

Distinguishing Between Intentional and Unintentional Sequences of Actions

Elisheva Bonchek Dokow (le7.bonchek.dokow@gmail.com)

Gonda Multidisciplinary Brain Research Center, Bar-Ilan University

Gal A. Kaminka (galk@cs.biu.ac.il)

Department of Computer Science and Gonda Multidisciplinary Brain Research Center, Bar-Ilan University

Carmel Domshlak (dcarmel@ie.technion.ac.il)

Faculty of Industrial Engineering and Management, Technion

Abstract

Human beings, from the very young age of 18 months, have been shown to be able to extrapolate intentions from actions. That is, upon viewing another human executing a series of actions, the observer can guess the underlying intention, even before the goal has been achieved, and even when the performer failed at achieving the goal. We identify an important preliminary stage in this process, that of determining whether or not an action stream exhibits any intentionality at all. We propose a model of this ability and evaluate it in several experiments.

Keywords: Intention; Cognitive Modeling.

Introduction

The topic of imitation has been the focus of much research in cognitive science and psychology (Meltzoff & Decety, 2003), neurophysiology (Rizzolatti, Fogassi, & Gallese, 2001), and artificial intelligence. Understanding the mechanisms underlying imitation and the time-line of their development is a part of understanding *Theory of Mind* and other aspects of social cognition. The AI community tries to model and implement this ability in software agents and robots, for the purpose of producing socially intelligent systems that can interact more meaningfully and usefully with humans.

Many different types of imitation exist, from the lower levels of gestural, facial and vocal mimicking to the higher level of goal imitation. The latter—the ability to understand the intention underlying a stream of actions, and reproduce the intended goal—is the type that we focus on here. How exactly this process takes place is yet an open question, and different researchers have addressed different aspects of it.

One of the more intriguing studies done in this area is by Meltzoff (1995), who has shown that 18-month old children are able to imitate the goal of an acting adult, *even when all they see is a series of failed attempts*. However, children are not able to do this when they observe arbitrary, intention-less, motions. These results, according to Meltzoff, assert the presence of some form of Theory of Mind at this young age.

Artificial systems have yet to reach a performance level comparable to that reported by Meltzoff. Much of the work on modeling this ability has focused on identifying the goal itself. Rao, Shon, and Meltzoff (2007) lay forth a Bayesian model for imitating goals that have been realized, and state that they intend to develop it in order to handle unrealized goals as well. Hongeng and Wyatt (2008) parse visual input and attempt to infer the goal before it is completed based on

visual cues such as color and shape. However, when dealing with intentions that have not been realized—i.e., when the acting agent failed at achieving its goal—the problem becomes much more challenging. Since the observed end-state in this case is not necessarily a goal, the observing agent must first determine whether or not there is anything worth imitating here, that is, if the actions were performed with a goal in mind, and only then can it proceed to attempt to infer what exactly that goal was.

Indeed, the open challenge we tackle in this work is that of identifying whether or not an action stream has any underlying intention at all. In Meltzoff's setup (described in more detail later), the behavior of the control groups has shown that when action streams did not have any underlying intention, the observing children did not attempt to imitate the acting adult. This is crucial, since before the observing agent embarks on the intimidating task of guessing what the goal actually is, it would be wise to first decide whether there is any goal to look for.

In this paper we model this ability of discerning intentional action from unintentional action. The key idea underlying our work is the principle of rational action, which states that an agent that has a goal will take actions to achieve this goal. Inspired by this principle, we determine the intentionality of observed sequences of actions by looking at whether they are *efficient*, i.e., they monotonically move the agent further away—in problem state space—from its initial state.

We evaluate the model in two very different environments. First, we reproduce two of Meltzoff's experiments in a discrete version, using STRIPS notation¹, and show that our method results are compatible with his. Second, we report on experiments in which our method results were contrasted with adult human judgment of surveillance videos. While we only have preliminary results in this environment, they are very promising and show that our method tends to evaluate motions similarly to humans.

Background and Related Work

There is a vast amount of literature on the general topic of imitation and on, specifically, goal imitation. We cannot hope to cover it all here. We note that throughout the paper, we use the terms "goal" and "intention" colloquially, while a

¹Formal language for describing states and actions in AI planning (Fikes & Nilsson, 1971).

clear distinction is sometimes made between them in previous work, e.g., (Tomasello, Carpenter, T. Behne, & Moll, 2005).

From the computational research, we refer here only to two of the more recent ones on goal inference. Meltzoff himself took a first step in this direction (Rao et al., 2007), by modeling the task in a Bayesian framework. They trained their model on several example trajectories leading to different goals, so that when given a test scenario the model could determine the goal, before it was reached. Hongeng and Wyatt (2008) analyze real-world video input, and use learning algorithms to determine higher-level goals from low level movement. Both these works build on past experience—multiple exposures to a limited set of possible goals, and learning actions that are associated with them. They also both assume intentionality, and therefore go directly to the task of inferring what that intentionality is. Thus our work on recognizing intentionality complements theirs.

Harui, Oka, and Yamada (2005) attempt to determine whether intentionality is present at all. However, their results are based mainly on vocal cues, such as "oops", to signal an accidental action as opposed to an intentional one. We ignore such features, since in Meltzoff (1995)'s paradigm they were neutralized. No one else, to the best of our knowledge, has attempted to computationally identify intentionality in action.

There are several psychological theories regarding the stance taken when dealing with intentionality. Meltzoff (2002) takes the mentalistic stance that infants' ability to interpret intentionality makes use of an existing theory of mind—reasoning about the intents, desires and beliefs of others. Gergely and Csibra (2003), on the other hand, take a teleological stance, that infants apply a non-mentalistic, reality-based action interpretation system to explain and predict goal-directed actions. As Gergely and Csibra say themselves, this teleological evaluation should provide the same results as the application of the mentalistic stance as long as the actor's actions are driven by true beliefs, as is our case.

The principle of rational action (Gergely & Csibra, 2003; Watson, 2005) plays a major role in intentional action. It states that intentional action functions to bring about future goal states by the most rational actions available to the actor within the constraints of the situation. In other words, intentional action is necessarily efficient and as such, proceeds monotonically away from the initial state.

A Method of Intentionality Recognition

We first describe briefly Meltzoff's 1995 experiments. We then present our technique for determining intentionality.

Motivation

In order to understand the motivation for our model, as well as the setup used to evaluate it, we briefly describe some details of Meltzoff's experiment. The purpose of his experiment was to test whether children of 18-months of age are able to understand the underlying intention of a sequence of actions, even when that intention was not realized (the acting agent failed to achieve the goal).

For five different novel toy objects, a target action was chosen. For example, for a two-piece dumbbell-shaped toy, the target action was pulling it apart. For a loop and prong device, the target action was to fit the loop onto the prong. The children were divided into four groups—"Demonstration Target", "Demonstration Intention", "Control Baseline" and "Control Manipulation". The children in the "Demonstration Target" group were shown three repetitions of a successfully completed act, such as pulling apart the dumbbell, or hanging the loop on the prong; their voluntary response was to reproduce the same act when the objects were handed to them. The children in the "Demonstration Intention" group were shown three *failed attempts* by the adult to produce the goal, where the adult (seemingly) failed at reaching it. These children's re-enactment of the goal reached a level comparable to that of the children who saw the successful attempts. This shows that children can see through the actions to the underlying intention, and extrapolate the goal from the actions. The children in the "Control Manipulation" group saw the object manipulated three times in ways that were not an attempt to reach the chosen target act. This was done in order to make sure that mere manipulation of the object is not enough for the children to reproduce the goal. The last control group—"Control Baseline"—had the children just see the object, without it being manipulated at all. Both control groups did not show significant success at reproducing the target act.

Meltzoff's experiment shows that when children discern an underlying intention, as in the two Demonstration groups, they attempt to imitate it. When they do not detect such an intention, as in the Control groups, they do nothing, or sometimes mimicked the arbitrary acts of the adult (in the "Control Manipulation" group; obviously, children were imitating *what they understood to be* the intention of the adult).

Thus a model of goal imitation must first be able to model the ability to discern whether there is an underlying intention. Only then is it relevant to attempt to discern what that intention is. This would explain why children in both "Demonstration" groups were motivated to look for an underlying intention, while children in the "Control Baseline" group were not. This also explains why children in the "Control Manipulation" group sometimes reproduced the actions of the adult, even when it was not exactly what the experimenter had in mind. As long as the trace exhibited some "rationality of action", or efficiency, the children concluded that there was an intention worth imitating.

Recognizing Intentionality

We denote the observation trace by $t = s_0, \dots, s_k$, i.e. a sequence of states, brought about by the actions of the demonstrating agent. s_0 is the initial state, and s_k is the terminal state. The task of the observing agent is to decide, given this trace, whether there was an underlying intention or whether the acting agent behaved unintentionally.

Inspired by the principle of rational action, we check for some form of efficiency in the trace. It is reasonable to expect that a trace with an underlying intention will exhibit a

clear progression from the initial state towards the goal state, which is the most efficient way to bring about that goal, starting from the initial state. Note that we do not know at this stage whether or not there is an underlying goal to the trace, and even if there is, if it is reached successfully. On the other hand, unintentional traces would not be driven by such efficiency, and would fluctuate towards and away from the initial state, without any clear directionality.

To do this, we define a distance measure *dist*. This distance measure is dependent on the nature of the world being modeled. For example, when dealing with geographical targets, the distance could simply be the Euclidean (and indeed it is, in one of our experiments). In a discrete state-space, defined by STRIPS notation, we use Bonet and Geffner (1999)'s Heuristic Search Planner to generate optimal plans from the initial state to every state in the trace, and the number of action steps in each generated plan is taken to be the distance to the respective state. If the demonstrating agent acts efficiently—taking only optimal action steps that bring it closer to the goal—then the distance will keep increasing. While if it acts randomly, executing various actions that do not necessarily lead anywhere, the distances will fluctuate.

There are a few requirements for the distance measure. We do not require this distance to obey symmetry ($d(s_1, s_2) = d(s_2, s_1)$). However, this distance should always be positive and equal 0 only from a state to itself. Using any such distance measure, we capture the notion of optimality, in the sense of a shortest path from one state to another.

Thus from the original state trace we induce a sequence of distance measurements $d_1 = \text{dist}(s_1, s_0), \dots, d_k = \text{dist}(s_k, s_0)$, measuring the *optimal (minimal) distance* between each state in the sequence, and the initial state. Thus, for every state, we have an indication of how much the demonstrating agent would have had to invest (in time, number of elemental actions, or any other resource, depending on how the distance is defined), had it been intending to reach that state. We argue that enough information is preserved in this sequence for our observing agent to come to a satisfying decision.

We want to calculate from this sequence a measure of intentionality, which we take to be the proportion of local increases in the sequence—at how many of the states along the trace has the distance from the initial state increased as compared to the previous state, out of the total number of states in the trace. This will give us an idea of how efficient the action sequence is. More formally,

$$u = |\{d_i > d_{i-1}\}_{i=1}^k| \quad (1)$$

is the number of states in the trace where the distance from the initial state increases, as compared to the distance at the previous state. Taking this number and dividing it by the total number of states in the trace,

$$p = \frac{u}{|\{d_k\}_{i=1}^k|} \quad (2)$$

gives us a measure of intentionality for the action sequence.

The higher the resulting p , the more intentionality is attributed to the action. If a binary answer is preferred, we can determine a cutoff level above which we conclude intentionality is present, and below which we conclude it is not.

For example, in the case of clear intentionality, we would expect a strictly monotonically increasing sequence of distances; the agent proceeds from the initial state, at each step moving farther and farther away from it, and closer and closer to the intended goal. At the other end, if the observed agent is not driven by an intention to reach any particular state, we would expect the sequence to fluctuate in a seemingly random fashion, with the agent sometimes moving away from the initial state and sometimes moving back towards it. Of course, this is merely a motivational argument. In the next section we show that this simple intuitive method does indeed produce the expected results.

Implementation and Evaluation

In order to evaluate the success of our proposed measure of intentionality, we implemented it in two different environments. The first uses a discrete abstraction of Meltzoff's experiments, modeled in standard AI planning problem description (STRIPS), and the second uses surveillance videos.

Discrete Versions of Meltzoff's Experiments

We model Meltzoff's experiment environment as an 8-by-8 grid, with several objects and several possible actions which the agent can execute with its hands, such as grasping and moving. We implemented two of the five object-manipulation experiments mentioned by Meltzoff: The dumbbell and the loop-and-prong. For the dumbbell, there is one object in the world, which consists of two separable parts. The dumbbell can be grasped by one or both hands, and can be pulled apart. For the loop-and-prong, there are two objects in the world, one stationary (the prong), and one that can be moved around (the loop). The loop can be grasped by the hand, and released on the prong or anywhere else on the grid. As previously described, we use Bonet and Geffner (1999)'s HSP to compute the distance measure.

We manually created several traces for the dumbbell and for the loop-and-prong scenarios, according to the descriptions found in Meltzoff's experiment, to fit the four different experimental groups. For example, a visual representation of the "Demonstration Target" trace is given for the dumbbell object in Figures 1(a)–1(i).

In addition, we created a random trace, which does not exhibit any regularity. We added this trace since the children in Meltzoff's "Control Manipulation" group were sometimes shown a sequence with underlying intention, albeit not the target one. For each trace we calculated the sequence of distances, using the above mentioned HSP algorithm, and then computed the proportion p .

Results

Figure 2 show some plots of the sequence of distances associated with the Dumbbell experiments. The step number in the sequence is measured in the X axis. The Y axis shows




Figure 1: Dumbbell Demonstration Target (left to right, top to bottom).




Figure 2: Distance as a function of state in sequence in the Dumbbell experiments.

the distance. Figure 2(a) shows an almost perfectly monotonically increasing distance trace for the "Demonstration Intention II" trace, where the right hand slips off the dumbbell, and so returns to the state it was at before it grasped it. Since only 10 out of 12 of the states showed an increase in the distance from the initial state, relative to the previous state, the intentionality score is 10/12. Figure 2(b) shows the distance sequence for the "Random" trace. Here the graph fluctuates, demonstrating the unintentionality of the trace.

Table 2 shows the calculated measure of intentionality, for each of the traces in the prong-and-loop experiment, and Table 1 shows the same for the dumbbell experiment. In both tables, each row corresponds to a different type of state sequence. The right column shows the measure of intentionality as computed by the method described above.

In Meltzoff's experiments, every child was shown three traces, and only then was handed the objects. There is certainly information in this seeming redundancy; see (Meltzoff, Gopnok, & Repacholi, 1999) who show that when only one

trace was shown to the "Demonstration Intention" group, the children were unable to reproduce the goal. However, we do not treat this at this stage in our model. So, while every child was shown three possibly different traces, we calculated our measure of intentionality separately for each of these traces, which is why we have more than one row in the table for some of the groups.

For example, the prong-and-loop procedure failed in two different ways in Meltzoff's "Demonstration Intention" experiment—either with the loop being placed too far to the right of the prong ("Demonstration Intention I" in Table 2), or too far to the left ("Demonstration Intention II"). Both these actions received an intentionality score of 1, since the end-state was reached in the most efficient possible way. In the discussion section we elaborate on the meaning of this.

The dumbbell procedure as well failed in two different ways—with the right hand "accidentally" slipping off the dumbbell while trying to pull it apart ("Demonstration Intention I" in Table 1), or with the left hand slipping off ("Demonstration Intention II"). When the right hand slipped off it ended up slightly closer to the point where it was before the action was initiated, as opposed to where the left hand ended up when it slipped off. For this reason, the intentionality measure for "Demonstration Intention I" is slightly lower than for "Demonstration Intention II".

Trace	Measure of Intentionality
Demonstration Target	1
Demonstration Intention I	0.8333
Demonstration Intention II	0.9166
Control Baseline	0
Control Manipulation	0.8333
Random	0.5384

Table 1: Calculated measure of intentionality for STRIPS implementation of the dumbbell experiment.

Trace	Measure of Intentionality
Demonstration Target	1
Demonstration Intention I	1
Demonstration Intention II	1
Control Baseline	0
Control Manipulation I	0.7777
Control Manipulation II	0.7777
Control Manipulation III	1
Random	0.5555

Table 2: Calculated measure of intentionality for STRIPS implementation of the prong-and-loop experiment.

In both experimental setups, the "Demonstration Target" trace received a clear score of 1, the highest possible intentionality. This happened because every step in the trace was necessary for bringing about the goal in the most efficient way—each and every state progressed away from the initial state and towards the goal state. The "Control Baseline" trace received a 0, since nothing at all happened in that trace—the

world remained static, at the initial state, without any change throughout the trace. The "Random" trace received a low score, just a bit above 0.5, since the number of states progressing away from the initial state was roughly equal to the number of states returning towards it. The "Demonstration Intention" traces exhibited a significant measure of intentionality, as did the "Control Manipulation". The latter can be explained by observing, as mentioned above, that even when the adults manipulated the objects in a way that was not the original intention of the experimenter, nevertheless the manipulation **did** exhibit an intentionality to reach **some** state, as opposed to just wandering about aimlessly in the space of possible states. For the dumbbell object, the arbitrary act was pushing the ends inwards (this same act was demonstrated three times). For the prong-and-loop object, the arbitrary acts were moving the loop along an imaginary line above the prong, from right to left ("Control Manipulation I"), from left to right ("Control Manipulation II"), and placing it just below the prong ("Control Manipulation III"). This last act received the ultimate intentionality score, since the end-state was reached by the most direct path.


Video Experiment

A second set of experiments was carried out in order to compare our model's results to those of human observers. In particular, we are interested in how human observers classify real-life human movement, and whether their judgment of intentionality correlates with those of our model. To test this, we used the CAVIAR video repository of surveillance videos. We selected a dozen movies from the repository. With respect to intentionality, these range from movies that show very deliberate movements (a person crossing a lobby towards an exit), to some that are less clear (a person walking to a paper stand and browsing, then moving leisurely to a different location, etc.). We compared human subjects' judgment of the intentionality of motions in these videos, to the predictions of our model.

Let us begin by describing how we measure intentionality using our model. The ground truth position data of the selected videos is a part of the repository, and we use it as a basis for our intentionality measurements. The planar coordinates of the filmed character in every frame in the video were taken as a state in the trace, and the distance measure we used was the Euclidean distance. As above, for every state we calculated the distance from the initial state, and then checked for how many of those states the distance increased, relative to the previous state.

Figure 3(a) shows a graph of the path of movement of the observed character, in planar coordinates, in one of the videos from the repository (video `bww1_gt`). Because we are plotting planar coordinates, the amount of time spent at each point is not represented here. Figure 3(b) shows a plot of the distances of each state in the path, from the initial state. The X axis measures the video frame number. The Y axis measures the distance from the initial location of the person in question. For example, the measure of intentionality for this movement

path was $p = 0.48133$. Using a cutoff value of 0.5, this movement was classified as non-intentional. The interested reader is invited to watch the video and compare it to the graphs presented here.



(a) Path of movement. (b) Distances of each state from initial state.

Figure 3: Examples from the `bww1_gt` video.

Those same videos were shown to human subjects who were asked to write down their opinion regarding the intentionality of the viewed character. They were given the option of segmenting the video if they thought the character changed its intention along the trace. Here we faced some difficulty in the experiment design. In pilot experiments, it became clear that asking the subjects to directly rank the "strength of intentionality" of a video segment leads to meaningless results. For instance, some subjects in pilot experiments chose to give high intentionality marks to a video segment showing a person seemingly walking around aimlessly. When we asked for an explanation, the answer was that the person in the video clearly intended to pass the time.

We thus needed to measure intentionality indirectly. To do this, subjects were requested to write down a sentence describing the intent of the person in the video, typically beginning with the words "The person intends to ...". The idea behind this is that in segments where there is clear intentionality, a clear answer would emerge (for instance, "The person intends to exit the room"); in other video segments, the unclear intentionality would result in more highly varied answers (e.g., some would write "intends to pass the time", while others would write "intends to walk", etc.). This divergence can be measured by various means; we chose the information entropy function as it is used in statistics to measure dispersion of categorical data.

Results

We unfortunately did not complete the final analysis of the results. However, preliminary results seem to indicate that our model's classification of the movement as intentional correlates with the results obtained from the human subjects. In particular, in videos showing clear goals the human subjects tend to agree on the way the intention is described. In videos that are less clear, there is indeed divergence of the answers. Moreover, the divergence is also temporal: In movies where the goal is unclear, subjects disagreed not only on the description, but also on the internal segmentation of the video clip into segments of changing intentions. Some subjects cut the movie into several segments, while others did not. They also did not agree on the timing of the segments. Such disagreement was not noticed in the clearer movie clips.

Discussion

This work measures intentionality using a very basic feature of the stream of action. We ignore other aspects of the dynamics of the movement that certainly contain information regarding intentionality. Moreover, we assume a state-space of sufficient resolution and detail. We find justification for this in the psychological literature. Blakemore and Decety (2001) quote several works on how static images convey dynamics. Meltzoff (2007) himself uses such a discretization in yet another variant of his original experiment. In this version, instead of showing the children the full dynamics of the action, he showed them three successive static states. This technique assumes that such a representation contains enough of the information regarding the intent of the actor. In the same paper, Meltzoff also describes the failed attempt to separate the dumbbell as "hold the dumbbell and then remove one hand quickly", which is again a very physical description, similar to the way we modeled the experiment. Although it does not convey the notion of "effort", this description is yet enough to give the children a sense of intentionality.

Another point worth addressing is the high intentionality scores that some of the demonstrations received—at times the highest possible ($p = 1$), equal to that of the "Demonstration Target" group. We stress again that we are dealing here with a preliminary stage in the process of goal imitation, that of intentionality detection. It would be wrong to conclude that a maximal score of intentionality indicates *success* at achieving the goals. Rather, we only conclude intentionality of the action and leave the question of whether the reached end-state was indeed the intended goal for a later stage.

Our model also does not deal with the fact that the demonstrations were repeated three times for every child. This information can also be used in determining intentionality (see, for example, Watson (2005) who mentions persistence as a sign of intentionality), as well as for the later stage of determining whether the reached end-state is the intended goal.

Future Work

Having only just touched the tip of the iceberg regarding the intriguing phenomena of intentionality detection and goal imitation, there is yet much work to be done. In addition to more rigorously testing and evaluating our current model, we intend to broaden it to deal with the notions of persistence and equifinality—information carried by the repetition of every demonstration three times. It would also be interesting to add the possibility of handling varying environmental constraints, such as obstacles, which affect the calculation of the distance measure, as well as treating false beliefs regarding those environmental constraints, and seeing how they affect the conclusion reached regarding intentionality.

Acknowledgments. Videos were taken from EC Funded CAVIAR project/IST 2001 37540. This research was partially supported by ISF grant #1357/07.

References

Blakemore, S. J., & Decety, J. (2001). From the perception of action to the understanding on intention. *Nature reviews, neuroscience*.

- Bonet, B., & Geffner, H. (1999). Planning as heuristic search: new results. In S. Biundo & M. Fox (Eds.), *Proceedings of the 5th european conference on planning*. Durham, UK: Springer: Lecture Notes on Computer Science.
- Fikes, R. E., & Nilsson, N. J. (1971). Strips: a new approach to the application of theorem proving to problem solving. *Artificial intelligence*.
- Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: the naive theory of rational action. *TRENDS in cognitive science*, 7(7).
- Harui, K., Oka, N., & Yamada, Y. (2005). Distinguishing intentional actions from accidental actions. In *Proceedings of the 4th IEEE international conference on development and learning*.
- Hongeng, S., & Wyatt, J. (2008). Learning causality and intentional actions. In E. Rome, J. Hertzberg, & G. Dorffner (Eds.), *Towards affordance-based robot control*. Springer.
- Meltzoff, A. N. (1995). Understanding the intentions of others: re-enactment of intended acts by 18-month-old children. *Developmental psychology*, 31(5).
- Meltzoff, A. N. (2002). Imitation as a mechanism of social cognition: origins of empathy, theory of mind, and representation of action. In U. Goswami (Ed.), *Blackwells handbook of childhood cognitive development*. Blackwell.
- Meltzoff, A. N. (2007). The "like-me" framework for recognizing and becoming an intentional agent. *Acta psychologica*.
- Meltzoff, A. N., & 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*.
- Meltzoff, A. N., Gopnok, A., & Repacholi, B. M. (1999). Toddlers' understanding of intentions, desires and emotions: exploration of the dark ages. In P. D. Zelazo, J. W. Astington, & D. R. Olson (Eds.), *Developing theories on intention: social understanding and self control*. Lawrence Erlbaum Associates.
- Rao, R. P. N., Shon, A. P., & Meltzoff, A. N. (2007). A bayesian model of imitation in infants and robots. In C. L. Nehaniv & K. Dautenhahn (Eds.), *Imitation and social learning in robots, humans, and animals: behavioural, social and communicative dimensions*. Cambridge University Press.
- Rizzolatti, G., Fogassi, L., & Gallese, V. (2001). Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature reviews in neuroscience*.
- Tomasello, M., Carpenter, M., T. Behne, J. C. abd, & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and brain sciences*.
- Watson, J. S. (2005). The elementary nature of purposive behavior: evolving minimal neural structures that display intrinsic intentionality. *Evolutionary psychology*.