Perspectives on Imitation: From Neuroscience to Social Science

Volume 2: Imitation, Human Development, and Culture

edited by Susan Hurley and Nick Chater

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

2005

5 Intention Reading and Imitative Learning

Michael Tomasello and Malinda Carpenter

5.1 Introduction

Imitation presents very difficult problems for mechanistic theories of human psychology. Behaviorism never knew what to do with imitation, since it represented a kind of learning not easily analyzable into stimuli and responses. Information-processing psychology has basically ignored imitation; virtually no textbook of cognitive psychology or cognitive science even mentions it. Perhaps the main problem is that social learning and imitation are very closely intertwined with processes of social cognition—understanding how other persons work and understanding what a particular person is doing on a specific occasion—and these also have not been given much attention by either behaviorists or mainstream cognitive scientists.

Tomasello et al. (1993a) provided an evolutionary and developmental account of the way human social learning relates to processes of social cognition more generally. The basic proposal was that what differentiates human social learning from that of other animals, including our nearest primate relatives, is that humans understand the behavior of others not just as body movements, but as intentional, goal-directed action. Simply put, humans perceive others not as moving their limbs in particular ways, but as doing such things as opening a drawer, giving a gift to someone, washing the dishes, telling a story, throwing a ball—each of which may be done with many different body movements so long as the same goal in the external world is reached. Thus, when they attempt to reproduce the actions of others, humans—at least in some circumstances—reproduce the actions as they have understood them from the point of view of the intentionality involved, that is, the intended effect on the external world, including the social world.

In the decade since the publication of the paper by Tomasello et al. (1993a), empirical studies of the social learning of various animal species

and human children have multiplied severalfold. Consequently, there have been important advances—empirically, theoretically, and methodologically—in the understanding of these topics. In this chapter we review the most important empirical and methodological contributions of the past decade in this field, and we then try to assess what they mean for our understanding of how processes of social learning are shaped, or perhaps even created, by different forms of social cognition—especially those we will call intention reading. First we deal briefly with a few important studies on nonhuman primates and other nonhuman animals, but our attention is primarily directed at the many new studies focused on human children, mainly in the second year of life.

5.2 Imitative Learning in Nonhuman Animals

Tomasello et al. (1993a) claimed that there was no convincing evidence that any nonhuman animals engaged in humanlike imitative learning. They explained humanlike imitative learning in the following way: A child, for example, observes an adult using a knife to open a bottle. The child understands the intentional structure of the action; the adult's goal is to open the bottle and she has chosen one behavioral means, among other possible means, for doing this. The child can then either adopt the adult's means or not, as she chooses. Other animals do not understand intentional action in this way, and so they cannot engage in this kind of social learning process, which Tomasello et al. (1993a) called imitative learning, one form of cultural learning. Around the same time, Whiten and Ham (1992) coined another term, "goal emulation," to describe the case in which the learner parses the observed action into end and means, but then chooses to ignore the means used by the demonstrator and employ a different one instead. Both imitative learning and goal emulation thus require some understanding of the demonstrator's goal or intention.

The claim was that other animals do not perceive or understand the distinction between goals as mental entities representing the desired state of affairs the organism is attempting to bring about, and behavioral means toward goals ("means-ends dissociation" in the terminology of Piaget, 1952/1963). They thus engage in two other forms of social learning instead. One is mimicking. Mimicking refers to the process in which an organism reproduces the body movements of another, without an understanding of any goal that might be structuring those movements. For example, Hayes and Hayes (1952) trained their human-raised chimpanzee Viki to reproduce various body movements and gestures that they performed, for example, blinking their eyes or clapping their hands. They trained her throughout her daily life in their home for a period of more than 17 months before systematic testing began. The training consisted of a human performing a behavior and then using various shaping and molding techniques, with rewards, to get Viki to repeat the behavior ("Do this"). After she had become skillful, some novel behaviors were systematically introduced. In general, she reproduced them faithfully and quickly; she had clearly "gotten the idea" of the mimicking game.

Later, Custance et al. (1995) demonstrated in a more rigorous fashion similar abilities in two nursery-reared chimpanzees after they were trained for 3.5 months in a manner similar to that for Viki. Of the forty-eight novel actions demonstrated after the training period, one subject correctly reproduced thirteen and the other correctly reproduced twenty. Whether the subjects could see their own responses or whether the response was "invisible" (e.g., facial expressions) was not a significant factor. Evidence that this is only mimicking comes from a study by Call and Tomasello (1995), who tested a human-raised orangutan who was trained in similar ways. He was successful in learning body movements, but then when shown how to solve a problem involving an object (a human used a tool and told the orangutan to "Do this"), he could not transfer his mimicking skills to this goal-directed action.

The best mimickers in the animal world are various species of birds (especially, for example, parrots). Birds' skills of vocal mimicry are well known, but recently they have also been shown to reproduce the body movements of others whom they observe solving a task. For example, if they observe a conspecific using its beak to open a container, they will do the same, whereas if they observe the conspecific using its foot to open a container, they will do that (F. Campbell et al., 1999; Zentall et al., 1996; see also a similar behavior by a marmoset species reported by Voelkl & Huber, 2000). But it is not currently known whether these intriguing observations represent cases of imitative learning, in which the learner discerns the intentional structure of the observer's behavior and chooses to reproduce its means to reach the goal, or instead mimicry, in which the observer simply copies body movements relatively blindly. Research with human infants has attempted to distinguish these two cases in various ways, as we will see later, and these could be used to good effect in the studies with birds.

The other major form of social learning that does not depend on reading intentions is called emulation learning. In emulation learning, an observer simply watches the objects with which a demonstrator is interacting and

learns something about them. When an organism observes another organism manipulate objects in this way, it learns a lot about those objects and their affordances for action—much more than when it observes those objects sitting idle. Thus, if we give naive chimpanzees a rock and a nut, they may not discover on their own how to crack the nut open. But if they see another do it, they might learn from this observation that nuts can be opened, which creates a new possibility for them, and they might even learn something about the rock's role in the process. In emulation learning, the organism learns new things, some of them quite complex; it is just that they are about the environment, not about behavior. We can say that the observer is attending to the end result in the environment produced by the other—which was, from the human point of view, the behaver's goal. But in emulation learning, as opposed to imitative and other forms of cultural learning, the observer does not attend at all to the behavior of the other, much less to the other's goal.

There are many studies that show that chimpanzees and other apes are good emulators (e.g., Nagell et al., 1993; Call & Tomasello, 1995; see Tomasello, 1996, for a review). One apparently contradictory study is that of Whiten et al. (1996), who presented chimpanzees with a transparent "foraging box" containing fruit. On any given trial, the box could be opened by one of two mechanisms, each of which could be operated in two ways. For each mechanism, a human experimenter demonstrated one way of opening the box to some subjects and the other way to other subjects (with the other mechanism being blocked). The subjects were then given the chance to open the box themselves. The results were that for one mechanism there was no effect of the observed demonstration. For the other mechanism there was some evidence that chimpanzees were more likely to use the manner of opening demonstrated by the experimenter. The authors claimed to have demonstrated imitative learning. However, in the analysis by Tomasello (1996), the chimpanzees could easily have learned to match the manner of opening via emulation. One group of chimpanzees saw that the stick afforded pushing through the clasp, while the other saw that it afforded twisting and pulling. A similar analysis applies to a recent study with marmosets, who were especially skillful at learning a box's affordances through emulation (Bugnyar & Huber, 1997).

Some researchers have claimed that emulation learning is not a useful concept because it can never be distinguished from imitative learning when an organism reproduces some action on an object because the object normally undergoes some kind of transformation from which the learner could be learning (R. Byrne & Russon, 1998). But in research with human

children, investigators have created some interesting new paradigms that are potentially capable of distinguishing imitative learning and goal emulation from emulation learning and mimicking. These are elaborated later, but they basically involve the demonstrator either trying unsuccessfully to do something or doing something by accident—the two main cases where the demonstrator's intention does not match the result produced on the object in the real world. In these cases, a learner engaged in imitative learning or goal emulation will reproduce the action she believes the demonstrator intended, whereas a mimicker will produce the body movements only (including the unsuccessful and accidental actions), and an emulator will reproduce the unintended result. There have been two attempts to do these kinds of studies with chimpanzees (Call et al., submitted; Myowa-Yamakoshi & Matsuzawa, 2000), but neither study produced clearly interpretable results.

5.3 Imitative Learning in Human Children

Human neonates can mimic some facial expressions of other persons from very soon after birth (Meltzoff & Moore, 1977, 1989). By the middle of the first year, infants also can learn new things about objects and their affordances via emulation learning (Barr et al., 1996; von Hoftsten & Siddiqui, 1993). But until recently, it was unclear at what age infants were capable of engaging in imitative learning and goal emulation with intention reading. Recently a number of novel experimental techniques have been devised in an effort to determine precisely what kinds of social learning children are employing. The basic problem is that if, for example, an adult takes the top off of a pen and a child then does the same, there are many possible explanations, including emulation, mimicking, and imitative learning, among others. Researchers have therefore employed some of the following techniques for dissociating different components of what the child perceives, understands, and reproduces of the demonstration:

- demonstrations with unusual means to an end that children would be unlikely to use on their own spontaneously, and monitoring children's looking behavior to see if they are concerned with the goal;
- demonstrations of trying or failing or having an accident, in which the adult's surface behavior is not an accurate reflection of what she intends to do; and
- demonstrations that can be interpreted in different ways, depending on the child's understanding of what the adult is intending to do—which is

manipulated in various ways by providing different interpretive contexts before or during the demonstration.

In what follows we look at examples of these different techniques in turn and then at some other interesting cases of intention reading and imitative learning that have interesting theoretical implications but that have been less well researched.

5.3.1 Unusual Means

Meltzoff (1988a) presented 14-month-old infants with a novel demonstration of an adult bending down and touching her head to the top of a box, which then lit up. Although the infants could more easily have solved this task by emulation (e.g., by touching the box with their hand instead of their head), they instead chose to use the same means as the adult, unusual as it was—they too touched their head to the top of the box. These infants could have been either mimicking the adult's unusual action without understanding her goal of turning on the light, or else they could have been copying this action with the same goal in mind—imitative learning. In order to determine which of these two social learning mechanisms the infants were using, Carpenter et al. (1998b) tested 9- to 15-month-old infants on the head touch and other similar tasks. However, they delayed the illumination of the light slightly after the infant's reproduction of the action, and coded whether infants looked in anticipation to the light. They found that 12-month-old and older infants, on average, looked to the light in anticipation (before it came on). If the light did not come on, these infants often repeated their action or looked quizzically at the adults. This suggests that the infants were adopting the adult's means in order to achieve the same goal as the adult—to turn on the light. The infants thus were not just mimicking the adult's action but instead were engaging in imitative learning of her novel action.

5.3.2 Accidents, Trying, and Failing

By 12 months of age it seems that infants are not just mimicking adults but instead can discern their goals and choose to use the same behavioral means the adults used to achieve the same goal. It has been assumed in these studies that the child understands the adult's goal mentalistically—she sees a kind of thought bubble coming from the adult's head containing a picture of the end state the adult is trying to bring about—but this is not necessarily the case. It is possible that what the infants were trying to reproduce was simply the observed result of the adult's action (the illuminated light—the external goal) with no understanding of the adult's mental goal or intention. Two further studies of infants in this age range—in which the observed result of the action differed from the adult's thought-bubble intention or goal—indicate that infants do indeed use their understanding of others' goals and intentions when learning novel actions from them.

In one study, infants were shown the same actions, for the same results, but the adult's intention varied across conditions. The infants could thus only use an understanding of the adult's intentions to solve this problem correctly. Carpenter et al. (1998a) showed 14- to 18-month-olds a series of two actions on objects, in counterbalanced order. For each object, the two actions were followed by an interesting result, for example, the sudden illumination of colored lights. In the key conditions, one of the demonstrator's actions was marked verbally as intentional ("There!") and one was marked verbally as accidental ("Whoops!"), but otherwise the actions looked very similar. Instead of mimicking both actions they saw, even the youngest infants reproduced the actions marked as intentional significantly more often than those marked as accidental. In a third condition, when both actions were marked verbally as intentional, the infants typically reproduced both actions.

In the other study, infants were able to go beyond filtering out unintended actions and achieve results that they never saw in their entirety. Meltzoff (1995) showed 18-month-olds an adult either successfully achieving a result on an object (e.g., pulling apart two halves of a dumbbell) or trying but failing to achieve that result (e.g., the adult's hands slipped off the ends of the dumbbell). The infants produced the completed result equally often in both conditions—whether they had seen the adult produce that result or had only seen the adult's intention but not the completed result. Various control conditions reported by Meltzoff et al. (1999) indicated that infants in the intention condition were not just absorbing lower-level information such as the first part of the action (e.g., grasp the toy here and make pulling movements). When the adult demonstrated the failed action on a small dumbbell and gave infants a giant dumbbell, the infants tried different means to pull the toy apart. Likewise, when infants were given a trick dumbbell that was glued together, their hands slipped off just like the adult's had, but they continued trying to separate it, again using different means. Meltzoff thus concluded that infants understood the adult's unfulfilled intention and produced the result that the adult meant to produce (instead of copying the adult's surface behavior).

Bellagamba and Tomasello (1999) replicated these findings with 18-month-olds, but found that 12-month-olds did not reproduce the adult's intended action when they only saw her trying unsuccessfully to perform it.

A related study with older children is that by Want and Harris (2001). In this study, the adult showed 2.5- and 3.5-year-old children a mistaken action before showing them the correct solution to a problem or showed them the correct solution alone. That is, in the mistake condition, the adult first inserted a tool into the incorrect side of an apparatus, saying "Oops," and then inserted it into the correct side. Three-year-olds but not 2-year-olds performed better in the mistake condition than in the condition in which they saw only the correct solution (although it is unclear whether the children benefited from watching the adult's mistake or simply gained more information than the other children from the extra highlighting of causal information in the incorrect demonstration).

So by 14 to 18 months of age, infants can distinguish between intentional actions and accidental or unfulfilled actions, and choose to reproduce (or produce) the intended actions or results. Instead of copying exactly what others do, infants do what others intend to do. They do this for human actors, but not for machines (Meltzoff, 1995). They thus are not mimicking or emulating in these situations, but instead are engaging in the imitative learning (or goal emulation) of adults' goal-directed actions.

5.3.3 Manipulating Children's Interpretation

Other studies have manipulated the social learning context in an effort to influence children's interpretation of adult intentional action, which should have an influence on what they reproduce if, and only if, they are interpreting the behavior intentionally. Gergely et al. (2002) showed 14-month-olds an adult touching her head to the top of a box to turn on a light (as in Meltzoff, 1988a). However, for half of the infants, the adult's hands were occupied (she was holding a blanket around her shoulders) and for half the adult's hands were free. Infants who saw the hands-free demonstration touched the box with their heads significantly more often than infants who saw the hands-occupied demonstration (all the infants also touched the top of the box with their hands). The infants thus used the context of the situation to interpret the adult's behavior, appearing to assume that if the adult's hands were free and she still chose to use her head, then there must be a good reason for this choice. However, if the adult's hands were occupied, then the use of her head was explained as necessary given her circumstances—not an essential part of her action (and thus the infants did not reproduce this action). These infants' interpretation of the

adult's goal probably differed across conditions, depending on the context of the situation (even though the adult's actions were identical). In the hands-occupied condition her goal was "turn on the light," and in the hands-free condition it was "turn on the light with your head." By 14 months, infants thus evidence a deeper understanding of intentional action and how it relates to the surrounding context and what this means for their own choice of a behavioral means in similar circumstances.

A series of studies of older children by Bekkering and colleagues (Bekkering et al., 2000; Gleissner et al., 2000; Wohlschläger et al., in press) extends these findings. For example, Bekkering et al. (2000) showed 3- to 6-year-old children an experimenter touching a table in one of two locations. In one condition, there were dots on the table in those locations and in another condition there were no dots. In the "no dot" condition, the children usually matched the adult's behavior exactly, even copying her crossed or straight arm positions—presumably because there was no other apparent goal to her actions than these arm movements. In the "dot" condition, however, the children touched the same locations as the experimenter, but often did not match her exact arm positions. This is presumably because when there were dots they interpreted the adult's goal as "touching the dots" (by whatever means), whereas when there were no dots the only possible goal was "moving one's arm like this." Bekkering and colleagues concluded that imitation in young children is guided by their understanding of adults' goals; that is, there is a hierarchy of goals and subgoals and children imitate what they perceive adults' main goal to be. Sometimes this involves matching others' actions; sometimes it does not.

A further study also shows that children's use of the context to interpret adults' actions influences what they learn from a demonstration. Carpenter et al. (2002) demonstrated to five groups of children how to pull out a pin and open a box. What differed among groups was what children experienced just prior to this demonstration, with some children receiving information about the adult's "prior intention" (i.e., what she intended to do with the box as she approached it). One group of children watched this demonstration alone; these children thus did not know what the adult's prior intention was. Three other groups received some information about the adult's prior intention before seeing this demonstration. Either the adult tugged unsuccessfully on the door of the box, or showed the box already open, or visited and opened three different boxes before demonstrating how to open the test box. Thus, all children in all four of these conditions saw a full demonstration of how to open the box, but only the children in the three prior intention conditions could know what the adult

was about to do before she began this demonstration. The children in a control demonstration group saw a goal-irrelevant action (the adult raked her fingers down the roof of the box) before the full demonstration, and a sixth baseline group received no demonstration at all. Two- and 2.5-year-old children were significantly better at opening the box themselves when they knew the adult's prior intention. This was the case even when the adult's actions on the test box were absolutely identical in the prior and no prior intention groups (i.e., when children gained information about the adult's prior intention through her actions on other boxes). It is interesting that the children who did not know what the adult was about to do performed just as poorly as the children who received no demonstration at all. This study demonstrates with special clarity the role of intentional understanding in children's imitative learning because a child that did not interpret the adult's behavior intentionally would have learned the same amount in all the different conditions. Instead, the way that the children interpreted the adult's prior intentions actually enabled them to imitatively learn something that they otherwise could not have learned.

What is remarkable about children's behavior in these studies is this: Children saw the same demonstration in different experimental conditions. What differed across conditions was various contextual factors that led the children to interpret the adult's actions in different ways, that is, under different intentional descriptions: touching a dot or not, trying to open a box or not, trying to turn on a light or trying to turn on a light with the head. As young as 14 months, the children's interpretation of the adult's action was then directly reflected in what they subsequently attempted to reproduce of the adult's behavior.

5.3.4 Imitating Reciprocal Behavior

There are some kinds of actions that children observe and attempt to imitate that have a special structure because they involve people having goals toward one another reciprocally. For example, a mother might blow a raspberry on her child's arm. If a child wants to imitate this behavior, she is faced with a choice, depending on her interpretation. Thus, she might blow a raspberry on her own arm, in exactly the same place the mother did, or alternatively, she might blow a raspberry back on her mother's arm—interpreting the behavior in this case reciprocally as "blowing on the partner's arm." Tomasello (1999) called this role reversal imitation. In a pilot study, Carpenter et al. (submitted) have found that 12- and 18-month-old children are able to employ this reciprocal interpretation in some cases. At both ages, children reciprocated in these kinds of dyadic, body-oriented

situations, but 18-month-olds were more likely than 12-month-olds to reciprocate in situations involving interactions around objects.

This same interpretation applies to the learning of a piece of language, since learning to use linguistic symbols is also reciprocal. Thus, when an adult uses a linguistic symbol in a communicative act, the adult intends things toward the child's attentional state; she wants the child to attend to something. Consequently, to learn to use a symbol like an adult, a child must learn to use it toward the adult in the same way the adult used it toward the child (Tomasello, 1999). It is interesting that Rakoczy et al. (2004) have provided evidence that something like this is also going on in children's early symbolic play. Before 2 years of age, children learn symbolic behaviors with objects by imitatively learning them from adults, in much the same way that they learn instrumental actions with artifacts. But from about 2 years of age on, they look to the adult more often, and in some cases smile more often, when producing the symbolic behaviors. This is evidence that children of this age are reproducing a special kind of intentionality, a kind of shared intentionality (mutually reciprocal) in which for the moment we agree to, for example, treat this pencil as if it were a horse.

5.3.5 Learning Words

Given the general ability to learn a linguistic symbol through role-reversal imitation, it is still the case that in learning particular words on particular occasions children often need to read the adult's intentions to connect the word appropriately to its intended referent. Several language acquisition studies show that children as young as 18 months can combine all of the types of intention reading we have discussed earlier while imitatively learning novel words. For example, in a study of 24-month-olds by Tomasello and Barton (1994), an adult announced her (prior) intention to find a target object by saying, "Let's go find the *toma*." She searched through several buckets, extracting and rejecting with a scowl the novel objects inside. She then extracted another novel object with an excited expression and stopped searching. In a later comprehension test, when asked to go get the *toma* themselves, children chose the object the adult had identified as fulfilling her intention. Akhtar and Tomasello (1996) used a modified procedure to show that 24-month-old children could identify the intended referent even when the adult was unable to open the container with the target object inside—that is, when she had an unfulfilled intention. Tomasello et al. (1996) replicated both these studies with 18-month-old children.

Another study investigated children's use of their understanding of intentional versus accidental actions when learning novel words. In a study of 24-month-olds by Tomasello and Barton (1994), the adult announced her (prior) intention to perform a target action by saying "I'm going to *meek* Big Bird!" She then performed one accidental action (saying "Whoops!") and one intentional action (saying "There!"), in counterbalanced order. Later, when the children were asked to *meek* a different character themselves, they performed the action the adult had marked as intentional. Finally, using a preferential-looking paradigm, Poulin-Dubois and Forbes (2002) found that 27- but not 21-month-old children could use an actor's eye-gaze and gestures to learn verbs that differed only in intention (e.g., topple versus knock over). These word-learning studies thus also provide evidence of children's understanding of accidental actions, unfulfilled intentions, and prior intentions by 18–24 months of age.

5.3.6 Children with Autism

Children with autism show mixed results on tests of imitation (see S. Rogers, 1999 and I. Smith & Bryson, 1994, for reviews). They are clearly very good at mimicking—their tendency to engage in echolalia is evidence of that. However, very few studies have attempted to determine what other social learning mechanisms these children are capable of, and in particular how their understanding of intentions (or the lack thereof) might affect the process. Carpenter et al. (2002) found that 3- to 4-year-old children with autism (with nonverbal mental ages ranging from 28 to 50 months) performed at near-ceiling levels on imitation tasks such as the head touch task used by Carpenter et al. (1998a). These children too looked to the light in anticipation (before it came on), as often as did a control group of children with other developmental delays—indicating their appreciation of the goal-directed nature of this action. Two studies that used versions of Meltzoff's (1995) test of understanding of others' unfulfilled intentions also found no impairment for children with autism (Aldridge et al., 2000; Carpenter et al., 2001), again suggesting their appreciation of the goal-directed nature of these unfulfilled actions.

However, there may still be important differences in the way in which children with autism and other children copy others' behavior. For example, Hobson and Lee (1999) found that adolescent children with autism copied the particular style of a demonstrator's actions less often than did developmentally delayed and typically developing children. They interpreted this deficit in terms of autistic children's problems in identifying with the "attitude" of the adult demonstrator, including her intention in using a particular behavioral style.

5.3.7 Criticisms

Recently there have been several studies that have attempted to show that intention understanding is unnecessary for children to show the same pattern of results as that taken as evidence of understanding of others' intentions in some imitation studies. For example, in a follow-up to Meltzoff's (1995) study of understanding of unfulfilled intentions, Huang et al. (2002) found that infants produced the complete target action as often in conditions in which they were provided with information about the objects' affordances or in which the relevant parts of the objects were moved in close proximity to each other as they did in the completed and unfulfilled conditions of Meltzoff (1995). Likewise, Thompson and Russell (in press) found that children who saw a "ghost" condition in which the experimental apparatus moved by itself performed as well as children who saw an adult demonstrate the same actions on the apparatus. These researchers concluded that children could be using nonimitative, nonsocial-cognitive understanding to succeed in the intention imitation (and other imitation) tasks.

However, there are several imitation studies in which these kinds of explanations do not hold because the actions are exactly the same (or very similar) across the different intention conditions. For example, in the study by Carpenter et al. (1998a) of children's understanding of accidental and intentional actions, the actions were performed as similarly as possible across conditions, with the main difference being the verbal labeling ("Whoops" or "There") of the adult's intention. Likewise, in the studies by Gergely et al. (2002) and by Bekkering and colleagues, the only difference between conditions was whether the adult had a blanket around her or whether there were dots on the table. Finally, in the study by Carpenter et al. (2002) of children's understanding of others' prior intentions, in the two most interesting conditions—when the children did not know the adult's prior intention and when they inferred her prior intention from her actions on other boxes—again, the actual demonstrations on the test box were absolutely identical and the children succeeded in imitating only when they knew the adult's prior intention. These studies provide particularly strong evidence of children's understanding of others' intentions because this is the only thing that varied across conditions.

Thus, given the evidence from all of the different imitation paradigms, these alternative explanations are not so likely across the board, and we thus think it is more parsimonious to assume that children do use their understanding of intentional action to reenact the behavior of others. They may also be capable of using other mechanisms, such as an understanding of the affordances of objects or spatial contiguity as well, but we believe that their understanding of the intentions of others is more basic and, in certain situations like language acquisition, is the predominant method used from early on.

5.4 Conclusion

In an admirable attempt at simplicity, some researchers have proposed that we define imitation as "doing what others do" (e.g., R. Byrne, 1995). This is fine as far as it goes, but it takes only a moment's reflection to see that "what others do" can be defined or interpreted in many different ways from an intentional point of view. Is the adult moving her hands around the box in certain ways or trying to open the box? Is the adult moving her arm in a certain way or touching the dot? Is the adult turning on the light or turning on the light with her head? Is an accidental act part of what someone is doing? The point is that social learning and imitation cannot be viewed in isolation from the child's other social-cognitive skills. Children reproduce what they understand others to be doing, and so their social learning and imitation depend crucially on their skills of social cognition, especially those involved in reading the intentions of other people. The different manifestations of intention reading in young children's imitative learning, as reviewed here, are summarized in table 5.1.

Overall, research over the past decade has supported the general view of Tomasello et al. (1993a) with regard to these matters. Although there is some controversy about the skills of great apes in some experiments, there is no doubt that human children are the planet's most skillful imitators (even in the 1996 study by Whiten et al., the children were much better than the chimpanzees). Great apes mostly focus on changes of state in the objects involved in a demonstration, and they pay much less attention to the actual behavior, much less intentions, of the demonstrator. Human children, in contrast, focus on the intentions of the demonstrator to give them a definition of what she is doing, and this manifests itself in the myriad ways reviewed here. However, this does not mean that our knowledge has not advanced. It has in many ways, the most important of which are listed here.

Table 5.1
Development of intention reading in imitation in typically developing children

	0–9 months	12 months	14–15 months	18 months	24 months	36 months
Level of development	No understanding of intentions	Understanding ▪ Goal-directed action	Understanding ▪ Intentions (rational)	Understanding ▪ Unfulfilled intentions ▪ Reciprocal intentions ▪ Communicative intentions	Understanding ▪ Prior intentions ▪ Symbolic intentions	Understanding ▪ Hierarchy of goals
Form of imitation	Mimicking or emulation	Imitative learning	Imitative learning or goal emulation	Role reversal imitation	All	All

1. We now know much more about how children interpret the intention of a behavior, based on such things as direct behavioral cues (e.g., those indicating effort, failure, or accidents), immediate context (e.g., the presence or absence of a concrete, perceptible goal), and preceding behavioral and perceptual context (e.g., what the demonstrator has been doing just prior to the demonstration).
2. We now know of the theoretical possibility of goal emulation in which the observer understands the demonstration in terms of both its ends and means, but chooses its own means (Whiten & Ham, 1992). And we have begun to investigate some of the factors that determine which behavioral means a child will choose on a given occasion (e.g., Gergely et al., 2002).
3. We also now have begun to think about hierarchies of goals, and the implication of this for imitative learning (Bekkering et al., 2000). For example, is the adult turning on the light or turning on the light with her head? Determining whether a behavior or an aspect of behavior is irrelevant or relevant to a demonstrator's achieving a goal is often far from straightforward.
4. Finally, it seems that we are learning more about the imitative learning skills, and intention-reading skills, of children with autism, who may be less disadvantaged in this domain than previously believed.

We conclude with a possible avenue for theoretical progress in a field that has sometimes been mired in terminological disputes. Call and Carpenter (2002) note that there is just so far one can go by identifying and naming new social learning processes haphazardly. A much more systematic approach is to break down social learning into the sources of information at the observer's disposal, for example, the actions of the demonstrator, the changes of state in the environment (result), the demonstrator's goal or intention, and facts about the immediate context and the immediately preceding behavioral context that affect the observer's interpretations. By breaking down the information in this way, and possibly in other ways, we can identify more precisely exactly what children and other animals are learning when they learn to do what others do.[1]

1. See the comments on this chapter by Claxton (vol. 2, ch. 8.7, p. 194). ED.

6 On Learning What Not to Do: The Emergence of Selective Imitation in Tool Use by Young Children

Paul L. Harris and Stephen Want

6.1 Introduction

Tool use by human beings depends on a process of social transmission. Children rarely invent new tools; rather, they see others use them, and they do likewise. In the study of nonhuman primates, the study of imitative tool use has been a major focus of research. Surprisingly, however, imitative tool use by young children has rarely been studied. Instead, developmental psychologists have focused on the imitation of simple facial gestures in the first year of life (Want & Harris, 2001). In this chapter, we seek to remedy that neglect. We describe an investigation of imitative tool use by 2- and 3-year-old children. To set these studies against a larger backdrop, we first consider the history of human toolmaking and distinguish between two broad types of social transmission: nonselective and selective imitation.

6.2 A Brief History of Human Toolmaking

The earliest-known stone tools date from the beginning of the Paleolithic period, more than 2 million years ago. Oldowon stone tools (named after the archaeological site of Olduvai Gorge in Tanzania) have been found in South and East Africa and are associated with *Homo habilis*, who lived from about 2.2 to 1.6 million years ago. Such tools were "manufactured" in a relatively simple fashion by using one rock as a hammer to strike another, breaking it into a larger core and a smaller, detached fragment or flake. The exact purpose of these tools remains uncertain; they may have been used for nut-cracking, woodworking, the scraping of hides, or animal butchery (Schick & Toth, 1993). Approximately 1.4 million years ago, tools with greater standardization appear. These are circular or oval flakes that are shaped around their entire edge to form a bilaterally symmetrical handaxe. This form of tool persisted for a very long time—it was being made