Robotic Grasping and Manipulation Competition: Competitor Feedback and Lessons Learned

Joe Falco[1], Yu Sun[2], and Maximo Roa[3]

[1]National Institute of Standards and Technology (NIST)
[2]University of South Florida
[3]German Aerospace Center (DLR)

Abstract. The First Robot Grasping and Manipulation Competition, held during IROS 2016, allowed researchers focused on the application of robot systems to compare the performance of hand designs as well as autonomous grasping and manipulation solutions across a common set of tasks. The competition was comprised of three tracks that included hand-in-hand grasping, fully autonomous grasping, and simulation. The hand-in-hand and fully autonomous tracks used 18 predefined manipulation tasks and 20 objects. Additionally, a bin picking operation was also performed within the hand-in-hand and fully autonomous tracks using a shopping basket and a subset of the objects. The simulation track included two parts. The first was a pick and place operation, where a simulated hand extracted as many objects as possible from a cluttered shelf and placed them randomly in a bin. The second part was a bin picking operation where a simulated robotic hand lifted as many balls as possible from a bin and deposited them into a second bin. This paper presents competitor feedback as well as an analysis of lessons learned towards improvements and advancements for the next competition at IROS 2017.

Keywords: robot, grasping, manipulation, competition, benchmarks

1 Introduction

The first Robot Grasping and Manipulation Competition, held during the 2016 International Conference on Intelligent Robots and Systems (IROS) in Daejeon, South Korea was sponsored by The IEEE Robotics and Automation Society (RAS) Technical Committee (TC) on Robotic Hands Grasping and Manipulation (RHGM) [1]. The goal of the competition was to bring together researchers focused on the application of robot systems to benchmark the performance of autonomous grasping and manipulation solutions across a variety of application spaces, including healthcare, manufacturing, and service robotics. Being the first of a planned series of competitions in the area of grasping and manipulation, this competition was designed to evaluate the performance of robot solutions that include grasp planning, end-effector design, perception, and manipulation control.

2 Competition Overview

The competition was comprised of three tracks; hand-in-hand grasping, fully autonomous grasping and manipulation, and simulation. The hand-in-hand and fully autonomous tracks used 18 predefined manipulation tasks and 20 objects that were readily obtainable through on-line retailers. Additionally, a bin picking operation was performed for these two tracks using a shopping basket and a subset of the objects. In order to help teams prepare their systems for the competitions, the rules, along with 10 randomly chosen predefined tasks and supporting objects, were provided one month prior to the event. The complete set of competition tasks (Figure 1) and supporting objects (Figure 2) were released to contestants one week before the competition and the actual IROS competition setup and objects were available for test two days before the competition. The competition design used many items from the Yale-CMU-Berkeley (YCB) Object and Model Set [2] and the 2015 Amazon Picking Challenge (APC2015) [3] object datasets[1]. The YCB dataset was designed for developing benchmarks in robotic grasping and manipulation research and the APC2015 dataset supports the Amazon Picking Challenge, a competition developed to spur advancement in fundamental technologies for automated picking in unstructured warehouse environments.

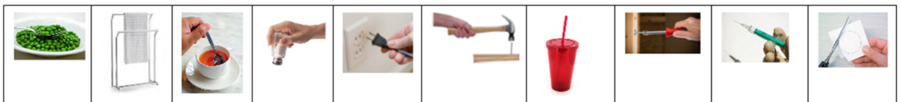

Pick-up peas with spoon	Hang towel on rack	Stir water with spoon	Shake salt shaker	Plug into a socket	Hammer a nail	Insert straw into cup with lid	Turning bolt with nut driver	Extend and press syringe	Cut Paper

Fig. 1. The ten tasks of the hand-in-hand and autonomous tracks used during the IROS 2016 Grasping and Manipulation Competition .

The hand-in-hand track enabled teams to compete based on the mechanical characteristics of their hand designs without the added requirements of an integrated robot system. It consisted of two stages. The first stage was pick and place where ten objects were removed from a shopping basket and placed within

[1] Certain commercial entities and items are identified in this paper to foster understanding. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

Fig. 2. Task and bin-picking objects used in the IROS 2016 Grasping and Manipulation Competition .

an identified area on a table top. The second stage was manipulation where a set of ten predefined tasks were performed. This track was carried out using the assistance of a human volunteer to support the hand through both stages as shown in Figure 3.

The fully autonomous track required a complete robot system containing hand, arm, and perception components to accomplish the same two stages as shown in Figure 3 using a locate object, plan, and execute grasp approach to solving the problem. The tasks contained within this track also tested a robot system's manipulation capabilities.

The simulation track, using the Kris Locomotion and Manipulation Planning Toolbox (Klampt) [4], consisted of two stages (Figure 4). The first stage was pick and place where the task was to extract as many objects as possible from a cluttered shelf and place them randomly in a bin. The second stage was bin picking where a simulated grasping and manipulation system lifted as many balls as possible from a bin and deposited them into a second bin.

Within each track, total scores earned in both stages were accumulated and used to rank performance. The time to complete each stage was also tracked to determine a winner in the case of a tie, where the contestant using the least amount of time had the advantage.

3 Competitor Feedback

Upon completion of the IROS 2016 RGMC competition, several of the organizers sponsored an interactive feedback session with competitors as an opportunity for both the administrators and teams to discuss and to identify opportunities for improvement. This one-hour feedback session was intentionally held following

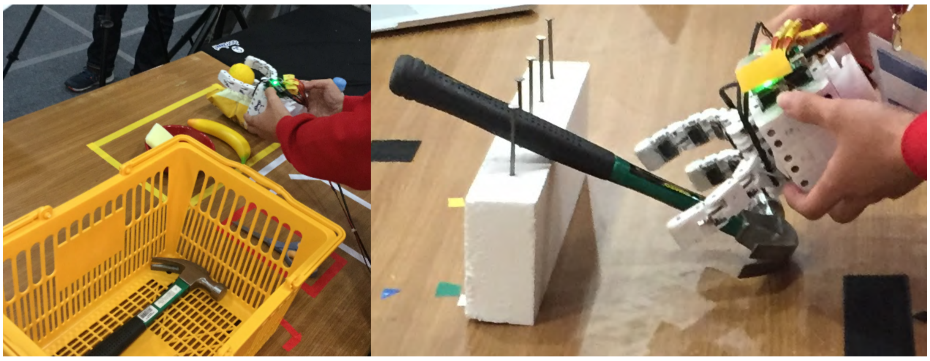

Fig. 3. Competitor hand design being used during the hand-in-hand track: pick and place (left) and hammer a nail task (right).

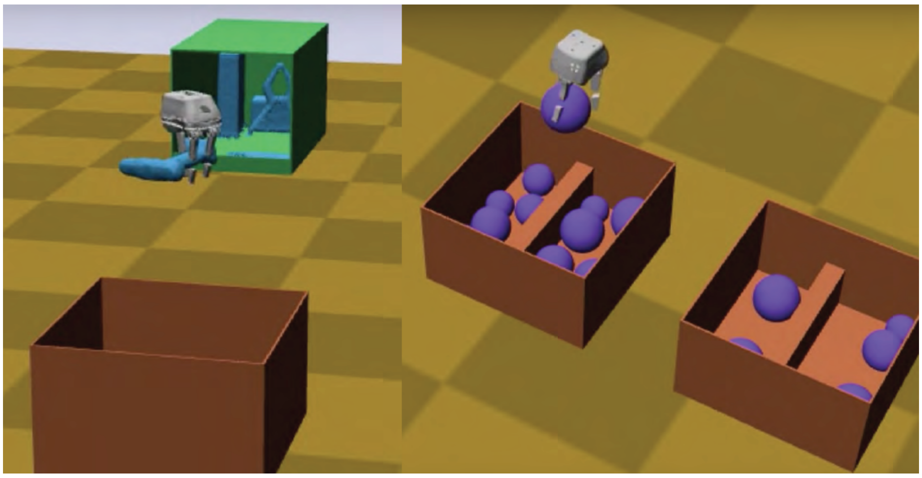

Fig. 4. Simulation task showing the extraction of as many objects as possible from a cluttered shelf and subsequent placement in a bin (left) and pick and placement of as many balls as possible from one bin to another (right).

the computation of final scores to ensure that competitors were candid about their competition experience.

The hand-in-hand track of the competition was developed to evaluate robotic hand hardware designs without the need for an integrated robotic system, eliminating the need of a robotic arm for part manipulation and a perception solution for part localization. This track relied on the use of volunteers with non-technical backgrounds to operate the robotic hand during the associated tracks based on instruction that was automatically generated by the team's computer with no input from competitors via audio or teleoperation. Contrary to the fact that each team was given a period of time to coach volunteers on the use of their hand at the competition prior to each track, some contestants were under the impression that this coaching was limited to an automated instructional mechanism such as a video or software interface and that interaction with volunteers prior to the competition was not permitted. While this did not seem to affect the competition, some contestants felt that the time spent on the details of the instructional video, details that could have been resolved through coaching, detracted from robot system development time. Others found it difficult to give volunteers the exact procedures for grasping objects using the coaching method. It was discussed that the use of inexperienced and disparate hand-in-hand operators leads to subjective based inconsistencies of results across the systems under test. A suggested solution to operator inconsistency was to allow each team to provide an expert operator of the hand which assumes that manual operation is optimized and test results more closely track the performance of the hand hardware. In such a test scenario, a robotic hand designed for use in human-robot interaction applications such as prosthetics may inherently outperform hands designed for autonomous operation when integrated as a robot system. It was noted that, despite the issues associated with the hand-in-hand track of the competition, there was a definite need for the competition to better benchmark capabilities of robotic hands without the need for an autonomous robotic solution.

The fully autonomous track required integrated robotic systems consisting of arm, end effector, and perception components. Discussions indicated that competition rules regarding autonomy were misleading. The confusion may have stemmed from the choice of wording “fully autonomous” by the organizing committee. Fully autonomous within the robotics community most often implies no human intervention on a robotic operation. In the case of the competition, although the term “fully autonomous” was used, the organizers intended to allow certain degrees of teleoperation and human intervention in order to reduce the difficulty associated with the manipulation component of the competition. It was noted that this first competition was purposely made less challenging by use of these leniencies in order to assess the readiness of technology without discouraging participation in future competitions. One contestant indicated that autonomous tasks should be made strictly autonomous with no teleoperation or human intervention.

Another area of discussion was the use of tools grasped by the robotic hand to acquire an object outside of any task-specific tools defined by the competition.

An example of a task specific tool defined by the competition involved the use of a predefined nut driver to tighten a bolt into a threaded hole. In one instance, a team made use of foam blocks containing an adhesive surface as a tool to be grasped by their robot. Here the strategy involved grasping the block using a conventional gripper and using adhesion to acquire an object that could not be easily grasped (e.g., a bag of potato chips). Such competitor-defined tools were valid based on competition rules which only penalized manual reconfiguration of end effectors. Another discussion regarding the autonomous track pertained to challenges posed by limitations in perception systems. Being color-based, many of the perception systems had difficulty in discriminating between objects in the pick-and-place event because so many of the them happened to be yellow or contain a lot of yellow. Examples include the shopping basket, lemon, banana, sponge, candy wrapper, potato chip bag, and scissor handle. These objects and the basket can be found on the right side of Figure 2.

Both the hand-in-hand and fully autonomous tracks used a shopping basket randomly filled with objects to be grasped. At the start of these tracks, competition administrators randomly filled each basket and the baskets were delivered to each team location. Competitors identified instances where the random distribution of objects presented a disadvantage in comparison to the random distribution presented to other teams. Instances were described where objects could not be grasped because they were located too close to basket walls for the planned grasp pose. Other instances were described where objects prevented access to several other underlying objects. It was suggested that the randomness of object placement should be predetermined and fixed across teams with defined levels of difficulty for each predetermined distribution. It was also suggested that competition tasks could be broken into steps, where if a particular step is unachievable, the contestant could skip this task forfeiting associated points and move on to the next. In summary, the discussion led to the conclusion that fixed data sets would allow more control of increasing levels of difficulty that can equally apply to the data given to all competitors. Additionally, it was noted that the shopping basket used was too small for some end-effector sizes.

With regard to the simulation track, contestants felt that future events should more carefully evaluate available simulation packages to determine which is best suited and the most reliable to support the competition tasks. In addition, it was suggested that in the event of another simulation track that the organizers should attempt to tie the simulation tasks to the real world competition tasks.

More general discussions indicated that more time was needed to complete competition tracks. Regarding future competitions, there were suggestions for additional tracks such as dynamic tracking of objects to be grasped and in-hand manipulation. Other suggestions included that future competitions provide a wider range of tasks to support a broader range of hand designs, that some tasks remain unknown until the competition, and that task instructions be more descriptive. Discussion also indicated that benchmarks are needed for the integrated systems in addition to those for individual hand performance. It was also noted that more logistics support funding was needed to support competitor travel as well as shipping costs for competition equipment.

4 Lessons Learned

The hand-in-hand track proved difficult to coordinate and score primarily due to the use of volunteers. The volunteers were chosen based on their educational background, being that of the non-engineering and science related fields. It was also apparent that there were differences in eye-hand coordination capabilities between volunteers and there was confusion regarding the training process for these individuals. This is in contrast with other evaluation approaches which rely on users who are already trained. One example of this is the evaluation of urban search and rescue teleoperated robots through the use of professional responders who are have taken training on how to operate the robots [5]. It is apparent that this methodology will yield a better evaluation of the robotic system under test, but would be too time consuming to support at a conference-sponsored competition. Other evaluations of search and rescue robots on a standard task set require the robot developer to supply the best operator of the system to conduct the testing. This method would be better suited for the hand-in-hand track, however the organizers have concluded that because of its subjective nature, the hand-in-hand track will not be included in subsequent competitions.

The organizer's observations of this competition with the added discussion from competitors have resulted in the decision that all tracks will be fully autonomous with no allowances for teleoperation or manual intervention of manipulator compliance. Objects included in the pick and place track will be carefully selected to ensure a range of object feature variability, including taking into consideration relative object-to-object and object-to-bin color and texture contrast. In addition, methods will be developed to stage object placement in bins to ensure that all participants compete using the same random bin complexity.

With respect to issues concerning the use of tools in addition to the end-effector, the organizers feel that competitors should be free to use any tooling (i.e., custom, hand tools, suction cups) with the continued provision that any manual reconfiguration of an end-effector across different tasks would result in a new robotic hand and a new score based on the tasks it performs. The organizers feel that a diverse set of objects and tasks will drive competitors towards end-effector designs that are adaptable to many object types and away from multiple customized designs which in the long run will increase development complexity, time, and costs.

To improve the instructions of future competitions, the organizers will consider the use of video to help better explain the caveats associated with the competition rule set. In addition, the competition venue will make tools available for those interested in evaluating the performance of their hands without the need of an autonomous robotic system through the use of the National Institute of Standards and Technology (NIST) robotic hand grasping and manipulation benchmarks for assessing hand characteristics such as grasp strength, slip resistance, grasp cycle time, touch sensitivity, and in-hand manipulation [6] [7]. These benchmarks include a set of physical measurements with supporting test methods and instrumented object artifacts that assess elemental performance of robotic hands through the use of external measurement devices. To address the stated needs for better benchmarking techniques within the competition for supporting performance measures of fully autonomous robotic systems, the organizers will investigate the tasks and measures that promote the use of unbiased evaluation methods to assess how well a robot system performs in a particular application space.

5 Conclusion

The first Robot Grasping and Manipulation Competition, held during IROS 2016, allowed researchers focused on the application of robot systems to compare the performance of hand designs as well as autonomous grasping and manipulation solutions across a common set of tasks. At the time of this publication, the 2nd RHGM-sponsored Grasping and Manipulation Competition has been approved for IROS 2017 in Vancouver, Canada which will build on the successes and make adjustments based on this analysis of the 2016 competition. In general, it was determined that the hand-in-hand track with manual operation is too subjective for the competition space and is omitted in order to allow for more development time to prepare for the 2017 tracks which all require autonomous robot systems. To assess hand designs as stand-alone robotic system components, competitors will be given the opportunity to quantify the basic performance of their hands using a set of grasping and in-hand manipulation benchmarking tools. In addition, the simulation track is also omitted from the 2017 event pending an analysis of simulation tools to identify those that are best-suited and the most reliable to support the competition tasks.

The competition will consist of two tracks: (1) Service Tasks, and (2) Manufacturing Tasks. The manufacturing track will include the added challenge of pick-and-place where parts to be assembled will be randomly located on a kit tray. The goal is to pick the objects up and assemble them per a set of instructions. Emphasis will be placed on designing a random scheme of objects that can be easily reproduced for each team in order to keep the level of difficulty of the random distribution the same across teams. In addition, the design of the distribution will ensure that objects in close proximity are reasonably contrasted for detection with a perception system. The manufacturing track will focus on small parts assembly where tasks will incorporate fastening methods such as threading, snap fits, and gear meshing using standard components such as screws, nuts, washers, gears, and electrical connectors [8]. Components to be assembled will be presented in a structured format to simplify the perception problem. The service track will consist of several daily living tasks (DLTs) similar to the tasks defined in the 2016 competition. The tasks are designed to fit within four levels of difficulty where more points are given the greater the level of difficulty.

All tracks will be designed to be fully-autonomous which means that once the timer starts for a given task or set of tasks, there can be no human intervention. Time is recorded for each task or task set in order to decide the winner in the case of point-based ties. There will be no restrictions for automatic end-effector reconfiguration or the use of a tool held by the end-effector; however, manual reconfiguration or changing of an end-effector across different tasks would be considered as a new robotic hand and a new score based on the tasks it performs following each manual process.

References

1. IEEE RAS TC Robotic Hands Grasping and Manipulation, http://www.rhgm.org, Access Date: 6/30/2017
2. Yale-CMU-Berkeley Object Dataset, http://rll.eecs.berkeley.edu/ycb/, Access Date: 6/30/2017
3. 2015 Amazon Picking Challenge, http://rll.berkeley.edu/amazon_picking_challenge/, Access Date: 6/30/2017
4. Klampt (Kris Locomotion and Manipulation Planning Toolbox), http://motion.pratt.duke.edu/klampt/, Access Date: 6/30/2017
5. Messina et al.,: Statement of Requirements for Urban Search and Rescue Robot Performance Standards, DHS and NIST Preliminary Report, pp. 27, 2005
6. J. Falco, K. Van Wyk: Grasping the Performance: Facilitating Replicable Performance Measures via Benchmarking and Standardized Methodologies, Robotics and Automation Magazine, Vol. 22, No.4, pp. 125-136, December 2015
7. Performance Metrics and Benchmarks to Advance the State of Robotic Grasping Web Portal, https://www.nist.gov/programs-projects/performance-metrics-and-benchmarks-advance-state-robotic-grasping, Access Date: 6/30/2017
8. Performance Metrics and Benchmarks to Advance the State of Robotic Assembly Web Portal, https://www.nist.gov/programs-projects/performance-metrics-and-benchmarks-advance-state-robotic-assembly, Access Date: 6/30/2017