Learning Object Affordances For Tool Use And Problem Solving In Cognitive Robots

Lorenzo Jamone, Giovanni Saponaro, Alexandre Antunes,
Rodrigo Ventura, Alexandre Bernardino, and Jose Santos-Victor

Instituto Superior Tecnico, Lisbon, Portugal
ljamone@isr.tecnico.ulisboa.pt

Abstract. One of the hallmarks of human intelligence is the ability of predicting the consequences of actions and efficiently plan behaviors based on such predictions. This ability is supported by internal models that human babies acquire incrementally during development through sensorimotor experience: i.e. by interacting with objects in the environment while being exposed to sensori perception. An elegant and powerful concept to represent these internal models has been proposed in developmental psychology under the name of *object affordances*: action possibilities that an object offers to an agent. Affordances are learned ecologically by the agent and exploited for action planning. Clearly, endowing artificial agents with such cognitive capabilities is a fundamental challenge both in artificial intelligence and robotics. We propose a learning framework in which an embodied agent (i.e. in our case, the humanoid robot iCub) autonomously explores the environment, and learns object affordances as probabilistic dependencies between actions, object visual properties and observed effects; we use Bayesian Networks to encode this probabilistic model. By making inferences across the learned dependencies a number of cognitive skills are enabled: e.g. i) predicting the effects of an action over an object, or ii) selecting the best action to obtain a desired effect. By exploring object-object interactions the robot can develop the concept of *tool* (i.e. a handheld object that allows to obtain a desired effect on another object), and eventually use the acquired knowledge to plan sequences of actions to attain a desired goal (i.e. problem solving).

Keywords: Cognitive humanoid robots, affordances, tool use, prediction, planning, problem solving, Bayesian Network

1 Introduction

Humans solve complex tasks on a routine basis, by choosing, amongst a vast repertoire, the most proper actions to apply onto objects in order to obtain certain effects. According to developmental psychology [1], the ability to predict the functional behavior of objects and their interaction with the body, simulating and evaluating the possible outcomes of actions before they are actually executed, is one of the purest signs of cognition, and it is acquired incrementally during development through the interaction with the environment. Neuroscientific evidence [2] supports the idea that, in the brain, these predictions happen during action planning through the activation of sensorimotor structures that couple sensory and motor signals. To reproduce such intelligent behavior in robots is an important, hard and ambitious task. One possible way to tackle this problem is to resort to the concept of affordances, introduced by Gibson in his seminal work [3]. He defines affordances as action possibilities available in the environment to an individual, thus depending on its action capabilities. From the perspective of robotics, affordances are powerful since they capture the essential world and object properties, in terms of the actions that a robot is able to perform. They can be used to predict the effects of an action, or to plan the actions to achieve a specific goal; by extending the concept further, they can facilitate action recognition and be exploited for robot imitation [4], they can be a basis to learn tool use [6, 5], and they can be used together with planning techniques to solve complex tasks. We propose a probabilistic model of affordances that relates the shape properties of a hand held object (intermediate) and an acted object (primary) with the effects of the motor actions of the agent, measured as relative displacements of the primary object. We performed experiments in which the iCub [7] humanoid robot learns these object affordances by performing numerous actions on a set of objects displaced on a table (see Fig. 1). The learned model can then be used to predict the consequences of actions, leading to behaviors such as tool use and problem solving.

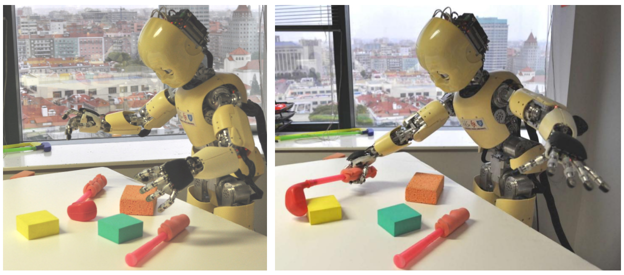

Fig. 1. The iCub humanoid robot standing in front of a table full of objects. The knowledge of the objects affordances can be exploited for problem solving.

2 Related work

Many computational models have been proposed in the literature in order to equip robots with the ability to learn affordances and use them for prediction and planning. The concept of affordances and its implications in robotics are discussed by Sahin et al. [8], who propose a formalism to use affordances at different levels of robot control; they apply one part of their formalism for the learning and perception of traversability affordances on a mobile robot equipped with range sensing ability [9]. In the framework presented by Montesano et al. [10], objects affordances are modeled with a Bayesian Network [11], a general probabilistic representation of dependencies between actions, objects and effects; they also describe how a robot can learn such a model from motor experience and use it for prediction, planning and imitation. Since learning is based on a probabilistic model, the approach is able to deal with uncertainty, redundancy and irrelevant information. The concept of affordances has also been formalized under the name of object-action complexes (OACs, [12]).

3 A computational model of affordances

We follow the framework of [10], where the relationships between an acted object, the applied action and the observed effect are encoded in a causal probabilistic model, a Bayesian Network (BN)whose expressive power allows the marginalization over any set of variables given any other set of variables. It considers that actions are applied to a single object using the robot hands, whereas we model interobject affordances, including new variables that represent the intermediate object as an individual entity, as depicted in Fig. 2 (left side). The BN of our approach explicitly models both primary (acted) and intermediate (held) objects, thus we can infer i) affordances of primary objects, ii) affordances of intermediate objects, and iii) affordances of the interaction between intermediate and primary objects. For example, our model can be used to predict effects given both objects and the performed action, or choose the best intermediate object (tool) to achieve a goal (effect to be produced on a primary object). Both objects are represented in the BN network as a set of basic shape features obtained by vision (e.g. convexity, eccentricity). Further details can be found in [5].

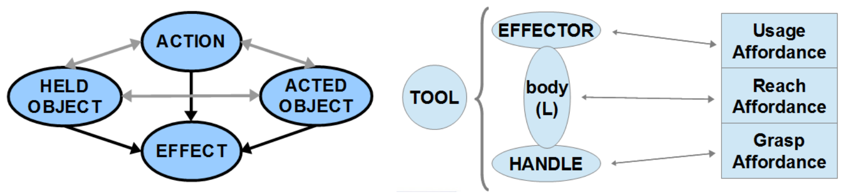

Fig. 2. On the left: Bayesian Network model of affordances, modeled as relations between actions, effects and objects (both held and acted). On the right: general model for tool use. The model on the left corresponds to the Usage Affordances part of the model on the right.

3.1 A model for tool use

The affordances of the intermediate hand-held object (i.e. the tool) can be incorporated in a more complete model for cognitive tool use. Tools can be typically described by three functional parts: a handle, an effector, and a body of a certain length L connecting the two (see right part of Fig. 2). These three parts are related to three different motor behaviors humans have to perform in order to successfully use a tool: grasping the handle, reaching for a desired pose with the effector and then executing an action over an affected object. Each of those behaviors requires some prior mental reasoning, first to estimate whether the behavior is feasible (e.g. is the handle graspable?) and then to plan the correct motion to be executed (e.g. determine the target hand pose and the finger movements to grasp the tool). We can therefore define three levels of tool affordances: i) *usage affordances*, ii) *reach affordances* and iii) *grasp affordances* (see right part of Fig. 2). These affordances relate to specific problems: i) what actions the tool affords, because of its effector properties, ii) what part of the workspace the tool affords to reach for, depending on its length, iii) what grasps the tool affords, based on the shape and size of the handle. The outcomes of these three reasoning processes are based on internal models that the robot can learn through motor exploration. The model of affordances in the left part of Fig. 2 represents the *usage affordances*. In previous work we proposed a learning framework that enables a robot to learn its own body schema [13–15], and to update it when tools are included [16, 17], and a representation of its own reachable space [18, 19]; these internal models are related to the *reach affordances*. Also, a number of models for learning and using *grasp affordances* have been proposed in the literature (e.g. [20, 21]).

3.2 Use affordances for problem solving

Since the early days of Artificial Intelligence (AI), planning techniques have been employed to allow agents to achieve complex tasks in closed and deterministic worlds. Every action has clearly defined pre-conditions, and generates deterministic post-conditions. However, these assumptions are not plausible if we consider a real robot acting in real unstructured environments, where the consequences of actions are not deterministic and the world is perceived through noisy sensing.
The affordance model (and more generally, the tool use model) depicted in Fig. 2 provide probabilistic predictions of actions consequences, that depend on the perceived visual features of the objects and on the robot sensorimotor abilities and previous experiences. Inspired by recent advances in AI, we can use these predictions within probabilistic planning algorithms, to achieve a grounding of the planning operators based on the robot sensorimotor knowledge. Through this computational machinery, the robot is able to plan the sequence of actions that has the higher probability to achieve the goals, given its own motor abilities and the perceived properties of the available objects.

References

1. C. von Hofsten, An action perspective on motor development. Trends in Cognitive Sciences, 8, 266272. 2004.
2. V. Gallese, L. Fadiga, L. Fogassi, and G. Rizzolatti, Action recognition in the premotor cortex. Brain, 119, 593609. 1996.
3. J. J. Gibson, The Ecological Approach to Visual Perception. Boston, MA: Houghton Mifflin. 1979.
4. M. Lopes, F. S. Melo and L. Montesano, Affordance-based imitation learning in robots. IEEE IROS. 2007.
5. A. Goncalvez, J. Abrantes, G. Saponaro, L. Jamone and A. Bernardino, Learning Intermediate Object Affordances: Towards the Development of a Tool Concept. IEEE ICDL/Epirob. 2014.
6. A. Goncalvez, G. Saponaro, L. Jamone and A. Bernardino, Learning Visual Affordances of Objects and Tools through Autonomous Robot Exploration. IEEE ICARSC. 2014.
7. G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. von Hofsten, K. Rosander, M. Lopes and J. Santos-Victor, The iCub humanoid robot: An open-systems platform for research in cognitive development, Neural Networks, 23, 11251134. 2010.
8. E. Sahin, M. Cakmak, M. R. Dogar, E. Ugur and G. Ucoluk, To Afford or Not To Afford: A New Formalization of Affordances Toward Affordance-Based Robot Control, Adaptive Behavior, 15(4), 447472, 2007.
9. E. Ugur and E. Sahin, Traversability: A Case Study for Learning and Perceiving Affordances in Robots. Adaptive Behavior, 18(3-4), 258284. 2010.
10. L. Montesano, M. Lopes, A. Bernardino and J. Santos-Victor, Learning Object Affordances: From Sensory Motor Coordination to Imitation, IEEE Transactions on Robotics, 24(1), 1526. 2008.
11. J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann. 1988.
12. N. Kruger, C. Geib, J. Piater, R. Petrick, M. Steedman, F. Worgotter, A. Ude, T. Asfour, D. Kraft, D. Omrcen, A. Agostini, and R. Dillmann. ObjectAction Complexes: Grounded abstractions of sensorymotor processes. Robotics and Autonomous Systems, 59(10), 740757. 2011.
13. L. Jamone, L. Natale, K. Hashimoto, G. Sandini and A. Takanishi, Learning Task Space Control through Goal Directed Exploration. IEEE ROBIO. 2011.
14. L. Jamone, L. Natale, G. Metta, F. Nori and G. Sandini, Autonomous Online Learning of Reaching Behavior in a Humanoid Robot. International Journal of Humanoid Robotics, 9(3). 2012.
15. B. Damas, L. Jamone and J. Santos-Victor, Open and Closed-Loop Task Space Trajectory Control of Redundant Robots Using Learned Models. IEEE IROS. 2013.
16. L. Jamone, B. Damas, J. Santos-Victor and A. Takanishi, Online Learning of Humanoid Robot Kinematics Under Switching Tools Contexts. IEEE ICRA. 2013.
17. L. Jamone, B. Damas, N. Endo, J. Santos-Victor and A. Takanishi, Incremental development of multiple tool models for robotic reaching through autonomous exploration. PALADYN Journal of Behavioral Robotics, 3(3), 113-127. 2013.
18. L. Jamone, L. Natale, G. Sandini and A. Takanishi, Interactive online learning of the kinematic workspace of a humanoid robot. IEEE IROS. 2012.

19. L. Jamone, M. Brandao, L. Natale, K. Hashimoto, G. Sandini and A. Takanishi, Autonomous online generation of a motor representation of the workspace for intelligent whole-body reaching. Robotics and Autonomous Systems, 64(4), 556-567. 2014.
20. L. Montesano and M. Lopes, Learning grasping affordances from local visual descriptors. IEEE ICDL. 2009.
21. R. Detry, D. Kraft, O. Kroemer, L. Bodenhagen, J. Peters, N. Kruger and J. Piater, Learning Grasp Affordance Densities. PALADYN Journal of Behavioral Robotics, 2(1), 117. 2011.