Toward Learning the Binding Affordances of Objects: A Behavior-Grounded Approach

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

Abstract

This paper introduces a developmental approach to learning the binding affordances of objects by a robot. A behavior-based framework is used to *ground* the affordance representation in the behavioral repertoire of the robot. The affordances are learned during a behavioral babbling stage in which the robot randomly chooses sequences of exploratory behaviors, applies them to the objects, and detects invariants in the resulting set of observations. The invariants are calculated relative to the robot's body. The approach was implemented and tested in a dynamics robot simulator.

Introduction

A simple object like a stick can be acted upon in various ways. For example, a stick can be grasped, pushed, thrown, broken, chewed, scratched, etc. It is still a mystery how animals and humans learn these affordances and what are the cognitive structures used to represent them.

The term *affordance* was first introduced by James Gibson (1979). Gibson defines affordances as "perceptual invariants" that are directly perceived by an organism and enable it to perform tasks (Gibson 1979). Gibson is not specific about the way in which affordances are learned, but he suggests that some affordances are learned in infancy when the child experiments with external objects. Furthermore, he suggests that object affordances are learned in relation to the capabilities of the learner's body. For example, an object might be graspable for an adult, but may not be graspable for a child. Therefore, Gibson suggests that a child learns "his scale of sizes as commensurate with his body, not with a measuring stick" (Gibson 1979, p. 235).

The autonomous exploration of external objects is also given a prominent role in Piaget's theory of child development (Piaget 1952). According to Piaget, intelligent behaviors are first developed in the process of interaction with objects. His theory divides the first two years of development into six stages; the role that external objects play in the development of the child increases with each additional stage. Piaget even suggests that the mathematical abilities of humans have their origins in object interaction as children first learn the concept of a number by counting external objects.

This paper describes preliminary work toward a developmental approach for learning the binding affordances of objects. The term *binding* is used to denote affordances that allow a robot to attach an object to its body so that the object's movements can be controlled reliably by the robot. Binding affordances should be distinguished from *output* affordances which allow a robot to use an object to act on another object, i.e., to use the first object a as a tool (Stoytchev 2005).

The affordance representation described here uses a behavior-based approach (Arkin 1998) to *ground* the affordances of objects in the existing behavioral repertoire of the robot. The affordances are learned during a behavioral babbling stage in which the robot randomly chooses different exploratory behaviors, applies them to the objects, and detects regularities (or invariants) in the resulting set of observations. The observations are expressed relative to a well known reference point: the robot's body. The body of the robot is represented using a *robot body schema* model (Stoytchev 2003) which is also learned developmentally.

The computational model described here was inspired by the theories of Gibson and Piaget. It should be noted, however, that it does not attempt to implement either of these theories. Instead, it is intended for use by autonomous robots. The model was implemented and tested in a dynamics robot simulator.

Related Work

Gibson

Gibson divides environmental objects into two main categories: attached and detached. *Attached objects* are defined as substances "partially or wholly surrounded by the medium" which cannot be displaced without becoming detached (Gibson 1979, p. 241). *Detached objects*, on the other hand, are objects that can be displaced; they are portable and afford carrying.

Detached objects must be comparable in size with the animal under consideration in order to afford behavior. For example, an object is graspable if it is approximately "hand size" (Gibson 1979, p. 234) or has opposable surfaces the distance between which is less than the span of the hand (Gibson 1979, p. 133).

Gibson also seems to suggest that affordances are learned by detecting perceptual invariants and linking them to the behaviors that were executed while the invariant was perceived. For example, an object affords throwing if it can be grasped and moved away from one's body with a swift action of the hand and then letting it go. The perceptual invariant in this case is the shrinking of the visual angle of the object as it is flying through the air. This highly interesting "zoom" effect will draw the attention of the child (Gibson 1979, p. 235). The behavioral sequence that reliably reproduces this invariant is a grasping behavior followed by a throwing behavior.

The computational model described in this paper relies on perceptual routines that detect invariants in the movements of objects relative to the body of the robot. Sequences of exploratory behaviors that can reliably reproduce a perceptual invariant are used to represent the affordances of objects.

Piaget

Piaget's theory divides the first two years of human life into six distinct stages (Piaget 1952). With each additional stage, the behaviors of the child progress from simple to more intelligent ones. The role that external objects play in the development of the child also increases with each additional stage. A complete review of Piaget's theory is beyond the scope of this article. However, a brief summary of the first three stages of development is provided below since our method for learning object affordances resembles the secondary circular reactions (stage III) in Piaget's theory.

Stage I: Reflex Structures (0-1 Month)

Piaget suggests that at birth children have no cognitive structures. Instead they have reflex structures for sucking, grasping, and crying. For example, newborn children close their hands when their palms are touched. Similarly, children start sucking any object that comes into contact with their lips (Piaget 1952, p. 89).

Stage II: Primary Circular Reactions (1-4 Months)

The infant's reflex structures are gradually transformed into sensorimotor action schemas, which Piaget calls *primary circular reactions*. This happens after repeated use of the reflex structures, which the baby would apply to any object. For example, babies would grasp blankets, pillows, fingers, etc. Stage II infants, however, are not concerned with the objects around them and would not pay attention to the effects of their actions on the external world. They would execute an action even if it is not applied to any object. It is not uncommon for them to open and close their hands in mid-air. The repeated use of the action forms the primary circular reaction.

Stage III: Secondary Circular Reactions (4-8 Months)

At the end of Stage II, infants are more capable of exploring their world. They can form associations between their actions and the results produced in the external environment. The child actively tries to reproduce or prolong these results. Through this repetition the child discovers and generalizes behavioral patterns that "produce and make interesting sights last" (Piaget 1952, p. 171). Piaget calls these behavioral patterns *secondary circular reactions*.

Related Work in Robotics and AI

Krotkov (Krotkov 1995) notes that relatively little robotics research has been geared towards discovering external objects' properties other than shape and position. Some of the exploration methods employed by the robot in Krotkov's work use robot behaviors coupled with sensory routines to discover object properties. For example, the "whack and watch" method uses a wooden pendulum to strike an object in order to estimate its mass and coefficient of sliding friction. The "hit and listen" method uses a blind person's cane to determine the acoustic properties of objects.

Fitzpatrick et al. (2003) used a similar approach to program a robot to poke objects with its arm and learn the rolling properties of the objects from the resulting displacements. They used a single poking behavior parameterized by four possible starting positions for the robot's arm. The robot learns a model of how each object slides (e.g., cars tend to slide in the direction of their elongated axis while balls can slide in any direction).

The most common form of binding is achieved through grasping behaviors. The robotics literature offers numerous examples of robotic grasping of objects (Cutkosky 1989; Stansfield 1991; Pollard 1996). However, most of these studies have approached the problem from an engineering perspective rather than a developmental perspective. It is not uncommon to find studies which attempt to find a formula for the best grasp points before the robot has even touched the object. While the merits of the two approaches can be debated it is clear that living organisms learn the best way to grasp objects through active trial and error.

Piaget's theory has also inspired the research of Drescher (1991) and Edelman (1987).

Behavior-Grounded Representation of Binding Affordances

Justification

The related work on animal object exploration indicates that animals use stereotyped exploratory behaviors when faced with a new object (Power 2000; Lorenz 1996). These behaviors are species specific and may be genetically predetermined. For some species of animals these exploratory behaviors include almost their entire behavioral repertoire: "A young corvide bird, confronted with an object it has never seen, runs through practically all of its behavioral patterns, except social and sexual ones." (Lorenz 1996, p. 44).

Thus, the properties of an object that an animal is likely to learn are directly related to the behavioral and perceptual repertoire of the animal. Furthermore, the learning of these properties should be relatively easy since the only requirement is to perform sequences of exploratory behaviors and observe their effects. Based on the results of these "experiments" the animal builds an internal representation of the object and the actions that it affords.

This paper takes a similar approach to learning the binding affordances of objects by a robot. A set of exploratory behaviors is used to *ground* the affordances of all objects to which the robot is exposed. Sequences of exploratory behaviors that can reliably reproduce a perceptual invariant are learned autonomously using a behavioral babbling technique (essentially a random walk through the set of exploratory behaviors). The perceptual invariants are expressed as visual functions that detect regularities in the movements of objects relative to the body of the robot. For example, the quality of grasps can be evaluated by shaking the robot's hand. A grasp is good if the grasped object moves in the same way as the wrist. A grasp is not good if the object's movements (or lack of movements) are not be correlated with the movements of the robot.

The main advantage of this approach is that objects' affordances are expressed in concrete terms (i.e., behaviors and object movements relative to the robot's body) that are directly available to the robot's controller. Because of this representational choice, the robot can autonomously learn the affordances of new objects.

Furthermore, the robot can autonomously verify the affordance representation of familiar objects in case some inconsistencies develop over time. For example, if a familiar object becomes deformed it may no longer be graspable by the robot. However, the robot can directly test the accuracy of its representation by executing the same set of exploratory behaviors that was used in the past. If any inconsistencies are detected in the resulting observations they can be used to update the object's representation. Thus, the accuracy of the representation can be directly tested by the robot.

Theoretical Formulation

The previous sub-section presented a justification for the *behavior-grounded* affordance representation. This section formulates these ideas using the following notation.

Let the body of the robot be represented as a set $\mathcal{B}$ of k rigid bodies, where $\mathcal{B} = \{B_1, B_2, \ldots, B_k\}$. The bodies are connected with a set of joints $\mathcal{J}$ which impose limits on their movements. A set $\mathcal{F} = \{F_1, F_2, \ldots, F_k\}$ of Cartesian frames is also defined such that every rigid body B_i has an associated frame F_i.

Also, let there be a set $\mathcal{L}$=$\{L_1, L_2, \ldots, L_l\}$ of l distinct body locations that are distributed along the surface of the robot's body. Each body location L_i has an associated visual signature S_i that can be identified by the robot's vision system. Furthermore, each body location is associated with one coordinate frame as given by the many-to-one mapping function $\pi : \mathcal{L} \rightarrow \mathcal{F}$. For each body frame F_i there is a function $\mathcal{C}_i$ which converts camera-centric coordinates into frame F_i coordinates. For example, $\mathcal{C}_i(S_j) \rightarrow [X, Y, Z]_{F_i}$ returns the coordinates of the visual signature S_j expressed in frame F_i coordinates.

Distinct locations are also defined for environmental objects. These locations represent local features of the object (e.g., corners) that can be identified easily. Let each object have at least w distinct locations $\mathcal{L}^o = \{L_1^o, L_2^o, \ldots, L_w^o\}$. Let each location L_j^o have a visual signature S_j^o. The positions of these locations can be tracked by the robot's vision system and their coordinates can be expressed in any body frame F_i using the conversion functions defined above, i.e., $\mathcal{C}_i(S_j^o) \rightarrow [X, Y, Z]_{F_i}$. Furthermore, let the identity of each environmental object be determined by a set of perceptual features $f_1, f_2, \ldots, f_h$. It is assumed that these features can be identified by the robot's vision system and can be used for object recognition.

The robot's perceptual routines continuously provide a stream of observations in the form of an observation vector $\mathcal{O}(t)$ = $[o_1(t), o_2(t), \ldots, o_n(t)]$. The observation vector contains information about the current positions of the robot's body locations, the positions of the object's locations, and information about the features of the object.

Let $\beta_1, \beta_2, \ldots, \beta_e$ be the set of exploratory behaviors available to the robot. Each behavior, has one or more parameters that modify its outcome. Let the parameters for behavior β_i be given as a real-valued parameter vector $P_i = [p_1^i, p_2^i, \ldots p_{q(i)}^i]$, where $q(i)$ is the number of parameters for this behavior. The behaviors, and their parameters, could be learned by imitation, programmed manually, or learned autonomously by the robot. For the purposes of this paper, however, the issue of how these behaviors are selected and/or learned will be ignored.

A set of invariant functions $\mathcal{I}_1, \mathcal{I}_2, \ldots, \mathcal{I}_r$ is also defined. These functions take as parameters the observation sequences generated by individual behaviors and try to detect regularities (or invariants) in the sensory data. The invariant functions relate the movements of the object's locations to the movements of the robot's body locations evaluated in a specific body frame.

For example, if behavior β_i starts at time t and ends at time $t + \tau$ it will generate the following sequence of observations: $\mathcal{O}(t), \mathcal{O}(t+1), \ldots, \mathcal{O}(t+\tau)$. The value of invariant function $\mathcal{I}_i$ calculated in frame F_j will be denoted with $\mathcal{I}_{i,j} = \mathcal{I}_i(\mathcal{C}_j(\mathcal{O}(t)), \mathcal{C}_j(\mathcal{O}(t+1)), \ldots, \mathcal{C}_j(\mathcal{O}(t+\tau)))$. In other words, all elements of the observation vector are converted to coordinates in frame F_j before the invariant function is calculated. The invariant functions return a value in the interval $[0, \infty)$. A value of 0 means that a perfect synchrony was observed between the movements of the robot's body locations and the movements of the object's locations. For each invariant function $\mathcal{I}_i$ a binary valued function Γ_i is also defined which outputs 1 if the value of the invariant is below a given threshold ϵ_i and 0 otherwise.

With this notation in mind, the task of learning the binding affordances of an object can be reduced to the task of populating the values of the following *Affordance Table* (see below). Where $s_1, \ldots, s_u$ are behavioral test sequences composed of exploratory behaviors as described below.

Behavioral Sequence	**Frame 1 Invariants**			. . .	**Frame k Invariants**		
s_1	$\mathcal{I}_{1,1}$	. . .	$\mathcal{I}_{r,1}$	. . .	$\mathcal{I}_{1,k}$	. . .	$\mathcal{I}_{r,k}$
s_2	$\mathcal{I}_{1,1}$	. . .	$\mathcal{I}_{r,1}$	. . .	$\mathcal{I}_{1,k}$	. . .	$\mathcal{I}_{r,k}$
. . .	. . .	. . .	. . .	. . .	. . .	. . .	. . .
s_u	$\mathcal{I}_{1,1}$	. . .	$\mathcal{I}_{r,1}$	. . .	$\mathcal{I}_{1,k}$	. . .	$\mathcal{I}_{r,k}$

The robot keeps a history of the executed behaviors and their parameters. Whenever it detects an invariant it queries the history and generates test sequences of increasing length. These sequences are then applied to the object to test which ones can reproduce the invariant. More specifically, let h be an entry in the history (containing both a behavior and its parameters). Furthermore, let the current history be $H = h_0, \ldots, h_{i-3}, h_{i-2}, h_{i-1}, h_i$. Where h_0 is a behavior that puts the robot in its start position (i.e., it resets the robot to a known configuration). If an invariant is detected after h_i is complete then the robot's controller generates the following test sequences: (h_0, h_i), (h_0, h_{i-1}, h_i), $(h_0, h_{i-2}, h_{i-1}, h_i)$, $(h_0, h_{i-3}, h_{i-2}, h_{i-1}, h_i)$, etc.

The shortest test sequence that reproduces the invariant is included in the affordance table. The invariant is considered reproduced if the last behavior in the test sequence generates at least one high interest reading (i.e., $\Gamma_i = 1$) for any of the invariant functions in any of the body frames. Also, the last behavior should have changed the position of the object as observed in camera coordinates.

Experimental Environment

Dynamics Simulator

All experiments were performed using a dynamics robot simulator developed in house. All objects in the simulator are modeled as rigid bodies. The bodies describing the robot are connected with joints which impose limits on their movement. The simulator calculates the friction and collision forces between the simulated objects and updates their positions and velocities at regular intervals (every 0.005 seconds). The dynamics are calculated using the Open Dynamics Engine library v 0.035 (Smith 2003). All experiments were run on a Pentium 4 machine (2.4 GHz, 512 MB RAM), running RedHat Linux 9.0.

Robot

The robot is a simulated model of a CRS+ A251 manipulator arm (see Figure 1). The robot has 5 degrees of freedom (waist translate, waist roll, shoulder pitch, elbow pitch, wrist pitch, wrist roll) plus a gripper. A simulated camera is mounted above the robot's working area. Vision routines are simulated by reading the positions of the objects from the internal data structures of the simulator.

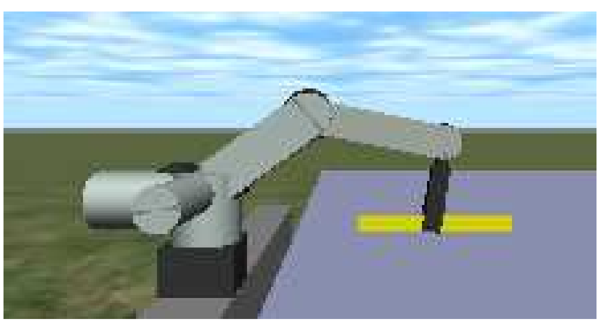

Figure 1: The simulated CRS+ A251 manipulator used in the experiments holding a stick object.

The robot has 23 markers located on its body that serve as the body locations $\mathcal{L}$. Figure 2 shows a close up of the arm and wrist and the positions of some of the markers. The robot also has 5 simulated tactile sensors located on its gripper. There are two touch sensors per finger (inner and outer surface). The fifth sensor is located on the surface of the wrist between the two fingers (See Figure 2).

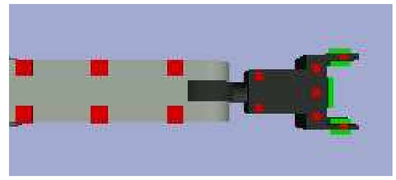

Figure 2: A closeup of the robot's arm and wrist. Red squares show the locations of some of the body markers. Green rectangles show the locations of the tactile sensors.

Objects

Seven different objects were used in the experiments: stick, spindle, mallet, dumbbell, beam, H-frame, and π-frame (see Figure 3). The objects have different shapes which makes them easier or harder to grasp. For example, the dumbbell can be grasped only in its middle section while the beam object is too thick to fit in the robot's gripper. All objects are color coded so they can be uniquely identified by the robot. Their colors are included in the observation vector (see below).

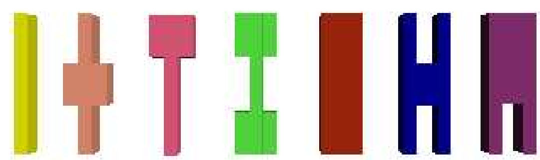

Figure 3: The seven objects used in the experiments. From left to right: stick, spindle, mallet, dumbbell, beam, H-frame, and π-frame.

Exploratory Behaviors

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.

Id	Exploratory Behavior	Parameters
β_0	*rotate-arm*	relative angle
β_1	*rotate-wrist*	relative angle
β_2	*pitch-wrist*	relative angle
β_3	*touch-object*	relative offset
β_4	*lift-arm*	offset from table
β_5	*lower-arm*	offset from table
β_6	*open-gripper*	
β_7	*close-gripper*	

The first three behaviors perform rotational movements with the arm and wrist respectively: *rotate-arm* moves the arm left or right; *shake-wrist* moves only the wrist left or right; and *pitch-wrist* moves only the wrist up or down. Each of these behaviors has one parameters which determines the direction of rotation and the magnitude of the movement in that direction. Two different values were used for the angle parameter : $+20$ and -20 degrees.

The *touch-object* behavior moves the wrist of the robot until it comes into contact with a specific location on the object. The position of this location is specified with an offset which gives the coordinates of a point along the major axis of the object. Three offset values were used (-10cm, 0cm, and +10cm) which correspond to the lower, middle, and upper part of the object's main axis.

The *lift-arm* and *lower-arm* behaviors move the arm of the robot vertically up and down respectively. They have only one parameter that specifies the end offset of the wrist relative to the table. The parameter for *lift-arm* has only one value of +20cm; *lower-arm* also has one value of 0cm.

The last two behaviors simply open and close the gripper. They have no parameters, but they check the values of the tactile sensors to determine when to stop the movements of the gripper if it comes into a contact with an object. Grasping is achieved through the friction force calculated by the simulator between the closed gripper and the object. Therefore, objects can fall off if they are not grasped properly. The last two objects shown in Figure 3 can be attached to the robot by opening the griper while it is inserted between the two vertical bars (i.e., this is a different form of binding).

Observation Vector

The observation vector has 55 real-value components as shown in the table below. The first 23 components hold the positions of the robot's body markers in camera coordinates. The next five elements represent the activation values of the five touch sensors. Elements 29,30, and 31 hold the RBG values of the object's color which are used to uniquely identify each object. The last 24 components give the positions of the object's corners in camera coordinates. The corners are used as object locations because they are easy to identify. Since different objects have different number of corners (between 8 and 24 for the objects shown in Figure 3) the last entries of the observation vector may be marked as unused.

Observation	Meaning
$o_1 - o_{23}$	X,Y,Z positions of body locations
$o_{24} - o_{28}$	grasp sensor activation values
$o_{29} - o_{31}$	R,G,B color components of the object
$o_{32} - o_{55}$	X,Y,Z positions of object locations

Body Frames and Invariant Function

Two body frames were defined for the robot. The first frame is attached to its arm. The second frame is attached to its wrist. Each frame is uniquely specified by three markers that always remain at equal distances relative to each other. The positions of the markers can be used to form three orthonormal vectors which define a frame. A Gram-Schmidt process was used to calculate the basis vectors.

Only one invariant function was used. It returns the sum of the standard deviations of the positions of all object corners observed during the last executed behavior. This sum is normalized by the number of corners so objects with large number of corners don't get artificially high scores, i.e.,

$$\mathcal{I}_1 = \frac{1}{c} \sum_{i=32}^{32+c-1} stdev(o_i(t), o_i(t+1), \ldots, o_i(t+\tau))$$

where c is the number of object corners. An empirically estimated threshold value ϵ_i=2_{cm} was used in all Γ_i functions.

Experimental Evaluation

For each of the seven objects shown in Figure 3 a separate learning trial was conducted. During the learning trials the robot was allowed to freely explore the objects. The exploration consists of trying different behaviors, observing their results, and filling up the affordance table using the history method described above. The learning trials were limited to thirty minutes of simulation time per object. The initial placement of the objects was random, but they were always placed within the robot's sphere of reach. If an object was pushed out of reach during its exploration then the simulation was restarted with a new random position.

The data from each learning trial was sufficient to populate the entries of the affordance table for each object. The table below shows the data gathered for the dumbbell object in a typical run. The behavior ids are the same as in the previous section. The first behavior in each sequence (not shown) is a reset behavior which puts the robot in its start position and also drops any objects that the robot might be holding. The maximum sequence length was set to four.

Learned Sequences (Dumbbell)				$\mathcal{I}_{1,1}$	$\mathcal{I}_{1,2}$
$\beta_3(0_{cm})$	$\beta_7()$	$\beta_0(+20°)$		0.1	0.1
$\beta_3(0_{cm})$	$\beta_6()$	$\beta_7()$	$\beta_4(+20_{cm})$	1.8	0.2
$\beta_3(0_{cm})$	$\beta_2(-20°)$	$\beta_7()$	$\beta_0(+20°)$	0.1	0.1
$\beta_2(-20°)$	$\beta_3(0_{cm})$	$\beta_5(0_{cm})$	$\beta_0(-20°)$	0.7	0.6

As the table shows, the *touch-object* and *close-gripper* behaviors (β_3 and β_7 respectively) appear very often in the learned sequences. Based on this frequency information these behaviors can be automatically selected as a prerequisite for binding with the object (however, this was not performed in this set of experiments). The sequences learned for the H-frame and the π-frame used the *open-gripper* behavior to achieve the binding.

Learned Sequences (π-frame)				$\mathcal{I}_{1,1}$	$\mathcal{I}_{1,2}$
$\beta_3(-10_{cm})$	$\beta_6()$	$\beta_0(-20°)$		0.4	0.4
$\beta_3(-10_{cm})$	$\beta_2(+20°)$	$\beta_6()$	$\beta_0(+20°)$	0.9	0.9
$\beta_3(-10_{cm})$	$\beta_1(-20°)$	$\beta_6()$	$\beta_0(-20°)$	1.3	0.9

Similar results were obtained for the other objects. It was somewhat surprising to find that the affordance table for the beam object was not empty. The beam is the only object that does not fit inside the robot's gripper. Nevertheless, the robot learned to slide it sideways by sticking its gripper into the object and rotating its arm.

Learned Sequences (Beam)			$\mathcal{I}_{1,1}$	$\mathcal{I}_{1,2}$
$\beta_6()$	$\beta_3(0_{cm})$	$\beta_0(-20^\circ)$	1.7	1.7
$\beta_7()$	$\beta_3(0_{cm})$	$\beta_0(+20^\circ)$	1.5	1.5

The tables also show that the observed object variations in body frame 2 (the one attached to the wrist) are usually equal or smaller than the variations in frame 1 (the one attached to the arm). This information can be exploited to compress the affordance table by considering only the wrist frame in any future object explorations.

It should also be noted that all seven objects are relatively thin and can be grasped sideways by the robot if they are first rotated by 90 degrees. However, the robot lacks the required exploratory behavior that can perform such a rotation. Therefore, it could not discover these other ways for binding with the objects. Adding the capability to learn new exploratory behaviors can resolve this problem.

MPEG movies from the experiments are available at: (http://www.cc.gatech.edu/˜saho/tooluse/DevRob2005).

Conclusions and Future Work

This paper introduced a novel approach for developmental learning of the binding affordances of objects by a robot. A behavior-based approach was used to *ground* the affordance representation in the behavioral repertoire of the robot. More specifically, the affordances of different objects were represented in terms of a set of exploratory behaviors and their observed effects. Simulation experiments were conducted in a dynamics robot simulator. Our initial experiments suggest that the behavior-grounded approach can be used by a robot to autonomously learn the binding affordances of different objects.

The affordances are learned during a behavioral babbling stage in which the robot randomly chooses sequences of exploratory behaviors, applies them to the objects, and detects regularities (or invariants) in the resulting set of observations. Only sequences that reliably reproduce an invariant are used to represent the affordances. The invariants are calculated relative to the robot's body.

A shortcoming of the behavior-grounded approach is that there are object affordances that are unlikely to be discovered since the required exploratory behavior is not available to the robot. This problem has also been observed in animals, e.g., macaque monkeys have significant difficulties learning to push an object away from their bodies because this movement is never performed in their normal daily routines (Ishibashi, Hihara, & Iriki 2000). This problem can be resolved, however, if the ability to learn new exploratory behaviors is added.

Future research can extend this approach in several ways. First, the number of invariant functions should be increased so that the robot can detect additional properties of the objects. Second, interpolation between behavioral parameters can be used to quickly come up with new variations of the exploratory behaviors (e.g., finding intermediate grasp points between the three ones used here). Third, generalization techniques can be used to learn the affordances of a new object by relating them to the affordances of familiar objects. And finally, the simulation results should be replicated on a real robot.

References

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

Cutkosky, M. R. 1989. On grasp choice, grasp models, and the design of hands for manufacturing tasks. *IEEE Transactions on Robotics and Automation* 5(3):269–279.

Drescher, G. L. 1991. *Made-up minds: a constructivist approach to artificial intelligence*. MIT Press.

Edelman, G. M. 1987. *Neural Darwinism: The Theory of Neuronal Group Selection*. Basic Books.

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

Ishibashi, H.; Hihara, S.; and Iriki, A. 2000. Acquisition and development of monkey tool-use: behavioral and kinematic analyses. *Canadian Journal of Physiology and Pharmacology* 78:958–966.

Krotkov, E. 1995. Perception of material properties by robotic probing: Preliminary investigations. In *IJCAI*, 88–94.

Lorenz, K. 1996. Innate bases of learning. In Pribram, K. H., and King, J., eds., *Learning as Self-Organization*. Mahwah, NJ: Lawrence Erlbaum Pub.

Piaget, J. 1952. *The origins of intelligence in children*. Intl. Univ. Press.

Pollard, N. S. 1996. Synthesizing grasps from generalized prototypes. In *IEEE ICRA*.

Power, T. G. 2000. *Play and Exploration in Children and Animals*. Mahwah, NJ: Lawrence Erlbaum Pub.

Smith, R. 2003. *The Open Dynamics Engine (ODE) User Guide*. http://ode.org/.

Stansfield, S. A. 1991. Robotic grasping of unknown objects: A knowledge-based approach. *International Journal of Robotics Research* 10(4):314– 326.

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

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