Autonomous Learning of Tool Affordances by a Robot

Alexander Stoytchev
Mobile Robot Laboratory
Georgia Institute of Technology
Atlanta, GA 30332, U.S.A.
saho@cc.gatech.edu

Introduction

The ability to use tools is one of the hallmarks of intelligence. Tool use is fundamental to human life and has been for at least the last two million years. We use tools to extend our reach, to amplify our physical strength, to transfer objects and liquids, and to achieve many other everyday tasks. A large number of animals have also been observed to use tools (Beck 1980). Some birds, for example, use twigs or cactus pines to probe for larvae in crevices which they cannot reach with their beaks. Sea otters use stones to open hard-shelled mussels. Chimpanzees use stones to crack nuts open and sticks to reach food, dig holes, or attack predators. These examples suggest that the ability to use tools is an adaptation mechanism used by many organisms to overcome the limitations imposed on them by their anatomy. Despite the widespread use of tools in the animal world, however, studies of *autonomous* robotic tool use are still rare.

This abstract describes the empirical evaluation of one specific way of representing and learning the functional properties or affordances (Gibson 1979) of tools. A longer version of this paper appears in (Stoytchev 2005).

Behavior-Grounded Affordance Representation

The tool representation described here uses a behavior-based approach (Arkin 1998) to *ground* the tool affordances in the existing behavioral repertoire of the robot. The representation is learned during a behavioral babbling stage in which the robot randomly chooses different exploratory behaviors, applies them to the tool, and observes their effects on environmental objects. The functionality of a tool is represented with an *Affordance Table* of the form:

Grasping Behavior and its Parameters	Exploratory Behavior and its Parameters	O_{start}	O_{end}	Replication Probability

In each row of the table, the first entry represents the grasping behavior that was used. The second entry represents the exploratory behavior and its parameters. The next two entries store the observation vector at the start and at the end of the exploratory behavior. The last entry estimates the probability of replicating the observation vectors when this sequence of behaviors is performed multiple times.

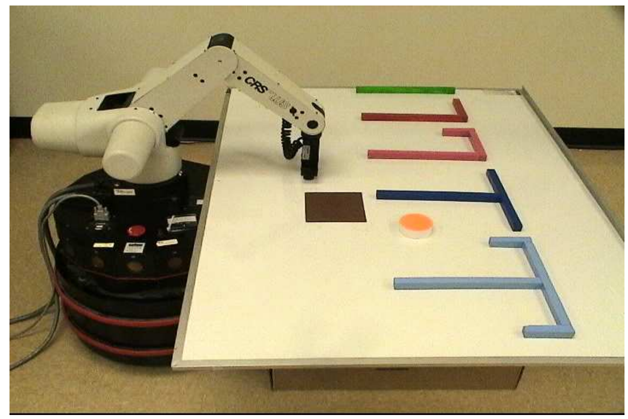

Figure 1: The figure shows the mobile manipulator and the five tools used in the extension of reach experiments.

Experimental Environment

All experiments were performed using a CRS+ A251 manipulator arm (Figure 1). The robot has 5 degrees of freedom plus a gripper. In addition to that, the arm was mounted on a Nomad 150 robot which allows the manipulator to move sideways. The robot's wrist, the tools, and the environmental object were color coded so that their positions can be uniquely identified and tracked using computer vision. The camera was mounted above the robot's working area.

Five tools were used: stick, L-stick, L-hook, T-stick, and T-hook (Figure 1). An orange hockey puck was used as the environmental object (also called attractor). The choice of tools was motivated by the similar tools that Köhler used in his experiments with chimpanzees (Köhler 1931).

Exploratory Behaviors

Five exploratory behaviors were used: *Extend arm*, *Contract arm*, *Slide arm left*, *Slide arm right*, and *Position wrist*. All behaviors used here were encoded manually from a library of *motor schemas* and *perceptual triggers* (Arkin 1998) developed for this specific robot. The behaviors result in different arm movement patterns as described below.

The first four behaviors move the arm in the indicated direction while keeping the wrist perpendicular to the table on which the tool slides. These behaviors have a single parameter which determines how far the arm will travel relative to its current position. Two different values for this parameter were used (2 and 5 inches). The *position wrist* behavior moves the manipulator such that the centroid of the attractor is at offset (x, y) relative to the wrist.

Learning Trials

During the learning trials the robot was allowed to freely explore the properties of the tools. The exploration consists of trying different behaviors, observing their results, and filling up the affordance table. The initial positions of the attractor and the tool were random. If the attractor was pushed out of tool reach then it was manually placed in a new random position. The learning time was limited to one hour per tool.

A good way to visualize what the robot learns is with graphs like the ones shown in Figure 2. The figures show the observed outcomes of the exploratory behaviors when the T-hook tool was applied randomly to the hockey puck.

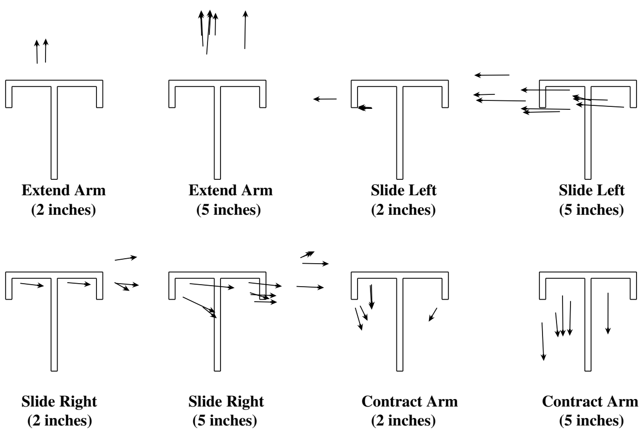

Figure 2: Visualizing the affordance table. Each figure shows the observed movements of the puck after a specific exploratory behavior was performed multiple times. The start of each arrow corresponds to the position of the puck in wrist-centered coordinates (i.e., relative to the tool's grasp point at the lower part of the handle) just prior to the start of the exploratory behavior. The arrow represents the total distance and direction of movement of the puck in camera coordinates at the end of the exploratory behavior.

Testing Trials

Two types of experiments were performed. They measured the quality of the learned representation and its adaptation abilities when the tool is deformed, respectively. During testing trials, the best affordance for a specific step in a tool task is selected using a greedy heuristic search. MPEG movies from the experiments are available at: (http://www.cc.gatech.edu/~saho/papers/AAAI_2005/).

Extension of Reach

In this experiment the robot was required to pull the attractor over a color coded goal region. Four different goal positions were defined. The first goal is shown in Figure 1 (the dark square in front of the robot). The second goal was located farther away from the robot. To achieve it the robot had to push the attractor away from its body. Goals 3 and 4 were placed along the mid-line of the table to the left and right of the robot's position shown in Figure 1.

In addition to that there were 4 initial attractor positions per goal (located along the mid-line of the table, 6 inches apart). The tool was always placed in the center of the table. A total of 80 experiments were performed (4 goals $\times$ 4 attractor positions $\times$ 5 tools). The table below summarizes the results. The values represent the number of successful solutions per goal, per tool. Four is the possible maximum.

Tool	Near Goal	Far Goal	Left Goal	Right Goal
Stick	0	2	4	4
L-stick	4	2	4	4
L-hook	4	3	4	4
T-stick	3	3	4	4
T-hook	4	4	4	4

As can be seen from the table, the robot was usually able to solve this task. The most common failure condition was due to pushing the attractor out of tool's reach. A notable exception is the *Stick* tool which could not be used to pull the object back to the near goal. The robot lacks the required exploratory behavior (*turn-the-wrist-at-an-angle-and-then-pull*) which can detect this affordance of the stick. Adding the capability to learn new exploratory behaviors can resolve this problem.

Adaptation After a Tool Breaks

The second experiment was designed to test the flexibility of the representation in the presence of uncertainties. The uncertainly in this case was a tool that can break.

To simulate a broken tool, the robot was presented with a tool that has the same color as another tool with a different shape. More specifically, the learning was performed with a T-hook which was then replaced with an L-hook. Because color is the only feature used to recognize tools the robot believes that it is still using the old tool. The task of the robot was the same as described in the previous subsection (i.e., 16 experiments = 4 goals $\times$ 4 attractor positions).

The two tools differ in their upper right sections. Whenever the robot tried to use affordances associated with the missing parts of the tool they did not produce the expected attractor movements. Thus, their probability of success was reduced and they were excluded from further consideration.

The robot was successful in all 16 experiments.

References

Arkin, R. 1998. *Behavior-based robotics*. MIT Press.

Beck, B. B. 1980. *Animal Tool behavior: The use and manufacture of tools by animals*. NY: Garland STMP Press.

Gibson, J. J. 1979. *The ecological approach to visual perception*. Boston: Houghton Mifflin.

Köhler, W. 1931. *The mentality of apes*. Harcourt, Brace, and Co.

Stoytchev, A. 2005. Behavior-grounded representation of tool affordances. In *IEEE ICRA*.