Elements of tool use for a simulated robot arm

Robert St. Amant

September 10, 2001

Our goal is to build an agent that has some limited ability for tool use. Although experimenting with a real robot would give the most compelling results, we believe that many important issues can be addressed with a sufficiently detailed physical simulation. Further, with a simulation we can explore some behaviors that may not be possible for a real robot, such as striking a surface or object with some force. Requirements for our simulated robot agent include the following:

- *The agent is embodied*—it interacts with its environment via sensors and effectors. In our case, the agent has a simulated robot arm and hand, and a very simple vision system. For our purposes, realistic vision will not be important, so we mave build a system that gives global, perfect information.
- *The agent exists in a physical environment.* The agent will be placed in a tabletop environment that contains objects and simple hand tools. Objects will have spatial extent (i.e., they take up space, rather than being points), mass, and other physical properties.
- *The agent has interoceptive (proprioceptive) sensors.* The agent can sense the force it applies to objects, as when it picks up an object or pushes against a wall. It can sense the position of its joints.
- *The agent has behaviors.* Given some directive, it can move its robot arm appropriately. It need not learn everything from scratch.
- *The agent has "reflexes."* We will build into the agent specific behaviors for when it encounters specific situations: on touching an object from a specific angle, for example, the hand should try to grasp the object.
- *The agent has preferred positions and movements.* Just as it requires less effort for us to hold our arms in some positions than others, or move along some paths than others, the robot arm should have comparable properties, whether effort is measured by cost or force or something more arbitrary.
- *The agent can take open-loop or closed-loop actions.* For example, the agent may be directed to move its arm to a particular position, using feedback to achieve the goal, or instead may be directed to achieve some minimum velocity in some direction without regard for the final position, ignoring feedback about intermediate positions.

Our experimentation with this agent will be influenced by work in ecological psychology; we concentrate on "interactional" properties, properties of the relationship between the agent and its environment that arise from their interaction. A useful class of interactional properties are those found in the affordance literature [St. Amant, 1999]: reachable (i.e., an object is within reach of the agent), graspable (i.e., an object is small enough and of a shape to be held in the hand), movable (i.e., an object is light enough to be lifted), and so forth. These can be combined to create new properties; in fact, there is some ambiguity in these concepts, for example, whether graspable implies reachable. These properties are interesting because an agent might plausibly learn them by exploration, and thus address part of the symbol-grounding problem, associating symbols in a reasoning system with objects and properties in the real world.

Some work in robotics has taken an ecological perspective [Arkin *et al.*, 1997; Arkin, 1998; Murphy, 1999], but these concentrate on issues other than manipulation, such as vision and navigation. Projects of Mataric, Steels, and Brooks are nevertheless relevant.

Here is the progression by which we expect the agent to learn interactional properties and eventually to use simple tools.

1. We will walk the agent will through an exploration of the environment, such that it moves its arm through through the simulated space (i.e., vary its joint angles through configuration space) in an exhaustive fashion, at some particular level of resolution. When the agent encounters an object, its reflexes may cause it to interact with the object, depending on the details of the encounter. (See Cohen et al.'s NEO system for a symbolic version of this process [Cohen *et al.*, 1996; Cohen *et al.*, 1997].)

2. We will record the actions of the agent and its interactions with the environment, again at some particular resolution. The result is a dataset of behavioral traces. The information recorded will include such things as the forces on the joints over time, the time at which an object or surface was encountered, the effect this has on the agent's movement, etc.

3. We will select and run an unsupervised clustering algorithm on the traces. (NEO used a contingency table learning method called MSDD; we'll have to look around for an appropriate clustering method. In the simplest case, we might just use a standard statistical techniques such as single linkage [Everitt, 1993; Sibson, 1973], but some numerical form of conceptual clustering would be better [Fisher, 1987; Michalski and Stepp, 1983].) Below is an incomplete list of the kinds of clusters we hope to identify. (We don't expect the algorithm to produce a symbolic representation of its clusters, of course; this is our external view of what it is doing.)

   - Unrestricted movement.
   - Hand touches/grazes a surface (light force).
   - Hand strikes a surface (moderate to strong force, open-loop).
   - Hand pushes a surface (light to strong force, closed-loop).
   - Hand touches an object.
   - Hand grasps an object (reflex).
   - Hand pushes an object (unlike a surface, an object moves when pushed.)
   - Hand grasps an immovable object.
   - Hand lifts a grasped object (reflex plus action.)

4. We will develop an algorithm to infer a parameterized, schema-based representation of the actions in each cluster. (This could be an extension of the clustering algorithm, but it's not clear yet; qualitative reasoning research may be relevant here.) The schemas will allow us to direct the agent, for example, to touch a given surface at a specific point. The representation need not be perfect, but it does need to be generalized directly from the clustered data and to reflect internal sensed data specialized to the clusters (e.g., force on the gripper for a grasped object, weight of a lifted object, etc.) The representation should also allow us (and the agent) to identify the most relevant sensory or interactional properties associated with individual schemas; e.g., striking an object is similar to unrestricted movement except for the force applied to the object at the end.

5. We will now introduce tools into the environment. We will put a tool in the agent's hand, and tell it to carry out one of the actions it has learned. For example, we will put a hammer in the agent's hand and direct it to strike a surface or object. The agent will be able to sense and record the effect of the action using the hammer just as it did without the hammer.

6. We will develop an algorithm to allow the agent to associate tools with specific schemas. For example, if we have a strikingasurface schema in the force at the end of the movement is its identifying characteristic, then the agent should detect that having a hammer in the hand increases this force. This association should allow us to ask the agent questions such as, "Can the object you are holding be used as a hammer?"

The step above concerning the introduction of tools is based on our current understanding of simple physical tools not as amplifiers of existing abilities [Hutchins, 1995] but as functional transformations of forces or structural constraints. Many simple tools *do* act as amplifiers; for example, a hammer increases the force of a blow, and a screwdriver increases the twisting force of unaided fingers by allowing different sets of muscles to apply to the problem. The effect of such tools can also be seen as a transformation from one range of outputs to another. Transformations work in the opposite direction as well, however. For example, a long-nosed, short-handled pair of forceps will exert less force than unaided fingers, but allow longer reach, which may be useful in surgical tasks. A pair of tweezers may reduce the force of unaided fingers, but can work in a tighter area with smaller objects. A socket wrench, as an alternative to a socket driver, translates approximately linear force (over a short range) into angular force. A torque wrench is another interesting case: it restricts the range of force applied to a bolt by imposing a maximum on it. What our agents are doing is using this basic insight (if that is what it is) to guide their inference about what constitutes a tool or not.

The question of "Can this object be used as a hammer?" is also interesting. It relies on several related entities: the agent, the object, and the environment. (It also implicitly relies on the task, but we've constrained that to be the strikingasurface task.) What will the agent have to do to see if an object can be used as a hammer?

- Grasp it. It may be too large or small. (We may forgo this test, by the way, if it's too difficult to arrange for sequences of actions by the robot arm.)

- Lift it. It may be too heavy.

- Swing it. It may produce no additional force. It may be too long to be moved freely in an enclosed environment.

- Strike a specific point. The point may not be reachable by the arm, or reachable only through a high-effort, low-force movement.

With this approach we are neglecting a number of important processes in the general use of tools. These include the following:

- Reasoning about causal relationships between tools and objects. These kinds of relationships will emerge out of their interaction, but the agent won't produce explicit representations of them or be able to reason about them.

- Reasoning about extended sequences of actions.

- Reasoning about relationships over time.

- Reasoning about finely detailed motions.

There are other areas as well. Nevertheless, if we can complete all of what is described above, it should be a compelling story.

References

Arkin, Ronald C.; Cervantes-Perez, F.; and Weitzenfeld, A. 1997. Ecological robotics: A schema-theoretic approach. In Bolles, R.C.; Bunke, H.; and Noltemeier, H., editors 1997, *Intelligent Robots: Sensing, Modelling and Planning, eds*. World Scientific.

Arkin, Ronald C. 1998. *Behavior-based Robotics*. MIT Press.

Cohen, Paul R.; Oates, Tim; Atkin, Marc; and Beal, Carole R. 1996. Building a baby. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*. 518–522.

Cohen, Paul R.; Atkin, Marc; Oates, Tim; and Beal, Carole R. 1997. Neo: Learning conceptual knowledge by sensorimotor interaction with an environment. In *Proceedings of the First International Conference on Autonomous Agents*. 170–177. To appear in Intelligence Magazine.

Everitt, Brian 1993. *Cluster Analysis*. John Wiley & Sons, Inc.

Fisher, D. H. 1987. Knowledge acquisition via incremental conceptual clustering. *Machine Learning* 2:139–172.

Hutchins, Edwin 1995. *Cognition in the Wild*. MIT Press, Cambridge, MA.

Michalski, R. S. and Stepp, R. E. 1983. Learning from observation: Conceptual clustering. In Michalski, R. S.; Carbonell, J. G.; and Mitchell, T. M., editors 1983, *Machine Learning: An artificial intelligence approach*. Morgan Kaufmann Publishers, Inc.

Murphy, Robin 1999. Case studies of Gibson's ecological approach to mobile robots. *IEEE Transactions on Systems, Man and Cybernetics* 29:1.

Sibson, R. 1973. Slink: An optimally efficient algorithm for the single-link cluster method. *Computer Journal* 16(1):30–34.

St. Amant, Robert 1999. User interface affordances in a planning representation. *Human Computer Interaction* 14(3):317–354.