

**Birkbeck ePrints: an open access repository of the
research output of Birkbeck College**

<http://eprints.bbk.ac.uk>

Csibra, Gergely; Gergely, Gyorgy (2007).
'Obsessed with goals': Functions and mechanisms
of teleological interpretation of actions in humans.
Acta Psychologica, **124** (1): pp. 60-78.

This is an author-produced version of a paper published in *Acta Psychologica* (ISSN 0001-6918). This version has been peer-reviewed but does not include the final publisher proof corrections, published layout or pagination. The published version as cited above is available to subscribers at <http://dx.doi.org/10.1016/j.actpsy.2006.09.007>. Copyright © 2007 Elsevier Science BV.

All articles available through Birkbeck ePrints are protected by intellectual property law, including copyright law. Any use made of the contents should comply with the relevant law.

Citation for this version:

Csibra, Gergely; Gergely, Gyorgy (2007). 'Obsessed with goals': Functions and mechanisms of teleological interpretation of actions in humans. *London: Birkbeck ePrints*. Available at: <http://eprints.bbk.ac.uk/archive/00000472>

Citation for the publisher's version:

Csibra, Gergely; Gergely, Gyorgy (2007). 'Obsessed with goals': Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica*, **124** (1): pp. 60-78.

'Obsessed with goals':

Functions and mechanisms of teleological interpretation of actions in humans

Gergely Csibra

Birkbeck, University of London

György Gergely

Institute for Psychology, Budapest

Humans show a strong and early inclination to interpret observed behaviours of others as goal-directed actions. We identify two main epistemic functions that this 'teleological obsession' serves: on-line prediction and social learning. We show how teleological action interpretations can serve these functions by drawing on two kinds of inference ('action-to-goal' or 'goal-to-action'), and argue that both types of teleological inference constitute inverse problems that can only be solved by further assumptions. We pinpoint the assumptions that three currently proposed mechanisms of goal attribution (action-effect associations, simulation procedures, and teleological reasoning) imply, and contrast them with the functions they are supposed to fulfil. We argue that while action-effect associations and simulation procedures are generally well suited to serve on-line action monitoring and prediction, social learning of new means actions and artefact functions requires the inferential productivity of teleological reasoning.

Human beings cannot help but interpret the actions they observe in *functional* terms. In fact, the very notion of 'action' implies a motor behaviour performed by an agent, which is conceived in relation to the end state it is destined to achieve, even if this end state is not attained in every particular instance of an action. Plenty of evidence shows that even young human infants are inclined to interpret at least certain types of behaviours as goal-directed actions. However, there is no consensus about how they achieve this, and what factors influence their interpretation of others' observed behaviours as goal-directed actions (Csibra, 2003; Gergely, 2003; Hauf & Prinz, 2005; Woodward, Sommerville, & Guajardo, 2001).

In this paper, we take the 'functional stance' to interpret people's 'functional stance' towards actions. In other words, we ask what *functions* humans' obsession with goals serve, why we prefer to see others' behaviours as means to ends. Having outlined potential answers to the question of functions, we shall survey the *mechanisms* that various current theories propose for goal attribution, and consider them in relation to the functions we identified. We shall argue that these different mechanisms of goal attribution differ in how well they can fulfil one or the other of the functions in question. We close our arguments with speculations about the developmental significance of early goal attribution in human infants.

I The functions of goal attribution

Teleological interpretations of actions in terms of goals tell us *why* the action has been performed or, in other words, it provides a special type of explanation for the action. What is the benefit that we gain from such an explanation?

Keil (2006) lists several functions of explanations in general, and most of them are applicable to teleological explanations of actions as well. Attributing

a goal to an observed action allows us to predict its further course, to evaluate the action's causal efficacy in terms of its outcomes, and to justify the action. (Other explanatory reasons, like providing diagnosis and gaining aesthetic pleasure are perhaps less applicable to teleological interpretations of actions.¹) While predicting the future course of an on-going action can be clearly locally advantageous to the observer, the other two functions provide mostly long-term benefits that arise as a result of the specific way they construe and represent the observed actions. In particular, the evaluation of causal efficacy of an action allows us to learn about novel means to achieve desirable goal states, and justifying an action in terms of a goal enables us to learn about its function or the function of the tools employed by the action. Below we will discuss these different functions of goal attribution in turn.

1.1 Goal prediction and action anticipation

There are at least two ways in which goal attribution allows predicting future events. The more trivial one is due to the fact that a goal represents a state or an event subsequent to the action it belongs to. Thus, as long as goal attribution is made on-line, concurrent with action observation, goals refer to and predict future states or events. Since goals are normally

¹ The remaining function in Keil's (2006) list, 'affixing blame', is often behind goal attributions in everyday life. People tend to think that actors are more responsible for the intended than the unintended outcomes of their actions, i.e., for those that served as goals rather than arose as by-products of a particular action. Human infants (Behne, Carpenter, Call, & Tomasello, 2005) and chimpanzees (Call, Hare, Carpenter, & Tomasello, 2004) are also sensitive to this distinction and are likely to rely on it in their appraisal of social partners. Goal attribution may also support joint collaborative actions, which allow cooperation between individuals (Tomasello et al., 2005). As here we are focusing on epistemic functions, we are not going to discuss these 'social' functions of goal attribution in the present paper.

accomplished, goal attribution to not-yet-completed actions implies, by definition, a specific kind of prediction.² For example, if you observe someone's behaviour in the kitchen and interpret it as 'coffee making', you will by virtue of this interpretation predict the appearance of a cup (or more cups) of coffee in the kitchen in the near future. Inferring goals from actions inherently involves a reference to some future state.

There is, however, a less trivial way as well by which goals help predicting future events. While goal attribution itself allows us to jump ahead in time and predict a hypothesized future state, it also enables us to fill up the intervening time by *action anticipation*. For example, interpreting an action as 'coffee making' would allow you to anticipate that the actor will take water from a source (tap or bottle), will apply heat to the water (in a kettle or on the hob), will put coffee into a container (cup, filter, or coffee machine), etc. The exact details of these predictions will depend on environmental circumstances, your knowledge about artefacts and procedures, about the actor's habits, abilities and preferences, and further (for example, verbal) information you may obtain in the particular situation. Clearly, however, the content of your action predictions will be most crucially determined by the specific goal implied by your interpretation of the observed action as an instance of 'coffee making'. As the example suggests, action predictions from goals can be made by breaking down the path towards the goal into sub-goals in a hierarchical fashion, eventually arriving at elementary motor acts, like grasping or stepping, at the finest resolution.

The benefit of generating action anticipation at such a fine-grained level is that it enables us to track and monitor others' actions, which in turn enables revising our action interpretation if our predictions are not confirmed, or intervening in relevant ways if we feel it necessary. If you anticipate the actor putting instant coffee into the cup, but you observe that she is reaching towards the box with tea bags in it, you will change your original goal attribution from 'coffee making' to 'tea making', or you can notify her that you do not like tea (assuming that you think that the actor's ultimate goal is to offer you a cup of hot beverage), or you can inform her about the actual contents of the box (tea bags) in case you suspect that she is acting on a false belief. Such on-line tracking and anticipating of ongoing events is not restricted to action interpretation, but is characteristic to all domains of dynamic event perception (Wilson & Knoblich, 2005).

The crucial difference between the two kinds of prediction that teleological action interpretation permits is in the *direction of inference* they involve (Table 1). Predicting a future goal state from

interpreting an ongoing action represents an 'action-to-goal' inference ('What is the function of this action?'), while during action anticipation the inference goes in the opposite direction from 'goal-to-action' ('What action would achieve that goal?'). As we have seen, action perception in real life involves both kinds of inference and both kinds of prediction intertwined, as we constantly revise goals and sub-goals attributed to actors on the basis of verified or disconfirmed action predictions. Nevertheless, it is important to maintain the distinction between these two kinds of inferences because, as we shall see, they require different types of computations.

Primary Function	Type of Inference	
	'Action-to-Goal'	'Goal-to-Action'
On-line Prediction	Goal prediction: Predicting the likely effect of an ongoing action	Action anticipation: Predictive tracking of dynamic actions in real time
Social Learning	Discovering novel goals and artefact functions	Acquiring novel means actions by evaluating their causal efficacy in bringing about the goal

Table 1. The functions of teleological interpretation of actions

1.2 Social learning of novel means actions and novel goals

When we interpret an action as goal-directed, we conceive of it as playing a causal role in bringing about its goal state, rather than just being coincidental with it. If the action, or some aspect of it, is novel to us, relating it to the goal will represent the action as a novel means that can be successfully used to achieve the goal state. Learning about a new action's causal potential to bring about a goal state is especially useful when we have not had a means to that end in our own action repertoire, or when the new action represents a causally more efficient way to achieve the goal than the means action we have formerly had access to. Therefore, interpreting novel actions in terms of goals can allow us to learn useful new means by observing others use them successfully to achieve their goals.

For example, assume you see someone perform a novel action sequence in the kitchen that involves operating a machine you have never seen before. You may be able to attribute to this activity the goal of 'coffee making' even if you are unable to extract and understand the relevant causal factors that makes the action sequence a successful procedure for achieving this goal. This can happen, for example, if the actor informs you about her goal before or while engaging in the particular action sequence, and/or if you see that the observed action reliably succeeds in realizing the

² Sometimes we predict not accomplishment but failure for observed actions. However, even when this is the case, we make our predictions with reference to the goal (the intended outcome) of the action.

goal state (as you soon end up with a cup of coffee in front of you as a result). Thus, assigning the goal of 'coffee making' to a new action sequence involving the use of a new artefact and seeing that the goal is successfully achieved by it allows you to learn both that the machine you have never seen before can be used 'for' making coffee, and that the particular way you observe the other operate the machine is a new means action that can be used to achieve the goal of coffee making.

In addition, observing successful goal-directed actions can also inform us about the *causally relevant properties* of actions, objects, and situations that enable goal attainment. For example, observing how the machine is operated and registering the manner in which the action sequence is performed can highlight and make the causally relevant information that leads to goal achievement more readily accessible for you to extract and represent.

Social learning of actions to achieve desirable ends is not restricted to humans. In fact, much of the literature on animal social learning focuses on transmission of innovative behaviour among individuals by observational means (Heyes & Galef, 1996). The vast majority of the reported instances of social learning via observation in non-human animal species represent behaviours that lead to instantaneous rewards (usually food) to the actor. According to some theorists, local adaptivity of the observed behaviour is indeed a precondition for social learning to occur in non-human animals (Galef, 1995; Heyes, 1993). Thus, individuals of various species seem to acquire means actions from each other as long as the end result of the action is familiar and desirable for them. Note that, similarly to humans, non-human animals may also gain knowledge about the causally relevant aspects of an observed action in relation to the goal state. They may rely on such knowledge when producing their own way to achieve the same outcome (i.e., when they 'emulate', see Tomasello, 1996). Alternatively, if the observed action does not allow them to extract such relevant causal knowledge, they may opt to blindly 'copy' the actual action sequence as faithfully as they can in order to secure the attainment of the same desired outcome (i.e., they 'imitate', see Horner & Whiten, 2005).³

Social learning of new means (as discussed so far) allows us to figure out what action is to be performed (or what tool is to be used) to achieve a certain goal. Thus, this kind of learning matches actions to goals, means to ends, artefacts to functions. Sometimes, however, the new information that we derive from the teleological interpretation of an observed action is not about new means actions but about new goals the actions can be used to achieve. For example, you see your friend pouring the freshly brewed coffee from one cup into another, then she washes up the first

cup, pours the coffee back into it, then she washes up the second cup, and repeats this sequence several times. None of these actions are novel to you; yet, the whole sequence does not make sense at first sight. Then you notice that by the end of this action sequence your cup of coffee that was too hot to drink to start with becomes cooler and ready to be consumed. This leads you realize that your friend was probably in a hurry and wanted to quickly cool down your coffee to a consumable temperature without diluting it. In other words, since this desirable outcome justifies the action sequence that brought it about so successfully, it will be attributed as its goal leading you to learn a new goal (as well as identifying the action sequence as a means to achieve this goal). Again, your learning of the new goal may be based just on observing its successful attainment by the perceived action sequence or you may be directly informed about it by the actor. Of course, you can also base your inference that identifies the new goal on your already existing causal knowledge that washing the cups in cold water cools them and allows them to quickly take up heat from the drink. In either case, by teleologically interpreting an observed event you may learn about a new goal (e.g., that one can quickly cool down hot coffee), while possibly also acquiring a new means action to achieve it. Observing such an event can also highlight and thereby facilitate learning about the causal properties of the objects used (cups, in this case) that are relevant to the achievement of the new goal and that your knowledge representation of the range of affordance properties of this particular artefact has not contained before.

In addition, if the action is performed by a novel tool, the goal attributed to the action can be attached to the tool as its function. Unlike the goal of an action, which is local to a particular incidence, functions 'stick' to tools even when they are not in use. One way to characterize the function of an artefact is that it is a 'frozen' goal, for which it can be used as a tool by a means action under a large variety of situational and person constraints. Consequently, one can infer artefact function from goal attribution, and such an inference constitutes social learning of artefact function (e.g., Casler and Kelemen, 2005). Suppose that, instead of pouring the drink from cup to cup, you see you friend discharge the coffee into some equipment that circulates it under running cold water. If you manage to figure out that the goal of this action is to cool the coffee quickly, you have also come up with a plausible hypothesis concerning the function of the circulating equipment. Thus, attributing goals to unfamiliar actions could assist social learning of artefact functions.

Unlike social learning of means actions, social learning of new goals and artefact functions is extremely rare in non-human animals: in fact, this kind of social learning may turn out to be a human-specific phenomenon (cf. Csibra & Gergely, 2006; Gergely & Csibra, in press). The reasons behind this are twofold.

³ For a more detailed discussion of this distinction, see Gergely & Csibra (in press).

First, although non-human animals, and especially non-human primates, are said to develop cultures that are qualitatively not different from those of humans (Laland & Hoppitt, 2003; Whiten, Horner, & Marshall-Pescini, 2003), they certainly lack the rich variety of artefacts that are a characteristic feature of all human cultures. In fact, the repertoire of objects that are used as tools by non-human species seems severely limited (similarly to the range and types of goals they pursue, see below). Thus, the need for specialized mechanisms for learning artefact function does not seem to arise in other species. Additionally, despite the ample evidence for tool use in primates (Boesch & Boesch, 1993) and birds (Weir, Chappel, & Kacelnik, 2002), there is no indication that individuals of any of these species would conceptualize the objects they use as tools as possessing permanent functions that would be represented as their essential property. Unlike humans, other animals do not routinely store their 'tools' for later use (but see Mulcahy & Call, 2006), neither do they invent, design, manufacture, and maintain tools, and they rarely transport them for long distances (see also Csibra & Gergely, 2006; Gergely & Csibra, in press). Second, and relatedly, the range of goals that non-human animals pursue differs enormously from that of humans. Because of the human capacity for recursive tool use (using tools to make other tools) and the human inclination to represent multi-level hierarchies of goals, humans pursue a wide variety of goals that very often do not serve their immediate adaptive needs, but function as local sub-goals to some distant end. For a human observer, this multitude of different types of goals that are nevertheless useful to acquire simply because they may function as valuable sub-goals for the observer's own goal-directed activities in the future. As non-human animals are normally not engaged in long sequences of novel actions that have their useful effect only in the distant future, but pursue a highly restricted – and mostly innately specified – set of local goals that are familiar to them and their conspecifics (food, mating partner, shelter, territorial protection etc.), they do not learn novel goals by observation.

2 Goal attribution and action prediction as inverse problems

So far we have concentrated on the functions of goal attribution (the *why* question), but we have systematically treated the issue of *how* goals are extracted from actions by observers. Before we turn to the candidate mechanisms of goal attribution, we have to discuss the nature of the task that these mechanisms are required to solve.

As we have seen, both the predictive and the social learning functions of goal attribution cover two kinds of inferences (Table 1). In computational terms, either the percept of an action provides the input of the system, activating an appropriate mechanism to identify and output the goal to be attributed to the actor (as in goal prediction and social learning of

functions), or it is the goal that serves as the input to the system for which the mechanism is expected to produce a (predicted or learnt) action as the output. The complexity of both types of computation depends on the range of possible outputs that a system should be able to choose from.

In case of a 'goal-to-action' type of computation (for example, action prediction), the observer is confronted with the problem of sorting and choosing from the myriad of potential actions that the actor may perform. Even if one considers only those actions that would eventually produce the attributed goal, the possibilities are countless. Coffee can be made many different ways, and each sub-goal towards the end could also be reached by countless routes. Intuitively we all know that many of these courses of action are rather unlikely and should not be considered seriously (cf. the 'frame problem', Dennett, 1978; Fodor, 1983). Your friend could start her task, for example, by designing and building a new coffee machine, or by planting a coffee plant in the garden, or by washing the coffee beans in the sink, but these actions would seem implausible for an observer. Notice that this assessment of relative plausibility is not the input, but the output of our action prediction system, which must have already judged these actions as improbable. However, if the sole input to the system is the goal to be achieved, these actions, which would eventually lead to goal achievement, cannot be a priori excluded from consideration.

Because the end of an action sequence does not in itself determine the details of the action unambiguously, inferring actions from goals is an *inverse problem*. Inverse problems are ill-defined problems that cannot be solved by analytic methods because the available information does not constrain the solution sufficiently. Inverse problems are well known in many disciplines like physics and neuroimaging (Schmidt, George, & Wood, 1999), and many tasks that the animal visual system has to solve can also be conceptualized as inverse problems (Scholl, 2005). For example, infinitely many different arrangements of objects in the 3D space could produce the same 2D projection on the retina, and the visual system is confronted with the job to find the most likely interpretation of this 2D image in the 3D space (Marr, 1982). Similarly, infinitely many different action sequences could produce the same end state, and when one wants to predict actions knowing the end result (the goal), her task is to predict the most likely course of action from among the many possible ones that could lead to the same end.

What about the other type of inference that we use in goal attribution? Is it also an inverse problem to figure out the goal of an action either online (to predict the end state) or a posteriori (to learn goals and artefact functions)? The answer again depends on the range of potential goals that the observer has to be able to consider when making this kind of inference. For creatures that are interested only in

goal states that carry immediate adaptive value, and that can be easily discriminated from each other, this range is small. When a non-human animal has observed an action, it only has to check its end result against a short list of known rewards (has it resulted in food?, has it provided shelter?, etc.), and a match will provide it with the goal to be attributed. If there is no match, the action may not be interpreted as goal-directed at all. Humans, however, are obsessed with goals and can assign a goal (by coming up with a functional explanation) for practically any action. As we discussed in the previous section, the likely function of this species-specific 'teleological obsession' is the need to learn a large variety of novel goals and artefact functions from others.

In humans, therefore, where almost anything can be a goal state, 'action-to-goal' inference also becomes an inverse problem. If you observe your friend in the kitchen pouring water from a bottle to a cup, what is the goal? To fill the cup? To empty the bottle? To cool (or warm up) the cup? To measure the volume of water? To store water? To play with wet things? Or is pouring water into the cup a sub-goal to enable one to drink water from the cup (a further goal higher up in the goal-hierarchy)? Our knowledge about people, kitchens, cups, and water, and further information about the actual situation will help us to find an answer, but perhaps nothing in the pouring action itself can serve as the basis for choosing among these possibilities. Just like a certain arrangement of edges on the retina can be projected by many arrangements of objects in the world, an action can be performed to achieve many different kinds of goals. The fact that our goal attribution system could answer this question easily in most situations is not an indication of the simplicity of the problem, but evidence for a very efficient goal extraction mechanism (cf. Baker, Tenenbaum, & Saxe, in press).

Inverse problems are unsolvable by analytical methods, but this does not entail that there is no feasible way to find the most likely solution for such a problem. To do this, one has to rely on further assumptions about the nature of the phenomena – assumptions that will constrain the number of potential solutions to be considered. The human visual system, for example, implicitly assumes that objects are illuminated from above and, if no other information resolves the ambiguities concerning the convexity of a surface, assigns depth values to points on the surface accordingly (Ramachandran, 1988)⁴. It is important to note that the assumptions that are used in solving inverse problems *are* assumptions; they do not have to be, and are not always, true. The validity of a solution to an inverse problem crucially depends on the probability of the assumptions that have been drawn upon being true. Since natural lighting on the Earth is, indeed, coming from above, the assumption

of 'illumination from above' is likely to lead to valid solutions, unless artificial lights have been applied in the situation. Also, the applicability of certain assumptions may depend on the actual circumstances, for example on the availability of some information in the input. The assumption on illumination direction is useless in an environment where only directionless ambient light is present.

3 Mechanisms of goal attribution

What cognitive mechanisms are available for humans for goal attribution and for utilizing functional understanding of behaviours for social learning? In the recent years, three distinct mechanisms have been proposed for this purpose: action-effect associations, simulation procedures, and teleological reasoning. Evidence has been published to support the functioning of all these mechanisms in early infancy. We shall briefly describe these mechanisms and their supporting evidence in turn, paying special attention to the question of what assumptions they make to solve the inverse problems discussed in the previous section.

3.1 Action-effect associations

One theoretical view on action planning, summarized by the ideomotor principle (James, 1890), emphasizes the role of goal representation in the generation of motor actions (Hommel, Müsseler, Aschersleben, & Prinz, 2001). According to this view, the representation of actions in the actor's cognitive system includes and is related to the representation of their desired distal effects, and these inherently related action-effect representations are linked to each other through bidirectional associations. Thus, the 'idea' of the goal state, i.e., the desired effect, automatically activates the corresponding action, while the activation of an action elicits the anticipation of the distal effect associated with it. These links are established by simple associations upon observing the effects that one's actions have produced, and these associations start to build up from early on in infancy.

An interesting extension of this approach of action control is the idea that perceiving and understanding other people's actions are also based on similar bidirectional associations between actions and their effects (e.g., Elsner, *present issue*). Thus, observed actions are proposed to be represented by being linked to the effects they have been seen to bring about (rather than, for example, just in terms of their antecedent conditions). As the effect of an action normally (though not necessarily) corresponds to the intended goal, this representation is akin to a kind of goal attribution. Moreover, such a stored action-effect association can be used in later occasions to predict the goal of an ongoing action, or to anticipate the behaviour of an observed individual with a known goal.

Human infants indeed tend to attend to the effects of actions they observe from very early on and expect the actor to produce the same effect again (e.g., Leslie,

⁴ Interestingly, even newborns are sensitive to the direction of illumination, and prefer to look at faces lit from above (Farroni et al., 2005).

1984). Woodward (1998) demonstrated that when 6-month-old infants see a hand repeatedly grasping one of two objects, they anticipate that the same object would be grasped again even when the spatial location of the objects are rearranged. It has also been shown that infants do not consider any end state as the effect of the action. For example, if the hand, instead of grasping it, just touches the object in an unfamiliar way, it does *not* generate expectation of the same ending in a new spatial arrangement (Woodward, 1999), but if the action generates a new and salient effect, infants anticipate the same effect in a new situation (Bíró & Leslie, *in press*, Király et al., 2003). Older infants have also been shown to associate actions and effects when the effects involved novel, unfamiliar artefacts (Elsner & Aschersleben, 2003).

When goals are attributed, or actions are predicted, on the basis of observed action-effect associations, the observer implicitly relies on the assumption that an action is directed towards the same goal state that has been produced by earlier, similar actions, and that the a goal state (an effect) will be achieved by an actor in a similar way, as it was achieved before. To the degree that many of the goal-directed actions that people perform are routine acts that we execute the same way every time, these assumptions will likely to lead to valid goal attributions and action predictions. The applicability of the action-effect associations for teleological interpretation of actions, however, is limited by several factors. First, its implicit assumptions do not offer a solution for either type of inverse problem if the observed action is novel, or if the situation does not afford the actor to reach its goal through the familiar action that has been associated with the goal state. Second, while an action can be associated with a number of different distal effects it has been observed to produce in the past, there is no further selection mechanism to identify the goal of the action in a particular situation. For example, if running fast makes one sweat regularly, when seeing someone running towards the bus-stop behind the corner, how could one decide to assign to the observed running action the goal of 'catching the bus', instead of 'producing sweat'?

3.2 Simulation procedures

According to an influential theory of mental state attribution, people understand each other's mind by imagining themselves in the other's position, and simulatively generating the mental states (beliefs, desires, intentions) that they would possess were they in the other's 'shoes' (Goldman, *in press*; Gordon, 1986). Beyond understanding the mind states of others, simulation procedures can also be utilized to generate predictions about their behaviour by asking, "What would I do in her situation?" Simulation theories of mindreading were recently given a new empirical boost by the discovery of mirror neurons in monkeys and mirror neuron areas in humans (for a review, see Rizzolatti & Craighero, 2004), which are

apparently similarly activated by executing and observing a particular action.

Simulation procedures would obviously be very useful for understanding goal-directed actions. A predictive simulation procedure takes the goal ascribed to the actor as its input and uses the observer's own decision system to generate (but not execute) the representation of the motor action that she herself would perform to achieve the goal in question. This simulated goal-directed action will then be predictively attributed to the actor (Gallese & Goldman, 1998). A 'retrodictive' (or 'postdictive') simulation procedure could be applied in the opposite way to recover the likely desire that generated an observed action ("What desire would make me do this?"), and the content of that desire would then be attributed as the 'goal' of the observed action (cf. Gergely & Csibra, 2003). The advantage of relying on such simulation procedures rather than invoking tacit laws, principles or models for action understanding is that the observer can exploit his/her own existing mental mechanisms that function to link his/her own goals and actions.

Several theoretical accounts link the early development of goal attribution to infants' own abilities to perform means actions. For example, Tomasello (1999) explains the sudden emergence of certain social cognitive skills around 9 months of age (the 'nine-month revolution') by the simultaneously occurring emergence of infants' understanding of their own selves and their own unfolding capacity for means-end actions, and by using this self-knowledge to interpret others' actions through simulation. Likewise, Meltzoff (2002) argues that ascribing goals to others' actions, and predicting their behaviour, is based on a projection of their own experiences to the observed actor, and the slow development of goal attribution in infancy is explained by the gradual nature of motor skill development. Indeed, Sommerville and Woodward (2005) have found that infants are more likely to attribute a goal to a means action that they are capable of performing than to actions that are not yet in their own action repertoire. Moreover, if they are provided with the opportunity to gain experience with successful goal attainment, even very young infants tend to use this new skill for interpreting similar observed actions of others (Sommerville, Woodward & Needham, 2005).

Simulation procedures tackle the inverse problems of teleological action interpretation by reducing the possible range of solutions through relying on the 'equivalence' assumption that the observed actor has the same motor constraints and preferences as the observer (the "like me" hypothesis, Meltzoff, 2005, *present issue*). Since this assumption is to a large extent true of conspecifics, an effective simulation may lead to a valid solution on many occasions. These procedures, however, cannot be used when this assumption is obviously not met (e.g., for a non-human actor), or when the observed action is completely

novel that the observer has no first-person experience with, or could not herself perform (due, say, to the immaturity of an infant's motor system, or to individual motor deficits).

3.3 Teleological reasoning

An alternative theoretical approach to action understanding emphasizes that, in order to interpret an action as goal-directed, one has to take into account the relevant constraints of the situation in which the action is performed. The outcome (the effect) of an action may, or may not, be seen as the goal, depending on whether the outcome is judged to *justify* the action in the given situation. Such normative evaluation of actions is based on the *principle of rational action* (Csibra & Gergely, 1998; Gergely & Csibra, 2003), which allows for the assessment of the relative efficiency of the action performed to achieve the goal within the situational constraints given.

Beyond contributing to the evaluation of whether some behaviour is a well-formed goal-directed action in relation to a certain end state, this principle can also be used productively (Csibra, Bíró, Koós, & Gergely, 2003). Thus, when observing an ongoing action, one can infer the likely goal of the action by assessing what end state would be efficiently brought about by the action given the particular situational constraints. This end state then provides the goal to be attributed to the actor. Conversely, if the goal is known, the efficiency principle helps to generate an action prediction by inferring what the most efficient course of action towards the goal state would be in the given situation. Teleological reasoning is a very flexible tool in action understanding, but it requires the recruitment of the relevant background knowledge that the observer accumulated about the physical constraints of the situation and of the actor.

We and others have collected ample evidence that infants do engage in teleological reasoning when they observe actions. They evaluate the efficiency of goal approach, and predict new actions in new situations accordingly (Gergely, Nádasdy, Csibra, Bíró, 1995). They apply the principle of rational action to the interpretation of behaviours of people (Sodian, Schoeppner, & Metz, 2004; Kamewari, Kato, Kanda, Ishiguro, & Hiraki, 2005), hands (Woodward & Sommerville, 2000; Phillips & Wellman, 2005), puppets (Sodian, Schoeppner & Metz, 2004), robots (Kamewari et al., 2005), geometric shapes (Gergely et al., 1995; Csibra et al., 2003; Wagner & Carey, 2005), even when these actors do not exhibit unambiguous agency or animacy cues (Csibra et al., 1999). In addition, infants take into account the efficiency of observed new actions when they decide what aspects of them they should imitate (Gergely et al., 2002).

The basic assumption that teleological reasoning applies for solving the inverse problems of goal attribution is the efficiency principle itself. Teleological reasoning will lead to legitimate conclusions only if the observed actor's behaviour approximates the ideal of

efficiency. Since biological systems tend conserve energy, this assumption is likely to produce valid predictions and goal attributions, and has been shown to be a computationally viable way of teleological action understanding (Baker, Tenenbaum, & Saxe, in press). However, insufficient knowledge about the constraints of the actor or the situation may produce wrong predictions or goal attribution by teleological reasoning.

4 Matching functions and mechanisms

As we have seen, all three proposed mechanisms of teleological interpretation of actions apply assumptions that are likely to be true in real life situations where the problem of goal attribution arises. If you see your friend loading coffee into the coffee machine, you may confidently guess that her goal is 'making coffee', because (1) this is the usual effect associated with this action, (2) you would typically do this when you wanted to make coffee, and (3) this end justifies the action, considering what you know about coffee and coffee machines. However, all three methods carry limitations as well. If you (1) have not yet seen anyone making coffee, or (2) have never made coffee yourself, or (3) have no idea about the function and relevant causal properties of the unfamiliar machine your friend is operating, then applying these assumptions would not help in resolving the inverse problem of goal attribution. Because of these limitations of applicability, we could expect that the goal attribution mechanisms listed above will have different strengths and weaknesses when applied to fulfil the various functions of teleological interpretation of actions.

4.1 Goal prediction and action anticipation

As the example above suggests, all three mechanisms are capable of finding a known goal for a familiar action. Teleological reasoning can search through the possible end states that an action could potentially lead to until it finds one that would be reasonably worth the effort represented by the action. An iterative retrodictive simulation procedure can also be applied for this purpose: potential goal conjectures can be fed into the observer's own action planning system, and if the output matches the perceived action, the likely goal of the actor has been found (Gallese & Goldman, 1998). However, the fastest response in this situation would be generated by recalling the effect that is most strongly associated with the perceived action, as it would not require searching many potential outcomes. Indeed, it may well be that the different mechanisms of goal attribution are all involved when one tries to find a goal for a familiar action. While the stored action-effect associations would come up with quick hypotheses about the goal, simulation procedures and teleological reasoning could verify these solutions by applying their own criteria.

Goal attribution assists action anticipation by allowing the observer to break down the path leading

to the goal state to sub-goals and anticipating their occurrence in turn. Unless the action sequence is a well-known routine that is always performed the same way, action-effect associations are of only limited use in this type of prediction because they do not take into account situational constraints. Even if a sub-goal has been identified (say, pouring water into the coffee machine), the action-effect mechanism will likely predict the action that is most strongly associated with that sub-goal (say, taking water from the tap), even when that action is unavailable in the given situation (say, there is no water supply). Action-effect associations cannot come up with novel actions (pouring water into the coffee machine from a bottle of mineral water), because this unfamiliar action has not been associated during prior experience with the desired outcome. The assumption that action-effect associations provide the basic mechanism for identifying goal-directed actions leads to significant limitations when it is used for action anticipation and seems to fail to capture the flexibility that characterizes human infants' early competence in this regard.

In contrast, simulation procedures work well when one has to generate action predictions for a known goal state. Considering the question, "What would I do in that situation to achieve my goal?" can generate novel predictions, and can flexibly adapt to the constraints of the actual situation. As long as the observed individual operates with a similar knowledge background and motor constraints, simulation is likely to produce valid action predictions. In fact, we hypothesize that action simulation and action mirroring evolved exactly for the purpose of action monitoring and flexible action anticipation for goal-directed actions of conspecifics (cf. Csibra, 2005).

Teleological reasoning can also provide novel action predictions, with the extra benefit of being able to do so for non-human actors as well. However, within the more restricted task-domain of predicting the actions of a conspecific individual, teleological reasoning can hardly be discriminated from simulation procedures, because in this case simulation is clearly a natural and effective way to find the most efficient action towards a goal state. (After all, the observer herself is also a biological agent, who is similar to the observed conspecific in that it normally performs the most efficient action available to her to achieve a goal within the constraints of the situation.)

4.2 Social learning

Learning a new action for a certain goal state requires two skills: finding the relevant element of the observed action sequence that has played a causal role in bringing about the desired effect, and storing the link between this action and the outcome for later use. The second task is rather trivial and accomplished easily by an action-effect association mechanism. Simulation procedures have also been suggested to play an important role in reproduction of actions

(Iacoboni, 2005; Meltzoff, 2005). For example, one function proposed for mirror neurons in humans is that they help us break down observed action sequences into elementary motor acts and reassemble them into novel actions (Rizzolatti & Craighero, 2004).

The first task in acquiring a means action by observation, however, is not trivial. Researchers who want to develop robots that imitate are routinely confronted with the problem of how to decide which aspect of the observed behaviour is to be reproduced (Breazeal & Scassellati, 2002; Gergely, 2003). Is the fact that your friend leans forward when she puts the coffee into the cup relevant? Do you have to reproduce her pleasurable expression upon smelling the coffee beans or the melody she hums while making coffee? Many aspects of observed novel behaviours are not relevant part of the actual goal-directed action but are simply superfluous. Neither stored action-effect associations nor simulation procedures are suitable to sort these elements apart, because neither of these mechanisms can be invoked to interpret unfamiliar actions. What is required for this task is a causal analysis of the affordance structure of the observed actions and the artefacts they involve in order to recover which elements of those actions (and artefact use) are necessary and sufficient for producing the desired effect. Teleological reasoning could provide exactly this kind of analysis.⁵ There is some evidence to suggest that chimpanzees seem to be able to apply some form of teleological reasoning when acquiring means actions for a desirable goal (Horner & Whiten, 2005).

The other kind of social learning task where teleological interpretation of actions is essential is when one has to infer the novel goal of an action (and/or the function of an artefact) from observation. The difficulty of this task is that the useful effect of the behaviour can be either in the far future hence not observable for the learner (e.g., planting a trap), or may be embedded in a set of different co-occurring outcomes of which many will be by-products rather than goals of the action (e.g., hitting an object with another object may result in bending one, making a dent on the other, while producing a sound effect and sparks, etc.). Action-effect associations would not help the observer to figure out the goal in these situations, because the effect (and many times the action as well) is novel. In fact, relying on learned associations would mislead the observer when the task is to figure out an unfamiliar goal state. Similarly, simulation procedures would either lead to no goal extraction, or would arrive at a wrong conclusion by asking, "What desire would make me do this?" Such procedures would defeat the purpose of social learning of novel goals and

⁵ In fact, when the causal properties of means actions are opaque for the observer, even teleological reasoning will fail in this task. Elsewhere we hypothesize that human pedagogy (manifestation of relevant knowledge embedded in means actions) evolved to assist human children to acquire opaque means actions (Csibra & Gergely, 2006; Gergely & Csibra, 2005; in press).

artefact functions, because they would apply assumptions (replication of known goals, or extension of own goals to others) that are not valid in the actual situation.

In contrast, teleological reasoning may provide useful cues for valid goal attribution in such cases. Its basic assumption, according to which actors achieve their goals efficiently, allows the observer to sort the observed effects in terms of whether there is an easier way to achieve them than the action that produced them. The assumption of efficiency also helps to verify hypotheses about possible future goal states: the more the action seems to be the best way to achieve the hypothesized goal, the more likely that it is the real goal of the action. In many cases, the observer's knowledge about the causal properties of the objects involved in the action and about the motor constraints and preferences of the actor will not be enough for unambiguous identification of a single goal, but teleological reasoning will still enable to constrain the set of possible goal states to be considered sufficiently for the observer to learn about the function of actions and artefacts. Relevant background knowledge about situational factors or higher-order goals pursued may also be relied on to choose the goal to be attributed from co-occurring outcomes that are locally equally efficiently produced by the action. For example, while producing sparks and driving a nail in may be brought about equally efficiently by hammering, knowledge of the larger context and ultimate goal such as building a wooden house may help teleological reasoning to choose the latter rather than the former consequence as the actual goal of the hammering action (as it can efficiently serve as a sub-goal to the distal goal of house building, while making sparks cannot).

5 Conclusions

In this paper, we have provided a theoretical overview on the epistemic functions and potential mechanisms of teleological interpretation of actions. We identified two basic functions of goal attribution: providing local predictions and enabling longer-term social learning. Both functions involve two kinds of inferences according to the direction of computation between actions and goals, and both kinds of inferences pose inverse problems for observers, at least in humans.

In the developmental literature on goal understanding, three mechanisms have been proposed to fulfil these functions. The fundamental difference between these mechanisms originates from the tacit assumptions that they apply to solve the inverse problems of goal attribution. When matching these mechanisms and their assumptions to the functions to be fulfilled, we found varying degrees of success. The least flexible mechanism, recalling action-effect associations, performs well in many cases of goal prediction because it provides a quick on-line solution that is likely to be valid. However, when productive solutions are required, action-effect associations are less useful. In contrast, simulation procedures provide

new, productive solutions in action prediction contexts. In other words, one can predict actions that have never been perceived before by applying the basic assumptions that simulation procedures presuppose.

The productivity of teleological reasoning is even broader, making it possible to attribute goals to non-human agents, and to come up with previously unknown functions as goals. This aspect of teleological reasoning makes it particularly suitable to support social learning of means actions and artefact functions. The price to be paid for this flexibility, however, is lower speed of processing, which is a considerable disadvantage in on-line tasks, like action and outcome prediction.

As this summary suggests, we believe that the three proposed mechanisms of goal attribution do not compete but complement each other. First, none of them is better in all situations than the others. Depending on the task demands and the information available, one or the other mechanism would provide faster or more valid answers. Second, these mechanisms could support each other during their implementation. For example, the fast effect prediction provided by action-effect associations can serve as a starting goal-hypothesis for the teleological reasoning or simulation systems for verification, and the solutions provided by teleological reasoning in social learning situations can be stored as action-effect associations for subsequent rapid recall.

It is also misleading to characterize one or the other of these mechanisms as 'innate' or 'learned'. The basic assumptions that these mechanisms rely upon are likely to be innate in the sense of not having been derived from experience. Infants may be innately biased to associate actions with their outcomes, to assume some identification with other human beings (the "like me" hypothesis), and to search for justification for actions in terms of efficiency. At the same time, all three mechanisms rely on previously accumulated experience, and all can support further learning. Action-effect associations can only work from the material provided by earlier experience of such pairings, and accumulate such associations by storing new ones. Simulation procedures crucially depend on first-person action experience and especially on the ability to perform means-end action sequences, and such motor learning may immediately allow children to interpret and predict others' actions as well. Teleological reasoning also requires a knowledge base to operate on. Evaluating efficiency, for example, is impossible without some basic knowledge of physical potentials of objects and agents, and can even lead to insights into new and relevant physical or biological constraints. Thus, all three mechanisms should predict some early successes in goal attribution, followed by a gradual improvement of performance during development.

Because the three mechanisms we have discussed rely on different, and sometimes even conflicting,

assumptions, they cannot be derived from each other. Thus, the question of primacy (i.e., which one provides the 'real' basis of understanding of goal-directed actions) is the wrong question, which cannot be answered on conceptual grounds. Just like neither the 'shape from shading' nor the 'depth from occlusion' mechanism of 3D vision is more fundamental than the other, the relative contribution of the mechanisms of teleological interpretation to the development of understanding goal-directed actions is an empirical, rather than theoretical, question.

If, however, we contrast the three mechanisms on a phylogenetic rather than an ontogenetic time scale, it is more plausible to assume that they are differently applied. Many species are likely to form action-effect associations that help them to anticipate the outcome of observed actions before they have been finished. The existence of mirror neurons in monkeys also indicates that simulation procedures are invoked in action prediction in primates. One study even suggests that teleological reasoning operates in chimpanzees to acquire means actions from others (Horner & Whiten, 2005). To our knowledge, however, only humans apply teleological reasoning to figure out novel goals and artefact functions from others' actions. We speculate that this ability represents a human-specific adaptation to learn actions whose goal is not directly adaptive and to acquire functions of artefacts, and this adaptation is also reflected by our obsession to interpret events in functional terms (cf. Kelemen, 1999).

In addition, the productivity of teleological reasoning opens up the possibility of stipulating fictional goal states as well. We have shown previously (Csibra et al., 2003) that 12-month-old infants infer the presence of occluded physical objects, or the achievement of unseen goals, in order to interpret an event as an efficient goal-directed action. In other words, if a relevant aspect of reality that would justify an observed behaviour as a well-formed goal-directed action is perceptually not directly accessible, infants can "fill in" the missing element by mentally stipulated contents and rely on the efficiency assumption to make such inferences. A recent study by Onishi, Baillargeon, and Leslie (*present issue*) suggests that 15-month-olds can even go one step further and posit states of the world that are counterfactual to the perceptual evidence but will allow an observed action to be interpreted as a goal-directed pretence act. Briefly, the study demonstrates that infants expect that sequentially organized actions are directed towards a specific goal state: a pouring action is anticipated to result in an outcome state that provides the enabling condition (liquid in the cup) for a subsequent drinking action from the same cup. Crucially, infants recognize the causal and teleological relatedness of these actions even if no liquid is present and therefore neither the subgoal nor the final goal is ever achieved. This result strongly suggests that infants' obsession with goals make them stipulate a fictional world (Csibra &

Gergely, 1998) in which the observed action would be an efficient goal-directed action, decouple this world from reality (Leslie, 1987) and attribute it to the actor's mind. Although mental state attribution clearly requires further cognitive mechanisms as well, this example nicely demonstrates how teleological obsession and teleological reasoning allow infants to figure out the contents of mental states that are necessary to understand pretence acts as well-formed goal-directed actions within a mentally stipulated representation of fictional reality. Thus, teleological reasoning functions not only to assist social learning of novel means and ends but also to infer mental state contents from actions.

References

- Baker, C. L., Tenenbaum, J. B., & Saxe, R. R. (in press). Bayesian models of human action understanding. *Advances in Neural Information Processing Systems*.
- Behne, T., Carpenter, M., Call, J., & Tomasello, M. (2005). Unwilling versus unable: infants' understanding of intentional action. *Developmental Psychology*, *41*, 328-337.
- Biró, S. & Leslie, A. M. (in press). Infants' perception of goal-directed actions: Development through cue-based bootstrapping. *Developmental Science*.
- Boesch, C. & Boesch, H. (1993). Diversity of tool use and tool-making in wild chimpanzees. In A. Berthelet & J. Chavaillon (Eds.), *The Use of Tools by Human and Non-Human Primates* (pp. 158-187). Oxford: Oxford University Press.
- Breazeal, C., & Scassellati, B. (2002). Robots that imitate humans. *Trends in Cognitive Sciences*, *6*, 481-487.
- Call, J., Hare, B., Carpenter, M., & Tomasello, M. (2004). Unwilling or unable: Chimpanzees' understanding of human intentional action. *Developmental Science*, *7*, 488-498.
- Casler, K. & Kelemen, D. (2005). Young children's rapid learning about artifacts. *Developmental Science*, *8*, 472-480.
- Csibra, G. (2003). Teleological and referential understanding of action in infancy. *Philosophical Transactions of the Royal Society, London B*, *358*, 447-458.
- Csibra, G. (2005). Mirror neurons and action understanding. Is simulation involved? <http://www.interdisciplines.org/mirror>.
- Csibra, G., Biró, S., Koós, S., & Gergely, G. (2003). One-year-old infants use teleological representations of actions productively. *Cognitive Science*, *27*, 111-133.
- Csibra, G., & Gergely, G. (1998). The teleological origins of mentalistic action explanations: A developmental hypothesis. *Developmental Science*, *1*, 255-259.
- Csibra, G. & Gergely, G. (2006). Social learning and social cognition: The case for pedagogy. In Y. Munakata & M. H. Johnson (Eds.), *Processes of Change in Brain and Cognitive Development. Attention and Performance XXI* (pp. 249-274). Oxford: Oxford University Press.
- Csibra, G., Gergely, G., Biró, S., Koós, O., & Brockbank, M. (1999). Goal attribution without agency cues: The perception of 'pure reason' in infancy. *Cognition*, *72*, 237-267.
- Dennett, D. (1978). *Brainstorms*. Cambridge: MIT Press.
- Elsner, B. (present issue). Infants' imitation of goal-directed actions: the role of movement and action effects.
- Elsner, B. & Aschersleben, G. (2003). Do I get what you get? Learning about effects of self-performed and observed actions in infants. *Consciousness & Cognition*, *12*, 733-752.
- Farroni, T., Johnson, M. H., Menon, E., Züljan, L., Faraguna, D., & Csibra, G. (2005). Newborns' preference for face-relevant stimuli: Effects of contrast polarity. *Proceedings of the National Academy of Sciences of the United States of America*, *102*, 17245-17250.
- Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge: MIT Press.
- Galef, B. G. Jr. (1995). Why behaviour patterns that animals learn socially are locally adaptive. *Animal Behaviour*, *49*, 1325-1334.
- Gallese, V. & Goldman, A. (1998) Mirror neurons and the simulation theory of mind reading. *Trends in Cognitive Sciences*, *2*, 493-501.

- Gergely, G. (2003). What should a robot learn from an infant? Mechanisms of action interpretation and observational learning in infancy. *Connection Science*, 15, 191-209.
- Gergely, G., Bekkering, H., & Király, I. (2002). Rational imitation in preverbal infants. *Nature*, 415, 755.
- Gergely, G. & Csibra, G. (2003). Teleological reasoning in infancy: the naïve theory of rational action. *Trends in Cognitive Sciences*, 7, 287-292.
- Gergely, G., & Csibra, G. (2005). The social construction of the cultural mind: Imitative learning as a mechanism of human pedagogy. *Interaction Studies*, 6, 463-481.
- Gergely, G. & Csibra, G. (in press). Sylvia's recipe: The role of imitation and pedagogy in the transmission of human culture. In N. J. Enfield & S. C. Levinson (Eds.), *Roots of Human Sociality: Culture, Cognition, and Human Interaction*. Oxford: Berg Publishers.
- Gergely, G., Nádasdy, Z., Csibra, G., & Biró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, 56, 165-193.
- Goldman, A. (in press). *The Simulating Mind*. New York: Oxford University Press.
- Gordon, R. (1986). Folk psychology as simulation. *Mind and Language*, 1, 158-171.
- Hauf, P., & Prinz, W. (2005). The understanding of own and others' actions during infancy: "You-like-me" or "Me-like-you". *Interaction Studies*, 6, 429-445.
- Heyes, C. M. (1993). Imitation, culture and cognition. *Animal Behaviour*, 46, 999-1010.
- Heyes, C. M. & Galef, B. G. Jr. (Eds.) (1996). *Social Learning in Animals: The Roots of Culture*. New York: Academic Press.
- Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001). The Theory of Event Coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, 24, 849-937.
- Horner, V. & Whiten, A. (2005). Causal knowledge and imitation/emulation switching in chimpanzees (*Pan troglodytes*) and children (*Homo sapiens*). *Animal Cognition*, 8, 164-181.
- Iacoboni, M. (2005). Mechanisms of imitation and imitation in animals. In S. Hurley & N. Chater (Eds.), *Perspectives on Imitation: From Mirror Neurons to Memes* (Vol. 1, pp. 77-100). Cambridge: MIT Press.
- James, W. (1890). *The Principles of Psychology*. New York: Henry Holt.
- Kamewari, K., Kato, M., Kanda, T., Ishiguro, H., & Hiraki, K. (2005). Six-and-a-half-month-old children positively attribute goals to human action and to humanoid-robot motion. *Cognitive Development*, 20, 303-320.
- Keil, F. C. (2006). Explanation and understanding. *Annual Review of Psychology*, 57, 227-254.
- Kelemen, D. (1999). Functions, goals, and intention: Children's teleological reasoning about objects. *Trends in Cognitive Sciences*, 3, 461-468.
- Király, I., Jovanovic, B., Prinz, W., Aschersleben, G., & Gergely, G. (2003). The early origins of goal attribution in infancy. *Consciousness and Cognition*, 12, 752-769.
- Laland, K. N. & Hoppitt, W. (2003). Do animals have culture? *Evolutionary Anthropology*, 12, 150-159.
- Leslie, A. M. (1984). Infant perception of a manual pick-up event. *British Journal of Developmental Psychology*, 2, 19-32.
- Leslie, A. M. (1987). Pretense and representation: the origins of 'theory-of-mind'. *Psychological Review*, 94, 412-426.
- Marr, D. (1982). *Vision: A Computational Investigation Into the Human Representation and Processing of Visual Information*. New York: Freeman.
- Meltzoff, A. N. (2002). Imitation as a mechanism of social cognition: Origins of empathy, theory of mind, and the representation of action. In U. Goswami (Ed.), *Blackwell Handbook of Childhood Cognitive Development* (pp. 6-25). Malden, MA: Blackwell.
- Meltzoff, A. N. (2005). Imitation and other minds: The "like me" hypothesis. In S. Hurley & N. Chater (Eds.), *Perspective on Imitation: From Neuroscience to Social Science* (Vol. 2, pp. 55-77). Cambridge, MA: MIT Press.
- Meltzoff, A. N. (present issue). The 'like me' framework for recognizing and becoming an intentional agent.
- Mulcahy, N. J. & Call, J. (2006). Apes save tools for future use. *Science*, 312, 1038-1040.
- Onishi, K. H., Baillargeon, R., & Leslie, A. M. (present issue). 15-month-old infants detect violations in pretend scenarios.
- Phillips, A. T. & Wellman, H. M. (2005). Infants' understanding of object-directed action. *Cognition*, 98, 137-155.
- Ramachandran, V. (1988). Perception of shape from shading. *Nature*, 331, 163-166.
- Rizzolatti, G. & Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience*, 27, 169-192.
- Schmidt, D. M., George, J. S., & Wood, C. C. (1999). Bayesian inference applied to the electromagnetic inverse problem. *Human Brain Mapping*, 7, 195-212.
- Scholl, B. J. (2005). Innateness and (Bayesian) visual perception: Reconciling nativism and development. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The Innate Mind: Structure and Contents* (pp. 34-52). New York: Oxford University Press.
- Sodian, B., Schoeppner, B., & Metz, U. (2004). Do infants apply the principle of rational action to human agents? *Infant Behavior and Development*, 27, 31-41.
- Sommerville, J. A., & Woodward, A. L. (2005). Pulling out the intentional structure of action: The relation between action processing and action production in infancy. *Cognition*, 95, 1-30.
- Somerville, J. A., Woodward, A. L., & Needham, A. (2005). Action experience alters 3-month-old infants' perception of others' actions. *Cognition*, 96, B1-B11.
- Tomasello, M. (1996). Do apes ape? In *Social learning in animals: The roots of culture*, edited by C. M. Heyes & B. G. Galef, 319-346. New York: Academic Press.
- Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Cambridge, MA: Harvard University Press.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28, 675-735.
- Wagner, L. & Carey, S. (2005). 12-month-old infants represent probable ending of motion events. *Infancy*, 7, 73-83.
- Weir, A. A. S., Chappel, J., & Kacelnik, A. (2002). Shaping of hooks in New Caledonian crows. *Science*, 297, 981.
- Whiten, A., Horner, V. & Marshall-Pescini, S. (2003). Cultural Panthropology. *Evolutionary Anthropology*, 12, 92-105.
- Wilson, M. & Knoblich, G. (2005). The case for motor involvement in perceiving conspecifics. *Psychological Bulletin*, 131, 460-473.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, 69, 1-34.
- Woodward, A. (1999). Infants' ability to distinguish between purposeful and non-purposeful behaviours. *Infant Behavior & Development*, 22, 145-160.
- Woodward, A. L., & Sommerville, J. A. (2000). Twelve-month-old infants interpret action in context. *Psychological Science*, 11, 73-77.
- Woodward, A. L., Sommerville, J. A., & Guajardo, J. J. (2001). How infants make sense of intentional action. In B. Malle, L. Moses & D. Baldwin (Eds.), *Intentions and Intentionality: Foundations of Social Cognition* (pp.149-169). Cambridge, MA: MIT Press.

Acknowledgements

We thank Alan Leslie and Victoria Southgate for their valuable comments on the draft of this paper. This work was partly supported by the UK MRC Programme Grant #G9715587.

Correspondence concerning this article should be addressed to Gergely Csibra, School of Psychology, Birkbeck, University of London, Malet Street, London WC1E 7HX, United Kingdom. E-mail: g.csibra@bbk.ac.uk