15

The Evolutionary Ancestry of Our Knowledge of Tools: From Percepts To Concepts

Marc D. Hauser and Laurie R. Santos

> To primitive man each thing says what it is and what he ought to do with it… a fruit says 'Eat me'; water says 'Drink me'; thunder says 'Fear me'.
>
> (Koffka 1935, 7)
>
> A lot of stuff that's domain specific or species specific or both has to be innate in order that we should come to have the concept DOORKNOB (or for that matter, the concept RED)… The issue is whether it requires a lot of innate intentional stuff, a lot of innate stuff that has content.
>
> (Fodor 1998, 143)

It is hard to imagine two more polar extremes than Koffka and Fodor. An attempt to unite them would surely create one of the most heinous intellectual marriages of all time. Koffka, a leading proponent of Gestalt psychology, suggested that when an organism perceives the world, the world spits back affordances, properties of the environment that tell the organism what to do. Under this view, there are no mental representations, no concepts. Fodor, on the other hand, one of the central architects of the representational theory of the mind, has argued that concepts do all the interesting work. In fact, for Fodor, all concepts are innate and without them organisms wouldn't be able to learn. This is Fodor's strongest challenge to those interested in concept acquisition.

Our goal in this chapter is not to arbitrate between these views. Rather, we explore one particular corner of our conceptual world—the domain of tools—and assess how such knowledge developed both phylogenetically and ontogenetically. In particular, we hope to convince the reader of three things. First, the comparative approach to conceptual representation is crucial to any theory of concepts as it forces one to entertain how concepts are acquired and managed in the absence of language. Second, although a number of animals naturally manufacture and use tools, such behavior does not entail evidence for a concept of tool or artifact more generally. Rather, tool-use may grow out of a particularly good sense of the functionally relevant features *associated* with tools, a sense driven by affordances or perceptual categories as opposed to conceptual knowledge. This claim does not deny the possibility of concepts in animals. It does, however, raise the criterial bar used by many researchers to assess conceptual representations in animals, as well as in human infants and toddlers. Third, there are serious methodological problems underlying most current attempts to uncover the content of animal representations and concepts more specifically. Some studies use training approaches, while others explore spontaneous abilities; some test animals in the wild, some in captivity; some use tasks which require action, while others require only perception; some consider discrimination among category exemplars sufficient, while others demand evidence for theory-like organization of knowledge that is separate from the perceptual input that might be used for discrimination. Given such variation, it is not always easy to draw straightforward conclusions about what any particular animal does or does not know. We hope to clarify some of these general problems by assessing what animals know about tools, and in particular, whether there is any evidence that animals have a concept of tool; this specific discussion should have significant implications for how we evaluate the evidence for concepts in non- or pre-linguistic creatures more generally.

1. CARVING UP OBJECT KNOWLEDGE

The world is filled with objects. All animals make discriminations among some of them. The kinds of discriminations they make, and the sensory resources they recruit, are often different. Sometimes, differences in discrimination are due to the kinds of selection pressures that have operated to favor either coarse- or fine-grained analyses of object categories. For example, animals in the wild are frequently concerned with the threat of predation. At the simplest level, prey must discriminate between predators and non-predators. In some species, however, there has been strong selection for individuals to make even more fine-grained discriminations among different types of predators, a situation that has been carefully described in vervet monkeys living on the savannas of Kenya (Cheney and Seyfarth 1990; Struhsaker 1967). Because vervets are preyed upon by a wide variety of predatory species, each with a unique hunting style, they have evolved a predator-specific alarm response, including acoustically distinctive calls and escape responses. Thus, seeing a leopard (often) causes vervet monkeys to give a call that is different from the one they give to eagles and, respectively, causes them to run up into a tree as opposed to running under a bush. For vervets, therefore, there is both the superordinate category of PREDATOR and the subordinate-level categories of LEOPARD, EAGLE, and SNAKE.

Differences in discrimination are also mediated by the kinds of perceptual and cognitive abilities that each organism brings to the task of categorizing objects in the world. Dogs make distinctions between the urine markings of other dogs, while humans do not. Humans draw distinctions between filet mignon, sirloin, and hamburger, while dogs do not. These are descriptions of performance, of recognizing objects as members of particular categories. But what do they reveal about either the perceptual or conceptual capacities of these species? On the one hand, we know that at both central and peripheral levels, the olfactory system of a dog is able to make finer discriminations than the olfactory system of a human. Differences in performance are not due to differences in one species being smarter than the other, at least not in any interesting sense of the word 'smarter'. It also says nothing at all about conceptual organization. When it comes to pieces of beef, dogs may well be able to discriminate between filet mignon, sirloin, and hamburger, but it is unlikely that they have a conceptual representation like ours that places beef as one of several food categories of meats, that rank-orders each type of meat in terms of quality, that places different cuts of beef in relationship to other kinds of food with different values, that sees these as the kinds of distinctions a meat-eater would make but that a vegetarian would care little about, and so on. For dogs, we presume that beef is beef is beef. It is an object that, in Koffka's phrasing, cries out 'Eat me' without playing any significant role in the animal's conceptual representations. Dogs may prefer filet mignon over sirloin and sirloin over hamburger, showing the same kinds of preferences as a human connoisseur of beef. And they might make these distinctions by attending to particular features of each cut of beef, using both smell and sight. But using particular perceptual features to make such distinctions is not sufficient to show that an organism, human or non-human, has a concept. Concepts are not mere collections of features, although featural distinctions certainly play a role. What makes the conceptual distinction more interesting, at least from our perspective, is that particular tokens of a class are situated *in relationship* to other tokens, and the organization of tokens is mediated by a particular theory of how they cohere. Given this perspective, it therefore becomes imperative to assess both how concepts are acquired and how they are modified.

With respect to conceptual acquisition, we see the field dividing into three core groups (see Fig. 15.1), with lots of bleeding at the edges. Group 1, championed by Fodor, posits that concepts are largely innate, and that without this starting point there would be no learning. On this view, there is massive continuity between the concepts of a child and those of an adult. Note, importantly, that Fodor's view does not deny learning. Rather, his point is that core concepts are already part of the mind's system of core knowledge and that our experience with the world can shape how this core knowledge is reorganized. Group 2 posits few or no innate representations, and sees concept acquisition as tapping quite domain-general learning mechanisms (for reviews of the latest thinking, see Mandler, this volume). Quine (1960), as one proponent of this view, argues that children acquire concepts such as DOG and HAMMER by simply building up a set of associations based on relevant features. What makes the varieties of DOG tokens cohere is that individual learners build up a similarity space of features based on associations. A different view, which also denies the significance of mental representations, is represented by the Gestalt psychologists such as Koffka, as well as their successors such as the ecological psychologist Gibson (Gibson 1979; 1950). On this view, as discussed above, organisms make distinctions between objects based on the perceptual properties that they afford. Perceptual features of an object are therefore like instructions for action. A HAMMER cries out 'Pick me up, and strike something'. Group 3 posits that there are innate concepts that guide the acquisition process, with domain-specific as opposed to domain-general learning mechanisms at the core of the ontogenetic process. This view is held by a number of developmental psychologists such as Carey (1996), Keil (1994), and Spelke (1994), as well as several evolutionary psychologists, including Cosmides and Tooby (1994) and Pinker (1997). Under this view, there is a system of core knowledge divided into domains such as folk physics, psychology, and biology, that guides how we experience and digest the world taken in by our sensory systems. This core knowledge is theory-like in that it is based on explanatory principles that are separate from the raw perceptual input, including such notions as agency and intentionality.

Group 1 and the Quinean version of Group 2 have relatively few supporters with respect to those pursuing the empirical evidence for concepts; we note, however, that Fodor's challenge is extremely important, and essential to the empirical research program we will defend. There is a slightly larger contingent supporting the Gibsonian version of Group 2, with interesting modifications on this view, especially with respect to how perceptual analyses may provide the necessary structure for building concepts (Mandler 2000). We take Group 3 to be the most serious contender (i.e. it's the one we believe and will defend; see also the chapter in this volume by Mahon and Caramazza), although there are disagreements concerning several pieces of the argument. Rather than debate these points here, we turn next to the target ontological problem of this essay—tools—and briefly discuss some of the relevant theoretical distinctions, especially as they bear on domain-general versus specific learning mechanisms.

2. USING TOOLS AND CONCEPTUALIZING THEM

When early anthropologists heralded our uniqueness by proclaiming us 'Man the toolmaker', they were referring to the hand-axes and cutting tools of some of the first hominids. Although relatively simple objects, these tools stood out on the comparative landscape as no other animal had, apparently, ever created or used a tool in nature. There were, of course, studies such as Kohler's showing that chimpanzees in captivity could use tools, including the famous observation of Sultan piling up boxes so that he could stand on top and then, with a stick, dislodging a banana hanging from the ceiling. These observations at least suggested that animals such as chimpanzees have the capacity to see one or more objects as a means to some end. When Jane Goodall famously described wild chimpanzees using sticks, stripped of their leaves, to extract termites from a termite mound, man the toolmaker no longer stood out from the landscape. Moreover, Goodall's observations opened the floodgates, and soon dozens of observations of tool-use in the wild emerged, spanning a wide range of species including other apes, monkeys, non-primate mammals, and birds. One could no longer claim that tool-use and manufacture were uniquely human characteristics (Griffin 2002; Hauser 1988); see also Gould, this volume), and recent studies on the New Caledonia crow suggest that the highest level of sophistication may not even rest with the primates (Chappell and Kacelnik 2002; Hunt 1996, 2004; Hunt and Gray 2003).

Until the 1990s no study of tool-use had bothered to ask either how animals recognize and represent tools (as opposed to other objects), or how these cognitive processes limit the range and qualities of the tools invented and used. As we discuss in section 3, work in the last ten years has begun to remedy this problem. Here, however, we would first like to sketch some of the relevant theoretical issues (for similar discussions, see the chapters in this volume by Kelemen and Carey, Mandler, and Mahon and Caramazza) underlying an empirical inquiry into tool-use in the animal kingdom.

As an attempt to simplify the theoretical landscape, we will discuss five views about the nature of tool representations (Fig. 15.1). The first of these views—the *Innateness* view—stems directly from the Group 1, Arch-Nativist stance on concept acquisition. Under this view, our concept of tools emerges in the same way as our other concepts: they are genetically endowed, appearing in the absence of experience or formal tutelage. If this characterization is correct, there should be continuity between the concepts of adults and infants, and this includes concepts of tools.

The next two hypotheses fall out of the Group 2 empiricist view of concept acquisition. The first is what we refer to as the *Affordances* perspective. Under this theoretical stance, humans represent artifacts only on the basis of their physical features. The properties of a hammer, its graspable shape and hard, pounding edge, are taken in by our perceptual systems and simply cry out for the action of hammering. Under this view, however, there is no HAMMER representation, no ontological category of an object with a particular shape that is used for a particular function that tends to be found in a toolbox, and certainly no organizing theoretical framework for organizing HAMMER into the more general concept of ARTIFACT. Our representation of a hammer-like object simply consists of a mapping between certain perceptual features (e.g. smoothness, hardness) and certain functional possibilities (e.g. graspability, poundability). As such, this view has a weak ontological structure.

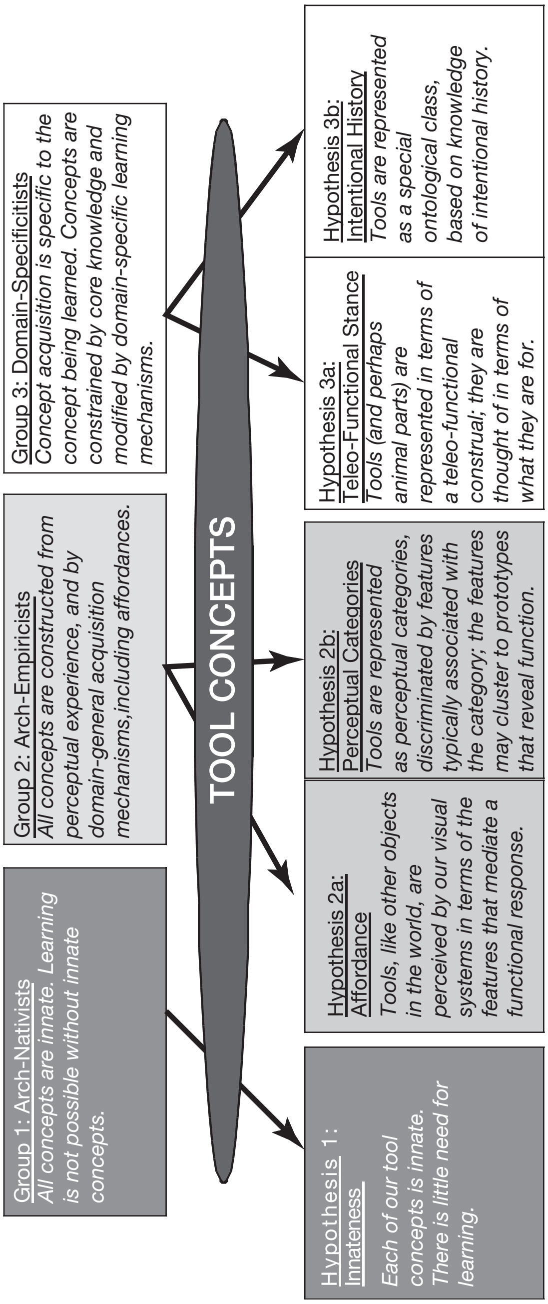

Figure 15.1. A rough guide to conceptual theories of tools

The other empiricist hypothesis of tools, what we call the *Perceptual Category*, is an ontologically richer proposal. Under this view, we possess ontological categories like HAMMER, CHAIR, and CLOCK that are distinguished from each other on the basis of different features. A HAMMER is distinguished from a CHAIR because the former has particular articulated parts (a rigid handle, a hard flat striking edge, etc.), whereas the latter does not. However, this difference is not definitional. Clearly, a hammer with a folding handle that locks into position before you strike something would still be a HAMMER. The prototype view, one variant of the Perceptual Category perspective, suggests that organisms represent different kinds of artifacts as constellations of features; an object with more hammer-like features will be considered a better HAMMER than objects with less hammer-like features. What is important about features from the perspective of distinguishing one object kind from another, however, is that some features are more salient when categorizing artifacts than others. For example, artifacts can readily change some of their perceptible properties without changing kinds; the hammer maintains its hammer-ness whether it is blue, red, green, or rainbow-colored. It can, however, lose its hammer-ness if it is a different shape. Similarly, a hammer must at least be harder than the object one wishes to strike, usually a nail. Clowns in a slapstick routine will, of course, use foam-rubber hammers in order to hit someone over the head. But this routine makes our point: it is because we expect a hammer to be harder than the object it strikes that we find humor in a clown striking a human over the head.

The last two hypotheses fit with a larger domain-specificity proposal. Again, the domain-specificity approach to conceptual knowledge posits that organisms are endowed with innate core knowledge that constrains how they divide up the world into ontological chunks and how experience modifies their ontologies. Each conceptual chunk or domain tends to reflect computational problems that are ecologically salient for a particular organism, both over developmental and evolutionary time. As many researchers have argued, a salient domain in human ecology is that of artifacts. As any new parent can attest, human infants are born into a world of artifacts that they quickly learn to use, categorize, and name. For these reasons, domain-specificitists tend to agree that humans possess specialized mechanisms for representing and categorizing objects in the artifactual domain. What these mechanisms are and how they operate, however, is open to much speculation.

One way in which artifacts are distinguished from many, if not most, other kinds of objects is that they have a particular function. When we see a tool, even if we have never used it before, we tend to think or ask 'What is it for?', even though we would never formulate a similar question for a rock. This type of question is reserved for certain types of entities—those with the characteristics of functional design. The *Teleo-Functional* hypothesis proposes that humans reason about tools in line with just such a teleological mode of construal. This hypothesis was originally put forward by Keil (1994), and has been extended both theoretically and empirically by Kelemen (1999*c*). Importantly, the Teleo-Functional view of artifacts is domain-specific—for adults, it tends to apply only to objects that are created for a specific purpose. Interestingly, it is a view that is part of theology and evolutionary biology. Theologians' answer to the question of design is, clearly, some divine entity. Evolutionists' answer is natural selection; as Darwin articulated, natural selection is the only evolutionary process that can create non-random design features, suited to social and ecological problems. When it comes to artifacts, their design is solely due to the mind of their creator. Of course, it is also possible for an object to be used for a particular function even though it was not designed for this function; a shoe can be used as a hammer and a clock can be used as a paperweight. Nevertheless, it is important to note that a shoe used as a hammer does not necessarily become a HAMMER. Instead, it remains a SHOE because it was designed for the purpose of being used as a shoe.

The importance of a tool's history and intended design leads to the *Intentional History* perspective. This hypothesis, originally championed by Bloom (1996), builds on an argument originally developed by Levinson with respect to art. For Levinson (see this volume), the only way to understand art is by understanding the historical-intentional aspects of the product. In contrast to RED or LEMON, art can only be evaluated based on the intention of the artist (why she painted or sculpted the piece) at the particular time she created the piece. If it was intended as art then, it should be considered as art today, even if we don't like it. Bloom argues, symmetrically, that all artifacts are represented in terms of their intentional history—the reason why they were made and what they were made to do. Thus, we call something a HAMMER if its creator designed it with the intention of hammering things. An important consequence of this view is that even if the end-product fails to satisfy the designer's original intent (e.g. a broken hammer), we can nonetheless call it a HAMMER, albeit a bad one. Importantly for this view, one should not only be able to extract the intentional history of an artifact by seeing the artifact made or being told about it with language, but also by inferring intent from design. Thus, it is not the case that surface design is irrelevant in distinguishing hammers from screwdrivers. Rather, what their design invokes in human minds (at least adults) is an inference about what an intentional designer initially created them for.

If our rough sketch of the theories of tools, and artifacts more generally, is reasonable, then a research program designed to uncover how this conceptual domain evolved and develops must contend with these different positions. Critically, we believe, in thinking about either tool-use or creation in non-linguistic organisms (animals and human infants), it is essential to distinguish between recognition and perceptual discrimination on the one hand, and conceptual knowledge on the other. Although animals, for example, make and use a range of artifacts (e.g. tools, nests, dams), and discriminate among them using salient perceptual features, other evidence is required to show that these animals have a domain-specific theory about why tools are different from animals or trees, and that this theory guides conceptual change. To be clear, if animals lack an understanding of teleology or intentional design, then this does not take away from the significance of their tool-using abilities. It does, however, mean that they may represent tools in a fundamentally different way than human adults do.

We turn next to a review of some of the relevant literature on non-human primate tool-use, with a specific focus on both methodological issues and theoretical issues that relate to the hypotheses reviewed above and sketched in Fig. 15.1. Concerning the latter, we are particularly interested in exploring the extent to which tool-use relies on domain-specific as opposed to domain-general mechanisms. Consequently, we focus on the use of tools both by natural tool-users and by species that never spontaneously use tools in nature or in captivity. If performance on tool-related tasks depends on evolved specializations for tool use, then natural tool users should show greater proclivities than species that never spontaneously use tools.

3. HOW APES AND MONKEYS THINK ABOUT TOOLS

3.1. Wild Chimpanzees

Any review of primate tool-use should begin with the most prolific non-human tool-user, the chimpanzee (Matsuzawa 1996; Matsuzawa and Yamakoshi 1996; McGrew 1992; Tomasello and Call 1997; Whiten *et al.* 1999). Chimpanzees use a variety of tools in a number of different situations, including twig sticks to probe termite mounds, stone hammers to crack nuts, crumpled leaves to sponge water, and bark sandals to climb up prickly trees (Alp 1997; Boesch and Boesch-Achermann 2000; Goodall 1986; Matsuzawa and Yamakoshi 1996; McGrew 1992). There is also much evidence to suggest that chimpanzees modify objects to create better tools; chimpanzees are known to chew leaves for better absorbency (Sugiyama 1995) and to bite the leaves off branches to allow them to fit in termite mounds more efficiently.

A number of recent observations suggest that in addition to using tools adeptly, chimpanzees may possess a sophisticated understanding of the functional properties of the objects they use. Studies by Boesch and colleagues (Boesch and Boesch-Achermann 2000) in Tai Forest, Ivory Coast, together with studies by Matsuzawa and colleagues (Matsuzawa and Yamakoshi 1996) in Bossou, Guinea, provide exquisitely detailed observations of nut-cracking behavior in chimpanzees. The Tai chimpanzees eat five different nut species, none of which can be opened with the bare hand. To crack them, chimpanzees secure the nuts using anvils (made of roots, branches, or rocks) and hit them with hammers (made of branches or stones). Boesch and colleagues suggest, based on their observations, that chimpanzees understand the properties that are necessary for an effective hammer. In particular, chimpanzees carry heavy stones further than wooden clubs (Boesch and Boesch 1984), suggesting that they understand the importance of the tool's material property; heavier hammers are more effective than lighter ones, and thus they are willing to carry them further. Matsuzawa's observations generally confirm the reports from Tai, but also demonstrate that chimpanzees appear to be sensitive to the physics of the tool-using task. In particular, and in contrast to juvenile chimpanzees, adults are sensitive to the stability of the anvil and its capacity to hold, in a stable position, the target nut. When adult chimpanzees find a suitable rock, one with a relatively flat surface or a depression that will envelop the nut, they assess whether the rock is level with the ground. If it is slanted, they will place a second rock under the first, in an attempt to provide a more level striking surface. Juveniles often place nuts on slanted surfaces, thereby leading to numerous failed attempts to crack the nut as it rolls off the anvil before they can strike. Suzuki *et al.* (1995) argue that chimpanzees in the Ndoki forest of Congo may possess a similar understanding of their probing tools. They report that chimpanzees use two different types of sticks while termite fishing: a perforating stick, a wide twig used for making deep holes into the termite nest, and a fishing probe, a thin stick, often with chewed-off, brush-like ends, that is used to extract the termites. That chimpanzees use two physically different objects as functionally different tools suggests that they are identifying different properties of objects and using those properties to choose which tools will be more effective for a given job.

The problem with the findings reviewed thus far is that we don't know how these chimpanzees developed an understanding of which tools are better than others. One possibility is that individual chimpanzees learn which tools are most effective through a process of trial and error. A Tai chimpanzee may try out many different kinds of objects as hammers, only to gradually learn that stones crack nuts most efficiently. If this is true, then an experienced chimpanzee only knows to carry a heavy stone further because his experiences with previous heavy stones taught him that these objects will crack nuts most effectively. Alternatively, chimpanzees may possess a richer understanding of the problem and may predict which features are important for a functional tool without trial-and-error learning. If this alternative is correct, then even inexperienced chimpanzees might know to carry stones long distances, having the foresight that stones are heavy and able to crack things better than softer objects like branches. Unfortunately, however, without data on an individual's previous experience with using a given tool, we cannot distinguish between these alternatives. And this point holds even when one considers the important developmental data collected by Matsuzawa and others, showing that it takes between eight and nine years before chimpanzees develop functionally appropriate, and problem-specific, tool-using techniques. Young animals could show sensitivity to some features but not others. These developmental data are thus mute with respect to this point.

3.2. Chimpanzees in the Laboratory

Faced with the problems of assessing the depth of a wild non-human primate's understanding of tools, many researchers have taken to investigating what captive primates know about tools. In captivity, as opposed to the wild, experimenters can selectively vary an individual's experiences with an object and document the process by which an individual comes to learn about a particular object and its function. In a recent book, Povinelli (2000) presents the most recent of such studies, focusing on a group of chimpanzees. Across a wide variety of paradigms and conceptual problems, Povinelli consistently finds that chimpanzees lack a sophisticated understanding of the physics of tool-use, focusing instead on perceptual features that in some cases provide correct answers to the task at hand, but then fail to generalize to the more common, and functionally relevant, components of the problem. Here we discuss a few examples to highlight the relevant dimensions of Povinelli's approach, and the importance of distinguishing between having conceptual knowledge of tools that is sensitive to functional design features, and using perceptual features that are only sometimes predictive (associated with) of success.

Povinelli and colleagues tested seven chimpanzees. Subjects started the experiments as young juveniles and ended as sub-adults. Because different experiments were run at different ages, it is not possible to assess, with complete accuracy, whether differences in performance across experiments is due to differences in age, and whether performance on any given experiment might have been different with older chimpanzees; there are only a few cases where a young chimpanzee was re-tested on the same experiment at an older age. We mention this point here both because of the developmental studies on tamarins described below, and because of Matsuzawa's field studies which show that competent tool-use emerges after the age of 9–10 years old. Many of Povinelli's experiments were run on far younger chimpanzees.

In one experiment, Povinelli and colleagues examined whether or not chimpanzees understand that a tool's material often affects its function. They trained chimpanzees on a task in which subjects were allowed to use one of two T-shaped pulling tools to obtain an out-of-reach food reward (for similar tasks, see Brown 1990; Hauser 1997). During training, these pulling tools were made out of rigid materials (PVC tubing and plywood). Once subjects mastered the pulling task, however, they were presented with trials in which the top of one of the two tools was changed. The top of this new tool was made out of a flimsy material (a thin strip of rubber), and thus was no longer able to bring the food within reach. Povinelli and colleagues found that only one of their seven chimpanzees succeeded on the flimsy tool problem, consistently choosing the rigid tool over the flimsy tool during the test trials. The authors conclude that chimpanzees fail to spontaneously take into account a pulling tool's material when choosing between two potential tools.

Povinelli and colleagues also examined what these chimpanzees understand about a tool's three-dimensional orientation. They presented the same chimpanzees with a choice between two rake-shaped pulling tools. When these tools were oriented with their tines upward and their bases placed flat on the tray, they served as functional pulling tools; a flat base efficiently pulled the food within reach. However, when the rake was oriented with its tines facing downward, it no longer served as an effective pulling tool; with the tines down, pieces of food readily slipped under the base of the tool and thus could not be retrieved. Povinelli and colleagues found that chimpanzees did not distinguish between these two orientations. Chimpanzees were as likely to choose rakes oriented with tines up as rakes oriented with tines down. Chimpanzees also failed to attend to the substrate on which tools operated, pulling a tool that caused the reward to drop into a trap as often as a tool that brought the food within reach (Povinelli 2000, ch. 15.5). Povinelli's work suggests that, despite their skillful use of tools in the wild, chimpanzees do not understand the physics or functionality of tools, at least under the conditions tested.

3.3. Capuchins

Like chimpanzees, capuchins spontaneously use tools in the wild and in captivity (Phillips 1998; Visalberghi 1990; Westergaard *et al.* 1998). Visalberghi and her colleagues should be credited with the first systematic attempt to investigate what captive primates understand about the functional properties of different objects, using experiments on captive capuchins. Visalberghi and Trinca (1989) presented four tufted capuchins (*Cebus apella*) with a task in which a piece of food was placed out of reach inside a clear tube. Capuchins spontaneously solved the task by inserting a stick into the tube to obtain the food. Visalberghi and Trinca then explored what these subjects learned about the important aspects of their pushing tool by changing the properties of the available sticks. For example, they presented subjects with a bundle of sticks that was too wide to fit inside the tube and a series of broken sticks which together could be pushed inside the tube to obtain the reward. Although all subjects eventually solved the task with these new tools, none of the subjects solved the task spontaneously. Instead, subjects seemed to discover the correct solution only through a long series of trial-and-error attempts with different tools. Similarly, the capuchins failed to recognize the important aspects of the substrate on which their pushing tool operated. Specifically, Visalberghi and colleagues (1995) presented subjects with a modified version of the trap task in which the trajectory of the food reward was impeded by the trap in the bottom of the tube; if the food reward was pushed over this trap, it would fall inside and become unattainable. Only one of four capuchins was able to solve this trap problem, but only through a series of trial-and-error attempts. Subsequent transfer tests revealed that this subject failed to understand the trap:

she persisted with the same strategy even when the trap was positioned at the top of the tube, a position that had no impact on the movement of the food reward through the tube. These results, and others presented by Visalberghi and her colleagues, lead to the conclusion that although capuchins naturally use a variety of tools, they seem to lack an understanding of the functional properties of these tools. Like chimpanzees, they lack an understanding of the relevant physics.

3.4. Non-Tool-User Tool Use

In contrast to the approach taken with captive capuchins and chimpanzees, we have focused on one species—the cotton-top tamarin—that never spontaneously uses tools in the wild or in captivity (Hauser 1997; Hauser, Kralik, and Botto-Mahan 1999; Hauser, Pearson, and Seelig 2002; Hauser *et al.* 2002; Santos, Miller, and Hauser 2003 Santos *et al.* 2006; Santos *et al.* 2005). As mentioned above, the primary goal of these studies has been to examine whether tool users are equipped with cognitive specializations for tool-use that non-tool-users lack. Said differently, can animals such as tamarins recognize the functionally relevant features of tools even if they never use such objects in captivity or in the wild?

In the first series of experiments, Hauser (1997) presented adult cotton-top tamarins with a task in which subjects had to pull one of two tools in order to obtain an out-of-reach food reward; this task was modeled on studies of human infants by Brown (1990). Subjects were initially presented with a tray holding two blue canes positioned near a small food reward. One of the two canes was hooked around the food reward, the other placed with the food reward outside the hook. With this set-up, only one cane was effective in pulling the out-of-reach food reward. Results showed that subjects with no prior experience with tools quickly learned to pull the correctly positioned cane. After subjects learned the initial task, Hauser tested subjects with a series of new canes that differed from the original on only a single dimension. For example, an experimenter presented subjects with a choice between a tool with a novel shape (e.g. a right-angled L-shape) and a tool with a novel texture (e.g. a bumpy cane). Similarly, subjects could choose between a tool with a new color (e.g. a red cane) and a tool with a new size (e.g. a cane that was twice as long). If tamarins rank-order the features with respect to their impact on the tool's functionality—how they affect the physics—then they should accept changes in color and texture, as these play no role in the tool's capacity to bring food forward; in contrast, changes in size and shape directly influence functionality, especially given the tamarins' dexterity, and capacity for altering the tool's properties. Subjects reliably chose differently colored and textured canes over differently shaped and sized canes, and did so on the first trial. These results suggest that tamarins spontaneously regard changes of a tool's form (its shape and size) to be more important to its function than changes to its surface features (color or texture).

• Q1 **Figure 15.2.** Photographs of two non-tool-users—a cotton-top tamarin (left) and a ring-tailed lemur (right)—using simple pulling tools

Similar patterns emerged in a more recent series of experiments aimed at investigating what tamarins understand about a tool's material (Santos *et al.*, 2006). In this experiment, a different group of inexperienced tamarins were given a choice between a tool with a different color (e.g. a pink cane) and a tool made of a different material (e.g. a flimsy piece of blue yarn shaped in a cane shape). Subjects rejected canes made of yarn, indicating that even in the absence of direct experience (i.e. given the nature of the task, subjects evaluated their options on each trial based on visual information, and then selected a tool; under these conditions, they lacked the opportunity to try out tools), tamarins understand rigidity to be an important property of a pulling tool.

To further pursue the importance of experience, Hauser, Pearson, and seelig (2002) extended the work on adult tamarins to infants. Thjey found that infant tamarins, with no experimental experience of any kind, showed the same pattern of results as adults. Remarkably, in the absence of any relevant experience (e.g. there were no opportunities to manipulate freely moving objects), tamarins as young as 4 months understand which properties are functionally relevant to the effectiveness of a pulling tool.

After presenting subjects with these single-dimension changes, Hauser and colleagues went on to present adult and infant subjects with tools that differed from the original along many dimensions (e.g. a green V-shaped tool with large bumps; Hauser 1997; Hauser, Pearson, *et al.* 2002). Subjects received a final condition in which these novel tools were pitted against the original tool positioned in an incorrect orientation. Both infant and adult tamarins spontaneously chose the novel but functionally correct tool over the familiar yet now incorrect tool, suggesting that tamarins of all ages understand that orientation is an important aspect of a successful pulling tool. Hauser, Kralik, and Botto-Mahan (1999) reported similar success using a cloth pulling tool, one in which the functional or physical task was support (i.e. the food reward was either on the cloth or off it) as opposed to containment in the canes task; across a variety of manipulations to the size, shape, color, and material of the cloth, subjects invariably picked the cloth that provided continuous access to the reward; these preferences often emerged on the first trial of a new condition involving a new featural transformation.

Taken together, these findings suggest that tamarins, even infants with no task-relevant experience, take into account the relevant features of a pulling tool (e.g. shape, size, material, and orientation), and also disregard those features that are irrelevant for effective pulling (e.g. color, texture). Nevertheless, more recent evidence suggests that although they understand the relevant features of pulling tools, tamarins may lack a more sophisticated understanding of the pulling task.

We (Santos, Pearson, *et al.* 2006) have explored whether or not tamarins are able to solve the more difficult types of problems that chimpanzees and capuchins cannot solve. Specifically, they presented tamarins with a trap task like the one used to test chimpanzees and capuchins, and then examined whether or not tamarins spontaneously attend to the position of the trap. Curiously, tamarins were as likely to pull a cane positioned over a trap as they were to pull a cane over a continuous surface. Like chimpanzees and capuchins, tamarins fail to understand the trap. Similarly, tamarins also fail a task involving rake-shaped tools. When an experimenter presented tamarins with a choice between one rake with tines up and one rake with tines down, subjects showed no preferences, even though the rake with tines up was functionally more efficient in bringing food forward. Like chimpanzees, the tamarins neglected to focus on this more subtle aspect of the pulling tool's orientation.

We have also recently completed a parallel sets of studies with other non-tool-using primates—vervet monkeys (Santos, Pearson *et al.*, 2006), rhesus macaques (Santos *et al.*, 2003), and lemurs (Santos, Mahajan, and Barnes, 2006); these species rarely, if ever, use tools in the wild or in captivity (e.g., Hauser 1988). In one study, we showed free-ranging rhesus macaques a purple L-shaped tool sitting on a flat stage with a ramp. We then habituated subjects to a human experimenter using the L-shaped tool to push a grape down the ramp. After habituation, we changed the properties of the tool. In one test condition we changed the color of the tool, showing subjects a pink L-shaped tool. In the other condition, we changed the shape of the tool, showing subjects a purple I-shaped tool; this newly shaped cane lacked the base necessary for pushing the grape. We reasoned that if subjects perceived the change as relevant to the tool's function then they should dishabituate to the test trial, evidenced by an increase in looking time. Results showed that rhesus dishabituated only to the new shape condition. Like tamarins tested on a different paradigm, rhesus macaques spontaneously pay more attention to a tool's shape than its color.

We have now also tested vervet monkeys (Santos, Pearson *et al.*, 2006) and lemurs (Santos, Mahajan and Barnes 2006) on the original canes task (Hauser 1997). Like tamarins, both vervets and lemurs show the same kinds of sensitivities to functionally relevant and irrelevant features. Thus, they attend to such features as size and orientation, while ignoring such features as color and texture. In contrast with tamarins, as well as chimpanzees and capuchins, vervets tested on a pulling/containment problem appeared more sensitive to other functional properties of the tool task. Specifically, a significant number of the vervets tested picked (1) a cane resting on a continuous surface over a cane resting on a trap; (2) a rake with tines up over a rake with tines down; and (3) a cane with the tip part of the hook broken (functionally irrelevant) over a cane with a break between the stem and hook (functionally relevant, as pulling the stem failed to advance the hook containing the food reward). These results are, of course, puzzling as they suggest that a non-tool-user—the vervet monkey—has a more sophisticated understanding of tools than two natural tool-users—chimpanzees and capuchins. Because there were some differences in the design of these experiments, as well as differences in the experimental histories of the test subjects, it is not possible at present to distinguish between significant differences among these species in their representation of tools and differences in testing procedures. Even with the reported differences, it must be further acknowledged that although the vervets performed above chance on these tasks, they made errors on approximately 30 per cent of all trials; whether this error rate is due to a relatively weak understanding of these functional problems or to something else (attention, motivation) is currently unclear.

4. HOW TO THINK ABOUT PRIMATES THINKING ABOUT TOOLS

In the last section we reviewed what non-human primates—both tool-users and non-tool-users—understand about tools based on observations of tool-use in five different primate species. Now we turn to the more pressing and difficult issue of how these animals represent tools. It is clear from the data reviewed that non-human primates represent tools in a fundamentally different way than they represent other kinds of objects. We base this claim on the observation that when monkeys and apes use an object as a tool they use different featural criteria than when they use an object as food or a landmark. As described above, a number of primates—tamarins, rhesus macaques, vervets, and lemurs—recognize that shape, size, material, and orientation are relevant featural dimensions for a functional tool, while color and texture are not. In contrast, when tamarins classify an object as a landmark, they use color and shape as relevant dimensions, but ignore orientation (deIpolyi, Santos and Hauser, 2001). Similarly, rhesus macaques, who also consider shape more important than color when reasoning about tools (Santos, Miller and Hauser 2003), nevertheless pay attention to a different set of features when reasoning about food objects. When categorizing novel foods, rhesus from the same population pay more attention to surface features like color than form features like shape (Santos, Hauser, and Spelke 2001). More importantly, these results hold true for infant macaques who, though lacking experience with solid foods, also recognize that color is more important than shape for edible objects. At a general level, these findings lend support to the domain-specific view of concept acquisition. Even at very early ages and in the absence of task-relevant experience (Hauser, Pearson, and Seelig 2002; Santos *et al.* 2002), non-human primates seem to parse the objects in their world into meaningful global categories—tools, foods, landmarks, and animals (Santos, Hauser, and Spelke 2002). Such evidence suggests that non-human primates may have innate biases to interpret their world in domain-specific ways. In addition, non-human primates seem to reason about different domains in ways that make ontological sense; their recognition of which features matter for different domains seem to map onto those of conceptually sophisticated human adults. For these reasons, we side with the domain-specificists and argue that both tool-using and non-tool-using primates are biased to distinguish tool-like objects from other ontological categories, and that these biases facilitate experience-based learning about different kinds.

Although it is clear that many primate species distinguish tools from other kinds of things, we are still left with the question of how primates represent objects within the category of tools and, perhaps more importantly, whether or not non-human primates represent *tools as* TOOLS. The fact that tamarins, a species that never spontaneously uses tools in the wild or in captivity, categorize artifacts using the same featural criteria as do tool-users, and do so as experientially naive infants, suggests an innate bias, but one that can not be interpreted as an innate concept of TOOL, *sensu strictu*. Given that tamarins do not naturally use tools, they simply can't have an innate understanding of tools. These data raise several significant methodological and theoretical issues. Methodologically, we argue that a description of featural biases is no longer sufficient for documenting how animals represent the tools that they use. Showing which features are relevant and which irrelevant is a first pass, but only that. Other experiments are necessary to show how non-linguistic animals think about tools. This methodological point leads to two theoretical points. If tamarins do not have a concept of TOOL, then how do they represent the objects they use in means–end tasks? And if it is not a concept of TOOL, then why, giving their lack of interest in tools under normal conditions, do they show such exquisite sensitivity to the functionally relevant and irrelevant features of tools when tested in the lab? At present, we don't have a satisfactory answer to the second question. Tamarins must face means–end problems in the wild: problems of connectivity in the canopy, problems involving the recognition of interconnected tree branches that must be used, and possibly tugged on, to get from one branch to another. These physical problems, involving means–end relationships, may fuel their capacity in captivity, tapping something like analogical reasoning: tree branch is to distant fruit as cane is to marshmallow. We leave this as a testable hypothesis, and turn next to the more general issue of how non-human primates represent tools, and how such representations may or may not differ from humans.

Another question that remains unanswered is whether tamarins represent objects in the domain of tools in the same way as conceptually sophisticated adult humans. Consider again the views of human tool representation that we put forth in Fig. 15.1. According to the Affordances view, for example, humans begin representing tools via their salient perceptual properties. Chair properties cry out to be sat on, whereas clock properties do not. As explained, this view of concept learning is ontologically weak; there are no theories of tools, how they function or what they are designed to do. Early concept learners simply perceive the salient affordances of different tool-like objects and use them accordingly. The current data on non-human animals is consistent with this type of categorization. From a very early age, tamarins recognize that a hard, cane-shaped object can be used as a pulling device. As such, they too seem to perceive tools in terms of salient, action-oriented properties like affordances.

Our data, however, are also consistent with the Perceptual Categories view. This perspective predicts that organisms learn about and classify things in the world on the basis of their perceptual features. A chair is considered a CHAIR because it possesses the type of features that CHAIRS tend to have: four legs, a flat seat, a back, and so on. These features are not definitional, but rather serve as a guide to classifying objects into particular categories. More importantly, however, the Perceptual Categories view is ontologically richer than the Affordances view. An animal with a capacity to form perceptual categories represents a tool as more than just the sum of its relevant features. A particular category of tool or artifact is represented as a kind. There are no sit-on-able objects and hammer-able objects, but CHAIRS and HAMMERS, meaningful ontological entities. Do non-human animals also represent tools as kinds? As reviewed earlier, non-human primates do take into account the types of features that a particular tool has. More importantly, they do not take into account any and all features; instead, they focus on only the particular features that are relevant for what a tool does (Hauser 1997; Hauser, Kralik, and Botto-Mahan 1999). This detail is important, as it implies that non-human primates might be doing more than just paying attention to features; they may instead understand tools in terms of their function, and therefore in terms of a rich theory of different kinds of things (Keil 1990, 1994).

This point brings us to the last two theories of tool concepts—those consistent with the domain-specificity position. In contrast to the domain-general empiricist position embodied in both the Affordances and Perceptual Category views, our tamarin data suggest that non-human primates may represent tool objects differently than other ontological kinds (e.g. landmarks), thereby implying that domain-specific learning mechanisms play a significant role in the acquisition of a tool concept. Unfortunately, however, the data we reviewed are silent on the question of whether tamarins and other primates reason about either the function of or intentions behind the objects they use in these tasks. For example, the data we reviewed could also be construed as consistent with the Teleo-Functional Concepts position. This perspective proposes that we represent tools as objects designed for a particular purpose. A chair is a CHAIR because it exists for the purpose of sitting. The Teleo-Functional perspective further posits that when we look at a chair, we infer the purpose of the object from the complexity of its features. If an object has many chair-like features, we assume that it is *for* a *CHAIR*-like purpose. Therefore, a teleological view of primate tool representations predicts that they should pay attention to the features of a tool that are relevant to its function, not because these properties are necessary and sufficient for their representation of the tool, but because these properties provide clues to the tool's purpose. Non-human primates clearly recognize the important features of tools. However, to date, the relevant data have been silent on whether or not primates use these features to infer some teleological origin.

A similar problem holds for the Intentional History position. Under this view, the intention of the designer determines its ontological kind. A chair is a CHAIR because it was made by a person who had the intention of turning several slabs of wood into a CHAIR. From this perspective, we would therefore expect primates to pay particular attention to a tool's causally relevant features because these features often serve as clues to the intentions of a designer. If an object has many chair-like features, we assume that the designer intended it to be a CHAIR. From the perspective of the Intentional History position, therefore, primates should pay attention to the features of a tool that are relevant to its function. This attentional focus should not arise because these are the necessary and sufficient properties for representing of the tool, but because these properties provide clues about the designer's intentions. The data reviewed above are again silent on the question of whether non-human primates use these features to infer the intentions of a designer; no study has yet examined whether primates (or other animals) draw inferences about the intentions of a tool- or artifact-maker. Fortuantely, a growing body of work in the field of primate social congnition suggests that it may be feasible to test the intentional history hypothesis. In contrast to what primate researchers thought a few years ago, primates (and other animals) appear to be capable of inferring the intentions of others under some conditions (for reviews, see Lyons and Santos 2006; Tomasello *et al.* 2005). Chimpanzees, for example, readily distinguish between the accidental and intentional actions of others (Call *et al.* 2004), particularly in tasks that involve competition (see Hare 2001, Lyons and Santos 2006). What is not yet known is whether this capacity translates to non-social domains, including especially the domain of artifacts, and whether primates use their intention-reading skills to attribute to others *an intent to create particular kinds of objects*, as the intentional history position might suggest. If primates do not use their intentional understanding when reasoning about tools, then this missing ingredient would make their use and thinking about tools quite different from our own. Taken together, the data reviewed here present a bit of a puzzle. Non-human primates seem to reason about tools as a specific domain of objects, but to date there is little direct evidence that they represent different tools as meaningful ontological entities. This lack of evidence is due in part to the fact that the ability to understand kinds is often defined by the capacity to label different types of objects with words. Since non-human primates lack the ability to use words, it is difficult to assess whether or not they have kind representations.

One move in the direction of establishing kinds in the absence of language comes from studies of property–kind individuation in human infants and non-human primates. The basic question is how and when infants come to appreciate that there are different kinds of objects and that such kind distinctions play a role in object individuation? To explore this question, Xu, Carey, and their colleagues (Van de Walle *et al.* 2000; Xu and Carey 1996, 2000; Xu, Carey, and Welch 1999) have run a series of experiments with infants between the ages of 10 and 12 months; the methods entail both expectancy-violation looking-time procedures as well as search assays. In one condition, an experimenter presented an infant with an empty stage and then lowered a screen. Next, the infant watched as a toy duck emerged and then returned on one side of the screen, followed by a truck emerging and returning behind the opposite side of the screen. The experimenter then lowered the screen, allowing the infant to look at a stage with a duck and a truck, or a stage with just the duck or just the truck. If infants use property–kind differences to individuate the number of objects present, then they should look longer at the outcome of one object than at the outcome of two. Ten-month-olds look as long at outcomes of one and two objects, whereas 12-month-olds look longer at outcomes of one object. Ten-month-olds apparently fail to use property–kind distinctions to individuate the number of objects hidden behind the screen, while 12-month-olds use such distinctions. Importantly, especially for the Xu and Carey position, it is infants with comprehension of words for such objects (e.g. 'duck', 'truck') that show the differences in looking times. This leads to the hypothesis that language may be necessary for the acquisition of kind concepts. Although this hypothesis has been challenged by data from Needham and Baillargeon (2000) showing that younger infants can make these distinctions, but in the absence of language, an even stronger test would come from studies of non-linguistic primates. Using both the looking-time method (Uller, Carey, and Hauser 1997; Munakata *et al* 2001), as well as the search procedure (Phillips and Santos, in press; Santos *et al.*2002), we have now demonstrated parallel effects with rhesus monkeys. For example, using the expectancy-violation procedure with a piece of carrot and squash, rhesus monkeys look longer when shown an outcome of one object than when shown an outcome of two(Uller *et al* 1997) Thus, language is clearly not necessary for kind distinctions of the kind described by Xu and Carey, although language may well play a role in changing the kind distinction. It is now essential that comparable studies be run with objects from other domains (i.e. not only food, but artifacts including tools) and other species, testing whether featural changes in some domains play a more significant role in individuation than do other features.

We have argued that non-human primates possess a domain-specific understanding of tools. They seem to understand which features are relevant to a tool's ability to function, and recognize that these features are different from those that are relevant for other objects including landmarks and food. Unfortunately, however, the data we have reviewed cannot determine the precise content of primate representations of tools. Consequently, we cannot yet say whether monkeys and apes simply use the affordances of tool-like objects (no representations at all), represent tools by highly associative perceptual features, or think about tools in a theory-laden fashion, considering either the function of a given artifact or the intentionality of the tool-maker. Nevertheless, the reader should certainly not come away feeling cynical. Over the past ten years primate researchers have gone beyond merely observing tool-use and have gained a much richer understanding of the way primates actually categorize the tools that they use. Such work has pushed us much closer to figuring out the representations that underlie primate tool-use. Perhaps more significantly, however, we have developed a suite of new empirical techniques that can be used to examine how primates represent both concrete objects such as tools, food, and landmarks as well as more abstract concepts such as number and mental states (Gallistel 1990; Hauser 2000; Heyes and Huber 2000; Shettleworth 1998; Tomasello and Call 1997; Tomasello *et al.* 2005)). Some of these techniques, which directly parallel those used with human infants (Hauser and Carey 1998), require no training, can be used across different ages and species, and allow for direct comparisons with data on human conceptual acquisition. Such techniques promise a bright future for work investigating concepts in non-human primates and may reveal the evolutionary roots of our conceptual knowledge.

Although the parallels between human and non-human tool use are striking, at least in those areas where specific observations and experiments have been brought to bear on the hypotheses, we must not lose sight of some equally striking differences, summarized in Table 15.1. Some of these proclaimed differences

Table 15.1. Comparisons between animal and human tools

Trait	Animal tools	Human tools
Sensitivity to relevant design features	yes	yes
Use of multiple materials for different tools	yes	yes
Use of multiple materials for a single tool	no	yes
Tool consists of articulated parts	no	yes
Tool consists of parts with different functions	no	yes
Tool kept with user for more than one application	no	yes
Tool designed for one function is used for another	no	yes

may, of course, turn out to be wrong, based at present on insufficient data. Others, however, are most likely correct and, we believe, should point the way to a better understanding of how and why our representations of tools are different from those of the other tools-users on the planet. Consider, for example, the observation that no non-human animal has ever created a tool with multiple, functional parts and nor has any animal designed a tool for one function and then used this object for a different purpose; chimpanzees use sticks for termite fishing and for throwing, but a stick that has had its leaves removed *for* fishing is not then used for throwing. This may appear trivial, but we believe it actually represents a profound difference, one that may tap our capacity to entertain multiple representations of the same object, thereby leading to a combinatorial explosion of possible functions. For example, although we know that a screwdriver was designed for tightening screws, if we ask someone to write down other possible functions, they will generate a long list once they get into thinking about problems in this way; a screwdriver can be used to keep a window open, as a paperweight, to punch a hole in a piece of paper, as a fork, as a 'brush' to paint with, and so on. When this capacity evolved is currently unclear. Looking at the archaeological record might suggest that it is recently evolved. Most of the early tools were made of one material (stone), with one functional part (the point of a handaxe), and one functional role. This view is, in our opinion, misguided, as it is based solely on the fossilizable artifacts. The early hominids most likely created tools out of wood, and may well have used vines to bind one object to another, creating a tool with multiple parts. Independently of its phylogenetic origin, once this capacity evolved it opened the door to a different kind of tool technology and a different way of thinking about tools. And once this kind of representation evolved, it could be transformed and passed on by two other, uniquely human, capacities: imitation and teaching. The synergy between our concepts of tools and our ability to copy and teach provided the foundation for cultural innovations and change, from the simple Swiss army knife to a computer with megagigahertz processing speed. But that is another story.