

Modeling Affordances using Bayesian networks

Luis Montesano Manuel Lopes Alexandre Bernardino Jose Santos-Victor

Abstract—Affordances represent the behavior of objects in terms of the robot’s motor and perceptual skills. This type of knowledge plays a crucial role in developmental robotic systems, since it is at the core of many higher level skills such as imitation. In this paper, we propose a general affordance model based on Bayesian networks linking actions, object features and action effects. The network is learnt by the robot through interaction with the surrounding objects. The resulting probabilistic model is able to deal with uncertainty, redundancy and irrelevant information. We evaluate the approach using a real humanoid robot that interacts with objects.

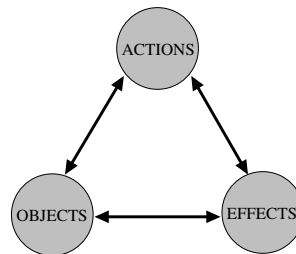
I. INTRODUCTION

The development of embodied systems able to act in a complex world and to interact with humans requires the definition of adequate knowledge representations to support the execution of a large number of tasks. In this paper, we study a concept in biological systems known as affordances [1] and propose a computational model to acquire and use them.

Gibson introduced the affordance concept to define the relation between an agent and its environment through its motor and sensing capabilities (e.g. graspable, movable or eatable) as illustrated in Fig. 1. For instance, humans can grasp a cup or sit on a sofa, but not vice versa. Dogs can sit on a sofa but cannot grasp a cup.

A robot can exploit this type of knowledge to understand the behavior of the world in terms of its own actions/perceptions. It can predict the effects of an action executed upon an object or select the most appropriate object for a given goal. These basic capabilities are at the core of higher level cognitive skills such as making decisions or interacting with other robots or humans. Indeed, they allow to recognize actions and intentions from other agents and can be used, for instance, in imitation [2], [3]. Thus, affordances are a suitable representation of the world in terms of the robot capabilities which may be used as common knowledge for higher development levels.

In order to model, acquire and use affordances, one must consider two important properties. First, affordances appear from the interaction between the robot and the environment and, hence, depend on the world, the agent’s motor and perceptual capabilities and its experience. For instance, the lack of sound perception prevents learning the affordances associated to a drum. Second, the robot needs to have ready a



inputs	outputs	function
(O, A)	E	Predict Effect
(O, E)	A	Recognize action & Planning
(A, E)	O	Object recognition & selection

Fig. 1. Affordances as relations between (A)ctions, (O)bjects and (E)ffects, that can be used to solve different questions: prediction, action selection or object selection.

set of basic capabilities and concepts that allow it to interact with the world. In other words, the robot should be at a developmental stage [3] where it can detect objects, perform basic manipulations and measure the resulting effects.

The contributions of this paper are threefold. Firstly, we propose a model for affordances based on Bayesian Networks (BN) [4] a general probabilistic representation for causal dependencies. Although BN have been used in many contexts to model statistical dependencies (e.g. for relational data [5]), it is, to our knowledge, the first time they are used to model affordances. Secondly, we show how this model can be learned from observations as a structural learning method, where affordances are encoded as probabilistic relations between actions, objects and action effects. Finally, we show how the learned affordance model can be used for inference or planning in several types of tasks.

The proposed model has been evaluated using the humanoid robot Baltazar (see Fig. 1). The robot used grasp, touch and tap actions on different objects and measured the effects as salient changes in the sensor measurements (e.g. object motion). The results suggest that the learned network correctly models the dependencies between actions, object features and the resulting effects and detects irrelevant features.

The paper is organized as follows. Next section discusses the related work. Section III describes our affordance learning using Bayesian networks. Section V validates the proposed model. In Section VI we draw the conclusions and comment future developments.

II. RELATED WORK

According to Gibson’s definition, affordances describe the relation between a living being and its environment [1].

L. Montesano, M. Lopes A. Bernardino J. Santos-Victor are with the Instituto de Sistemas e Robótica, Instituto Superior Técnico, Lisboa, Portugal. {lmontesano,macl,alex,jasv}@isr.ist.utl.pt

This work was (partially) supported by the FCT Programa Operacional Sociedade de Informação (POSC) in the frame of QCA III and the project PTDC/EEA-ACR/70174/2006, and by the EU Proj. (IST-004370) RobotCub.

Gibson argues that this relation is shaped by the perceptual and motor abilities of the agent. Hence, the word affordance represents what the elements of the environment afford to the agent. In ecological psychology, there is still an ongoing discussion to establish a definition or model for the affordances (see [6] for a brief review). Some authors have shown different examples of affordance in humans comparing the perceptions among different people [7], measuring response times to task elicited by specific orientations [8] or perceiving heaviness [9]. Unfortunately, there is few evidence of how humans learn affordances.

In a robotic context, we can distinguish two different approaches. The first one simply uses affordances as prior knowledge linking objects and actions. In [10] a computational cognitive model of infant grasp learning is proposed, where affordances model information that helps the modules performing the action. Affordances have also been used as prior distributions for action recognition in a Bayesian framework [11] or to perform selective attention in obstacle avoidance tasks [12].

The second type of approaches focuses on the learning of emergent affordances and their application to different tasks. In [13], a robot learned the direction of motions of different types of objects after poking and used this information to recognize actions performed by other agents. Then, the robot used these maps to obtain the same motion of an observed action. In [14] imitation is also driven by the effects. However, they focus on the interaction aspects and do not consider a general model for affordances. The biologically inspired behavior selection mechanism of [15] uses clustering and self organizing feature maps to relate object invariants to the success or failure of an action. All the previous approaches learn a specific type of affordance using the relevant information extracted from their sensor channels. A more complete solution has been recently proposed in [16] where a supervised learning procedure also selects the appropriate features from a set of visual SIFT descriptors. The work in [17] focuses on the importance of sequences of actions and invariant perceptions to discover affordances in a behavioral framework.

III. MODELING AFFORDANCES

In this section, we first describe the assumptions in terms of robot capabilities required to start learning the affordances. Then, we present the affordance model and how it can be used to make inference for prediction, recognition and planning purposes.

A. Robot skills and notation

We assume that the robot is endowed with a set of skills (either programmed or learned through experience) that support reasoning at a more abstract level than joint positions or raw percepts. We assume that the robot is able to detect nearby objects and measure basic features like position, color, shape or size. In addition we also hypothesize that the robot can interact with objects in a simple manner through a pre-existing action repertoire and can extract some

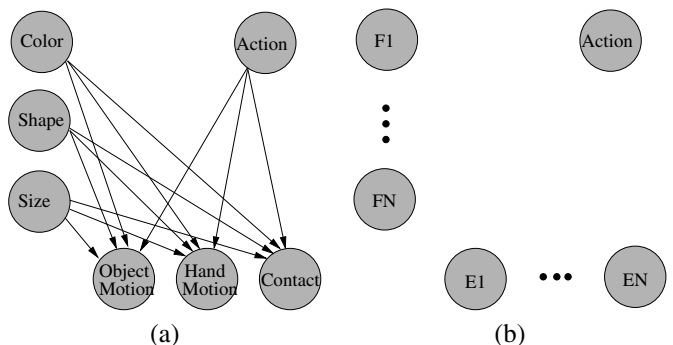


Fig. 2. Bayesian network model to represent the affordances. (a) An example of the previous model using color, shape and size information for the object features and motion and contact information as effects. (b) Generic model where the nodes represent the actions A , the object features available to the robot $F(1)...F(n)$ and the effects obtained through the actions $E(1)...E(m)$.

features from its own body such as its hand position and velocity in the image. We address the affordance learning problem at this level of abstraction where the main entities are the robot actions, the features of the objects and the resulting effects. Details on the implementation of these skills are given in Section V.

More formally, let the discrete random variable $A = \{a_i\}$ represent the execution of a robot action. Object properties and effects are also modeled using discrete random variables corresponding to the classes detected by the robot. We denote $F_r = \{F_r(1), \dots, F_r(n_r)\}$ the robot own features extracted by the pre-processing modules (for instance, hand position in the image, proprioception,...) and $F_o = \{F_o(1), \dots, F_o(n_o)\}$ object features extracted for object o . Finally, let $E = \{E(1), \dots, E(n_e)\}$ be the effects detected by the robot after executing an action. The set of nodes X is formed by the discrete variables A, F_r, F_o and E , $X = \{A, F_r, F_o, E\}^1$.

B. Affordance model

This section introduces Bayesian networks as the framework to represent the relations between actions A , robot features F_r , object features F_o and effects E . We use a Bayesian network (BN) to encode the dependencies among these variables (see Fig. 2).

A Bayesian network (BN) $B = (G, \Theta)$ over a set of variables $X = \{X_1, \dots, X_n\}$ is a probabilistic directed graphical model where the nodes of the graph G represent the random variables X and the (lack of) arcs represent conditional independence assumptions. $\Theta = \{\theta_i\}$ represents the set of parameters defining the conditional probability distributions $p(X_i | X_{Pa(X_i)}, \theta_i)$ of each node in the graph depending on the parents $X_{Pa(X_i)}$.

Affordances can be represented using BNs in a very natural way, since an arc from $X_i \rightarrow X_j$ can be interpreted as X_i causes X_j (see [18]). In this way, affordances are actually represented by the arrows between nodes and the parameters of the resulting conditional distributions. This representation

¹We represent random variables by capital letters X and its realizations as x .

has several advantages. It allows to take the uncertainty of the features and effects perception into account and provides a unified framework for learning and using affordances.

C. Inference

Since the structure of the BN encodes the relations between actions, object features and effects, we can now compute the distribution of a (group of) variable(s) given the values of the others. The most common way to do this is to convert the BN into a tree and then apply the junction tree algorithm [19] to compute the distributions of interest. It is important to note that it is not necessary to know the values of all the variables to perform inference.

Based on these probabilistic queries, we are able to use the affordance knowledge to answer the questions of Fig. 1 simply by computing the appropriate distributions. For instance, the prediction of the effects when observing an action a_i given the observed object features f_j is just $p(E | A = a_i, F = f_j)$. It is important to note that the query can combine features, actions and effects both as observed information or as the desired output.

IV. LEARNING AFFORDANCES

In this section, we address the affordance learning task. As mentioned before, affordances arise from the interaction between the agent and the surrounding objects. Thus, the robot has to apply its actions to different objects and observe the resulting effects. Let $D = X^{1:N}$ be the set of such trials. We want to discover the network structure and the parameters that best model the interaction between the robot and the objects.

We are interested in learning the structure G , which is actually an instance of a model selection problem. In a Bayesian framework, this can be formalized as estimating the distribution on the possible network structures $G \in \mathcal{G}$ given the data. Using the Bayes rule, we can express this distribution as the product of the marginal likelihood and the prior over graphs,

$$p(G | D) = \eta p(D | G) p(G) \quad (1)$$

where $\eta = p(D)^{-1}$ is the normalization constant. The prior term $p(G)$ allows to incorporate prior knowledge on possible structures.

Given the discrete representation of actions, features and effects, we use multinomial distributions and their corresponding conjugate, the Dirichlet distribution, to model the CPDs $p(X_i | X_{Pa(X_i)}, \theta_i)$ and the corresponding parameter priors $p(\theta_i)$. Using this representation, the marginal likelihood for a node X_i and its parents given the trials D is [20]:

$$\begin{aligned} p(x_i^{1:N} | x_{Pa(x_i)}^{1:N}) &= \int \left[\prod_{n=1}^N p(x_i^n | x_{Pa(x_i)}^n, \theta_i) \right] p(\theta_i) d\theta_i \\ &= \prod_{j=1}^{|X_i|} \frac{\Gamma(\alpha_{ij})}{\Gamma(\alpha_{ij} + N_{ij})} \prod_{k=1}^{|X_{Pa(x_i)}|} \frac{\Gamma(\alpha_{ijk} + N_{ijk})}{\Gamma(\alpha_{ijk})} \end{aligned}$$

where N_{ijk} counts the number of datasets with $X_i = j$ and $X_{Pa(X_i)} = k$ and $N_{ij} = \sum_k N_{ijk}$ and Γ represents the gamma function. The pseudo-counts α_{ijk} are the Dirichlet hyper parameters of the selected prior distribution on θ_i and $\alpha_{ij} = \sum_k \alpha_{ijk}$. The marginal likelihood of the data is simply the product of the marginal likelihood of each node,

$$p(D | G) = p(X^{1:N} | G) = \prod_i p(x_i^{1:N} | x_{Pa(x_i)}^{1:N}) \quad (2)$$

where we have made explicit the dependency on the graph structure G .

Unfortunately, the number of BNs is super exponential with the number of nodes [21]. Thus, it is infeasible to explore all the possible graphs and one has to approximate the full solution. This can be done using Markov Chain Monte Carlo (MCMC) [22] to construct a chain whose stationary distribution is $p(G | D)$.

The previous learning scheme is able to distinguish among equivalence classes². So as to be able to infer the correct causal dependency, it is necessary to use interventional data where we fixed some of the variables to a specific value to disambiguate between graphs in the same equivalence class.

In the case of a robot interacting with its environment, there are several variables that are actively chosen by the robot: the action and the object. In other words, they are interventional data set by the robot to some specific value at each trial. Interventional data is currently an important research topic within BN learning algorithms (see [23]). Under the assumption of a perfect intervention of node i , the value of $X_i = x_i^*$ is set to the desired value and its CPD is just an indicator function with all the probability mass assigned to this value $p(X_i | X_{Pa(X_i)}, \theta_i) = I(X_i = x_i^*)$. As a result, the variable X_i is effectively cut off from its parents $X_{Pa(X_i)}$. More sophisticated models allow for uncertain interventions (see [23]).

Once the structure of the network has been established, the parameters θ_i of each node can be estimated using a Bayesian approach [24]. The estimated parameters can still be sequentially updated on-line allowing to incorporate the information provided by new trials.

V. EXPERIMENTS

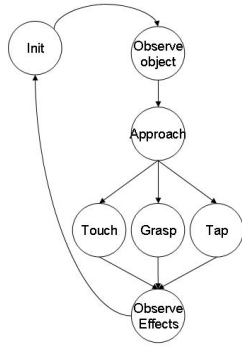
In this section we present a set of experimental results to illustrate the acquisition and usage of affordance knowledge. We used Baltazar, a 14 degrees of freedom humanoid torso composed by a binocular head and an arm. Baltazar is able to perform three different actions $A = \{a_1 = \textit{grasp}, a_2 = \textit{tap}, a_3 = \textit{touch}\}$. The robot applies its actions on a set of different objects with two shapes (box and ball), with different colors and three sizes (see Fig. 3)(a).

We recorded a set of 250 trials following the protocol depicted in Fig. 3(b). At each trial, the robot was presented with an object. Baltazar randomly selected an action and

²Two directed acyclic graphs G and G' are equivalent, if for every BN $B = (G, \Theta)$ there exist another network $B' = (G', \Theta')$ such that they define the same probability distribution.



(a)



(b)

Fig. 3. (a) The playground for the robot contains objects of several sizes, colors and shapes. (b) Experiments protocol. The object to interact is selected manually and the action is randomly selected. Object properties are recorded in the INIT to APPROACH transition when the hand is not occluding the object. The effects are recorded in the OBSERVE state. INIT moves the hand to a predefined position in open-loop.

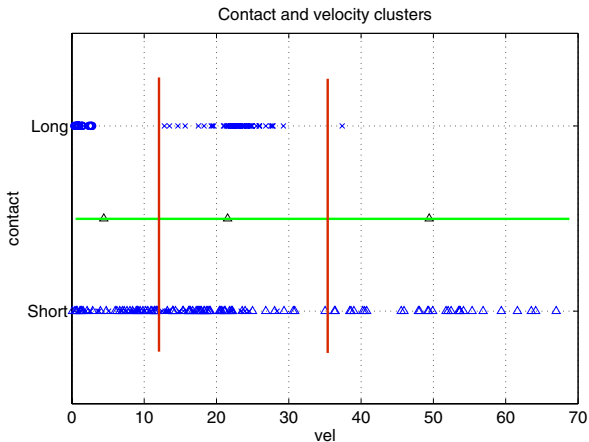


Fig. 4. Clustering of velocity and contact. The figure show for each trial (grasp is represented by x , tap by Δ and touch by o) the measured velocity (x -axis) and the contact information (y -axis). The vertical lines show the cluster boundaries for velocity and the horizontal line for contact.

approached its hand to the object. When the reaching phase is completed, it performed the selected action (*tap*, *grasp* or *touch*) and finally returns the hand to the initial location. During the action, both object features and resulting effects are recorded.

A. Robot skills

The robot motor skills have been learned following the framework presented in [11]. Each action is executed in three steps. First, the hand is moved into the field of view of the robot using a static sensory-motor map. Then, the hand is moved toward the object in closed-loop using visual servoing and finally the robot performs the action.

Regarding perceptual skills, we have selected three features to describe an object: color, shape and size. Each one is modeled as a n -dimensional vector space: a 16 bin hue histogram for color; a 6- d vector for shape based on geometric properties such as circleness or eccentricity; and

the area in the image for size. The effect detectors measure the object velocity and the hand velocity in the image and the distance between them. A simple contact sensor measures the persistent contact between object and hand. Since our setup is clearly discrete, we applied the X -means algorithm [25] to form clusters in the space of each object feature and effect. Table I summarizes the clustering results for the different variables and provides the notation used in the remainder of this section. The clusters are computed prior to the learning of the affordance networks and they represent the domain of each node of the network.

TABLE I

SUMMARY OF VARIABLES AND VALUES.

Symbol	Description	Values
A	Action	<i>grasp, tap, touch</i>
C	Color	clustered in <i>green₁, green₂, yellow, blue</i>
Sh	Shape	clustered in <i>ball, box</i>
S	Size	clustered in <i>small, medium, big</i>
V	Object velocity	clustered in <i>small, medium, big</i>
HV	Hand velocity	clustered in <i>small, big</i>
Di	Object Hand Distance	clustered in <i>small, medium, big</i>
Ct	Contact	clustered in <i>short, long</i>

We would like to stress that the final purpose is to learn the affordances given a set of available motor and perceptual skills, not to make a perfect object classification. Due to the simplicity of the perception process and different working conditions, clustering errors occur. The learning task has to cope with object and effect misclassification.

Figure 4 shows the result of the clustering algorithm for the object velocities and contact information. Roughly, a grasp action resulted in medium velocity (except in one case where the ball fell down the table) while tap produced different velocity patterns depending on the shape and size of the object. The combination of different features produces different patterns that can be used to infer statistical dependencies and causation.

B. Affordances

Based on the previous descriptors of actions, object features and effects, we present some experiments to illustrate the ability of the proposed model to represent and learn affordances. We would like to remark that the robot does not receive any information about the success or not of the actions. The interest is in understanding the effects obtained by the actions in an unsupervised manner.

As mentioned before, affordance knowledge depends on the robot skills. We first show how the robot is able to distinguish the effects of different actions and simultaneously perform feature selection discarding irrelevant features. We present two cases. In the first one, the robot does not have contact information and relies only on its visual features. The corresponding network fails to distinguish between touch and tap on a box since the effects are equal (boxes do not roll and the robot's hand positions w.r.t. the object are almost equal). In the second one, we add contact information and show how the resulting network disambiguates these actions using the

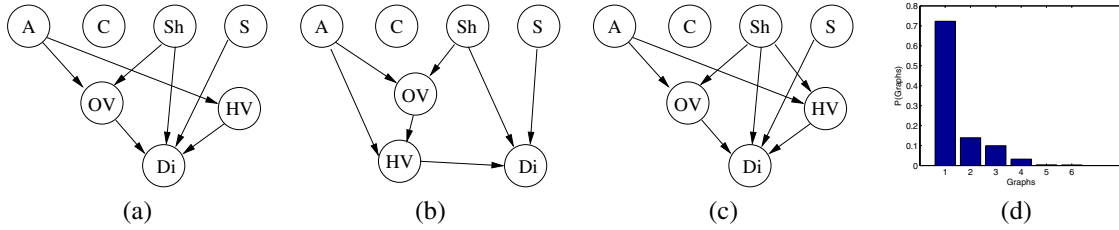


Fig. 5. Affordances network learned without contact information using the MCMC algorithm. (a-c) Three most likely graphs and (d) the distribution of the equivalent class graphs.

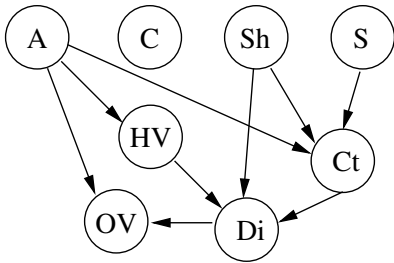


Fig. 6. Maximum likelihood graph using contact information.

new information. Note that tap actions produce only a brief contact with the object.

We used the MCMC algorithm with random initialization and BDeu priors to give uniform priors to different equivalence classes [20]. The MCMC stopped after 5000 samples, generated with a burn-in period of 500^3 .

Figure 5 shows the most likely networks and the distribution of graphs obtained with the MCMC algorithm without using contact. As shown in the histogram of Fig. 5(d), most of the probability mass is assigned to the network depicted in Fig. 5 (a). Although there is not a ground truth to compare the estimated networks, we see that color has been detected as irrelevant when performing any action. Shape and size are important for grasp and tap since they will have an impact on observed object velocity, contact and distance between the hand and the object. Hand velocity depends mainly on the action, since the robot performs the same movements independently of the success or failure of the intended action. The other two networks (Fig. 5 (b) and (c)) are slight variations of the previous one, but still capture the main properties of the world presented to the robot.

Figure 6 shows the maximum likelihood graph using contact information. The resulting network captures a similar structure and dependencies as in the previous case.

Figure 7 shows the weighted marginal likelihood of the estimated networks as the number of robot trials increases for both cases. Notice that the network with contact information is more complex. Therefore, a similar marginal likelihood indicate a better modeling of the data as expected due to the extra information.

Recall that the actual dependencies are encoded in the multinomial CPDs of each node given the graph. In order to

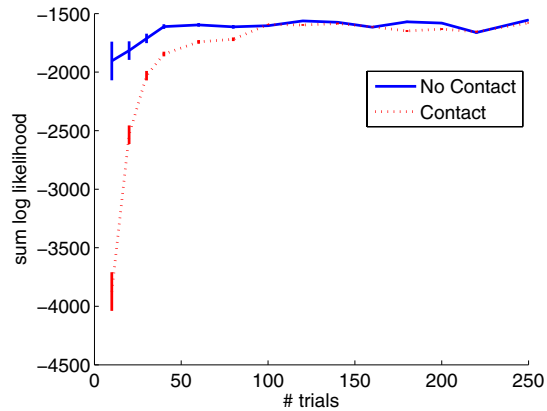


Fig. 7. Marginal likelihood of the data given the learned as the number of robot trials increases. The vertical bars show the variance of the likelihood.

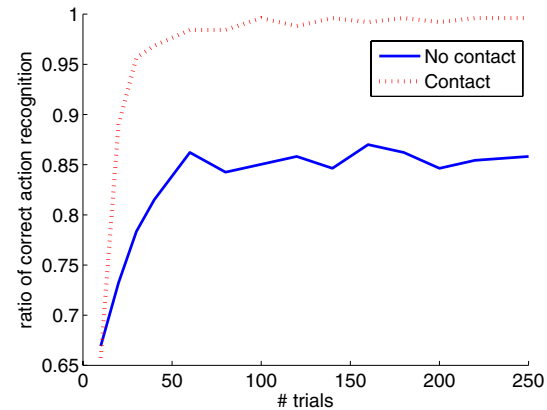


Fig. 8. Action recognition. Ratio of correct classifications for different number of trials.

evaluate the quality of the network we also used a *leave one out* validation technique to classify the robot action given the object features and the perceived effects. Figure 8 shows the ratio of correct classifications. Due to the lack of contact information, the robot cannot distinguish between touch and tap actions executed upon a box. The action recognition ratio converges to 0.85, where the errors are due to the unobservability between touch and tap. If the robot has this information available, the network is able to distinguish between them and this ratio achieves 0.98. Notice that when recognizing actions performed by other people,

³The implementation of the algorithms is based on the BNT toolbox for Matlab, <http://bnt.sourceforge.net/>.

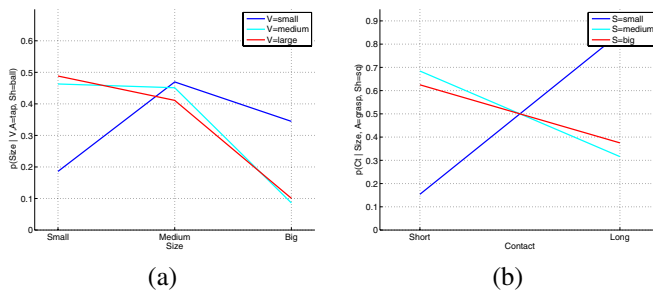


Fig. 9. Examples of CPD for the learned network: (a) shows $p(S | V = v_i, A = \text{tap}, Sh = \text{ball})$, the CPD of the shape of a ball given the action was a grasp for every possible value of the velocity. (b) shows $p(Ct | S = s_i, A = \text{grasp}, Sh = \text{sq})$, the CPD of contact given a grasp was performed on a box for every value of size.

contact information is usually not available. In this case, the recognition ratio for the network with contact information also degrades to 0.85.

To further validate that the network actually captures the correct dependencies, we show next some conditional probability distributions. Figure 9(a) shows the distribution of size after performing a tap on a ball for different velocities. According to it, small balls move faster than bigger ones which was the actual behavior during the experiments. The predicted contact of a grasp action for different sizes is present in Figure 9(b). It basically states that grasps succeed more often with small objects than with bigger ones.

In summary, we have shown that the proposed model allows to represent affordance knowledge and learn it through experience. This is done based on the motor and perceptual skills in a completely unsupervised manner. There is no notion of success or failure and the network may not be able to distinguish between non separable objects given the used descriptors. However, it constructs a plausible model of the behavior of the different objects under different actions that can readily be used for prediction and planning.

VI. CONCLUSIONS

This paper addresses the unsupervised learning and usage of affordances, i.e. the relations between actions, objects and effects. We have proposed a general model based on Bayesian networks and presented a particular implementation based on discrete random variables. The network does not assume prior knowledge on the relations between the objects and effects and implicitly performs feature selection during the learning phase.

Based on the previous framework, there are plenty of opportunities for future research. Biological systems develop many of their different skills in parallel. We are now investigating how to dynamically incorporate new robot capabilities (actions) or world knowledge in the learning algorithms. This is important to deal with higher dimensional spaces. Although the proposed model can directly learn through observation of other agents', it is necessary to develop mechanisms to update the knowledge sequentially and to deal with new actions or effects. Finally, active exploration strategies may also allow to reduce the number of trials

required to learn the affordances in more complex and realistic situations.

REFERENCES

- [1] J. J. Gibson, *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin, 1979.
- [2] S. Schaal, "Is imitation learning the route to humanoid robots," *Trends in Cognitive Sciences*, vol. 3(6), 1999.
- [3] L. Montesano, M. Lopes, A. Bernardino, and J. Santos-Victor, "Affordance, development and imitation," in *International Conference on Development and Learning*, 2007.
- [4] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [5] D. K. N. Friedman, L. Getoor and A. Pfeffer, "Learning probabilistic relational models," in *International Joint Conferences on Artificial Intelligence*, 1999.
- [6] A. Chemero, "An outline of a theory of affordances," *Ecological Psychology*, vol. 15, no. 2, pp. 181–195, 2003.
- [7] J. Konczak, H. Meeuwssen, and M. Cress, "Changing affordances in stair climbing: The perception of maximum climbability in young and older adults," *Journal of Experimental Psychology: Human Perception & Performance*, vol. 19, pp. 691–7, 1992.
- [8] E. Symes, R. Ellis, and M. Tucker, "Visual object affordances: Object orientation," *Acta Psychologica*, vol. 124, no. 2, pp. 238–255, 2007.
- [9] M. Turvey, K. Shockley, and C. Carello, "Affordance, proper function, and the physical basis of perceived heaviness," *Cognition*, vol. 73, 1999.
- [10] E. Oztop, N. Bradley, and M. Arbib, "Infant grasp learning: a computational model," *Experimental Brain Research*, vol. 158, pp. 480–503, 2004.
- [11] M. Lopes and J. Santos-Victor, "A developmental roadmap for learning by imitation in robots," *IEEE Transactions on Systems, Man, and Cybernetics - Part B: Cybernetics*, vol. 37, no. 2, April 2007.
- [12] A. Slocum, D. Downey, and R. Beer, "Further experiments in the evolution of minimally cognitive behavior: from perceiving affordances to selective attention," in *Conference on simulation of adaptive behavior*, Paris, France, 2000.
- [13] P. Fitzpatrick, G. Metta, L. Natale, S. Rao, and G. Sandini, "Learning about objects through action: Initial steps towards artificial cognition," in *IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, 2003.
- [14] H. Kozima, C. Nakagawa, and H. Yano, "Emergence of imitation mediated by objects," in *Second International Workshop on Epigenetic Robotics*, Edinburgh, Scotland, 2002.
- [15] I. Cos-Aguilera, L. Cañamero, and G. Hayes, "Using a SOFM to learn object affordances," in *Workshop of Physical Agents (WAF)*, Girona, Spain, 2004.
- [16] G. Fritz, L. Paletta, R. Breithaupt, E. Rome, and G. Dorffner, "Learning predictive features in affordance based robotic perception systems," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Beijing, China, 2006.
- [17] A. Stoytchev, "Behavior-grounded representation of tool affordances," in *International Conference on Robotics and Automation*, Barcelona, Spain, 2005.
- [18] J. Pearl, *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2000.
- [19] C. Huang and A. Darwiche, "Inference in belief networks: A procedural guide," *International Journal of Approximate Reasoning*, vol. 15, no. 3, pp. 225–263, 1996.
- [20] D. Heckerman, D. Geiger, and M. Chickering, "Learning bayesian networks: the combination of knowledge and statistical data," *Machine Learning*, 1995.
- [21] G. Cooper and E. Herskovits, "A bayesian method for the induction of probabilistic networks from data," *Machine Learning*, vol. 9, no. 4, pp. 309–347, 1992.
- [22] D. Madigan and J. York, "Bayesian graphical models for discrete data," *Intl. Statistical Review*, vol. 63, pp. 215–232, 1995.
- [23] D. Eaton and K. Murphy, "Exact bayesian structure learning from uncertain interventions," *AI & Statistics*, In Press.
- [24] D. Heckerman, "A tutorial on learning with bayesian networks," in *In M. Jordan, editor, Learning in graphical models*. MIT Press, 1998.
- [25] D. Pelleg and A. W. Moore, "X-means: Extending k-means with efficient estimation of the number of clusters," in *ICML '00*, San Francisco, CA, USA, 2000.