

DexWrist: A Robotic Wrist for Constrained and Dynamic Manipulation

Martin Peticco*, Gabriella Ulloa*, John Marangola, Pulkit Agrawal

Improbable AI Lab

Massachusetts Institute of Technology

{mpettico, gulloa, jmgola, pulkitag}@mit.edu

Abstract—We present the *DexWrist*, a compliant robotic wrist designed to advance robotic manipulation in highly-constrained environments, enable dynamic tasks, and speed up data collection. *DexWrist* is designed to be close to the functional capabilities of the human wrist and achieves mechanical compliance and a greater workspace as compared to existing robotic wrist designs. The *DexWrist* can supercharge policy learning by (i) enabling faster teleoperation and therefore making data collection more scalable; (ii) completing tasks in fewer steps which reduces trajectory lengths and therefore can ease policy learning; (iii) *DexWrist* is designed to be torque transparent with easily simulatable kinematics for simulated data collection; and (iv) most importantly expands the workspace of manipulation for approaching highly cluttered scenes and tasks. More details about the wrist can be found at: dexwrist.csail.mit.edu.

I. INTRODUCTION

Significant advances have been made in robotic manipulation with large-scale data collected using bimanual systems such as the ALOHA [9] and Franka Panda Dual Arm Setup [6]. Despite the progress, current manipulation remains limited to clean environments with minimal clutter and open workspaces, unlike real-world household environments with highly cluttered environments that impose significant workspace constraints on the robot. Imagine tasks like fetching an item from a packed refrigerator, plugging cables into a desktop computer kept under the desk, or placing a jar in a kitchen closet. A majority of off-the-shelf robotic arms are insufficient at operating in such workspaces. Furthermore, teleoperating robots in such environments is tedious due to the need to avoid contact with surrounding objects and to navigate the robot arm through narrow spaces.

Humans effortlessly operate in cluttered/constrained environments to perform fine manipulation. The general principle for such skill is for the robotic hand to approach and move the object in as many configurations as possible while minimizing the motion of the robot itself. One way to achieve this is to have more degrees of motion, be closer to the object being manipulated, and to co-locate joints. For humans, the wrist and a multi-finger hand provide such an ability. While the topic of multi-finger hands is an exciting one, it is beyond the scope of our paper, and we instead focus on the arm’s morphology and its effect on constrained manipulation and speed of data collection – the two significant problems to overcome for robotic operation in messy real-world environments.

Many of the most popular robotic arms, such as the UR and

Franka Panda, have 6-7 degrees of freedom and were historically designed for top-down manipulation. They have a wrist with three serially located joints that are unable to move in a constrained space (e.g., cluttered refrigerators), nor can they perform dynamic moves inherent to human daily activities. Such serial design also makes teleoperation less intuitive due to the kinematic differences from the human wrist, which has co-located joints, and the serial joints also reduce the number of kinematic solutions [16, 22, 43]. Additionally, commonly used coupled kinematic chain designs like differential joint wrists are complex to model and simulate [31]. Lastly, to meet high torque specifications, robotic systems typically employ large gear reductions, which result in insufficient compliance for reacting to external forces [16].

We propose *DexWrist* that overcomes these limitations by using decoupled parallel actuation and compliant actuators to improve manipulation in constrained and dynamic scenarios and speed up data collection. *DexWrist* is capable of activities of daily living (ADLs), which impose functional constraints on the robotic wrist’s size, strength, compliance, speed, bandwidth, precision, and range of motion (ROM). Torque is critical for all manipulation, from opening jars to lifting, and the wrist must sustain loads during full arm motion. Compliance is essential for adapting to unexpected environmental factors like collisions or uneven surfaces, preventing damage to the robotic system and its surroundings. ADL completion at human-like speeds and the ability to rapidly vary torques (bandwidth) are necessary for dynamic tasks. Lastly, a wide range of motion is paramount for operations in highly confined workspaces, such as in human households.

Why does this matter for learning policies? *DexWrist* is compatible with a large number of robotic arms. We use the combination of a low-cost AgileX arm and the *DexWrist* as our experimental platform. Our results show that *DexWrist* enables data collection for highly-constrained manipulation in cluttered and hard-to-reach spaces. Furthermore, teleoperation with *DexWrist* is substantially more intuitive and results in faster data collection with shorter trajectories. While our focus in this work is the design of *DexWrist* – a tool for robot learning – we hope that *DexWrist* will attract wide interest from the robot learning community interested in going after cluttered manipulation and scaling up data collection.




Fig. 1: We present *DexWrist*, a robotic wrist that allows for constrained and dynamic manipulation and speeds up teleoperated data collection and makes teleoperation more intuitive.

II. PRIOR WORKS

A. Serial Wrists.

Most commercially available robotic arms have an integrated serial wrist, such as the UR5e [12], Franka Panda [6], and AgileX PiPER [11]. There is also a variety of standalone serial wrists like the Montagnani switchable stiffness wrist [32] and the Chirikjian spherical stepper motor [20]. The main issue with many serial wrists is that they are not only large and non-back-driveable, but their kinematic differences from the human wrist complicate constrained manipulation and limit the number of feasible solving methods for teleoperation demos.

B. Coupled Parallel Wrists.

The Omni-Wrist [37], Carpal Robotic wrist [14], and Damerla prosthetic wrist [21] have made significant size improvements. However, modeling and simulation complexities are still pertinent due to their coupled nature.

C. Decoupled Parallel Wrists.

The Agile Eye [25] and Negrello soft wrists [33] are two examples of decoupled parallel wrists. While their kinematics closely emulate human wrist function, sizing presents itself as a limitation with this kinematic structure. The *DexWrist* aims to minimize the size gap while maintaining a decoupled kinematic configuration.

D. U-Joint Style Wrists.

Robotic wrists on commercial platforms such as the Unitree H1-2 and the GALAXEA R1 have their first motor stationary relative to the base of the wrist, and the second motor in its entirety rotated by the first motor. The end effector is mounted to the output of the second motor. This has the downside of a higher moment of inertia due to the first motor moving the weight of the second motor, limiting the range of dynamic tasks. *DexWrist*, however, has both motors mounted stationary relative to the base of the wrist, greatly reducing the inertia of the end effector and allowing for more dynamic tasks.

III. FUNCTIONAL REQUIREMENTS

Our first goal is to characterize functional requirements and the form factor of a robotic wrist that can perform small workspace manipulation while retaining daily dynamic task capabilities. Identifying average human exertion attributes for these characteristics, which are summarized in Table I, provides a path for developing hardware that can achieve ADL manipulation.

A. Torque, Load Capacity, and Compliance

Two similar studies indicated that the maximum wrist flexion/extension (F/E) and radial/ulnar (R/U) deviation (Figure 2) torques are 4.6-11.9 Nm and 4.7-10.8 Nm, respectively [42, 40]. The F/E and R/U torques cover the same range of values, allowing the critical assumption that these degrees of freedom may be designated similar desired torque requirements. However, while both studies provide relevant insight into human wrist maximum performance, we aim to focus on **average** human wrist torque outputs required for ADLs. Another study indicates that a torque of 3 Nm was sufficient to complete 93% of ADL tasks [34, 15], and already takes into account the inertial torques of the human hand and the object of manipulation. We will be using objects of manipulation and an end effector (AgileX PiPER Gripper [2]) with similar inertial torques, allowing us to use the same final value of at least 3 Nm for both F/E and R/U. Torques for P/S were not investigated as it will be in series with F/E and R/U, and most traditional robot arms include this DOF.

Static strength is especially important for tasks requiring a locked wrist during full arm motion, such as lifting full grocery bags or a gallon of milk, each weighing roughly 4 kg. Including a maximum 1 kg end effector, a robotic wrist must sustain 5 kg of load in each axial direction.

Compliance integration may be accomplished by minimizing the torque needed to back-drive the actuators. At rest, human wrists require low forces to reorient the hand. For this, the robotic wrist must also allow the end effector to be reoriented with ease, warranting no more than 5 N of force. This translates to a back-driveability torque of at most 0.4 Nm.




Fig. 2: **Left:** Human wrists have 3 degrees of freedom: flexion/extension (F/E), radial/ulnar (R/U) deviation, and pronation/supination (P/S). Anatomically, the F/E and R/U degrees of freedom are in parallel and are preceded by P/S in series. **Right:** *DexWrist* DOFs mirroring human wrist.

B. Speed, Bandwidth, Kinematics, and Precision

To properly characterize wrist speed, past work designed a game where human subjects use isolated wrist motion to accomplish tasks [39]. The study revealed peak wrist movement speeds were 10-53.3 RPM, providing a target range for robotic wrist speed quantities.

Studies examining human wrist responses to force perturbations and visual position targets provide useful information for characterizing bandwidth [29, 23]. Specifically, the long-latency reflex refers to a conscious wrist response to external stimuli. These findings indicate that the long-latency reflex occurs within a range of 50-100 ms, correlating to bandwidth frequencies of 10-20 Hz.

A study recorded human F/E and R/U angles using an electrogoniometer during ADL completion [36]. For tasks involving personal hygiene and food preparation, it was found that 40° each of flexion and extension, 10° of radial deviation, and 30° of ulnar deviation is a reasonable representation of the range necessary for ADL completion as it was sufficient to complete 22 of the 24 tasks.

To get wrist angular precision, [28] and [27] both indicated minimum wrist angular precisions of 3.47° and 4.58°, respectively.

C. Size and Weight

The benefits of human wrist compactness are maintained in our size constraints defined by anthropometric data conducted by NASA [3]. The 95th percentile value of male human wrist measurements provides a wrist width and height of 61.4 mm. The actuators and power sources of our design are to be located within the forearm as it typically is unoccupied space in commercially available humanoid robots. Thus, the length of this robotic wrist is dictated by the maximum expected male forearm length of 349 mm [3]. To allow space for the elbow joint, wrist length is limited to 174.5 mm.

This robotic wrist would be attached to a robot arm in the same fashion any end effector would. We would like to model our desired weight using a typical example of a smaller gripper, such as the Robotiq gripper [1]. This allows us to arrive to a final desired weight of approximately 1 kg. This is a standard weight that typical robot arms, such as the UR5e [12], can handle.

IV. WRIST DESIGN

A 2-DOF robotic wrist was designed according to the aforementioned functional requirements. Two custom stepped planet compound planetary gearboxes separately drive each independent DOF of a decoupled parallel kinematic mechanism (PKM).

A. Quasi-Direct Drive Stepped Planet Compound Planetary Gearbox

The actuator selection and transmission design processes are driven by the torque, speed, compliance, and size requirements. Actuator selection must be optimized for power density in order to minimize transmission size and gear reduction, allowing us to satisfy the driving requirements simultaneously. BLDC motors are among the most power-dense actuators available and have been widely used in the development of robotic and prosthetic wrists [21]. The torque, speed, and transmission size requirements for this application indicate that the minimum power density required is 0.405 W/cm³. The CubeMars GL40 motor [7], with a power density of 1.50 W/cm³, was selected as it best minimized the projected transmission ratio while falling within size requirements compared to other commercially available BLDC motors.

A stepped planet compound planetary gear transmission was compared alongside ball screw, 2-stage planetary gearbox, and linear spur gear transmissions. The unique combination of load transparency, back-driveability with low gear reductions, and size efficiency finalized our choice of pursuing the stepped planet compound planetary gearbox transmission option. A custom transmission, depicted in Figure 3, was designed and optimized for this application since commercially available planetary gearboxes are too large. The smallest commercially produced planetary gearbox from Maxon [5] providing the minimum gear reduction is 31.7 mm in length, which exceeds the size requirements by 12.8 mm.

A 13:1 gear reduction is implemented to clear the 3 Nm torque requirement with the selected motor's rated torque of 0.25 Nm. With a specific load inertia over 100,000 times greater than the reflected motor inertia, this gearbox is considered quasi-direct drive (QDD) [41]. The Lewis Equation (eq.1)[8, 13] was used to determine tooth bending stress σ_t , taking into account tangential load W_t , gear module m_n , face width b , Lewis Factor Y (from engineering charts), and a de-rating factor K_d . The de-rating factor K_d quantifies elements that may make the load higher than it should be for the load being transmitted, taking into account application factor K_a , load distribution factor K_m , size factor K_s , and dynamics factor K_v . These constants address the physical imperfections and effects of driving equipment, actuator type, shaft misalignment, material properties, manufacturing processes, etc. The tooth bending stress must be within a factor of safety (FOS) of 3 relative to its yield strength σ_y , allowing us to optimize for face width with this constraint. 1045 Carbon Steel was selected as the gear material because of its strength ($\sigma_y = 450$ MPa), machinability, and wide use in general gear applications.

TABLE I: Comparison of Desired and Achieved Functional Requirements.

Functional Requirement	Desired	Range	Ours	Pass
Rated Active Torque (Nm)	3	≥ 3 [34, 15]	3.75 ± 0.05	✓
Back-driveability Torque (Nm)	0.4	< 0.4	0.33 ± 0.06	✓
Load Capacity X/Y/Z (kg)	5	≥ 5	5	✓
Rated Active Speed (RPM)	50	10 to 53.3 [39]	96.6 ± 9.4	✓
Bandwidth (Hz) @ 3.75Nm	20	10 to 20 [29, 23]	10.15 ± 1.34	✓
Angular Precision (°)	3.5	0 to 3.5 [28, 27]	1.65	✓
F/E ROM (°)	80	-40 to 40 [36]	-40 to 40	✓
R/U ROM (°)	40	-10 to 30 [36]	-40 to 40	✓
Width (mm)	61.4	51.5 to 61.4 [3]	64	~
Height (mm)	61.4	51.5 to 61.4 [3]	66.5	✗
Length (mm)	174.5	± 5 [3]	178.2	✓
Weight (kg)	1	0 to 1 [1]	0.97	✓




Fig. 3: QDD Stepped Planetary Compound Gearbox with 13:1 transmission.

$$\sigma_t = \frac{W_t m_n}{Y_b} \cdot K_d, \quad K_d = \frac{K_a K_m K_s}{K_v} \quad (1)$$

The AS5047P adapter board encoder [4] measures the motor input position while the Mjbots Moteus-n1 brushless motor controller [10] communicates with the CubeMars GL40. The motor controller boards for both degrees of freedom are daisy-chained together via CAN and connect to an external power supply providing 12V for the actuators.

B. 2-(R, RR) Parallel Kinematic Mechanism

PKMs best emulate the F/E and R/U human wrist degrees of freedom and provide an efficient way of localizing degrees of freedom. This claim is directly investigated through the workspace comparison between the AgileX PiPER Arm with its serial wrist and the AgileX PiPER with a parallel wrist. One example of a task involving small workspace manipulation is reaching into a spice cabinet, where some kitchen configurations have deep angled cabinets. To represent this scenario in simulation, each arm was imported into the PyBullet environment along with a rendering of a deep angled cabinet. Inverse kinematics was performed for uniformly distributed target points across the inner cabinet workspace. If a joint pose exists for a tested target point without collisions with itself or the cabinet, the target point is considered reachable. This simulation found that the number of reachable points increased by 88% when the parallel wrist is used instead of the serial wrist. Figure 4 illustrates that the parallel wrist workspace

encapsulates the serial wrist workspace and is also capable of reaching more surrounding points.

There are three main driving requirements in designing a decoupled 2-DOF PKM having high torque capabilities within human form factor: ROM, size, and load capacity. Figure 5 shows the 2-(R, RR) PKM designed for this robotic wrist, an adaptation of the 2-DOF Agile Eye [24], with a 1:1 transmission ratio. To ensure the linkage geometry has sufficient strength, four load cases were analyzed: 5 kg load in each axial direction (X,Y, and Z independently) and full 3 Nm exertion by both degrees of freedom. A minimum shaft diameter of 6mm is necessary for FOS = 3 and the smallest readily available off-the-shelf bushings have a 9 mm outer diameter and are 6 mm thick. Using these hardware specifications, linkage geometries were optimized to achieve the desired ROM without self-collisions and to fit within the prescribed size envelope while maintaining strength. The chosen linkage material is 17-4 PH Stainless Steel ($\sigma_y = 1345$ MPa), resulting in a minimum FOS of 3.26 for the linkages under the four load cases.

V. TELEOPERATION FRAMEWORK AND SYSTEM INTEGRATION

A. Integration Setup

We use the AgileX PiPER as an experimental platform to evaluate the overall dexterity and performance of our wrist design for manipulation. The PiPER is a robotic arm with highly back-driveable joints and six degrees of freedom, but it lacks orthogonal roll-pitch movement for the last two degrees of freedom, making human-like wrist circumduction difficult. In light of this limitation, we remove the last two joints of the AgileX and replace them with our 2-DOF wrist design, maintaining a total of six degrees of freedom. The AgileX PiPER gripper, an ALOHA-style gripper, is mounted onto the wrist as the end effector.

B. Controller and Pipeline Details

We developed a comprehensive policy learning pipeline that enables seamless demonstration collection, training, and deployment across diverse teleoperation controllers, wrist configurations, and end-of-arm tooling. During teleoperation, absolute end-effector pose targets $\mathcal{T}_w^{ee} \in \text{SE}(3)$ are obtained




Fig. 4: **Left:** *DexWrist* wrist simulation. **Middle:** Serial wrist simulation. **Right:** Workspace comparison plot illustrating reachability of each arm configuration within the deep angled cabinet. The *DexWrist* improved constrained workspace reachability by 88% when compared to the AgileX wrist.




Fig. 5: **Left:** 2-(R, RR) PKM isometric view with red dotted lines depicting DOF rotational axes and pivot point. **Middle:** Side view highlighting kinematic chain in dark blue (RR). **Right:** Front view highlighting kinematic chain in gray (R).



Fig. 6: The three tasks shown to evaluate the performance of the wrist as compared to the stock AgileX arm.

from a teleoperation controller which are converted to joint position targets via inverse kinematics and sent to a lower level joint stiffness controller operating at a higher control frequency. More details regarding the policy learning pipeline can be found in Appendix A.

C. User Study Task Descriptions

A user study is conducted to measure the performance at recording demonstrations on three separate tasks using both the original AgileX PiPER arm and the modified arm with *DexWrist*. Images of the tasks being performed are shown in Figure 6.

- 1) *Picking from a Cluttered Refrigerator*: Pick up a highly occluded cup from deep inside a fridge while ensuring surrounding objects are not knocked over.
- 2) *Unplugging*: Reach through the narrow gap between a monitor and a desktop computer to unplug a USB cable.

- 3) *Picking from a drawer*: Pick up a cup from deep inside a drawer.

D. Scene Randomization for Cluttered Refrigerator

The scene and robot configuration are systematically varied during both teleoperation and evaluation. Initial position of the robot end-effector is fixed across all trajectories. All objects inside the fridge were subject to relatively small amounts of randomization every reset with the target object being subject to slightly more variation in initial position and rotation relative to the other objects.

VI. EXPERIMENTAL RESULTS

A. Characterization Experiments

1) *Torque and Bandwidth*: We placed a Vernier Go Direct Force and Acceleration Sensor 70 mm away from the pivot point to measure the force the *DexWrist* could exert. The rated torque output was calculated to be 3.75 ± 0.05 Nm (10 trials were collected for all experiments). We simultaneously calculated the bandwidth using the rising time (t_r), the time taken to reach 90% torque from 10% torque, with the $B(\text{Hz}) = 0.35/t_r(\text{s})$ relationship. The resulting bandwidth is 10.15 ± 1.34 Hz, which is on the lower end of the desired range.

2) *Back-driveability Torque*: In a similar setup as before, the Vernier Go Direct Force and Acceleration Sensor was used to measure the force required to back-drive the *DexWrist* as the wrist was lowered onto the sensor. The torque necessary to back-drive the robotic wrist is 0.33 ± 0.06 Nm.

3) *Load Capacity*: To validate the strength of the *DexWrist*, the Vernier Go Direct Force and Acceleration Sensor was used to push against the wrist hard stops with force equivalent to 5 kg. The structural skeleton of the wrist was able to sustain the required load capacity without damage.

4) *Speed, Angular Precision, and Range of Motion*: We recorded the end effector motion of the *DexWrist* when moved between its motion limits. The Vernier Video Analysis software was used to track the end effector and calculate its speed, final position, and range of motion. The resulting rated speed greatly surpassed our requirements at 96.6 ± 9.4 RPM, the angular precision was 1.65 degrees, and the range of motion for both F/E and R/U were confirmed to be -40 to 40 degrees.

5) *Size and Weight*: The *DexWrist* length fits within the designated requirements. Due to the spherical bearings necessary for the driving links, the height and width were above the required size. However, the width is only 4% larger than the target value. The assembly weighs 0.97 kg.

B. Teleoperation in Constrained Environments

The robot was teleoperated to perform the task both with *DexWrist* and the stock PiPER wrist in the three constrained manipulation settings as outlined in V-C. As shown in Figure 7, using *DexWrist* to record demonstrations in constrained spaces significantly reduced the trajectory length and resets per successful trajectory compared to the stock PiPER Arm with the default wrist.

Across all tasks, the *DexWrist* greatly reduced both the average number of resets and the average operator time. Task videos are available here.

C. Behavioral Cloning

1) *Method*: To evaluate the impact of wrist design on manipulation performance, we trained diffusion policies [18] on 141 demonstrations recorded separately on the AgileX PiPER equipped with both the default wrist and the *DexWrist* (282 demonstrations total). For both system configurations, we train CNN-based diffusion policies with identical hyperparameters, operating at 20Hz using a DDIM sampler [38] to perform the task. More details on the BC methodology and hyperparameters can be found in Appendix B.

2) *Task Specification*: Retrieve an occluded flattened soda can from deep within a cluttered refrigerator and place it on the table. Failure occurs if any object is knocked over, the camera disconnects, or the refrigerator is displaced. Specifically, the task attempts to retrieve a flattened soda can which requires reaching deep into the back of the refrigerator, positioning the gripper such that the fingers are parallel to the back wall of the fridge (i.e., the gripper’s opening axis is nearly orthogonal to the refrigerator’s rear wall). This awkward orientation makes it difficult to avoid disturbing nearby objects during retrieval. Limited vertical clearance between shelves, dividing rack, and bottom tray creates minimal tolerance for arm movement. Initial object occlusion further increases task difficulty.

3) *Evaluation*: We evaluated the diffusion policies trained for each respective system at the same six epochs (75, 150, 225, 300, 375, 750). For each checkpoint, we collect 15 roll-outs in the real environment (90 trials per system), randomly resetting the objects in the scene before each iteration. More details can be found in Table IV

4) *Results*: The best performing policy trained for the AgileX + *DexWrist* combination exhibited a 50% relative improvement in success rate over the policy trained for the default AgileX system. Qualitatively, it was observed that actions taken by the system with the default wrist led to a higher frequency of catastrophic failures, characterized as failure events in which multiple objects were violently displaced, including the refrigerator itself, in some cases. The time taken to complete the task was recorded for successful trials. We

found that the AgileX + *DexWrist* completed the task **3.24x** faster than the default configuration, on average.

TABLE III: Autonomous task completion time statistics for successful trials using the best checkpoint for each respective system. $N = 15$ for both configurations.

System	Policy Task Completion Time (s)			
	Mean	Min	Max	
AgileX + Default Wrist	91.0 \pm 7.9	55.2	134.2	
AgileX + DexWrist (Ours)	28.1 \pm 2.2	20.5	49.0	

VII. DISCUSSION AND FUTURE WORK

We empirically show that *DexWrist* reduces both the total time required (i.e., more intuitive teleoperation) and the average trajectory length of successful demonstrations provided by human teleoperators in constrained spaces. Reduction in the number of environment resets and total teleoperation time leads to a more efficient and therefore more scalable data collection process in constrained spaces, which is critical for scaling up data collection in consumer settings. The increased workspace of the robot with *DexWrist* enables completion of tasks that were not previously possible with the default wrist configuration in constrained environments. Lastly, torque transparency and backdrivability makes *DexWrist* capable of performing dynamic tasks via torque control. Videos of *DexWrist* performing highly dynamic tasks with human-level wrist dexterity can be found here.

The behavioral cloning experiments conducted in the cluttered refrigerator setting demonstrate the advantage of using the *DexWrist* for policy learning in confined spaces where traditional top-down manipulation approaches are infeasible or require task-specific robot integration. As shown in Figure 8, the *DexWrist* demonstrated a modest but consistent improvement in policy success rate compared to the default wrist configuration in the cluttered refrigerator. Furthermore, the *DexWrist* policies completed the task **3.24x** faster than the default wrist configuration. While this performance improvement can be partially attributed to the observed reduction in trajectory length during teleoperation, we attribute much of this performance improvement to the human-like kinematics of *DexWrist*.

The *DexWrist*’s performance advantage in constrained environments stems from its anthropomorphic kinematic design creating solutions that are inherently more robust and natural. In human-designed spaces like kitchens and refrigerators, the *DexWrist*’s human-like joint constraints naturally generate inverse kinematics solutions that closely match the control envelope of human wrist configurations, leading to a structured action space that aligns with human demonstration patterns, and environments making it a natural choice for task space control in constrained spaces. On the other hand, the serial kinematic chain of the AgileX with the default wrist forces the policy to more carefully plan points in the task space as it must orchestrate all the DoFs of the robot in a manner which




Fig. 7: Demonstrations recorded from successful trajectories. Resets were performed in the event of a severe robot collision, surrounding objects being knocked over, or a failed grasp. For each configuration, $N \geq 40$.




Fig. 8: Task success rates for the AgileX PiPER robot with stock wrist (40%) versus DexWrist (60%). For each system, we report the highest success rate among all evaluated checkpoints. Error bars show the binomial standard error based on 15 trials of the best checkpoint only.

does not knock over objects in the vicinity of the end effector, causing a failure.

Unlike conventional robotic wrists, such as the AgileX, each DoF of the DexWrist is independently controlled by one actuator. The DexWrist’s decoupled parallel kinematic chain eliminates the need for complex joint coordination. Consequently, in policy rollout, we observe that the DexWrist takes more direct actions towards the can without needing to coordinate a serial kinematic chain through inverse kinematics; this is in contrast to the original arm, which tends to pause and get stuck during this pre-approach step. Finally, since large-scale human behavior data is more readily available and scalable than collecting hardware-platform-specific demon-

strations, improving a robotic arm’s ability to successfully complete tasks using an end-effector action space is critical for scaling up policy learning in human-centric environments.

Future research directions include investigating synergies between our wrist design and reinforcement learning approaches for dynamic manipulation tasks via torque control. Additionally, early experiments have signaled that the human-like wrist design of the DexWrist closes the embodiment gap to the human wrist morphology better than traditional robotic wrists do. We hypothesize that this applies to end-effector action spaces, both in constrained and unconstrained environments. In the future, we hope to systematically evaluate this hypothesis by training a single policy on demonstrations collected naturally by humans using the UMI [19] and transferring it to both traditional robotic wrists and the DexWrist.

ACKNOWLEDGMENTS

We thank the members of the Improbable AI lab—especially Nolan Fey and Branden Romero—for the helpful discussions and feedback on the paper. We are grateful to MIT Supercloud and the Lincoln Laboratory Supercomputing Center for providing HPC resources. We acknowledge AgileX for providing us a custom version of their PiPER firmware that allowed operation of the arm with the final two DOFs removed. We are grateful to MIT Supercloud and the Lincoln Laboratory Supercomputing Center for providing HPC resources. We acknowledge support from ONR MURI under grant number N00014-22-1-2740. Finally, we acknowledge support from Sony and Toyota Research Institute. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized

to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

REFERENCES

- [1] 2F-85 and 2F-140 Robot Grippers from Robotiq | Electromate Inc, . URL <https://www.electromate.com/2f-85-and-2f-140-grippers/>.
- [2] AgileX Piper Gripper | MYBOTSHOP.DE, € 285,95, . URL <https://www.mybotshop.de/AgileX-Piper-Gripper>.
- [3] ANTHROPOMETRY AND BIOMECHANICS, . URL <https://msis.jsc.nasa.gov/sections/section03.htm>.
- [4] AS5047P ADAPTERBOARD, . URL <https://www.digikey.com/en/products/detail/ams-osram-usa-inc/AS5047P-ADAPTERBOARD/5452344>.
- [5] Compact Drives, Motors, Gears, Sensors | maxon group, . URL https://www.maxongroup.us/maxon/view/category/gear?etcc_cu=onsite&etcc_med_onsite=Product&etcc_cmp_onsite=GPX+Planetengetriebe&etcc_plc=Overview-Page-Gears&etcc_var=%5bch%5d%23de%23_d_&target=filter&filterCategory=planetary&q=GPX.
- [6] Franka Emika Panda robot - RoboDK, . URL <https://robodk.com/robot/Franka/Emika-Panda>.
- [7] GL40 KV70 - GL - GL Series - Gimbal Motor - Product Series, . URL <https://www.cubemars.com/goods-1132-GL40+KV70.html>.
- [8] Lewis Factor Equation for Gear Tooth Calculations, . URL <https://www.engineersedge.com/gears/lewis-factor.htm>.
- [9] Mobile ALOHA: Learning Bimanual Mobile Manipulation with Low-Cost Whole-Body Teleoperation, . URL <https://mobile-aloha.github.io/>.
- [10] moteus-n1, . URL <https://mjbots.com/products/moteus-n1>.
- [11] PiPER, . URL <https://global.agilex.ai/products/piper>.
- [12] UR5e Lightweight, versatile cobot, . URL <https://www.universal-robots.com/products/ur5e/>.
- [13] Saif Ali. *DUDLEY'S HANDBOOK OF PRACTICAL GEAR DESIGN and MANUFACTURE*. URL https://www.academia.edu/45138343/DUDLEYs_HANDBOOK_OF_PRACTICAL_GEAR_DESIGN_and_MANUFACTURE.
- [14] James Bryan R. Alvarez, Miguel Antonio G. Apolinar, Gerardo L. Augusto, and Laurence A. Gan Lim. Design and Development of a Carpal Wrist Robotic Manipulator. In *2019 IEEE 11th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment, and Management (HNICEM)*, pages 1–5, November 2019. doi: 10.1109/HNICEM48295.2019.9073534. URL <https://ieeexplore.ieee.org/document/9073534>.
- [15] Will Anderton, Scott Tew, Spencer Ferguson, Joshua Hernandez, and Steven K. Charles. Movement preferences of the wrist and forearm during activities of daily living. *Journal of Hand Therapy: Official Journal of the American Society of Hand Therapists*, 36(3):580–592, 2023. ISSN 1545-004X. doi: 10.1016/j.jht.2022.07.003.
- [16] Neil M. Bajaj, Adam J. Spiers, and Aaron M. Dollar. State of the Art in Artificial Wrists: A Review of Prosthetic and Robotic Wrist Design. *IEEE Transactions on Robotics*, 35(1):261–277, February 2019. ISSN 1941-0468. doi: 10.1109/TRO.2018.2865890. URL <https://ieeexplore.ieee.org/document/8624352/?arnumber=8624352>. Conference Name: IEEE Transactions on Robotics.
- [17] Remi Cadene, Simon Alibert, Alexander Soare, Quentin Gallouedec, Adil Zouitine, and Thomas Wolf. Lerobot: State-of-the-art machine learning for real-world robotics in pytorch. <https://github.com/huggingface/lerobot>, 2024.
- [18] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 2024.
- [19] Cheng Chi, Zhenjia Xu, Chuer Pan, Eric Cousineau, Benjamin Burchfiel, Siyuan Feng, Russ Tedrake, and Shuran Song. Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots, 2024. URL <https://arxiv.org/abs/2402.10329>.
- [20] G.S. Chirikjian and D. Stein. Kinematic design and commutation of a spherical stepper motor. *IEEE/ASME Transactions on Mechatronics*, 4(4):342–353, December 1999. ISSN 1941-014X. doi: 10.1109/3516.809513. URL <https://ieeexplore.ieee.org/document/809513>.
- [21] Revanth Damerla, Kevin Rice, Daniel Rubio-Ejchel, Maurice Miro, Enrico Braucher, Juliet Foote, Issam Bourai, Aaryan Singhal, Kang Yang, Hongju Guang, Vasil Iakimovitch, Evelyn Sorgenfrei, and Shorya Awatar. Design and Testing of a Novel, High-Performance Two DoF Prosthetic Wrist. *IEEE Transactions on Medical Robotics and Bionics*, 4(2):502–519, May 2022. ISSN 2576-3202. doi: 10.1109/TMRB.2022.3155279. URL <https://ieeexplore.ieee.org/document/9722901/?arnumber=9722901>. Conference Name: IEEE Transactions on Medical Robotics and Bionics.
- [22] Yu-Heng Deng and Jen-Yuan (James) Chang. Human-Like Posture Correction for Seven-Degree-of-Freedom Robotic Arm. *Journal of Mechanisms and Robotics*, 14(024501), September 2021. ISSN 1942-4302. doi: 10.1115/1.4051842. URL <https://doi.org/10.1115/1.4051842>.
- [23] Christopher J. Forgaard, Ian M. Franks, Dana Maslovat, Laurence Chin, and Romeo Chua. Voluntary reaction time and long-latency reflex modulation. *Journal of Neurophysiology*, 114(6):3386–3399, December 2015. ISSN 0022-3077. doi: 10.1152/jn.00648.2015. URL <https://journals.physiology.org/doi/full/10.1152/jn.00648.2015>. Publisher: American Physiological Society.
- [24] Clement M. Gosselin and François Caron. Two degree-of-freedom spherical orienting device, October 1999. URL <https://patents.google.com/patent/US5966991A/en>.
- [25] C.M. Gosselin, E. St. Pierre, and M. Gagne. On the development of the Agile Eye. *IEEE Robotics*

- & *Automation Magazine*, 3(4):29–37, December 1996. ISSN 1558-223X. doi: 10.1109/100.556480. URL <https://ieeexplore.ieee.org/document/556480>.
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. URL <http://arxiv.org/abs/1512.03385>.
- [27] Matthew E. Holman, Gary Goldberg, and Benjamin J. Darter. Accuracy and precision of a wrist movement when vibrotactile prompts inform movement speed. *Somatosensory & Motor Research*, 37(3):165–171, July 2020. ISSN 0899-0220, 1369-1651. doi: 10.1080/08990220.2020.1765766. URL <https://www.tandfonline.com/doi/full/10.1080/08990220.2020.1765766>.
- [28] Jonggi Hong, Lee Stearns, Jon Froehlich, David Ross, and Leah Findlater. Evaluating angular accuracy of wrist-based haptic directional guidance for hand movement. In *Graphics Interface*, pages 195–200, 2016. URL <https://graphicsinterface.org/wp-content/uploads/gi2016-25.pdf>.
- [29] Sho Ito and Hiroaki Gomi. Visually-updated hand state estimates modulate the proprioceptive reflex independently of motor task requirements. *eLife*, 9:e52380. ISSN 2050-084X. doi: 10.7554/eLife.52380. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7108863/>.
- [30] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [31] Shihua Li, Sen Wang, Haoran Li, Yongjie Wang, and Shuang Chen. Type Synthesis of Fully Decoupled Three Translational Parallel Mechanism with Closed-Loop Units and High Stiffness. *Chinese Journal of Mechanical Engineering*, 36(1):113, October 2023. ISSN 2192-8258. doi: 10.1186/s10033-023-00908-3. URL <https://doi.org/10.1186/s10033-023-00908-3>.
- [32] Federico Montagnani, Marco Controzzi, and Christian Cipriani. Preliminary design and development of a two degrees of freedom passive compliant prosthetic wrist with switchable stiffness. In *2013 IEEE International Conference on Robotics and Biomimetics (RO-BIO)*, pages 310–315, December 2013. doi: 10.1109/ROBIO.2013.6739477. URL <https://ieeexplore.ieee.org/document/6739477>.
- [33] Francesca Negrello, Sariah Mghames, Giorgio Grioli, Manolo Garabini, and Manuel Giuseppe Catalano. A Compact Soft Articulated Parallel Wrist for Grasping in Narrow Spaces. *IEEE Robotics and Automation Letters*, 4(4):3161–3168, October 2019. ISSN 2377-3766. doi: 10.1109/LRA.2019.2925304. URL <https://ieeexplore.ieee.org/document/8746558>.
- [34] Autumn Pando and Dr Steven Charles. Characterization of Wrist Kinetics during Activities of Daily Living. *Journal of Undergraduate Research*, 2013(1), September 2013. URL <https://scholarsarchive.byu.edu/jur/vol2013/iss1/1966>.
- [35] Younghyo Park and Pulkit Agrawal. Using apple vision pro to train and control robots, 2024. URL <https://github.com/Improbable-AI/VisionProTeleop>.
- [36] Jaiyoung Ryu, William P. Cooney, Linda J. Askew, Kai-Nan An, and Edmund Y. S. Chao. Functional ranges of motion of the wrist joint. *The Journal of Hand Surgery*, 16(3):409–419, May 1991. ISSN 0363-5023. doi: 10.1016/0363-5023(91)90006-W. URL <https://www.sciencedirect.com/science/article/pii/036350239190006W>.
- [37] J. Sofka, V. Skormin, V. Nikulin, and D.J. Nicholson. Omni-Wrist III - a new generation of pointing devices. Part II. Gimbals systems - control. *IEEE Transactions on Aerospace and Electronic Systems*, 42(2):726–734, April 2006. ISSN 1557-9603. doi: 10.1109/TAES.2006.1642585. URL <https://ieeexplore.ieee.org/document/1642585>.
- [38] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *CoRR*, abs/2010.02502, 2020. URL <https://arxiv.org/abs/2010.02502>.
- [39] Lev Vaisman, Laura Dipietro, and Hermano Igo Krebs. A Comparative Analysis of Speed Profile Models for Wrist Pointing Movements. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*, 21(5):756–766, September 2013. ISSN 1534-4320. doi: 10.1109/TNSRE.2012.2231943. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4689593/>.
- [40] Ting Xia and Laura A. Frey Law. Wrist joint torque-angle-velocity performance capacity envelope evaluation and modelling. *International Journal of Human Factors Modelling and Simulation*, 5(1):33, 2015. ISSN 1742-5549, 1742-5557. doi: 10.1504/IJHFMS.2015.068120. URL <http://www.inderscience.com/link.php?id=68120>.
- [41] Jaeha Yang, Junyoung Moon, Jaewook Ryu, Jehyeok Kim, Kimoon Nam, Sungjin Park, Yoosun Kim, and Giuk Lee. Design of a Quasi-Direct Drive Actuator with Embedded Pulley for a Compact, Lightweight, and High-Bandwidth Exosuit. *Actuators*, 12(1):21, January 2023. ISSN 2076-0825. doi: 10.3390/act12010021. URL <https://www.mdpi.com/2076-0825/12/1/21>. Number: 1 Publisher: Multidisciplinary Digital Publishing Institute.
- [42] Yuichi Yoshii, Hiroshi Yuine, Ohashi Kazuki, Wen-lin Tung, and Tomoo Ishii. Measurement of wrist flexion and extension torques in different forearm positions. *BioMedical Engineering OnLine*, 14:115, December 2015. ISSN 1475-925X. doi: 10.1186/s12938-015-0110-9. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4676844/>.
- [43] Chao Yu, Minghe Jin, and Hong Liu. An analytical solution for inverse kinematic of 7-DOF redundant manipulators with offset-wrist. In *2012 IEEE International Conference on Mechatronics and Automation*, pages 92–97, August 2012. doi: 10.1109/ICMA.2012.6282813. URL <https://ieeexplore.ieee.org/abstract/document/6282813>. ISSN: 2152-744X.
- [44] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in

neural networks. *CoRR*, abs/1812.07035, 2018. URL <http://arxiv.org/abs/1812.07035>.

APPENDIX A CONTROL AND PIPELINE DETAILS

Given an SE(3) pose target generated either from a teleoperation controller or a policy $\pi_\theta(\cdot|o_{t-T_p:t})$, a differential inverse kinematics problem, formulated as a constrained quadratic program (QP) is solved to obtain the desired joint velocities \dot{q}_d . The resulting velocities are then Euler-integrated to generate joint position setpoints for a low-level joint-stiffness PD controller operating at 1kHz.

Our teleoperation framework supports multiple input modalities to accommodate different user preferences and operational contexts. We support the 3DConnexion SpaceMouse for precise desktop control, iPhone ARKit for mobile spatial tracking, direct manual jogging with gravity compensation for intuitive physical interaction, and immersive control through the Apple Vision Pro [35]. The system leverages the standardized LeRobot dataset format [17] paired with Hugging Face Hub’s cloud infrastructure, providing robust data storage, version control, and visualization capabilities that facilitate collaborative development and reproducible research.

APPENDIX B POLICY LEARNING IMPLEMENTATION DETAILS

In this appendix, we provide the implementation details and set of hyperparameters used for training diffusion policies, as shown in Table IV.

Behavioral cloning (BC) is a direct, supervised learning approach in which a policy is trained on expert demonstrations to learn a mapping from observations to actions. More precisely, BC optimizes the parameters of the policy by maximizing the likelihood of expert demonstrations:

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{o_t \sim \mathcal{D}} [\log \pi_\theta(a_t|o_t)]$$

\mathcal{D} denotes a dataset of expert demonstrations. Sequences of predicted actions (a_t, \dots, a_{t+T_p}) and sequences of observed states (o_{t-T_0}, \dots, o_t) at a given timestep t are denoted by a_t and o_t , respectively, for brevity, with observation horizon T_o and prediction horizon T_p .

Identical parameters were used for diffusion policies trained for both platforms. Both policies take RGB images from a wrist-mounted camera in addition to the proprioceptive state as input to generate predicted action sequences over a horizon

of T_p timesteps. We use absolute end-effector position control and a discrete gripper action $g_{\text{action}} \in \{\text{open}, \text{close}, \text{no-op}\}$ as our action space where rotations are represented using the continuous 6D rotation representation proposed by [44]. With the exception of the gripper state which is represented continuously, the proprioceptive state of the robot is represented as absolute end-effector pose with the same SE(3) representation used for action targets. A circular ring buffer is updated with proprioceptive state at 200Hz and used to synchronize RGB frames from the wrist camera with the robot’s proprioceptive state using hardware timestamps.

TABLE IV: Hyperparameters used for all diffusion policies.

Parameter	Value
<i>Architecture</i>	
Vision encoder	ResNet18 [26]
Input image size (N, H, W, C)	(1, 240, 320, 3)
Kernel size	5
U-Net down dims	(256, 512, 1024)
N group norm groups	8
Diffusion step embedding dim	128
Action dim	10
Observation dim	10
Prediction horizon, T_p	16
Observation horizon, T_o	2
Action horizon, T_a	8
<i>Diffusion Process</i>	
DDIM training steps	100
DDIM inference steps	16
β_{start}	1e-4
β_{end}	0.02
<i>Training</i>	
Batch size	128
Learning rate	1e-4
Learning rate scheduler	Cosine
Warmup steps	500
Optimizer	Adam [30]
β_1, β_2	0.95, 0.999
Weight decay	1e-6
Training iterations	500K
Gradient clipping	10.0
Loss	MSE
<i>Normalization</i>	
Proprioceptive State	Min/Max
Action	Min/Max
Wrist RGB	[0, 1], Z-Score
<i>Image Augmentations</i>	
Random crop (H, W)	(216, 288)
Brightness jitter	(0.9, 1.1)
Contrast jitter	(0.9, 1.1)
Saturation jitter	(0.9, 1.1)
Sharpness adjustment factor	1.5
Sharpness adjustment probability	0.5
Noise	$\mathcal{N}(0, 0.1)$
Max no. augmentations	3