Robotic tool use and problem solving based on probabilistic planning and learned affordances

Alexandre Antunes, Giovanni Saponaro, Atabak Dehban, Lorenzo Jamone,
Rodrigo Ventura, Alexandre Bernardino, José Santos-Victor
Institute for Systems and Robotics, Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal
{aantunes, gsaponaro, adehban, ljamone, rodrigo.ventura, alex, jasv}@isr.tecnico.ulisboa.pt

Abstract*— The ability of predicting the consequences of actions is a distinctive sign of human intelligence. Predictions are used to efficiently plan actions sequences that allow problem solving. This ability is acquired incrementally by humans during development through sensorimotor experience: i.e. interacting with objects while perceiving salient events. An appealing concept to represent this sensorimotor knowledge has been proposed in psychology under the name of *object affordances*: action possibilities that an object offers to an agent. Affordances are learned by the agent (animal or robot) 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 framework in which an embodied agent (i.e. in our case, the humanoid robot iCub) explores the environment and learns object affordances as relationships between actions, object visual properties and observed effects. The learned affordances enable a number of cognitive skills: 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 hand-held 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).

I. 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 [5], [6], and they can be used together with planning techniques to solve complex tasks. We propose a 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 have explored two learning approaches: probabilistic learning (Bayesian Networks, BN [11]) and neural networks (Denoising Auto-encoders, dA [22]). The iCub [7] humanoid robot learns these affordances by performing numerous actions on a set of objects displaced on a table (see Fig. 1). The learned models can then be used to predict the consequences of actions, leading to behaviors such as tool use and problem solving.

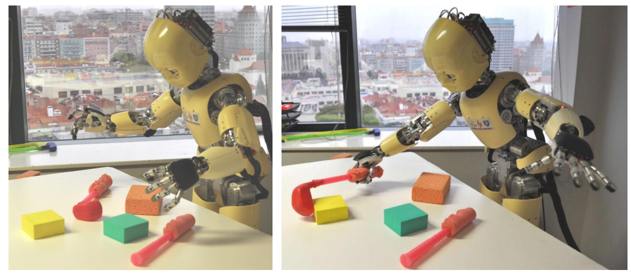

Fig. 1: The iCub humanoid robot standing in front of a table full of objects. The knowledge of the objects affordances can be exploited for problem solving.

II. 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]).

III. Computational models of affordances

In recent work [5], [6] we proposed a model in which the relationships between a hand-held object, 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 another object held by the robot (i.e. an intermediate object, a tool), as depicted in Fig. 2 (top image). 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], [6].
Recently we also explored a different learning approach to estimate these sensorimotor relationships, based on Denoising Auto-encoders (dA [22]). This approach has a number of appealing features. First, since dA are a variant of neural networks, they can be applied in online learning scenarios. Second, they impose no constraint on the type of learning data (continuous or discrete). Third, they have been successfully applied to multi-modal information fusion, allowing to retrieve one modality given the others [23]. Therefore, by using this approach we can obtain the same expressive power of the BNs (i.e. the possibility of infer any set of variables given any other set of variables), but with the advantage of using continuous data and online incremental estimation of the model. This is not possible with classic BNs, that require data to be discretized/clustered, and are learned in batch mode.

A. 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 bottom 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

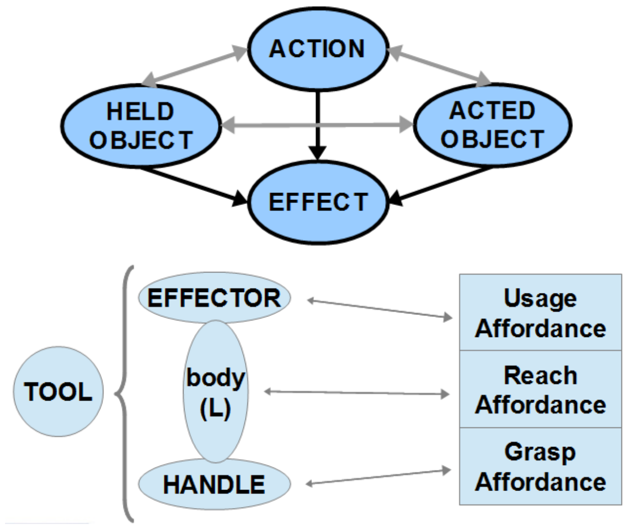

Fig. 2: Top image: Bayesian Network model of affordances, modeled as relations between actions, effects and objects (both held and acted). Bottom image: general model for tool use. The model in the top image corresponds to the Usage Affordances part of the model in the bottom image.

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]).

B. 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. More specifically, within the EU Project Poeticon++, we address the problem of having a robot executing motor tasks requested by a human through spoken language (i.e., verbal instructions). Verbal instructions do not typically have a one-to-one mapping to robot actions, due to various reasons: economy of spoken language, e.g., one short instruction might indeed correspond to a complex sequence of robot actions, and details about action execution might be omitted; grounding, e.g., some actions might need to be added or adapted due to environmental contingencies; embodiment, e.g., a robot might have different means than the human ones to obtain the goals that the instruction refers to. Therefore, we propose a general cognitive architecture to deal with these issues, based on three steps: i) language-based semantic reasoning on the instruction (high level), ii) formulation of goals in robot-symbols and probabilistic planning to achieve them (mid-level), iii) action execution (low-level). The mid-level is the layer in which we take in consideration the robot prior sensorimotor experience (i.e. the learned affordances) to ground the planning operators. Our system merges semantic knowledge from high-level with robot prior knowledge coming from the affordances and perceptual knowledge about the current context coming from the robot sensors, by using probabilistic planning; the output of the mid-level is the motor actions to be executed by the low-level, which in our case is represented by the iCub robot motor control routines.

IV. Conclusions

We present here our overall effort in modeling object affordances and use such knowledge to enable cognitive tool use and problem solving in an embodied artificial agent (i.e. the iCub humanoid robot). A number of recent [5], [6] and forthcoming papers describe our approaches and results in more detail.
All the software we develop is publicly available (`https://github.com/robotology/poeticon`).

Acknowledgments

This work was partially supported by the Portuguese Government – Fundação para a Ciência e a Tecnologia (UID/EEA/50009/2013) and by the European Commission under projects POETICON++ (FP7-ICT-288382) and LIMOMAN (PIEF-GA-2013-628315).

References

[1] C. von Hofsten, An action perspective on motor development. Trends in Cognitive Sciences, 8, 266–272. 2004.
[2] V. Gallese, L. Fadiga, L. Fogassi, and G. Rizzolatti, Action recognition in the premotor cortex. Brain, 119, 593–609. 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. Gonçalves, 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. Gonçalves, 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, 1125–1134. 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), 447–472, 2007.
[9] E. Ugur and E. Sahin, Traversability: A Case Study for Learning and Perceiving Affordances in Robots. Adaptive Behavior, 18(3-4), 258–284. 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), 15–26. 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. Object–Action Complexes: Grounded abstractions of sensory–motor processes. Robotics and Autonomous Systems, 59(10), 740–757. 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), 1–17. 2011.
[22] P. Vincent, H. Larochelle, Y. Bengio and P.A. Manzagol, Extracting and composing robust features with denoising autoencoders. Proceedings of the 25th international conference on Machine learning, 1096-1103. 2008.
[23] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee and A. Y. Ng, Multimodal deep learning. Proceedings of the 28th international conference on Machine learning, 689-696. 2011.