# Adaptive Mobile Manipulation for Articulated Objects In the Open World

Anonymous CVPR submission

Paper ID

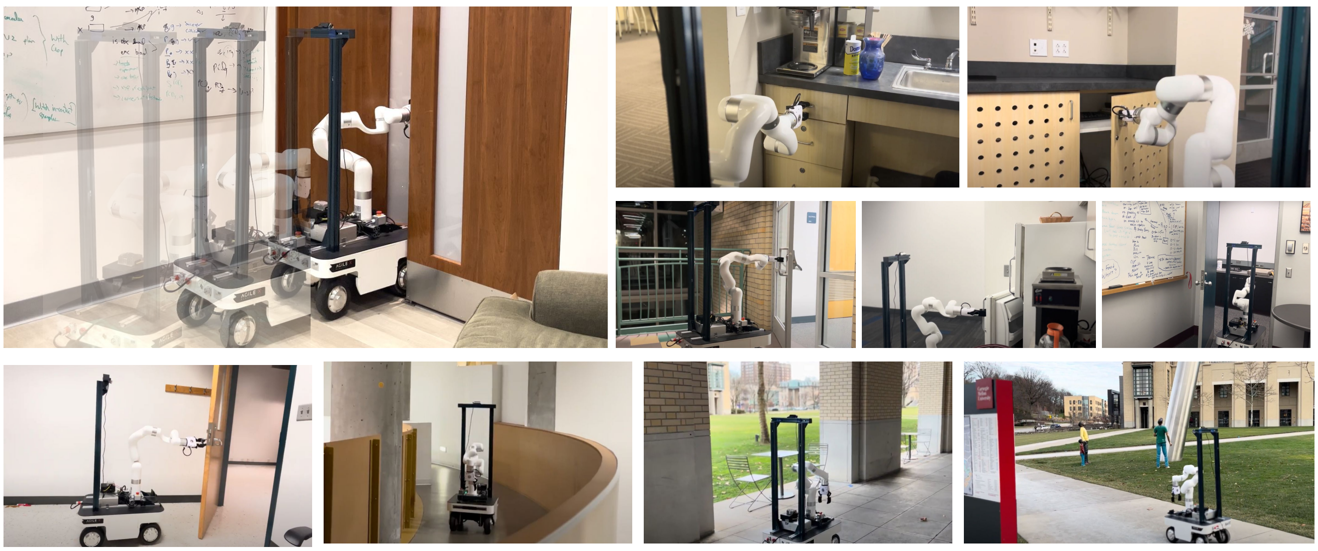

Figure 1. **Open-World Mobile Manipulation System**: We use a **full-stack** approach to operate articulated objects such as real-world doors, cabinets, drawers, and refrigerators in open-ended unstructured environments.

## Abstract

*Deploying robots in open-ended unstructured environments such as homes has been a long-standing research problem. However, robots are often studied only in closed-off lab settings, and prior mobile manipulation work is restricted to pick-move-place, which is arguably just the tip of the iceberg in this area. In this paper, we introduce Open-World Mobile Manipulation System, a **full-stack** approach to tackle realistic articulated object operation, e.g. real-world doors, cabinets, drawers, and refrigerators in open-ended unstructured environments. The robot utilizes an adaptive learning framework to initially learns from a small set of data through behavior cloning, followed by learning from online practice on novel objects that fall outside the training distribution. We also develop a low-cost mobile manipulation hardware platform capable of safe and autonomous online adaptation in unstructured environments with a cost of around* $25,000$ *USD. In our experiments we utilize 20 articulate objects across 4 buildings in the CMU campus. With less than an hour of online learning for each object, the system is able to increase success rate from* $50\%$ *of BC pre-training to* $95\%$ *using online adaptation. Video results at* https://open-world-mobilemanip.github.io/.

## 1. Introduction

Deploying robotic systems in unstructured environments such as homes has been a long-standing research problem. In recent years, significant progress has been made in deploying learning-based approaches [2, 5, 11, 16] towards this goal. However, this progress has been largely made independently either in mobility or in manipulation, while a wide range of practical robotic tasks require dealing with both aspects [4, 8, 15, 18]. The joint study of mobile manipulation paves the way for generalist robots which can perform useful tasks in open-ended unstructured environments, as opposed to being restricted to controlled laboratory settings focused primarily on tabletop manipulation.

However, developing and deploying such robot systems in the *open-world* with the capability of handling unseen objects is challenging for a variety of reasons, ranging from the lack of capable mobile manipulator hardware systems to the difficulty of operating in diverse scenarios. Consequently, most of the recent mobile manipulation results end up being limited to pick-move-place tasks[9, 12, 17, 19], which is arguably representative of only a small fraction of problems in this space. Since learning for general-purpose mobile manipulation is challenging, we focus on a restricted class of problems, involving the operation of articulated objects, such as doors, drawers, refrigerators, or cabinets in open-world environments. This is a common and essential task encountered in everyday life, and is a long-standing problem in the community [1, 3, 6, 7, 10, 13, 14]. The primary challenge is generalizing effectively across the diverse variety of such objects in unstructured real-world environments rather than manipulating a single object in a constrained lab setup. Furthermore, we also need capable hardware, as opening a door not only requires a powerful and dexterous manipulator, but the base has to be stable enough to balance while the door is being opened and agile enough to walk through.

We take a **full-stack** approach to address the above challenges. In order to effectively manipulate objects in open-world settings, we adopt a *adaptive learning* approach, where the robot keeps learning from online samples collected during interaction. Hence even if the robot encounters a new door with a different mode of articulation, or with different physical parameters like weight or friction, it can keep adapting by learning from its interactions. For such a system to be effective, it is critical to be able to learn efficiently, since it is expensive to collect real world samples. The mobile manipulator we use as shown in Figure. 1 has a very large number of degrees of freedom, corresponding to the base as well as the arm. A conventional approach for the action space of the robot could be regular end-effector control for the arm and SE2 control for the base to move in the plane. While this is very expressive and can cover many potential behaviors for the robot to perform, we will need to collect a very large amount of data to learn control policies in this space. Given that our focus is on operating articulated objects, can we structure the action space so that we can get away with needing fewer samples for learning?

Consider the manner in which people typically approach operating articulated objects such as doors. This generally first involves reaching towards a part of the object (such as a handle) and establishing a grasp. We then execute constrained manipulation like rotating, unlatching, or unhooking, where we apply arm or body movement to manipulate the object. In addition to this high-level strategy, there are also lower-level decisions made at each step regarding exact direction of movement, extent of perturbation and amount of force applied. Inspired by this, we use a hierarchical action space for our controller, where the high-level action sequence follows the grasp, constrained manipulation strategy. These primitives are parameterized by learned low-level continuous values, which needs to be adapted to operate diverse articulated objects. To further bias the exploration of the system towards reasonable actions and avoid unsafe actions during online sampling, we collect a dataset of expert demonstrations on 12 training objects, including doors, drawers and cabinets to train an initial policy via behavior cloning. While this is not very performant on new unseen doors (getting around 50% accuracy), starting from this policy allows subsequent learning to be faster and safer.

Learning via repeated online interaction also requires capable hardware. As shown in Figure 1, we provide a simple and intuitive solution to build a mobile manipulation hardware platform, followed by two main principles: (1) Versatility and agility - this is essential to effectively operate diverse objects with different physical properties in potentially challenging environments, for instance a cluttered office. (2) Affordabiluty and Rapid-prototyping - Assembled with off the shelf components, the system is accessible and can be readily be used by most research labs.

In this paper, we present **Open-World Mobile Manipulation System**, a **full stack** approach to tackle the problem of mobile manipulation of realistic articulated objects in the open world. Efficient learning is enabled by a structured action space with parametric primitives, and by pretraining the policy on a demonstration dataset using imitation learning. Adaptive learning allows the robot to keep learning from self-practice data via online RL. Repeated interaction for autonomous learning requires capable hardware, for which we propose a versatile, agile, low-cost easy to build system. We introduce a low-cost mobile manipulation hardware platform that offers a high payload, making it capable of repeated interaction with objects, e.g. a heavy, spring-loaded door, and a human-size, capable of maneuvering across various doors and navigating around narrow and cluttered spaces in the open world. We conducted a field test of 8 novel objects ranging across 4 buildings on a university campus to test the effectiveness of our system, and found adaptive earning boosts success rate from 50% from the pre-trained policy to 95% after adaptation.

## References

[1] Christopher G Atkeson, PW Babu Benzun, Nandan Banerjee, Dmitry Berenson, Christoper P Bove, Xiongyi Cui, Mathew DeDonato, Ruixiang Du, Siyuan Feng, Perry Franklin, et al. What happened at the darpa robotics challenge finals. *The DARPA robotics challenge finals: Humanoid robots to the rescue*, pages 667–684, 2018. 2

[2] Shikhar Bahl, Abhinav Gupta, and Deepak Pathak. Human-

to-robot imitation in the wild. *RSS*, 2022. 1

[3] Nandan Banerjee, Xianchao Long, Ruixiang Du, Felipe Polido, Siyuan Feng, Christopher G Atkeson, Michael Gennert, and Taskin Padir. Human-supervised control of the atlas humanoid robot for traversing doors. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, pages 722–729. IEEE, 2015. 2

[4] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil J Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla, Deeksha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez, Karl Pertsch, Jornell Quiambao, Kanishka Rao, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Kevin Sayed, Jaspiar Singh, Sumedh Sontakke, Austin Stone, Clayton Tan, Huong Tran, Vincent Vanhoucke, Steve Vega, Quan Vuong, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. Rt-1: Robotics transformer for real-world control at scale, 2023. 1

[5] Matthew Chang, Theophile Gervet, Mukul Khanna, Sriram Yenamandra, Dhruv Shah, So Yeon Min, Kavit Shah, Chris Paxton, Saurabh Gupta, Dhruv Batra, et al. Goat: Go to any thing. *arXiv preprint arXiv:2311.06430*, 2023. 1

[6] Sachin Chitta, Benjamin Cohen, and Maxim Likhachev. Planning for autonomous door opening with a mobile manipulator. In *2010 IEEE International Conference on Robotics and Automation*, pages 1799–1806. IEEE, 2010. 2

[7] Mathew DeDonato, Felipe Polido, Kevin Knoedler, Benzun PW Babu, Nandan Banerjee, Christoper P Bove, Xiongyi Cui, Ruixiang Du, Perry Franklin, Joshua P Graff, et al. Team wpi-cmu: achieving reliable humanoid behavior in the darpa robotics challenge. *Journal of Field Robotics*, 34(2):381–399, 2017. 2

[8] Zipeng Fu, Tony Z. Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. In *arXiv*, 2024. 1

[9] Alexander Herzog, Kanishka Rao, Karol Hausman, Yao Lu, Paul Wohlhart, Mengyuan Yan, Jessica Lin, Montserrat Gonzalez Arenas, Ted Xiao, Daniel Kappler, et al. Deep rl at scale: Sorting waste in office buildings with a fleet of mobile manipulators. *arXiv preprint arXiv:2305.03270*, 2023. 2

[10] Advait Jain and Charles C Kemp. Behaviors for robust door opening and doorway traversal with a force-sensing mobile manipulator. Georgia Institute of Technology, 2008. 2

[11] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots. *arXiv preprint arXiv:2107.04034*, 2021. 1

[12] Peiqi Liu, Yaswanth Orru, Chris Paxton, Nur Muhammad Mahi Shafiullah, and Lerrel Pinto. Ok-robot: What really matters in integrating open-knowledge models for robotics. *arXiv preprint arXiv:2401.12202*, 2024. 2

[13] K. Nagatani and S.I. Yuta. An experiment on opening-door-behavior by an autonomous mobile robot with a manipulator. In *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*, pages 45–50 vol.2, 1995. 2

[14] L. Peterson, D. Austin, and D. Kragic. High-level control of a mobile manipulator for door opening. In *Proceedings. 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2000) (Cat. No.00CH37113)*, pages 2333–2338 vol.3, 2000. 2

[15] Nur Muhammad Mahi Shafiullah, Anant Rai, Haritheja Etukuru, Yiqian Liu, Ishan Misra, Soumith Chintala, and Lerrel Pinto. On bringing robots home, 2023. 1

[16] Dhruv Shah, Ajay Sridhar, Nitish Dashora, Kyle Stachowicz, Kevin Black, Noriaki Hirose, and Sergey Levine. Vint: A foundation model for visual navigation, 2023. 1

[17] Charles Sun, Jedrzej Orbik, Coline Devin, Brian Yang, Abhishek Gupta, Glen Berseth, and Sergey Levine. Fully autonomous real-world reinforcement learning with applications to mobile manipulation, 2021. 2

[18] Ruihan Yang, Yejin Kim, Aniruddha Kembhavi, Xiaolong Wang, and Kiana Ehsani. Harmonic mobile manipulation, 2023. 1

[19] Sriram Yenamandra, Arun Ramachandran, Karmesh Yadav, Austin Wang, Mukul Khanna, Theophile Gervet, Tsung-Yen Yang, Vidhi Jain, Alexander William Clegg, John Turner, et al. Homerobot: Open-vocabulary mobile manipulation. *arXiv preprint arXiv:2306.11565*, 2023. 2