Learning to act on objects

Lorenzo Natale*, Sajit Rao*, and Giulio Sandini

[1] * indicates joint first authors
[2] LIRA Lab, DIST. Univ of Genova, Italy Viale Francesco Causa, 13
Genova 16145, Italy
{sajit, nat, sandini}@dist.unige.it

Abstract. In biological systems vision is always in the context of a particular body and tightly coupled to action. Therefore it is natural to consider visuo-motor methods (rather than vision alone) for learning about objects in the world. Indeed, initially it may be necessary to act on something to learn that it is an object! Learning to act involves not only learning the visual consequences of performing a motor action, but also the other direction, i.e. using the learned mapping to determine which motor action will bring about a desired visual condition. Learning to act may be an important precursor to "event-interpretation", even when the events are object-object events that don't involve the motor system in any way.
In this paper we show how a humanoid robot uses its arm to try some simple pushing actions on an object, while using vision and proprioception to learn the effects of its actions. We show how the robot learns a mapping between the initial position of its arm and the direction the object moves in when pushed, and then how this learned mapping is used to successfully position the arm to push/pull the target object in a desired direction.

1 Introduction

All biological systems are embodied systems, and an important way they have for recognizing and differentiating between objects in the environment is by simply acting on them. Only repeated interactions (play!) with objects can reveal how they move when pushed (eg sliding vs rolling), how the size of the object correlates with how much force is required to move it, e.t.c. In a discovery mode, the visual system learns about the consequences of motor acts in terms of such features, and in planning mode the mapping may be inverted to select the motor act that causes a particular change. These two modes of learning the consequences of a motor act, and selecting a motor act to achieve a certain result, are obviously intertwined, and together are what we mean by "learning to act".

Learning to act is important not only to guide motor behavior but may also be a necessary step for event-interpretation in general, even if the motor system is not involved in any way. For instance, by the age of 6 months children can predict that in a collision with a stationery object, the size of a moving object is related to how far the stationary object moves [1]. This is just one of several things that children appear to learn from experience about their physical environment [2] [3]. What is the source of this knowledge? and how can we build systems that learn to interpret events in the physical world? Computer vision approaches to "event-interpretation" have naturally tried to solve this problem in the domain of vision alone. However given that vision does not exist independently of other modalities in biological systems, and knowledge about the world is acquired incrementaly in a developmental process we are taking a somewhat different approach. We assume that *it may be necessary to learn to act on objects first before we can learn to interpret more complicated events involving object-object interactions*. One source of evidence in support for this approach comes from the body of work about mirror neurons [4]. These are neurons in motor area F5 of the rhesus monkey that fire when the monkey performs a particular goal-directed action, but which also fire if it just sees another agent perform a similar action. While the mechanisms of this mapping are still far from clear the fact that the events are mapped to the monkey's existing motor repertoire gives a strong hint that the ability to visually interpret the motor-goal or behavioral/purpose of the action may be helped by the monkey's ability to perform that action (and achieve a similar motor goal) itself.

The focus of this paper, therefore is on *learning to act on objects*, not only because in itself it's a vital skill to understand the consequences of actions, and plan future actions, but also because it could be a necessary precursor to event-interpretation of other object-object interactions.

2 Learning the effect of pushing/pulling actions

We show how a humanoid robot [5] that has already learned to saccade, and reach towards points in space with its arm, now pushes/pulls an object around in front of it, and learns the effect of its actions on the object, and thereafter uses this knowledge to drive motor-planning.

It is important to note that by "effect" we mean not only the effect on the object, but also the effect on the robot - the force felt by the robot, or the amount it had to move its head to continue tracking, for example. In this initial experiment we consider only one effect - the direction the object moves as a result of the action. There are naturally many other effects that one could also pay attention to: how far the object moves, how long the object continued moving after the initial touch, e.t.c. These are on the agenda for continuing experiments. However, in the experiment described here the robot attends only to the instantaneous direction of motion of the target just after it has been pushed/pulled by the robot. The goal of the experiment is to learn the instantaneous direction of motion of the target object for each of several different approach motions of the hand from different directions. This learned knowledge is then later used by the robot to select the appropriate motor action to move an object in a desired direction.

3 Description of the experiment

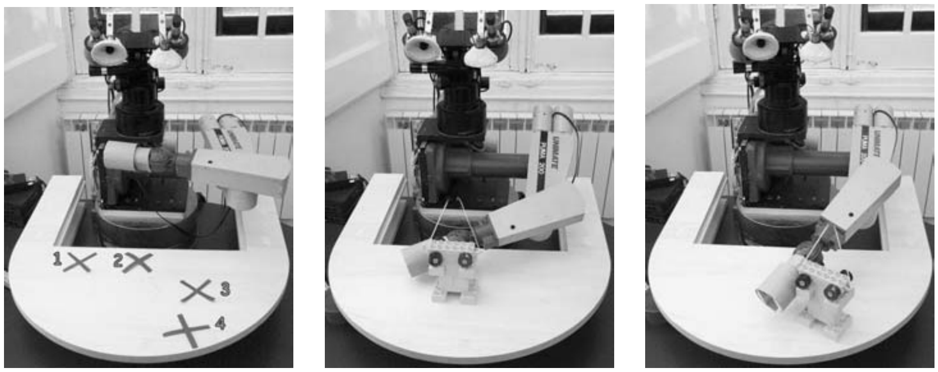

(a) initial arm-positions for target approach

(b) At the beginning of a pushing movement

(c) At the end of the movement

Fig. 1. The experimental setup

Figure 1 A Shows the experimental setup. The humanoid robot "Babybot" has a 5 DOF head, and a 6 DOF arm, and 2 cameras whose cartesian images are mapped to a log-polar format [6]. The robot also has a force sensitive wrist, and a simple piece of metal for a hand. The target is placed directly in front of the robot on the play-table. The robot starts from any of four different starting positions (shown in the figure) at the beginning of a trial run.

3.1 A single trial

In a typical trial run the robot continuously tracks the target while reaching for it. The target (even if it is moving) is thus ideally always centered on the fovea, while the moving hand is tracked in peripheral vision. Figure 1 (B) shows the arm at one of its initial positions and (C) shows the end of the trial with the target having been pushed to one side.

The moment of impact - when the hand first touches the object is an important event and its localization in time is critical. A sharp increase in the magnitude of retinal target position (caused by the instantaneous error in tracking) is used to localize the instant, and the direction of the displacement vector is extracted[3]

[3] Another source of information that also carries information about the moment of impact are the force values - a sharp discontinuity in the force profiles marks the

After the initial impact the system continues to try to reach for the centroid of the target and therefore ends up smoothly pushing the target in a particular direction. This continues until it either loses track of the target, which may fall off the table for example, or go outside the workspace. During each such trial run, the time evolution of several state variables are continuously monitored:

1. Vision: Position of the hand in Retinal coordinates - extracted from color segmentation of the hand.
2. Vision: Position of the target object in Retinal coordinates - extracted from color segmentation of the object.
3. Proprioception: 3 Joint coordinates of the arm (we fix the wrist for this experiment..thus eliminating 3 other degrees of freedom)
4. Proprioception: 5 Joint coordinates of the head
5. Proprioception: 3 Force components $[F_x F_y F_z]$ at the wrist.
6. Proprioception: 3 Torque components $[T_x T_y T_z]$ at the wrist

For the purpose of this experiment however we extract only two instantaneous values from this wealth of available data: one is the initial joint position of the arm (only the initial position, not the entire trajectory!), and the other is the instantaneous direction of target displacement at the moment of impact.

3.2 The target for learning

The goal of this experiment is to learn the effect of a set of simple pushing/pulling actions from different directions on a toy object. As we mentioned earlier in 1 there are many effects, both on the object and the robot, that could be attended to. But here we focus on only one effect, namely the direction of motion of the target. This is an effect to learn because, as we show, it can be used in motor-planning to move an object in a goal-directed mode. The target for learning (given a fixed target position directly in front of the robot) is a mapping from the initial position of the hand to the direction of target motion. Note that the initial hand-position uniquely determines the trajectory to the target. This trajectory could be different in different parts of the workspace, and is dependent on the kind of control used, (e.g. equilibrium point control) to generate the dynamics. However, because it is unique given the initial position of the hand and the (fixed) end position of the target, there is no need to remember the entire trajectory - the initial hand-position is sufficient.

The target for learning therefore is a mapping from initial hand position to direction of target movement. So, Associated with each initial hand position is a direction map (a circular histogram) that summarizes the directions that the target moved in when approached from that position. After each trial the appropriate direction map is updated with the target motion for that particular trial.

moment of impact. This information is not used at the moment but could be used to make the localization more robust, when the target is being obscured by the hand for instance.

Why map initial arm-position to target motion direction rather than say angle of approach of the hand at the moment of impact? The angle of approach of the hand would certainly correlate well with the direction of motion of the target. The reason we prefer the arm position instead is that the association lets us easily look up the answer to the inverse problem of motor planning, namely given a desired direction of motion of the target we can just lookup the position(s) where the arm should be initially positioned.

4 Results

Figure 2 shows the four direction maps learned, one for each initial arm position considered. The maps plot the frequency with which the target moved in a particular direction at the moment of impact. Therefore longer radial lines in the plot point towards the most common direction of movement. As we can see, the maps are sharply tuned towards a dominant direction.

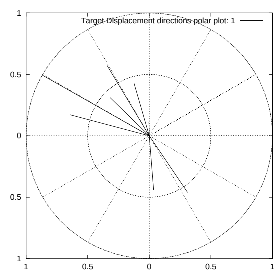

(a) Map for hand-position 1

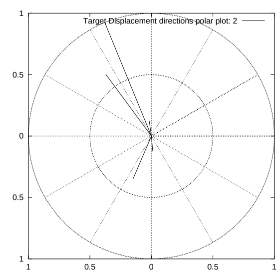

(b) Map for hand-position 2

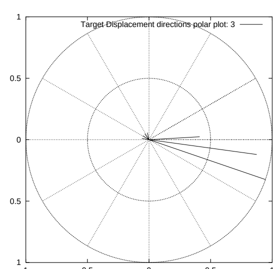

(c) Map for hand-position 3

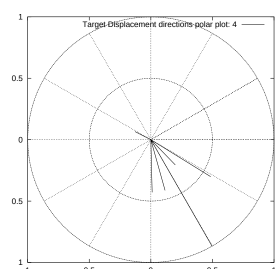

(d) Map for hand-position 4

Fig. 2. The learned target-motion direction maps, one for each initial hand-position

These maps are used to drive motor planning in a straightforward manner;

- The system is presented with the usual target as before, but this time also with another toy nearby. The goal is to push the target towards the new toy. The system first foveates on the target, while also locating the new toy in its peripheral vision. The retinal displacement of the toy is used as the the desired position $\boldsymbol{r_d}$,
- the direction of this displacement vector Θ, is taken to be the direction of desired motion of the target and is used to find the direction map M_Θ with the closest matching dominant direction.
- The robot first moves its hand to the hand-position associated with map M_Θ and then begins its motion towards the target. The dynamics takes care of the rest, resulting in the motion of the target towards the desired direction.

Figure 3 shows one example of the learned maps being used to drive goal-directed action. The round toy is the new desired position towards which the target must be pushed. Note that initially in (A) the arm is in an inconvenient position to achieve the goal of pushing the target in the desired direction, but the prior learning selects a better starting position (B), leading to a successful action (C).

A quantitative measure of the improvement brought about by learning is to look at the error of target motion (towards a goal) when the hand-position is randomly chooses among the 4 starting points (Figure 3(C)) and compare it to the error after learning (Figure 3(D)) when the initial hand position is picked based on the learning. The Figure shows that the error drops from about 90deg to about 35deg. The error would be even lower if more than four starting hand-positions were considered, as would be the case if we were running the experiment in continuous mode where we would uniformly sample the space of all hand-positions.

5 Discussion

The experiment discussed here is just the very first step towards "learning to act". It makes several simplifying assumptions to test the basic idea of learning the effects of motor acts, and then driving motor-planning with the learned knowledge. There are therefore several directions for improvement, of which the major ones are:

- *Moving to a continuous space of hand-positions:* We have considered only four initial hand-positions in this experiment. To cover the whole space of initial hand positions however, a more natural approach is to pick hand-positions randomly during trials while building up a table with the hand-positions actually visited..and learning a target motion map for each those hand-positions. Another approach is of course to train a neural network with the target motion directions as inputs and the hand-positions as outputs.

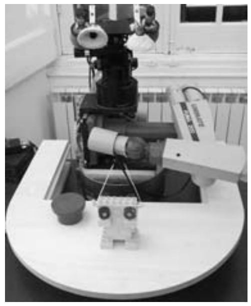

(a) The round toy is the new desired target position

(b) The learned maps are used to re-position the arm in preparation for the pushing movement

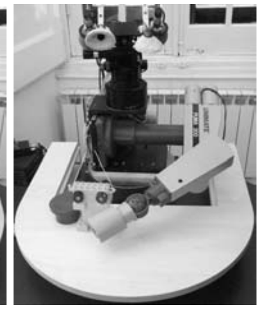

(c) At the end of the movement

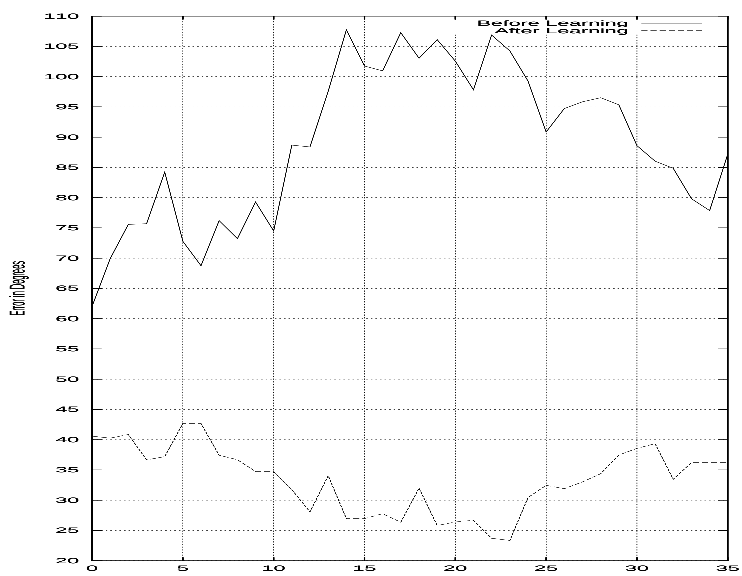

(d) Error in degrees, before(top) and after (bottom) learning, across several trials

Fig. 3. The learned direction maps are used to drive goal-directed action

- *Interleaving the learning with the planning:* At present for simplicity we first have the learning/discovery phase and then the motor planning phase. But in principle there is no need for this separation and we intend to move to a more continuous mode where both learning and planning are happening continuously.
- *Increasing the number of visual variables:* In this particular experiment the same target and hand speed were used throughout and the only variable varied was initial-hand position. However the speed of the hand and the type of target could be varied too in future experiments. This would require paying attention to a much larger set of event features: the size of the target, the distance moved, the force profile on the hand, for instance to discover useful regularity.

6 Conclusion

We have shown a system that "learns to act" on a target object. In a play/discovery phase it pushes/pulls the target from several different directions while learning about the effect of the action. In another goal-directed play phase it uses its learned maps to select the initial arm position that will enable it to push a target toy towards another toy.

The work described here makes a novel contribution towards the area of "event-interpretation" because the constraints imposed by the combined modalities of vison, motor, and proprioception may make it easier to interpret certain self-generated events than with vision alone. Furthermore, interpreting self-generated events may be a necessary first step to interpret more complex object-object events.

References

1. Kovotsky, L., Baillargeon, R.: The development of calibration-based reasoning about collision events in young infants. Cognition, Vol 67 (1998), 311–351. Elsevier
2. Spelke, E.S.: Initial Knowledge: six suggestions. Cognition, Vol 50 (1994), 431–445. Elsevier
3. Spelke, E.S., Breinlinger, K., Macomber, J., Jacobsen, K.: Origins of Knowledge. Psychological Review, Vol 99 (1992), 605–632.
4. Rizzolatti, G., Fadiga, L., Gallese, V., Fogassi, L.: Premotor cortex and the recognition of motor actions. Cognitive Brain Research, Vol 3 (1996), 131–141. Elsevier
5. Metta, Panerai, Manzotti, Sandini: Babybot: an artificial developing robotic agent. From Animals to Animats: The sixth International Conference on the simulation of Adaptive Behavior. (2000)
6. Sandini, G., Tagliasco, V.: An Anthropomorphic Retina-like Structure for Scene Analysis. Computer Vision, Graphics and Image Processing, Vol 14(3) (1980), 365-372