ORCA: An Open-Source, Reliable, Cost-Effective, Anthropomorphic Robotic Hand for Uninterrupted Dexterous Task Learning

Clemens C. Christoph[1†], Maximilian Eberlein[1†], Filippos Katsimalis[1†], Arturo Roberti[1†], Aristotelis Sympetheros[1†], Michel R. Vogt[1†], Davide Liconti[1], Chenyu Yang[1], Barnabas Gavin Cangan[1], Ronan J. Hinchet[1], Robert K. Katzschmann[1*]

Abstract— **General-purpose robots must possess human-like dexterity and agility to match the versatility of humans. A human-like form factor further enables leveraging the wealth of data available from human hand interactions. However, the bottleneck in dexterous manipulation lies not only in software but arguably even more in hardware. Robotic hands matching human capabilities are often prohibitively expensive, bulky, or require enterprise-level maintenance, making them inaccessible for broader research and applications. What if the research community could get started with reliable dexterous hands within a day? We present the ORCA hand, a 17-DoF tendon-driven robotic hand with fully integrated tactile sensors that can be assembled in under 8 hours and has a BOM of $2k. We showcase design features such as popping joints, auto-calibration, and tensioning systems that significantly reduce complexity while increasing reliability, accuracy, and robustness. We benchmark the ORCA hand across a variety of tasks, ranging from teleoperation and imitation learning to zero-shot sim-to-real reinforcement learning. Furthermore, we demonstrate that our hand is capable of withstanding over 10k cycles, or approximately 20 hours of continuous operation. CAD files, source code, and documentation will be made available at `orcahand.com`.**

I. Introduction

Reproducing the intricate dexterity of the human hand has long been a central challenge in robotics [1], [2]. While robotic grippers excel in industrial automation, their limited versatility makes them unsuitable for interacting with tools and objects designed for human hands [3], [4]. Consequently, extensive research has been devoted to developing anthropomorphic robotic hands and training them to solve complex manipulation tasks [5], [6]. However, compared to grippers, anthropomorphic hands require significantly more actuators, increasing complexity in their assembly and control. Additionally, good anthropomorphic hand hardware must be durable, repeatable, and versatile to be used in machine learning applications. Whether it is the sim2real gap in reinforcement learning (RL) or teleoperation accuracy in imitation learning (IL), bottlenecks in dexterous manipulation stem not only from software but also, and maybe more importantly, from hardware limitations [7], [8].

Existing tendon-driven robotic hands, such as the Shadow Hand [10], have demonstrated impressive capabilities in dexterous manipulation tasks [11]. However, these hands cost over $100,000 USD, require substantial maintenance [12] and are difficult to repair due to their proprietary and highly integrated designs. Other tendon-driven hands, such as the InMoov hand [13] and the DexHand [14], offer advantages in being open-source and low-cost. However, the InMoov hand is limited in dexterity, while the DexHand is challenging to assemble, and neither has demonstrated real-world applicability in autonomous manipulation tasks. Alternatives to tendon-driven hands are direct-driven hands, such as the Allegro Hand [15], priced at around $15,000,

[1]Soft Robotics Lab, IRIS, D-MAVT, ETH Zurich, Switzerland
† Equal contribution.
*Corresponding author: `rkk@ethz.ch`

Fig. 1: (A) The ORCA hand closely mimics its human counterpart with the same form factor, a bony structure, and silicone-cast skin. The ORCA hand is 3D printed but incorporates joints designed to pop before breaking, making it resistant to overload breaks while retaining the advantages of bearing pinhole joints, such as stability and simple kinematics. (A1) Just before the joint pops. (A2) Applying pressure pops it into place and keeps it secure. (A3) depicts our spool system, which enables manual retention without unscrewing the spools or tendons. (B) We show that our hand can be deployed in real-world settings by running our self-resetting imitation learning policy for over 7 hours. (C) Our reliability test reveals our hand's robustness and the high repeatability of joint movements.

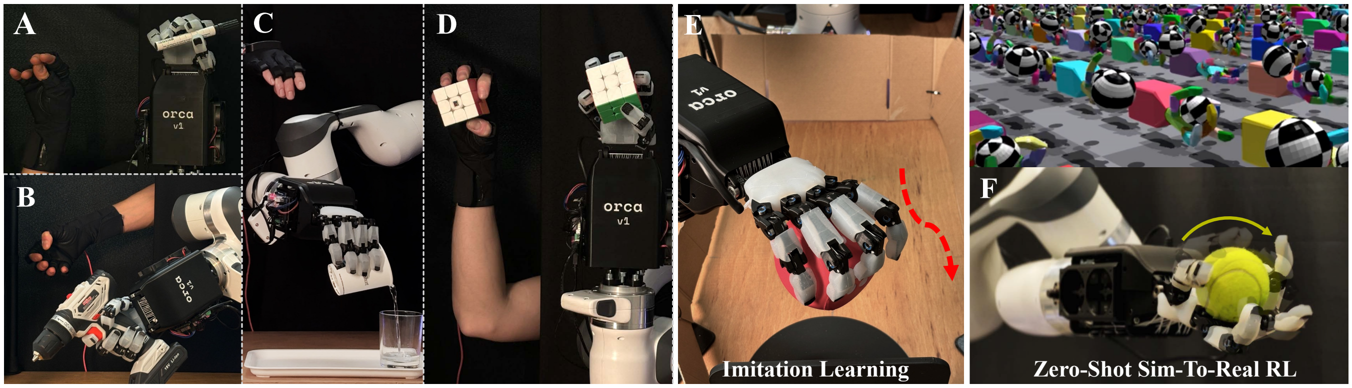

Fig. 2: Versatility of the ORCA hand: (A)-(D) Teleoperation with ROKOKO [9] gloves. (A) Holding a pen (B) Using a drill, showing high dexterity. (C) Liquid pouring. (D) Grasping a cube: This picture illustrates how closely the ORCA hand resembles a human hand. (E) IL with walls and a slider for self-resetting. (F) Policies in simulation, such as rolling a ball, can be deployed zero-shot to the real world due to the ORCA hand's low joint errors.

and the LEAP hand [16], where the latter is an open-source, cost-effective solution (about $2,000) that requires only three hours of assembly and provides unprecedented reliability. Nonetheless, all direct-driven hand designs share limitations such as bulkiness, restricted form factors, and an inability to match the softness and agility of a human hand. In fact, placing motors within the fingers creates inertia issues, which prevents the quick and dynamic movements that are essential for human-like motion.

In this paper, we present the ORCA hand: a tendon-driven, dexterous, and anthropomorphic robotic hand with fully integrated tactile sensors. The ORCA hand is designed for reliability, simplicity, and versatility across a wide range of tasks. Key contributions of our integrated system include:

- An open-source, 3D-printable design with a cost of less than $2000, which can be assembled by a single person without prior experience in under 8 hours.
- A joint design that pops before it breaks, enhancing the durability of 3D-printed components and streamlining the assembly process.
- Auto-calibration enabled by tendon routing through the center of rotation. This minimizes joint position errors and increases repeatability.
- Fully integrated tactile sensors and sensor wiring, which can be produced in-house, offering a compact and modular solution.

We demonstrate the hand's dexterity by teleoperating it to perform complex tasks that traditional robotic grippers cannot accomplish. Through a variety of reliability tests we demonstrate the ORCA hand's capability to offer exceptional reliability, durability and consistent performance during tens of hours of operation. To showcase the ORCA hand's dexterity and accuracy we implement fine motor control tasks like in-hand object orientation. Leveraging the anthropomorphic design we are also able to implement imitation learning tasks like the picking and placement of cubes and execute these autonomous tasks continuously for multiple hours without any human intervention on the hand hardware.

II. System Design

The ORCA hand design follows the requirements set above, namely dexterity and reliability at a low cost and low complexity. It comprises five fingers, including an opposable thumb and an actuated wrist (Fig. 1A), and is similar in size to an average human hand [17]. The fingers are mounted on a base that resembles the human palm containing the carpal and metacarpal bones. The palm is connected to the wrist mechanism which is mounted on the *tower*. The tower contains the motors and all auxiliary electronics, and is enclosed in a protective casing. The rest of this chapter describes the most important design features of the ORCA hand in detail.

A. Tendon Actuation for Agility

Each joint, except for the wrist joint, is actuated using two fishing lines (Nylon fibers braided into a 1 mm diameter rope) under tension, here referred to as tendons. One tendon is responsible for flexion (flexor), and the other for extension (extensor). This decision was made based on the fact that a smaller form factor and lower finger inertia more closely mimic the nimbleness and dexterity of human hands compared to direct-driven hands. Moreover, tendon actuation makes the ORCA hand independent of the choice of actuator, enabling actuation technologies other than electric motors to be used in the future, such as artificial muscles.

While tendon actuation offers many advantages, it also presents challenges such as friction buildup, wear, and slack over time, which can affect movement precision and longevity. We mitigate those challenges as follows:

- We avoid direct contact of tendons with PLA by deflecting the tendons around smooth metal pins and rods, as depicted in Fig. 3C.
- We use Teflon tubes for non-linear routing, e.g. from the bottom of the thumb to the wrist.
- Finally, tendons can be manually re-tensioned using a ratchet spool mechanism mounted on the motors, as shown in Fig. 1A3. The ratchet is attached to the top spool, allowing rotation in one direction while locking movement in the other. This design makes the ORCA hand user-friendly, as re-tensioning can be done

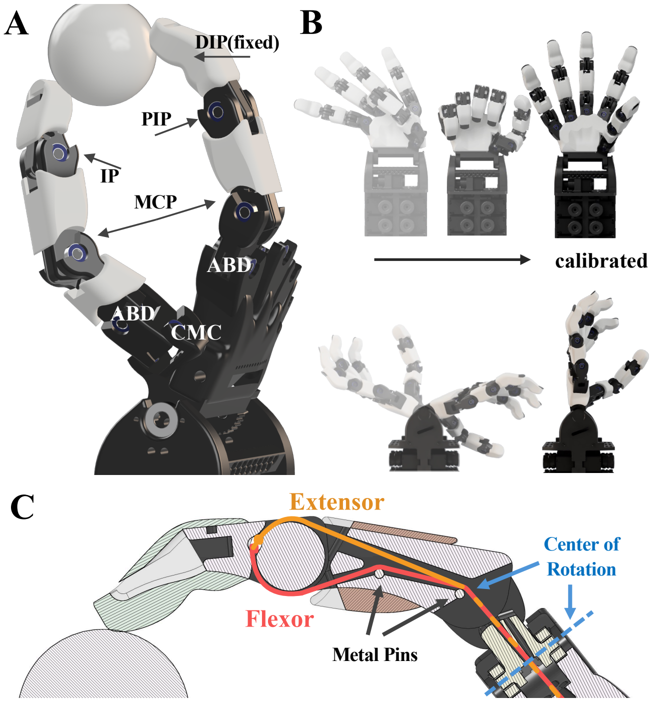

Fig. 3: (A) Naming convention of the joints. The thumb includes an additional degree of freedom. (B) Auto-calibration: The three-step process moves all joints to their respective limits and determines a mapping between motor and joint angles without any external sensors. (C) Routing of the DIP joint. The tendons are guided around metal pins, to reduce friction and eliminate wear over time. Moreover, they always are guided through the center of rotation for straightforward control.

in seconds without the need to unscrew the spool or tendon. The spool system quickly eliminates any slack that accumulates. Additionally, the tendons can be easily loosened, allowing joints to be popped out for quick replacement of broken parts.

B. Poppable Pin Joints

Rolling contact joints [18], [14] have become a popular alternative to pinhole joints for 3D-printed hands due to their ability to dislocate instead of breaking. However, these mechanisms require ligaments, which can loosen over time and increase complexity. We introduce a pin joint design that allows them to "pop" out of place and dislocate instead of breaking when excessive radial and axial loads are applied (Fig. 1A1). This is achieved by placing the bearings in circular arc-shaped grooves, which hold them tightly under normal operating conditions but allow them to dislocate in the event of a collision. This mechanism combines the advantages of pinhole joints, such as axial stability and straightforward kinematics, with the robustness of ligament-based rolling contact joints, while being extremely quick and easy to assemble.

C. Finger and Palm Design

Fingers two to five (index to pinky) have three actuated joints that mimic those of the human finger: the Proximal Interphalangeal (PIP), the Metacarpophalangeal (MCP) and the Abduction (ABD) joint (Fig. 3A). The ORCA PIP joint corresponds to the human PIP joint, while the MCP and ABD joints together correspond to the human MCP joint which can perform both flexion/extension and abduction/adduction, which is necessary in various dexterous manipulation tasks [19]. The range of motion (RoM) for each joint is shown in Table I, based on the human anatomy [17].

Joint Name	Fingers 2 to 5		Thumb	
	Flexion	Extension	Flexion	Extension
IP	-	-	100°	20°
PIP	130°	20°	-	-
MCP	110°	20°	115°	20°
ABD	30°	30°	45°	45°
CMC	-	-	48°	53°

TABLE I: Range of motion of joints of ORCA hand.

Omitting the DIP joint on fingers two to five was an intentional choice with the benefits of greatly simplifying assembly and reliability. The DIP joint is not directly actuated in the human hand, and as such it is of less importance compared to the other hand joints in dexterous manipulation tasks. A positive side-effect of the removal of the DIP joint is the significant increase in the space available for tactile sensing integration on the fingertips.

The thumb differs from other fingers and has four instead of three DoFs, i.e. four joints, the Interphalangeal (IP), MCP, ABD and Carpometacarpal (CMC) joint. Additionally, it is placed on the palm in an opposable manner to other fingers with a supination of 15° [17].

D. Wrist Design

One of the biggest limitations of hands missing the wrist joint is that the palm cannot be oriented parallel to a surface, such as a table, because the robotic arm would block the way. This significantly reduces grasping performance. Motivated by this, we added one rotational DoF around the transverse (radioulnar) axis of the hand. The human wrist can flex and extend about 80° [17], while the ORCA wrist mechanism is capable of achieving 60° of flexion/extension. Instead of using tendon actuation at the wrist, we opted for belt drive to account for the increased loads the wrist experiences. We use a standard GT2 timing belt with fiberglass reinforcement, which exhibits negligible slack buildup over time. We decided not to add a second DoF at the wrist to represent the radial/ulnar deviation of the human wrist, in order to avoid unnecessary complexity and increased cost, since the human wrist is significantly limited in this type of motion compared to flexion/extension [17].

E. Integrated tactile sensing

Many works have also focused on enhancing the manipulation abilities of robotic hands by integrating tactile sensors, particularly on the fingertips. These sensors include force sensors [20], [21], piezoresistive pressure sensors [22], capacitive pressure sensors [23], and Hall effect-based sensors [24]. In the ORCA hand, we utilize Force Sensing Resistors (FSR) (RP-C7.6-ST Thin-Film Pressure Sensor) mounted onto a solid FDM-printed PLA backplate, and covered by silicone-molded skin, to provide us with binary tactile feedback on all five fingertips. While FSR sensors are in theory capable of providing feedback on the magnitude of applied force, a binary interpretation was chosen due to the compliant skin, combined with its irregular surface, dampening external forces to a different degree depending on the location on the fingertip where the force is applied. As such, the force magnitude cannot be estimated without additional information about the location of indentation, which the FSR sensors by themselves are incapable of providing. All sensors are connected to the external electronics through thin copper wires (∅ 0.2mm), which are routed through the internal finger and palm structure through PTFE tubing, to protect them from the environment and external forces, while furthermore providing a clean visual appearance. To read the output of the FSR sensors, each sensor forms a voltage-divider with a $10k\Omega$ resistor, the three nodes of which are connected to a 5V input, an analog input, and ground respectively, all provided by an Arduino Nano Every.

F. Self-Calibration for Accurate Control

In machine learning settings, such as data collection for imitation learning, consistency in recordings is critical. If there is high variance in the gap between human teleoperation and the robotic hand across different runs, it becomes more difficult for policies to accurately imitate the human [25]. Moreover, this consistency is not only crucial during data collection but also during policy deployment. Any offset between the training data and the deployed system can result in poor policy performance, as the learned model may fail to generalize effectively to scenarios with mismatched conditions. This is especially important when attempting to reduce the sim-to-real gap. While it is possible to ensure reliable control consistency using proprioception sensors, as demonstrated in [10], this adds significant cost and complexity. We therefore propose an alternative, more cost-effective approach:

Let θ_i be the true angle, $\tilde{\theta}_i$ the commanded angle of joint $i, \forall i \in \{1, \ldots, 17\}$, and ϕ_i the motor's angular displacement for the same joint. The ideal goal is to always have $\theta_i = \tilde{\theta}_i$. However, in practice, especially with tendon-driven hands, there is a discrepancy because motor-to-joint mapping relies on a model where:

$$\begin{aligned}\theta_i &= \tilde{\theta}_i + \epsilon(r_j, r_m, s, d, m) \\ &= f(\phi_1, \ldots, \phi_{17}) + \epsilon(r_j, r_m, s, d, m)\end{aligned}$$

where $f(\phi_1, \ldots, \phi_{17})$ is generally non-linear, and ϵ represents model errors including tendon radius of joint and motor (r_j, r_m), built-up slack s, servo drift d and measurement errors m in case of manual calibration which can introduce substantial offsets.

Referring to the routing of the ORCA hand, as depicted in Fig. 3, each tendon passes through or near the center of rotation (CoR). This design ensures that joint positions are approximately decoupled and can be actuated linearly and independently. Due to the system's linearity and our focus on controlling the motor angles ϕ_i, we can express the desired angles as $\tilde{\theta}_i = R'_i \cdot \phi_i$, with $R'_i \in \mathbb{R}^+$. Consequently, the relationship between the actual angles θ_i and the motor angles yields

$$\theta_i = R'_i \cdot \phi_i + \epsilon(r_M, r_J, s, l, d).$$

Since we can directly control the motor angles, we define ϕ_i in terms of $\tilde{\theta}_i$ as $\phi_i = \frac{1}{R'_i} \cdot \tilde{\theta}_i = R_i \cdot \tilde{\theta}_i$. To determine R_i, we use a procedure referred to as **auto-calibration**, which estimates R_i and minimizes the error terms. The process is visually depicted in Fig. 3B and works as follows:

1. Let $|RoM_i|$ denote the absolute range of motion of joint i. Using CAD models, the angular distance between the physical stops of each joint is precisely determined, providing $|RoM_i|$ with minimal uncertainty. 2. During calibration, each joint i is moved to its extreme mechanical limits (fully flexed $\phi_{i,\text{max}}$, fully extended $\phi_{i,\text{min}}$). 3. The total servo rotation over the joint's range of motion is computed as:

$$\Delta\phi_i = \phi_{i,\text{max}} - \phi_{i,\text{min}}.$$

4. The joint-to-motor transmission ratio is then calculated using linear interpolation:

$$R_i = \frac{\Delta\phi_i}{|RoM_i|}.$$

and thus we get

$$\theta_i = \frac{1}{R_i} \cdot \phi_i + \epsilon(|RoM_i|, s, d) = R'_i \cdot \phi_i + \epsilon(|RoM_i|, s, d)$$

where $\epsilon(|RoM_i|, s, d) \ll \epsilon(r_M, r_J, s, d, m)$.

III. Hardware Performance Tests

A. Reliability and Robustness

To evaluate the ORCA hand's reliability and robustness in long-duration tasks, we conduct an experiment in which we actuate the hand's joints continuously for 2.5 hours (Fig. 1C). We attach a plush animal to the palm of the hand and have it grasp it with all fingers every four seconds. This setup, using a compliant object, is similar to the repeatability test performed by [16] and allows us to assess behavior under increased stress over a greater range of motion. Moreover, to test the wrist joint's durability, we flex and extend the wrist to 40 degrees at one-fourth of the finger frequency, so every 16 seconds. The hand reliably performs the same grasping movement for all 2,250 grasping cycles without breaking, motor shutdown, or excessive tendon slack buildup. In Fig. 4C, we display the maximum current in the motors actuating the middle finger's MCP and PIP joints for each grasping cycle, as well as the maximum current of the wrist joint motor for each wrist extension-flexion cycle. The current used by each motor provides an estimate of how much external torque and friction must be overcome for a given joint angle change. The fact that the maximum current for each cycle remains roughly the same over the 2.5 hours of uninterrupted joint movement demonstrates the hand's robustness, the high repeatability of joint movements, and its capacity for long-duration operation. The fans on the side of the hand's tower prevent the motors from overheating (Fig. 4C), allowing the hand to function almost indefinitely.

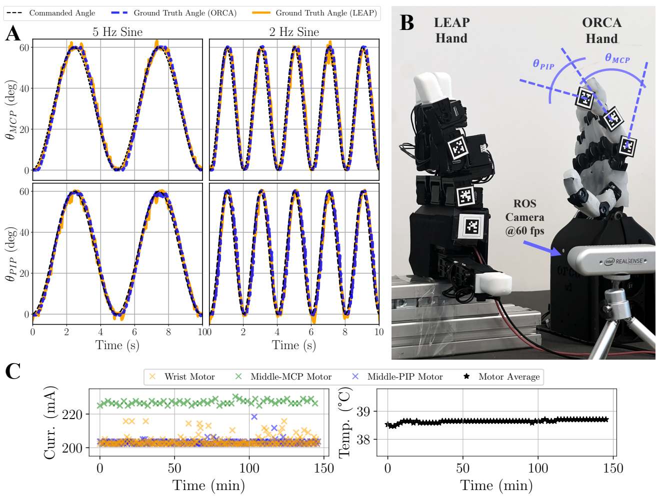

Fig. 4: (A) Comparison of joint responses to sine wave signals at 2 Hz and 5 Hz for the ORCA and LEAP hands: Despite the tendon-driven design, the ORCA hand achieves similar accuracy and latency as the LEAP hand. (B) Experimental setup to benchmark accuracy and latency of ORCA and LEAP hand actuation responses: AprilTags are recorded with a camera and used to infer the ground truth joint angles, which are then synchronized with the commanded angles. (C) Reliability test: Uninterrupted and consistent object grasping (2200+ cycles) and wrist motion (550+ cycles) for 2.5 hours without any breakdown, motor overheating, or performance degradation.

B. Accuracy and Latency

Accurate control of joint motion is crucial for reliable performance across any dexterous manipulation task. Additionally, excessive latency between action commands and their execution in the real world can severely hinder the implementation of closed-loop behaviors. Incorporating latency into the hand's model for training in simulation can be helpful to bridge the sim2real gap. To further demonstrate the reliability of the ORCA hand, we, therefore, benchmark the accuracy and latency with which the hand's joints can follow diverse action commands. We propose the following experimental setup to benchmark the accuracy and latency of robotic hands:

By attaching distinct AprilTags along one finger and measuring their relative orientations to each other, each joint angle on the finger can be calculated with relatively low error ($\sigma = 0.08°$ for our experimental setup). We place three tag36h11 AprilTags across the MCP and PIP joints of the ORCA hand's index finger and align them in a single plane (Fig. 4B). We then position a RealSense D435i camera in front of the AprilTags, using a custom script to ensure that the camera faces the tags at an angle of $90° \pm 5°$ for more precise angle measurements. We record image frames at 60 fps and save them to a ROS2 bag file. Simultaneously, we publish commanded angles that actuate the index finger's MCP and PIP joints in a sine wave pattern and log them into the same ROS2 bag file. This allows us to synchronize the commanded angles with the real angles obtained from the image frames and evaluate the system's latency and accuracy offline.

We compare the accuracy and latency of the ORCA hand with the LEAP [16] hand to benchmark our system's tendon-driven dynamics against the dynamics of a direct-driven robot hand (Fig. 4A). We actuate the MCP and PIP joints of both hands' index fingers with 2 Hz and 5 Hz sine wave patterns. Before each test, we leverage the auto-calibration mechanism of the ORCA hand to account for any changes in tendon length or slack that might have accumulated beforehand. We demonstrate that, through auto-calibration, the finger joints of the ORCA hand accurately follow the commanded sine input. In fact, we achieve similar accuracy to the LEAP hand while being far less bulky thanks to our tendon-driven actuation design. Additionally, we observe that the ORCA hand's joints rotate more smoothly than those of the LEAP hand, which exhibit more jerks, presumably due to the cables and possibly increased inertia interfering with the LEAP hand's finger joint motion. Both hands exhibit average latencies of less than 0.2 seconds, most of which comes from the software process. However, slack in the ORCA hand may introduce additional latency, which is why retensioning the spools periodically is important for robust performance.

One limitation of our benchmarking experiment is that we cannot evaluate the accuracy and latency of the hands for faster movements, as the AprilTags can not be tracked at sine wave frequencies below 2 Hz or for step signals due to motion blur at 60 fps. Recording the finger joints with a ROS2-compatible camera capable of 240 fps could enable more thorough system identification in the future.

C. Reinforcement Learning

Reinforcement learning is commonly used to learn dexterous tasks that are challenging to demonstrate or require fine motor control, such as in-hand object reorientation [26]. A key challenge in applying RL to dexterous manipulation tasks is that policies learned in simulation often perform poorly when transferred to the real robotic hand. We use the IsaacGymEnvs wrapper from [18] to train 4096 ORCA hand models in parallel with an advantage actor-critic architecture to learn in-hand ball reorientation. We demonstrate that after 1 hour of training with domain randomization, we can deploy a robust policy on the physical ORCA hand (Fig. 6B), that can successfully reorient a tennis ball along a given rotation axis.

D. Imitation Learning

Imitation learning has become another predominant approach in the manipulation community, as it enables learning tasks from a set of demonstrations without requiring task-specific rewards or simulation environments. Various architectures have been proposed to extract meaningful representations of observations and map them to the correct actions [27]–[29]. However, the application of imitation learning to dexterous platforms presents additional challenges mainly due to their higher-dimensional action spaces [30], [31].

To demonstrate autonomous task execution with the Orca Hand, we employed the state-of-the-art architecture from [27]. Our setup consists of the hand mounted on a Franka Panda robotic arm, equipped with two external cameras and one wrist-mounted camera.

Demonstrations were collected using motion capture gloves [9], which provided absolute wrist tracking and finger pose estimation. We retargeted these demonstrations into the robot's state space using an energy-based minimization objective, similar to [32]. The wrist pose was used to control the robotic arm in Cartesian end-effector space. This teleoperation method allowed us to showcase the versatility of the hand across a wide range of tasks (Fig. 2) and facilitated the rapid and intuitive collection of demonstrations, even from non-trained operators.

The policy takes as input three camera images along with proprioceptive data from both the robotic arm's end-effector and the hand. The output consists of an action chunk predicting future actions.

Over approximately 2h30m, 214 video samples were collected for training the policy on an *NVIDIA GeForce RTX 4090* GPU for 500 epochs, which took approximately 4h.

For the proposed task, we conducted ablation studies on image preprocessing. Specifically, we compared three policy variations: (1) a baseline policy trained on raw RGB inputs, (2) a policy incorporating segmentation of the cube's color, and (3) a hybrid approach trained on both datasets.

For a comprehensive description of the network and color masking parameters, please refer to the appendix.

E. Tactile Sensing

To determine the AT (absolute threshold) of the tactile sensors, a controlled orthogonal force was applied to the front surface of a fingertip using a cylindrical indenter with a diameter of $2cm$ and a flat contact surface. The applied force was varied by placing calibrated weights on top of the indenter. Registered touch was classified as any output reading above $0.01V$ on the respective analog input on the Arduino Nano Every. Despite the used FSR (Force Sensing Resistor) sensors being rated for a minimum trigger force of $0.29N$, the fingertip was capable of registering forces as low as $0.05N$ with perfect accuracy over 10 cycles. It is unknown if this is due to a lower than rated minimum trigger force of the commercial FSR sensors, or if the silicone skin had an unintentional pre-loading effect on the sensor. However, wear and tear in the silicone skin is able to drastically affect the AT of the sensors. While the sensor mounted on the ring finger of the hand showed no degradation after thousands of grasp cycles as part of the experiments outlined in section III-A, the AT for the sensor mounted on the pinkie finger increased to $6.38N$ due to degradation of the silicone skin creating an air gap between the silicone skin and the mounted FSR sensor. Furthermore, after around 4500 to 7000 grasp cycles, the thin copper wires connecting the sensors to the external electronics snapped on the thumb, index, and middle finger. The snapping points occurred at different heights but were all concentrated in the area of the MCP and ABD joints.

IV. Additional Results and Discussion

In this section, we show the results of both human teleoperation and imitation learning with the ORCA hand and discuss additional hardware-related findings.

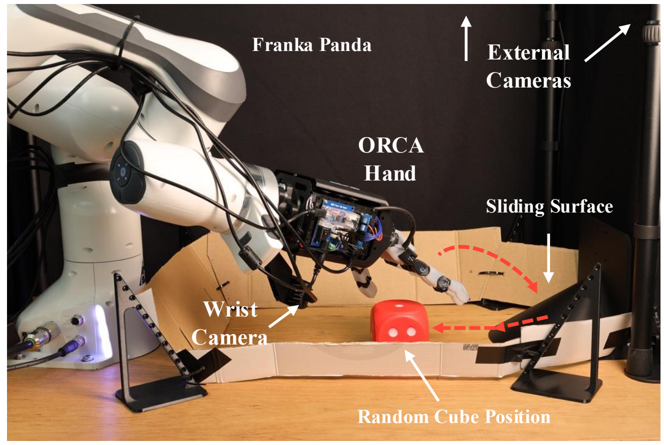

Fig. 5: Experimental setup for repeated pick & place: The cardboard serves as a fence, preventing the cube from rolling out of the testing area and enabling uninterrupted, long-duration policy deployment.

A. Teleoperation - Dexterous Manipulation

The dexterity and stability of the ORCA hand was successfully evaluated by picking up and interacting with a variety of objects:

- Stacking of 3 small and large cubes (same as for IL). The cubes were first placed singularly on the table.
- Grabbing a plush toy (about the size of 3 large cubes) lying on the table.
- Grabbing a tennis ball lying on the table
- Twisting open the cap of a *Nutella* jar ($\varnothing 8\,cm$). The jar itself is screwed on.
- Spinning a fidget spinner for $2\,s$. The fidget spinner is placed on a finger by hand, but the grasping and spinning is done purely by teleoperation.
- Picking up a pen and writing "Hello" on a fastened piece of paper (font size $\sim$200).
- Picking up a piece of paper lying on a fastened box.
- Picking up a cup and pouring its contents ($50\,ml$ water) into another cup.

B. Imitation Learning - Repeated Pick & Place

To highlight the reliability of both the hardware platform and the trained IL policy, we designed a continuous pick-and-place evaluation task. The robotic hand is required to pick up a cube (6 cm side length) from a table and place it onto a sliding surface, which then causes the cube to fall back onto a random location on the table (Fig. 5, 6A).

To evaluate different policies, we collected the ratio of failure positions to success positions (regarding picking up the cube) within a testing area for 60 iterations (10 per sub-area), as shown in Fig. 7. The most successful policy, using only masked images of the cube, is deployed for 7h 17min (approx. 2'000 grasping cycles) with no human intervention on the ORCA hand's hardware and minimal intervention in aiding in the pick-and-place task (Fig. 1B). The policy has consistent output over the entire duration of the test, showcasing the hand's durability in its lack of tendon slack or tendon rupture.

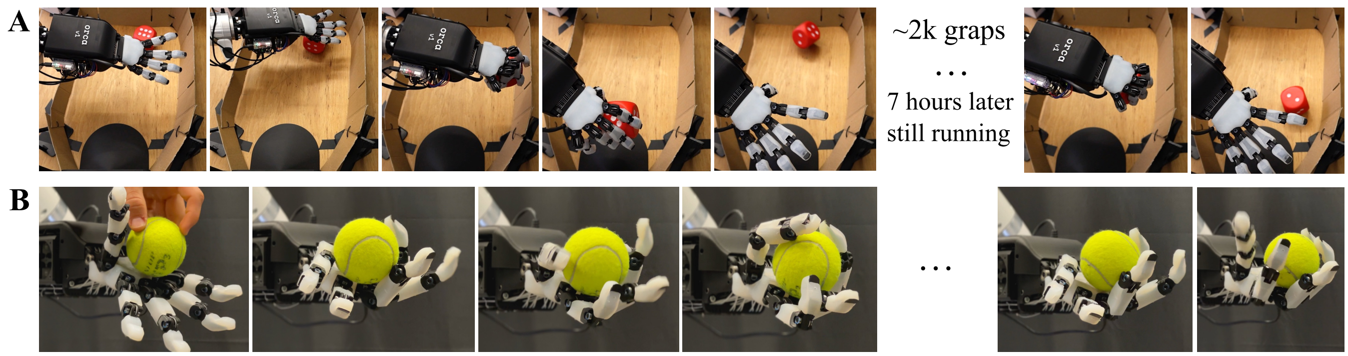

Fig. 6: (A) Prolonged imitation learning experiments over several hours. Walls and a sliding surface allow for self-resetting of the experiment. (B) Policies built in simulation are deployed zero-shot to the real world due to the ORCA hand's low joint errors.

For extended videos and time-lapses of the reliability test, please refer to the ORCA project website.

C. Tactile Sensing

While integrated tactile sensing provides a cost-effective and low barrier-of-entry approach to providing tactile feedback at the fingertips, the current design still shows a variety of limitations regarding reliability over thousands of grasp cycles. These limitations are namely the degradation of the silicone skin that is vital to reliably transmit the contact forces to the sensors (occurrence in 2 fingertips after approx. 2000 to 4000 grasp cycles) as well as the snapping of the thin copper wires used for signal transmission (occurrence in 3 fingertips after approx. 4500 to 7000 grasp cycles). Both of these limitations will be addressed in future work to allow for reliable tactile sensing integration into autonomous tasks.

D. Ease of Assembly

The ORCA hand is designed to be easily assembled and repaired. Given all the necessary mechanical parts, it takes one person less than 8 hours to assemble one entire ORCA hand. A detailed step-by-step instruction on how to assemble and repair the hand will be made available on the project website.

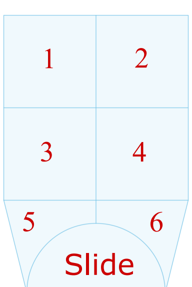

	RGB	Masked	Mixed
Area 1	50%	80%	80%
Area 2	10%	60%	20%
Area 3	80%	100%	100%
Area 4	60%	100%	80%
Area 5	60%	100%	80%
Area 6	40%	80%	60%
Total	50%	86.$\bar{6}$%	70%

Fig. 7: Testing area with respective policy success rates (% out of 10)

V. Conclusion

The ORCA hand offers an accessible solution for advancing robotic manipulation, and we hope to have laid the foundation for real-world tasks and cutting-edge manipulation research. A great deal of work has gone into making this human-like, compliant hand as robust and reliable as possible while remaining versatile, easy to control, and cost-effective.

Nevertheless, our study has limitations that suggest directions for future research. Prolonged use of the ORCA hand necessitates manual re-tensioning to sustain optimal performance. To address this, we aim to develop a re-tensioning mechanism that autonomously reduces tendon slack without the need of human intervention. This advancement could enable even longer operation times, potentially accelerating real-world RL applications for dexterous manipulation research. Future work of our hand may also include full integration of the sensors into the learning pipeline, as well as more advanced deep learning algorithms to learn complex interactions with the environment more effectively.

Acknowledgment

Thanks to Robert Jomar Malate, who led the Real World Robotics course at ETH Zurich and helped initialize this project. We also thank Carlo Sferrazza and the Berkeley Robot Learning Lab for providing us access to the LEAP Hand. Additionally, we thank Jeannette Bohg and the Interactive Perception and Robot Learning Lab for giving us access to the Shadow Hand for comparison, which, despite the company's technical support, we were not able to get to work.

We are grateful for the funding through a research collaboration with armasuisse. Filippos Katsimalis thanks the Bodossaki Foundation, SYN-ENOSIS, and the Onassis Foundation for their financial assistance. Aristotelis Sympetheros thanks the John S. Latsis Public Benefit Foundation and SYN-ENOSIS for supporting his postgraduate studies.

References

[1] A. Bicchi and V. Kumar, "Robotic grasping and contact: a review," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, vol. 1, 2000, pp. 348–353 vol.1.

[2] A. Okamura, N. Smaby, and M. Cutkosky, "An overview of dexterous manipulation," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, vol. 1, 2000, pp. 255–262 vol.1.

[3] C. Yu and P. Wang, "Dexterous manipulation for multi-fingered robotic hands with reinforcement learning: A review," *Frontiers in Neurorobotics*, vol. 16, 2022. [Online]. Available: https://www.frontiersin.org/journals/neurorobotics/articles/10.3389/fnbot.2022.861825

[4] P. Seguin, C. Preault, P. Bidaud, and J.-P. Gazeau, "From specialized industrial grippers to flexible grippers: Issues for grasping and dexterous manipulation," *Foundations and Trends® in Robotics*, vol. 11, no. 1, pp. 1–89, 2023. [Online]. Available: http://dx.doi.org/10.1561/2300000074

[5] S. Kadalagere Sampath, N. Wang, H. Wu, and C. Yang, "Review on human-like robot manipulation using dexterous hands," *Cognitive Computation and Systems*, vol. 5, no. 1, pp. 14–29, 2023. [Online]. Available: https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/ccs2.12073

[6] Y. Huang, D. Fan, H. Duan, D. Yan, W. Qi, J. Sun, Q. Liu, and P. Wang, "Human-like dexterous manipulation for anthropomorphic five-fingered hands: A review," *Biomimetic Intelligence and Robotics*, p. 100212, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2667379725000038

[7] A. Billard and D. Kragic, "Trends and challenges in robot manipulation," *Science*, vol. 364, no. 6446, p. eaat8414, 2019. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.aat8414

[8] T. Feix, J. Romero, C. H. Ek, H.-B. Schmiedmayer, and D. Kragic, "A metric for comparing the anthropomorphic motion capability of artificial hands," *IEEE Transactions on Robotics*, vol. 29, no. 1, pp. 82–93, 2013.

[9] Rokoko, "Rokoko official website," 2025, accessed: 2025-01-16. [Online]. Available: https://www.rokoko.com

[10] Shadow Robot Company, "Dexterous hand series," 2024, accessed: 2024-12-26. [Online]. Available: https://www.shadowrobot.com/dexterous-hand-series/

[11] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, *et al.*, "Learning dexterous in-hand manipulation," *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.

[12] OpenAI, I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, J. Schneider, N. Tezak, J. Tworek, P. Welinder, L. Weng, Q. Yuan, W. Zaremba, and L. Zhang, "Solving rubik's cube with a robot hand," 2019. [Online]. Available: https://arxiv.org/abs/1910.07113

[13] Inmoov, "Inmoov Hand," 2024, accessed: 2024-12-26. [Online]. Available: https://inmoov.fr/

[14] DexHand Project, "DexHand," 2024, accessed: 2024-12-26. [Online]. Available: https://www.dexhand.org

[15] SimLab Co., Ltd., "Allegro Hand," 2024, accessed: 2024-12-26. [Online]. Available: https://www.allegrohand.com

[16] K. Shaw, A. Agarwal, and D. Pathak, "Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning," 2023. [Online]. Available: https://arxiv.org/abs/2309.06440

[17] B. Hirt, H. Seyhan, M. Wagner, and R. Zumhasch, *Hand and Wrist Anatomy and Biomechanics*, 2017th ed. Thieme Verlag, 2017, publication Title: Hand and Wrist Anatomy and Biomechanics. [Online]. Available: https://www.thieme-connect.de/products/ebooks/lookinside/10.1055/b-0036-140288

[18] Y. Toshimitsu, B. Forrai, B. G. Cangan, U. Steger, M. Knecht, S. Weirich, and R. K. Katzschmann, "Getting the ball rolling: Learning a dexterous policy for a biomimetic tendon-driven hand with rolling contact joints," in *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*, 2023, pp. 1–7.

[19] T. C. Pataky, M. L. Latash, and V. M. Zatsiorsky, "Multifinger ab- and adduction strength and coordination," *Journal of Hand Therapy*, vol. 21, no. 4, pp. 377–385, 2008. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0894113008000239

[20] Z.-H. Yin, B. Huang, Y. Qin, Q. Chen, and X. Wang, "Rotating without seeing: Towards in-hand dexterity through touch," 2023. [Online]. Available: https://arxiv.org/abs/2303.10880

[21] P. Weiner, C. Neef, Y. Shibata, Y. Nakamura, and T. Asfour, "An embedded, multi-modal sensor system for scalable robotic and prosthetic hand fingers," *Sensors*, vol. 20, no. 1, 2020. [Online]. Available: https://www.mdpi.com/1424-8220/20/1/101

[22] J. Egli, B. Forrai, T. Buchner, J. Su, X. Chen, and R. K. Katzschmann, "Sensorized soft skin for dexterous robotic hands," 2024. [Online]. Available: https://arxiv.org/abs/2404.19448

[23] A. Schmitz, M. Maggiali, L. Natale, B. Bonino, and G. Metta, "A tactile sensor for the fingertips of the humanoid robot icub," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 2212–2217.

[24] M.-J. Seo and J. C. Yoo, "Omnidirectional fingertip pressure sensor using hall effect," *Sensors*, vol. 21, p. 7072, 10 2021.

[25] S. Belkhale, Y. Cui, and D. Sadigh, "Data quality in imitation learning," 2023. [Online]. Available: https://arxiv.org/abs/2306.02437

[26] A. Handa, A. Allshire, V. Makoviychuk, A. Petrenko, R. Singh, J. Liu, D. Makoviichuk, K. V. Wyk, A. Zhurkevich, B. Sundaralingam, Y. Narang, J.-F. Lafleche, D. Fox, and G. State, "Dextreme: Transfer of agile in-hand manipulation from simulation to reality," 2024. [Online]. Available: https://arxiv.org/abs/2210.13702

[27] S. Dasari, O. Mees, S. Zhao, M. K. Srirama, and S. Levine, "The ingredients for robotic diffusion transformers," 2024. [Online]. Available: https://arxiv.org/abs/2410.10088

[28] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," 2024. [Online]. Available: https://arxiv.org/abs/2303.04137

[29] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," 2023. [Online]. Available: https://arxiv.org/abs/2304.13705

[30] Y. Qin, Y.-H. Wu, S. Liu, H. Jiang, R. Yang, Y. Fu, and X. Wang, "Dexmv: Imitation learning for dexterous manipulation from human videos," 2022. [Online]. Available: https://arxiv.org/abs/2108.05877

[31] D. Liconti, Y. Toshimitsu, and R. Katzschmann, "Leveraging pre-trained latent representations for few-shot imitation learning on an anthropomorphic robotic hand," in *2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids)*, 2024, pp. 181–188.

[32] A. Sivakumar, K. Shaw, and D. Pathak, "Robotic telekinesis: Learning a robotic hand imitator by watching humans on youtube," 2022. [Online]. Available: https://arxiv.org/abs/2202.10448