Maximally Informative Interaction Learning for Scene Exploration

Herke van Hoof[1], Oliver Kroemer[1], Heni Ben Amor[1] and Jan Peters[1,2]
{hoof,kroemer,amor,peters}@ias.tu-darmstadt.de

I. Introduction

Recognizing and manipulating objects is an essential ability for many robots. Robots in dynamic, unstructured environments will frequently encounter novel objects. Hence, pre-defined object information cannot always be available.

Instead, a robot can autonomously acquire knowledge through interaction with its environment. Interaction can be used to learn, among others, the appearance and shape of objects, as well as how to manipulate them.

As object segmentation is a prerequisite for learning properties and affordances of individual object, we focus on obtaining such a segmentation. In this task, passive observation is not sufficient as static scenes with novel objects are often ambiguous with respect to the true segmentation. In our approach, a robot resolves such segmentation ambiguities by pushing regions to test whether they are physically connected.

Interactive segmentation approaches have, so far, interacted with objects by performing fixed actions, random actions, or actions selected using a heuristic. In real world domains, most robots can execute a large variety of actions. Existing action selection strategies do not address the issue of how to choose the best action in a principled way.

We argue that the robot should choose the action it expects to be the most informative. We propose an autonomous system that selects maximally informative actions in a principled manner. The expected informativeness of actions is quantified using the information-theoretic measure of information gain.

II. Object segmentation using maximally informative interaction

The starting point of our segmentation approach is an over-segmentation of the visual scene into coherent segments. Subsequently, the robot actively explores the scene using pushes directed at segments. Segment movement is tracked to determine which segments constitute coherent objects.

After executing a push, the robot observes which segments moved as a result. For each pair of segments ij we define a random variable g_{ij} indicates whether these segments belong to the same object. The probability $p(g_{ij} = 1)$ is estimated given the observed data $\mathcal{D}$. Free parameters Θ are estimated using expectation-maximization to handle the presence of latent variables g_{ij}.

[1]TU Darmstadt, FB Informatik, FG IAS, Darmstadt, Germany
[2]Max Planck Institute for Intelligent Systems, Tübingen, Germany

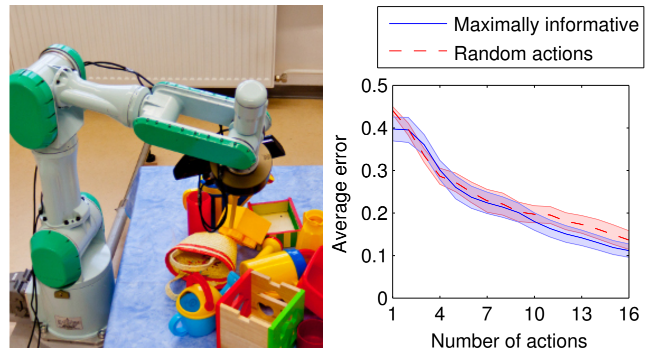

(a) (b)

Fig. 1: a) Experimental set-up. Our robot is equipped with an RGBD sensor unit and a rod to push objects. b) Mean and standard error of the per-edge error over 10 trials on the robot for both action selection methods.

Initially, no matter which segment is pushed, all pushes are expected to provide the same quantity of information. However, any obtained knowledge can be exploited to choose the most informative action. First, a probability distribution over outcomes $\mathbf{o}$ is estimated for every action v. Then, the action v is selected that maximizes the expected information gain

$$\mathbb{E}_{\mathbf{o}}[\mathrm{KL}(p(G|\mathcal{D}, \mathbf{o}, \Theta)||p(G|\mathcal{D}, \Theta))|\mathcal{D}, v, \Theta] ,$$

i.e., the expected Kullback-Leibler divergence between the joint distribution G over the pairwise connections g_{ij} before and after observing the outcome $\mathbf{o}$.

III. Experiments and results

An idealized simulation showed a significant advantage of choosing maximally informative actions over random action selection. An implementation on a real robot showed a modest increase in learning speed when maximally informative actions were used, as shown in Fig. 1b.

IV. Summary

We have proposed a bottom-up approach to object learning using interaction with the robot's environment. This approach allowed our robot to discover how to decompose a scene into objects. The task was done completely autonomously, with minimal prior knowledge. Learning was faster when maximally informative actions were used. Discovering the segmentation will be a first step in autonomously aquiring knowledge about objects for manipulation.