

Guest Editorial: Introduction to the Special Issue on Benchmarking Protocols for Robotic Manipulation

I. INTRODUCTION

Benchmarks are crucial for analyzing the effectiveness of an approach against a common basis, providing a quantitative means for interpreting performance. Carefully designed and widely recognized benchmarks encourage the research community to focus on certain key research challenges, promote competition, foster a climate for novel solutions, and, therefore, contribute dramatically to the advancement of a field. While some robotics-related fields (such as object recognition and segmentation) actively utilize benchmarks, there are essentially no robotic manipulation benchmarks that are widely adopted by the research community despite their highly acknowledged necessity.

Discussions within the robotic manipulation research community via a number of workshops and similar meetings have identified some primary obstacles to the development and adoption of benchmarking procedures in our field, including:

- Lack of communication and agreement between researchers for the standards and characteristics of a benchmark
- Lack of widely utilized data sets that target manipulation research
- Lack of a reputable and central venue to distribute the benchmarks
- Lack of professional rewards to encourage researchers to develop and utilize benchmarks

This special issue seeks to help break some of these barriers by encouraging collaborations among different research groups, encouraging the use of existing data sets, and boosting the visibility and dissemination of benchmarking procedures via a reputable publishing venue. The special issue is dedicated to letters that propose and demonstrate novel and widely useful benchmarking protocols for robotic manipulation research. The letters focus on describing well-defined experimental procedures that are ready to be applied by other researchers in similar topic areas for quantifying performance of research approaches in robotic manipulation and its various sub-fields such as:

- Manipulation planning (e.g., performance of grasp planners)
- Mechanism design (e.g., performance of robotic hands)
- Machine learning (e.g., learning manipulation abilities)

- Cognitive robotics (e.g., task representations)

One of the primary challenges in developing effective benchmarking procedures relates to the balance of specificity versus generality. For instance, high-level system performance metrics (such as that done within the Amazon Picking Challenges) can be used by the widest range of research groups, but tell little about the performance of the specifics of the approaches being used. For instance, was the good or bad performance due to the hardware design, the perception system, or the planning approach? On the opposite side of the spectrum, a very narrowly designed evaluation procedure that specifies the hardware platforms and many software subsystems might be able to speak very specifically to the effectiveness of a grasp planner, for instance, but might not generalize, and would only be used by researchers with that particular combination of subsystems available to them. It is therefore left to the authors of proposed benchmarking procedures to find a suitable middle ground to provide sufficient quantitative evaluation of specific research approaches while enabling as many researchers as possible to implement them.

In addition, authors of submitted letters were highly encouraged to:

- work in collaboration with multiple research groups from similar areas to develop benchmarking procedures (to avoid overfitting to a particular approach and boost the overall impact),
- make use of existing published data sets when possible (e.g., standard objects and models), unless they are specifically inappropriate
- utilize a standard template that summarize the protocol (procedure and constraints) and benchmark (reporting of results)
- provide multimedia files that illustrate and demonstrate the protocol
- report baseline experimental results obtained with authors' own setups.

The guest editors strongly encouraged the prospective authors to submit a 1-2 pages long letter of intent regarding the scope of their benchmark to receive feedback prior to full submission. These letters were also asked to include descriptions regarding the purpose of the benchmark, the overall task protocol, and how performance will be quantified and reported. Among the received letters, it is identified that several different research groups were targeting the same/similar robotic manipulation

problem. In order to avoid duplicated benchmarks and achieve a unified assessment standard via the combined expertise of different research groups, the guest editors established a collaboration mechanism: The researchers were notified that another research group is targeting a similar problem, and were asked, without knowing others' identity, if they wanted to collaborate with them. As a result of this overall process, the majority of the letters in this special issue were written as a collaboration of two or more research groups (some of which had not collaborated before), which, the guest editors believe, prevented overfitting of the benchmarks at large and boosted their impact.

II. GUIDE TO THE SPECIAL ISSUE

A. Benchmarks on Robotic Hand/Gripper Assessment

Two letters in the special issue focus on robotic hand/gripper assessment. In [1] various robotic hand properties are assessed such as grasp strength, grasp cycle time, and repeatability. Whereas, the work [2] focuses on more extreme cases, and provides experimental protocols and assessment metrics to quantify the resilience of robotics hands. These methods identify the limits of a robotic hand by investigating the dynamic loads that make the hand lose its grasp on an object and that might physically damage the hand.

B. Benchmarks on Robotic Grasping

Two robotic grasping benchmarks are provided in this special issue. In [3], the benchmark identifies the limits and capabilities of the robotic system (e.g., workspace limits, payload limits) and allows to normalize the results accordingly. The benchmark in [4] provides a rigorous procedure to assess the performance of grasp planning algorithms, while minimizing the effects of other elements in the grasping pipeline. While [3], [4] are focused on grasping from table top, [5] provides a benchmark for a more industrial application, i.e., bin-picking, considering pick-and-place of fruit and vegetables.

C. Benchmarks With Tests

Three letters adapt existing manipulation tests to the robotics domain. The work in [6] provides experimental protocols for robotic implementation of the Box and Blocks test, which is a common assessment procedure in rehabilitation. The work in [7] utilizes Rubik's cube to assess the system's precision manipulation capabilities as well as its ability to conduct sequential manipulation. The work in [8] presents a benchmark for robotic assembly using a standard assembly board (which is available to purchase on demand) that includes tasks such as insertion, cable routing and screwing.

D. Benchmarks on Deformable Object Manipulation

Two letters in the special issue focus on deformable object manipulation, both of which require bimanual operation. In [9], a cloth manipulation benchmark is presented which targets folding, spreading and dressing. The work in [10] focuses on

intricate bimanual manipulation tasks, i.e., watchmaking and belt assembly.

E. Other Benchmarks

Four other letters in the special issue focus on more specific robotic manipulation tasks. In [11], a benchmark is provided for assessing in-hand manipulation accuracy. In [12], a benchmark is provided for assessing aerial manipulation capabilities, e.g., aerial grasping, force control etc. The work in [13] focuses on dynamic human-to-robot handovers, and provides experimental procedures for transferring cups with different fullness levels and transparencies. The work in [14] provides a framework to assess the performance of robotic manipulation simulators via a real-world dataset.

BERK CALLI, *Guest Editor*
Worcester Polytechnic Institute
Worcester, MA 01609, USA

AARON DOLLAR, *Guest Editor*
Yale University
New Haven, CT 06520, USA

MAXIMO A. ROA, *Guest Editor*
German Aerospace Center - DLR
Weßling, 82234, Germany

SIDDHARTHA SRINIVASA, *Guest Editor*
University of Washington
Seattle, WA 98195, USA

YU SUN, *Guest Editor*
University of South Florida
Tampa, FL 33620, USA

REFERENCES

- [1] J. Falco *et al.*, "Benchmarking protocols for evaluating grasp strength, grasp cycle time, finger strength, and finger repeatability of robot end-effectors," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 644–651, Apr. 2020, doi: [10.1109/LRA.2020.2964164](https://doi.org/10.1109/LRA.2020.2964164).
- [2] F. Negrello *et al.*, "Benchmarking hand and grasp resilience to dynamic loads," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 1780–1787, 2020, doi: [10.1109/LRA.2020.2969180](https://doi.org/10.1109/LRA.2020.2969180).
- [3] F. Bottarel, G. Vezzani, U. Pattacini, and L. Natale, "GRASPA 1.0: GRASPA is a robot arm grasping performance benchmArk," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 836–843, 2020, doi: [10.1109/LRA.2020.2965865](https://doi.org/10.1109/LRA.2020.2965865).
- [4] Y. Bekiroglu *et al.*, "Benchmarking protocol for grasp planning algorithms," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 315–322, Apr. 2020, doi: [10.1109/LRA.2019.2956411](https://doi.org/10.1109/LRA.2019.2956411).

- [5] H. Mnyusiwalla *et al.*, “A bin-picking benchmark for systematic evaluation of robotic Pick-and-Place systems,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 1389–1396, Apr. 2020, doi: [10.1109/LRA.2020.2965076](https://doi.org/10.1109/LRA.2020.2965076).
- [6] A. S. Morgan *et al.*, “Benchmarking cluttered robot pick-and-place manipulation with the box and blocks test,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 454–461, Apr. 2020, doi: [10.1109/LRA.2019.2961053](https://doi.org/10.1109/LRA.2019.2961053).
- [7] B. Yang, P. E. Lancaster, S. S. Srinivasa, and J. R. Smith, “Benchmarking robot manipulation with the rubik’s cube,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 2094–2099, Apr. 2020, doi: [10.1109/LRA.2020.2969912](https://doi.org/10.1109/LRA.2020.2969912).
- [8] K. Kimble *et al.*, “Benchmarking protocols for evaluating small parts robotic assembly systems,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 883–889, Apr. 2020, doi: [10.1109/LRA.2020.2965869](https://doi.org/10.1109/LRA.2020.2965869).
- [9] I. Garcia-Camacho *et al.*, “Benchmarking bimanual cloth manipulation,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 1111–1118, Apr. 2020, doi: [10.1109/LRA.2020.2965891](https://doi.org/10.1109/LRA.2020.2965891).
- [10] K. Chatzilygeroudis *et al.*, “Benchmark for bimanual robotic manipulation of semi-deformable objects,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 2443–2450, Apr. 2020, doi: [10.1109/LRA.2020.2972837](https://doi.org/10.1109/LRA.2020.2972837).
- [11] S. Cruciani, B. Sundaralingam, K. Hang, V. Kumar, T. Hermans, and D. Kragic, “Benchmarking in-hand manipulation,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 588–595, Apr. 2020, doi: [10.1109/LRA.2020.2964160](https://doi.org/10.1109/LRA.2020.2964160).
- [12] A. Suarez, V. M. Vega, M. Fernandez, G. Heredia, and A. Ollero, “Benchmarks for aerial manipulation,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 2650–2657, Apr. 2020, doi: [10.1109/LRA.2020.2972870](https://doi.org/10.1109/LRA.2020.2972870).
- [13] R. Sanchez-Matilla *et al.*, “Benchmark for Human-to-Robot handovers of unseen containers with unknown filling,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 1642–1649, Apr. 2020, doi: [10.1109/LRA.2020.2969200](https://doi.org/10.1109/LRA.2020.2969200).
- [14] J. Collins, J. McVicar, D. Wedlock, R. Brown, D. Howard, and J. Leitner, “Benchmarking simulated robotic manipulation through a real world dataset,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 1, pp. 250–257, Jan. 2020, doi: [10.1109/LRA.2019.2953663](https://doi.org/10.1109/LRA.2019.2953663).