Tool Use for Autonomous Agents

Robert St. Amant*
Information Sciences Institute
University of Southern California
Marina del Rey, CA 90292
stamant@isi.edu

Alexander B. Wood
Department of Computer Science
North Carolina State University
Raleigh, NC 27695
abwood@eos.ncsu.edu

Abstract

The intelligent use of tools is a general and important human competence that AI research has not yet examined in depth. Other fields have studied the topic, however, with results we can compile into a broad characterization of habile (tool-using) agents. In this paper we give an overview of research on the use of physical tools, using this information to motivate the development of artificial habile agents. Specifically, we describe how research goals and methods in animal cognition overlap with those in artificial intelligence. We argue that analysis of activities of tool-using agents offers an informative way to evaluate intelligence.

Introduction

The use of tools is a hallmark of intelligent behavior. It would be hard to describe modern human life without mentioning tools of one sort or another. For example, before writing this paragraph, I brushed my teeth with a tool, cooked and ate breakfast with tools, unlocked my office door with a tool, and jotted down notes with a tool. At my desk I am surrounded by tools that support cognitive as well as physical activities. A calendar aids my memory, a calculator improves my arithmetic, a whiteboard allows me to apply my visualization abilities, and of course my computer facilitates an even wider range of cognitive activities. In philosophical circles, some hold that tool use is central to intelligent behavior and that it should rival language in the study of cognitive phenomena (Preston 1998).

Tool use can be thought of as a specialized case of problem solving, but it can also be viewed in more general terms. Nilsson (1995) writes, "[Intelligent programs] need to be able to find out about what knowledge and tools are available to match the problems they face and to learn how to use them," arguing that systems with such capabilities may gain general human competence. This challenge has been taken up, in part, by researchers in robotics (Bluethmann *et al.* 2003; Stoytchev 2003) and intelligent user interfaces. It also drives the narrative in this paper; we believe that physical tool use can act as a precursor to cognitive tool use and thus more intelligent effective behavior.

*On sabbatical from North Carolina State University.

Copyright © 2005, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

The overall goal of our research is to characterize the class of habile (tool-using) agents in computational terms. Broadly speaking, a sophisticated habile agent can reason about how to exploit objects and features in its environment to reach its goals. A physical habile agent can perform spatial and physical reasoning (both qualitative and quantitative) about objects in relation to the agent's own location and dimensions. Most importantly, a habile agent can reason about the relationships between its capabilities, the problem at hand, and the tools that it can bring to bear in order to transform the problem into one that is more easily solved. While specialized physical habile agents already exist (e.g., some factory robots and most prominently Robonaut (Bluethmann *et al.* 2003)), the goal of our research is to develop agents with much more general physical tool-using abilities—such agents may not even explicitly incorporate the concept of tools, but can opportunistically exploit objects with appropriate properties.

This paper gives an overview of research on the use of physical tools, using this information to motivate the development of artificial habile agents. In the following section, we outline research on natural tool use in fields outside AI, focusing on work in animal cognition. We use our brief survey to motivate a set of desirable properties in an artificial habile agent. We then describe an approach to defining general tests of agent intelligence based on tool-using ability. We conclude with a discussion of how consideration of tool use may provide useful challenges for AI research.

Research on tool use

Accounts of tool use have appeared in the AI literature (e.g., (Agre & Horswill 1997; Bluethmann *et al.* 2003; Brady *et al.* 1984)), but the most extensive analyses are found elsewhere, in particular in the field of animal cognition. The most widely accepted definition of tool use, in any field, is due to Beck (1980):

> Thus tool use is the external employment of an unattached environmental object to alter more efficiently the form, position or condition of another object, another organism, or the user itself when the user holds or carries the tool during or just prior to use and is responsible for the proper and effective orientation of the tool.

a. Diagram of food tube apparatus

b. Diagram of rod for reaching through a narrow opening

Figure 1: Objects for animal tool use experiments

Although this definition may appear unwieldy, all of its conditions turn out to be needed to distinguish tool use in the animal kingdom from other activities. Three examples illustrate the insights we can gain from animal tool use.

Some of the simplest animals to use tools are wasps (*Ammophila urnari* and *A. Yarrowi*) that pound earth down into a nest with the help of a pebble (Oswalt 1973). Thus activities recognizable as tool use can be found in agents with limited manipulation, no capacity for learning, and only the simplest processing capabilities. This indicates that with careful programming it is possible to build a range of tool-using behaviors into an artificial agent (e.g., (Bluethmann *et al.* 2003).)

In laboratory experiments, the tool-using capabilities of several species of primates have been tested using a transparent tube in which an item of food is placed, as shown in Figure 1a. The end of the tube is too narrow for the subject's hand to enter, but tools such as reeds are made available for the subject to use. The results of experiments with capuchin monkeys (*Cebus apella*) are especially suggestive. Visalberghi and Limongelli (1996) observe that capuchins often succeed in this task by rapid-fire selection of different strategies, though trying out many inappropriate strategies along the way. For example, given a bundle of reeds too bulky to fit through the tube in aggregate, one monkey hit upon the step of unwrapping tape from around the reeds, but then tried to push through the tube with the tape. For AI researchers, this example is interesting less because of the success or failure of the monkeys but instead because it is reminiscent of an automated search process driven by poor heuristics for evaluation or operator selection.

Povinelli (2000) has carried out more detailed experiments along comparable lines. In one, chimpanzees (*Pan troglodyte*) were tasked with obtaining an out-of-reach piece of fruit visible through a small hole. Solving the problem involves grasping a rod, as shown in Figure 1b, and poking the narrow end of the rod through the hole to push the fruit out. The chimpanzees generally picked up the rod by the straight end (on the right side in the diagram), rather than the more difficult-to-grasp end, and thus failed in the task. This example shows where difficulties may arise in an agent's understanding of the effect of a tool and points to the importance of affordances (Gibson 1979) in tool use (this also leads to an interesting distinction between *proximal* and *distal* affordances—those associated with grasping tools and with applying them to objects.) Povinelli's experiments drive his speculation about why orangutans and chimpanzees are much more capable tool users than gorillas: arboreal environments have produced in orangutans and chimpanzees (gorillas are ground-dwellers) a flexible kinematic body image that can be adapted to tool-using activities.

Experimentation with tool use in animals at the macro-level is complemented by results in cognitive neuroscience. In one of the best-known studies, Iriki et al. (1996) monitored the neuronal activity of Japanese monkeys (*Macaca fuscata*), in an examination of the internal representation of the body during tool use. The monkeys were required to use a rake tool to reach for food pellets placed at the end of a table. Results demonstrate convincingly that the body schema of a tool user extends beyond the boundaries of the physical body to include the tool. One important qualifier in their findings is that the body schema is only modified when the monkey intends to use the rake as a tool, not during passive grasping of the tool. In other words, when the rake is held in the hand during tool use, the body schema is subject to expansion, whereas afterwards, even with the tool still in hand, the body schema reverts to its usual state.

As other studies show, for humans as well as other animals, an internal body schema is not constrained to modeling morphological boundaries, but can be extended to incorporate objects attached to bodies, such as clothes, ornaments, and tools (Berlucchi & Aglioti 1997). For instance, when one becomes skilled with a hammer, the hammer feels like part of one's hand or arm—this turns out to be true from a neural perspective as well as a macro-level behavioral perspective. Baber (2003) describes this phenomenon in terms of tools being "forgotten" in use, when expert users are involved; Clark (2002) refers to "transparent equipment."

This is just a small sampling of research on animal cognition and tool use; Beck (1980) and Tomasello and Call (1997) give excellent introductions to animal and primate tool use respectively.[1] Our purpose in reviewing such results is to highlight the common concerns of animal cognition and AI researchers, and to suggest the benefits of cross-fertilization for evaluating the intelligence of artificial agents. In addition to the AI-centric observations made about animal cognition work above, there are more general areas of overlap.

- Animal cognition researchers are acutely aware of the pitfalls of wishful thinking. Much of Povinelli's research, for example, is motivated by the possibility that some general problem-solving skills attributed to chimpanzees can

[1] Work on the evolution of human cognition and tool use is also relevant (Deacon 1998; Gibson & Ingold 1993; Mithin 1996; Oswalt 1973; Sterelny 2003). Much of this work can contribute to AI research, though we find very few references to it outside the philosophy of AI (e.g. (Chalmers 2001, parts 6.4c and 6.4d)).

be explained by simpler mechanisms (by analogy, in the same way that stimulus response agents are simpler than planning agents.) His experiments are designed to tease out and eliminate various potential explanations for sophisticated surface behavior.

- Animal cognition research is also careful in distinguishing different capabilities in non-human intelligence. For example, while chimpanzees are generally smarter than dogs, dogs outperform chimpanzees in some tasks that for humans require the mental capacity to track invisible displacement of objects (Tomasello & Call 1997). That is, human development is a useful guide for understanding the scaffolding necessary for human intelligence, but some level of intelligence can be reached by taking different paths.
- Animal cognition research has produced a wide range of tasks that are common across the field. Food tube experiments are one example. Working with such common evaluation environments is now commonplace in subfields of AI such as planning, machine learning, and message understanding. A lesson for AI is the depth to which experimentation is pursued on a single evaluation apparatus. This supports generalization over subjects with comparable physical capabilities and allows discrimination between similar behaviors at a fine level of detail. Evaluation environments for artificial agents should provide a comparably rich set of related, incrementally modifiable tests.
- As suggested by the points above, research on animal cognition is strongly theory-driven. For example, a refinement of the food tube adds a hole through which treats drop if they are pushed from the wrong direction—solving the problem requires a grasp of the causality of directional pushing. A further refinement inverts the tube so that the hole is on top, rending it irrelevant—how the problem is solved (or not) can indicate a causal rather than merely associational understanding of the presence of the hole. The match with AI research is obvious.

As valuable as the results of animal cognition research are in understanding naturally arising intelligence, they provide only limited insight into the context of tool use. Most tool-using ability in animals is targeted at the acquisition of food, though more intelligent animals such as chimpanzees and elephants also use tools for grooming and self-defense. For artificial tool-using agents, we are interested in a more general characterization of behavior independent of what is sometimes termed the "material culture" of particular species (McGrew 1992). In general, how do agents apply their capabilities, in the context of tool use?

We find an answer in taxonomies of tool-based action. There exist extensive taxonomies of tools based on application or engineering principles, but these do not capture important aspects of tools in use. For example, the U.S. Patent Office describes wrenches as "engaging a work part and exerting or transmitting a twisting strain thereto, or means for imparting or transmitting an actuating force to such a tool"(as quoted in (Hohn 2005, p. 60).) Taxonomies of human tools for subsistence (e.g., spears and fishing implements) due to Oswalt (1973) come closer, but are too abstract. Baber gives an extensive overview of tool use and cognition from a human factors perspective (Baber 2003), but no general taxonomies of action. Perhaps most promising is a small taxonomy that describes software tools as metaphors for physical tools (St. Amant & Horton 2002):

- *Effective* tools produce a persistent effect on materials or the environment, e.g., hammers, screwdrivers, and saws.
- *Instruments* provide information about materials or the environment, e.g., calipers and magnifying glasses.
- *Constraining* tools constrain or stabilize materials for the application of effective tools, e.g, clamps and straight edges.
- *Demarcating* tools impose perceivable structure the environment or materials, e.g., a carpenter's pencil or pushpins.

To shift perspective to the actions of the agent, the agent uses effective tools to transform its actions on the environment and uses instruments to transform its perception of the environment. The agent uses constraining tools to stabilize the environment for its effective actions and uses demarcating tools to add information to the environment in aid of perception, by creating external representations. These categories of action can guide the selection of tasks by which we can evaluate the competence of a given agent and judge whether its design is functionally adequate.

A "tooling test"

We can take the taxonomy given above to describe the tool-using behavior of a sophisticated autonomous agent. Consider an urban search and rescue robot some years from now:

> In the course of exploring a ruined building, the robot stops in the entrance to a room, blocked by a gap in the flooring too wide for its wheels to cross. It retreats to the hallway, where it has previously seen loose panels of wall board. It pushes a panel over the gap and crosses into the room. The robot's first target is a pile of rubble near the entryway. The pieces are heavy, but a nearby length of pipe serves as a good lever to lift some of the debris aside. The robot continues its search when nothing of interest is found. Because lighting is variable and navigation uncertain, the robot marks each region of the room it has examined with a splash of fluorescent paint; this makes it easy for the robot to tell what remains to be done, and it informs other searchers (human or robotic) that the room has been explored. On completing its inspection, the robot sprays a fine mist of particles into the air, looking for are air currents that may indicate otherwise hidden passages. No movement. The robot continues to another area of the building.

This somewhat artificial example illustrates use of each of the classes of tools given in the previous section. While perhaps useful as a challenge problem, the scenario given above is far too specific (and too difficult given the current

state of the art in robotics) to guide research toward intelligent tool-using ability. Instead, we a propose a multi-stage test for intelligence, in which each stage that is passed gives evidence for greater intelligence than needed in the previous stage. The test will involve the subject manipulating physical tools in order to accomplish specific tasks; obstacles will be set in the way that can be overcome by the use of these tools. We call this test the tooling test.[2]

Just as with Turing's imitation game (Turing 1950), we need to establish a few basic assumptions about the subject of the tooling test, whom for convenience we will refer to as T. First, we need to have some understanding of T's physical capabilities. If T were a small child, for example, it would not be reasonable to give him a forty pound sledge hammer and the task of driving railroad spikes. If T is nonanthropomorphic in its physical or sensory capacities, we need to be even more careful about making inferences based on what T can do (sometimes described as the effectivities, or action capabilities, of an agent (Turvey & Shaw 1979).) Tools reflect the capacities and limitations of their users as much as they do the functions they are designed to facilitate. Second, we need to know something about the goals that T might plausibly have. These might involve a desire to gain food or other sustenance, desire for a reward of some type, or perhaps a desire to leave the experimental setup. This assumption is almost universal in animal cognition experiments, and even if T is, say, a mechanical robot, it will eventually be necessary for T to recharge itself. In our discussion, for convenience, we will use acquisition of an object as a generic placeholder for the actual goals that T may have. Third, as with the imitation game, we assume that T is basically cooperative, willing to go along with our experiments. We can draw no conclusions from a test subject that sits like a stone in the middle of a room.

Given these assumptions, we can define a sequence of incremental tests of T's tool-using capabilities.

Simple tool use. In this stage, T must apply simple tools to reach its goal. These simple tools are physical extensions of T's unaided effectivities. For example, the object to be acquired might be placed out of reach, but it can be retrieved by poking it with a rod. We refer to this as an example of simple tool use only partly because the tool itself is simple; more importantly, the behavior of the tool is transparent to the tool user. Other scenarios that entail simple tool use might involve placing the target object in a transparent container that can only be opened by using a lever to lift a heavy lid, or by using a hammer to force a peg out of position, or by using forceps for manipulating an otherwise inaccessible knob. In each case, T is prevented from carrying out a task because of physical constraints, but an appropriate tool is available to overcome these constraints, and the mechanics of the tool-based activity are observable as the tool is used.

These sorts of tool-using tests are a staple in animal cognitive research, as described above. The classic monkey-and-bananas problem is probably the best known of these. Results have helped researchers to better understand the extent to which animals can reason about spatial relationships, kinematics and dynamics, qualitative physics, and so forth, in comparison with humans. By increasing the complexity of the tasks to be performed (we can imagine arbitrarily complicated Rube Goldberg devices that must be worked through) T has the opportunity to display more and more sophisticated tool-using behavior.

Tool construction. Construction of novel tools and refinement of existing tools is only rarely seen among animals, and on an evolutionary scale it is a relatively recent capacity for human beings. We can test T's capabilities in this regard by providing it with tools suitable for making or modifying other tools. For an example of tool modification, suppose that the object acquisition task involves turning a bolt. T tries to turn the bolt with a wrench, but the wrench has too short a handle to provide sufficient leverage. A length of pipe is also nearby, which T uses as a sleeve to extend the wrench handle and thus completes the task. Examples of tool construction are comparable, working from raw materials rather than tool components. (Note that one of the earliest examples of tool construction, rock knapping, is a demanding skill for modern humans to acquire, but more sophisticated tool-making tools can make such tasks easier.)

Surprisingly, chimpanzees and humans are not the only creatures capable of tool modification. For example, Sterelny (2003) describes crows able to bend a piece of wire into a shape appropriate to fish out food from a test apparatus. In the end, it is not clear that there is a conceptual difference between tool use directly on materials and tool use for making other tools, but this turns out to be a soft boundary between human and animal tool users.

Non-transparent tool use. A more difficult set of problems arises in scenarios in which the mechanisms of a tool, in its operation or its effect, are partially or wholly hidden. For example, imagine picking a simple lock, or working the levers of a complex machine with an opaque housing, or even fishing for a wire inside a wall by feeding an electrician's snake through a small hole.

What we gain from posing such problems is the opportunity to observe T's exploratory strategies. Search is often held to be a core aspect of intelligent behavior. Comparable tasks are used in studies of human development and animal cognition. For example, one experiment design involves hiding an object in a container, and then moving this container between a series of other containers in different locations, with the object being deposited surreptitiously in one of those containers. An unsophisticated search strategy will examine containers in locations that have not been visited; this strategy is followed by most primates (but, surprisingly, not domesticated dogs, as mentioned above.) Piaget (1955) has described solving such problems as Stage 6 behavior, associated with mental representations. In these cases, T must reason about what is happening out of sight; behavior that relies purely on observations, without inference, will be less successful, or at least less efficient, than behavior that relies on some understanding of what can possibly be happening behind the scenes.

[2]The coinage and the idea of casting the evaluation of tool-using agents in terms of a general test of intelligence are due to Ronald Endicott (personal communication).

There are two extensions possible here. First, let's suppose that the problem is so difficult that blind search will take T quite a long time before reaching success. Enter a collaborator, C. C carries out the appropriate action or sequence of actions to accomplish the task. In our tool-using examples above, C might move the lock pick into a specific sequence of configurations, or move specific levers in order, or move the electrician's snake along a particular path. If T can solve comparable problems after having observed C's action, we count this as an example of learning, a key facet of intelligence. Again, we can vary the internals of a given apparatus to test the extent to which T has generalized its observations, and we can make the learned procedure as complex as desired.

Second, we can introduce uncertainty to the problems. Suppose that the tool or environment responds only probabilistically to the correct manipulation. Does T learn despite failure? Such learning again indicates some level of intelligence.

Collaborative tool use. A sophisticated ability to solve problems using tools, either alone or after instruction, captures many of the characteristics we associate with intelligence: the ability to reason about causality, physics, spatial relationships; planning and learning; even creativity and opportunism. One aspect of intelligence that is missing, however, is the ability to communicate. This is addressed in a stage that involves our collaborator as a part of all tasks. C attempts to solve a task but cannot succeed, lacking some crucial physical aid or tool. For example, acquiring a target object might involve prying up a lid that is too heavy or large for one to lift alone, turning a nut on a bolt that rotates freely unless held in place, and so forth. T passes this test by observing C's actions, recognizing C's implicit goal, and contributing appropriately to reach the goal.

There are three variations of interest in these scenarios. In one, C works toward one of T's goals; in anoother, C pursues a goal to which T can be assumed to be indifferent. Both of these variations require T to effectively assume C's perspective for the purpose of goal recognition, but the latter case is more difficult in that C's goal is one that T would not ordinarily have. The third variation is the most interesting: some problems may only be solvable, given T's limitations, if C directly controls T's actions (with T acting partially autonomously.) This variation is subject to the most difficulty in assigning credit for intelligent behavior—a difficult but manageable problem for researchers in variable autonomy and mixed-initiative systems.

Discussion

We can make the relationship between the Turing Test and the Tooling Test explicit as follows: in each of the above scenarios, we can substitute for T a robot remotely controlled by a human operator. Given sufficient practice, an experienced operator will be able to reach some level of competence in solving the problems posed, even given the limitations of the robot's hardware and the communication link. An independent observer, or even the collaborator in the experiments, could then decide whether the robot were acting on its own or under the complete or partial control of another human.

Viewed in this way, the tooling test is clearly derivative of the Turing Test and its variants. It has some similarities to Harnad's Total Turing Test (Harnad 1992), in that it involves physical interaction with the environment, but it is more restrictive than the Turing Test. The activities required of the agent are well circumscribed: solve this test, and then solve another test. Although some tests involve cooperation, none requires language.

Despite the restrictions on generality and the use of language, the tooling test offers some advantages over the Turing Test. Based on Paul Cohen's criteria for a good test of intelligence (Cohen 2004), we can say the following:

- The tooling test is a reasonable proxy function. To pick a fanciful example, if aliens descended on Earth in a space ship, we would have strong evidence for their intelligence even if we were unable to communicate with them or even to observe them communicating with each other. Their mastery over their physical environment would be sufficient to indicate much broader competence.
- The tooling test is specific. Different instances of the test require the exercise of different types of intelligent behavior, as discussed in the previous section. Further, the tests are independent of culture. A conversant may fail the Turing Test by being unaware of current events, the connotations of words (even syllables), the subtleties of human relationships, and so forth, but this is hardly convincing evidence for lack of intelligence. While Turing expected machines successful in the imitation game to be reasoning the way that humans do, in these tests we are willing to accept completely unexpected solution paths as evidence of intelligence.
- Perhaps most importantly, the tooling test is diagnostic:
  - Each trial has an explicit success criterion: has a goal been achieved? This obviates some traditional concerns about the validity of the Turing Test (e.g., Is ten minutes enough for a conversation? How sophisticated must the human conversant be?)
  - Like the Turing Test, the tooling test can be extended to as great a sophistication as desired. The equivalent of a deep discussion of King Lear might be the construction and operation of a complex piece of equipment, based on exploration, observation, and learning.
  - The result of a sequence of tests is not a binary intelligent/non-intelligent judgment. Rather, just as in research on animal tool use, we can say that a tool-using agent is intelligent in this way or that, based on the sophistication of their actions and our theories about their underlying reasoning processes.
- A simple version of the tooling test can be solved today, with current technology, as demonstrated in existing systems (Bluethmann *et al.* 2003).

Is the tooling test practical? One potential barrier is that problems must be constructed such that the subject cannot solve them without the use of tools. Analyzing the physical capabilities of a habile agent, even a designed agent, is

non-trivial. In animal tool use, this leads to experiments that are specialized to classes of subjects (food tube experiments would be inappropriate for, say, a hypothetical intelligent mouse.) Nevertheless even simple tasks can be informative with respect to the basic functional components required in a habile agent; for example, our experiments with a tool-using Aibo (which we call *Canis habilis*) make use of an out-of-reach object and a graspable rod, which should generalize to a wide range of agent platforms. The broader question of identifying a comprehensive set of specific tool-using tasks, however, remains open.

From a philosophical standpoint, the tooling test is further subject to some of the same objections that the Turing Test has raised, most obviously that it takes a behaviorist approach to the attribution of intelligence to an agent. We believe that this concern is alleviated by the use of comparable tests in animal cognition. Researchers have long since abandoned the notion that animals are automata, and assess different aspects of the intelligence of non-humans without qualifying their observations as being of the mere imitation or simulation of intelligence. Chimpanzees, elephants, and so forth have some measurable degree of intelligence; we can apply the same metrics to artificial agents.

Acknowledgments

Thanks to Paul Cohen and Adele Howe for their critiques of this research, to Ronald Endicott for his collaboration during its formative stage, and to Daniel Dennett for pointing us to work on the evolution of human cognition. This effort was partially supported by the National Science Foundation under award IIS-0083281. The information in this paper does not necessarily reflect the position or policies of the U.S. government, and no official endorsement should be inferred.

References

Agre, P., and Horswill, I. 1997. Lifeworld analysis. *Journal of Artificial Intelligence Research* 6:111–145.

Baber, C. 2003. *Cognition and Tool Use: Forms of Engagement in Human and Animal Use of Tools*. Taylor and Francis.

Beck, B. B. 1980. *Animal Tool Behavior: The Use and Manufacture of Tools*. New York, NY: Garland Press.

Berlucchi, G., and Aglioti, S. 1997. The body in the brain: neural bases of corporeal awareness. *Trends in Neuroscience* 20(12):560–564.

Bluethmann, W.; Ambrose, R.; Diftler, M.; Askew, S.; Huber, E.; Goza, M.; Rehnmark, F.; Lovchik, C.; and Magruder, D. 2003. Robonaut: A robot designed to work with humans in space. *Autonomous Robots* 14(2–3):179–197.

Brady, J. M.; Agre, P. E.; Braunegg, D. J.; and Connell, J. H. 1984. The mechanic's mate. *Artificial Intelligence* 72(1–2):173–215.

Chalmers, D. 2001. Contemporary philosophy of mind: An annotated bibliography [parts 6.4c and 6.4d]. WWW document: http://jamaica.u.arizona.edu/˜chalmers/biblio.html.

Clark, A. 2002. Minds, brains and tools. In Clapin, H., ed., *Philosophy of Mental Representation*. Clarendon Press.

Cohen, P. R. 2004. If not the Turing Test, then what? In *Proceedings of AAAI (invited talk)*, xvi.

Deacon, T. W. 1998. *The symbolic species: The coevolution of language and the brain*. W. W. Norton & Company.

Gibson, K. R., and Ingold, T., eds. 1993. *Tools, language, and cognition in human evolution*. Cambridge, UK: Cambridge University Press.

Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Houghton Mifflin.

Harnad, S. 1992. The Turing Test is not a trick: Turing indistinguishability is a scientific criterion. *SIGART Bulletin* 3(4):9–10.

Hohn, D. 2005. A romance of rust: Nostalgia, progress, and the meaning of tools. *Harper's Magazine* 310(1856):45–63.

Iriki, A.; Tanaka, M.; and Iwamura, Y. 1996. Coding of modified body schema during tool use by macaque postcentral neurons. *Neuroreport* 7:2325–2330.

McGrew, W. C. 1992. *Chimpanzee Material Culture*. Cambridge, UK: Cambridge University Press.

Mithin, S. 1996. *The Prehistory of the Mind: The Cognitive Origins of Art, Religion and Science*. London: Thames & Hudson.

Nilsson, N. 1995. Eye on the prize. *AI Magazine* 16(2):9–16.

Oswalt, W. H. 1973. *Habitat and Technology: The Evolution of Hunting*. New York, NY: Holt, Rinehart, & Winston.

Piaget, J. 1955. *The Construction of Reality in the Child*. Routledge and Kegan Paul.

Povinelli, D. 2000. *Folk Physics for Apes*. NY: Oxford University Press.

Preston, B. 1998. Cognition and tool use. *Mind and Language* 13(4):513–547.

St. Amant, R., and Horton, T. E. 2002. Characterizing tool use in an interactive drawing environment. In *Second International Symposium on Smart Graphics*, 86–93.

Sterelny, K. 2003. *Thought in a Hostile World: The Evolution of Human Cognition*. Oxford: Blackwell Publishers.

Stoytchev, A. 2003. Computational model for an extendable robot body schema. Technical Report GIT-CC-03-44, Georgia Institute of Technology, College of Computing.

Tomasello, M., and Call, J. 1997. *Primate Cognition*. New York, NY: Oxford University Press.

Turing, A. 1950. Computing machinery and intelligence. *Mind* 59:433–460.

Turvey, M. T., and Shaw, R. E. 1979. The primacy of perceiving: An ecological reformulation of perception for understanding memory. In Nilsson, L. G., ed., *Perspectives on memory research*. Lawrence Erlbaum. 167–222.

Visalberghi, E., and Limongelli, L. 1996. Acting and understanding: Tool use revisited through the minds of capuchin monkeys. In Russon, A. E.; Bard, K. A.; and Parker, S. T., eds., *Reaching into Thought: The minds of the great apes*.