DexWrist: A Robotic Wrist for Constrained and Dynamic Manipulation

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

Improbable AI Lab
Massachusetts Institute of Technology, Cambridge, MA 02139

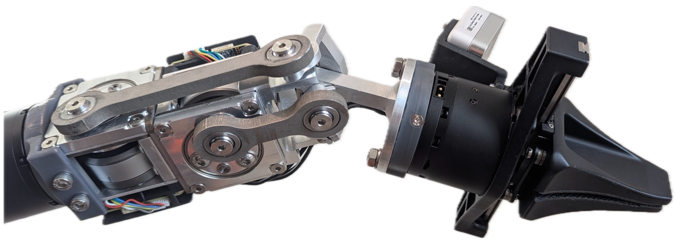

Fig. 1: We present *DexWrist*, a robotic wrist that allows for constrained and dynamic manipulation and speeds up teleoperated data collection and makes teleoperation more intuitive.

***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 simulateable 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 progress has been made in robotic manipulation with large-scale data from bimanual systems such as ALOHA [5] and the Franka Panda Dual Arm Setup [4]. However, most robotic manipulation still operates under simplified conditions, unsuitable for real-world tasks in cluttered spaces, such as retrieving items from crowded refrigerators or plugging cables into obstructed outlets. Conventional robotic arms struggle to perform tasks in these environments, and teleoperation is often tedious due to the need to avoid collisions and maneuver through tight spaces.

Humans navigate such environments with ease largely due to joint co-location of degrees of freedom and high dexterity in the wrist and hand. This paper focuses on addressing two major challenges in robotic manipulation: limited motion in constrained spaces, and slow data collection. Current robotic wrists, such as those on UR and Franka arms, rely on serial joint arrangements designed for open workspaces. These designs limit their reachability and reduce the number of viable motion solutions [10, 16, 35]. Moreover, common differential joints are difficult to model [24], and high-torque gear systems reduce compliance, leading to poor adaptability to external forces [10].

We introduce DexWrist, a parallel-actuated, compliant robotic wrist that addresses these limitations. It supports activities of daily living (ADLs), meeting strict requirements for torque, compliance, speed, bandwidth, precision, and range of motion. DexWrist is designed to handle high loads and environmental disturbances while enabling rapid, intuitive teleoperation. Combined with a low-cost Agile-X arm, DexWrist enables efficient data collection in cluttered settings and improves task execution in constrained environments. Our results show significant improvement in manipulation capabilities and data collection efficiency, making DexWrist a valuable tool for advancing robot learning in real-world scenarios.

II. Prior Works

A. Serial Wrists.

Most commercially available robotic arms have an integrated serial wrist, such as the UR5e [7], Franka Panda [4], and AgileX PiPER [6]. There is also a variety of standalone serial wrists like the Montagnani switchable stiffness wrist [25] and the Chirikjian spherical stepper motor [14]. 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 [30], Carpal Robotic wrist [8], and Damerla prosthetic wrist [15] 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 [18] and Negrello soft wrists [26] 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.

TABLE I: Comparison of Desired and Achieved Functional Requirements

Functional Requirement	Desired	Range	Ours	Pass
Rated Active Torque (Nm)	3	≥ 3	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	10–53.3	96.6 ± 9.4	✓
Bandwidth (Hz) @ 3.75Nm	20	10–20	10.15 ± 1.34	✓
Angular Precision (°)	3.5	0–3.5	1.65	✓
F/E ROM (°)	80	–40–40	–40–40	✓
R/U ROM (°)	40	–10–30	–40–40	✓
Width (mm)	61.4	51.5–61.4	64	∼
Height (mm)	61.4	51.5–61.4	66.5	✗
Length (mm)	174.5	±5	178.2	✓
Weight (kg)	1	0–1	0.97	✓

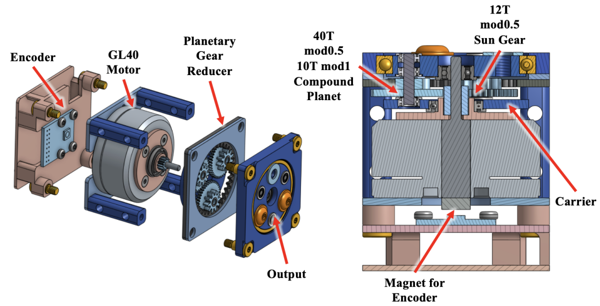

Fig. 2: QDD Stepped Planetary Compound Gearbox with 13:1 transmission.

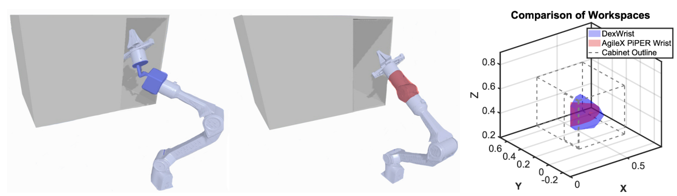

Fig. 3: **Left:** *DexWrist* wrist simulation. **Middle:** Serial wrist simulation. **Right:** Workspace comparison plot illustrating reachability of each arm configuration within the deep angled cabinet. The *DexWrist* improved constrained workspace reachability by 88% when compared to the AgileX wrist.

III. Wrist Design

A 2-DOF robotic wrist was designed according to the functional requirements outlined in D. 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 Gearbox

To meet the 3 Nm, 50 rpm, and back-driveability targets within a 61 mm cube, we pair a high-density BLDC (CubeMars GL40, 1.50 W cm^{-3}) with a custom 13:1 stepped-planet compound gearbox (Fig. 2). Commercial units (e.g. Maxon, 31.7 mm long) were too large, so we designed a quasi-direct-drive stage whose reflected inertia is $< 10^{-5}$ of the load. Gear tooth bending stress is checked with the Lewis equation, using K_d factors and a 1045-steel rim to keep FOS ≥ 3. An AS5047P encoder and Moteus-n1 controller close the 1 kHz torque loop over CAN.

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

The decoupled 2-DOF PKM (Fig. 4) mirrors human F/E and R/U axes while localising inertia at the hand. Simulated reach inside a deep, angled cabinet (PyBullet) shows an 88 % increase in collision-free targets over the AgileX serial wrist (Fig. 3). Four load cases—5 kg in $\pm X, Y, Z$ and 3 Nm active torque—set a 6 mm shaft and 9 mm bushing minimum; 17-4 PH stainless links give FOS = 3.3 without exceeding the 64 mm envelope.

IV. 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 $\mathcal{T}_w^{ee} \in \mathrm{SE}(3)$ are obtained from a teleoperation controller. Using this pose target, 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. More details on the policy learning pipeline can be found in Appendix B.

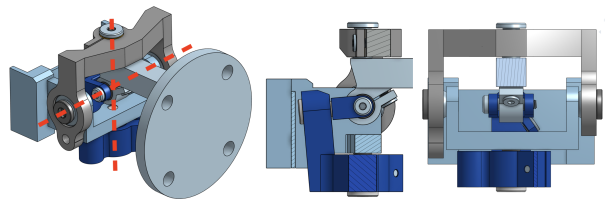

Fig. 4: **Left:** 2-(R, RR) PKM isometric view with red dotted lines depicting DOF rotational axes and pivot point. **Middle:** Side view highlighting kinematic chain in dark blue (RR). **Right:** Front view highlighting kinematic chain in gray (R).

Fig. 5: The three tasks shown to evaluate the performance of the wrist as compared to the stock AgileX arm.

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 are shown in Figure 5.

1) *Picking from a Cluttered Refrigerator*: Pickup 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.

V. 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 $B(Hz) = 0.35/t_r(s)$ 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.

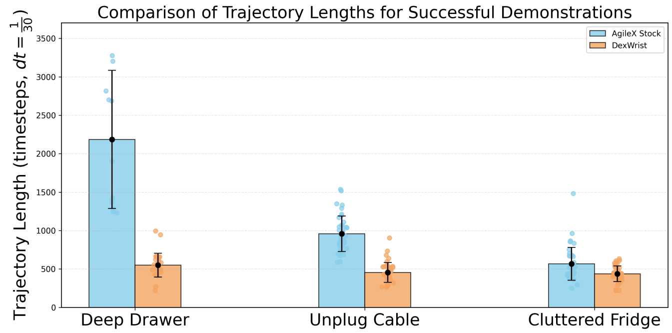

Fig. 6: Teleoperation demonstrations recorded from successful trajectories. Resets were performed in the event of a severe robot collision, surrounding objects being knocked over, or a failed grasp. For each configuration, $N \geq 40$.

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 spaces outlined in IV-C. The results in Table III show that trajectories collected in constrained spaces using DexWrist required significantly less timesteps and environment resets per successful trajectory compared to the stock PiPER Arm with the default wrist.

C. Behavioral Cloning

1) Method: To evaluate the impact of wrist design on manipulation performance, we trained diffusion policies [12] 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 [31] to perform the task. Refer to Appendix A for the implementation details.

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. The target requires challenging end-effector positioning: the gripper must reach the back of the refrigerator with fingers parallel to the rear wall (opening axis nearly orthogonal to the back wall). This awkward orientation complicates collision avoidance with nearby objects. Limited vertical clearance between shelves, dividing rack, and bottom tray creates minimal tolerance for arm movement. Initial target occlusion further increases task difficulty.

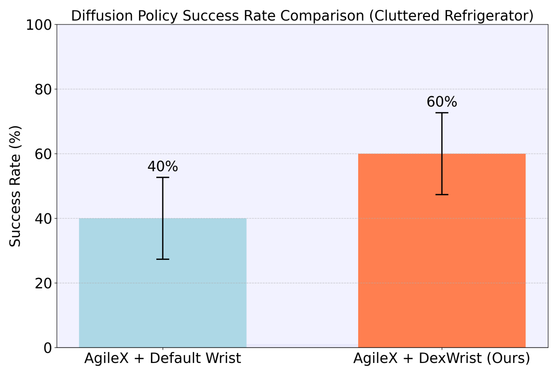

Fig. 7: For each system, we report the highest success rate among all evaluated checkpoints. Error bars show the binomial standard error based on 15 trials of the best checkpoint only. Evaluation details in appendix.

The checkpoint with the best success rate trained for the AgileX + DexWrist combination exhibited a 50% relative improvement in success rate compared to the policy trained for the AgileX + default wrist system. On successful trials, the DexWrist was observed to be **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 ± 7.9	55.2	134.2
AgileX + **DexWrist (Ours)**	$\mathbf{28.1 \pm 2.2}$	**20.5**	**49.0**

VI. 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 7, 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 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 demonstrations, 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.

Directions for future work includes investigating reinforcement learning approaches for performing dynamic tasks and transferring policies trained on demonstrations from the UMI [13] to systematically investigate whether the Dexwrist more effectively closes the human-robot wrist embodiment gap compared to traditional wrists.

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] Franka Emika Panda robot - RoboDK, . URL https://robodk.com/robot/Franka/Emika-Panda.

[5] Mobile ALOHA: Learning Bimanual Mobile Manipulation with Low-Cost Whole-Body Teleoperation, . URL https://mobile-aloha.github.io/.

[6] PiPER, . URL https://global.agilex.ai/products/piper.

[7] UR5e Lightweight, versatile cobot, . URL https://www.universal-robots.com/products/ur5e/.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] 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.

[15] 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 Awtar. 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.

[16] 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.

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[21] 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.

[22] 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/.

[23] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[24] 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.

[25] 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 (ROBIO)*, pages 310–315, December 2013. doi: 10.1109/ROBIO.2013.6739477. URL https://ieeexplore.ieee.org/document/6739477/.

[26] 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.

[27] 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.

[28] Younghyo Park and Pulkit Agrawal. Using apple vision pro to train and control robots, 2024. URL https://github.com/Improbable-AI/VisionProTeleop.

[29] 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.

[30] 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.

[31] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *CoRR*, abs/2010.02502, 2020. URL https://arxiv.org/abs/2010.02502.

[32] 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/.

[33] 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.

[34] 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/.

[35] 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.

[36] 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
POLICY LEARNING DETAILS AND HYPERPARAMETERS

In this appendix, we provide the background, implementation details and complete set of hyperparameters used to training both diffusion policies. Detailed hyperparamters can be found 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:

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

$\mathcal{D}$ denotes a dataset of expert demonstrations. Sequences of predicted actions $(a_t, \ldots, a_{t+T_p})$ and sequences of observed states $(o_{t-T_0}, \ldots, o_t)$ 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 [36]. 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.

1) Evaluation Details: 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 rollouts in the real environment (90 trials per system), randomly resetting the objects in the scene before each iteration. More details can be found in Appendix C

APPENDIX B
PIPELINE DETAILS

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 [28]. The system leverages the standardized LeRobot dataset format [11] paired with Hugging Face Hub's cloud infrastructure, providing robust data storage, version control, and visualization capabilities that facilitate collaborative development and reproducible research.

TABLE IV: Hyperparameters used for all diffusion policies.

Parameter	**Value**
Architecture	
Vision encoder	ResNet18 [19]
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
Diffusion Process	
DDIM training steps	100
DDIM inference steps	16
β_{start}	1e-4
β_{end}	0.02
Training	
Batch size	128
Learning rate	1e-4
Learning rate scheduler	Cosine
Warmup steps	500
Optimizer	Adam [23]
β_1, β_2	0.95, 0.999
Weight decay	1e-6
Training iterations	500K
Gradient clipping	10.0
Loss	MSE
Normalization	
Proprioceptive State	Min/Max
Action	Min/Max
Wrist RGB	[0, 1], Z-Score
Image Augmentations	
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

APPENDIX C
SCENE RANDOMIZATION: CLUTTERED REFRIGERATOR SETTING

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.

APPENDIX D
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.

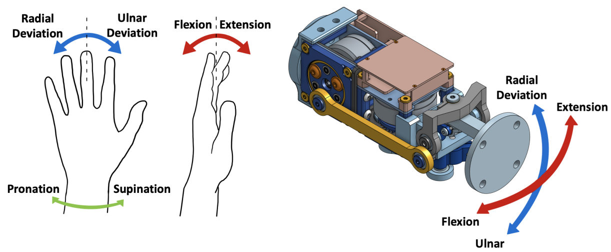

Fig. 8: **Left:** Human wrists have 3 degrees of freedom: flexion/extension (F/E), radial/ulnar (R/U) deviation, and pronation/supination (P/S). Anatomically, the F/E and R/U degrees of freedom are in parallel and are preceded by P/S in series. **Right:** *DexWrist* DOFs mirroring human wrist.

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 8) torques are 4.6-11.9 Nm and 4.7-10.8 Nm, respectively [34, 33]. 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 [27, 9], 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.

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 [32]. 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 [22, 17]. 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 [29]. 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: [21] and [20] 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 [7], can handle.