Toward Autonomous Learning of an Ontology of Tool Affordances by a Robot

Jivko Sinapov
Department of Computer Science
Iowa State University
jsinapov@iastate.edu

Alexander Stoytchev
Department of Electrical and Computer Engineering
Iowa State University
alexs@iastate.edu

Introduction

The problem of autonomous learning of affordances typically requires a robot to learn the types of changes it can induce and detect in its environment (Sahin *et al.* 2007; Sinapov & Stoytchev 2007). In sufficiently complex environments, however, it is impossible to know in advance the exact nature and number of possible environmental outcomes that the robot can induce through its behaviors. In addition, the changes that the robot can detect are often high-dimensional, making it difficult to use standard machine learning algorithms. This work addresses this problem by proposing a framework in which the robot learns a *taxonomy* for the types of perceivable changes produced by its own behaviors. The proposed method also allows the robot to incrementally update the taxonomy and to conceptualize new types of observed outcomes. In addition, the robot solves a *hierarchical classification* task by learning a model that predicts the future outcome of its behaviors in relation to the learned taxonomy. Thus, the robot builds an affordance ontology consisting of an outcome class taxonomy and a predictive model grounded in the robot's perceptual and behavioral repertoire.

Experimental Setup

The theoretical framework described below was tested on a tool manipulation task. All experiments were performed using the open-source dynamic robot simulator BREVE. The robot is a simulated arm with a gripper attached to the wrist, as shown in Fig. 1. Two tools were used: a T-Stick tool and an L-Stick tool. The last object in the simulation is a small cylindrical puck which can be moved by the tool when the robot performs an action. Videos of the simulation are available at http://www.cs.iastate.edu/˜jsinapov/AAAI08/.

The robot's set of behaviors, $\mathcal{B}$, consists of 6 possible behaviors with the tool: *push, pull, slide left, slide right, rotate left*, and *rotate right*. The robot's cue, C_i, is a 30 by 30 retinal image centered on the puck, as shown in Fig. 2 c). The observed outcome O_i is 12-dimensional and describes the puck's vertical and horizontal displacement (relative to the puck's starting position) in the robot's field of view over the course of 60 simulator time steps, and sampled every 10 frames, as shown in Fig. 2 d).

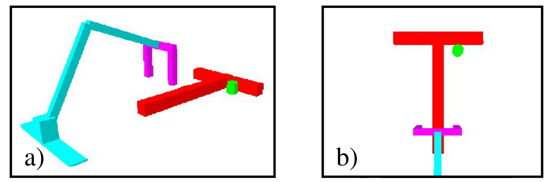

Figure 1: a) Snapshot of the robot arm in the dynamics simulator; b) View from the robot's simulated camera.

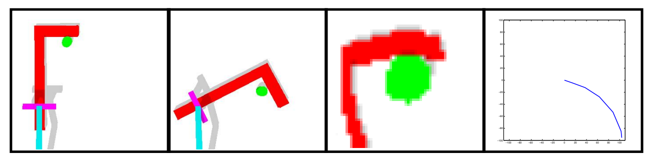

a) b) c) d)

Figure 2: A trial with the L-Stick tool and *rotate right* behavior; a) view at beginning of trial; b) view at end of trial; c) retinal mapping image used as cue at beginning of trial; d) observed outcome plotted as a trajectory at end of trial.

Theoretical Model

Let $X_i = (B_i, C_i)$ be an input data point indicating that the robot is executing behavior $B_i \in \mathcal{B}$ while detecting a cue $C_i \in \mathbb{R}^n$. Let $O_i \in \mathbb{R}^m$ describe the detected outcome (e.g., visual movement) after the behavior B_i has been executed.

Because O_i can be high-dimensional, it is difficult to learn a model to predict O_i given X_i using standard machine learning algorithms designed for output classes with a single label. To address this problem, the robot learns a hierarchical taxonomy of outcomes, $\mathcal{T}$, which is a tree defined over outcome classes (i.e., nodes) $v_0, \ldots, v_M$. Let $O_j^{mean} \in \mathbb{R}^m$ denote the outcome *prototype* for the observed outcomes that belong to node v_j. Outcome classes can vary from specific (near the bottom of the tree) to more general (near the top of the tree). In the experiments conducted for this study, the taxonomy is learned using an incremental hierarchical clustering framework in which an outcome class v_j in $\mathcal{T}$ is split into sub-classes once the number of observed outcomes that fall into it reaches a threshold. The split operation adds child nodes to v_j resulting in more refined outcome classes. The split is performed using the X-Means clustering algorithm.

The robot incrementally updates the taxonomy $\mathcal{T}$ while learning a model $\mathcal{M}(X_i) \rightarrow \hat{P}_i$ that predicts a path, $\hat{P}_i = [v_0, \ldots, v_l]$, from the root node v_0 to some leaf node v_l in $\mathcal{T}$ which describes how the predicted outcome relates to the learned taxonomy. To solve the hierarchical classification problem required for this, each non-leaf node j has an associated model M_j which is trained to predict the child outcome class of the (yet unobserved) outcome O_i associated with X_i. Formally, $M_j(X_i) \rightarrow \hat{c_k}$ where $\hat{c_k} \in children(v_j)$. For example, the root node in the taxonomy shown in Fig. 3 contains a model M_0 which given an input X_i predicts which child outcome class (v_{01}, v_{02}, v_{03} or v_{04}) the future outcome O_i belongs to. Thus, applying a recursive top-down prediction routine results in a predicted path, $\hat{P}_i$, from the root node to a leaf node in the tree. Each model M_j is realized by an incremental ensemble classifier framework.

In summary, the robot learns an affordance ontology $\mathcal{A} = \{\mathcal{T}, \mathcal{M}\}$, where $\mathcal{T}$ is the hierarchical taxonomy of outcomes induced by the robot's behaviors $\mathcal{B}$, and $\mathcal{M}$ is the set of predictive models contained in the non-leaf nodes of $\mathcal{T}$.

Experiments and Results

Experiment 1: In the first experiment, the robot conducts 1200 trials with the L-Stick. During each trial, the puck is randomly placed near the tool and a random behavior is selected for execution. After the 200^{th} trial, the first split of the root node in $\mathcal{T}$ occurs, resulting in four child outcome classes. After all 1200 trials, the robot has formed 13 leaf outcome classes. A partial visualization of the learned taxonomy is shown in Fig. 3. The root outcome class, v_0, contains all observed trajectories of the puck as a result of performing a behavior with the L-Stick. When evaluated on a

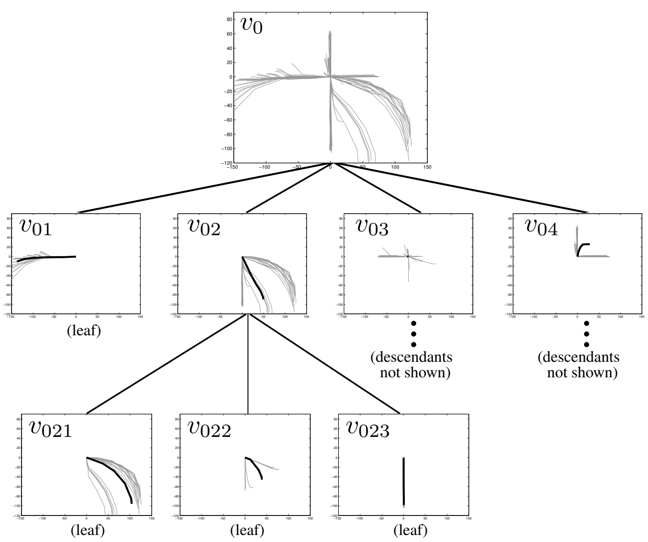

Figure 3: A partial visualization of the learned outcome taxonomy, $\mathcal{T}$, after 1200 trials with the L-Stick tool. For each outcome class v_j the darker trajectory denotes the outcome prototype O_j^{mean}, while the lighter trajectories visualize the observed outcomes that fall within v_j. The full taxonomy is shown in the accompanying web page.

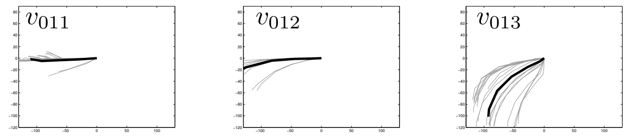

Figure 4: New leaf outcome classes formed after interaction with the T-Stick and inserted into the taxonomy built through interaction with the L-Stick (see Fig. 3)

set of 600 test trials, all predictive models in $\mathcal{M}$ achieve an Area Under ROC Curve (AUC) (averaged over class labels) between 0.82 and 1.0 with the majority higher than 0.9, indicating that the robot is capable of predicting the outcome classes given its cues and behavior.

Experiment 2: The second experiment tests how the robot can adapt an already learned affordance ontology when a new tool is introduced. To do that, 1200 trials were performed with the L-Stick (resulting in the outcome taxonomy shown in Fig. 3), followed by 1200 trials with the T-Stick tool. Unlike the L-Stick, the T-Stick allows the robot to bring the puck closer by also rotating the tool to the left.

After the 920^{th} trial with the T-Stick tool, one of the leaf outcome classes, v_{01} (see Fig. 3) is split and three child nodes are added (shown in Fig. 4). The new leaf node v_{013} is the outcome class that is formed to describe the novel outcomes observed only with the T-Stick. The newly formed model M_{01} achieves an AUC of 0.98 when evaluated on novel trials with both tools, indicating that the robot has learned to accurately distinguish between the novel outcome classes $v_{011}, v_{012}, v_{013}$ based solely on its sensory input, without explicitly knowing which of the two tools it is currently using. Additional details and full visualization of the learned outcome taxonomies are available at http://www.cs.iastate.edu/~jsinapov/AAAI08/.

Conclusion and Future Work

The experiments show that by learning an affordance ontology with an adaptive taxonomy, the robot can form hierarchically structured outcome classes that accurately describe the changes it can induce and detect in its environment. The robot also learns predictive models that allow it to anticipate the outcomes of its actions in advance. In future work, the framework can be improved through the use of more advanced taxonomy learning algorithms as well as dealing with outcomes that describe how the robot's action impacts more than a single object. The framework will also be implemented and evaluated on a 7-DOF Barrett WAM robot arm.

References

Sahin, E.; Cakmak, M.; Dogar, M.; Ugur, E.; and Ucoluk, G. 2007. To afford or not to afford: A new formalization of affordances toward affordance-based robot control. *Adaptive Behavior* 15(4):447–472.

Sinapov, J., and Stoytchev, A. 2007. Learning and generalization of behavior-grounded tool affordances. In *Proc. 7th IEEE Int. Conf. on Development and Learning (ICDL).*