Learning to Solve Insertion Tasks Using Exploratory Behaviors

Philip Koonce, Vasha Dutell, José Farrington,
Vladimir Sukhoy, Shane Griffith, and Alexander Stoytchev
pkoonc1@swarthmore.edu, vdutell@uoregon.edu, jose.farringtonzapata@uprrp.edu,
{sukhoy, shaneg, alexs}@iastate.edu

***Abstract*— The shape-sorter puzzle is a quintessential toy that children play with to learn how to fit objects in holes. Psychology research shows that human children learn object fitting during an early developmental period. We implemented similar developmental learning in our robot using random exploratory behaviors and proprioceptive feedback with the shape sorter puzzle. Analysis shows that successful fits can be detected from proprioceptive feedback streams. Ultimately, we hope to show that this developmental approach can enable systems to be more robust and capable of accurately detecting and eventually predicting successful insertions in new, more complicated peg-in-hole situations.**

I. INTRODUCTION

Human children undergo a developmental period during which they explore their environment and capabilities [6]. This developmental period is crucial for learning to operate in the real world. Developmental robotics seeks to emulate this period in robots. Robots that learn developmentally are programmed to construct their own categories and concepts of their environment by interacting with it. Developmental solutions to problems in robotics have the potential to be more useful because they generalize better. If we are successful in teaching the robot to solve the shape-sorter puzzle using these principles and can generalize the solution to more complicated peg-in-hole problems, then we can further validate developmental robotics as a more robust way of programming.

There are three basic developmental principles [11] on which this work is based: 1) a robot must be able to verify all knowledge it gains, 2) the robot must have a physical body in order to interact with its environment and physically verify anything it learns, and 3) the robot must be able to create its own representations of its environment as it perceives them through its sensors. In response to many failures throughout the history of AI to create generalizable intelligence, Brooks wrote that systems must be able to construct their own representation of the environment because their perception of the world is different from the programmer's perception [1]. With these ideas in mind, we developed a system that records proprioceptive feedback as it attempts to fit blocks into holes by playing with a shape sorter puzzle.

II. RELATED WORK

A. Developmental Psychology

Research in developmental psychology describes the developmental period during which human children learn, among other things, to insert objects into holes [5]. Given a circular peg and a box with a corresponding hole and without demonstration of insertion, children aged 10 months only rub the rod against the box. Only at 44 months do children begin to learn to insert a rod into a hole [5]. Furthermore, not for another year (56 months) do children release the rod into the hole. This suggests that at this stage they are able to detect that they have achieved a fit. In the same study, children were given multiple shapes to fit [5]. Children are first able to complete insertion of a block at 12 months if given a demonstration by adult, and at 44 months without demonstration. This is the period of learning that this study worked toward.

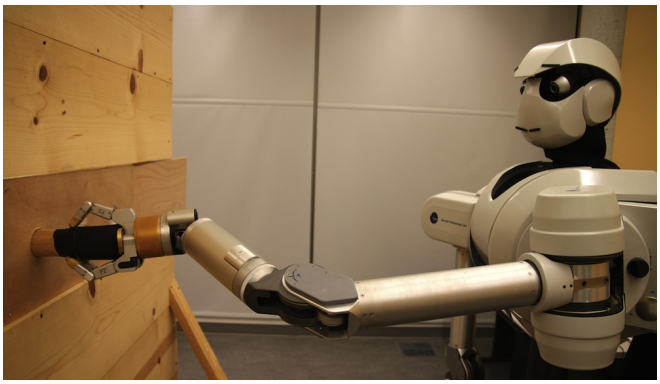

Fig. 1. The upper-torso humanoid robot used in the experiments attempting an insertion.

B. Robotics

The peg-in-hole problem has historically been approached using geometrical representations and simulations [9]. Some researchers have extended their systems to include physical implementations [2] [7] [13]. These pre-programmed solutions require high accuracy and are not often generalizable to new pegs and holes in uncontrolled environments. Some work has been done approaching the problem using developmental learning to improve fitting [8], and others have used proprioception to detect fits more generally [9], but in all these cases the goal was to decrease time and increase accuracy in completing an insertion, not to detecting a successful fit. Furthermore, these solutions were often restricted to a cylindrical peg in a circular hole. The act of insertion and the fit itself will "feel" different for different blocks, so these solutions cannot generalize well. For example, a circle will rotate within the hole both during the fitting process and

Fig. 2. The robot's environment.

during the fit itself, while all other shapes will be constrained. Different shapes will also have different clearances within the hole, therefore movement within the fit will be constrained to a different degree among hole and block pairs.

The shape-sorter puzzle can be solved in a simulation using machine learning algorithms and a visual model that does not include explicit modeling of shapes [3]. This approach has also been successfully implemented with a physical robot that can solve the block game given a bootstrapping phase during which a teacher performs a successful trial, and the system learns by observing this act. [4]. In neither of these approaches was the robot able to detect a sucessful fit from experience.

III. EXPERIMENTAL SETUP

A. Robot

All block fitting experiments were performed with the upper-torso humanoid robot shown in Fig. 2. This robot had two 7-DOF Barrett Whole Arm Manipulators (WAMs) as its arms. The end effectors of both WAMs were two BarrettHands BH8-262. Only the left arm was used in the experiments. The head contained two Logitech webcams for recording vision and two Audio-Technica U853AW UniPoint cardioid condenser microphones to record audio.

B. Experimental Fixture with Blocks and Holes

Three different shapes of blocks and holes were used in the experiments: the circle, the cross, and the hexagon, (see Fig. 3(a)). Each block could fit only one hole. The blocks and the holes were cut from wood. To make grasping more reliable and to use the same grasping strategy for all three blocks, they were placed in foam can coolers. These coolers

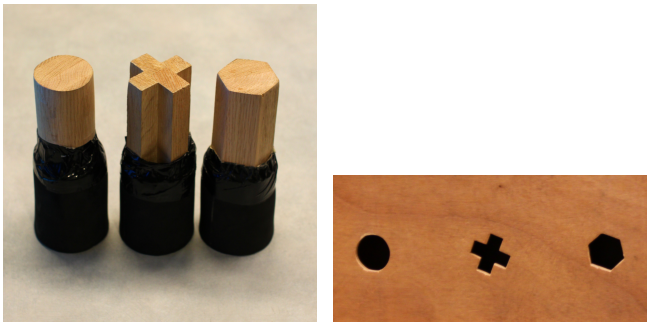

(a) Blocks. (b) Holes.

Fig. 3. Wood blocks and holes used in the experiments. The blocks are inserted into foam can coolers and sealed with duct tape to make robotic grasping more robust.

prevented the wood blocks from sliding in the grip of the robot's aluminum fingers.

The three holes were cut as shown in Fig. 3(b). The plate with holes was mounted on a wood fixture located in front of the robot, as shown in Fig. 2. The plate could slide left and right in the fixture, so that a specific hole could be presented to the robot for each experimental trial.

C. Experimental Trials

For each of the 9 combinations of the 3 blocks and 3 holes, the robot performed 20 trials, for a total of $3 \times 3 \times 20 = 180$ trials. A typical trials lasted for 90-120 seconds.

The robot started each trial from a pre-defined position in joint space. Each trial consisted of: 1) grasping a block, 2) pushing randomly against the fixture in the region around the hole, 3) performing exploratory behaviors, and 4) retrieving the arm to start another trial. Blocks for grasping were provided by the human experimenter so that the block faced the experimental fixture after the robot grasped it as shown in Fig. 1.

To push to the hole, a joint configuration was sampled randomly in the pre-defined region around the hole. The robot then moved from the initial position in the direction of the randomly sampled joint configuration until it hit the board. The robot detected collisions with the board using joint torques: when the torque in any joint exceeded its limit, the robot stopped its current movement and started a new one. The robot performed at least five exploratory behaviors to attempt to fit the block. Each of these behaviors consisted of pushing, sliding, and rotating the block in the vicinity of the hole. More specifically, a behavior was performed as follows:

1) Generate a random point P_1 in the exploratory region around the hole.
2) Calculate a target destination T_1 by scaling the vector from the initial position to P_1 by 1.2.
3) Move arm to T_1.
4) Generate a random point P_2 in the exploratory region.
5) Calculate a target destination T_2 by scaling the vector from the initial position to P_2 by 0.9.
6) Move arm to T_2.

When the robot successfully inserted the block into the hole, the experimenter signaled the robot with a keyboard command. Upon receiving this signal, the robot began five more exploratory behaviors and recorded the timestamp of the beginning of the "fit segment" in a log file. These additional behaviors were performed exactly like the behavoirs used to fit the block, except that the block was inside the hole.

After all explorations, the robot returned its arm to the initial position and started a new trial.

D. Sensory Data

The sensory modalities were recorded as follows: vision at 15 fps, 640x480, audio at 44.1 kHz, 16-bit stereo, and proprioceptive data in the form of joint positions, velocities, and torques from the WAM at 500 Hz.

IV. METHODOLOGY

The goal of this study was to investigate if a robot can detect when it has fit a block into a hole. In other words, the question is: "Is there a pattern in the robot's sensorimotor experience that corresponds to sensations after it fit a block in a hole?"

A. Data Partitioning

To answer this question, the proprioceptive data recorded by the robot during exploration was partitioned into 10 second segments that overlapped by 50%. The total number of these segments in the dataset was 2020. These segments were assigned to two categories: "fit" and "no fit", according to human experimenter labels described in Section III-C. Joint torque data for the example segments in these two categories is shown in Fig. 4.

B. Torque Correlation Patterns

Visual inspection of the joint torque data in Fig. 4 suggests that torques in "fit" segments are more correlated than in "no fit" segments. This suggests that torque correlation patterns can separate segments into two categories.

To quantify correlations between torque values in data segments across different joints, the Pearson correlation coefficient was used. That is, given the two vectors of joint torque readings X and Y, the correlation coefficient is

$$\rho_{X,Y} = \frac{\mathrm{cov}(X,Y)}{\sigma_X \sigma_Y},$$

where $\mathrm{cov}(X,Y)$ is the covariance of the two samples, σ_X and σ_Y are the standard deviations. The covariance is defined as

$$\mathrm{cov}(X,Y) = E[(X-\mu_X)(Y-\mu_Y)],$$

where μ_X and μ_Y are the mean values of X and Y, and E is the operator that computes the expected value. The value of the correlation coefficient $\rho_{X,Y}$ lies between -1 and 1. Values close to -1 correspond to high negative correlation (as one value increases, another decreases). Values close to zero mean that the two samples are not correlated. Values close to 1 correspond to high positive correlation (i.e. both values increase simultaneously).

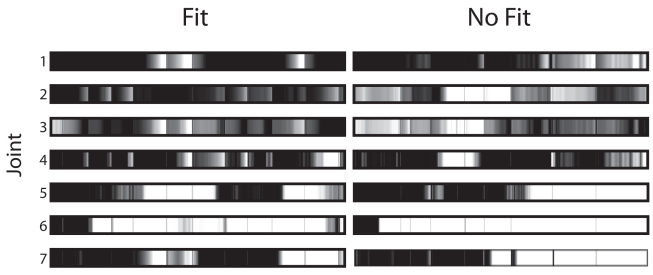

Fig. 4. Raw proprioceptive feedback from the seven joints.

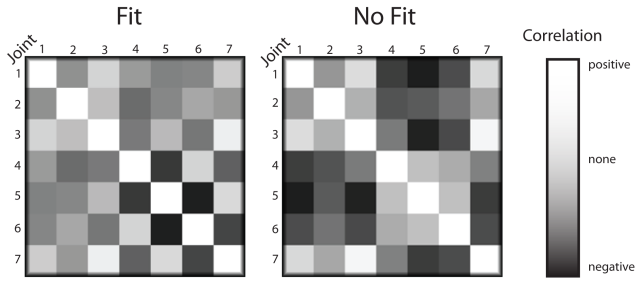

Fig. 5. Correlation matrices for the seven joints.

Each data segment contained the torque data for the 7 joints of the WAM. The correlation coefficients were computed for each of the 49 possible joint combinations and grouped into 7×7 correlation matrices. Two example correlation matrices: one for a "fit" segment, and another for a "no fit" segment from the dataset are shown in Fig. 5.

V. RESULTS

A. Distance Metric Based on Torque Correlations

The 7×7 torque correlation matrices for the data segments, described in Section IV-B, were used to compute the distances between these segments. More formally, given the two segments A and B and their corresponding torque correlation matrices $M^{(A)}$ and $M^{(B)}$, the distance between A and B was set to

$$d_{AB} = \sqrt{\sum_{i=1}^{7}\sum_{j=1}^{7}(M_{ij}^{(A)} - M_{ij}^{(B)})^2}.$$

These distances were computed between all 2020 segments in the dataset, resulting in a distance matrix D with $2020 \times 2020 = 4080400$ entries. This distance matrix quantifies similarities and differences in torque correlation patterns for all segments in the dataset.

B. Isomap Embeddings for the Distance Matrix

To see if it is possible to detect "fit" data segments in proprioceptive stream based on the joint torque correlation patterns, the 2020×2020 distance matrix D was embedded in $\mathbb{R}^n$, where $n = 1, \ldots, 10$. These embeddings were

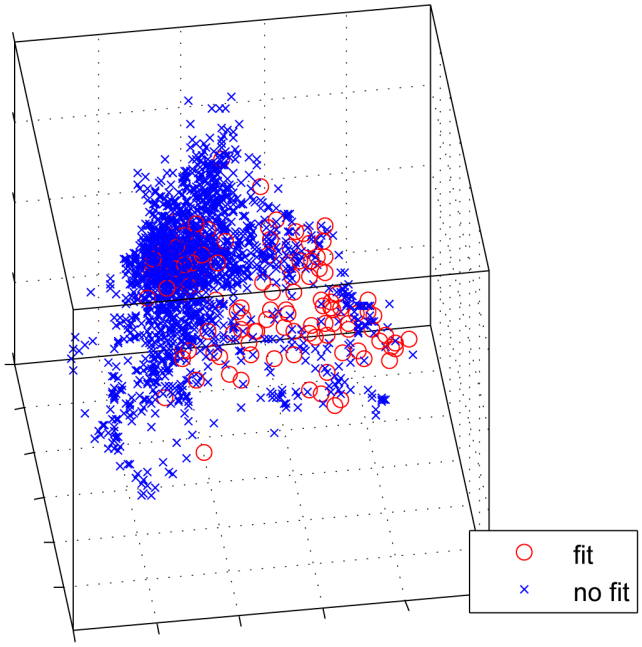

Fig. 6. Isomap embedding of correlation metrics

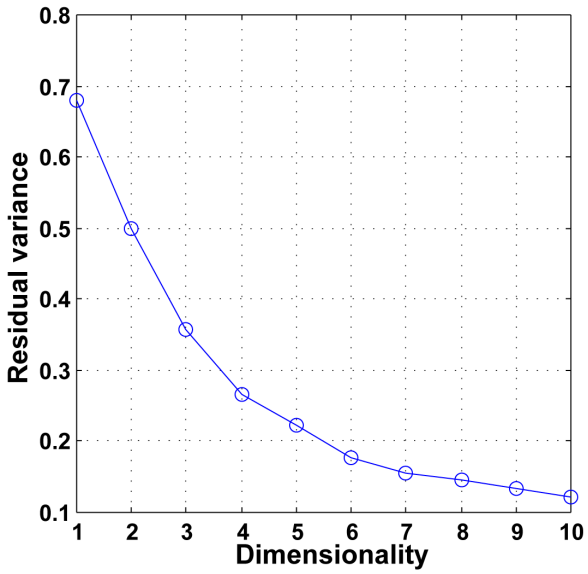

Fig. 7. Residual variance in Isomap embeddings for different dimensions

performed using the Isomap algorithm [12]. The residual errors for different values of n are shown in Fig. 7. The 2D embedding was not useful, as its residual error was too high. The 3D embedding is shown in Fig. 6.

In the 3D embedding, the data points for "fit" data segments are clearly separated from the mass of "no fit" data segments. This suggests that there is a pattern of torque correlations in the "fit" segments. This pattern can be used to separate the two categories.

VI. CONCLUSIONS AND FUTURE WORK

The results of this study show that it is possible for a robot to detect a fit from proprioception alone. These results also show that random explorations are effective for fitting blocks into holes. Also, exploratory behaviors proved to generalize well for other fitting tasks: the robot successfully fit a 240 volt electric plug into a socket in only minutes of exploration using the same behaviors that were used to fit blocks into holes.

The study did not implement any learning or extend the data analysis to yield any improvement in fitting performance; it only demonstrated that there is a salient difference in proprioception between fits and non-fits for a block, such that it is possible to detect fits using proprioceptive data alone. The results also show that a predictive model could easily be implemented to accurately detect successful fits Such results justify further work in this direction.

This study can be extended in two steps: 1) implement an effective unsupervised clustering algorithm that accurately separates fits from non-fits, and 2) perform data analysis and learning in real time as the robot explores. With these two contributions, the robot can test the accuracy of its fitting and learn incrementally.

Finally, it was shown that multiple modalities improve machine learning performance [10]. In this experiment, audio and video were recorded as the robot fit blocks into holes. This study used neither audio nor video in the analysis, but it can be used to train an unsupervised visual model that improves accuracy of fit detection. Accuracy can also be increased with a larger training set, i.e. more differently shaped block and hole pairs.

REFERENCES

[1] R. Brooks. Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159, 1991.

[2] H. Bruyninckx, S. Dutre, and J. De Schutter. Peg-on-hole, a model based solution to peg and hole alignment. In *Proceedings of ICRA*, pages 1919–1924, 1995.

[3] M. Felsberg, P. Forssén, A. Moe, and G. Granlund. A cospal subsystem: Solving a shape-sorter puzzle. In *AAAI Fall Symposium: From Reactive to Anticipatory Cognitive Embedded Systems*, pages 65–69, 2005.

[4] M. Felsberg, J. Wiklund, and G. Granlund. Exploratory learning structures in artificial cognitive systems. *Image and Vision Computing*, 27(11):1671–1687, 2009.

[5] A. Gesell. Infant Behavior: Its Genesis and Growth. *Behavior*, 49(6), 1935.

[6] E.J. Gibson. Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge. *Annual review of psychology*, 39(1):1–42, 1988.

[7] S. Gordon. Automated assembly using feature localization. *AITR-932*, 1986.

[8] V. Gullapalli, J. Franklin, and H. Benbrahim. Acquiring robot skills via reinforcement learning. *Control Systems Magazine, IEEE*, 14(1):13 –24, feb 1994.

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

[10] J. Sinapov and A. Stoytchev. The Boosting Effect of Exploratory Behaviors. In *Proceedings of 24th AAAI Conference on Artificial Intelligence*, pages 1613–1618, 2010.

[11] A. Stoytchev. Some basic principles of developmental robotics. *IEEE Transactions on Autonomous Mental Development*, 1(2):122–130, 2009.

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

[13] D. Whitney. Quasi-static assembly of compliantly supported rigid parts. *Journal of Dynamic Systems, Measurement, and Control*, 104:65–78, 1982.