Toward learning to solve the peg-in-hole problem

Vasha Dutell, Philip Koonce, José Farrington
Vlad Sukhoy, Shane Griffith, Alexander Stoytchev

Motivation

The peg-in-hole problem is solvable in restricted situations. The goal of this research is to create a generalizable system that fits objects effectively and efficiently using exploratory behaviors and learns to detect these fit events.

Developmental Robotics

Humans go through a developmental period during which they explore their environment and capabilities. Developmental robotics seeks to implement a similar period for robots. A robot is programmed to explore its environment, and then develops its own categories and concepts of the environment based on that exploration.

Related Work

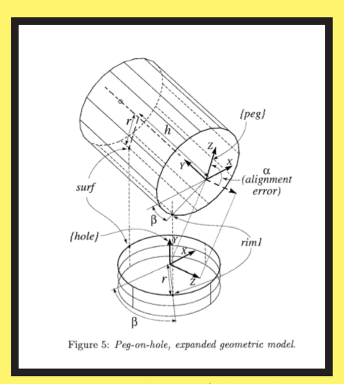

[Bruyninckx et al. (1995)]

An example of a traditional pre-programmed geometric approach to solving the peg-in-hole problem.

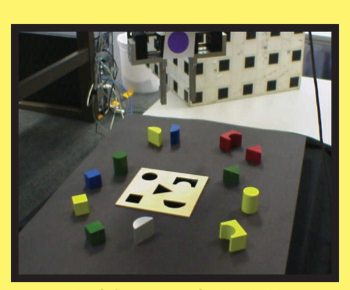

[Felsberg et al. (2005)]

COSPAL - Solved the shape sorter puzzle visually by the robot observing a teacher grasp and place blocks.

Methods

Experimental Setup → Exploratory Behaviors → Multimodal Feedback → Analysis

Experimental Setup

Shapes: circle, cross, and hexagon. Rubber cup holders used to simplify grasping.

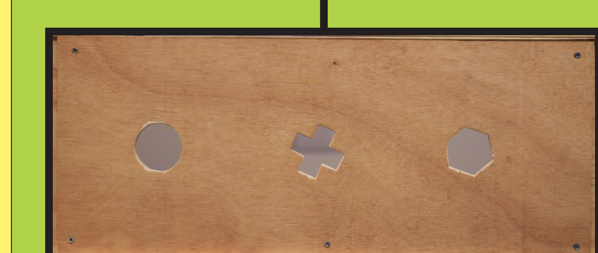

Board with holes is mounted on wall in front of the robot.

Exploratory Behaviors

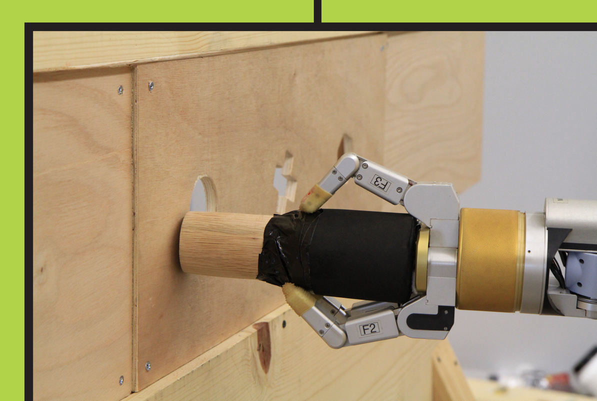

The robot pushes the board with a block until it is stuck. The robot then explores the region around the hole using pushing behaviors. The robot makes 5 attempts to fit the block. If it achieves a fit, 5 additional pushing behaviors are performed to explore the 'fit' state. 20 of these trials are performed for every block and hole combination.

3 blocks X 3 holes X 20 trials = 180 trials

Multimodal Feedback

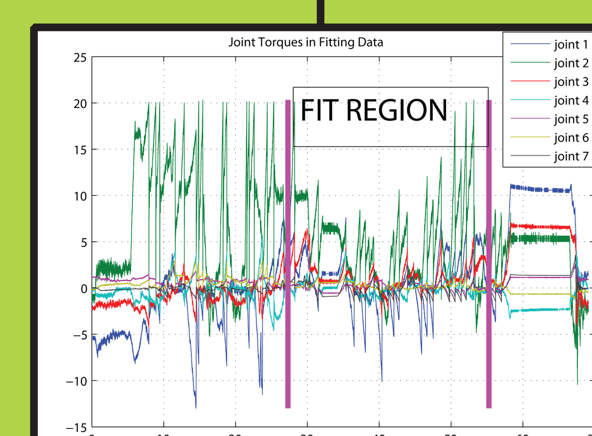

During exploration, robot records joint positions, joint torques, audio, and video. This data is time-stamped. The human experimenter detects when the robot has fit the block, signals the robot to perform additional pushing behaviors, and marks the data accordingly.

Analysis

Segment proprioceptive stream into 10 second slices with 50% overlap. Mark slice as 'fit' if it falls into temporal interval of a fit. Compute joint torque correlations from these slices. Use these correlations to quantify similarity between slices. Embed the similarity in low-dimensional space (3D) using Isomap. Detect patterns in embedded 'fit' slices.

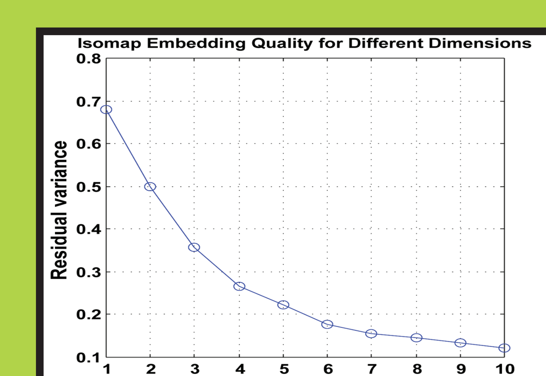

Results

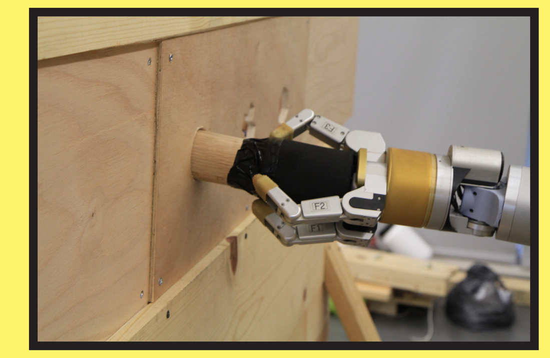

The robot fit the block with 15%-50% success, depending on the shape.

Human labels of 'fit' and 'non-fit' for 10 second slices. Two sample slices are shown.

Fit

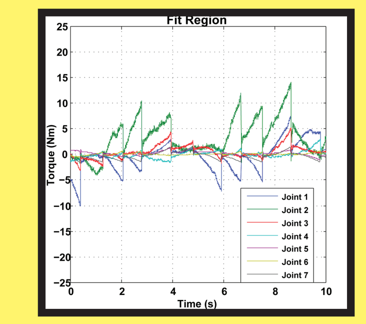

No Fit

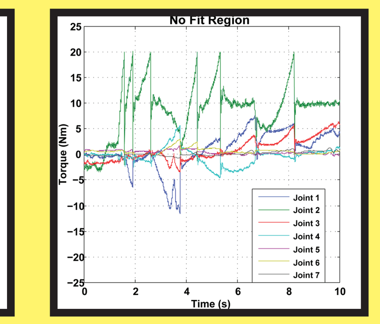

Joint torque is extracted from proprioceptive stream

Fit Correlation Matrix

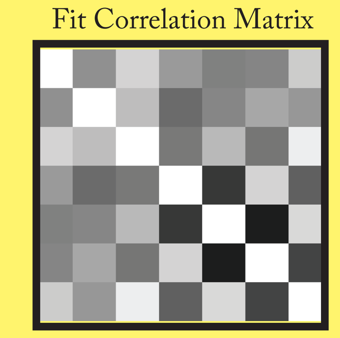

No Fit Correlation Matrix

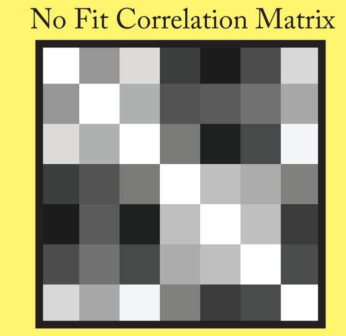

Joint torque correlations are computed.

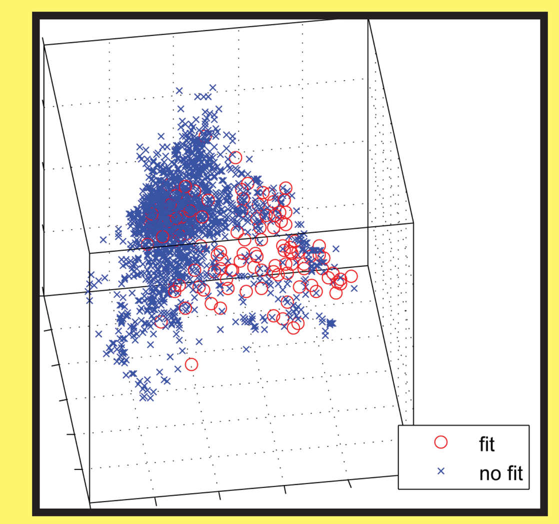

3D Isomap Embedding of similarities between correlation matrices

Isomap embedding shows that fits and non-fits are clustered. This research shows that it is possible to distinguish fits from non-fits using the robot's experience.

Future Work

- Distinguish fits from non-fits from robot's experience
- More blocks and holes
- Use audio and video to better detect fits
- Learn visual model to detect novel holes
- Learn to use electrical plugs with wall sockets

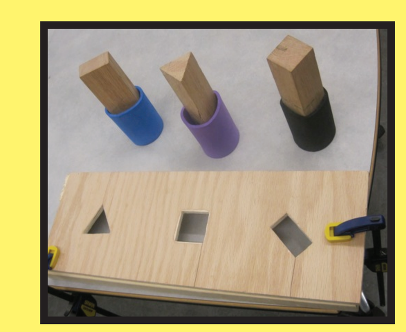

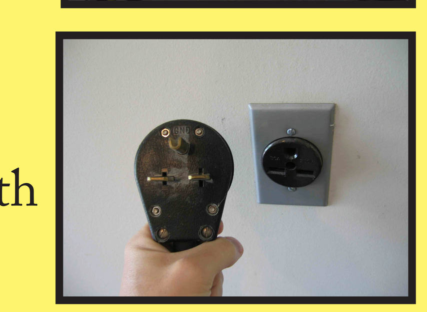

Funded by NSF Grant IIS-0851976

IOWA STATE UNIVERSITY