Towards Fully Automated Learning of Grasping Affordances

Christian Wressnegger, Ashish Jain, Misel Batmendijn,
Manuel Lopes, Luis Montesano, Alexandre Bernardino and Jose Santos-Victor

Instituto Superior Técnico
Computer and Robot Vision Laboratory
{chwress,ajain,mbatmend,macl,lmontesano,alex,jasv}@isr.ist.utl.pt

Abstract

In this paper we address how a robot can learn to grasp objects from experience with a focus on open-ended learning. We aim at a system that gets along without a person assisting in interacting with the objects during its whole lifespan. The implementation was done based on the humanoid robotic platform iCub. In particular we dealt with: i) policies for the selection of promising training data ii) automated labeling of training samples and, iii) on-line learning. The presented system will serve as the basis for the evaluation of our research on cognitive learning.

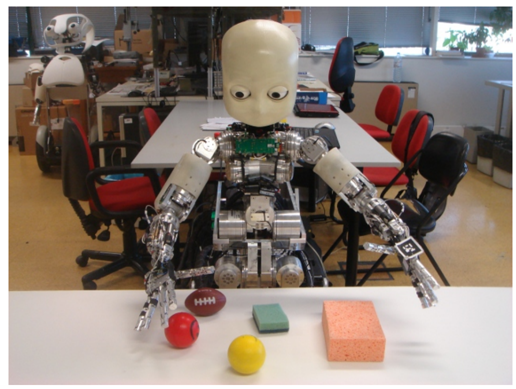

Figure 1. The setup of the experiment.

1. Introduction

To enable autonomous world interaction a robotic system has to reason about individual affordances [1]. A particular important object affordance is grasping. Learning how to grasp is driven by a large number of trials to enrich the robot's experience. Nowadays many approaches utilize simulators to automate such experiments in order to be able to automatically reset the environment over and over again. Doing this on real hardware usually requires significant assistance from the outside. In this paper we present a system to overcome this in order to establish learning in a open-ended way.

For this we build upon the work presented in [4]. It predicts the best suitable grasping points of an object based on a set of local visual descriptors of a single image. Figure 2 shows an example of such a probability map. This approach uses Binomial-Beta distributions to model the probability of successful and failed experiments. Based on this a non-parametric kernel approach is employed to reason about unseen objects based on previous observations. Hence, this already makes itself out to be largely independent of the action of a programmer (parameter tuning, etc.). The only point remaining is the supply of the labeled training data (the outcome of concrete experiments). Since the experiment's result is considered to be a boolean value (failure or success) this can easily be left to the robot.

In fact the used approach isn't limited to grasping but is suitable for affordances in general. The actual type and application is defined by the experiment that provides the success or failure state or in our case, by what the robot considers as success. The overall (long-term) objective is to carry on to a level where this is also learned as part of the robot's experience to enable unsupervised learning, as implied in [5].

2. System Overview

Due to the progress of the ROBOTCUB project[1] we increasingly use the iCub for our experiments on learning cognitive skills. The iCub is a humanoid robot designed to replicate the physical capabilities of a three and a half year old child. All in all it has 53 degrees of freedom (DOF) including 9 DOF for each hand [3].

[1] http://www.robotcub.org/

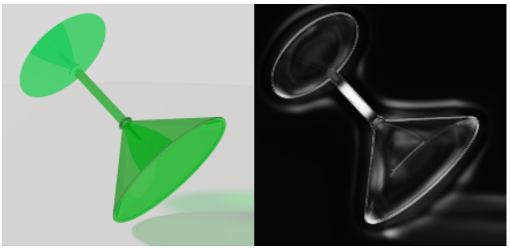

Figure 2. The predicted probability map for an upside down Martini glass. The brighter an area is in the right image, the higher is the probability for a successful grasp.

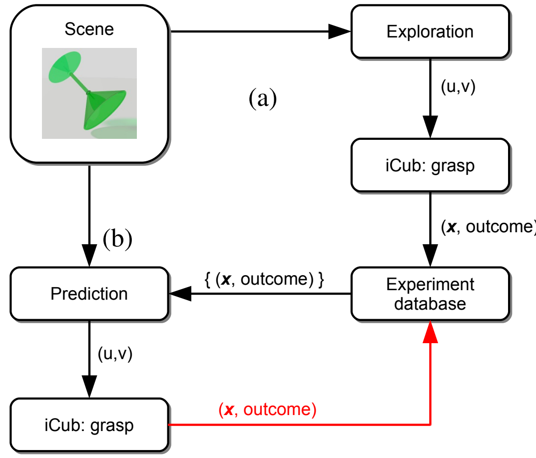

Figure 3. System bootstrapping (a) and application (b) schemes of the presented system for a single image. The red arrow indicates the extended ability of on-line learning.

Figure 1 shows the setup consisting out of the iCub and a simple table with objects on it. The robot is supposed to perceive an object on the table, grasp and lift it and put it back (drop it) afterwards. In this way we are able restart a grasping experiment on the same object but every time with a different orientation and viewing angle. The setup divides into two disjoint modes of operation (cf. Figure 3): a) the automated collection of training data for bootstrapping the system and b) the application.

2.1. System bootstrapping

In this mode (cf. Figure 3a) we do not factor in learning or prediction of grasping points as described in [4] at all. We rather make the robot explore the scene by grasping multiple times at various positions on the table. Hereto, we aim at strategically selecting points that are the most expressive ones. This mainly applies for points on the object or in its close surrounding.

To map a specific grasping point in the image (u, v) to the actual object on the table $(x, y, z)_{table}$ we use a 2D projective transformation (homography).The homography is fully defined by the kinematics of the robot and the relative position of the table to the robot's base coordinates. Once the 3D coordinates of the object are determined we are able to reach it and perform a grasping experiment.

For the classification of the current experiment we have to determine if the robot holds (successfully grasped) an object or not. The hands of the iCub lack tactile sensors which would allow us to directly reason about the result. To overcome this limitation we analyze the power consumption of the individual joints when manipulating objects. The increased consumption that a specific load implies gives us a good index of if the hand envelopes an object. However, on its own this doesn't enable us to infer a reliable metric that tells us how good or stable a grasp is. Therefore, once the robot successfully grasped an object, we perform a dynamic movement of the arm and check if the hand afterwards still holds the object.

2.2. Application

For the application setup we use the previously collected data to predict good grasping points as described in [4]. Furthermore we attempt to perform this analysis in an on-line manner. Instead of assuming a fixed database of labeled data we enable to extend this database over time and update the prediction accordingly (cf. Figure 3b). This requires us to derive the kernel parameters at every update of the database. Since this is computational intensive we try to keep the amount of training data small but expressive.

To do so we perform the following evaluation after every grasping experiment for the feature vector (obtained by generic filters applied to the saturation channel of the image) at that particular point: In a first step we determine if a feature $\mathbf{x}$ improves our current prediction of the probability distribution: $\lambda \cdot (p - p_*)^2 > (p - p'_*)^2$ with $0 < \lambda \leq 1$. Whereas p is the observed outcome, p_* the prediction based on the previous state of the database and p'_* the prediction taking feature $\mathbf{x}$ into account. In case it actually improves our prediction we search the database for similar features in a second step. Given the database already contains a feature $\mathbf{x}'$ that is "identical" $d(\mathbf{x}, \mathbf{x}') < \varepsilon_1$ only the *observed* probability of feature $\mathbf{x}'$ is adjusted. Otherwise we weight the adjustment of the observations for all features within a distance ε_2. Only if there are no similar features the new feature $\mathbf{x}$ is included without modifications.

3. Discussion and future work

Due to the use of YARP [2] for the implementation of the software components, the system is also applicable to other robotic systems than the iCub. However, the homography for projecting grasping points is bound to a robot's kinematics and needs to be adapted accordingly. The second component of the homography refers to the projection plane and it's position relative to the robot. Therefore, it is possible to vary the plane of operation (table tops, floor, etc.) by adjusting this transformation. However, this represents another constraint to the experiment's current environment. To overcome this limitation we consider an extension to multiple predictions from different camera positions and/or the use of stereo cameras as motivated in [8] in order to replace the homography.

4. Conclusions

We outlined a system capable of independently performing the necessary experiments for learning object affordances. Moreover we discussed strategies of selecting expressive features with respect to the approach presented in [4]. The thereby achieved reduction of the required training set for the statistical model enables an on-line extension of the approach.

5. Acknowledgments

This work was supported by Fundação para a Ciência e a Tecnologia (ISR/IST pluriannual funding). For our initial experiments on on-line learning we used the *"Stanford synthetic object grasping point data"* [6, 7].

References

[1] James J. Gibson. *The Ecological Approach to Visual Perception*. Lawrence Erlbaum Associates, 1 edition, October 1979.

[2] Giorgio Metta, Paul Fitzpatrick, and Lorenzo Natale. YARP: Yet Another Robot Platform. *International Journal on Advanced Robotics Systems, Special Issue on Software Development and Integration in Robotics*, 3(1):43–48, March 2006.

[3] Giorgio Metta, Giulio Sandini, David Vernon, Lorenzo Natale, and Francesco Nori. The icub humanoid robot: an open platform for research in embodied cognition. In *PerMIS: Performance Metrics for Intelligent Systems Workshop*, August 2008.

[4] Luis Montesano and Manuel Lopes. Learning grasping affordances from local visual descriptors. In *IEEE 8th International Conference on Development and Learning*, pages 1–6, June 2009.

[5] Luis Montesano, Manuel Lopes, Alexandre Bernardino, and José Santos-Victor. Learning object affordances: From sensory – motor coordination to imitation. *IEEE Transactions on Robotics and Automation*, 24(1):15–26, February 2008.

[6] Ashutosh Saxena, Justin Driemeyer, Justin Kearns, and Andrew Y. Ng. Robotic grasping of novel objects. In *Neural Information Processing Systems 19*, December 2006.

[7] Ashutosh Saxena, Justin Driemeyer, Justin Kearns, Chioma Osondu, and Andrew Y. Ng. Learning to grasp novel objects using vision. In *10th International Symposium of Experimental Robotics*, July 2006.

[8] Ashutosh Saxena, Justin Driemeyer, and Andrew Y. Ng. Robotic grasping of novel objects using vision. *The International Journal of Robotics Research*, 27(2):157–173, February 2008.