# 1

# Tool Use and Causal Cognition

## An Introduction

*Teresa McCormack, Christoph Hoerl, and Stephen Butterfill*

Why study tool use if you are interested in causal cognition? Take an everyday example of a tool, such as a spoon, a hammer, or even a coin used to loosen a screw because no screwdriver is to hand (all examples taken from chapters in this volume). Generally, whether a tool is useful for a given end, and how it should be used to reach that end effectively, will depend on particular physical properties of the tool and the targets of its application, as well as on basic causal-mechanical principles connecting these properties. A question running through the chapters in this volume is to what extent, and in what senses, tool users themselves need to be sensitive to, or engage in causal reasoning about, these properties and principles. One of the hard issues on causal cognition concerns identifying its varieties and how different varieties of causal cognition interact in enabling individuals or groups to manipulate and explain their environment. Studies of tool use provide some new hypotheses and fresh directions on issues such as this, as this volume demonstrates.

One specific way of bringing out what is special about approaching causal cognition through the study of tool use is by looking at two other ways in which aspects of our ability to learn about or become aware of causal relations have been studied empirically.

It is a striking fact about the literature in experimental psychology on causal learning that many of the experiments are deliberately set in contexts in which substantive background knowledge about how the world works is not relevant or useful. There is a good reason for this: often researchers have been concerned with testing models of causal learning that assume such learning in one way or another involves the detection of statistical information about the covariation of putative causes and effects. To test such models, the experimenter needs to control the nature of the statistical information that participants are exposed to, and he or she will usually remove other factors that may impact on causal judgments over and above such information. Of course, some researchers do consider how sensitivity to statistical information may interact with knowledge of particular substantive causal principles regarding, e.g., interactions between physical objects (or indeed lead to the acquisition of such knowledge). However, the underlying assumption in this area of research often seems to be that if we can appropriately model the detection of the relevant statistical information, we will have got to the heart of how causal relationships are learned.

There is another, quite different, strand of research in experimental psychology that focuses on the perception of causation. This research follows the Michottean tradition of establishing the circumstances in which participants will report perceiving two events to be causally related to each other (Michotte, 1946). Participants in such experiments are shown events in which objects move in spatiotemporal patterns highly controlled by the experimenter and asked to judge, for example, whether the movement of one object was caused by the movement of another (e.g., Choi & Scholl, 2005; White & Milne, 1997). Typically, experimenters have examined in detail how such judgments are affected by spatiotemporal parameters of the events that are observed. Research in this tradition has primarily focussed on simple collision events in which no information about the mechanics of the scenario other than the effects of object collision is relevant.

Studies of tool use stand in stark contrast to both of these research traditions. Rather than focusing on the detection of statistical information, the tasks involved in studying tool use are not designed to exclude substantive causal knowledge regarding the behavior of physical objects. Instead of stripping the physics out of the research paradigms, tool use studies potentially put knowledge of causal-mechanical principles at centre stage. Tool use studies also differ from studies of perceptual causation in that the participants (be they adults, children, or non-human animals) are not passively observing objects interact, but handling objects themselves. Moreover, objects are being used in such studies, for the most part, in the service of the participants' goals: participants are trying to do things with the tools they are using. These facts about tool use studies mean that the sets of issues they generate are in some cases quite different from the sets of issues generated by other sorts of studies of causal cognition. For a start, the issue of whether tool use requires sensitivity to physical principles, and of what sort, is at the centre of debates about the extent to which tool use requires genuine causal understanding or reasoning (see contributions in this volume by Goldenberg, Seed and colleagues, Peacocke, and Povinelli and Penn). Moreover, if causal cognition is involved in tool use, it is, literally in most cases, a "hands-on" sort of causal cognition. As Greif and Needham highlight in their chapter, it is common in the psychological literature to actually define a tool as something that is an "extension of the body that expands the functional range of a limb." This in itself raises interesting further issues that don't arise in other paradigms studying causal cognition, such as how tool use may affect representations of the body and personal space (see chapters by Spence and by Cardinali et al.).

On a theoretical level, the specific significance that research on tool use might have within the context of studies of causal cognition might helpfully be connected to a distinction between two notions of causation, which Woodward, in his chapter, calls the difference-making and the geometrical-mechanical notion, respectively (see Peacocke's chapter for related discussion). Put in those terms, the distinction, which has its origins in discussions about the metaphysics of causation, may not be immediately familiar to psychologists. However, as Woodward points out, we can in fact also conceive of it as a distinction on the level of what it is to represent or understand causality, which connects more directly with work in psychology. Geometrical-mechanical accounts describe causation in terms of a physical process connecting cause and effect. Related ideas have been cashed out in various different ways within psychology, with perhaps the most salient recent example being Wolff's (2007) characterization of causal representation in terms of force dynamics. Shultz's (1982) seminal work on children's causal reasoning also stems from an approach of this type, with Shultz arguing that children make causal judgments on the basis of a grasp of the notion of "force transmission." On this type of account, causal representation puts to work a grasp of basic principles of mechanical interactions, and indeed it has been explicitly claimed that the notion of a mechanism connecting cause and effect is itself at the core of causal cognition. Moreover, as Woodward points out, it has been claimed that "the concepts and principles deployed in recognition and understanding of mechanical interactions serve as the basis for more general notions of causation and causal mechanism." For example, Wolff (2007) argues that psychological causation is, at its most basic level, understood in terms of principles of force dynamics that stem from how mechanical interactions are understood. If this is one's general approach to causal representation it is easy to see that tool use studies might be thought of as an ideal context in which to examine causal abilities. Put simply, one might think that tool use is the basic situation in which this sort of understanding of mechanical interactions is being put to work, and that studies of tool use should focus on teasing out the precise way in which it is underpinned by or involves a grasp of mechanical principles. Thus, on this view, tool use studies reveal the most basic type of causal understanding being put to use, in a way that studies that focus on learning statistical relationships between cause and effect or studies of perceptual causation do not.

The contrasting approach to causation is what Woodward terms a difference-making approach. On this approach, a cause is understood to be something that makes a difference to whether or not an effect occurs, with the notion of difference-making being spelt out in a variety of ways. On his own, interventionist, approach, A causes B if there is a relationship between A and B that remains invariant under interventions. Clearly, what exactly this comes to, and in particular how exactly the notion of an intervention is to be understood, needs to be spelled out in much more detail (see Woodward, 2003). But, as Woodward puts it, "interventionist accounts attempt to capture in this way the commonsense idea that causes can be thought of as 'handles' for manipulating and controlling their effects." Indeed, recently, some psychologists have also suggested that what it is to represent a relationship as causal is not a matter of assuming that there is a mechanism connecting cause and effect, but a matter of representing what would happen to the effect if certain interventions were to be carried out on the cause (see in particular, Schulz et al., 2007). Moreover, the general idea suggested by difference-making accounts—that learning whether a relationship is causal may often involve using information about the likelihood of the effect occurring in the presence and absence of the purported cause—is, as Woodward points out, at the heart of many traditional studies of contingency learning. In such studies, experimental psychologists deliberately present participants with exactly this sort of statistical information. Yet, in these types of studies participants are typically passive observers of events and do not themselves intervene to fix the values of variables (there are some notable exceptions to this; see e.g., Sobel & Kushnir, 2006; Steyvers et al., 2003). Studies of tool use might, by contrast, be thought to be of particular interest in the context of difference-making approaches to causation such as interventionism, because they provide circumstances in which participants are actively trying to intervene on the world in the service of their particular goals. Tool use could potentially be seen as making use in action of the information or knowledge one holds about difference-making relations involving physical objects.

There are, then, a variety of reasons for thinking that studies of tool use should be of particular interest within the context of research on causal cognition. Yet, at the same time, there might also be grounds for arguing that such studies are in fact only of limited value in assessing the extent to which an individual can actually engage in anything that deserves to be called genuine causal reasoning or understanding. For instance, the very fact that tool use studies directly assess forms of practical, goal-directed behavior raises issues about the sophistication of the cognitive abilities underpinning tool use. Thus, one can query whether this sort of purely practical ability necessarily requires any type of reasoning about the causal relations involved (e.g., see discussion in this volume by Woodward, Peacocke, Penn and Povinelli, and Seed and colleagues).

Edwards et al. also argue in their chapter that tool use studies are potentially less useful to animal researchers interested in causal cognition than studies employing methodologies that have been used to examine causal structure learning in humans (and in young children in particular, e.g., Gopnik et al., 2001; Kushnir & Gopnik, 2007). Moreover, they argue that it is the latter methodologies, rather than those involved in tool use studies, that level the playing field between humans and animals in the domain of research on causal cognition. In one of their studies, monkeys are presented with sets of objects and shown patterns of covariation between the presence or absence of individual objects or pairs of objects and the presence or absence of a certain outcome, i.e., a detector only activates when certain objects are placed on it. (They also compare the monkeys' performance on this task with that of children on a similar one.) The measure of interest in Edwards et al.'s studies is whether the monkeys will arrive at the appropriate judgments as to which objects possess the causal power to set off the detector, as indexed by their tendency to place the relevant objects on the detector to receive a food reward. It is perhaps useful to spell out one very obvious but crucial difference between such methodologies and those of tool use studies. In the studies described by Edwards et al., monkeys have a goal (to obtain one or more grapes) and are sensitive in some way both to the causal powers of objects and to the fact that placing them on a piece of apparatus is necessary for achieving their goals. Thus, these studies assess goal-directed non-verbal behavior. However, although obtaining the reward depends on using an object appropriately, it seems intuitively wrong to classify this as an instance of tool use. One particular background intuition that may be at work here is reflected in some existing definitions of tool use (e.g., Goldenberg & Iriki, 2007) that restrict talk about tool use to cases in which an object is used to alter the spatial location or arrangement of a target, as when a hook is used to drag an object, or a hammer to drive in a nail (see Campbell's chapter for more detailed discussion).

Perhaps even more to the point, though, in Edwards et al.'s study, sensitivity to the causal powers of the objects in question would seem to be completely detached from any substantive knowledge about the nature of the objects and the detector itself. There are no apparent physical principles, over and above the basic physical contact between object and detector, that are relevant to the causal status of the objects in the experiment. Thus, the monkeys can be seen as in a situation parallel to that described by Campbell in his chapter, in which someone learns that a gadget opens a set of curtains and thus can operate the curtains, but has no idea as to the underlying mechanics of the situation that underpin the pattern of covariation that they have observed. Of course, whether this is important depends critically on how one wishes to characterize the nature of causal cognition. It is interesting to note that in the parallel studies with children in which this type of paradigm has been used, researchers have then gone on to examine whether children seem to have assumptions about whether the causal powers of the objects are related to their *internal* (in contrast to external) appearance. For example, David Sobel has examined whether children assume that objects with similarly colored insides, rather than outsides, are more likely to have similar causal status with regards to whether they set off a detector (Sobel et al., 2007). The underlying notion of a causal mechanism guiding such studies is the idea that causal mechanisms are tied to object properties that are typically "hidden" from view and must be discovered or inferred indirectly. It can be seen immediately that given this guiding conception of causal mechanisms, tool use is not an interesting context in which to study causal cognition. This is because in nearly all studies of tool use, the causally relevant spatial and physical properties of tools and targets are not concealed from perception (though this leaves scope for further debate as to whether they should be classified as directly observable properties, which is an issue we will return to below in the context of discussing the chapters by Povinelli and Penn, and by Seed and colleagues).

In fact, Edwards et al. argue against tool use as a marker for causal cognition for a different reason: what they take to be important to demonstrate is that participants must figure out sets of causal relationships in a system of causally related variables. Thus, in line with a number of other theorists, they see the ability to sort out causes from effects and figure out the (potentially complex) nature of such relationships in order to form a *causal model* of a system as being at the heart of causal cognition (see also Gopnik et al., 2004; Sloman, 2005). Tool use studies are difficult to interpret in this framework for the reason Edwards et al. suggest: because it is hard to clearly summarize the relevant sets of causal relationships in terms of the sort of structures specified in the causal models approach. We might again try to tease out what might be the sort of guiding assumption behind Edwards et al.'s skepticism about tool use studies in this respect. Although not explicitly discussed by Edwards et al., one possible candidate may be that possessing a causal model of a system involves representing causes as distinct from their effects, and understanding the directionality of the relationships and the dependencies and interdependencies in the system. One could align this requirement with Woodward's suggestion that explicit causal representations "decouple" means and ends and moreover separate out the representations of means themselves into the sub-component links in a causal chain (i.e., the intermediate steps in a causal sequence). A potential worry about tool use studies may be that they do not, or do straightforwardly, demonstrate that the tool user can separate out and represent the components of the causal system with which they are operating in this way. (However, see Campbell's chapter for a characterization of what he calls "intelligent tool use" that recruits at least some of the theoretical elements in terms of which Woodward's interventionist approach analyses causal systems.)

Facts about how tool use abilities are acquired may also raise similar issues about the sophistication of the abilities involved in successful tool use. First, developmentally, use of many familiar tools seems to be acquired gradually with repeated practice with objects (see Greif and Needham's chapter), and indeed may in some cases emerge from repeating action patterns with objects in which the objects are not initially treated as tools. Second, as Greif and Needham also emphasize, tool use often takes place in the context of particular social practices: there are social norms for how a fork or a hammer is used and a social consensus as to what such implements usually look like. Indeed, studies of tool use have extensively informed debates about the cognitive underpinnings of imitation and the social acquisition of knowledge in animals (e.g., Whiten et al., 2005). This suggests that there may be a distinctive way in which causal knowledge about tools can be acquired that is quite different from the sort of causal cognition assessed in standard paradigms of causal cognition. It may be tempting to suggest that if tool use is acquired in either of these ways—due to repeated practice with objects or due to imitation/social learning—what may be acquired is exactly the sort of undifferentiated behavioral routine that Woodward argues falls short of explicit causal representation. On this picture, the obvious challenge to researchers is to show that tool users can do more than simply reproduce such routines.

## 1. The cognitive basis of tool use: Mechanical reasoning versus manipulation knowledge

A common theme across several chapters is that we can distinguish between different types of tool use in terms of the cognitive resources that they draw upon, although the authors make this distinction in different ways. The general idea that there are different types of cognitive skills underpinning tool use is a source of considerable debate within the comparative literature, with different researchers interpreting animals' tool use in ways that vary dramatically in terms of the richness of the cognitive resources posited to underpin the behavior. This is also true in the developmental literature (see Lockman, 2000; Greif and Needham's chapter in this volume). Given this, neuropsychological findings that show dissociations in different aspects of tool use are particularly interesting, because they could be interpreted as providing hard evidence that there are genuinely different ways in which cognitive resources contribute to tool use. In his chapter, Goldenberg makes a key distinction between tool use that may be underpinned by manipulation knowledge—roughly speaking, information about how a particular tool is typically put to use—and tool use that may be underpinned by what he terms mechanical reasoning. Tool use underpinned by manipulation knowledge draws on a type of memory that, amongst other things, specifies behavioral routines about how an individual tool is used, whereas tool use underpinned by mechanical reasoning involves what Goldenberg views to be a type of problem-solving in which the tool user figures out how the tool should be used given the mechanics of the situation the user finds themselves in. Both of these types of cognitive resources may come into play in everyday tool use. However, Goldenberg believes that there is evidence for a distinction between them: the existence of double dissociations between, on the one hand, patients who seem to have preserved manipulation knowledge, and thus have some ability to use familiar objects, but impaired mechanical reasoning, as shown by their lack of insight into how novel tools could be used, and, on the other hand, patients who seem to be most severely impaired in manipulation knowledge but may retain mechanical reasoning abilities.

01-McCormack-Ch-01.indd 6 8/9/2011 9:33:02 AM

The import of Goldenberg's use of the term "mechanical reasoning" is that human tool use usually goes beyond simply bringing to bear well-practiced behavioral routines that could potentially have been honed through a process of trial and error. Indeed, Goldenberg believes that novel tool use studies, in which patients are faced with a situation in which they have to select appropriately from a range of unfamiliar tools the correct one to perform a task that they have also not faced previously, are particularly interesting because we can examine whether patients select the correct tool before they have ever had an opportunity to try it out. Goldenberg's distinction between manipulation knowledge and mechanical reasoning might be thought to provide a starting point for considering how we should interpret the tool use abilities of non-human primates and young children. For instance, we might want to consider the simple hypothesis that these populations possess manipulation knowledge, usually through a process of trial and error, rather than mechanical reasoning abilities, and then consider what types of evidence would support such a position. Moreover, we could give this hypothesis more theoretical bite by arguing that only in the case of mechanical reasoning is genuine causal cognition involved.

However, one reading of Greif and Needham's chapter is that such a distinction on its own does not capture the complex ways in which the learning of tool-specific actions and the sort of understanding that may underpin mechanical reasoning interact with each other in the acquisition of tool use. They portray the acquisition of tool use as involving an interplay between action patterns, knowledge about how a tool is usually used, and conceptual knowledge about object properties that could be used to reason about tools. As they point out, through repeated action with a tool children can actually generate data that are suitable for serving as the input to such conceptual knowledge. That is, they argue that it is through hands-on experience with tools that children become sensitive to the particular physical properties that underpin successful tool use. However, Greif and Needham also argue that the relationship between such knowledge and action is more complex and bidirectional than this simple picture. For example, they observe that it is only with development that children seem to reliably and appropriately repeat a successful action with a tool, and suggest that even generating appropriate repetitions efficiently (which would presumably provide further useful data for the child) may already depend upon representing the relationship between tool and object in a way that is not available to younger children. Moreover, they argue that although this efficient reproduction of successful actions may lead to a rigidity that initially constrains tool use (as evidenced, for example, in the child's reluctance to use a spoon in a novel way), this constrained use of tools leads to children producing more selective and focused actions, which in turn allows them to pay attention to and filter out the exact properties of the tool and target necessary for successful tool use. Thus, on Greif and Needham's picture, this initial rigid use of behavioral routines ultimately facilitates the conceptual knowledge that they believe allows generalization of tool use to novel contexts and the ability to use novel tools. This is not to say that Greif and Needham do not want to characterize the developmental end-point of tool use as involving something similar to what Goldenberg describes as mechanical reasoning. However, they do not see a simple one-way developmental progression from the learning of action routines to mechanical reasoning.

One implication of the arguments in Peacocke's chapter is also that it may be misleading to think of the crucial psychological distinction in terms of a distinction between rigid behavioral routines and mechanical reasoning, as described so far. The key distinction that Peacocke himself argues for is one between tool use that may be underpinned by a grasp of sets of conditional goal–action pairs (i.e., rules such as "To get G, do A") and tool use that may involve a genuine grasp of causal relationships. As he argues, the sort of behavior that may be explicable by possession of such rules goes far beyond what we would consider as inflexible behavioral routines. This is because he thinks of such rules as potentially embedded in conditionals (e.g., "When C holds, to get G, do A; when D holds, to get G, do E"), and indeed these conditions could even make reference to object properties (e.g., be of the form "When C is rigid, to get G, do A"). He suggests that such rules could be combined and used to reach solutions in a way that would appear to be creative, without yet being underpinned by any genuine grasp of causal notions. Thus, on his picture, a simple distinction between manipulation knowledge and mechanical reasoning does not seem to capture the possible range of explanations of behavior that we may wish to make use of in describing the cognitive abilities underpinning tool use. For a start, on his picture, manipulation knowledge could vary dramatically in terms of the complexity of the rules that underpin it (one issue that may be of relevance here is that some developmental psychologists argue that children are not capable of dealing with embedded conditional rules until around four years of age; Frye et al., 1996). Moreover, Peacocke's position also allows for a sense in which tool use, especially innovative tool use, might be said to involve reasoning but without necessarily involving causal cognition; we discuss in the next section what he takes to be necessary for causal cognition.

The general idea that we can distinguish between types of tool use that involve different levels of cognitive sophistication appears again in Campbell's chapter, which attempts to characterize intelligent tool use. The key distinction he makes is between tool use that is grounded in an awareness of what he terms the relevant standing properties of the tool at hand, and tool use that may involve a sensitivity to the covariation between one's actions with the tool and outcomes, but is not grounded in such an awareness. By "the relevant standing properties" here, Campbell means the physical properties of the tool by virtue of which it can bring about the desired changes in the object or surface that one is trying to transform by using the tool, properties such as its weight, solidity, and sharpness. We could immediately hypothesize that it is the kinds of awareness Campbell has in mind here that seems to be missing in some of the most severely impaired patients Goldenberg describes: not only do they not seem to know, e.g., how a knife is used (manipulation knowledge), or what it is usually for (functional knowledge), they don't seem to have any sense of what could be done with an object of this sort. As he also puts it, even saying that those patients act as if they had never used the relevant tools before does not get to the bottom of the issue. A person who had never used a knife before but had some awareness of its causally relevant properties would not try to cut a loaf of bread by pressing the knife into it without moving it to and fro, as those patients will do.

Campbell points out that intelligent tool use does not just "demand that you somehow have an internalized manual for the correct use of the tool, guiding your actions," and Goldenberg could agree on this point. Where the two accounts appear to differ is on exactly how to characterize what more is involved in more sophisticated tool use. In his use of the term 'mechanical reasoning', Goldenberg seems to be suggesting something akin to an inferential process, whereby the tool user is not just aware of the properties of the tool, but engages in a chain of reasoning about the implications of these properties. Campbell explicitly denies that we need assume any type of reasoning process here, or indeed that the intelligent tool user need to be in possession of a "theory" about the causal significance of the properties. On his account, what makes tool use intelligent is whether in actually making use of the tool, the user has the appropriate awareness of its properties. This awareness will typically entail that various types of behavior, beyond rigid behavioral routines, are possible for the tool user (we will return to this issue below). However, for Campbell, in exploiting this awareness the tool user is emphatically not reasoning. Rather, he argues, there are direct-wired connections between use of the tool and awareness of the relevant object properties: there is a direct dependency between this awareness and the pattern of tool use itself in that one simply causes the other. This raises interesting developmental issues about how such a direct dependency may arise that link with those discussed by Greif and Needham. Note specifically that Campbell uses the term "direct-wired" rather than "hard-wired," which allows for the possibility that the direct dependency between awareness of the relevant properties and patterns of tool use is something that can emerge, perhaps as a result of the sort of interplay between acquisition of action patterns and sensitivity to functionally relevant properties that Greif and Needham describe.

To sum up, although the neuropsychological literature described by Goldenberg could be used as a basis for suggesting a core distinction between manipulation knowledge and mechanical reasoning, where only the latter is viewed as involving causal cognition, whether this is the only or the most useful way to distinguish between two cognitive bases for tool use is a matter of debate amongst the chapters' authors. Moreover, these considerations raise the general issue of whether it is potentially unhelpful to think of there just being one sense in which tool use may involve causal cognition.

## 2. Causal cognition: Unitary or multidimensional?

Campbell himself is trying to draw out the difference between intelligent and non-intelligent tool use, and does not address the question as to in what sense specifically causal cognition is involved in tool use. The hypothesis that tool use may be guided by a grasp of rules describing links between goals and actions rather than genuine causal understanding is what Peacocke terms the "austere hypothesis," and, as we have mentioned, he gives a number of examples of what might look like fairly sophisticated behavior that may nevertheless be describable in terms of conditional rule use. On Peacocke's account, one basic way in which tool use may go beyond merely exploiting such rules is if it is informed by beliefs that have a counterfactual flavor, i.e., if the agent does A to get G, in part because of a belief that he will not get G if he does not do A. Moreover, though, what is critically important for Peacocke is that the tool user who has a grasp of causal notions will understand *why* such beliefs hold, i.e., by virtue of which explanatory properties it is the case that doing A is followed by G and not doing A is not followed by G. That is, where tool use involves causal understanding, the tool user is committed to there being an explanation of the connection between their action and the outcome in terms of properties that in fact underpin the connection.

Thus, although he believes that some types of tool use are indeed underpinned by causal understanding, Peacocke argues against the assumption, discussed extensively by many researchers (e.g., Penn & Povinelli, 2007; Taylor et al., 2009; Visalberghi & Tomasello, 1998), that tool use behavior can serve as evidence for a grasp of causality. As he puts it, "[e]ven creative and ingenious uses of tools can be explained without any attribution of grasp of causality." Rather, demonstrating a grasp of causality, on Peacocke's view, is, at least in part, a matter of demonstrating that the subject has particular beliefs as to why there is a connection between how it acts and whether it achieves its goals. For Peacocke, such beliefs turn on a grasp of the causal role of properties such as heaviness, resistance, and force, as provided by an intuitive mechanics. In other words, demonstrating that the subject holds such beliefs is a matter of showing that it can engage in a particular kind of explanation, e.g., as to why use of a given tool can bring about the intended effect, rather than simply the fact of tool use (or even manufacture; e.g., Kenward et al., 2005; Weir & Kacelnik, 2006) itself. Connectedly, Peacocke takes as being particularly revealing of a grasp of causality a circumstance in which a subject would be surprised on being shown that the mechanism underpinning the connection between two events was other than might be expected, thus showing that the subject did work under the assumption of an underlying mechanism of a particular sort. By contrast, a subject who only possessed goal–action rules would have no grounds for being surprised in this type of circumstance.

The idea that, in interpreting tool use, it is important to consider what, if any, type of explanatory reasoning the tool user is capable of engaging in also appears in Povinelli and Penn's chapter. These authors claim that non-human primates may not just fail to be committed to particular explanations of the success of their actions, but moreover that they differ from humans fundamentally in that they simply do not consider the issue of why some actions are successful and other are not—an ability Povinelli and Penn refer to as diagnostic causal reasoning.[1] For example, they mention one study in which human children but not chimpanzees try hard to find out why a particular wooden block (covertly weighted off-center by the experimenter) would not stand up. In this study, chimpanzees simply continued to try to make the block stand up, whereas 3-year-olds explored the block in various ways to try to diagnose the problem.

As we have seen, Peacocke describes this kind of grasp of the explanatory value of properties as part of an intuitive mechanics. His claim that the counterfactual beliefs associated with a grasp of causality must be rooted in possession of an intuitive mechanics can be contrasted with the view Woodward explores in his chapter. Interestingly, Woodward argues that it is not helpful to assume that there is a single sense in which cognition can be described as causal. Rather, he suggests that a geometrical-mechanical grasp of causation (which we might spell out in terms of possession of an intuitive mechanics) and a difference-making grasp of causation (with which we might align with a grasp of certain counterfactuals) may each capture quite different aspects of causal cognition. Indeed, he argues that the two are distinct from one another in that it may be

[1] Indeed, we note that Edwards et al. also believe this to be a core human/animal difference.

perfectly possible to possess appropriate difference-making knowledge without it being underpinned by geometrical-mechanical knowledge, or vice versa, and in either of these cases it would still be correct to describe such knowledge as causal. One important idea that he elaborates in his chapter is that although these two types of causal cognition may be well integrated and meshed in mature humans—i.e., adult humans typically appreciate the relevance of geometrical-mechanical information for beliefs about difference-making relationships, and the relevance of information about difference-making relationships for beliefs about the existence of mechanisms—it is perfectly possible that this is not the case in all species. Thus, he argues that although it may seem initially plausible to explain any deficiencies observed in animal tool use in terms of, for example, a lack of the appropriate geometrical-mechanical understanding, it is equally plausible to assume that these deficiencies may result from a failure to properly integrate difference-making and geometrical-mechanical knowledge. He argues that such a failure to properly integrate these types of knowledge could also show up in contexts in which animals or human infants seem able to detect the relevant geometrical-mechanical properties, as measured in, for example, looking-time tasks, but unable to use those properties to appropriately guide their actions.

For Woodward, not only may different groups differ in terms of how well integrated these two strands of knowledge may be, but the sophistication of each of these strands of knowledge could also vary across groups: for example, geometrical-mechanical knowledge may vary in how complete it is, and difference-making knowledge may vary in terms of which sources the subject is able to acquire it from (e.g., just from their own actions, or from observation). Thus, on Woodward's picture, causal cognition is genuinely multidimensional, and the challenge to researchers conducting tool use studies is to figure out the status of tool use behaviors in terms of the various dimensions he specifies. One straightforward implication of his multidimensional approach is that in painting a picture of the causal competence of any group or species, we would need to consider the findings of tool use studies alongside the findings of studies using a wide variety of other paradigms.

## 3. Generalizing and generalizability

Animal researchers, such as those working with non-human primates, may find Woodward's approach daunting because it suggests that it may not be advisable to try to identify a single general difference between human and animal cognition. To some extent, this runs contrary to the approach taken by many comparative psychologists. For instance, Povinelli and Penn believe that non-human primates differ fundamentally in terms of the repertoire of concepts they possess, arguing that only humans possess concepts sufficient for underpinning diagnostic causal reasoning—causal reasoning in which the reasoner seeks explanations as to why the regularities that they encounter occur. Seed and colleagues, in their chapter, discuss empirical evidence that they believe contradicts Povinelli and Penn's claims about the conceptual repertoire of animals and put forward an alternative, but similarly global, view of what they see as the one key difference between human and animal cognition.

The positions put forward by Povinelli and Penn and by Seed and colleagues differ sharply; thus it may be useful to try to describe in neutral and general terms what the two sets of authors appear to agree on:[2]

(1) Engaging in genuinely causal reasoning about the physical world necessarily requires possession of certain types of concepts about objects and their interactions.
(2) These concepts are in some sense "abstractions" of perceptual properties.
(3) These concepts enable the subject to produce a range of tool use behaviors in a way that generalizes across contexts in an important sense; behavior not based on such concepts is not generalized in this way.

What seems to make the debate between these authors difficult is that what each takes to be the lower level alternative explanation of animal behavior is different. On Seed and colleagues's picture, the relevant distinction is between possessing arbitrary links between responses and perceptual properties of a scenario that are not themselves functionally relevant and possessing what they call structural knowledge, in which action can be based on the encoding of functionally relevant properties (such as weight or length). On this picture, animals operating at the lower level would fail to be sensitive to these functionally relevant properties, and Seed and colleagues argue that there is an abundance of evidence from their own studies that there are at least some animals, including both non-human primates and corvids, that do show such sensitivity and indeed distinguish between arbitrary and functionally relevant properties. In other words, animals operating at the lower level, as Seed and colleagues understand it, could not possess representations of object properties such as heaviness or floppiness. However, Seed and colleagues also believe that animals are not in fact restricted to that lower level.

Povinelli and Penn, on the other hand, take the lower level explanation to be one in which animals can do more than simply pick up arbitrary perceptual regularities. Rather, they can indeed be sensitive to properties such as heaviness that are functionally relevant, and can treat functionally relevant properties as different from functionally irrelevant ones, although they may require appropriate repeated experience to tune into the functionally relevant properties. Moreover they believe that it would not be incoherent to describe the animals as reasoning with these representations rather than simply producing inflexible behavioral routines linked to perceptual features of scenarios. However, Povinelli and Penn believe that there is nevertheless still an important sense in which the representations animals are using are more perceptually bound and less abstract than those of humans. Seed and colleagues believe this too, but they believe that animals' difficulties lie with understanding symbolic representations (i.e., that one stimulus can stand for or symbolize something else). Thus, the two sets of authors differ in terms of what they take to be the correct distinctions to make between the alternative types of knowledge and thus not surprisingly differ in terms of whether they wish to describe non-human primates as operating at the lower level.

[2] As we have seen, a view such as Woodward's might be seen as presenting a way of challenging the very way in which the debate is being set up here.

A large part of the difficulty here lies in the inherent vagueness of the notion of an abstract concept, with Povinelli and Penn, for example, arguing that there could be a way of representing and reasoning about a notion such as weight that is not as abstract as that used by humans, and Seed and colleagues arguing that being able to represent and reason about weight already involves abstracting beyond perceptual cues. Thus, at its most basic level, resolving their debate is a matter of dealing with very hard questions about what it is to possess concepts of certain types and in what sense concept possession involves generalizing beyond what is given in perception. Nevertheless, it seems to us that these authors need not necessarily disagree on every aspect of their characterization of non-human primate abilities. Specifically, Povinelli and Penn argue not just that animals do not possess a certain repertoire of concepts. As we saw already, they also argue for a key difference in how humans and animals exploit their knowledge of the world: animals are concerned with successful action, but they do not engage in diagnostic causal reasoning, reasoning aimed at uncovering why some actions are successful and others are not. Povinelli and Penn assume that these two deficits are in some sense two sides of the same coin: The assumption seems to be that uncovering explanations for successful or unsuccessful action is a matter of discovering exactly the sort of properties underpinning these actions that Povinelli and Penn believe animals have no concepts of. However, it is at least possible that Seed and colleagues could agree with Povinelli and Penn, if not in terms of what concepts animals do or do not possess, but in terms of animals' basic mode of applying their knowledge. That is, they could potentially agree that animals do not engage in diagnostic causal reasoning: in Povinelli and Penn's terms, they could agree that chimps never ask "why?," even if they think that chimps do possess some sort of abstract concepts.

As mentioned earlier, both sets of authors also take it to be important to examine empirically the extent to which the behavior an animal acquires is appropriately generalized to new contexts (see also Greif and Needham's chapter for a discussion of this issue in a developmental context), and here they seem to at least implicitly assume that generalizability may be an all-or-nothing ability that hinges again on either possessing or not possessing appropriately abstract representations of object properties. Roughly, the idea seems to be that the relevant concepts themselves generalize across situations in which quite different perceptual features are present, and because of this, an animal needs to possess such concepts to generalize their behavior across these contexts. Woodward's chapter already provides some alternative ways of thinking about generalizability in animal behavior: for example, in terms of how well difference-making and geometrical-mechanical knowledge are integrated, or in terms of whether animals can generalize across not just their own behavior over time but also across their own actions and actions they have observed others perform. If we turn to Campbell's chapter, we can also see a discussion of generalizability that allows for different degrees of generalizability and different ways in which generalizability may be limited.

Recall that Campbell's question is what constitutes intelligent tool use. For him, this question turns on whether tool use is underpinned by an awareness of the standing properties of the tool and its target (hardness, rigidity, etc.), and he points out that we should expect tool use behavior to display a certain amount of systematicity if it is underpinned by such awareness. On one way of reading Campbell, what he has in mind here is basically the same as what we have called generalizability in the above discussion of Seed and colleagues, and of Povinelli and Penn. Yet, Campbell's position is very different from theirs in at least two quite crucial respects.

First, on Campbell's view, there may be a linear spectrum of tool use behaviors that vary according to the extent to which they involve and are grounded in awareness of particular standing properties of the tool and the target. In fact, Campbell goes on to distinguish between three different dimensions along which we could distinguish different levels of generality in an animal's grasp of the bearing such properties might have on tool use:

(1) Animals may vary in the *breadth* of properties of a system that they are aware of and that underpin their tool use;
(2) There may be insufficient *spread* in the range of behaviors that are underpinned by such awareness, so that although the animal may appreciate the importance of all the relevant standing properties in one task, they do not invariably appropriately generalize across tasks contexts; and
(3) They may not be *systematic* in how they make use of their awareness of standing properties, such that although they may appreciate the significance of a property of a tool for successful action, they may fail to realize that systematic variations of the value of this property may be required if the value of a dimension of the problem context also varies.

Thus, like Woodward, Campbell's approach differs from that found both in Seed and colleagues and in Povinelli and Penn, in that Campbell argues for a multidimensional approach to considering the sophistication of tool use.

The second important point of divergence is one we have come across previously. Recall that Campbell does not think that the awareness of the standing properties of a tool and its target on which intelligent tool use turns, and which grounds the generalizability of tool use along different dimensions, must involve the ability to engage in reasoning, e.g., about how the standing properties of a tool are connected to what can be done with it. Thus, not only is his account multidimensional in the way just sketched, it also differs from the approach adopted by both Povinelli and Penn and Seed and colleagues in separating out the question as to whether a subject can generalize tool use, e.g., to new contexts and the question as to whether that subject can engage in diagnostic or explanatory reasoning.

## 4. Tool use and the body

We finish this introduction by briefly considering how the issues discussed in the chapters by Spence and by Cardinali et al. may link to the issues discussed in the other chapters. As we mentioned earlier, many researchers take as their working definition of a tool as being an "extension of the body that expands the functional range of a limb" (see Greif and Needham; it is highly unlikely that these authors would deny there are complex tools that do not easily fall under this definition; however, they take the interesting developmental questions to arise from examining the more basic cases in which this definition holds). We can look at Spence's and Cardinali et al.'s chapters to examine whether thinking about tools as extensions of the body is not just useful for definitional purposes or as a metaphor, but has a concrete psychological and neurological truth to it.

These two chapters discuss a wide variety of evidence from neuropsychology, neurophysiology, and experimental psychology for the hypothesis that tool use can lead to alterations in how the body is represented or in how attention is distributed relative to the body (or both). These alterations are specific to tool use: such alterations are not produced by moving an object that isn't a tool or by moving a tool without intending to use it. As Cardinali et al. discuss, one way of interpreting these findings is in terms of the idea that using a tool alters the body schema: put simply, the tool temporarily becomes represented as part of the body during its use. However, Spence argues that the findings are potentially consistent with an alternative view which he refers to as an attentional account: rather than modifying how space is represented, tool use modifies how likely a region of space is to attract attention or the extent to which a specific region of space is prioritized for attention.

Why might the question as to which interpretation of the experimental findings is appropriate be relevant to the debates about tool use we have described so far? One answer to this question is hinted at by Cardinali et al. toward the end of their chapter. They sketch the suggestion that, in evolutionary terms, tool use was initially underpinned by a process whereby tools became incorporated into the body schema. They portray this process as operating at a level below what we would normally describe as involving thinking or conscious cognition, and argue that with evolution the use of tools became underpinned by more sophisticated cognitive processes. Thus, although they believe that both human and animal tool use may be associated with changes to the body schema, they argue that human tool use typically involves additional cognitive processes. Although this proposal is sketched only briefly, it suggests an intriguing possible depiction of animal tool use: that animal tool use differs from that of mature humans in that it *relies* on a process whereby tools become incorporated into the body schema (see Johnson-Frey, 2003, for the related suggestion, based on neurological evidence, that animal tool use relies on purely sensorimotor representations). What is appealing about this notion is that it suggests that tool use may be primitive in that it simply piggy-backs on the animal's pre-existing motor skills—its ability to manipulate its own limbs in effective ways. This would clearly fall short of what Woodward, for instance, describes as the sort of tool use that involves causal cognition: tool use that involves representations in which means are decoupled from ends and in which intermediate links in the causal chain of means to ends are also potentially separated out. Moreover, we might expect animals to have difficulty generalizing tool use behaviors appropriately: for example, generalization to a new context may involve a different way of using the tool as a bodily extension and thus not be straightforward. And there would also be no reason to suppose that animals that are treating a tool as an extension of their body could easily recognize which properties of the tool itself are necessary for it to function appropriately. If this depiction of more primitive tool use were correct—i.e., as reliant on the tool becoming incorporated into the body schema—then it might suggest that there is nothing special about tool use per se that means we should treat it as an indicator of causal cognition. Tool use would thus potentially involve nothing more complex than is required for any of the other ways of negotiating the physical world that the animal has at its disposal.

Clearly, more work is required in spelling out what exactly this claim would amount to, and in establishing whether it is credible. For a start, as already mentioned, Spence argues in his chapter that the relevant findings do not establish that tool use alters the body schema itself. Instead tool use may alter how we represent the space around us (e.g., as being near or far from our bodies). Thus, for example, he points to evidence that suggests that tool use results not in peripersonal space being extended (because a limb is now represented as being longer, as a body schema account would hold) but rather as projected to far space (because a region of distant space is selectively re-coded as being nearer to one's body). As he discusses, resolving this issue empirically is not straightforward.

Moreover, even if we accept that tool use is associated with changes to the body schema itself, to our knowledge none of the experimental findings discussed in the chapters prove that changes in the body schema play a causal role in tool use. That is, it is a further question whether tool use *depends* upon or is facilitated by such changes in spatial representation. Indeed, the research that originally sparked much of the recent interest in this topic merely suggested that such changes might occur as *consequence* of tool use rather than underpinning it (e.g., see Spence's discussion of Irki et al.'s 1996 study). Yet, even if it is still an open question as to what sort of empirical evidence would clearly demonstrate that tool use is associated with changes in the body schema, and that these actually play a key role in tool use, the research reviewed by Cardinali et al. and by Spence clearly points to an important direction for further debate about which cognitive resources underpin tool use and what forms of causal cognition, if any, different types of tool use require.

## References

Choi, H., and Scholl, B. J. (2005). Perceiving causality after the fact: postdiction in the temporal dynamics of causal perception. *Perception*, **35**, 385–99.

Frye, D., Zelazo, P. D., and Brooks, P. J. (1996). Inference and action in early causal reasoning. *Developmental Psychology*, **32**, 120–31.

Goldenberg, G., and Iriki, A. (2007). From sticks to coffee-maker: Mastery of tools and technology by human and non-human primates. *Cortex*, **43**, 285–8.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, D. E., Kushnir, T., and Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, **111**, 1–31.

Gopnik, A., Sobel, D. M., Schulz, L. E., and Glymour, C. (2001). Causal learning mechanisms in very young children: Two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Developmental Psychology*, **37**, 620–9.

Iriki, A., Tanaka, M., and Iwamura, Y. (1996). Coding of modified body schema during tool use by macaque postcentral neurons. *Neuroreport*, **7**, 2325–30.

Johnson-Frey, S. H. (2003). What's so special about human tool use? *Neuron*, **39**, 201–4.

Kenward, B., Weir, A. A. S., Rutz, C., and Kacelnik, A. (2005). Tool manufacture by naïve juvenile crows. *Nature*, **433**, 121.

Kushnir, T., and Gopnik, A. (2007). Conditional probability versus spatial contiguity in causal learning: preschoolers use new contingency evidence to overcome prior spatial assumptions. *Developmental Psychology*, **43**, 186–96.

Lockman, J. J. (2000). A perception-action perspective on tool use development. *Child Development*, **71**, 137–44.

Michotte, A. [1946 (1963)]. *The perception of causality*. London: Meuthen.

Penn, D. C., and Povinelli, D. J. (2007). Causal cognition in human and nonhuman animals: A comparative, critical review. *Annual Review of Psychology*, **58**, 97–118.

Schulz, L., Kushnir, T., and Gopnik, A. (2007). Learning from doing: Intervention and causal inference in children. In A. Gopnik and L. Schulz (eds), *Causal learning: Psychology, philosophy, and computation* (pp. 67–85). Oxford: Oxford University Press.

Shultz, T. R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development*, **47**, 1–51.

Sloman, S. (2005). *Causal Models: How People Think about the World and Its Alternatives*. Oxford: Oxford University Press.

Sobel, D. M., and Kushnir, T. (2006). The importance of decision demands in causal learning from intervention. *Memory and Cognition*, **34**, 411–19.

Sobel, D. M., Yoachim, C. M., Gopnik, A., Meltzoff, A. N., and Blumenthal, E. J. (2007). The blicket within: Preschoolers' inferences about insides and causes. *Journal of Cognition and Development*, **8**, 159–82.

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., and Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, **27**, 453–89.

Taylor, A. H., Hunt, G. R., Medina, F. S., and Gray, R. D. (2009). Do New Caledonian crows solve physical problems through causal reasoning? *Proceedings of the Royal Society B: Biological Sciences*, **276**, 247–54.

Visalberghi, E., and Tomasello, M. (1998). Primate causal understanding in the physical and psychological domains. *Behavioural Processes*, **42**, 189–203.

Weir, A. A. S., and Kacelnik, A. (2006). A New Caledonian crow (*Corvus moneduloides*) creatively re-designs tools by bending or unbending aluminium strips. *Animal Cognition*, **9**, 317–34.

White, P., and Milne, E. (1997). Phenomenal causality: impressions of pulling in the visual perception of objects in motion. *American Journal of Psychology*, **110**, 573–602.

Whiten, A., Horner, V., de Waal, F. B. M. (2005). Conformity to cultural norms of tool use in chimpanzees. *Nature*, **437**, 737–40.

Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology: General*, **136**, 82–111.

Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.