Artificial Intelligence for Human-Robot Interaction
Papers from the AAAI 2015 Fall Symposium

Exploring Affordances Using Human-Guidance and Self-Exploration*

Vivian Chu and **Andrea L. Thomaz**
School of Interactive Computing
Georgia Institute of Technology
Atlanta, Georgia 30332-0250
Email: vchu@gatech.edu, athomaz@cc.gatech.edu

Abstract

Our work is aimed at service robots deployed in human environments that will need many specialized object manipulation skill. We believe robots should leverage end-users to quickly and efficiently learn the affordances of objects in their environment. Prior work has shown that this approach is promising because people naturally focus on showing salient rare aspects of the objects (Thomaz and Cakmak 2009). We replicate these prior results and build on them to create a semi-supervised combination of self and guided learning. We compare three conditions: (1) learning through self-exploration, (2) learning from demonstrations provided by 10 naïve users, and (3) self-exploration seeded with the user demonstrations. Initial results suggests benefits of a mixed initiative approach.

Introduction

Robots have begun the transition from factories to homes and to deal with the uncertainties that the real world holds, robots need to learn about the environment it is placed in. Luckily, robots can draw on input and feedback from humans in these new social settings. This work explores how to utilize human teachers to guide a robot's exploration when learning about the environment and how this guidance allows the robot to learn more efficiently about the world.

To understand how to best make use of humans to help teach robots about the environment, we look at representing the world as affordances and consequently, we explore the problem of affordance learning for robots. The term affordance, was first introduced by J.J. Gibson in 1977 (Gibson 1977) and we use the ecological definition from (Chemero 2003) of "action possibilities" that appear between an agent and the environment. More concretely, we represent affordances as the relationship between effects and a set of actions performed by an agent on an object. This representation is commonly used in robotics (Şahin et al. 2007; Montesano et al. 2008).

We present three strategies to tackle the affordance learning problem (1) the traditional self-exploration strategy where robots exhaustively interact with the workspace (2) human-guided exploration based on prior work from (Thomaz and Cakmak 2009) where humans provide examples interactions for the robot to learn from and (3) a mixed approach that combines self-exploration with information provided from human teachers. We compare these three strategies by learning five affordances across four different objects.

*This research is supported by NSF CAREER grant IIS-1032254

Related Work

The research area for understanding how robots can explore the world has been looked by a wide variety of researchers. One specific area that relates directly to this work is the concept of intrinsic motivation and curiosity driven exploration. Some early work in this area (Oudeyer, Kaplan, and Hafner 2007; Vigorito and Barto 2010; Schmidhuber 1991) looked at using rewards and expectations to guide the exploration as opposed to any human-guidance.

More recently, (Ivaldi et al. 2014; 2012) and (Nguyen and Oudeyer 2014) explored the idea of combining intrinsic exploration with human-guidance. The difference in this work lies in the combination of social and self-exploration. Both assume that there exists a reward signal that can be easily characterized. However, such reward does not exist for complex affordances such as the ones explored in this paper.

Finally, this work aims to build on the previous findings from (Thomaz and Cakmak 2009). However, this work not only applies human-guided affordance to more complex affordances, but looks at the combination of human-guidance and self-exploration, a task the prior work did not do.

Affordance Learning

To learn affordances, we need an agent to interact with the environment and observe the effects of the interaction. From these interactions and observations, the agent can then learn about what the environment affords for it. In the simplest case, if the agent's actions are discrete, then it could just try all the actions on all of the objects and model the outcomes. However, with any real object manipulation skill, the space of actions that the robot could try to perform on the object becomes so large that we require a method for efficiently sampling this space in order to build an accurate model.

In our approach, we assume the robot has a set of parameterized primitive actions (e.g. position of the end-effector

(a) Bread box (b) Drawer (c) Pasta Jar (d) Lamp

Figure 1: Objects that Curi explored

Table 1: Affordances

Object	Action	Effect	Affordance
Bread box	Move	Moves up	open-able
Pasta jar	Move	Moves	push-able
Drawer	Move	Moves	push-able
Drawer	Pick	Pulls out	open-able
Lamp	Pick	Pulls down	turn-on-able

with respect to the object in a pick-up action). Then the exploration space is defined by this space of these continuous-valued parameters on the primitive actions. We compare three different approaches for efficiently sampling this space to collect examples to build a model of particular action-object affordances.

Objects and Actions

For this work, we selected four household objects for the robot to interact with and can be seen in Fig. 1. Each of these objects are tracked using the RGB-D sensor and visual object information commonly used in affordance learning (Thomaz and Cakmak 2009; Montesano et al. 2008) were recorded throughout the entire interaction, but converted to 3D space rather than 2D images. Specifically, we record the object point cloud centroid (relative to the robot torso), color, orientation, volume of the bounding box, the dimensions of the bounding box (x,y,z), and the squareness of the object (the ratio of the number of points in the object to the area of the bounding box). We also store information from the 6-axis F/T sensor in the wrist (F_x, F_y, F_z, T_x, T_y, T_z) and the robot end-effector (EEF) relative to the robot torso.

The robot can perform two parameterized action primitives: **move** and **pick**. Each action is composed of a sequence of robot end-effector (EEF) poses relative to the centroid of the object point cloud. An EEF pose describes the position and orientation of the robot hand for all 6 degrees-of-freedom (DOF). A **move** action is comprised of two poses, the start and end pose of the EEF. The **pick** action has three poses: a start pose, a pose where the robot should close its hand, and an end pose. For both primitives, the resulting trajectory of the EEF is created by performing a quintic spline between the poses with an average velocity of 1 cm/second.

Affordances

This paper explores five specific affordances summarized in Table 1. The affordances selected, range from simpler affordances (push-able) to affordances that are more complex (turn-on-able). We define complex in terms of the action required to find the affordance. For "simple" affordances, the affordance can be found in a larger part of the space during exploration whereas a "complex" affordance requires movement or sampling along a specific portion of the space. For example, there are many directions to move relative to the object to find the "simple" affordance of push-able. However, to discover that the drawer is "open-able", the robot must pull in a specific direction to find the affordance.

Affordance Representation

To model the affordances, we selected Hidden Markov Models (HMMs) (Rabiner and Juang 1986). We chose this representation because of the ability to capture time varying information of the various affordances. For example, pulling the drawer requires specific forces to be felt throughout the interaction. Furthermore, HMMs allow us to develop a model of affordances that we can later sample from to generate new interactions with novel objects. The HMMs state-space contains the visual and haptic features described in the prior sections.

Exploration Strategies

The main question we ask is how to best navigate this search space and interact with objects to learn affordances. One of the major challenges of learning the affordances of an object rests largely on how to *efficiently* explore an object to produce the interesting effects. We present three strategies: self-exploration, human-guided exploration, and human-seeded self-exploration.

Self-Exploration

The self-exploration strategy exhaustively searches the workspace with little knowledge about where to search aside from knowing that it should perform actions around the object. This is the typical strategy taken for learning affordances (Fitzpatrick et al. 2003; Stoytchev 2005; Hermans et al. 2013) and the main decisions needed to discretize the workspace relate to (1) what range the robot should explore around the object and (2) the resolution (step-size) for each interaction.

To perform a simple exploration in all six dimensions of the EEF, quickly results in an exploration size that is practically infeasible. To reduce the search space, we provided starting points for self-exploration provided by an expert (one of the authors) as well as selecting the best orientation that would find the affordance. This is a reasonable assumption because many state-of-the-art algorithms focus at determining the best grasp points or points in which a robot EEF should interact with on the object (e.g. the handle on the bread box or the ball on the chain for the lamp). With these two assumptions, we now exhaustively explore only in the end position (x,y,z) of the action. However, exploration in these three dimensions can easily explode if the resolution of the search space is small enough and so we also limit the number of total explorations to 100 interactions per object.

Note, that to even to make the self-exploration tractable, we have to provide some human knowledge about the object and scene. Specifically, we provided the start position and

Table 2: Total Number of Interactions Per Strategy

Object	Action	Self	Humans	Guided
Bread box	Move	100	57	31
Pasta jar	Move	100	49	30
Drawer	Move	100	39	31
Drawer	Pick	100	45	31
Lamp	Pick	100	49	31

orientation of the EEF as well as the maximum distance that the EEF has to explore to find the affordance.

Human-Guided Self-Exploration

The next approach looks at how people can guide the exploration of objects to sample this space of affordances. Specifically, we look at the types of interactions naive users provide as examples of affordances in objects.

In order for naive users to control the robot, users were given the same action primitives (**move** and **pick**) that the robot had access to during self-exploration. Users generated these primitives by kinesthetically showing the robot the poses of each action (i.e. start, end, close hand pose). We conducted a user study with 10 naive users (5 male, 5 female). The participants were instructed to teach the robot about the 5 affordances over the 4 objects described in Table 1. For each object, they were told the specific action (**move** or **pick**) to use and the effect to show the robot. At the beginning of each session, the participants were instructed briefly on the definition of affordances. They also were given a brief training time with the robot on how to verbally command and move the robot for kinesthetic teaching. For practice, they taught two actions on two separate objects. These affordances were not included in the experiment.

For the "complex" affordances (i.e. bread box, open-able drawer, and lamp), users were given 10 minutes to explore the object and for the "simple" affordances (pasta jar and push-able drawer), users were given 5 minutes. At the end of the experiment, participants answered a survey question about their teaching strategy. The total number of interactions aggregated across all 10 users can be seen in Table 2.

Human-Seeded Self-Exploration

Our final strategy looks at a combination of the self and guided approaches. While users provide directed exploration strategies, it is cumbersome to have people provide an exhaustive set of interactions for each affordance. During self-exploration, the robot can easily generate an exhaustive area to search, but has no concept of where the exploration should be focused. Combining the strengths of both approaches should yield the best of both worlds.

For Human-Seeded Self-Exploration (HSSE), we use human demonstrations to constrain the search space during self-exploration. On a high-level, we used the variance of all user demonstrations to generate a set of explorations from the original demos. We do this by taking the first demonstration from each user for each affordance and compute the mean and standard deviation of this set. Specifically, for

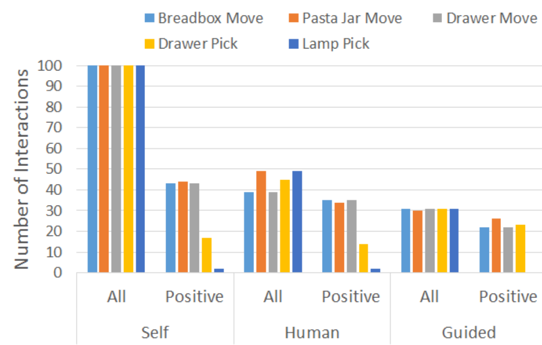

Figure 2: Number of total and positive examples for each affordance

each affordance (e.g., breadboxMove, drawerPick), there exists a mean ($\mu_{breadboxMove}$ and variance $\sigma^2_{breadboxMove}$) of the EEF relative to the object.

During self-exploration, we searched in a circular range around the start/close position of each affordance. During HSSE, we instead search around the $\mu_{affordance}$, with $\sigma^2_{affordance}$ as the step size. For each affordance, we have 27 search locations and use the same EEF orientation that was used during self-exploration. This search allows the robot to explore dimensions (x,y,z) that do not matter (higher variance) and focus on the axes that are highly constrained (low variance).

Furthermore, to focus the exploration on the direction of change during each demonstration, we do a second search of the area where the EEF explores this dimension. To keep things fair, the step-size for this search is selected based on the resolution during self-exploration.

Initial Results and Discussion

We used all three strategies to collect interactions with all 5 affordances. The number of interactions per object per approach can be seen in Table 2. Each interaction was hand labeled with the ground truth label of "Success" and "Failure" depending on if the interaction successfully found the affordance. The number of successful interactions vs. failed interactions can be seen in Figure 2.

The first thing that we noticed while collecting the data was the shear magnitude of explorations required to fully explore an area when no initial information is given. This is clearly seen in Table 2 where self-exploration requires nearly double the amount of explorations as the aggregate exploration for all 10 users. This suggests that human guidance can provide crucial information for robot exploration.

We then looked at the quality of explorations for each search method. We used a rough metric of the number of successes vs. failures. This can be seen in Figure 2. The graph shows the number of positive interactions for each affordance and each method. For self-exploration, less than half of the interactions resulted in the robot finding the affordance. This percentage drastically goes up during human-guidance and carries over to HSSE.

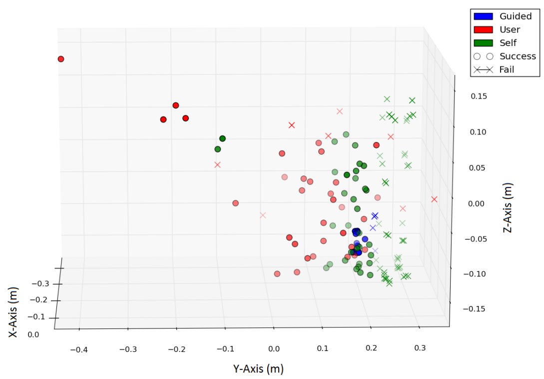

Figure 3: Exploration space of EEF relative to the drawer during the action move.

We also looked that physical space that the EEF explored. To visualize this space, we plotted the position of the EEF during successful and failed interactions relative to the object. One example can be seen in Figure 3. Interestingly, we can see that the successful interactions with the object have a high dependency on the y-axis. We can also see that self-exploration looks at a much larger area of the object space for both successful and failed interactions. The human-guided exploration has far more positive examples over a larger area than failed interactions. Finally, HSSE is highly concentrated and focused on the boundary between success and failure.

To verify that the increase of exploration space improves the affordance models, we trained two separate HMMs (one for the failed interactions and one for the successful interactions). We build two models because it allows us to use the relative likelihood to compare if an interaction successfully found the affordance without fine-tuning a threshold. This also emphasizes the importance of covering the entire task space of the affordance rather than focus on just the positive examples. To evaluate the models we performed cross validation on the collected interactions from each exploration method and compared them to each other. Initial results are positive and suggest that there is a direct relationship between the quality of model to the coverage of the task space.

In conclusion, the initial results show that there is much to be learned from human input to guide robot exploration for affordances. Future work will focus on obtaining precision and recall scores during cross validation to evaluate the quality of the learned models and their direct relation to the amount of task space coverage.

References

Chemero, A. 2003. An outline of a theory of affordances. *Ecological psychology* 15(2):181–195.

Fitzpatrick, P.; Metta, G.; Natale, L.; Rao, S.; and Sandini, G. 2003. Learning about objects through action - initial steps towards artificial cognition. In *International Conference on Robotics and Automation (ICRA)*, 3140–3145.

Gibson, J. 1977. The concept of affordances. *Perceiving, acting, and knowing* 67–82.

Hermans, T.; Li, F.; Rehg, J. M.; and Bobick, A. F. 2013. Learning stable pushing locations. In *Development and Learning and Epigenetic Robotics (ICDL) 2013*, 1–7.

Ivaldi, S.; Lyubova, N.; Gerardeaux-Viret, D.; Droniou, A.; Anzalone, S.; Chetouani, M.; Filliat, D.; and Sigaud, O. 2012. Perception and human interaction for developmental learning of objects and affordances. In *Humanoid Robots (Humanoids), 2012 12th IEEE-RAS International Conference on*, 248–254.

Ivaldi, S.; Nguyen, S. M.; Lyubova, N.; Droniou, A.; Padois, V.; Filliat, D.; Oudeyer, P.-Y.; and Sigaud, O. 2014. Object learning through active exploration. *Autonomous Mental Development, IEEE Transactions on* 6(1):56–72.

Montesano, L.; Lopes, M.; Bernardino, A.; and Santos-Victor, J. 2008. Learning object affordances: From sensory–motor coordination to imitation. *Transactions on Robotics* 24(1):15–26.

Nguyen, S., and Oudeyer, P.-Y. 2014. Socially guided intrinsic motivation for robot learning of motor skills. *Autonomous Robots* 36(3):273–294.

Oudeyer, P.-Y.; Kaplan, F.; and Hafner, V. 2007. Intrinsic motivation systems for autonomous mental development. *Evolutionary Computation, IEEE Transactions on* 11(2):265–286.

Rabiner, L., and Juang, B.-H. 1986. An introduction to hidden markov models. *ASSP Magazine, IEEE* 3(1):4–16.

Şahin, E.; Çakmak, M.; Doğar, M. R.; Uğur, E.; and Üçoluk, G. 2007. To afford or not to afford: A new formalization of affordances toward affordance-based robot control. *Adaptive Behavior* 15(4):447–472.

Schmidhuber, J. 1991. Curious model-building control systems. In *Neural Networks, 1991. 1991 IEEE International Joint Conference on*, 1458–1463 vol.2.

Stoytchev, A. 2005. Behavior-grounded representation of tool affordances. In *International Conference on Robotics and Automation (ICRA)*, 3060–3065.

Thomaz, A. L., and Cakmak, M. 2009. Learning about objects with human teachers. In *International Conference on Human Robot Interaction (HRI)*, 15–22.

Vigorito, C., and Barto, A. 2010. Intrinsically motivated hierarchical skill learning in structured environments. *Autonomous Mental Development, IEEE Transactions on* 2(2):132–143.