

Affordances as a Framework for Robot Control

Maya Çakmak Mehmet R. Doğar Emre Uğur Erol Şahin

Kovan Research Lab.

Dept. of Computer Engineering
Middle East Technical University, Turkey

Abstract

The concept of affordances, with its emphasis on the interaction between the organism and the environment, is highly relevant for robotics. In this paper, we present a new formalization of affordances that provides a framework for robot control, learning and planning. We argue that affordances, as relations within the robot-environment system, can be learned by the robot from its interactions with the environment. The interaction is represented as a nested triple of the form (*effect*, (*entity*, *behavior*)) indicating that the *behavior* applied in an environment perceived as the *entity*, would produce a perceivable *effect*. The robot accumulates such triples of raw sensory-motor data from different interactions and learns the affordance relations through the formation of *equivalence classes*. We present three studies that implement certain aspects of the formalism on a mobile robot moving in an environment that contains different types of objects. These studies respectively explore (1) the perceptual learning of affordances, (2) the development of goal-directed behaviors from a set of primitive ones through the learning of affordances, and (3) the use of learned affordance relations in planning.

1. Introduction

J.J. Gibson (Gibson, 1986) introduced the concept of *affordances* to refer to the action possibilities offered to the organism by its environment. For instance, a horizontal and rigid surface affords walk-ability, a small object below a certain weight affords throw-ability, for a human. He argued that affordances point both to the environment and the organism implying their complementarity. Although J.J. Gibson conceived the concept in his quest to develop a “theory of information pick-up” as a new theory of perception, it has influenced studies ranging from Human-Computer Interaction to Autonomous Robotics.

The concept of affordances has been a misty one since its conception (which may have contributed positively to its influence over a wide-range of fields). A number of formalizations have been proposed to clarify its meaning. To summarize briefly, Turvey (Turvey, 1992)

defined affordances as “dispositions” in the environment that get actualized with the interaction of the organism and the environment. Different from Turvey’s formalism, which attached affordances to the environment, Stoffregen (Stoffregen, 2003) and Chemero (Chemero, 2003) defined affordances as relations within the organism-environment system. Independent from these formalizations in Ecological Psychology, Steedman (Steedman, 2002) formalized affordances in Linguistics by providing an explicit link to action possibilities offered by the environment, and by proposing the use of the concept in planning.

The concept of affordances, with its implicit but central emphasis to the interactions between the organism and the environment, is highly relevant for *developmental/epigenetic robotics* as has already been noted (Lungarella et al., 2003). Developmental robotics treats affordances as a higher level concept, which a developing cognitive agent learns by interacting with its environment. There are studies that exploit how affordances reflect to learning (MacDorman, 2000), tool-use (Stoytchev, 2005), or decision-making (Cos-Aguilera et al., 2003). The studies that focus on learning mainly tackle two major aspects. In one aspect, the learning of consequences of an action in a given situation (Stoytchev, 2005) is referred to as affordance learning. In the other, studies focus on the learning of invariant properties of different environments that afford a certain action (MacDorman, 2000, Fritz et al., 2006). Studies in this latter group also relate these properties to the consequences of applying an action, but these consequences are in terms of internal values of the agent, rather than physical changes in the environment.

In (Fitzpatrick et al., 2003), learning of object affordances in a developmental framework is studied. The main vision set forth in this work is that a robot can learn what it can do with an object only by acting on it, “playing” with it, and observing the effects. After applying each action on different objects several times, the robot learns the roll-ability affordance of the objects, by observing the changes in the environment inflicted by the actions. In a recent study (Papudesi and Huber, 2006), *behavioral affordances* and *goals* are used for internally representing the state of the world within an artificial agent. For each action in its repertoire, the agent has *outcome*

predictors and *outcome indicators* that correspond to pre-conditions and effects of the action.

Despite the interest, the use of affordances in robotics is mostly confined to a source of inspiration. Moreover, a closer look to these studies reveals that their use are based on different and sometimes contradictory façades of this concept and that most studies cite only J.J. Gibson’s studies published in 70’s and 80’s. In the MACS project¹, we, as roboticists, are interested in how the concept of affordances can change our views towards the control of autonomous robots. Towards this end, we formalized the concept, outlined its implications towards robot control (Şahin et al., 2007) and started evaluating these implication on real robots.

2. Formalizing Affordances

In the existing formalisms proposed in Ecological Psychology, affordances are either placed in the environment as extended properties that are perceivable by the agent (Turvey, 1992), or, they are said to be a properties of the agent-environment system (Stoffregen, 2003, Chemero, 2003) (see (Şahin et al., 2007) for a complete review). Although helpful in partially clarifying this misty concept, these formalisms provide little guidance towards the use of the concept in robotics.

We propose that affordances are relations within the agent-environment system. Our formalization is based on relation instances of the form (*effect*, (*entity*, *behavior*)), meaning that there exists a potential to generate a certain *effect* when the *behavior* is applied on the *entity* by the agent. The *entity* represents the state of the environment as well as the state of the agent, as perceived by the agent. The *behavior* represents the physical embodiment of the interaction of the agent with the environment, and the *effect* is the result of such an interaction. For instance, the *lift-ability* affordance implicitly assumes that, when the *lift* behavior is applied on the entity *stone*, it produces the effect *lifted*, meaning that the stone’s position, as perceived by the agent, is elevated.

A single (*effect*, (*entity*, *behavior*)) relation instance is acquired through a single interaction with the environment. But this single instance does not constitute an affordance relation by itself, since it does not have any predictive ability over future interactions. Affordances should be relations with predictive abilities. This is achieved by building four types of equivalence classes.

Entity equivalence: The class of *entities* which support the generation of the same *effect* upon the application of a certain *behavior* is called an *entity equivalence class*. For instance, our robot can achieve the effect *lifted*, by applying the *lift* behavior on a *black-can*, or a *blue-can*. These relation instances can then be compacted by a mechanism that operates on the class to determine the perceptual invariants of the en-

tity equivalence class as:

$$(\textit{lifted}, (< *-\textit{can}>, \textit{lift}))$$

where $< *-\textit{can}>$ denotes the derived invariants of the entity equivalence class.

In this particular example, $< *-\textit{can}>$ means “cans of any color” that can be *lifted* upon the application of *lift*. Such invariants enable the robot to predict the *effect* of the *lift* behavior applied on a novel object, like a *green-can*. Such a capability offers great flexibility to a robot. When in need, the robot can search for entities with a desired affordance.

Behavior equivalence: Maintaining a fair treatment of the action aspect of affordances, the equivalence concept can be generalized to the *behavior*. For instance, our robot can lift a can using its *lift-with-right-hand* behavior. If the same effect can be achieved with its *lift-with-left-hand* behavior, then these two behaviors are said to be *behaviorally equivalent*. This relation can be represented as:

$$(\textit{lifted}, (< *-\textit{can}>, < \textit{lift-with-*hand}>))$$

where $< \textit{lift-with-*hand}>$ denotes the invariants of the behavior equivalence class.

Similar to the *entity equivalence*, the use of *behavioral equivalence* will bring in a flexibility for the agent. For instance, a humanoid robot which lifted a can with one of its arms, loses its ability to lift another can. However, through *behavioral equivalence*, it can immediately have a “change of plan” and accomplish lifting using its other hand.

Affordance equivalence: Taking the discussion one step further, we come to the concept of *affordance equivalence*. Affordances like traversability, are obtainable by “walking across a road” or “swimming across a river” as:

$$(\textit{traversed}, \left\{ \begin{array}{l} (< \textit{road}>, < \textit{walk}>) \\ (< \textit{river}>, < \textit{swim}>) \end{array} \right\})$$

That is, a desired effect can be accomplished through different (*entity*, *behavior*) pairs.

Effect equivalence: The notions of entity, behavior and affordance equivalence implicitly rely on the assumption that the agent, somehow, has the notion of *effect equivalence*. For instance, applying the *lift* behavior on a *blue-can* would generate the effect of “a blue blob rising in view”. If the robot applies the same behavior to a *red-can*, then the generated effect will be “a red blob rising in view”. If the two effects are considered to be equivalent, the two instances obtained from these experiments can be joined, to have the common effect of “blob rising in view”.

Effect equivalence is closely related to the goals of the agent. The effects generated in two different interactions are considered equivalent if they achieve the same goal. For instance, given the particular goal of

¹URL: <http://www.macs-eu.org>, (FP6-IST-004381).

having a blue-blob rise in view, the effects of lifting a blue-can or a red-can cannot be considered equivalent.

Finally, based on the discussion presented above, we propose a formal definition of an affordance as follows: *An affordance is an acquired relation between a certain <effect> and a certain <(entity, behavior)> tuple such that when the agent applies a (entity, behavior) within <(entity, behavior)>, an effect within <effect> is generated.*

3. Experiments towards Affordance-based Robot Control

The proposed formalism lays a framework over which the concept of affordances can be utilized in robot control. In this section, we present three experiments that implement different aspects of the formalism and discuss their implications.

The first experiment explores the formation of *entity equivalence classes*. The basic idea in the experiment stems from E.J. Gibson’s studies on perceptual learning. She suggests that learning of affordances is “discovering *distinctive* features and *invariant* properties of things and events” (Gibson, 2000), “discovering the information that specifies an affordance” (Gibson, 2003). She describes this method as “narrowing down from a vast manifold of (perceptual) information to the minimal and optimal information that specifies the affordance of an event, object or layout” (Gibson, 2003). In our framework, finding relevant features and invariant properties that specify whether a behavior will succeed or not in an environment, corresponds to building *entity equivalence classes*. In the experiment, the robot interacts with the environment by executing its behaviors and checks whether the execution of the behavior succeeds or not. Based on these experiences, it determines the features in the environment that are useful in predicting its behaviors’ success. In other words, the robot learns to perceive the environment in terms of the features that predict whether its behaviors will succeed or not. The idea is similar to *function-based-object-recognition*, however in this study the features that specify the functionality of the entities in the environment are learned by the robot through interaction rather than hard-coded into the robot by an expert.

The second experiment extends the first by adding the formation of *effect equivalence classes*. The robot achieves this by randomly performing unintentional primitive behaviors and discovering the changes it can consistently create in the environment. These changes are then associated with the executed behaviors and the situation in which the behavior is executed. This corresponds to linking *effect equivalence classes* with *behavior* and *entity equivalence classes*, which is the formation of affordance relations. Using these relations the robot can execute its primitive behaviors purposefully, to achieve a goal. This approach can also be related

to E.J. Gibson’s discussion on child development and affordances. She points out that babies use exploratory activities, such as mouthing, reaching, shaking to gain the perceptual data needed to learn the affordances in the environment, and that these activities bring about “information about changes in the world that the action produces” (Gibson, 2000). As development proceeds, exploratory activities become performatory and executed with a goal. Likewise, our robot develops purposeful goal-directed behaviors from unintentional primitive behaviors.

In the third experiment, the learned affordance relations are used in planning. Once the robot has acquired a representation of its capabilities in the form of affordance relations it can make inferences about the effect of action sequences and their practicability. This corresponds to using learned affordance relations as plan operators in the classical planning problem. The *<entity>* and *<behavior>* components in the learned relations, can be considered to correspond to the pre-condition and action components in classical planning systems. This link between affordances and the planning problem was noted earlier (Amant, 1999, Steedman, 2002), however, these studies assumed the existence of symbols. Opposing to this, we show that planning can be performed using learned relations that are mappings from low-level perceptual features to self-acquired effect categories.

3.1 Experimental framework

In our experiments we investigate the interactions of a wheeled robot moving in an environment cluttered with different objects. The environment contains four types of simple objects: rectangular boxes () , spherical objects () and cylindrical objects, either in upright position () or lying on the ground () . When contacted by the robot, these objects either roll away or block the robot’s motion.

The robotic platform used in this study is Kurt3D, a medium-sized, differential drive mobile robot, and its physics-based simulator MACSim whose sensor and actuator models are calibrated against their real counterparts.

In all experiments the robot has a repertoire of primitive behaviors each generating a certain displacement or rotation, unless the motion is obstructed by an obstacle. The robot interacts with the environment by performing one of its primitive behaviors and perceiving the environment both before and after the execution of each behavior.

The robot perceives its environment through its 3D scanner, which is based on a SICK LMS 200 2D scanner, rotated vertically with an RC-servo motor. It uses the range images from the scanner to extract a set of features which consists the robot’s perception of the environment. The feature set also contains two features obtained from its encoders. To obtain the scanner

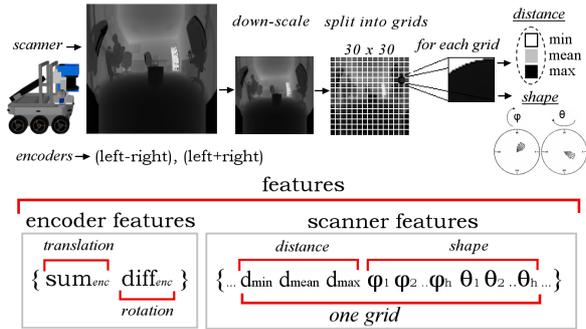


Figure 1: Phases of perception and content of the feature vector. Distance and shape features are extracted from the scanner range image. Also two displacement values, translation and orientation, are extracted from the encoders.

features, the range image is down-scaled to reduce the noise and it is split into uniform grids. For each grid cell, a number of distance and shape related features are extracted. The distance related features are the closest, furthest, and mean distances within the grid cell. The shape related features are computed from the normal vectors in the grid cell. The direction of each normal vector is represented using two angle channels φ and θ , in latitude and longitude respectively and two angular histograms are computed. The frequency values of these histograms are used as the shape related features (Figure 1).

3.2 Perceptual Learning

In the first experiment we investigate how the perceptual features that specify an affordance can be learned by the robot from interactions with the environment. Specifically, we study how a mobile robot can learn to perceive the traversability affordance in a room filled with different objects. We define traversability as “the ability to pass or move over, along, or through”. Hence, the environment is said to be traversable in a certain direction if the robot moving in that direction is not enforced to stop as a result of contact with an obstacle. Here we present an overview of the results obtained in the study. For more details on the results, please see (Uğur et al., 2007b).

The environment typically contains one or more objects, with arbitrary size, orientation and placement, in front of the robot. The process consists of three phases: an interaction phase, during which the robot accumulates a number of relation instances, a learning phase in which entity equivalence classes are learned from these instances, and an execution phase for testing. In order to collect instances, the robot perceives the initial environment and executes one of the seven pre-coded movement behaviors, ranging from *turn-sharp-right* to *turn-sharp-left*. It records whether it was able to successfully traverse or not, based on the change in its en-

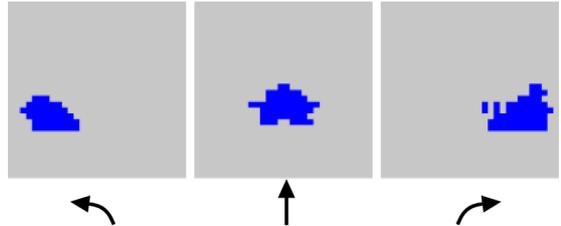


Figure 2: The relevant grids in the range image for three of the actions. A grid is marked as relevant if any of the features extracted from it were learned to be relevant.

coder values. The robot collects the relation instances, where the *entity* is the initially perceived feature vector, the *behavior* is the index of the executed behavior(1-7), and the *effect* is 1 or 0 indicating success or failure.

In the learning phase the robot first selects the relevant perceptual features using the ReliefF algorithm (Kira and Rendell, 1992). Using these relevant features, for each behavior, a Support Vector Machine (SVM) classifier (Vapnik, 1995) is trained, to learn the mapping from feature space to the effects (success/fail). After learning, the robot can predict whether the environment affords traversability for a given behavior, with around 95% success. As a result of learning a *perceptual economy* is achieved. Our analysis show that only 1% of the raw feature vector is relevant for perceiving traversability and that these relevant features are grouped on the range image with respect to the direction of the movement as shown in Figure 2.

In this experiment, entity equivalence classes are discovered by the trained classifiers whereas behavior and effect equivalences are assumed to be pre-coded.

In a different setup, the trained robot is tested in an environment inspired from Warren and Whang’s study (Warren and Whang, 1987) on walking through apertures. In this study, the subjects are faced with an aperture of varying width, and they are asked whether the aperture affords walking through or not. The results showed that the *aperture-to-shoulder-width ratio* is a body-scaled constant for this affordance, and that a *critical point* existed for the subject’s decision. In a similar vein to these experiments, we placed two box-shaped objects in front of the robot, and tested the robot’s predictions of traversability affordance for apertures with different widths. As shown in Figure 3, the robot is able to correctly perceive the affordances of pass-through-able apertures, where the *critical passable width* is clearly related to the robot’s width.

To investigate the generalization capability of the perceptual learning approach, in another setup, we restrict the types of objects in the interaction environment and perform testing with novel objects. The robot interacts only with lying cylinders, which may or may not afford traversability to the robot depending on

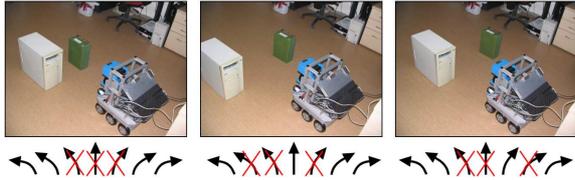


Figure 3: Three experiments for evaluating pass-through-ability for the robot. In (a) the width of the aperture is too narrow whereas in (b) it is wide enough to support the pass-through-ability. (c) shows the case where the aperture is slightly towards the right of the robot.

their relative orientation. After learning, the robot is tested with spheres, boxes and upright cylinders, objects that it has not interacted with before. Yet the robot is able to predict that boxes and upright cylinders were non-traversable (both 100% success), and that spheres are traversible (83% success). We claim that, in this study, the robot learns “general relations” that pertain to its physical interaction with the environment and that these relations are used in making successful predictions about the traversability of novel objects.

The learning of affordances in these experiments typically requires a large set of training data obtained from the interactions of the robot with its environment. Therefore, the learning process is not only time-consuming and costly but it is also risky since some of the interactions may inflict damage on the robot. To overcome this issue, in a recent work (Uğur et al., 2007a), we extended this learning system with two ideas. First, learning is conducted as an on-line process rather than a batch process. It is clear that a developing agent must be able to update its knowledge about its interaction with the environment continuously. Second, we used a curiosity measure to assess whether a given interaction opportunity is worth exploring or not so that the robot can select the most interesting interactions in the environment. Hence, the developing agent does not perform unnecessary interactions when it is confident that it will not bring about new information. The robot goes ahead with the interaction only when the classifier has a low confidence about its effect. This approach results in a substantial speed-up for the learning system (see (Uğur et al., 2007a) for more details on the results).

3.3 Development of goal-directed behaviors

In this experiment we used the concept of affordances in making the robot learn about its own capabilities. As in E.J. Gibson’s account of behavioral development in infants, we investigate the question of how goal-directed behaviors can be achieved starting from unintentional primitive behaviors.

Differing from the previous, in this experiment the interaction environment contains a single object and

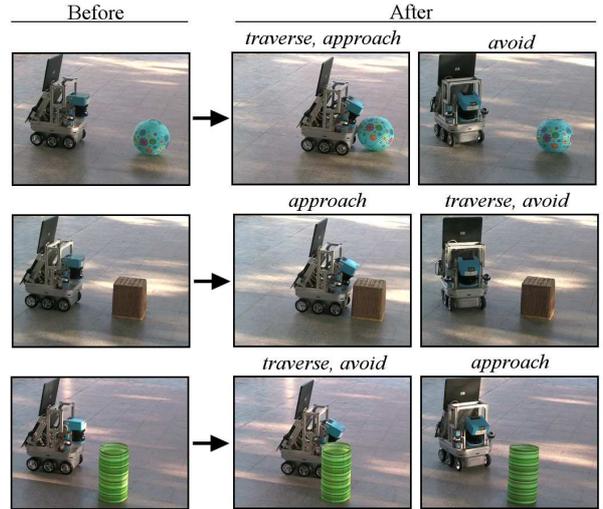


Figure 4: Three cases in which different goal-directed behaviors (*traverse*, *avoid*, *approach*) make use of different primitive behaviors (move-forward, turn-right, turn-left).

the robot has three primitive behaviors: *move-forward*, *turn-left*, *turn-right*. Learning differs in that *effects* are not represented as success/fail values, but instead, the actual perceivable change created by the behavior is discovered by the robot as the *effect*.

In the interaction phase, the robot perceives the environment before and after executing one of its primitive behaviors, to collect *relation instances*. The initial feature vector is the *entity* and the vectorial difference between the final and initial features is the *effect*.

Learning consists of three steps. First, within the set of *relation instances* of a behavior, similar effects are grouped together to get a more general description of the effects that the particular behavior can create. This is achieved through a k-means clustering of the effect instances of that primitive behavior and corresponds to obtaining the *effect equivalence classes* in the formalism. After clustering, each *effect class* is assigned an *effect-id* and the *effect prototype* of the class is calculated. Next, the relevant perceptual features are selected using the ReliefF algorithm and then an SVM for each behavior is trained using these relevant features, as to learn the mapping from the feature space to the effect-ids.

Goal-directed behaviors are achieved using the learned relations as follows. Given the perception of the environment, the trained classifiers can predict the effect class that the behavior will produce. By comparing the *effect prototype* of the predicted classes for each behavior, the robot can select the behavior that will produce the most useful effect in achieving its goal. We specify the goal as a criteria according to which all effect classes are sorted. The robot executes the behavior for which the predicted effect has higher priority according to the goal.

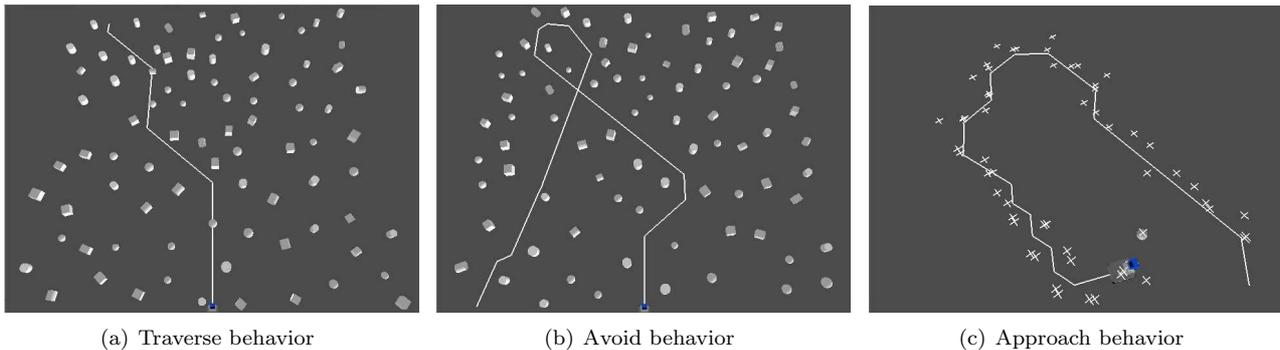


Figure 5: Three different behaviors achieved using the same primitive behaviors. In (a), the robot wanders around perceiving the traversability affordance in the environment. In (b), the robot displays typical obstacle-avoidance behavior, where it avoids all the objects. In (c), an example path where the robot follows an object using its *approach* behavior is shown. The plus signs mark the places that objects appear.

Three different goal-directed behaviors (*traverse*, *avoid* and *approach*) are obtained in this way. The first is a *traverse* behavior, which is achieved by giving higher priority to the effect classes whose prototypes have a greater forward-displacement. We achieved the *avoid* behavior by specifying the desired effect as having a high increase in the mean distance in the middle portion of the range image. This results in a behavior where the robot avoids contact with any object by turning away whenever something appears on its front. When the desired effect is changed to a high decrease in the mean distance, an *approach* behavior emerges. The robot moves forward towards an object on its front, and turns towards an object on its side, to obtain the desired decrease. Figure 4 shows how the goal-directed behaviors react in different environments.

We have also tested the *traverse* and *avoid* behaviors by placing the robot in an environment randomly filled with multiple objects, and the *approach* behavior by making the robot follow an object. In the *traverse* and *avoid* cases, the robot successfully explored the environment. For the *traverse* behavior, the robot also used the traversability affordance of the objects by rolling away the traversable objects on its way, and avoiding the non-traversable ones. Examples of these trials can be seen in Figure 5. For more details please see (Doğar et al., 2007).

In the described experiment, goal directed behaviors are obtained by sequencing primitive behaviors based on an action selection mechanism that accounts for the goal. Primitive behaviors can also be blended to obtain new behaviors that are better in terms of achieving the goal. This corresponds to using primitive behaviors simultaneously. More efficient behaviors are obtained with weighted sums of motor control parameters of the pre-coded primitive behaviors. The weight given to each primitive behavior is proportional to the similarity of its genuine effect to the desired effect specified by the goal. This idea has been studied in detail

in (Doğar, 2007).

3.4 Planning

In the second experiment, learned affordance relations were used in predicting the effects of primitive behaviors, so that the appropriate behavior could be selected in different situations to obtain an overall goal-directed behavior. These predictions can also be used to estimate the future entities that the robot will perceive after the execution of different behaviors, simply by adding the prototype of the predicted effect to the currently perceived entity. It is then possible to predict the effects of behaviors over the estimated future environments, again using the learned relations. The robot can estimate the total effect that a sequence of behaviors will create and it can predict the entity that it will perceive after the execution of the sequence. This constitutes the basic idea for using learned affordance relations in planning sequences of behaviors that lead to a desired goal. Note that, the goal can either be specified as a total effect to be obtained or a desired future state. The approach of using forward chaining in affordance-based planning was proposed by Steedman (Steedman, 2002).

We have tested the described method in the framework presented in the previous section. The robot starts by perceiving the present entity, and predicts the effects that each of its primitive behaviors (*move-forward*, *turn-left*, *turn-right*) will create. It estimates the three future entities and proceeds by predicting the effects of behaviors on those future entities and estimating the next entities. This process can be viewed as the breadth-first construction of a plan tree where the branching factor is the number of possible primitive behaviors. Meanwhile, the robot tests whether the goal is satisfied by the entities in the attained states or by the total effect of the sequence of behaviors that leads to those states. Planning stops when a sequence satisfies the goal.

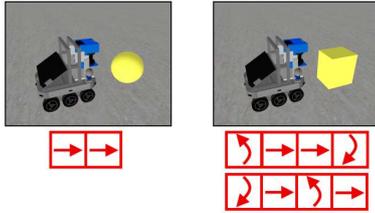


Figure 6: Two cases in which the robot generates different plans, given the goal of achieving a total translational displacement of a certain amount.

In the example presented in Figure 6 the robot is initially faced with two different situations. Its goal is specified as obtaining a positive change in the translation-related encoder feature, corresponding to a forward displacement of approximately 1 meter. In the case where the robot is faced with a spherical object, the prediction for the *move-forward* is an effect class with a high forward translation. The effect prototype also reflects the change in the position of the spherical object which either rolls away from the robot’s path or remains on its front. The estimated future entity is therefore one in which a similar effect prototype will be predicted for the *move-forward* behavior. Among other paths in planning tree, “three times *move-forward*” is the first that sums up to the desired change in the encoder feature. In the case where the robot is faced with a box object, the prediction for *move-forward* is an effect class with a low forward translation. The predicted effects for turning actions have no translation at all, however they imply the change in the position of the box in robot’s perceptive field. For instance, the estimated future entity after two *turn-lefts*, is one in which an object appears on the right of the range image and a *move-forward* now predicts an effect with a good forward translation. The obtained plan therefore consists of two turns and three forward moves.

Although there is a clear resemblance between the triples in the formalism and the plan operators used in classical planning, the basic difference in the operator structure can be noted. In a classical planning system, such as STRIPS (Fikes and Nilsson, 1971), the plan operators are mappings from a bounded set of pre-conditions to a bounded set of effects, indexed with an action. Both pre-conditions and effects are specified using a set of propositions or first-order predicates which are assumed to be detectable by the agent. In our approach the plan operators are affordance relations which are mappings from low-level perceptual features to an effect category identified by an id and a prototype. The relation is expressed with the decision surface of the trained classifier. Although the effect of the operator is explicit as in the classical operator, the pre-condition is inherent in the classifier that maps an entity to one of the effects. Despite these differences the

nature of the planning problem in our approach is not different from classical planning. The major strength of the described approach is that planning and learning are achieved within a common framework. The effect categories are formed by the robot with unsupervised learning and are based on its own perception of the consequences of performed actions. The trained classifiers whose predictions are used in planning, are based on the robot’s own interaction experiences. The acquisition of affordance relations can therefore be considered as a symbol formation and grounding process, that serves planning. A more detailed account of our planning approach and examples of different plans generated by the robot for different goals are presented in (Çakmak, 2007).

4. Conclusion

We argue that the concept of affordances provides a general framework for Epigenetic Robotics. To this end, we presented a new formalization of the concept that we have developed for robot control. Our formalization is based on relations between *equivalence classes*, that are formed using the interaction experiences of the robot with the environment. We proposed a new formalism, instead of using prior formalisms, because they lack certain aspects that are critical in building a robotic system that use affordances. Most of the prior formalisms (Turvey, 1992, Stoffregen, 2003, Chemero, 2003) are interpretations of the concept from a psychological-philosophical perspective, and they do not propose -or discuss- the concept from the perspective of building agents that use affordances. Steedman’s formalism (Steedman, 2002) takes a computational perspective, but it skips the perceptual aspect of affordances.

We presented three studies towards the use of the formalization on robots. In the first study, the robot learned the perceptual invariants of the environment that were required for the actualization of an affordance. Through learning, the robot achieved *perceptual economy*, using only 1% of the perceptual feature vector, and to *directly perceive* (that is, without going through a modelling of the environment) the affordances available in its environment. In the second study, we showed that starting from a number of primitive and exploratory behaviors, the robot can successfully develop goal-directed behaviors. Finally, in the third study, we have shown that the affordance relations, learned by the robot, can be used for planning.

These studies have provided preliminary results towards the implications put forward by the formalism and need to be extended. In particular, the formation of behavioral equivalence classes, as well as the concurrent formation of multiple equivalence classes need to be studied. We have shown that the concept provides a good framework for developing a symbol system for planning and we believe that this implies a potential

for its use in communication among robots with similar capabilities as well as for human-robot interaction. Another challenge ahead of us is the incorporation of the different ideas studied separately, into one robot control architecture.

This paper presents a framework and several studies that used this framework in different problems. We aimed at showing the variety of aspects that the formalism applies to, by including as many experiments as possible, which in turn limited the amount of detail that we could accommodate within this paper. We encourage the reader to refer to related articles for further technical detail and experimental results for these experiments.

References

- Amant, R. (1999). Planning and user interface affordances. In *Intelligent User Interfaces*, pages 135–142.
- Çakmak, M. (2007). Robot planning based on learned affordances. M.Sc. Thesis, Middle East Technical University, Turkey.
- Chemero, A. (2003). An outline of a theory of affordances. *Ecological Psychology*, 15(2):181–195.
- Cos-Aguilera, I., Canamero, L., and Hayes, G. (2003). Motivation-driven learning of object affordances: first experiments using a simulated Khepera robot. In *Proc. of the 9. Int. Conf. in Cognitive Modelling*.
- Şahin, E., Çakmak, M., Doğar, M. R., Uğur, E., and Ücoluk, G. (2007). To afford or not to afford: A new formalization of affordances towards affordance-based robot control. *Adaptive Behavior*. (in press).
- Doğar, M. (2007). Using learned affordances for robotic behavior development. M.Sc. Thesis, Middle East Technical University, Turkey.
- Doğar, M. R., Çakmak, M., Uğur, E., and Şahin, E. (2007). From primitive actions to goal-directed behavior using affordances. In *Proc. of IROS'07 (to appear)*.
- Fikes, R. and Nilsson, N. (1971). Strips: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3/4):189–208.
- Fitzpatrick, P., Metta, G., Natale, L., Rao, A., and Sandini, G. (2003). Learning about objects through action -initial steps towards artificial cognition. In *Proc. of ICRA '03*, pages 3140–3145.
- Fritz, G., Paletta, L., Kumar, M., Dorffner, G., Breithaupt, R., and Rome, E. (2006). Visual learning of affordance based cues. In *Proc. of the 9. Int. Conf. on SAB*, LNAI. Volume 4095., pages 52–64.
- Gibson, E. (2000). Perceptual learning in development: Some basic concepts. *Ecological Psy.*, 12(4):295–302.
- Gibson, E. (2003). The world is so full of a number of things: on specification and perceptual learning. *Ecological Psy.*, 15(4):283–288.
- Gibson, J. (1986). *The Ecological Approach to Visual Perception*. Lawrence Erlbaum Associates. Originally published in 1979.
- Kira, K. and Rendell, L. A. (1992). A practical approach to feature selection. In *ML92: Proc. of the Ninth international workshop on Machine learning*, pages 249–256. Morgan Kaufmann Publishers Inc.
- Lungarella, M., Metta, G., Pfeifer, R., and Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15(4):151–190.
- MacDorman, K. (2000). Responding to affordances: Learning and projecting a sensorimotor mapping. In *Proc. of ICRA '00*, pages 3253–3259.
- Papudesi, V. N. and Huber, M. (2006). Learning behaviorally grounded state representations for reinforcement learning agents. In *Proc. of EpiRob'06*.
- Steedman, M. (2002). Plans, affordances, and combinatorial grammar. *Linguistics and Philosophy*, 25.
- Stoffregen, T. (2003). Affordances as properties of the animal environment system. *Ecological Psychology*, 15(2):115–134.
- Stoytchev, A. (2005). Behavior-grounded representation of tool affordances. In *Proc. of ICRA '05*.
- Turvey, M. (1992). Affordances and prospective control: an outline of the ontology. *Ecological Psychology*, 4(3):173–187.
- Uğur, E., Doğar, M. R., Çakmak, M., and Şahin, E. (2007a). Curiosity-driven learning of traversability affordance on a mobile robot. In *Proc. of ICDL'07*.
- Uğur, E., Doğar, M. R., Çakmak, M., and Şahin, E. (2007b). The learning and use of traversability affordance using range images on a mobile robot. In *Proc. of ICRA '07*.
- Vapnik, V. N. (1995). *The Nature of Statistical Learning Theory*. Springer, New York, NY, USA.
- Warren, W. and Whang, S. (1987). Visual guidance of walking through apertures: body-scaled information for affordances. *Journal of Experimental Psychology*, 13(3):371–383.