Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence

Toward Learning to Solve Insertion Tasks: A Developmental Approach Using Exploratory Behaviors and Proprioception

Philip Koonce, Vasha Dutell, José Farrington, Vladimir Sukhoy, and Alexander Stoytchev
Developmental Robotics Laboratory
Iowa State University
pkoonc1@swarthmore.edu, vdutell@uoregon.edu, jose.farringtonzapata@uprrp.edu, {sukhoy, alexs}@iastate.edu

Introduction

This paper describes an approach to solving insertion tasks by a robot that uses exploratory behaviors and proprioceptive feedback. The approach was inspired by the developmental progression of insertion abilities in both chimpanzees and humans (Hayashi et al. 2006). Before mastering insertions, the infants of the two species undergo a stage where they only press objects against other objects without releasing them. Our goal was to emulate this developmental stage on a robot to see if it may lead to simpler representations for insertion tasks. Experiments were performed using a shape-sorter puzzle with three different blocks and holes.

Prior work on insertion tasks in Robotics (also known as peg-in-hole tasks) showed that proprioceptive feedback was effective in improving performance. Learning from this feedback allowed a robot to complete insertions more often (Lee and Kim 1988) and eliminated the need for very precise interaction models (Gullapalli 1995). In contrast, very accurate geometric and kinematic models were required to plan insertions without proprioception (Bruyninckx et al. 1995).

The closest analog to our approach to insertion tasks is the work under the COSPAL project, in which individual stages of the solution were verified using visual feedback (Felsberg et al. 2005). Our work differs from the COSPAL project in two ways: 1) our robot was fully autonomous, while in COSPAL it was bootstrapped by demonstration; and 2) our robot did not use vision, relying only on proprioception.

Experimental Setup

All experiments were performed using the upper-torso humanoid robot shown in Fig. 1. The robot's arms are two backdrivable 7-dof Barrett Whole Arm Manipulators (WAMs), each equipped with a 3-finger Barrett Hand. Only the left arm was used in the experiments.

Three different blocks and holes were used: circle, cross, and hexagon (see Fig. 1). Each block could fit into only one hole, for which there was less than 1 mm of free space between the block and the hole boundaries. The board with the three holes was mounted on a wooden fixture as shown in Fig. 1.a. It could slide left and right to present a different hole in front of the robot, which was stationary.

The robot performed 20 trials for each of the 9 combinations of 3 blocks and 3 holes for a total of 180 trials. A typical trial lasted about 50 seconds. Each trial was started from a fixed position in joint space with the block already grasped by the robot. During each trial the robot performed five exploratory behaviors, which consisted of pushing, sliding, and rotating the block in the vicinity of the hole. All behavioral parameters were generated at random, so each behavior was unique. If the robot inserted the block successfully, then the human experimenter signaled the robot to try to perform five more exploratory behaviors in the same trial, while the block was in the hole.

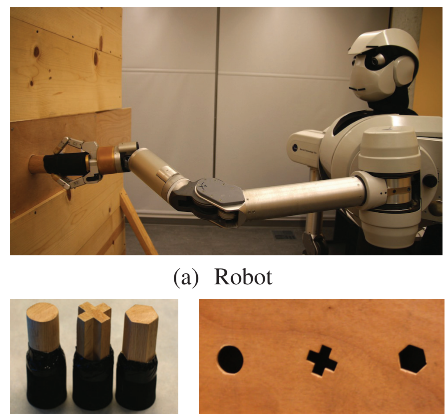

(a) Robot

(b) Blocks (c) Holes

Figure 1: Experimental setup: (a) the upper-torso humanoid robot used in the experiments, shown here while inserting a block; (b) the 3 blocks – circle, cross, and hexagon; (c) the 3 holes.

Proprioceptive data was recorded at 500 Hz in the form of joint torques and positions for the left arm. Audio and video were also recorded, but were not analyzed in this paper.

Methodology

The proprioceptive data for each trial was partitioned into two categories. All data recorded before an insertion was assigned to one category. All remaining data (if any) was assigned into another category. The partitions for one trial are shown in Fig. 2. Visual inspection shows that there are differences between the categories. For a more principled analysis, the joint torques recorded by the robot were split into 8-second segments with 50% overlap. For each trial there were about 11-12 segments. For all 180 trials there were 2020 segments. These segments were also partitioned into the *before* and *after* insertion categories.

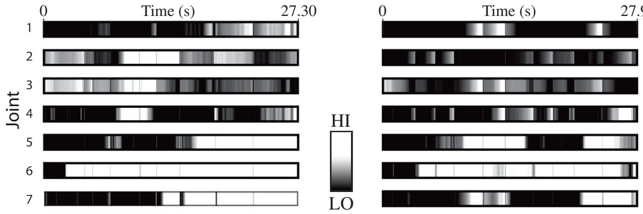

(a) Before insertion (b) After insertion

Figure 2: Sample joint torque data recorded from the WAM for eight seconds *before* (a) and *after* (b) the robot successfully inserted the cross-shaped block during one of the trials.

For every 8-second segment, a 7×7 correlation matrix was calculated to quantify the pairwise correlations between the joint torques. To measure the differences between different 8-second segments, the Euclidean distance between their correlation matrices was computed. This process resulted in a 2020×2020 distance matrix D.

Results

To quantify the ability of the robot to solve insertion tasks using exploratory behaviors, the percentage of trials with insertions was computed for all three matching block-hole combinations (see Fig. 3). The robot inserted the circle 55% of the time (11 out of 20 trials). The cross block was inserted 25% of the time and the hexagon 20%. The success rate was unexpectedly high, given that the explorations were performed randomly, without explicitly focusing on the insertion task. This suggests that insertion tasks may not be as hard as the complex topological contact models utilized by previous approaches seem to imply. If the robot is allowed to collide the peg with the walls and use random exploration to solve the rest of the task, then these detailed models might be unnecessary. Even more surprisingly, the results show that random exploration can solve insertion tasks using proprioception alone, completely ignoring visual feedback.

The data was analyzed to find differences between segments recorded before and after insertions. For all seven joints, the variance of the torque increased after insertions. In addition, the mean torque increased for all joints above the robot's wrist (joints 1, 2, 3, and 4). When the robot performed exploratory behaviors after an insertion, the torque safety limits were often exceeded because the robot's movements were constrained. If torque in any joint reached its safety limit, then all joint torques were reset to zero, the current behavior was interrupted, and the next one was started.

Additional analysis was performed to detect differences in joint torque correlations. The distance matrix D described above was embedded in lower dimensions using Isomap (Tenenbaum et al. 2000). In the 3D embedding, shown in Fig 4.a, there is a separation between the data points for the segments recorded before and after insertions. In particular, the sample means and standard deviations for the coordinates of the two groups of points (shown in Fig. 4.b) indicate significant differences between the categories. Even the data points recorded for the circle block were separated, despite the fact that this block affords relatively unconstrained wrist rotations to the robot. This suggests that the joint torques are correlated differently before and after insertions. The 2D embedding was not useful because its residual error was too high (see Fig. 4.c).

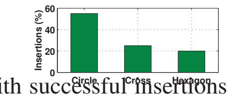

Figure 3: Trials with successful insertions as a percentage of the trials with matching block-hole combinations. In each trial the robot performed five random exploratory behaviors that consisted of pushing, sliding and rotating the block in the vicinity of the hole.

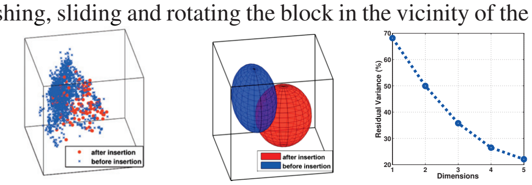

(a) 3D Isomap (b) Ellipsoids (c) Isomap error

Figure 4: Isomap embedding of the distance matrix D: (a) the 3D embedding; (b) ellipsoids fitted to the data points shown in (a); (c) residual error of the embedding as a function of dimensionality.

Conclusion and Future Work

This paper showed that a robot can use exploratory behaviors combined with proprioceptive feedback to perform insertion tasks. Even though the exploration was not focused, i.e., the robot was not actively trying to insert the blocks, the three blocks were inserted in at least 20% of trials.

It was demonstrated that proprioception can be used to verify insertions. In particular, it was shown that joint torques recorded before and after insertions were statistically different. Future work can use this to let robots verify insertions autonomously (e.g., a robot can release the peg after it reaches, say, 95% confidence of a successful insertion).

Future work can evaluate how the proposed approach generalizes. In another experiment, the robot was able to insert a 240V electric plug into its socket using the same exploratory behaviors. This suggests that the approach can solve other insertion tasks. In other words, it can scale up.

For more details and a video of the experiments see: http://www.ece.iastate.edu/~alexs/lab/projects/reu2010/.

This work was funded in part by NSF Research Experience for Undergraduates (REU) Grant IIS-0851976.

References

Bruyninckx, H.; Dutre, S.; and De Schutter, J. 1995. Peg-on-Hole: A model based solution to peg and hole alignment. In *Proceedings of IEEE ICRA*, volume 2, 1919–1924.

Felsberg, M.; Forssén, P.; Moe, A.; and Granlund, G. 2005. A COSPAL subsystem: Solving a shape-sorter puzzle. In *AAAI Fall 2005 Symposium*, 65–69.

Gullapalli, V. 1995. Skillful control under uncertainty via direct reinforcement learning. *Robotics and Autonomous Systems* 15(4):237–246.

Hayashi, M.; Takeshita, H.; and Matsuzawa, T. 2006. Cognitive development in apes and humans assessed by object manipulation. In *Cognitive Development in Chimpanzees*. Springer Tokyo. 395–410.

Lee, S., and Kim, M. 1988. Learning expert systems for robot fine motion control. In *Proceedings of the IEEE International Symposium on Intelligent Control*, 534–544.

Tenenbaum, J.; Silva, V.; and Langford, J. 2000. A global geometric framework for nonlinear dimensionality reduction. *Science* 290(5500):2319.