What should a robot learn from an infant?

Mechanisms of action interpretation and observational learning in infancy

György Gergely
Institute for Psychology Hungarian Academy of Sciences
Victor Hugo u. 18-22
1132 BUDAPEST
gergelyg@mtapi.hu

Abstract

The paper provides a summary of our recent research on preverbal infants (using violation-of-expectation and observational learning paradigms) demonstrating that one-year-olds interpret and draw systematic inferences about other's goal-directed actions, and can rely on such inferences when imitating other's actions or emulating their goals. To account for these findings it is proposed that one-year-olds apply a non-mentalistic action interpretational system, the 'teleological stance' that represents actions by relating relevant aspects of reality (action, goal-state, and situational constraints) through the principle of rational action, which assumes that actions function to realize goal-states by the most efficient means available in the actor's situation. The relevance of these research findings and the proposed theoretical model for how to realize the goal of epigenetic robotics of building a 'socially relevant' humanoid robot is discussed.

1. Introduction: Can infancy research inform robotics about how to interpret and learn from intentional actions of other agents?

Having been asked to talk about my infancy research at this workshop on Epigenetic Robotics, I feel I must start by making it clear that I can by no means be considered to be an artificial intelligence researcher. In fact, I am just an "old fashioned" cognitive developmentalist using experimental behavioral techniques such as "violation-of-expectation" or imitative learning paradigms to study the interpretative mechanisms and inferential capacities of infants' early understanding of intentional actions, and their ability to learn novel means through observation and imitation of the goal-directed actions of others. While in my work I have been developing abstract conceptual models intended to capture the remarkable complexities of these early competences, I have never written a connectionist program that would simulate these abilities, not to speak about actually building a robot that would do so.[1]

Given this background, you can imagine that I was slightly (though admittedly, pleasantly) surprised (but also somewhat incredulous at first) when I received the invitation to give a keynote address to the 3rd International Conference of Epigenetic Robotics. After some checking to make certain that the invitation wasn't sent by chance to the wrong e-mail address, my initial (and admittedly somewhat narcissistic) joyful reaction soon turned into a slight but distinctly unpleasant feeling of anxiety: Is there really anything instructive or relevant that our own research programme can offer to experts in AI and AL who are trying to build humanoid robots capable of acting in a goal-directed manner and of learning new ways of acting so by imitating other agents? So I did a little bit of home-work to see what kind of questions are being currently pursued by researchers in epigenetic robotics and how these questions flare with the type of problems that our developmental and experimental approach focuses on. I must admit that I was pleasantly surprised and intellectually quite intrigued to find an unexpected amount of significant convergence, but I also noted some important differences in the dominant focus of current research that characterizes the two approaches.

Briefly, my impression is that recent research in epigenetic robotics has been strongly preoccupied with and made significant advances towards modeling the "lower level" mechanisms involved in action perception and production and the ways in which these mechanisms maybe inherently related. The basic issues

[1] I did engage though in challenging arguments with 'real' AI researchers who attempted to replace our conceptual models with specific connectionist simulations that, they claimed, were able to account for our results without making use of our model's abstract constructs of folk psychology such as 'goals' and 'principles of rationality or efficiency'.

addressed include questions about how simple and more complex motor actions can be generated, what is the nature and role of "motor primitives" and forward models in action planning and control, or how to design engineering solutions to the thorny "correspondence problem" of establishing the perceptual-motor mapping or equivalence between the same or similar actions when they are observed and when they are being executed, to name just a few of the focal research questions of current-day robotics (e. g., Breazeal and Scassellati, 2002; Dautenhahn and Nehaniv, 2002; Schaal, 1999).

In these areas there has clearly been a significant amount of healthy interdisciplinary cross-fertilization between AI and robotics, and important advances in recent computational models (e. g., forward models) of intentional action planning and control (e.g., Wolpert et al., 2001), and the new discoveries made in cognitive neuroscience such as the identification of specialized neurons in the superior temporal sulcus (ST) of macaques that are sensitive to highly specific actions of body parts (e.g., Perrett et al., 1985) or of mirror neurons in the F5 region of the macaque brain that are sensitive to both the perception and the motor execution of specific goal-directed object manipulative actions (e. g., Rizolatti et al., 1988).

In contrast, when it comes to our own research on infants' understanding of intentional actions, the questions addressed are at a qualitatively different level, and the types of, admittedly extremely hard and exciting, research issues listed above concerning the basic mechanisms that mediate action perception, generation, and matching, are not directly dealt with or investigated, rather, they are mostly simply presupposed by our approach. Instead, our research programme has focused on "higher-order" cognitive processes in an attempt to characterize the *representational and inferential mechanisms* and *built-in "top-down" architectural constraints* of the *domain-specific action interpretational system* that, we believe, underlies young infants' demonstrated (and, as I hope to convince you, truly remarkable) early capacity to identify, infer, and attribute goals to observed actions of others, and their ability to infer which (or which aspects of) perceived intentional actions of other agents they should imitate during observational learning. It is for this reason (of qualitatively different levels of investigation) that I first felt rather uncertain as to whether and to what degree our demonstrations and models may prove to be informative to the central concerns of researchers in epigenetic robotics.

Eventually, however, I have become cautiously confident that the time may, in fact, have come for our infancy work on these "higher-level" cognitive constraints and interpretive mechanisms to provide useful information and possibly even to suggest new directions for future research for epigenetic robotics. This more optimistic diagnosis certainly seems warranted by such reoccurring programmatic statements that have recently been made by leading robotics researchers about the anticipated new directions and questions that they feel their research domain needs to turn to and tackle in the near future to make significant further progress possible. Just to give you a few examples: Cynthia Breazeal and Brian Scassellati (2002) in a recent review in Trends in Cognitive Sciences raise as their focal issue the question: "How does a robot know *what* to imitate?". Furthermore, in their "Questions for future research" they ask: "Just as children develop the ability to imitate the goal of an action rather than a specific act, can we construct robots that are capable of making this inference? Today's robots – they say - respond only to the observable behavior without any understanding of the intent of an action." (p. 486). Then they go on to ask: "Who should the robot learn from, and when is imitative learning appropriate? Robots that imitate humans today are programmed to imitate any human within view" (p. 486). In a similar vein, one of the "Outstanding question" of Stefan Schaal's (1999) excellent recent review in Trends in Cognitive Sciences is this: "*Understanding task goals*: How can the intention of a demonstrated movement be recognized and converted to the imitator's goal?" (p. 240). These new questions clearly signal a growing need within epigenetic robotics to move towards the "higher order" cognitive issues that are at the very center of our own research inquiries about infants' interpretative capacities within the domain of action understanding.

Therefore, I hope that by summarizing our major empirical findings of preverbal infants' capacity to interpret, reason about, and learn from the observed intentional actions of other agents (Csibra et al., 1999; 2003; Gergely et al., 1995; 2001, 2002) and by outlining our theoretical account of this early competence that we call the one-year-old's *'naïve theory of rational action'* or the *'teleological stance'* (Gergely and Csibra, 1997, 2003; Csibra and Gergely, 1998), I'll be able to make a useful contribution to the newly forming interdisciplinary dialogue between cognitive infancy research and epigenetic robotics.

2. What one-year-old infants understand about the goal-directed actions of other agents.

Before presenting supporting empirical evidence, let me first provide you with a brief list of those capacities that our research has demonstrated in one-year-old infants in the domain of action interpretation and observational learning and that seems to me to coincide with (or even go beyond) the new and ambitious research targets that the epigenetic robotic movement has set for itself. In other words, I suggest that these are the basic competencies that one-year-old human infants possess in this domain and that humanoid robots (or, if not them, then at least their designers) should learn from them:

(1) ATTRIBUTING GOALS TO ACTIONS: One-year-olds can interpret other agents' actions as goal-directed;

(2) EVALUATING THE RATIONALITY OR EFFICIENCY OF ACTIONS AS MEANS TO GOALS: One-year-olds can compare and evaluate which of the

alternative actions available to an agent within the physical constraints of a given situation is the most efficient means to the goal;

(3) PREDICTING NOVEL MEANS ACTIONS IN NEW SITUATIONS: By one year of age infants can form an active expectation that in order to realize her goal an agent 'ought to' perform the most rational or efficient means action available to her within the particular situation;

(4) INFERRING (NON-VISIBLE) ASPECTS OF GOAL-DIRECTED ACTIONS: Going "beyond the information given" – as Bruner famously said in 1957 (Bruner, 1957) - , one-year-olds can draw systematic and productive inferences to identify and (representationally "fill in") any one of the three basic aspects (Goal, Means action, and relevant Situational Constraints, see Figure 1) of the representation of an intentional action when that aspect is perceptually inaccessible to them, as long as they have direct access to the other two relevant aspects of the goal-directed action, on which they can base such an inference. In particular:

(4a) INFERRING A (NEW) *MEANS ACTION*: Given perceptual information about the Goal and the (changed) Situational Constraints, infants can infer (and predict) what novel Means action the agent 'ought to' perform to achieve its goal in the most efficient manner given the changed constraints of the situation;

(4b): INFERRING AN (UNSEEN) *GOAL*: Given perceptual information about the Situational Constrains and about the initial part of an Action (whose end-state is occluded from them), they can infer an (unseen) Goal that would justify the action as a rational or efficient means to the goal in the given situation; and

(4c): INFERRING (UNSEEN) *SITUATIONAL CONSTRAINTS*: Given perceptual information about the Action and about the Goal state, infants can infer a (non-visible) physical Situational Constraint (such as an obstacle that is occluded from their view) whose presence would justify the action as a rational or efficient means to the goal.

(5): TELEOLOGICAL EMULATION OF NEW GOALS AND RATIONAL IMITATION OF NOVEL MEANS IN OBSERVATIONAL LEARNING: When learning new actions to achieve a novel goal from observing an unfamiliar means action demonstrated by another agent, 14-month-olds can evaluate the rationality or efficiency of the observed means *both* in relation to the situational constraints of the model *and* in relation of their *own* situational constraints, and can use this information to decide whether to imitate the demonstrated novel means or to achieve the new goal through emulation.

3. The one-year-old's 'teleological stance' and the inferential 'principle of rational action'.

What makes these remarkable inferential feats possible for one-year-old infants who are, arguably, still lacking the metarepresentational means to attribute abstract and invisible causal intentional mental states (such as intentions, desires, and beliefs) to the agent's mind? (For contrary views that assume the early availability of at least some mentalistic representational aspects of a theory of mind by the end of the first year, see e.g., Tomasello, (1999); Kelemen, (1999).) To answer this question, Gergő Csibra and myself have proposed (Gergely and Csibra, 1997; 2003; Csibra and Gergely, 1998) that one-year-olds possess a *non-mentalistic (reality-based) teleological action interpretational system* or strategy, that we call the *'teleological stance'* (Figure 1) that establishes a teleological (rather than causal) explanatory relation among three relevant aspects of (current and future) reality: the observed behavior, a future state of reality (future in relation to the behavior), and the relevant aspects of physical reality that constrain possible actions in the particular situation in which the observed behavior

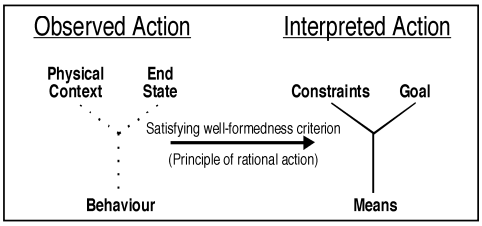

Figure 1: Teleological representation of goal-directed actions

unfolds. This action interpretational schema provides a *well-formed* (and thus acceptable) teleological representation of the observed behavior as an *efficient goal-directed action* only if the behavior can be evaluated as an effective (rational) way to bring about the future state given the physical constraints of the situation. If this well-formedness condition (that is articulated by the 'principle of rational action', see below) is satisfied by the representation in question, the future state will become encoded as the Goal, the behavior as a Means to the goal, and the relevant aspects of physical reality as the Situational Constraints on action (Figure 1).

We propose that such teleological action interpretations are driven by the core *'principle of rational action'* that captures our normative assumptions about the essentially functional nature of intentional actions (see Dennett, 1987; Gergely and Csibra, 2003). The rationality principle serves both as a *criterion of well-formedness* for teleological action interpretations and as an *inferential principle* guiding and constraining the construction of such action interpretations. In particular, the principle of rational action presupposes that a) *actions* function to bring about future *goal states*, and b) goal states are realized by the most rational (or efficient) action available to the actor within the *constraints of the situation*. Thus, the principle asserts that a teleological action explanation is well-formed (and therefore acceptable) if, and only if, the *action* realizes the *goal state* in a rational (or efficient) manner within the particular *situational* *constraints* (Figure 1).

4. Empirical evidence supporting the inferential productivity of the rationality principle in teleological action interpretations of one-year-olds.

Early understanding of goal-directed actions has been demonstrated using a variety of paradigms such as imitation (Carpenter et al., 1998; Gergely et al., 2001, 2002; Meltzoff, 1988; 1995), joint attention (Carpenter et al., 1998; Tomasello, 1999), and violation-of-expectation looking time studies (Csibra et al., 1999, 2003; Gergely et al., 1995; Király et al., 2003; Woodward, 1998; Woodward and Sommerville, 2000). Let me illustrate the complex nature of this understanding by one of our violation-of expectation studies (Gergely et al., 1995) (Figure 2).

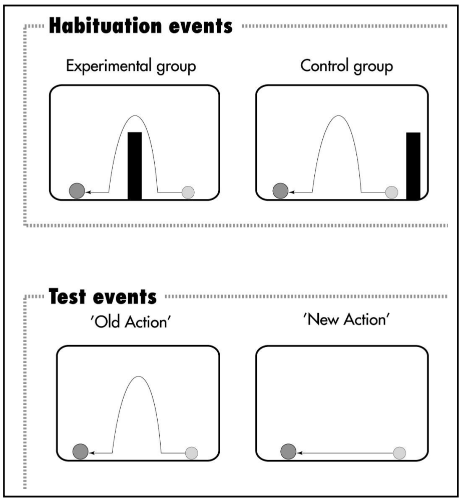

Figure. 2. Experimental and Control conditions of Gergely et al., (1995)

Twelve-month-olds were habituated to a computer-animated goal-directed action (Figure 2A) in which a small circle repeatedly approached and contacted a large circle (goal) by jumping over (means act) an obstacle separating them (situational constraint). Even though these 2D shapes had no human-like features, when we asked adults to describe what they see, they immediately interpreted this visual event as depicting an efficient means action to achieve a goal state (that of contacting the large circle), because they could justify the jumping approach as the most rational action available to realize that goal given the physical constraints of the situation (i. e., the presence of the 'obstacle' separating the two circles) (cf. Heider and Simmel, 1944). To test whether or not one-year-olds would also interpret this event in the same manner, we presented them, following habituation, with two types of test events (with their order of presentation being randomized across subjects) in which the obstacle was no longer present. In one of the test displays (Figure 2C) they saw again the already familiar 'jumping approach', which, however, could no longer be justified (by the presence of an obstacle) as a rational action to achieve the goal (the small circle jumped over empty space during its approach of the large circle). In contrast, in the other test event (Figure 2D) infants were presented with a perceptually novel, but sensible 'straight-line goal-approach' (that has become an available action alternative to get to the goal after the removal of the obstacle).

In spite of the fact that the old 'jumping approach' (2C) was perceptually similar to the means action presented during the habituation event (2A), subjects looked at it significantly longer (indicating violation-of-expectation) than at the novel 'straight-line approach' (2D) to which (even though it was perceptually novel) they showed no dishabituation at all. This suggests that the infants found the old 'jumping action' (2C) test event unexpected, because it seemed to them an inefficient way to reach the goal in the new situation where there was no obstacle to justify the jumping action as a rational means. In contrast, the fact that they did not dishabituate to the novel 'straight-line goal-approach' (2D) indicates that this action was expectable for them (in spite of its perceptual novelty) as it appeared to be the most efficient means to the goal that has become available after the disappearance of the obstacle.

Appropriate control conditions (see Csibra et al., 1999; Gergely et al., 1995) ruled out obvious alternative explanations. In the control study, during the habituation phase the rectangular object appeared *behind* the small circle (Figure 2B) and so it did not form an obstacle towards the goal object (making the more efficient straight-line goal-approach available already during the habituation event). In spite of this, the small circle approached the large circle through the same jumping action as in the experimental condition. Note, however, that in the Control condition this behavior could not be represented in a well-formed teleological representation as a goal-directed action, as there was an obviously more rational alternative means to the goal available, but not realized (the straight-line approach). Therefore, the infants could not generate any specific expectation about what type of goal-approach the small circle would follow in a changed situation. As a result, when the very same two test events that were shown in the experimental condition (see Figures 2C and 2D) were presented to the infants in the control study, the differential looking pattern found in the experimental condition has disappeared: the subjects looked equally at the old 'jumping approach' and the new 'straight-line approach' test events.

These results indicate that by 12 months infants can (a) interpret an other agent's action as goal-directed, (b) evaluate which one of the alternative actions available within the constraints of the situation is the most efficient means to the goal and (c) expect the agent to perform the most efficient means available in the given situation to realize the goal.

Above I argued that by applying the inferential principle of rational action one-year-olds can “go beyond the information perceptually given” and can productively infer any one of the three representational aspects of the teleological representation of a goal-directed action (Means act, Goal state, or relevant Situational constraints) if that aspect is perceptually inaccessible to them, as long as they have direct perceptual information about the contents of the other two representational elements. To demonstrate this property of inferential systematicity and productivity of the rationality principle, we habituated 12-month-old infants to computer-animated goal-directed actions in three types of situations (Gergely et al., 1995; Csibra et al., 1999; Csibra et al., 2003) (Figure 3). The different event displays were designed so that in each case one of the three basic elements necessary for a well-formed teleological action interpretation was made visually inaccessible. To interpret the action as an efficient and justifiable goal-approach, the infants had to use the rationality principle to infer and “fill in” the content of the relevant missing element of the representation.

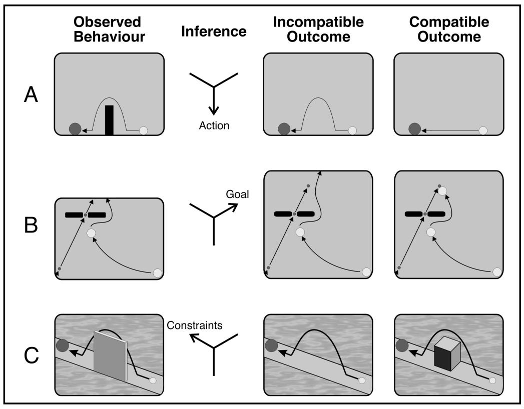

Figure 3: Three types of inference drawn from the teleological stance

Figure 3A (which depicts the Gergely et al. (1995) study that was discussed above) exemplifies the first type of teleological inference where infants have to infer the particular *means action* that is congruent with (i.e., can be seen as an efficient goal-approach in relation to) the visually specified *goal state* and *situational constraints*. As described above, the finding that infants looked significantly longer at the incongruent test display (old jumping approach) than at the congruent one (novel straight-line goal-approach) is evidence that they could draw the type of inference in question.

Figure 3B illustrates the second type of teleological inference where the infants had to infer a (non-visible) *goal state* to rationalize the incomplete action whose end state was occluded from them, as an efficient ‘chasing’ action (Csibra et al., 2003). During habituation a large ball was approaching a moving small ball until the latter passed through a small aperture between two obstacles and left the screen. The large ball, being too big to get through the aperture, had to make a detour around the obstacles before it also disappeared from view. In the two test events the upper part of the screen was opened up revealing one of two different end states: one congruent with the inferred *goal state* of an efficient ‘chasing’ action (the small circle stopped, at which point the large circle changed its course so that it ‘caught up with’ the small circle and contacted it), and one that was incongruent with the inferred goal (when the small circle stopped, the large one, without modifying its direction, passed by it leaving the screen without ever ‘catching’ the small circle). Twelve-month-olds looked significantly longer at the incongruent than at the congruent test display, suggesting that the incongruent outcome violated their expectation about the goal state that they had inferred to rationalize the incomplete action as an efficient ‘chasing’ event (for appropriate controls, see Csibra et al., 2003).

Finally, Figure 3C provides an example of the third kind of teleological inference to specify the particular *situational constraints* (occluded from view by a screen) in order to rationalize the small circle’s visible *action* (jumping approach) as an efficient means to realize the visible *goal state* (contacting the large circle) (Csibra et al., 2003). In the two test displays the screen was lifted either revealing an obstacle whose presence justified the jumping approach (congruent display) or revealing no such obstacle (incongruent display). Twelve-month-olds again looked significantly longer at the incongruent than at the congruent display, indicating that they inferred the presence of the occluded obstacle to justify the jumping approach as an efficient means to the goal (again, for appropriate controls see the original study reported in Csibra et al., 2003).

In sum, these results provide converging evidence indicating that by 12 months infants can take the *teleological stance* to interpret actions as means to goals, can evaluate the relative efficiency of means by applying the principle of rational action, and can generate systematic inferences to identify relevant aspects of the situation to justify the action as an efficient means even when these aspects are not directly visible to them.

5. Beyond the shortest pathway: The generality of the rationality principle.

At this point I anticipate strong resistance (see e.g., the controversy between Premack and Premack (1997) and Gergely and Csibra (1997)) and at least one specific objection against our – maybe, at first, radically sounding - theoretical proposal that one-year-olds, who may still lack the mentalistic competence to infer, represent and attribute intentional mental states (such as beliefs and desires) to other agents, nevertheless, already possess and apply inferentially productively *the general and abstract principle of rationality* that philosophers consider to be the central inferential component of mature theory of mind (Dennett, 1987; Fodor, 1987, 1992). The concrete objection is a straightforward one (communicated to me first by Paul Harris, pers. com.): looking at our habituation studies summarized above, one could justifiably suggest that in each case the action that, according to our theory, infants evaluate as the most rational means available to the goal, in fact, always coincides with the *shortest approach route* to the target object. It may be objected, therefore, that instead of relying on the general principle of rationality, infants may apply a *simpler and more specific criterion* of expecting the actor to always take the *shortest pathway* available to reach the target location.

In contrast, if infants, similarly to adults, employ the more general principle of rational action, they should be able to apply *other kinds of criteria* as well that could, under some circumstances, override the 'shortest pathway' criterion when interpreting behavior as an efficient goal-directed action. To demonstrate that this is, indeed, the case, we designed a study (Csibra et al., 1998) that pitted against each other 'shortest pathway' *versus* 'least effort' to see if the latter could also be applied as one of the criteria for evaluating the rationality of a goal-approach.

In a violation-of-expectation paradigm we habituated two groups of 12-month-olds to one or the other of two different versions of a 2D computer-animated event in which a rectangle approached a circle performing a worm-like motion pattern (see Figures 4A and 4B). The rectangle passed through a gap on a wall that separated it

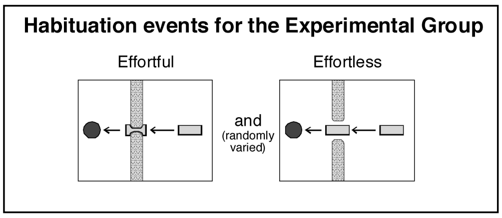

Figure 4A: The „Squeeze" study: Experimental group

from the stationary circle on the other side. The gap was either *wide*, allowing for an 'effortless goal-approach', or it was *narrow*, requiring the rectangle to squeeze through it exhibiting effortful movements ('effortful goal-approach'). The gap was positioned in such a way that passing through it corresponded to the *shortest, straight-line pathway* in between the rectangle and the circle. The two types of displays were randomly varied during habituation.

In the habituation events presented to the Experimental group (Figure 4A) the rectangle didn't have a choice of alternative routes to get to the circle on the other side of the wall: the only pathway it could take was through the one single gap on the wall. In contrast, in the habituation events presented to the Control group, the upper part of the wall was absent and so an alternative route to the goal

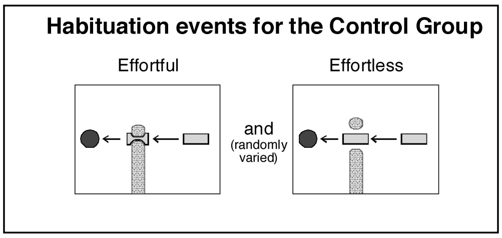

Figure 4B: The „Squeeze" study: Control group

(apart from the pathway through the gap) was also available which, though longer and involving a spatial detour, would not have made it necessary for the rectangle to engage in effortful squeezing (Figure 4B). In spite of this, similarly to the Experimental condition, the rectangle in the Control condition always took the 'shortest pathway' even when it had to squeeze through the narrow gap ('effortful goal-approach').

After habituation, the Experimental and Control groups were presented with the same two test events (Figure 5). The rectangle was again separated from the circle by a wall that had *two* gaps in it this time. The *narrow* gap, similarly to the habituation events, allowed for the shortest approach route to the goal, but it required effortful squeezing to get through ('shortest pathway/more effort'). In contrast, the position of the new (but *wider*) gap that was opened in the wall some distance below the

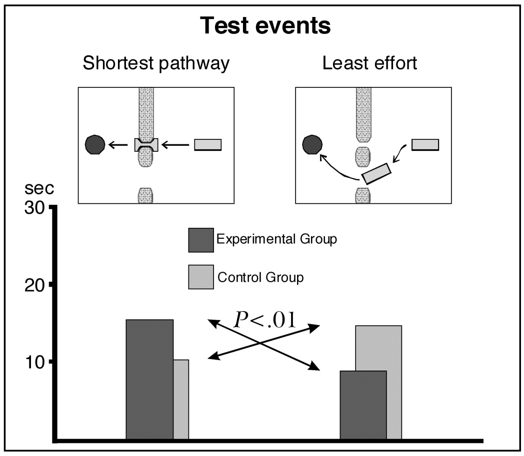

Figure 5: The "Squeeze" study: Looking times for the two types of test events in the Experimental *vs*. the Control conditions

narrow gap required a longer approach route involving a spatial detour, however, without making effortful squeezing necessary to get through it ('longer pathway/less effort'). Both test events were presented to each subject with their order of presentation being randomized across subjects. In the 'shortest pathway/more effort' test event the rectangle approached the target circle through the *narrow* gap, which allowed for the shortest approach route to the goal, but required effortful squeezing to get through it. In the ‘longer pathway/less effort’ test event the rectangle approached the target object through the *wider* gap that required a longer route through a spatial detour to get to the goal, but without the need to engage in effortful squeezing. The duration of the two test events was exactly the same (7.5 sec).

The results (see Figure 5) showed that subjects in the Experimental group looked longer at the ‘shortest pathway/more effort’ test event than the ‘longer pathway/less effort’ test event, while the Control group exhibited the *opposite* looking pattern. This crossover was significant as evidenced by a significant Condition X Test event interaction in a two-way ANOVA (F1, 38=8.35, P<.01) of looking times. Non-parametric statistics also confirmed this result (see Csibra et al., 1998).

We can, therefore, conclude that one-year-olds do not always expect an agent to approach its goal through the shortest path available. This suggests that the simpler ‘shortest pathway’ criterion is not a viable alternative to the more general rationality principle as the basis for judging what the most efficient goal-approach is within the constraints of a given situation. As the results for the Experimental group suggest, subjects expected the agent to take the *longer pathway that required less effort* when such an alternative to the goal became available during the test event. This indicates that one-year-olds are not restricted to the single criterion of expecting the ‘shortest pathway’ to the goal when evaluating the rationality of a goal-directed means action: in fact, they can clearly rely on other criteria as well (in particular, the criterion of ‘least effort’) when making such a judgment.

Finally, the results of the Control group, where already during the habituation event the agent did not follow the available alternative route to the goal that – at least, under some criteria such as ‘least effort’ – may have seemed more rational to the infant, seem to allow for two alternative interpretations. First, it is possible that the one-year-olds inferred and attributed a *specific disposition* to the agent (to always take the shortest path, or to squeeze whenever possible), and so they expected her to act according to this disposition even under the changed situational constraints of the test events. Second, it seems also possible that the infants reasoned that there may have been some further condition or aspect of the habituation situation (that they did not notice or were ignorant about) that must have justified the agent’s choice to take the shortest pathway even though it apparently required more effort. Therefore, on this ground they may have simply assumed the agent’s going through the shortest pathway must have been rational after all, and so they expected her to take the same approach route (that they have come to consider to be rational) in the changed situation of the test events as well.

To sum up: the results of the squeezing study clearly indicates that the principle of rationality that the one-year-olds rely on when evaluating the relative efficiency of alternative means to a goal is a general principle that allows for the application of multiple criteria and cannot be reduced to a single and more simple spatial criterion of always preferring the ‘shortest path’ to the goal.

6. Imitative learning of novel means in infants and robots: The problem of selective imitation.

Up till now I have only provided evidence for the teleological stance and its core inferential principle of rational action as a mechanism specialized for interpreting the goal-directed actions of other agents as those are perceived by the infant. Clearly, however, one of the most significant evolutionary advantage that the ability to interpret other’s actions as goal-directed provides for humans has to do with the vital functional role it plays in the *social transmission of culturally relevant new goals* and *new ways of acting to efficiently achieve such goals* from observing other agents’ novel intentional actions. The – possibly human-specific (Tomasello et al., 1993; Tomasello, 1999) and innate (Meltzoff and Moore, 1977, 1989) – ability to *imitate human actions* has been proposed by many as the basic mechanism that makes observational social learning of novel means from the action demonstrations of other human agents possible for our species. No wonder that one of the most cherished ambitions of epigenetic robotics has become to equip humanoid robots with the basic competence to imitate the observed actions of other agents (be it humans or other robots) and to use this ability to imitate for acquiring novel goal-directed actions (see Breazeal and Scassellati, 2002; Dautenhahn and Nehaniv, 2002; Schaal, 1999). Naturally, one of the major engineering hurdle towards achieving this aim was (and still is) to find efficient and generative solutions to the “correspondence problem” of mapping perceived movements of others onto the robot’s corresponding motor programs whose execution produces equivalent actions either by pre-wiring such a mapping or by designing learning solutions employing different sophisticated versions of supervised learning methods, forward modeling, “motor primitives” and connectionist learning nets (see Schaal, 1999, for a review). There has been clear progress in this area that was reinforced and informed by the recent discoveries of biological analogue mechanisms in the nervous system such as the mirror neurons (e.g., Fadiga et al., 1995; Rizolatti et al., 1996) and by psychological models such as Meltzoff and Moore’s (1997) ‘Active Intermodal Mapping’ mechanism to account for the phenomenon of neonatal imitation (Meltzoff and Moore, 1977, 1989).

However, it should be realized that no matter how far we advance in discovering and understanding the neural mechanisms that mediate the perceptual-motor mapping of actions or in finding engineering solutions to equip robots with analogous mapping mechanisms, such progress will at best provide us with some of the necessary, but never the sufficient conditions to fully understand or model the functionally more significant aspects of the human competence for imitative and observational learning. This is so because by exclusively relying on such automatic mechanisms that allow for action imitation, we would remain stuck forever with the pervasive problem of *how to avoid indiscriminate and automatic imitation* of anything (or at least any human or robot) that moves. Adults (and, as we shall see, even 14-month-old infants) are rather *selective in what human action they imitate and under what conditions they do so*. In fact, automatically imitating every human action that one is perceptually exposed to is a seriously dysfunctional pathological condition observable in patients with prefrontal lesions who cannot inhibit the tendency to compulsively imitate gestures or even complex actions performed in front of them by an experimenter (Lhermitte et al., 1986).

Clearly, what is needed is an additional account of *the inferential capacities that constrain and guide the imitative mechanism to be functionally selective*, a most significant problem that Breazeal and Scassellati (2002) have clearly put their fingers on when they raised as outstanding future problems for epigenetic robotics such questions as: "How does a robot know *what* to imitate?" or "Just as children develop the ability to imitate the goal of an action rather than a specific act, can we construct robots that are capable of making this inference? Today's robots respond only to the observable behavior without any understanding of the intent of an action." (p. 486).

In fact, until quite recently (see Bekkering et al., 2000; Gergely et al., 2001, 2002) *the problem of selective imitation* has not been fully recognized in developmental approaches to imitative learning either. Let us take as an example one of the ingenious and highly influential imitation studies by Meltzoff (1988, 1995) that has demonstrated that infants as young as 14 months of age can learn a novel means by imitation from observing an adult model's demonstration. The infants observed the model illuminate a box by leaning forward from waist and touching its top panel with her forehead. After a week, 67% of infants re-enacted this novel 'head-action', while no infant performed it spontaneously in a base-line control group that had not seen the action demonstrated. According to Meltzoff's (1995) own interpretation „infants do more than retrieve general goal or end state information („the panel can be lit"), *which would not necessarily mandate use of the head* [emphasis added]. They can remember the specific way something was done; they imitate the means used, not solely the general ends achieved." (p. 509).

Tomasello (1999) proposed that such *imitative learning* is human-specific as primates have been shown not to be able to imitatively copy specific novel means acts demonstrated to them. Instead, apes try to bring about the observed new outcome by performing motor actions already in their repertoire in a 'trial-and-error' manner (that actually often leads to the eventual (re)discovery of the demonstrated means or some other functional action with which they succeed in achieving the observed outcome). Tomasello has named this kind of observational learning *"emulation"* to distinguish it from true *"imitative learning"* that involves the faithful and automatic copying of the observed means. He argued that if infants used emulation in the Meltzoff study, one would have expected them to simply touch the box with their hand, instead of imitating the unfamiliar 'head-action'. Meltzoff (1988, 1995), however, did not report such 'hand-actions'.

Tomasello (1999) argued that „imitative learning of this type thus relies fundamentally on infants' tendency to *identify* with adults..." (p. 82) who are perceived by them (through Meltzoff and Moore's (1997) proposed innate mechanism of 'Active Intermodal Mapping) as "just-like-me". Tomasello further proposed that *identification* is a human-specific innate capacity that is lacking in apes as shown by the fact that apes can only emulate rather than being able to engage in imitative copying of observed means actions.

In sum: currently dominant models of imitative learning of novel means in developmental psychology (represented by the work of such influential researchers as Meltzoff (1988, 1995) or Tomasello (1999)) are characterized by two basic assumptions: 1. Re-enactment of novel means is due to an *automatic* tendency to *copy* the goal-directed action of a human model, and 2. This tendency is due to a human-specific drive for *identification* with human actors who are perceived through an innate perceptual-motor action mapping mechanism as *similar* ("just-like-me") by the infant.

As I suggested above, however, I think that this type of theory suffers from a serious shortcoming in so far as it cannot account for the functionally selective nature of human imitative learning that is arguably a significant adaptive feature of this possibly human-specific capacity. One piece of highly suggestive evidence indicating that young children's imitation of novel goal-directed actions does not necessarily and automatically involve the re-enactment of the specific means action modeled comes from a simple but clever set of studies designed by Harold Bekkering and Andi Wholschlager (e. g., Bekkering et al., 2000). In one condition they asked children between 3 and

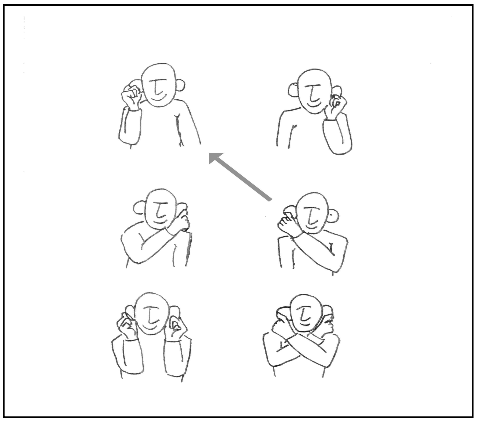

Figure 6: Ipsi-lateral error in imitation contra-lateral action in the Bekkering et al., (2000)

6 years of age to imitate an adult model's goal-directed target actions that involved touching either their left or their right ear with either an ipsi-lateral or a contra-lateral hand movement (see Figure 6). They found that while the children were practically errorless in reproducing the *goal* of the demonstrated actions (always touching their correct ear that corresponded to that of the adult's demonstration), they were, nevertheless, rather susceptible to commit a specific type of error: when contra-lateral hand actions were demonstrated to them (for example, when the adult touched his left ear with his right hand reaching across his body), they very often touched their corresponding ear (correct goal imitation) with an ipsi-lateral rather than a contra-lateral hand movement (failing to imitate the specific means action). In short, in their attempt to realize the same goal state as the adult, they tended to substitute for the modeled contra-lateral means action a simpler, more familiar, and thus more rational alternative means (the ipsi-lateral hand action) when such an alternative action was available to them.

7. Inferential constraints on observational learning in preverbal infants: Teleological emulation of new goals *versus* rational imitation of new means.

Let me also point out that the intriguing results of Meltzoff's (1988, 1995) "magic box" study, showing a rather automatic readiness on the part of 14-month-olds to faithfully imitate the unfamiliar 'head-action' to illuminate the box, is also hard to reconcile with the findings of the series violation-of-expectation studies (Csibra et al., 1999, 2003; Gergely et al., 1995; Gergely and Csibra, 2003) with 12-month-olds that I have reviewed above. This is so because those studies provided converging evidence that a) by taking the teleological stance one-year-olds – relying on the rationality principle - can evaluate which of the alternative actions available to an agent is the most efficient means to the goal within the constraints of the given situation and b) expect the agent to perform that particular means action that was judged to be the most rational alternative to realize its goal. Based on the teleological stance one would therefore predict that infants should re-enact the demonstrated action only if it seemed to them to be the most efficient alternative available to achieve the goal within the situation of the actor. One may then ask: why did Meltzoff's subjects re-enact the novel 'head-action' so faithfully, *when they could have simply touched the box with their hands*, an alternative action available to them that is simpler, more familiar, easier to perform, and so overall clearly a more rational means to the goal than the novel 'head-action'?

To solve this riddle, we speculated that Meltzoff's situation must have contained some situational features that actually allowed infants to 'rationalize' the 'head-action': in particular, we hypothesized that they may have noticed and interpreted the fact that *even though the model's hands were free*, she nevertheless did not use them, but touched the box with her forehead instead. This observation may have led the infants to hypothesize that there may be some aspect of the situation that they haven't noticed or are ignorant about, but due to which the novel 'head-action' must have some advantage in comparison to the – seemingly more rational – 'hand-action' in achieving the goal. Maybe then it was in order to figure out (and learn about) the nature of this assumed advantage, that, since their hands were also free (and so their situational constraints were identical to those of the adult), they decided to re-enact the novel 'head-action' themselves.[2]

To test this hypothesis, we replicated Meltzoff's study (Gergely et al., 2001, 2002) with one single modification using two conditions: in the *'Hands-occupied' condition* we changed the situational constraints of Meltzoff's original situation by arranging that *the model's hands were visibly occupied* when performing the 'head-action' (she, pretending to be chilly, wrapped a blanket around her shoulders holding it tightly with both hands, see Figure 7A). In contrast, the situational constraints remained the same as in Meltzoff's study in the '*Hands-free' condition* (where the model also pretended to be chilly and wrapped a blanket around her shoulders, but then put both of her hands on the table next to the box so that they were visibly free to be used, see Figure 7B).

Figure 7A: Hands-occupied

Figure 7B: Hands-free

In the 'Hands-free' condition (which, as pointed out in footnote 2 above, was structurally analogous to the Control condition of the "Squeeze" study), 69% of infants re-enacted the 'head-action', replicating Meltzoff's (1988) original finding. By contrast, in the 'Hands-occupied' condition, imitation of the novel

[2] Note the interesting analogy between this situation and the stimulus event of the Control condition of the "Squeeze" study discussed earlier (Figure 4B). In the latter the agent also had a potential choice of an apparently more efficient alternative route to the goal (involving a longer pathway which, however, required no effortful squeezing) that it, nevertheless, did not take, but approached the goal through the narrow gap that required more effort. In that case, during the test events (Figure 5) infants' relative looking times indicated that they formed an expectation that the agent would continue to perform the effortful squeezing action even when a different alternative route also requiring less effort has again become available to it. This finding seems analogous to the faithful imitation of the 'head-action' by Meltzoff's subjects in a situation where – just as their model – they also could have opted to touch the box with their free hand, but they still decided to re-enact the 'head-action' that originally must have appeared to them as the less rational alternative means.

‘head-action’ dropped significantly to only 21% ($p<.02$) (Figure 8). Thus, while it must have seemed sensible to the infants that the model whose hands were occupied performed the ‘head-action’ to illuminate the box (goal), 79% of the 14-month-olds decided *not* to imitate the ‘head-action’, because for *them*, whose hands were *free* (acting under *different situational constraints* than the model), the ‘head-action’ did *not* appear to be the most rational means available. In fact, all of these subjects illuminated the box by touching it with their hands, a non-imitative means action that was clearly the most rational alternative under *their* situational constraints.[3]

Finally, (and admittedly unexpectedly) we found that, whether the subjects re-enacted the ’head-action’ or not, *all* infants in *both* conditions performed the *’hand-action’* at least once (but often more than once: Mean= 2.1) within the 20 sec time-window of testing. This suggests that 14-month-olds are subject to an automatic emulation-like process whereby the memory of the effect (illumination-upon-contact) activates the response most strongly associated with establishing contact (hand-action). These emulative ‘hand-actions’, in fact, always preceded the imitative ‘head-action’ response

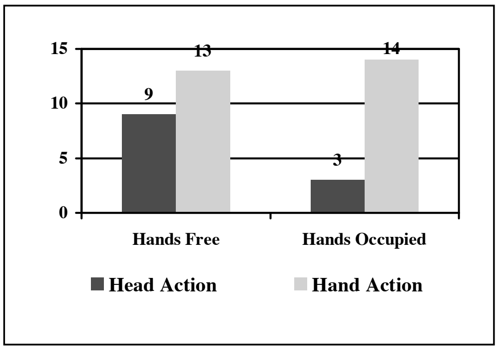

Figure 8: Number of subjects performing the ‘head-action’ *vs*. the ‘hand-action’ in the two conditions

[3] Note again the structural analogy, this time between the “Hands-occupied” condition and the Experimental condition (Figure 4A) of the “Squeeze” study. In the habituation event of the latter, the agent had no choice, but to squeeze through the only gap available on the wall to achieve its goal. This is analogous to the demonstrator’s situation in the “Hands-occupied” condition, who had no choice but to use her head to touch and illuminate the box. In the test events of the “Squeeze” study (Figure 5) when the situational constraints have changed and an apparently more rational alternative means action (requiring no effortful squeezing) became available, infants, as indicated by their relative looking times, formed an expectation for the agent to take the more rational alternative route to the goal that has become available. This is analogous to the finding that in the “Hands-occupied” condition infants chose *not* to imitate the model’s demonstrated ‘head-action’, but rather chose to emulate the goal in a rational manner by performing the ‘hand-action’ that in *their own* situation (whose situational constraints were different from that of the model) seemed the most rational means available to achieve the goal.

(where there was one) and were always successful in achieving the goal (illuminating the light-box). This makes it even more remarkable that the novel ‘head-action’ *was* imitated, even though only *selectively* (and therefore clearly not automatically) and only in the “Hands-free” condition. In that condition infants seemed to have interpreted the demonstrator’s choice to perform the ‘head-action’ rather than the also available – and, at least, apparently more rational – ‘hand action’, to indicate that there must have been some aspect of the situation (that the infants didn’t notice or understand) that, after all, justified the demonstrator’s choice of the ‘head-action’ as more rational in achieving the goal suggesting to the infants that the head-action’ must have some advantage over the ‘hand-action’ after all. It may be hypothesized that infants selectively imitated the ‘head-action’ in this condition driven by their ‘epistemic hunger’ to discover and learn about the nature of this advantage by comparing it to the alternative ‘hand-action’ (which they also performed).

In conclusion: these results strongly indicate that early imitation of goal-directed actions is *not an automatic response evoked by identification* with the observed agent, rather, it is the result of a *selective inferential process* guided and constrained by the evaluation of the rationality of alternative means available *both* in relation to the situational constraints of the model *and* in relation to the situational constraints of the infant herself.

8. Conclusions

In conclusion let us ask: What is there to be learned from the infancy studies and our theory of the teleological stance summarized above from the point of view of the specific concerns and aims of epigenetic robotics? In brief, I think that the main message is that in order to build an even remotely *‘socially relevant’ humanoid robot* it will not suffice to construct a machine that can produce actions, can perceive and imitate the actions of others, or can even learn to produce new actions from observing and imitating actions of other agents. To be able to equip robots with these capacities is, of course, a highly relevant (and obviously hard-won) achievement towards the realization of the ambitious aims of epigenetic robotics, but, in themselves, they amount to no more than fulfilling (some of) the necessary preconditions for constructing a ‘socially relevant’ humanoid robot. In order to even approximate the competence of preverbal human infants in the domain of action interpretation and production, epigenetic robotics will have to turn to the hard questions of how to construct mechanisms that implement *“top-down” constraints* that can make decisions guiding the action perception and production system about *what* action to produce and *when* (as well as what action *not* to generate under specific conditions), or *what* action to *imitate* to reach a goal, and *what* goal should be *emulated* rather that imitated under certain situational constraints.

I believe the most important lesson that can be derived from our research on the one-year-old's competence of action understanding is that such a "top-down" system should be conceived to be *inferential* in nature practically all the way down and even at the level when representing the causal mental states of other agents may not yet be present.

The typical first reaction of connectionist researchers in AI to our proposal that a teleological action interpretational system involving such abstract constructs as 'goals' and 'rationality' are present (and likely to be hard-wired) already in preverbal infants is to try to construct connectionist learning nets that, given certain input conditions, will be able to simulate the performance of our infants in our experimental situations, but will do so without building the abstract categories (such as goals and the principle of rationality or efficiency) into the connectionist net in any form. While I must admit that I am doubtful that such attempts would eventually succeed in eliminating the abstract representational concepts in question (certainly, the specific simulations proposed up till now did not manage to do so), I think running such simulations is certainly a worthwhile exercise as they will show us how far one can get with a purely "bottom-up" approach: an empirical issue one should not prejudge.

However, if the goal is to construct 'socially relevant' humanoid robots, I see no reason why researchers in robotics and AL should not pursue this goal also by designing forward engineering solutions (cf. Dennett, 1994) that would equip robots with *representational and inferential mechanisms* (and the relevant knowledge structures that these mechanisms could access) of the kind that are formalized in our teleological model and that could implement the "top-down" constraints necessary to guide the action production and perception systems to 'socially relevant' "choices" about when and what kind of action is adaptive to execute, imitate, or emulate. I can only hope that our experimental demonstrations of the *essentially inferential nature* of early action understanding and our formalization of the teleological interpretational system guiding such inferences may succeed in specifying useful directions for future research to be pursued in epigenetic robotics and AL providing "a little help to our friends" in these neighboring disciplines to realize their ambitious goal of creating 'socially relevant' humanoid robots that they have so bravely set for themselves.

References

Bekkering, H. et al. (2000). Imitation is goal-directed. *QuarterlyJournal of Experimental Psychology, 53A*, 153-64.

Breazeal, C. and Scassellati, B. (2002). Robots that imitate humans. *Trends in Cognitive Sciences, Vol.16*, No.11, 481-487.

Bruner, J. (1957). Beyond the information given. In Bruner, J. (1973) *Going Beyond the Information Given*, New York: Norton.

Carpenter, M.,et al. (1998). Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the Society for Research in Child Development, Serial No. 255, Vol. 63,* No. 4.

Csibra G., and Gergely, G. (1998). The teleological origins of mentalistic action explanations: A developmental hypothesis. *Developmental Science*, *1*:2, 255-259.

Csibra, G. et al. (1998). Beyond least effort: The principle of rationality in teleological interpretation of action in one-year-olds. Poster presented at the Swansong Conference of the Medical Research Council Cognitive Development Unit, London, September 1998.

Csibra G. et al., (1999). Goal-attribution without agency cues: The perception of 'pure reason' in infancy. *Cognition, 72*, 237-267.

Csibra, G. et al. (2003). One-year-old infants use teleological representations of actions productively. *Cognitive Science, vol. 27*(1),111-133.

Dautenhauhn, K. and Nehaniv, C. L. (Eds.) (2002). *Imitation in Animals and Artifacts*. Cambridge MA: MIT Press.

Dennett, D. C. (1987). *The Intentional Stance*. Cambridge MA: Bradford Books, MIT Press.

Dennett, D. C. (1994). Cognitive science as reverse engineering: Several meanings of „top-down" and „bottom-up" (pp. 249-259) In Dennett, D. C. *Brainchildren: Essays on Designing Minds*. Cambridge MA: Bradford Books, MIT Press.

Fadiga, L. et al. (1995). Motor facilitation during action observation: a magnetic stimulation study. *Journal of Neurophysiology, 73,* 2608-2611.

Fodor, J. (1987). *Psychosemantics*. Cambridge MA: MIT Press.

Fodor, J. (1992). A theory of the child's theory of mind. *Cognition 44*:283–296.

Gergely, G., & Csibra, G. (1997). Teleological reasoning in infancy: The infant's naive theory of rational action. A reply to Premack and Premack. *Cognition, 63,* 227-233.

Gergely, G., & Csibra, G. (2003). Teleological reasoning about actions: The naïve theory of rational action. *Trends in Cognitive Science*, July, (in press).

Gergely, G. et al. (1995). Taking the intentional stance at 12 months of age. *Cognition, 56*:2., 165-193.

Gergely, G.,et al. (2001). Rational imitation of goal-directed actions in 14-month-olds. (pp. 309-315), In J. D. Moore & K. Stenning (Eds.), *Proceedings of Cogsci 2001,* Edinburgh, August 1-5. LEA: London.

Gergely, G. et al. (2002). Rational imitation in preverbal infants. *Nature, Vol. 415*, p. 755.

Heider, F. & Simmel, S. (1944). An experimental study of apparent behavior. *American Journal of Psychology, 57,* 243-259.

Kelemen, D. (1999) Function, goals, and intention: children's teleological reasoning about objects. *Trends in Cognitive Sciences, 3, 461-468.*

Király, I. et al. (2003). Generality and perceptual constraints in understanding goal-directed actions in young infants. *Consciousness and Cognition*, (submitted).

Lhermitte, F. et al. (1986). Human autonomy and the frontal lobes: I. Imitation and utilization behavior: a neuropsychological study of 75 patients. *Annual Review of Neurology, 19,* 326-334.

Meltzoff, A.N. (1988) Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli. *Developmental Psychology, 24*, 470-476.

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

Meltzoff, A. N., and Moore, M. K. (1977). Imitation of facial and manual gestures by human neonates. *Science, 198*, 75-8.

Meltzoff, A. N., and Moore, M. K. (1989). Imitation in newborn infants: Exploring the range of gestures imitated and the underlying mechanisms. *Developmental Psychology, 25*, 954-62.

Meltzoff, A. N. and Moore, M. K. (1997). Explaining facial imitation: theoretical model. *Early Development and Parenting, 6*, 179-92.

Perrett, D. I. et al. (1985). Visual analysis of body movements by neurones in the temporal cortex of the macaque monkey: a preliminary report. *Behavioral and Brain Sciences, 16,*153-170.

Premack, D. and Premack, A. J. (1997). Motor competence as integral to attribution of goal. *Cognition*. *Vol 63*(2) May 1997, 235-242.

Rizzolatti, G. et al. (1988). Functional organization of inferior area 6 in the macaque monkey: II. Area F5 and the control of distal movements. *Experimental Brain Research, 71,* 491-507.

Rizzolatti, G. et al. (1996). Localization of grasp representations in humans by PET: 1. Observation versus execution. *Experimental Brain Research, 111,* 246-252.

Schaal, S. (1999). Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences, 3*:6, 233-242.

Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Cambridge, MA: Harvard University Press.

Tomasello et al. (1993). Cultural learning. *Behavioral and Brain Sciences, 16,* 495-552.

Wolpert, D. M. et al. (2001). Perspectives and problems in motor learning. *Trends in Cognitive Sciences, 5,* 487-494.

Woodward, A. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition, 69,* 1-34.

Woodward, A. L., and Sommerville, J. A. (2000). Twelve-month-old infants interpret action in context. *Psychological Science*, *11,* 73-77.