Construction of a 3D Object Recognition and Manipulation Database from Grasp Demonstrations

David Kent (RBE)
Advisor: Professor Sonia Chernova (CS)

Robot Autonomy and Interactive Learning

Abstract

Object recognition and manipulation are critical in enabling robots to operate within a household environment. There are many grasp planners that can estimate grasps based on object shape, but these approaches often perform poorly because they miss key information about non-visual object characteristics. Object model databases can account for this information, but existing methods for constructing 3D object recognition databases are time and resource intensive. We present an easy-to-use system for constructing object models for 3D object recognition and manipulation made possible by advances in web robotics. The database consists of point clouds generated using a novel iterative point cloud registration algorithm, which includes the encoding of manipulation data and usability characteristics. The system requires no additional equipment other than the robot itself, and non-expert users can demonstrate grasps through an intuitive web interface with virtually no training. We validate the system with data collected from both a crowdsourcing user study and a set of grasps demonstrated by an expert user. We show that the crowdsourced grasps can be just as effective as expert-demonstrated grasps, and furthermore the demonstration approach outperforms purely vision-based grasp planning approaches for a wide variety of object classes.

Project Goals

Develop a system for constructing an object recognition and grasping database with the following features:

- Database allows for novel objects to be added easily
- Data collection requires no further equipment than the robot itself
- Grasps can be demonstrated by non-expert users

Model Construction

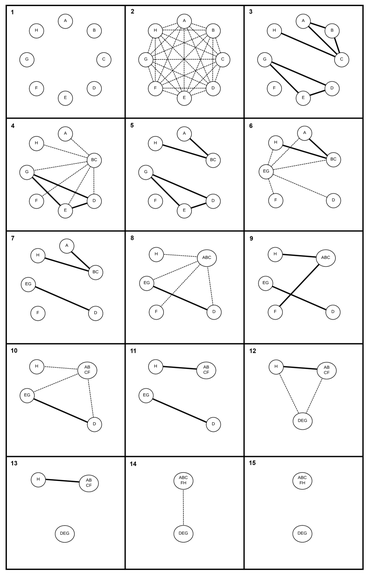

- Graph-based algorithm for iterative pairwise registration on a set of unordered object point clouds
- Nodes represent point clouds, dashed edges represent untested merges, solid edges represent potentially successful merges
- Algorithm evaluates candidate pairwise merges based on set of metrics such as percentage of overlapping points, distance error, and color difference
- Algorithm transforms grasp pose associated with each point cloud
- Results in object models with sets of example grasp poses within their reference frames

Experiment

- Conducted remote user study with 42 participants
- Participants controlled PR2 through browser-based interface
- Participants picked up as many objects as possible within 20 minutes
- Logged point cloud data and grasp poses

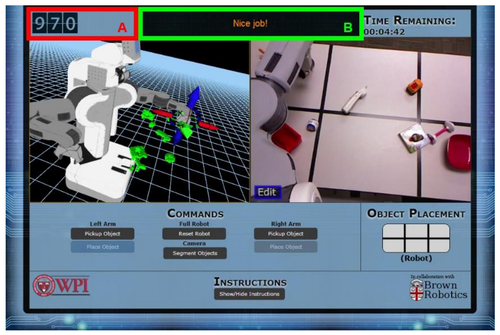

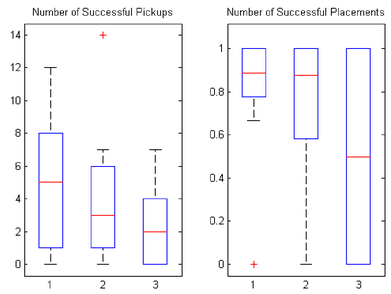

Study Conditions:

1. Full Feedback – interface provided score and comments
2. Score Only – interface provided only score
3. No Feedback – participants received no real-time feedback

Grasp Learning

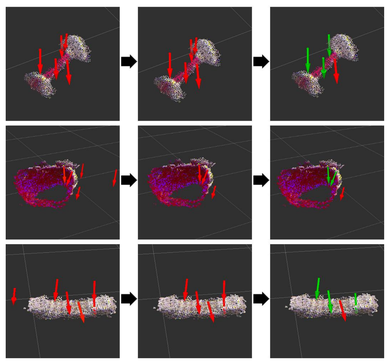

- Designed system to filter out unsuccessful grasps and learn grasp probabilities
- Takes output of model construction system as input
- Filters unsuccessful grasps by distance from the object
- Experimentally calculates grasp rate with online epsilon-greedy exploration algorithm
- Results in filtered set of grasps with high-probability grasps identified, shown in green

Object Recognition Results

Object \ Classification	Ball	Basket	Bone	Book	Bowl	Cup	Dragon	Duck	Phone	Truck
Ball	0.680	0	0	0	0	0	0.040	0	0.280	0
Basket	0	1.000	0	0	0	0	0	0	0	0
Bone	0	0	0.743	0	0	0	0.029	0	0.229	0
Book	0	0	0	1.000	0	0	0	0	0	0
Bowl	0	0.053	0	0	0.905	0.042	0	0	0	0
Cup	0	0	0	0	0	1.000	0	0	0	0
Dragon	0	0	0	0	0	0	1.000	0	0	0
Duck	0	0	0	0	0	0	0.400	0.600	0	0
Phone	0	0	0	0	0	0	0.050	0	0.950	0
Truck	0	0	0	0	0	0	0.300	0	0	0.700

Confusion matrix showing the classification rate of each object in the user study set

Grasping Results

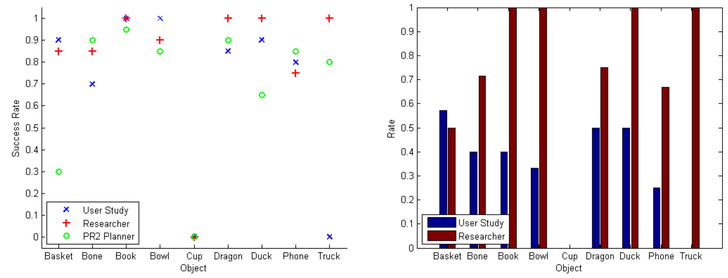

Grasp success rate (left) and high-probability grasp rate (right) per object

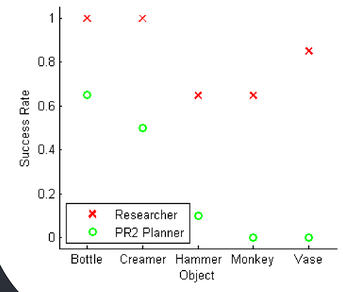

Above: The above figure shows example grasps calculated by a geometric grasp planner (top) and demonstrated by an expert user (bottom)
Left: Comparison of the grasp success rate from the developed algorithm and the PR2's off-the-shelf grasp planner