

# TidyBot++: An Open-Source Holonomic Mobile Manipulator for Robot Learning

Jimmy Wu<sup>1</sup>, William Chong<sup>2</sup>, Robert Holmberg<sup>3</sup>, Aaditya Prasad<sup>2</sup>, Yihuai Gao<sup>2</sup>,  
Oussama Khatib<sup>2</sup>, Shuran Song<sup>2</sup>, Szymon Rusinkiewicz<sup>1</sup>, Jeannette Bohg<sup>2</sup>  
<sup>1</sup>Princeton University <sup>2</sup>Stanford University <sup>3</sup>Dexterity

<http://tidybot2.github.io>

**Abstract:** Exploiting the promise of recent advances in imitation learning for mobile manipulation will require the collection of large numbers of human-guided demonstrations. This paper proposes an open-source design for an inexpensive, robust, and flexible mobile manipulator that can support arbitrary arms, enabling a wide range of real-world household mobile manipulation tasks. Crucially, our design uses powered casters to enable the mobile base to be fully *holonomic*, able to control all planar degrees of freedom independently and simultaneously. This feature makes the base more maneuverable and simplifies many mobile manipulation tasks, eliminating the kinematic constraints that create complex and time-consuming motions in nonholonomic bases. We equip our robot with an intuitive mobile phone teleoperation interface to enable easy data acquisition for imitation learning. In our experiments, we use this interface to collect data and show that the resulting learned policies can successfully perform a variety of common household mobile manipulation tasks.

**Keywords:** mobile manipulation, imitation learning, holonomic drive




Figure 1: We develop an open-source mobile manipulator with a holonomic base (left), and show that it can perform a variety of household tasks in a real apartment home (right).

## 1 Introduction

Imitation learning from real-world data is starting to show very promising results in robotics for both fixed-arm robots [1, 2, 3, 4, 5, 6, 7] and mobile manipulators [8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21]. However, one key bottleneck is the availability of data. Unlike in natural language processing, which can train on readily-available data from the internet, real-world data for training robot policies is much harder to come by. As a result, scaling data collection of robotic tasks has

become a high-interest research direction. Recent efforts have collected large robot learning datasets to address this [22, 23, 24, 25, 26, 27]. These datasets were largely collected on fixed-arm robot setups. However, to bring robots to their full potential, mobility is important as it will enable many more tasks in realistic household settings [8].

We believe that one reason there are so few data collection and learning efforts in mobile manipulation is the lack of suitable research hardware. Existing commercial options for mobile bases are often tailored towards industrial or warehouse use cases, and may be ill-suited for household environments due to their large size. They are also often expensive and are typically subject to kinematic constraints.

In this work, we propose an open-source design for a mobile base designed to carry a robot arm sized for use in household environments. In addition to being inexpensive, flexible, and easy to assemble, our base is *holonomic*: able to independently move in any of the three degrees of freedom (DoFs) on the ground plane— $(x, y, \theta)$ —at any time. We argue that this is an important advantage for more intuitive teleoperation, and greatly increases the ease of acquiring large amounts of training data for real-world imitation learning.

Nonholonomic robots, such as differential drive (wheelchair-like) or Ackermann drive (car-like) platforms, have constrained degrees of freedom in their motion. The most notable consequence of this is that they cannot move sideways. For example, cars cannot directly drive sideways into a street-side parking spot and must execute a multi-step parallel-parking maneuver.

In contrast, holonomic robots have no kinematic constraints and can simultaneously and independently control all three degrees of freedom. An example of a holonomic vehicle common in everyday life is an office chair, which can be smoothly pushed or rotated in arbitrary directions. This is enabled by the design of the caster wheels (Fig. 2), which have an offset between the vertical axis of the swivel mechanism and the roll axis of the wheel. This offset is a crucial design feature of casters and is what makes the office chair fully holonomic. It creates a lever arm that causes the wheel to trail behind the vertical axis of the swivel as the chair moves, automatically aligning the wheels to the direction of movement. Without a caster offset, the vehicle would be omnidirectional (capable of moving in any direction) but still nonholonomic, as the wheels have to be manually aligned to face the direction of desired motion before the vehicle can begin moving. Overall, holonomic drive is preferred for maximum maneuverability.




Figure 2: A simplified illustration of caster wheels on a holonomic base.

Our holonomic base uses a powered-caster drive mechanism [28]. It is driven by four motorized casters, and can be thought of as a motorized office chair. The ability to steer all four wheels makes the base omnidirectional, and the caster offset makes the base holonomic, allowing it to instantaneously accelerate in any direction as it does not need to first align the wheels to the direction of motion.

A holonomic mobile base enables easier teleoperation and kinesthetic teaching for collecting imitation learning data. Everyday tasks such as opening doors and cabinets often require sideways motions of the mobile base to improve the workspace of the arm during execution. This useful motion is not immediately available with a differential drive base. Instead, the robot has to replan vehicle trajectories to satisfy nonholonomic constraints, which costs extra motion and time with no added value to the task. A holonomic mobile base, on the other hand, can be much more reactive. It can be moved arbitrarily in any direction no matter the current configuration, allowing an operator to make fine adjustments to the positioning of the base.

A holonomic base is also useful for policy learning and inference. Recent real-robot imitation learning works have converged on the use of position representations, as they are more stable and less

noisy compared to velocities. However, a nonholonomic mobile base can only be controlled in velocity mode [8, 16]. A holonomic base, on the other hand, can be directly commanded to go to a task space position  $(x, y, \theta)$  in a repeatable manner, as it can independently control all DOFs with no constraints. In our experiments, we show that we can indeed train high-performing policies for our robot across several mobile manipulation tasks in a real apartment home. Additionally, we show that policies can be learned more easily with data collected from a holonomic base compared to a nonholonomic one.

To facilitate easy data collection with our new mobile manipulator, we also develop a mobile phone teleoperation interface. The interface uses the WebXR API [29] to stream the real-time 6-DoF pose of the mobile phone to a computer, which maps the phone motion to corresponding motions of the mobile base or arm via low-level control. WebXR is supported on most modern Android and iOS phones, so our interface does not require purchase of a separate teleoperation device. In our experiments, we use this teleoperation interface to collect data for training our policies.

Our holonomic mobile base is low-cost (\$5–6k USD) and designed from the ground up to optimize for mobile manipulation research productivity. We will fully open source all aspects of this system, including the hardware design, mobile phone teleoperation interface, policy learning setup, and low-level controller. We will also create a documentation webpage for the mobile base, including bill of materials (BOM), a hardware assembly guide with videos, and 3D CAD files. We believe these components can help democratize access to highly maneuverable mobile manipulators, increase ease and practicality of mobile manipulation data collection, and improve research reproducibility by providing a standardized and reusable robot platform.

Our key contributions in this work are thus: (1) an open-source design for a holonomic mobile manipulator, (2) a mobile phone teleoperation interface for easily collecting data with the mobile manipulator, and (3) demonstration that our system is capable of learning policies. Please see our project page at <http://tidybot2.github.io> for documentation, code, and qualitative videos of our robot in action.

## 2 Related Work

**Mobile manipulation hardware platforms.** Recently, there has been an increased interest in equipping manipulation robots with mobility to demonstrate their full potential for a variety of tasks—especially in domestic settings [8, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41]. These works utilize a number of different mobile bases which we compare to our base in more detail in Tab. 1. For example, several projects use the Tiago mobile manipulator, and learn policies with a combination of RL, simulation, and human motion transfer [42, 43, 44, 30, 45]. Tiago is holonomic through its use of mecanum wheels, but the wheel comes with a few drawbacks. Due to the multiple small rollers on the wheel, the contact point with the ground is discontinuous leading to vibrations during operation. They also provide less traction and have poorer ability to traverse door thresholds and other inevitable household clutter compared to a conventional wheel. Another common category of bases uses differential drive, which prevents them from freely navigating every degree of freedom [46, 36, 8]. An example of such a robot is Everyday Robots’ mobile manipulator, which often appears in Google’s robot learning papers as a scalable mobile manipulator platform [36, 47, 48]. This robot is neither open source nor purchasable by the public.

While the above works use wheeled bases, several works attach a robot arm to a quadruped and utilize the combined system as a mobile manipulator capable of traversing terrains that wheeled robots cannot access [35, 37, 49, 50]. Though this combination expands the set of navigable environments, these quadrupeds are nonholonomic since they are kinematically constrained by the orientation and placement of their limbs and feet.

**Data collection for mobile manipulation.** To address the lack of robotics data for learning manipulation policies, several works have developed data collection platforms. The majority of these platforms are built for fixed-arm setups [51, 52, 25, 53]. For example, the DROID dataset [25] was collected on a standardized setup with an arm mounted on a portable table. The authors use an




Figure 3: Our mobile base is designed to be modular and easily reconfigurable. It has very few components and can be assembled in 1 to 2 days.

Oculus controller to teleoperate the robot. However, this controller must remain in view of four IR receivers which can lead to unexpected motion if the controller moves out of view and back. Similar to our work, RoboTurk [51, 54] used a mobile phone to teleoperate fixed-base robot arms which is a much more flexible solution and does not require purchasing a dedicated teleoperation device. However, their system suffers from drift as they only rely on IMU measurements and do not use the camera. MART [53] and MOMART [55], which extend RoboTurk to multi-arm and mobile manipulation, respectively, suffer from similar shortcomings and have not been demonstrated on real robots. In our work, we use the WebXR API [29], which combines IMU data with visual odometry based on the phone’s camera to mitigate drift. TeleMoMa [45] is a teleoperation framework that supports multiple teleoperation interfaces and three commercially-available, high-cost robots (for a detailed comparison to our low-cost base, see Tab. 1). One of the supported teleoperation devices is a mobile phone app based on ARKit, similar to what we use. Our interface is based on WebXR, which leverages ARKit on iPhone but works on Android as well.

There are several works that propose data collection devices that are hand-held by the human demonstrator [52, 11] but in this case the demonstrator does not get direct feedback on whether a demonstration is kinematically feasible by the robot. Of those approaches, Dobb-E [11] proposes a low-cost reacher-grabber stick with a mounted iPhone to record data. The authors then train visuomotor policies on this data that are deployed on a differential drive Stretch robot [46]. Mobile ALOHA [8] is a dual-arm mobile manipulation platform capable of performing an impressive array of household tasks. However, the robot’s differential drive base and large footprint limits its maneuverability, and the arms are not able to reach the ground. Furthermore, the teleoperator is strapped to the back of the platform far away from the end effectors, which can make it hard to teleoperate precise actions. For our system, the teleoperator can freely walk around the scene and get very close when precision is required.

### 3 Hardware Design

We designed this mobile base concept from the ground up to optimize for mobile manipulation research productivity. It is simple, low-cost (\$5–6k), and modular (Fig. 3). The core is the drive system, which is based on readily available components from the FIRST Robotics Competition (FRC) [56] ecosystem. A basic frame built out of aluminum T-slot extrusions carries the four motorized caster modules that are powered through a fused power distribution panel by a sealed lead acid (SLA) battery. There are many similar components in the FRC ecosystem that could be used to build similar systems. The large and active community of FRC users and vendors ensures that components are well-documented and readily available.

Our drive system is derived from a set of SDS MK4 swerve modules [57] widely used in FRC. The swerve modules are very similar to a caster, with a wheel that can be actively steered and driven, except they have no caster offset and thus create a nonholonomic base. The modules use two motors with integrated encoders and CAN bus controllers, one for steering and one for driving. Additionally, there is an absolute encoder mounted directly to the steering axis to measure the steer

position, and thereby eliminate the need for a startup homing motion. A USB-to-CAN adapter is used to communicate with the motors and encoders through CAN bus.

We modify the MK4 swerve module to create a holonomic base design by introducing a caster offset using a minimal number of modifications to the stock swerve module: 2 custom 3D-printed wheel mounts and a custom machined shaft. The wheel mounts can be printed with PLA filament on a standard FDM 3D printer, and the shaft can be easily ordered from online machining services like Xometry using our provided CAD file and specifications. All other parts of the off-the-shelf kit are directly reused.

To complete our mobile manipulator, we add a mini PC (Intel NUC) and the Kinova Gen3 arm. Power for compute, manipulation, and peripherals is provided by a high-capacity (768 Wh), fast-charging (0–100% in 70 minutes) portable power station (camping battery) with AC power outlets. The portable power station (8.6 kg) and SLA battery (6 kg) serve as counterweights to prevent the base from tipping over. Note that it would be possible to build electrical circuitry to use only one battery for more space efficiency, but we opted for separate batteries as it provides more flexibility and makes setup much easier.

Our open-source design is highly customizable: the frame can be easily modified to support mounting of different arms or even multiple arms, and many additional sensors such as cameras, microphones, etc., can be easily mounted and powered, to suit the research being done.

### 3.1 Powered-caster vehicle kinematics

For our low-level controller, we model the kinematics of the mobile base largely following the formulation of the powered-caster vehicle (PCV) in Holmberg and Khatib [28], which describes the mapping between the joint space and the operational space  $(x, y, \theta)$  of the base.

Each caster module is modeled with two revolute joints: a steer joint and a roll joint. The steer joint determines the wheel’s steering angle, while the roll joint measures the wheel’s rotational movement. We have an incremental encoder in each motor, as well as an absolute encoder on the steer axis of each caster. With these encoder readings, we can calculate the joint positions and velocities using gearbox kinematics.

The main difference from the original PCV formulation [28] is that instead of a single caster offset  $b$ , our caster module has a two-dimensional offset: a traditional longitudinal  $b_x$  offset as well as a small lateral  $b_y$  offset (Fig. 4). This mainly comes as a byproduct of our design criteria to minimize the number of custom parts needed. It would be possible to have no lateral offset  $b_y$  by designing more custom parts, or designing a new caster from scratch, but this would significantly lower the accessibility of our design, as creating custom parts requires detailed design work and is very costly to manufacture at low quantities.

### 3.2 Design principles

Our holonomic mobile base is specifically designed and optimized for robot learning research. This objective shapes our key design principles:

**Research flexibility.** We would like users to be able to customize the robot to meet their own specific needs. The frame is designed using standard aluminum T-slot extrusions, providing flexibility to easily adjust the dimensions and shape of the frame. We use a portable power station (camping battery) with four AC power outlets, providing flexibility to power other kinds of robot arms as well as computers of various operating voltages. Components can be simply plugged in directly, with-




Figure 4: Isometric and top views of a simplified caster, showing the caster offsets  $b_x$  and  $b_y$ , wheel radius  $r$ , steer and roll joints  $\phi$  and  $\rho$ , and caster module placement  $(h, \beta)$  from the base origin.

Table 1: Mobile base and mobile manipulator comparison

| Specification   | Ours   | Stretch | Tracer  | Ranger Mini | Husky | Fetch  | Tiago  |
|-----------------|--------|---------|---------|-------------|-------|--------|--------|
| Holonomic       | Yes    | No      | No      | No          | No    | No     | Yes    |
| Omnidirectional | Yes    | No      | No      | Yes         | No    | No     | Yes    |
| Swappable arm   | Yes    | No      | Yes     | Yes         | Yes   | No     | No     |
| Footprint (cm)  | 50x54  | 33x34   | 57x69   | 50x74       | 67x99 | 51x56  | 54x54  |
| Weight          | 34 kg  | 24.5 kg | 30 kg   | 63 kg       | 50 kg | 113 kg | 70 kg  |
| Payload         | 60 kg  | 10 kg   | 100 kg  | 80 kg       | 75 kg | —      | —      |
| Maximum speed   | 1 m/s  | —       | 1.6 m/s | 1.5 m/s     | 1 m/s | 1 m/s  | 1 m/s  |
| Runtime         | 8 h    | 2–5 h   | 4 h     | 7–8 h       | 3 h   | 9 h    | 8–10 h |
| Cost            | \$5.4k | \$25k   | \$7.6k  | \$13k       | \$20k | \$100k | \$100k |

out needing to set up new circuitry for power supply and voltage conversion. Additionally, we open source the entire control stack, all the way down to low-level control with motor velocity commands. This means researchers can have full control and are not limited to the API functionality exposed by a manufacturer’s proprietary software stack.

**Reliable and easily-sourced parts.** One challenge with building a mobile robot for research is the sourcing of parts such as motors, encoders, gears, and custom machined parts, which can be time consuming and costly. We design our drive system mainly using parts sourced from suppliers of the FIRST robotics competition (FRC) [56] for high schoolers. These are the same parts used by over 80,000 competition participants each year. Due to the popularity of the competition, the components are standardized and readily available for purchase from online vendors, and can be ordered online and typically delivered within a week. These parts are very reliable, as they have been battle-tested in the strenuous conditions of competition (125 lb robots moving at high speeds with frequent collisions), and core software components such as CAN drivers, motor control, and battery monitoring are all included. Our caster module design contains only three easily-obtained custom parts as described in Sec. 3.1. The top and bottom plates of the frame are laser-cut acrylic, which can be made with a laser cutter or directly ordered using an online service. All other components of the mobile base can be readily purchased from online retailers such as Amazon.

**Easy assembly and repair.** Our holonomic mobile base design is easy to assemble and can be put together in 1 to 2 days. Putting together the T-slot extrusion frame is the most time-consuming part of the process and takes around 6 hours. Each of the four caster modules can be assembled in less than 30 minutes using only hand tools, and subsequently installed on the frame. The custom wheel mounts for the casters can be 3D printed across 2 days. All electrical wiring can be done in less than 30 minutes and requires no soldering. For repairs, all parts can be easily removed from the frame in a modular fashion, including the caster modules. The robot does not need to be shipped back to a manufacturer when parts break. Instead, replacement parts can be purchased online and directly swapped out for the broken ones.

### 3.3 Specifications

In Tab. 1, we compare our holonomic base with other mobile bases and mobile manipulators commonly used in research. This includes the Stretch mobile manipulator from Hello Robot, the Tracer and Ranger Mini 2.0 AGVs from AgileX, the Husky AGV from Clearpath, and the Fetch and Tiago mobile manipulators. Our mobile base is maximally maneuverable, performant, and flexible, while also having the lowest cost. Note that while the Tiago is holonomic, it uses mecanum wheels which make the robot vibrate as it moves.

**Dimensions.** Our robot has a small footprint (54 x 50 cm), enabling it to navigate in household environments, including narrow doorways and hallways. The top of the mobile base is approximately 37.5 cm above the ground. This height allows the mounted arm to comfortably reach down towards the ground and up towards tabletops and countertops. Note that these dimensions can be easily customized to accommodate other arms.

**Weight.** The mobile base weighs 75 lb (34 kg). We mount a Kinova Gen3 7-DoF arm on top that weighs 27 lb (12 kg) including its mounting plate and power supply unit.

**Payload.** To get a rough idea of the robot’s maximum payload, we loaded 270 lb (122 kg) of weight plates onto the mobile base, for a total weight of 345 lb (156 kg). Even with this amount of weight, we found that the motors showed no signs of struggle. For a conservative estimate of a maximum payload suitable for long-term use, we halve the weight and use 60 kg as our estimate. This indicates that our design can comfortably support the weight of many other arms, as long as the frame is redesigned appropriately for securely mounting the arm. Note also that the payload may vary depending on the terrain (we conducted our test on a hard floor).

**Runtime.** We power the robot arm and compute using a portable power station with 768 Wh of capacity, which we found can handle 8 hours of continuous teleoperation runtime. For the motors, since they boot nearly instantly upon connection to power, the SLA battery powering the motors can be easily hotswapped during use if the voltage gets low.

**Traversability.** Our mobile base works on various indoor floor surfaces ranging from hard floor to high pile carpet, and can traverse many common floor obstacles such as door thresholds, floor mats, and elevator gaps. While not intended for outdoor use, we found during transport of our robot that it can handle many outdoor floor obstacles as well, such as bumpy sidewalks, steel construction plates, loading ramps (inclination 6.5 degrees), and speed bumps.

**Odometry.** To evaluate the odometry accuracy of the mobile base, we use motion capture with submillimeter accuracy to measure the pose of the base while driving it in several path shapes. Overall, we find the odometry quality to be quite high, with translation drift of less than 1 cm per meter of distance traveled, and rotation drift of less than 1 deg per 360 deg of rotation. This means that base movements are highly repeatable, as our holonomic base can be controlled in position mode to accurately reach goal poses  $(x, y, \theta)$  using odometry feedback.

## 4 Experiments

In these experiments, we aim to show (i) that our teleoperation interface can collect useful demonstration data to successfully train policies for a variety of household mobile manipulation tasks, and (ii) that holonomic drive offers advantages over differential drive for both teleoperation and policy learning. For qualitative videos of autonomous policy rollouts and teleoperation, please see our project page at <http://tidybot2.github.io>.

### 4.1 Imitation learning

We used our phone teleoperation interface to collect demonstrations for the 6 tasks shown in Tab. 2. We collected 100 demonstrations for the shorter *open fridge* task and 50 for all others. Data collection for each task took between 1 and 2 hours for 50 episodes, including overhead for environment resets.

We then used the data to train a diffusion policy [1] for each task. We trained each policy for 500 epochs and evaluated them by running 10 episodes of policy rollouts.

Success rates are shown in Tab. 2. Note that while diffusion policies are typically trained using 200 to 300 demonstrations, we found that 50 was already sufficient for the robot to learn to complete the task successfully. Performance can likely be further improved with more data. These results show that our system is capable of learning high-performing policies for useful tasks in real homes.

Table 2: Imitation learning results

| Task            | Success rate |
|-----------------|--------------|
| Open fridge     | 10/10        |
| Wipe countertop | 9/10         |
| Load dishwasher | 7/10         |
| Take out trash  | 10/10        |
| Load laundry    | 7/10         |
| Water plant     | 6/10         |

### 4.2 Differential drive comparison

To further highlight the advantages of holonomic drive over nonholonomic, we conduct a head-to-head comparison on the *wipe countertop* task. We collect 50 demonstrations of the task with our base

operating in differential drive mode. We implement this by applying the appropriate nonholonomic constraints to the desired base pose and then computing a velocity command to send to the base.

Fig. 5 shows a representative pair of paths from the demonstration data (path begins at the origin), illustrating the greater efficiency enabled by a holonomic base when compared to differential drive. Across all demonstrations for this task, the differential drive base travels on average 4.03 m (average episode duration 65.2 s), whereas the holonomic base travels only 2.03 m (average episode duration 27.4 s).

We then train a diffusion policy using the differential drive data and compare the resulting policy with the holonomic one from before. To ensure a fair comparison, the differential drive policy was trained under identical conditions (50 demonstrations and 500 epochs). While the holonomic base can perform the task with 9/10 success rate, we find that the differential drive policy only achieves 4/10 success rate.

Qualitatively, we find that the differential drive policy does not learn the task as effectively, even though it was given the same number of demonstrations. The main failure mode is that it frequently skips over portions of the countertop rather than wiping. We believe this is because the learning problem is strictly harder in the case of differential drive. The policy not only has to learn the wiping action, it also has to learn a complicated sideways moving strategy similar to parallel parking, whereas the holonomic policy can directly move the base sideways. Additionally, the differential drive maneuvers cause the camera’s view to swerve from side to side, reducing its quality, while the holonomic base can maintain a stable, forward-facing camera view.




Figure 5: In the *wipe countertop* task, the differential drive robot is forced to take a less efficient path as it is subject to nonholonomic constraints.

## 5 Limitations

One limitation of our mobile base is that it does not backdrive very well due to high steering friction in the caster modules, which can be attributed to the combination of high steer gear ratio (12.8) and small caster offset (14 mm). We confirmed that with the steer gearing removed, the base can indeed backdrive smoothly. The base could be made more easily backdrivable with further mechanical modifications using custom machined parts, but this would come at the cost of reducing the accessibility of the open-source design.

## 6 Conclusion

In this paper, we proposed an open-source design for a mobile manipulator with an inexpensive, robust, and flexible holonomic base. The design uses powered-caster wheel modules, which makes the base holonomic and therefore more maneuverable than commonly used nonholonomic designs. With our easy-to-use mobile phone teleoperation interface, we showed that many household mobile manipulation tasks become easy to demonstrate on our robot. Additionally, we showed that the data collected with our interface can be used to train high-performing policies on a variety of tasks. We hope that the open-source release of this mobile base design and teleoperation interface will enable the robot learning community to easily collect large quantities of mobile manipulation data that will form the basis for policy learning.



## Acknowledgments

We would like to thank Kevin Zakka, Yixuan Huang, Kevin Lin, Zi-ang Cao, Jingyun Yang, Rika Antonova, Marion Lepert, Sophie Lueth, Haoyu Xiong, Huy Ha, Cheng Chi, Philipp Wu, Fred Shentu, and Zhongke Yi for fruitful technical discussions. We thank Rajat Kumar Jenamani, Rishabh Madan, and the Cornell EmPRISE Lab for open sourcing their compliant controller for the Kinova arm. We would also like to thank Zi-ang Cao, Rika Antonova, and Haoyu Xiong for extensive hardware assistance, as well as Rika Antonova and Zipeng Fu for lending hardware. We are especially grateful towards the FIRST Robotics Competition (FRC) community for developing the rich ecosystem that made this project feasible. For extensive product and logistics support, we would like to thank Mandy Gove, Cory Ness, Omar Zrien, Jacob Caporuscio, and Dalton Smith from CTR Electronics; Ranjit Chahal and Harvey Rico from WestCoast Products; and John Rigsby from Swerve Drive Specialties. The work was supported in part by the Stanford Institute for Human-Centered Artificial Intelligence (HAI), Princeton School of Engineering, the Sloan Foundation, and the National Science Foundation under ECCS-2143601.

## References

- [1] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- [2] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- [3] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel. Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators. *arXiv preprint arXiv:2309.13037*, 2023.
- [4] C. Wang, L. Fan, J. Sun, R. Zhang, L. Fei-Fei, D. Xu, Y. Zhu, and A. Anandkumar. Mimicplay: Long-horizon imitation learning by watching human play. *arXiv preprint arXiv:2302.12422*, 2023.
- [5] H. Ha, P. Florence, and S. Song. Scaling up and distilling down: Language-guided robot skill acquisition. In *Conference on Robot Learning*, pages 3766–3777. PMLR, 2023.
- [6] J. W. Kim, T. Z. Zhao, S. Schmidgall, A. Deguet, M. Kobilarov, C. Finn, and A. Krieger. Surgical robot transformer (srt): Imitation learning for surgical tasks. *arXiv preprint arXiv:2407.12998*, 2024.
- [7] A. Z. Ren, J. Lidard, L. L. Ankile, A. Simeonov, P. Agrawal, A. Majumdar, B. Burchfiel, H. Dai, and M. Simchowitz. Diffusion policy policy optimization. *arXiv preprint arXiv:2409.00588*, 2024.
- [8] Z. Fu, T. Z. Zhao, and C. Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. In *arXiv*, 2024.
- [9] A. Prasad, K. Lin, J. Wu, L. Zhou, and J. Bohg. Consistency policy: Accelerated visuomotor policies via consistency distillation. In *Robotics: Science and Systems*, 2024.
- [10] S. Lee, Y. Wang, H. Etukuru, H. J. Kim, N. M. M. Shafiullah, and L. Pinto. Behavior generation with latent actions. *arXiv preprint arXiv:2403.03181*, 2024.
- [11] N. M. M. Shafiullah, A. Rai, H. Etukuru, Y. Liu, I. Misra, S. Chintala, and L. Pinto. On bringing robots home. *arXiv preprint arXiv:2311.16098*, 2023.
- [12] R. Yang, Y. Kim, A. Kembhavi, X. Wang, and K. Ehsani. Harmonic mobile manipulation. *arXiv preprint arXiv:2312.06639*, 2023.

- [13] H. Etukuru, N. Naka, Z. Hu, S. Lee, J. Mehu, A. Edsinger, C. Paxton, S. Chintala, L. Pinto, and N. M. M. Shafiullah. Robot utility models: General policies for zero-shot deployment in new environments. *arXiv preprint arXiv:2409.05865*, 2024.
- [14] A. Sridhar, D. Shah, C. Glossop, and S. Levine. Nomad: Goal masked diffusion policies for navigation and exploration. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 63–70. IEEE, 2024.
- [15] K. Ehsani, T. Gupta, R. Hendrix, J. Salvador, L. Weihs, K.-H. Zeng, K. P. Singh, Y. Kim, W. Han, A. Herrasti, et al. Spoc: Imitating shortest paths in simulation enables effective navigation and manipulation in the real world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16238–16250, 2024.
- [16] H. Xiong, R. Mendonca, K. Shaw, and D. Pathak. Adaptive mobile manipulation for articulated objects in the open world. *arXiv preprint arXiv:2401.14403*, 2024.
- [17] J. Yang, C. Glossop, A. Bhorkar, D. Shah, Q. Vuong, C. Finn, D. Sadigh, and S. Levine. Pushing the limits of cross-embodiment learning for manipulation and navigation. *arXiv preprint arXiv:2402.19432*, 2024.
- [18] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn. Humanplus: Humanoid shadowing and imitation from humans. *arXiv preprint arXiv:2406.10454*, 2024.
- [19] X. Cheng, J. Li, S. Yang, G. Yang, and X. Wang. Open-television: Teleoperation with immersive active visual feedback. *arXiv preprint arXiv:2407.01512*, 2024.
- [20] P. Sundaresan, Q. Vuong, J. Gu, P. Xu, T. Xiao, S. Kirmani, T. Yu, M. Stark, A. Jain, K. Hausman, et al. Rt-sketch: Goal-conditioned imitation learning from hand-drawn sketches. 2024.
- [21] H. Ha, Y. Gao, Z. Fu, J. Tan, and S. Song. Umi on legs: Making manipulation policies mobile with manipulation-centric whole-body controllers. *arXiv preprint arXiv:2407.10353*, 2024.
- [22] H. Walke, K. Black, A. Lee, M. J. Kim, M. Du, C. Zheng, T. Zhao, P. Hansen-Estruch, Q. Vuong, A. He, V. Myers, K. Fang, C. Finn, and S. Levine. Bridgedata v2: A dataset for robot learning at scale. In *Conference on Robot Learning (CoRL)*, 2023.
- [23] H.-S. Fang, H. Fang, Z. Tang, J. Liu, J. Wang, H. Zhu, and C. Lu. Rh20t: A robotic dataset for learning diverse skills in one-shot. In *RSS 2023 Workshop on Learning for Task and Motion Planning*, 2023.
- [24] O. X.-E. Collaboration, A. O’Neill, A. Rehman, A. Maddukuri, A. Gupta, A. Padalkar, A. Lee, A. Pooley, A. Gupta, A. Mandlekar, A. Jain, A. Tung, A. Bewley, A. Herzog, A. Irpan, A. Khazatsky, A. Rai, A. Gupta, A. Wang, A. Kolobov, A. Singh, A. Garg, A. Kembhavi, A. Xie, A. Brohan, A. Raffin, A. Sharma, A. Yavary, A. Jain, A. Balakrishna, A. Wahid, B. Burgess-Limerick, B. Kim, B. Schölkopf, B. Wulfe, B. Ichter, C. Lu, C. Xu, C. Le, C. Finn, C. Wang, C. Xu, C. Chi, C. Huang, C. Chan, C. Agia, C. Pan, C. Fu, C. Devin, D. Xu, D. Morton, D. Driess, D. Chen, D. Pathak, D. Shah, D. Büchler, D. Jayaraman, D. Kalashnikov, D. Sadigh, E. Johns, E. Foster, F. Liu, F. Ceola, F. Xia, F. Zhao, F. V. Frujeri, F. Stulp, G. Zhou, G. S. Sukhatme, G. Salhotra, G. Yan, G. Feng, G. Schiavi, G. Berseth, G. Kahn, G. Yang, G. Wang, H. Su, H.-S. Fang, H. Shi, H. Bao, H. B. Amor, H. I. Christensen, H. Furuta, H. Walke, H. Fang, H. Ha, I. Mordatch, I. Radosavovic, I. Leal, J. Liang, J. Abou-Chakra, J. Kim, J. Drake, J. Peters, J. Schneider, J. Hsu, J. Bohg, J. Bingham, J. Wu, J. Gao, J. Hu, J. Wu, J. Wu, J. Sun, J. Luo, J. Gu, J. Tan, J. Oh, J. Wu, J. Lu, J. Yang, J. Malik, J. Silvério, J. Hejna, J. Booher, J. Tompson, J. Yang, J. Salvador, J. J. Lim, J. Han, K. Wang, K. Rao, K. Pertsch, K. Hausman, K. Go, K. Gopalakrishnan, K. Goldberg, K. Byrne, K. Oslund, K. Kawaharazuka, K. Black, K. Lin, K. Zhang, K. Ehsani, K. Lekkala, K. Ellis, K. Rana, K. Srinivasan, K. Fang, K. P. Singh, K.-H. Zeng, K. Hatch, K. Hsu, L. Itti, L. Y. Chen, L. Pinto, L. Fei-Fei, L. Tan, L. J. Fan, L. Ott, L. Lee, L. Weihs, M. Chen, M. Lepert, M. Memmel,

- M. Tomizuka, M. Itkina, M. G. Castro, M. Spero, M. Du, M. Ahn, M. C. Yip, M. Zhang, M. Ding, M. Heo, M. K. Srirama, M. Sharma, M. J. Kim, N. Kanazawa, N. Hansen, N. Heess, N. J. Joshi, N. Suenderhauf, N. Liu, N. D. Palo, N. M. M. Shafiullah, O. Mees, O. Kroemer, O. Bastani, P. R. Sanketi, P. T. Miller, P. Yin, P. Wohlhart, P. Xu, P. D. Fagan, P. Mitran, P. Sermanet, P. Abbeel, P. Sundaresan, Q. Chen, Q. Vuong, R. Rafailov, R. Tian, R. Doshi, R. Mart'in-Mart'in, R. Baijal, R. Scalise, R. Hendrix, R. Lin, R. Qian, R. Zhang, R. Mendonca, R. Shah, R. Hoque, R. Julian, S. Bustamante, S. Kirmani, S. Levine, S. Lin, S. Moore, S. Bahl, S. Dass, S. Sonawani, S. Song, S. Xu, S. Haldar, S. Karamcheti, S. Adebola, S. Guist, S. Nasiriany, S. Schaal, S. Welker, S. Tian, S. Ramamoorthy, S. Dasari, S. Belkhale, S. Park, S. Nair, S. Mirchandani, T. Osa, T. Gupta, T. Harada, T. Matsushima, T. Xiao, T. Kollar, T. Yu, T. Ding, T. Davchev, T. Z. Zhao, T. Armstrong, T. Darrell, T. Chung, V. Jain, V. Vanhoucke, W. Zhan, W. Zhou, W. Burgard, X. Chen, X. Chen, X. Wang, X. Zhu, X. Geng, X. Liu, X. Liangwei, X. Li, Y. Pang, Y. Lu, Y. J. Ma, Y. Kim, Y. Chebotar, Y. Zhou, Y. Zhu, Y. Wu, Y. Xu, Y. Wang, Y. Bisk, Y. Dou, Y. Cho, Y. Lee, Y. Cui, Y. Cao, Y.-H. Wu, Y. Tang, Y. Zhu, Y. Zhang, Y. Jiang, Y. Li, Y. Li, Y. Iwasawa, Y. Matsuo, Z. Ma, Z. Xu, Z. J. Cui, Z. Zhang, Z. Fu, and Z. Lin. Open X-Embodiment: Robotic learning datasets and RT-X models. <https://arxiv.org/abs/2310.08864>, 2023.
- [25] A. Khazatsky, K. Pertsch, S. Nair, A. Balakrishna, S. Dasari, S. Karamcheti, S. Nasiriany, M. K. Srirama, L. Y. Chen, K. Ellis, P. D. Fagan, J. Hejna, M. Itkina, M. Lepert, Y. J. Ma, P. T. Miller, J. Wu, S. Belkhale, S. Dass, H. Ha, A. Jain, A. Lee, Y. Lee, M. Memmel, S. Park, I. Radosavovic, K. Wang, A. Zhan, K. Black, C. Chi, K. B. Hatch, S. Lin, J. Lu, J. Mercat, A. Rehman, P. R. Sanketi, A. Sharma, C. Simpson, Q. Vuong, H. R. Walke, B. Wulfe, T. Xiao, J. H. Yang, A. Yavary, T. Z. Zhao, C. Agia, R. Baijal, M. G. Castro, D. Chen, Q. Chen, T. Chung, J. Drake, E. P. Foster, J. Gao, D. A. Herrera, M. Heo, K. Hsu, J. Hu, D. Jackson, C. Le, Y. Li, K. Lin, R. Lin, Z. Ma, A. Maddukuri, S. Mirchandani, D. Morton, T. Nguyen, A. O'Neill, R. Scalise, D. Seale, V. Son, S. Tian, E. Tran, A. E. Wang, Y. Wu, A. Xie, J. Yang, P. Yin, Y. Zhang, O. Bastani, G. Berseth, J. Bohg, K. Goldberg, A. Gupta, A. Gupta, D. Jayaraman, J. J. Lim, J. Malik, R. Mart'in-Mart'in, S. Ramamoorthy, D. Sadigh, S. Song, J. Wu, M. C. Yip, Y. Zhu, T. Kollar, S. Levine, and C. Finn. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint arXiv:2403.12945*, 2024.
- [26] O. M. Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu, et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024.
- [27] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- [28] R. Holmberg and O. Khatib. Development and control of a holonomic mobile robot for mobile manipulation tasks. *The International Journal of Robotics Research*, 19(11):1066–1074, 2000.
- [29] I. W. W. Group. Webxr device api, 2024. URL <https://www.w3.org/TR/webxr/>.
- [30] S. Jauhri, J. Peters, and G. Chalvatzaki. Robot learning of mobile manipulation with reachability behavior priors. *IEEE Robotics and Automation Letters*, 7(3):8399–8406, 2022.
- [31] C. Sun, J. Orbik, C. M. Devin, B. H. Yang, A. Gupta, G. Berseth, and S. Levine. Fully autonomous real-world reinforcement learning with applications to mobile manipulation. In *5th Annual Conference on Robot Learning*, 2021.
- [32] N. Yokoyama, A. Clegg, J. Truong, E. Undersander, T.-Y. Yang, S. Arnaud, S. Ha, D. Batra, and A. Rai. Asc: Adaptive skill coordination for robotic mobile manipulation. *arXiv preprint arXiv:2304.00410*, 2023.

- [33] J. Wu, R. Antonova, A. Kan, M. Lepert, A. Zeng, S. Song, J. Bohg, S. Rusinkiewicz, and T. Funkhouser. Tidybot: Personalized robot assistance with large language models. *Autonomous Robots*, 47(8):1087–1102, 2023.
- [34] R. Ni and A. H. Qureshi. Ntfields: Neural time fields for physics-informed robot motion planning. In *The Eleventh International Conference on Learning Representations*, 2022.
- [35] R.-Z. Qiu, Y. Hu, G. Yang, Y. Song, Y. Fu, J. Ye, J. Mu, R. Yang, N. Atanasov, S. Scherer, et al. Learning generalizable feature fields for mobile manipulation. *arXiv preprint arXiv:2403.07563*, 2024.
- [36] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- [37] M. Liu, Z. Chen, X. Cheng, Y. Ji, R. Yang, and X. Wang. Visual whole-body control for legged loco-manipulation. *arXiv preprint arXiv:2403.16967*, 2024.
- [38] P. Liu, Y. Orru, C. Paxton, N. M. M. Shafiqullah, and L. Pinto. Ok-robot: What really matters in integrating open-knowledge models for robotics. *arXiv preprint arXiv:2401.12202*, 2024.
- [39] A. Kumar, Z. Fu, D. Pathak, and J. Malik. Rma: Rapid motor adaptation for legged robots. *Robotics: Science and Systems XVII*, 2021.
- [40] J. Yang, C. Deng, J. Wu, R. Antonova, L. Guibas, and J. Bohg. Equivact: Sim (3)-equivariant visuomotor policies beyond rigid object manipulation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9249–9255. IEEE, 2024.
- [41] J. Yang, Z.-a. Cao, C. Deng, R. Antonova, S. Song, and J. Bohg. Equibot: Sim (3)-equivariant diffusion policy for generalizable and data efficient learning. *arXiv preprint arXiv:2407.01479*, 2024.
- [42] J. Albardaner, A. S. Miguel, N. García, and M. Dalmau. Sim-to-real gap in rl: Use case with tiago and isaac sim/gym. *arXiv preprint arXiv:2403.07091*, 2024.
- [43] M. Arduengo, A. Arduengo, A. Colomé, J. Lobo-Prat, and C. Torras. Human to robot whole-body motion transfer. In *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*, pages 299–305. IEEE, 2021.
- [44] R. K. Megalingam, V. S. Naick, S. K. Manoharan, and V. Sivananthan. Analysis of tiago robot for autonomous navigation applications. In *2021 Second International Conference on Electronics and Sustainable Communication Systems (ICESC)*, pages 257–261, 2021. doi: [10.1109/ICESC51422.2021.9532735](https://doi.org/10.1109/ICESC51422.2021.9532735).
- [45] S. Dass, W. Ai, Y. Jiang, S. Singh, J. Hu, R. Zhang, P. Stone, B. Abbatematteo, and R. Martin-Martin. Telemoma: A modular and versatile teleoperation system for mobile manipulation. *arXiv preprint arXiv:2403.07869*, 2024.
- [46] H. Robot. Stretch open-source mobile manipulator, 2024. URL <https://hello-robot.com/stretch-3-product>.
- [47] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.
- [48] D. Driess, F. Xia, M. S. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023.

- [49] Z. Fu, X. Cheng, and D. Pathak. Deep whole-body control: Learning a unified policy for manipulation and locomotion. In *6th Annual Conference on Robot Learning*, 2022.
- [50] J. Zhang, N. Gireesh, J. Wang, X. Fang, C. Xu, W. Chen, L. Dai, and H. Wang. Gamma: Graspability-aware mobile manipulation policy learning based on online grasping pose fusion. *arXiv preprint arXiv:2309.15459*, 2023.
- [51] A. Mandlekar, Y. Zhu, A. Garg, J. Booher, M. Spero, A. Tung, J. Gao, J. Emmons, A. Gupta, E. Orbay, et al. Roboturk: A crowdsourcing platform for robotic skill learning through imitation. In *Conference on Robot Learning*, pages 879–893. PMLR, 2018.
- [52] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song. Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024.
- [53] A. Tung, J. Wong, A. Mandlekar, R. Martín-Martín, Y. Zhu, L. Fei-Fei, and S. Savarese. Learning multi-arm manipulation through collaborative teleoperation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9212–9219, 2021. doi: [10.1109/ICRA48506.2021.9561491](https://doi.org/10.1109/ICRA48506.2021.9561491).
- [54] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. *arXiv preprint arXiv:2108.03298*, 2021.
- [55] J. Wong, A. Tung, A. Kurenkov, A. Mandlekar, L. Fei-Fei, S. Savarese, and R. Martín-Martín. Error-aware imitation learning from teleoperation data for mobile manipulation. In *Conference on Robot Learning*, pages 1367–1378. PMLR, 2022.
- [56] FIRST (For Inspiration and Recognition of Science and Technology). FIRST robotics competition, 2024. URL <https://www.firstinspires.org/robotics/frc>.
- [57] S. D. Specialties. Mk4 swerve module, 2024. URL <https://www.swervedrivespecialties.com/products/mk4-swerve-module>.