most computational approaches for intention recognition rely on statistical or probabilistic methods that do not scale well with increasing actions and intentions. Our model is a first step toward comprehensive understanding of intentionality that may lead to more complete and flexible computational models.

If cognitive differences underlie the different kinds of intentional attributions, these differences can be manipulated in interesting ways for human-robot interactions. Because robots are programmable, they can be made to display only cues from one type of intention, leading an observer to cognitively categorize them in a particular way. For instance, suppose researchers establish that unfamiliar L-intentional agents do not invoke as much shyness or fear as H-intentional agents, by virtue of their apparently less complicated mental structure. Robots that interact with children can then be manipulated to display only L-intentionality cues, reducing the likelihood that children will be afraid of them by controlling how they are perceived. The ability to craft human-robot interactions in a completely unprecedented way becomes possible if the categorical difference between intention types extends to cognitive processing levels.

## Discussion

Based on our distinction between H- and L-intentionality, some models are more likely than others. Clearly, a model in which the feature sets for both intentionality types are identical is impossible if we are to maintain the distinction between the two. We previously defined attributing H-intentionality to an agent as believing that the agent has a theory of mind for some other target. Inherent in this definition is the idea that we also believe the agent in question has its own goals and is capable of intentional actions to achieve those goals; in other words, that we have a theory of mind for that agent. Therefore, L-intentionality seems to be inherent in attribu-

tions of H-intentionality. For this reason, the disjoint model and the overlapping model seem unlikely candidates for our purposes. If L-intentionality is inherent in H-intentionality, then it would be impossible for the latter to be perceived without the former, which rules out the top-down model. In fact, we believe that the bottom-up model (Figure 2c), in which all features of H-intentionality are also features that cue L-intentionality, has the most promise as a model of intention attributions. In this model, L-intentionality can be present on its own—as supported by the many intention-from-motion studies described in the Introduction—but the presence of H-intentionality presupposes the presence of L-intentionality. For completeness, we have listed all possible models in this paper, but only the bottom-up model really serves the distinction we draw between H- and L-intentionality.

Though we distinguish between H- and L-intentionality, we have not made any claims about how these types of intentionality might be treated differently once they are perceived. Because our distinction is novel, we do not yet have evidence to ascertain whether or not H-intentionality affects peoples' reasoning differently from L-intentionality. The characteristics of intentionality described here might vary based on the type of intentionality that is perceived.

One characteristic that has yet to be explored is whether intentionality is revokable. Though many experiments identify the presence of intentionality, very few explore conditions under which intentionality disappears. It is possible that once intentionality is perceived, it remains for the duration of the exposure; on the other hand, perception of intentionality might also fade or be removed through some mechanism, such as time since the last intentional action. Anecdotal evidence from our lab suggests that attributions of H-intentionality are hard to remove, while attributions of L-intentionality may not be. In the rock-paper-scissors experiment described earlier (Short et al., 2010), once participants made attributions of H-intentionality to the robot, those attributions persisted despite the absence of additional H-intentional behavior. On the other hand, in the chasing studies (Gao et al., 2009), it seems easy to imagine that an agent that no longer chases would stop appearing L-intentional after some time.

Philosophers have identified yet another interesting characteristic of intentionality attributions: actions with harmful side-effects seem to be perceived as more intentional than actions with beneficial side-effects. For instance, if a company owner institutes a new manufacturing process that increases profits but damages the environment, the owner is seen as intentionally harming the environment. On the other hand, if the owner institutes a process that increases profits and helps the environment, the owner is not seen as intentionally helping the environment (Knobe, 2005). Does this type of disparity hold for both types of intentionality? Would a side-effect from an L-intentional action be seen as equally harmful to the same side effect from an H-intentional action? Or does the complexity of reasoning perceived in H-intentional agents endow them with more responsibility for side effects?

Along the same vein, observing intentional actions from an outsider's perspective may be different from identifying intentionality based on actions that personally affect one's self. Are attributions of intentionality different based on whether the effect of the action is personally relevant? Does it matter if the personally relevant effect is positive or negative? Are there differences between H-intentionality and L-intentionality in these cases?

Animals, like robots, present an interesting boundary condition for intentionality. Animals are clearly sentient and are generally attributed beliefs and goals. Whether animals can be said to be intentional, however, is open to debate (Heyes & Dickinson, 1990). It would be interesting to determine whether animals fall into the L-intentional or H-intentional category, and whether this separation of intention types might make it easier to attribute intentionality to animals.

Descriptions of models in the previous section refer to features that cue intentionality, but the precise nature of these features is still to be determined. These features need not be exclusively visual; they might include auditory features like prosody or kinesthetic features like heat. Some features, such as “approaches a target with a velocity between  $x$  and  $y$ ” may be quite specific; these are often used in rule-based intention detectors, but lack flexibility to account for novel or stochastic situations. Other features might be more general, such as “orients eyes toward the target.” Enumerating and sorting these features into H- and L-intentionality cues is a significant task, but it has the potential to dramatically increase understanding about intention perception.

Identifying and evaluating features of intentionality will require carefully designed experiments, and robots as intentional actors may be useful in teasing apart attributions of intentionality. Establishing features of intentionality will also help those who design interactions between people and objects, such as those working with social robots. By understanding features of intentionality, roboticists will be able to design robots that detect and exhibit relevant intentional behaviors, which would strengthen non-verbal communication in human-robot interactions.

## Conclusions

This paper presents a new representation and vocabulary for classifying different types of intentionality. Using examples from the extensive psychology literature on intention recognition, we hypothesize that intention attributions can be categorized into two types, L-intentionality and H-intentionality, based on the kinds of perceptions that cue those attributions. We describe the benefit of social robotics as a platform for experimenting with intentionality perceptions, and we mention some past research from robotics that explores intention attributions under various conditions. We then outline possible models for the relationship between H- and L-intentionality based on the set of features that elicit the perception of each: completely disjoint, partially overlapping, nested with H-intentionality as a proper subset, nested with

L-intentionality as a proper subset, and identical. Along with each model description, we specify how that model could be falsified. We posit some important consequences of proving our intention duality hypothesis, and we discuss the potential validity of these models, identifying that bottom-up processing is the most likely model given our separation of H- and L-intentionality. We discuss characteristics of intentionality that might vary between H- and L-intentionality, and we pose questions for future exploration of this area of research.

## Acknowledgments

Thanks to Greg Trafton for inspiring conversations, to Brad Hayes for help with diagrams, and to anonymous reviewers for helpful comments. This material is based upon work supported by grants from the National Science Foundation under contracts No. 1139078, No. 1117801, and No. 0835767. The first author is supported by an National Science Foundation Graduate Research Fellowship.

## References

Aarno, D., & Kragic, D. (2006). Layered HMM for motion intention recognition. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems (IROS 2006)* (pp. 5130–5135). Beijing, China.

Bainbridge, W. A., Hart, J. W., Kim, E. S., & Scassellati, B. (2011). The benefits of interactions with physically present robots over video-displayed agents. *International Journal of Social Robotics*, 3, 41–52.

Baker, C. L., Tenenbaum, J. B., & Saxe, R. R. (2006). Bayesian models of human action understanding. In Y. Weiss, B. Schölkopf, & J. Platt (Eds.), *Advances in neural information processing systems 18*. MIT Press.

Baron-Cohen, S. (1995). *Mindblindness: An essay on autism and theory of mind*. Cambridge, MA: MIT Press.

Breazeal, C., & Scassellati, B. (1999). How to build robots that make friends and influence people. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems (IROS '99)* (Vol. 2, pp. 858–863). Kyongju, South Korea.

Feldman, J., & Tremoulet, P. D. (2008). *The attribution of mental architecture from motion: Towards a computational theory* (Tech. Rep. No. 87). Rutgers University Center for Cognitive Science (RuCCS).

Gao, T., Newman, G. E., & Scholl, B. J. (2009). The psychophysics of chasing: A case study in the perception of animacy. *Cognitive Psychology*, 59, 154–179.

Heider, F., & Simmel, M. (1944, April). An experimental study of apparent behavior. *The American Journal of Psychology*, 57(2), 243–259.

Heyes, C., & Dickinson, A. (1990). The intentionality of animal action. *Mind and Language*, 5(1).

Knobe, J. (2005, August). Theory of mind and moral cognition: exploring the connections. *Trends in Cognitive Sciences*, 9(8).

Leslie, A. M. (1987). Pretense and representation: The origins of “theory of mind”. *Psychological Review*, 94(4), 412–426.

Meltzoff, A. N. (1995). Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, 31(5), 838–850.

Michotte, A. (1963). *The perception of causality*. Oxford, England: Basic Books.

Mutlu, B., Yamaoka, F., Kanda, T., Ishiguro, H., & Hagita, N. (2009, March). Nonverbal leakage in robots: Communication of intentions through seemingly unintentional behavior. In *Human robot interactions (HRI'09)*. La Jolla, California: ACM.

Nehaniv, C. L., Dautenhahn, K., Kubacki, J., Haegele, M., & Parlitz, C. (2005). A methodological approach relating the classification of gesture to identification of human intent in the context of human-robot interaction. In *Proceedings of the IEEE international workshop on robot and human interactive communication (ROMAN 2005)* (pp. 371–377).

Pantelis, P. C., Cholewiak, S., Ringstad, P., Sanik, K., Weinstein, A., Wu, C.-C., et al. (2011). Perceptions of intentions and mental states in autonomous virtual agents. *Journal of Vision*, 11(11), 1990–1995.

Pelphrey, K. A., Morris, J. P., & McCarthy, G. (2004). Grasping the intentions of others: The perceived intentionality of an action influences activity in the superior temporal sulcus during social perception. *Journal of Cognitive Neuroscience*, 16(10), 1706–1716.

Phillips, W., Baron-Cohen, S., & Rutter, M. (1998). Understanding intention in normal development and in autism. *British Journal of Developmental Psychology*, 16, 337–348.

Scholl, B. J., & Tremoulet, P. D. (2000, August). Perceptual causality and animacy. *Trends in Cognitive Sciences*, 4(8).

Schrempf, O. C., Albrecht, D., & Hanbeck, U. D. (2007). Tractable probabilistic models for intention recognition based on expert knowledge. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems (IROS 2007)* (pp. 1429–1434). San Diego, CA.

Short, E., Hart, J., Vu, M., & Scassellati, B. (2010). No fair!! an interaction with a cheating robot. In *5th ACM/IEEE international conference on human-robot interaction* (pp. 219–226).

Tremoulet, P. D., & Feldman, J. (2000). Perception of animacy from the motion of a single object. *Perception*, 29, 943–951.

Tremoulet, P. D., & Feldman, J. (2006). The influence of spatial context and the role of intentionality in the interpretation of animacy from motion. *Perception and Psychophysics*, 68(6), 1047–1058.

Wellman, H. M., Cross, D., & Watson, J. (2001, May/June). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, 72(3), 655–684.