

CloudGripper: An Open Source Cloud Robotics Testbed for Robotic Manipulation Research, Benchmarking and Data Collection at Scale

Muhammad Zahid and Florian T. Pokorny

Abstract—We present CloudGripper, an open source cloud robotics testbed, consisting of a scalable, space and cost-efficient design constructed as a rack of 32 small robot arm work cells. Each robot work cell is fully enclosed and features individual lighting, a low-cost Cartesian robot arm with an attached rotatable parallel jaw gripper and a dual camera setup for experimentation. The system design is focused on continuous operation and features a 10 Gbit/s network connectivity allowing for high throughput remote-controlled experimentation and data collection for robotic manipulation. Furthermore, CloudGripper is intended to form a community testbed to study the challenges of large scale machine learning and cloud and edge-computing in the context of robotic manipulation. In this work, we describe the mechanical design of the system, its initial software stack and evaluate the repeatability of motions executed by the proposed robot arm design. A local network API throughput and latency analysis is also provided. CloudGripper-Rope-100, a dataset of more than a hundred hours of randomized rope pushing interactions and approximately 4 million camera images is collected and serves as a proof of concept demonstrating data collection capabilities. A project website with more information is available at <https://cloudgripper.org>.

I. INTRODUCTION

With recent progress in deep learning models for robotics [1], [2], [3], [4] and the use of large transformer architectures [5], the robotics community is increasingly faced with a training data acquisition bottleneck in exploring the full potential of these models. While application fields such as computer vision and natural language processing benefited from the availability of large scale image datasets [6] and internet-scale text datasets [7], robotic manipulation research is in comparison still data starved with individual authors typically relying on real-world-executions of only a few hours of robotic experiments. In addition, robotic manipulation work cell and robot setups as well as objects used for research still vary widely between research groups making benchmarking and replication of research challenging.

In this work, we present CloudGripper (see Fig. 1), the first rack-based and fully scalable cloud robotics system designed to significantly push the boundaries of parallelism and experimental data throughput for robotic manipulation and to serve as a community testbed for large-scale robotic manipulation research and dataset collection at multi terabyte scale. Secondly, CloudGripper offers a first reference implementation of a cloud robotics system of 32 robotic arms using which distributed robotic manipulation and machine learning algorithms can be evaluated and tested.

The authors are with the School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology mzmi@kth.se; fpokorny@kth.se.




Fig. 1: Top left: A CloudGripper rack of 32 robotic arms with front-panels removed for visibility, top right: individual robot cell; middle and bottom left: internal camera images of the system with white floor inlay, QR-marker and 3D-printed cubes and rope manipulation setup; middle and bottom right image: internal top and bottom camera images with transparent base plate to capture multi-view images of a rope manipulation scenario.

The core contributions of the present paper are:

- 1) We present the system and hardware design of CloudGripper, a first practical and remotely accessible cloud robotics test bed for robotic manipulation with a space-efficient and scalable rack-based design.
- 2) We validate our design by means of repeatability tests and provide a throughput analysis based on a single CloudGripper rack of 32 robots.
- 3) We present CloudGripper-Rope-100, a dataset consisting of more than a hundred hours of randomized rope manipulations generated using CloudGripper.
- 4) We discuss sample use cases of the system and present our vision for an open source community-based de-

velopment approach of CloudGripper and its software ecosystem with participation of the international community going forward.

II. RELATED WORK

Testbeds in Robotics: The need to improve repeatability of experiments and to find common benchmarking approaches has long been recognized by the robotic manipulation community, with multiple datasets, challenges and workshops over the years focussing on this topic [8], [9], [10], [11], [12]. The time limited nature of competitions such as the Amazon Picking Challenge [9] and the lack of shared identical experimental setups that comprise not just object datasets such as [8], but complete work cells, software stacks and configurations however still inhibits easy replication and comparison of results. In robotics research beyond robotic manipulation more broadly, several testbeds have been proposed to serve as common experimentation platforms [13]. Perhaps most well-known is the Robotarium testbed [14] featuring remote-controlled mobile robots. Similar to our approach, this testbed in particular also opted for low-cost components which ensures the scalability and maintainability of the system. As far as we are aware, CloudGripper however constitutes the first large scale testbed with a particular focus on remote robotic manipulation at scale.

Cloud Robotics: CloudGripper is designed as a reference cloud robotics research platform. The term cloud robotics was coined by James Kuffner in 2010 [15], and proposes the use of cloud computing, big data, collaborative learning and crowd sourcing (including open source development) to approach some of the most challenging problems in robotics. Early work focusing on networked and tele-operated robotics, such as the tele-garden project [16] and pioneering efforts such as the RoboEarth research project [17] have laid the foundations for cloud robotics and related commercial applications, such as the deployment of autonomous vehicles that are now becoming a reality. Recently, the development of ROS 2.0 [18] and efforts such as FogROS2 [19] furthermore continue to push the boundaries of how an edge/cloud approach to robotics can be implemented, but no public robotic manipulation-focused deployments that fully embrace the cloud robotics paradigm are currently publicly available for research end-users to the best of our knowledge.

Scale of Experimentation in Robotics Research: State of the art robotic manipulation research labs today require a significant capital, space and staffing investment and are typically equipped with a small number of robotic single and dual arm systems. Systems such as the ABB YuMi [20] dual arm robot and the Universal Robots UR5e [21] provide current robotic platform examples in particular. Since large-scale deployments of such systems are currently beyond the budgets of typical research groups, data collection and experimentation on these systems has also remained relatively small in scale in comparison to web-scale datasets being investigated in large language models [7], for example. Furthermore, data collection remains a labor intensive process

even on a single robot. Part of the reason for this is that to ensure safety of both equipment and personnel and due to the need to reset experiments, researchers today still commonly have to physically watch over ongoing experiments on a single robotic system at a time and manually supervise robotic experiments in real-time which is in stark contrast to today's machine learning work-flows, where researchers regularly schedule multiple compute jobs on remote compute clusters in parallel. Due to these obstacles, progress on truly large-scale real-world experimentation in robotic manipulation has been rare and limited to corporate-supported efforts such as Google's pioneering work [22], Amazon's ARMBench [23] and OpenAI's efforts such as [24]. Works that some of the authors of this work also co-authored, such as [25] and Dex-Net 1.0 [26] have shown the benefits to generalizability and performance when utilizing larger and larger training datasets - even if based on analytic or simulated scenarios, but scaling laws and exact training data size requirements are still poorly understood. Simulation and sim-to-real transfer [27] has been proposed as a solution to the lack of large scale real-world training data in particular, but significant challenges remain since simulating complex physical interactions remains challenging. This is in particular the case for the manipulation of deformable objects such as cloth and wires as well as fragile objects. In these cases, simulation still cannot fully match the richness of data collected from real-world robotic experimentation.

III. SYSTEM DESIGN

CloudGripper is designed as a space-optimized system similar to rack-based servers in a data center. Each rack, as displayed in Fig. 1, is constructed from extruded aluminium frames and carries 32 individual robot arms and work cells that are stacked to ensure space-efficiency. In the following, we discuss the mechanical and software design of the system, starting with the robot arm and work cell setup used in the system.

A. Integrated Robot Arm Work Cells

Work Cell design: Each CloudGripper robot is housed in a fully enclosed individual work cell featuring 5V and 12V power connections, a 1 Gbit/s Ethernet network and a USB-C cable connection. Each work cell has a dimension of $274mm \times 356mm \times 400mm$ and features individual LED lights and a detachable front-panel. For maintenance, the front panel can be removed and custom work cell components, such as 3D-printed parts, for example for a pick-and-place task, can be easily inserted or replaced.

Two RGB cameras are provided in each work cell - a Raspberry Pi 3 wide angle camera providing a top-view of the work environment and a 140 degree wide-angle camera providing a bottom view through a transparent plexiglass floor plate (see Fig. 1). The transparent floor plate is designed for experiments where an additional camera view from below can be beneficial, such as planar pushing tasks. Optionally, a non-transparent inlay can be inserted as shown in the bottom-left of Fig. 1, in which case the bottom camera is obstructed




Fig. 2: Rendering of an individual robot cell




Fig. 3: Gripper design with inbuilt servo motors

and not used. The plexi-glass plate also features screw-holes to fixate custom work environments that may be 3D-printed. Each robot cell is constructed predominantly from extruded $1 \times 1\text{ cm}$ profile aluminum beams and 3D printed parts. CAD designs are available via the project website cloudgripper.org.

Per-Robot Microcontrollers and Compute: For image processing and networking, each robot features a Raspberry Pi 4B which is powered through a USB-C connection. The Raspberry Pi 4B utilizes a Quadcore Cortex-A72 processor with 1.8GHz and is connected to a Teensy 4.1 microcontroller [28] based on the ARM Cortex-M7 processor. The microcontroller serves as a low level interface to the robot's sensors and actuators. The Teensy 4.1 architecture was selected due to its processing power and ability to provide high frequency control for the motors.

Actuation and Kinematic Structure: Each robot has a total of 5 degrees of freedom (x, y, z, r, d). A Cartesian base that uses two Nema 17 48Ncm stepper motors with 4000 CPR optical encoders controls the planar (x, y)-position using a CoreXY [29] motion platform which employs belts as displayed in Fig. 2. The CoreXY platform has in particular been used extensively for low-cost 3D-printer designs [29] where it has proven its effectiveness. The remaining three degrees of freedom (gripper Z-axis z , gripper rotation r and parallel jaw gripper opening d) are actuated by MG90S micro servos as shown in Fig. 3. A prismatic joint, driven by one of the servo motors accomplishes the z-axis motion in particular while a revolute joint is used to rotate the gripper around its vertical axis and the gripper is opened and closed using a further servo motor as shown in Fig. 3. The main parts of the gripper are 3D-printed from PETG while the gears are lasercut from plexiglas. The low-level controller is implemented on the Teensy 4.1 microcontroller. The current implementation allows arbitrary piecewise linear movements

in XY direction as well as separate rotation, Z-motions and gripper opening/closing and employs a preset velocity profile along a given trajectory. Calibration switches are added to re-calibrate the stepper motors and encoders.

Servo actuation and collision detection: Micro servos provide a suitable choice for the chosen small form factor gripper, but require calibration to ensure uniform performance. Using reference positions from a potentiometer, each servo was therefore calibrated to achieve consistent performance. Since the chosen servos operate via pulse width signals and do not offer feedback on their position, we furthermore integrated current sensors allowing for the detection of collisions. This strategy was used to enable the prevention of motor overheating and potential damage to the robot.

B. Software-Stack and Component Communication

CloudGripper's current software stack is focused on simplicity and will be extended significantly in future work. In its current form, each robot's Raspberry Pi 4 runs Raspberry Pi OS and administrators can access each robot via SSH for maintenance. Furthermore, a lean software stack was developed for the initial version of CloudGripper presented here. Each Raspberry runs a REST API server with token-based authentication that is implemented using the python Flask library [30] and allows remote end-users to request information about a robot's state, to obtain camera images at up to 30fps and to send motion commands that are then forwarded by the API to the microcontroller for execution. Additionally, remote end-users can request livestreams to view the robot in operation. The use of a standard REST API approach proposed here also enables simple development of applications such as teleoperation web-interfaces with buttons for remote control. Each of the 32 robots in a given CloudGripper rack features an individual 1 Gbit/s Ethernet

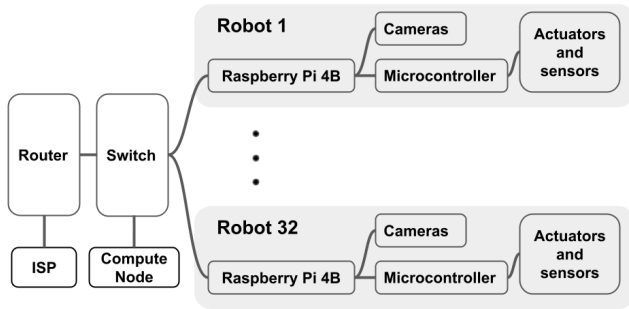


Fig. 4: CloudGripper rack networking with 32 robots connected to a switch which enables further communication to the internet via a router as well as access to a local area network server for local computation.

cable connection to a 10 Gbit/s switch that is then connected to both a router with a 10 Gbit/s internet connection and a 10 Gbit/s connection to a local storage and compute node, see Fig. 4 for details on CloudGripper’s networking layout.

C. CloudGripper as an Open Source Ecosystem

CloudGripper is released as an Open Source project with hardware and software released under permissive licensing. Beyond software and hardware, it is however our hope that the project will furthermore evolve into a living testbed with many international users where novel algorithms can be tested against open reference-implementations. The project website cloudgripper.org will serve as a central resource as the project develops and we intend to also host a growing selection of datasets there, starting with the CloudGripper-Rope-100 dataset which is described in this paper.


IV. SYSTEM DESIGN VALIDATION

To validate the design of CloudGripper, we present a sequence of experimental evaluations testing both individual robot repeatability and behaviour under extended load.

A. Repeatability analysis


For this evaluation, the repeatability of the robots’ motions was measured on a sample of 23 robots in a CloudGripper rack using a digital dial indicator in a fixed sample location for Z-axis measurements, representing the vertical displacement by the Z-axis servo motor. Furthermore XY-axis measurements, capturing the lateral displacement caused by the stepper motors were collected. For each, the robot was moved linearly between a set of five uniformly sampled waypoints in Z and XY dimension respectively before returning to its initial configuration where data from the dial indicator was recorded. This test was performed 10 times for each sample location and robot.

The distribution of XY-axis measurements is shown in Fig. 5 while the distribution of Z-axis measurements is displayed in Fig. 6. For the XY-axis measurements, the range of deviation in position was between $-0.084mm$ and $0.070mm$ and the standard deviation across all tests was approximately $0.027mm$, highlighting the consistency of the robot’s XY-axis movements. For the Z-axis measurements, the range of



Individual Robots Tested

Fig. 5: Box-plots of repeatability measurements in XY axis for 23 CloudGripper robots. Observed measurement ranges are displayed together with lower and upper quartiles which are indicated by blue boxes. The median of the measurements for each robot is shown as an orange line.



Individual Robots Tested

Fig. 6: Box-plots of repeatability measurements in Z axis for 23 CloudGripper robots. Observed measurement ranges are displayed together with lower and upper quartiles which are indicated by blue boxes. The median of the measurements for each robot is shown as an orange line.

deviation was higher and between $-0.419mm$ and $0.376mm$ with a standard deviation of approximately $0.109mm$. While still relatively minor in absolute scale given the low cost nature of the system components, some of the tested robots exhibited variability in repeatability in Z-axis measurements in particular. We hypothesize that variability of the servo motors used in these robots as well as variations in the 3D-printed linear Z-axis gear may be contributing factors to these differences. Despite these variations, we conclude that the system has a suitable repeatability to offer a viable option for a wide range of practical robotic experimentation tasks.

B. Network Stress-Test

One of the core objectives of our design is to facilitate scalability of experimentation by allowing users to control multiple robots simultaneously via API requests. We performed a network stress test where the local compute server, connected via LAN to CloudGripper, sent API requests for

still images from both cameras of a given robot at 10 fps. Starting with a single robot, we incrementally add an additional robot to these concurrent queries with one new thread per query process until all 32 robots are receiving concurrent api requests. As shown in Fig. 7, CPU and memory utilization on the querying server steadily increase with each new query thread.

As expected, network bandwidth usage increases linearly to approximately 80 megabytes per second, utilizing only less than a tenth of the theoretical maximal network bandwidth of 10 Gbit/s and thus illustrating the potential of the system to handle high concurrent loads. This result also indicates that a 10 Gbit/s network is likely able to accommodate a significantly larger number of robots.

Additionally, the maximal latency for all requests and robots remains below 100ms during these tests with a 95th percentile latency of approximately 70ms. We note that these performance metrics can likely be improved significantly in future since both the API server and requesting code used during these tests are based on simple python scripts written without any particular focus on optimization.

V. CLOUDGRIPPER-ROPE-100 DATASET

To demonstrate data collection capabilities and to stress-test the system with continuous use, we collected a dataset of randomized rope manipulations on 19 robots that were available for data collection.

Rope manipulation was selected as an example of a challenging manipulation problem due the inherent complexities of state estimation and prediction when manipulating deformable objects [31].

To collect data, 19 robots were fitted with a custom work environment that featured a rope affixed to one side of the workspace as seen in Fig. 8. The robot gripper is lowered to the glass plate and remains in this Z-position during the data collection throughout. Only the XY-coordinate is modified to execute pushing of the rope.

Data generation is initialized with the rope in a straightened position as seen in the top left of Fig. 8. On each robot candidate planar linear robot pushing motions are sampled uniformly. Using a simple color segmentation of the rope, we reject candidate linear XY-motion trajectories that do not result in a predicted collision with the rope. Candidate motions that result in intersection are executed, resulting in up to 10 consecutive random perturbations of the rope by linear pushing actions per episode. The rope is approximately reset using an automated resetting procedure that attempts to straighten the rope again. During the rope pushing process, camera images are captured at 10 fps and linear pushing actions are executed at a frequency of approximately 0.5Hz.

The CloudGripper-Rope-100 dataset contains approximately 110 hours of robot motion data and approximately 4 million individual camera images of pushing episodes. The combined dataset requires approximately half a terabyte of storage space in uncompressed form and contains the following components:




Fig. 7: CloudGripper network stress test evaluation conducted by sending concurrent API camera image requests. Here, the number of concurrent robots queried increases at each vertical dashed line by one, starting with one robot on the left and terminating with all 32 robots in the rightmost part of the plot. We observe near-linear increases in CPU and memory usage on the querying server while API request latency remains stable and only a small fraction of the total network bandwidth of 10 Gbit/s is utilized even when all robots are queried simultaneously.




Fig. 8: Images from the CloudGripper-Rope-100 manipulation dataset. The top two rows display sample top camera views while bottom rows display bottom camera views. In all cases a rope segmentation using Segment Anything [32] is applied to highlight the rope in the figures.

- **Metadata:** Robot identifiers, description of the robot environment, total duration of data collection, dataset size, cumulative number motion commands sent and camera calibration matrices are provided.

- **Robot Actions:** The time stamped API motion commands for each robot are stored to allow for replay of episodes on both real and simulated environments.
- **Camera Images:** Both camera views are stored as color images at 10 fps per robot with a resolution of 1280×720 pixels for the top camera and 640×480 pixels for the bottom camera.

VI. DISCUSSION OF FUTURE USE CASES

Training Datasets: The generation of rich large scale training datasets is an obvious use case for CloudGripper that we have started to demonstrate in the presented rope manipulation dataset. Future datasets focused on specific object types or task settings including cloth folding, pick-and-place tasks as well as simplified assembly tasks could be considered. The collection of human demonstration data by means of teleoperation by multiple remote users could also provide a valuable direction.

Iterative Algorithm Development: CloudGripper presents an opportunity for researchers to test robotics algorithms in a real-world context. A key question going forward will be to identify robot workspace setups and interaction scenarios that are of most interest to potential collaborators.

Benchmarking: By offering a standardized robot environment, CloudGripper may be used as a benchmarking and reference implementation platform on which common perception, planning and manipulation algorithms can be compared.

Transfer Learning: The availability of several near-identical robot cells offers a research opportunity to reason about transfer learning between tasks and physical workspace setups in a parallelized manner.

VII. SYSTEM LIMITATIONS AND FUTURE WORK

The CloudGripper system proposed here provides a first reference design of a large scale robotic manipulation cloud robotics research platform. The design is optimized for cost-efficiency, rapid prototyping and scalability instead of featuring more complex 6 or 7 degree of freedom robot arms made from precision machined components. As a result, the current robotic arm system is of course limited in the type of manipulations it can perform. We believe however that even in its current form, the unique scale offered by this approach will unlock novel insights in data-driven robotics research that were simply not accessible previously. A further aspect is that CloudGripper enables researchers to test novel cloud and edge computing strategies for robotics which we shall explore further in future. We also intend to provide ROS2 [18] wrappers and a simulation environment for Sim-To-Real experimentation in the near future.

VIII. ACKNOWLEDGEMENTS

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

REFERENCES

- [1] M. Shridhar, L. Manuelli, and D. Fox, "Perceiver-actor: A multi-task transformer for robotic manipulation," in *Conference on Robot Learning*. PMLR, 2023, pp. 785–799.
- [2] Y. Han, R. Batra, N. Boyd, T. Zhao, Y. She, S. Hutchinson, and Y. Zhao, "Learning generalizable vision-tactile robotic grasping strategy for deformable objects via transformer," *arXiv preprint arXiv:2112.06374*, 2021.
- [3] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich, "Rt-1: Robotics transformer for real-world control at scale," 2022.
- [4] A. Gupta, L. Fan, S. Ganguli, and L. Fei-Fei, "Metamorph: Learning universal controllers with transformers," *arXiv preprint arXiv:2203.11931*, 2022.
- [5] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [6] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. IEEE, 2009, pp. 248–255.
- [7] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "Gpt-4 technical report," 2023.
- [8] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The ycb object and model set: Towards common benchmarks for manipulation research," in *2015 international conference on advanced robotics (ICAR)*. IEEE, 2015, pp. 510–517.
- [9] C. Eppner, S. Höfer, R. Jonschkowski, R. Martín-Martín, A. Sieverling, V. Wall, and O. Brock, "Lessons from the amazon picking challenge: Four aspects of building robotic systems," in *Robotics: science and systems*, 2016, pp. 4831–4835.
- [10] "Ieee icra 2021 workshop on cloud-based competitions and benchmarks for robotic manipulation and grasping," 2021. [Online]. Available: <https://sites.google.com/view/icra2021-workshop/home>
- [11] Z. Liu, W. Liu, Y. Qin, F. Xiang, M. Gou, S. Xin, M. A. Roa, B. Calli, H. Su, Y. Sun, and P. Tan, "Ocrto: A cloud-based competition and benchmark for robotic grasping and manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 486–493, 2022.
- [12] S. Bauer, M. Wüthrich, F. Widmaier, A. Buchholz, S. Stark, A. Goyal, T. Steinbrenner, J. Akpo, S. Joshi, V. Berenz, V. Agrawal, N. Funk, J. Urain De Jesus, J. Peters, J. Watson, C. Chen, K. Srinivasan, J. Zhang, J. Zhang, M. Walter, R. Madan, T. Yoneda, D. Yarats, A. Allshire, E. Gordon, T. Bhattacharjee, S. Srinivasa, A. Garg, T. Maeda, H. Sikchi, J. Wang, Q. Yao, S. Yang, R. McCarthy, F. Sanchez, Q. Wang, D. Bulens, K. McGuinness, N. O'Connor, R. Stephen, and B. Schölkopf, "Real robot challenge: A robotics competition in the cloud," in *Proceedings of the NeurIPS 2021 Competitions and Demonstrations Track*, ser. Proceedings of Machine Learning Research, D. Kiela, M. Ciccone, and B. Caputo, Eds., vol. 176. PMLR, 06–14 Dec 2022, pp. 190–204.
- [13] A. Jiménez-González, J. R. Martínez-De Dios, and A. Ollero, "Testbeds for ubiquitous robotics: A survey," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1487–1501, 2013.
- [14] D. Pickem, P. Glotfelter, L. Wang, M. Mote, A. Ames, E. Feron, and M. Egerstedt, "The robotarium: A remotely accessible swarm robotics research testbed," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1699–1706.
- [15] B. Kehoe, S. Patil, P. Abbeel, and K. Goldberg, "A survey of research on cloud robotics and automation," *IEEE Transactions on automation science and engineering*, vol. 12, no. 2, pp. 398–409, 2015.
- [16] K. Goldberg, *The Robot in the Garden: Telerobotics and Telepistemology in the Age of the Internet*. MIT Press, 2001.
- [17] M. Waibel, M. Beetz, J. Civera, R. D'Andrea, J. Elfring, D. Gálvez-López, K. Häussermann, R. Janssen, J. Montiel, A. Perzylo, B. Schießle, M. Tenorth, O. Zweigle, and R. V. De Molengraaf, "Roboearth," *IEEE Robotics Automation Magazine*, vol. 18, no. 2, pp. 69–82, 2011.
- [18] S. Macenski, T. Foote, B. Gerkey, C. Lalancette, and W. Woodall, "Robot operating system 2: Design, architecture, and uses in the wild," *Science Robotics*, vol. 7, no. 66, p. eabm6074, 2022. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abm6074>
- [19] J. Ichnowski, K. Chen, K. Dharmarajan, S. Adebola, M. Danielczuk, V. Mayoral-Vilches, N. Jha, H. Zhan, E. Llontop, D. Xu, C. Buscaron, J. Kubiawicz, I. Stoica, J. Gonzalez, and K. Goldberg, "Fogros2: An adaptive platform for cloud and fog robotics using ros 2," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 5493–5500.
- [20] "Dual-arm yumi@ - irb 14000." [Online]. Available: <https://new.abb.com/products/robotics/robots/collaborative-robots/yumi/dual-arm>
- [21] "Ur5e lightweight, versatile cobot." [Online]. Available: <https://www.universal-robots.com/products/ur5-robot/>
- [22] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *The International journal of robotics research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [23] C. Mitash, F. Wang, S. Lu, V. Terhuja, T. Garaas, F. Polido, and M. Nambi, "Armbench: An object-centric benchmark dataset for robotic manipulation," *arXiv preprint arXiv:2303.16382*, 2023.
- [24] M. Andrychowicz, B. Baker, M. Chociej, R. Józefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Wang, and W. Zaremba, "Learning dexterous in-hand manipulation," *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.
- [25] F. T. Pokorny, K. Hang, and D. Kragic, "Grasp moduli spaces." in *Robotics: Science and Systems*, 2013.
- [26] J. Mahler, F. T. Pokorny, B. Hou, M. Roderick, M. Laskey, M. Aubry, K. Kohlhoff, T. Kröger, J. Kuffner, and K. Goldberg, "Dex-net 1.0: A cloud-based network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 1957–1964.
- [27] R. Jeong, Y. Aytar, D. Khosid, Y. Zhou, J. Kay, T. Lampe, K. Bousmalis, and F. Nori, "Self-supervised sim-to-real adaptation for visual robotic manipulation," in *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 2718–2724.
- [28] "Teensy@ usb development board." [Online]. Available: <https://www.pjrc.com/teensy/>
- [29] "Cartesian motion platform," 2012. [Online]. Available: <https://corexy.com/>
- [30] M. Grinberg, *Flask web development: developing web applications with python*. O'Reilly Media, Inc., 2018.
- [31] H. Yin, A. Varava, and D. Kragic, "Modeling, learning, perception, and control methods for deformable object manipulation," *Science Robotics*, vol. 6, no. 54, 2021.
- [32] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, "Segment anything," 2023.