

DOGlove: Dexterous Manipulation with a Low-Cost Open-Source Haptic Force Feedback Glove

Han Zhang^{1,2}, Songbo Hu¹, Zhecheng Yuan^{1,2,3}, Huazhe Xu^{1,2,3}

¹ Tsinghua University, ² Shanghai Qi Zhi Institute, ³ Shanghai AI Lab

<https://do-glove.github.io/>


Motion Capture
21 DoFs



DOGlove




Shadow Hand




LEAP Hand


Haptic+Force Feedback
5 DoFs



a) Teleoperation without Visual Input



b) Regulating the Sauce Flow



c) In-Hand Rotation with Force-Based Extrinsic Dexterity

Fig. 1: **DOGlove**, a haptic+force feedback glove designed for precise teleoperation and dexterous manipulation. It features 21-DoF motion capture and 5-DoF haptic+force feedback. By leveraging action and force retargeting, it enables the teleoperation of dexterous hands for complex, contact-rich tasks, including: a) without visual input, adjusting contact force with a bottle during teleoperation, b) regulating the flow of condensed milk, and c) performing in-hand rotation by using haptic+force feedback to adjust friction.

Abstract—Dexterous hand teleoperation plays a pivotal role in enabling robots to achieve human-level manipulation dexterity. However, current teleoperation systems often rely on expensive equipment and lack multi-modal sensory feedback, restricting human operators’ ability to perceive object properties and perform complex manipulation tasks. To address these limitations, we present **DOGlove**, a low-cost, precise, and haptic+force feedback glove system for teleoperation and manipulation. **DOGlove** can be assembled in hours at a cost under 600 USD. It features a customized joint structure for 21-DoF motion capture, a compact cable-driven torque transmission mechanism for 5-DoF multidirectional force feedback, and a linear resonate actuator for 5-DoF fingertip haptic feedback. Leveraging action and haptic+force retargeting, **DOGlove** enables precise and immersive teleoperation of dexterous robotic hands, achieving high success rates in complex, contact-rich tasks. We further evaluate **DOGlove** in scenarios without visual input, demonstrating the critical role of haptic+force feedback in task performance. In addition, we utilize the collected demonstrations to train imitation learning policies, highlighting the potential and effectiveness of **DOGlove**. **DOGlove**’s hardware and software system are fully open-sourced at <https://do-glove.github.io/>.

I. INTRODUCTION

Imitation learning (IL) has shown significant promise in addressing complex manipulation tasks [7, 8, 50, 49]. However, it often necessitates a substantial amount of task-specific data to train a generalizable learning policy. Efficiently collecting and ensuring the high quality of such demonstrations remains a persistent and challenging problem for the robotic community.

Teleoperation is among the most commonly used methods for collecting demonstrations, often involving the development of a wide range of devices tailored to meet diverse data acquisition requirements. These devices enable human manipulation behaviors’ transfer to various robotic platforms [13, 12, 6, 10, 16]. However, when it comes to dexterous hands, their high degrees of freedom (DoFs) and inherent complexity impose even stricter demands on operational precision and the accuracy of human motion capture. Hence, it is crucial to design an intuitive, responsive, and highly precise device specifically suited for dexterous hand teleoperation applications.




Fig. 2: **Teleoperation demos.** a) While squeezing condensed milk, the operator regulates the flow using haptic+force feedback from DOGlove. b) The operator grasps a slipping bottle without visual input. c) The user identifies object pairs solely through haptic+force feedback.

Vision-based methods are primarily used for tracking the human hand in dexterous hand teleoperation. A simple approach utilizes RGB cameras [27, 39], but the accuracy of hand gesture capture is questioned and further limited by visual obstacles during hand-object interactions. Motion capture (MoCap) systems [24, 40, 29, 42] provide stable hand tracking, while relying solely on visual input for teleoperation makes intuitive control challenging for the human operator.

Haptic+force feedback offers additional sensory information beyond vision alone, such as an object’s texture, friction, and softness. Integrating haptic+force feedback into teleoperation can enrich the feedback available during interaction and enable the completion of more challenging tasks. Recently, some commercialized haptic+force feedback gloves [9, 32, 22, 15] have shown promise for enabling intuitive teleoperation. However, these solutions are often prohibitively expensive and require significant integration efforts to work with existing robot learning frameworks.

In this paper, we introduce **DOGlove**, a low-cost, fully open-sourced, and easy-to-manufacture haptic+force feedback glove for dexterous manipulation. The glove can be assembled in hours for a total cost under 600 USD. Key features of DOGlove include:

21-DoF motion capture: DOGlove features an anthropomorphic design resembling the human hand, providing precise motion capture and a comfortable wearing experience. Moreover, we propose a customized joint structure that integrates a compact, low-cost yet accurate joint encoder, with the entire assembly measuring less than 15 mm in thickness.

5-DoF haptic+force feedback: DOGlove leverages a cable-driven mechanism to deliver force feedback to each finger while maintaining a compact and cost-effective design. Additionally, each fingertip is also equipped with a linear resonant actuator (LRA) to provide realistic haptic feedback. This integration of force and haptic feedback creates an immersive

and responsive interface for dexterous manipulation.

Action and haptic+force retargeting: We propose a general retargeting framework. For action retargeting, the rigid constraints of the glove allow fingertip positions to be mapped from the human hand to the target robotic hand. For haptic+force retargeting, the combination strategy enables users to perceive contact information during teleoperation.

The resulting system, DOGlove, provides precise hand pose motion capture and the ability to sense interactions with manipulated objects. This enables human operators to intuitively and efficiently teleoperate dexterous hands. As shown in Fig. 2, it further supports the completion of complex manipulation tasks. We evaluate the necessity of haptic+force feedback through a user study and further assess the teleoperation efficiency and data accuracy of DOGlove in several quantitative experiments.

Finally, we demonstrate that DOGlove seamlessly integrates with existing methods in robot learning. We use DOGlove to teleoperate the LEAP Hand mounted on a Franka robot arm, collecting data to train imitation learning policies. To foster further research, we **open-source** the mechanical designs, circuit designs, embedded code, assembly instructions, URDF models, retargeting methods, and MuJoCo simulation environment at <https://do-glove.github.io/>.

II. RELATED WORK

A. Data Collection from Human Demonstrations

A substantial amount of task-specific data is essential for imitation learning. In dexterous manipulation, obtaining high-quality hand motion data is critical for training effective policies. Prior work includes extracting demonstration data from human videos [35, 44, 45, 2] and hand trajectories [41, 46]. While these approaches are accessible and have shown promising results, the significant visual gap between recorded human

demonstrations and the robot's perception often makes real-world transfer challenging. An alternative is using dedicated hardware for data collection to bridge this gap. Hand-held grippers [36, 8, 28] have proven effective in capturing robot manipulation data. However, these systems are primarily designed for parallel grippers. Another widely used approach is MoCap systems, which record human demonstrations and extract hand motion data. These systems include camera-based methods [27, 53, 25], glove-based tracking systems [14, 42, 20, 21, 37, 30], marker-based tracking [52], and commercial MoCap solutions [38, 11]. While MoCap offers high-precision tracking, bridging the embodiment gap between human and robotic hands remains a persistent challenge.

B. Dexterous Hand Teleoperation

Collecting high-quality human demonstrations through robotic teleoperation systems [12, 6, 10, 16] also plays a critical role for advancing dexterous manipulation. Existing research has explored teleoperation from various perspectives, including leader-follower setups such as ALOHA [50, 51, 13, 1]. However, teleoperating dexterous hands remains a significant challenge. OpenTelevision [6] leverages VR devices to capture hand poses and streams the pose information for retargeting to robotic hands. BiDex [34], on the other hand, implements a teleoperation system based on commercial motion capture gloves [22] and leader arms. Compared to these frameworks and other glove-based systems [20, 21], DOGlove offers distinct advantages. It eliminates the need for expensive equipment while precisely capturing fingertip positions and delivering richer haptic+force feedback to the operator. This system achieves accurate dexterous hand teleoperation with a low-cost setup, making it an efficient alternative.

C. Teleoperation with Haptic+Force Feedback

While recent studies rely on visual information to capture environmental characteristics, vision alone inherently limits the richness of available sensory data. In contrast, haptic+force feedback enhances the teleoperation experience by providing greater immersion and improving perception of the robot's status and movement compared to vision-based methods. Bunny-VisionPro [10] and Liu et al. [21] apply real-time haptic feedback to enable more accurate manipulation. Xu et al. [43] build a bilateral isomorphic bimanual telerobotic system using a commercial force feedback glove [9] to enhance perception and improve performance in complex tasks. NimbRo-Avatar [31] and Mosbach et al. [23] integrate commercial force feedback glove [32] into dexterous teleoperation systems. However, these approaches rely on specialized or expensive equipment. In contrast, DOGlove provides a highly accurate teleoperation system with integrated haptic+force feedback at a significantly lower cost and can be widely used in dexterous manipulation.

III. GLOVE DESIGN OBJECTIVES

DOGlove is designed to precisely capture human hand poses and provide haptic+force feedback for intuitive teleoperation.

While ensuring these functionalities, the glove is optimized for accessibility by the research community, focusing on low cost, ease of manufacturing, and high performance. To achieve these goals, DOGlove incorporates the following design principles:

A. Low cost

Commercial products such as the SenseGlove Nova [32] and Manus VR [22] cost more than 5,000 USD, making them prohibitively expensive for many researchers. In contrast, DOGlove provides a low-cost solution under 600 USD.

B. Ease of manufacturing

All parts of DOGlove are either readily available for purchase online or manufacturable using standard methods. The glove's main body can be 3D-printed using a commodity 3D printer, while the remaining electronics and servos are easily sourced. The entire glove can be assembled within 6 hours.

C. Performance Sufficiency

To ensure precise fingertip position tracking, the glove's encoders deliver joint angle data with an error range of $\pm 7.2^\circ$, which can be further minimized through careful calibration. For intuitive haptic+force feedback, the servos provide sufficient stall torque to halt human finger movement, while the haptic engine supports multiple haptic waveforms to enhance tactile sensations.

D. Low latency

The MoCap system operates at a maximum frequency of 120 Hz, while the haptic+force feedback system achieves a maximum frequency of 30 Hz. Together with the retargeting algorithm, the system ensures seamless operation at a minimum frequency of 30 Hz, providing a smooth and responsive teleoperation experience.

IV. HARDWARE SYSTEM

A. Kinematic Design

The kinematic design of DOGlove refers to the use of constraints to achieve desired movements, emulating the natural motion of a human hand. To ensure precise MoCap capabilities and a comfortable wearing experience, DOGlove are designed to closely resemble the anthropomorphic structure of the human hand.

Several studies [3, 4] model the hand skeleton as a kinematic chain, represented by a hierarchical structure of rigidly connected joints. As shown in Fig. 3, the kinematic structure of the human hand primarily consists of two types of joints: hinge joints and ball joints.

In the index, middle, ring, and pinky finger, distal interphalangeal (DIP) and proximal interphalangeal (PIP) joints are hinge joints with 1-DoF, allowing only flexion-extension movements. In contrast, the metacarpophalangeal (MCP) joint is a ball joint with 2-DoF, allowing both flexion-extension and adduction-abduction movements.

The thumb differs slightly in its structure. The interphalangeal (IP) and metaphalangeal (MCP) joints are hinge joints

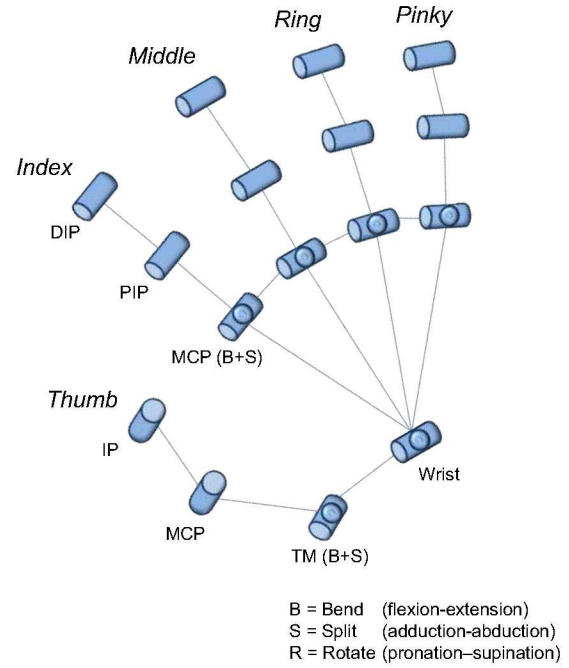



Fig. 3: **The kinematic structure of DOGlove**, designed to replicate the kinematics of a human hand. The MCP (B+S) and TM (B+S) joints are modeled as ball joints using a combination of two rotary joints. The right figure from [3] illustrates the simplified human hand kinematics.

with 1 DoF, while the additional trapeziometacarpal (TM) joint is a ball joint that supports both flexion-extension and adduction-abduction movements. To further enhance dexterity, an additional DoF at the wrist allows the thumb to perform pronation-supination movements.

To implement these joints in DOGlove, hinge joints are modeled as two linkages connected by a rotary joint, with a joint encoder installed on the rotary axis to capture motion. Ball joints are designed as a combination of two orthogonal rotary joints, each equipped with a joint encoder on its respective rotary axis.

The linkage lengths in DOGlove are designed to accommodate the majority of adult human sizes. To achieve this, a standard human finger length was first modeled, and the glove's linkage parameters were simulated to ensure an optimal range of motion. As shown in Fig. 4, improper linkage lengths can obstruct the natural flexion-extension of the fingers, leading to discomfort and reduced MoCap performance. Furthermore, DOGlove features a modular design where all fingers share a common structural framework. This modularity enables users to replace linkages with customized sizes as needed, enhancing both adaptability and usability.

B. Finger Design

As shown in the exploded-view in Fig. 3, DOGlove is composed of the thumb, index, middle, ring, and pinky finger assemblies, along with the palm base structure. The design of each finger assembly follows a modular approach, ensuring consistent structural elements across all fingers.

The exploded view of a single finger is illustrated in Fig. 5. The highlighted area indicates the basic components of a rotary joint. Each rotary joint is constructed using an M4×15 shoulder




Fig. 4: **Improper linkage lengths can cause collisions** between the human finger (link f, g) and the glove (link m), restricting finger movements and leading to discomfort and poor MoCap performance.

screw to connect the finger linkages, ball bearing, and joint encoder, secured with an M3 locknut. This design ensures smooth and reliable joint rotation. The main body of the finger, colored white and gold, is 3D printed using PETG material for ease of fabrication and durability.

Given the limited space on the back of the human hand, the finger assembly's width is constrained to less than 26




Fig. 5: Exploded view of the finger assembly, with the highlighted area indicating the basic components of a rotary joint.

mm. Simultaneously, to provide effective force feedback, the actuator must deliver a stall torque of at least 0.5 N·m. Additionally, adjustable stiffness requires the actuator's current to be regulated. Since the actuator is directly connected to the pulley system as a rotary joint for MCP_B , it is essential to measure its rotary position in real time to achieve precise joint angle control. Considering these design requirements, the Dynamixel XC/XL330 servo motors were selected as the actuators for force feedback. It fulfills the torque, size, and real-time position measurement needs, making it a suitable choice for DOGlove.

1) Joint Encoders:

To integrate joint encoders into the finger linkages, the encoders must be compact while maintaining high precision. Additionally, as 16 encoders are required in combination with 5 servo motors to achieve 21-DoF MoCap capabilities, the cost of each encoder needs to be affordable. Considering these constraints, we selected the Alps RDC506018A rotary sensor as the joint encoder. This compact encoder (W11 mm × L14.9 mm × H2.2 mm) is easily integrated into the 3D-printed joint structures. The encoder operates as a variable resistor, changing its resistance as the shaft rotates.

The resistance changes are converted into voltage signals using a simple voltage divider circuit. Due to the encoder's linear response, the voltage output is proportional to the actual joint angle. These voltage signals are read by an Analog-Digital Converter (ADC) module. For precise conversion, we use the TI ADS1256, a low-noise 24-bit ADC operating at 30k samples per second. The converted signals are then sent to a microcontroller unit (MCU), the ST Electronics STM32F042K6T6, which operates at a clock speed of 48 MHz. To optimize system performance and reduce OS scheduling overhead, the STM32's Direct Memory

Access (DMA) feature is utilized to accelerate joint encoder readings. Finally, the processed joint data is transmitted to the host machine via a serial port on the STM32.

The voltage readings are mapped directly to joint angles under the assumption that the supply voltage of the STM32 (approximately 3.3 V) corresponds to 360°, while the ground voltage (0 V) corresponds to 0°. Using the ADC output voltage, the joint angle is calculated as:

$$\alpha_{\text{joint}} = \frac{V_{\text{ADC}}}{V_{\text{CC}}} \cdot 360 \quad (1)$$

The primary error in this conversion comes from the linearity error of the encoder which is $\pm 2\%$ according to its datasheet. This results in an angular error of $\pm 7.2^\circ$ when measuring joint angles. To mitigate this, we employ a calibration process. Using an external high-precision joint encoder, we map the voltage reading to an actual joint angle, creating a correction table for each encoder. With this calibration, the error can be reduced to within $\pm 1^\circ$.

2) Cable-Driven Force Feedback Structure:

Force feedback refers to the feedback derived from applied forces and torques. To provide force feedback on the human fingers, the output torque of the Dynamixel servo must be transmitted to the glove's finger linkage system. As illustrated in Fig. 5, the rotary axis of the servo and the rotary axis of the MCP_B joint are misaligned. Consequently, a transmission mechanism is required to transfer the torque effectively.

Although a bevel gear system could serve as a potential solution, its implementation would require significant space to accommodate the two orthogonal gears. Additionally, transmitting large torque through gears can cause deformation in the gear shaft, leading to gear slippage. In contrast, a cable-driven

mechanism offers a more compact design while ensuring stable torque transmission.

Traditional cable-driven systems typically provide uni-directional force transmission on the tension side, relying on a spring to generate force in the opposite direction. However, this approach introduces unrealistic feedback sensations. While using two servos per finger could resolve this issue, it would significantly increase the glove's weight and cost.




Fig. 6: **Pulley system** of the cable-driven mechanism.

To address these challenges, DOGlove utilizes a pulley system to provide the bi-directional force feedback, as shown in Fig. 6. DOGlove uses a 0.6 mm stainless steel braided wire as the cable, chosen for its strength and durability. The *Servo Pulley* connects the servo to the finger linkage (*Finger Middle*) via the *Finger Pulley*, maintaining a 1:1 transmission ratio. To minimize friction during transmission, the *Fixed Pulley* is used to redirect the cable's path. When the *Servo Pulley* rotates clockwise, the tension on *Cable B* increases, causing *Finger Pulley* to rotate clockwise. The extra slack on the *Cable A* side is taken up by the *Servo Pulley A*. Since the finger linkage is fixed to the *Finger Pulley*, it also rotates clockwise, resulting in the extension movement of the MCP_B joint. Similarly, when the *Servo Pulley* rotates counterclockwise, the tension shifts to *Cable A*, producing a flexion movement of the MCP_B joint.

This configuration enables bi-directional torque transmission with a simple, compact, and cost-effective design.

3) **Fingertip Haptic Feedback:**

Haptic feedback refers to the feedback of tactile or kinesthetic properties. To further enhance the operator's tactile experience, each fingertip in DOGlove is equipped with a tactile actuator.

Traditional haptic actuators include eccentric rotating mass (ERM) motors and linear resonate actuators (LRAs). Limited by the inertia of the rotating mass, ERM motors are slow to start and stop, making it challenging to produce complex waveforms needed for subtle tactile sensations. On the contrary, LRAs offer linear motion, resulting in a cleaner and more precise tactile output.

In DOGlove, we use LRAs with a diameter of 8 mm and a height of 2.5 mm, installed close to the fingertips. These

LRAs provide vibration stimuli by resonating at approximately 240 Hz along Z axis, which is orthogonal to the fingertip surface. Operating at 1.2 V_{rms} , the LRAs generate high-quality haptic waveforms. To fully leverage the potential of the LRA, we employ the TI DRV2605L motor driver, which includes the licensed Immersion TouchSense® 2200 haptic library. This driver supports over 100 pre-programmed waveforms, allowing DOGlove to deliver realistic and refined haptic feedback.

C. **Wrist Localization**

In DOGlove, we design a shell with a 1/4 inch screw connector to accommodate external wrist localization devices. For our experiments, we use the HTC Vive Tracker for real-time wrist position tracking. However, the design is compatible with other solutions, depending on the user's requirement.

V. **RETARGETING**

A. **Action Retargeting**

To map human hand gestures to a robotic hand, it is essential to perform action retargeting, which converts motion data from the glove into robotic hand movements. This process addresses both the embodiment gap and motion discrepancies. Previous studies [42, 11, 38] highlight the significance of fingertips, as they are the primary contact area during object interactions. Building on this insight, we apply the 5-DoF haptic+force feedback to the human operators' fingertips and adopt a retargeting method focused on fingertip positions.

Our approach combines Forward Kinematics (FK) to compute human fingertip positions and Inverse Kinematics (IK) to calculate the corresponding robotic hand positions. When wearing DOGlove, the human operator secures their fingertips inside the finger caps. Since the glove acts as a rigid body, the relative positions of the fingertips with respect to the glove's origin can be accurately calculated. With DOGlove's anthropomorphic kinematic design and precise MoCap capabilities, fingertip positions are effortlessly determined using the glove's built-in FK. To map these positions to a robotic hand, we utilize Mink [48], a differential inverse kinematics library, to generate smooth and feasible motions for the robotic hand.

A size discrepancy often exists between the human hand and the target robotic hand. To address this, we introduce a scaling factor when calculating IK, allowing adaptation to different robotic hand sizes. This ensures an intuitive teleoperation experience where the robotic hand naturally mirrors the human hand's gestures. For instance, when the human operator opens their hand, the robotic hand open proportionally. Similarly, when the human operator brings their thumb and index finger together, the robotic hand's thumb and index finger also touch. This ability for precise fingertip alignment is critical for tasks like grasping small objects.

In our experiments, we deploy the system on the LEAP hand [33] in real-world scenarios and test it with various robotic hands in the MuJoCo simulator. Fig. 7 presents the

retargeting results of DOGlove in both simulation and real-world environments.




Fig. 7: **Action retargeting results:** Teleoperating the LEAP Hand to grasp a toy in the real world and teleoperating the Shadow Hand, Inspire Hand, and Allegro Hand in simulation.

B. Haptic+Force Retargeting

To provide the haptic+force feedback, it is first necessary to sense tactile or force information at the robotic hand’s fingertips. This can be achieved using a simple tactile sensor, such as the force sensing resistor (FSR) sensor [10, 17], or for better performance, by utilizing an F/T sensor [18, 5] or a vision-based tactile sensor [47, 19].

In our experimental setup, we install a 1-D force sensor on each fingertip of the LEAP Hand, with a measurement range of 3 kg and a precision of 1 g. During our quantitative experiments (Section VI), we identify a combination strategy for integrating haptic and force feedback that optimizes performance. This strategy along with the corresponding thresholds and feedback patterns is summarized in Table I. These three thresholds are set with values chosen based on optimal performance observed across all test trials.

Force Sensor Readings (g)	Haptic Feedback	Force Feedback
<10	×	×
10–50	✓	×
50–100	✓	✓
>100	×	✓

Table I: **The combination strategy** for haptic+force feedback in DOGlove.

The first threshold (10g): This threshold is used to distinguish actual contact forces from sensor noise. An incorrect setting of this threshold would either introduce noise or lose meaningful force information, both of which can reduce the success rate.

The second threshold (50g): This threshold determines when to introduce force feedback. During a user study w/o visual input (Section VI-A), we observe that human operators are highly sensitive to force feedback. To create a more realistic experience, force feedback is applied only after the force sensor readings exceed 50 g, which serves as the second threshold. A lower value would lead to premature activation, interfering with the human operator’s perception and typically lowering the success rate. Conversely, a higher value increases reliance on haptic feedback, making success dependent on its performance.

The third threshold (100g): This threshold is used to stop haptic feedback. Observations from the bottle-slipping

experiment (Section VI-B) reveal that continuous haptic feedback during teleoperation can create a misleading sensation, interfering with the operator’s ability to perceive the subtle properties of the object’s surface. To address this, a third threshold is set at 100 g. A lower value shifts reliance toward force feedback, while a higher value creates more overlap between simultaneous haptic and force feedback. This overlap can interfere with human perception and usually results in a lower success rate.

For force feedback, the Dynamixel servos operate in current-based position control mode. The force readings from the LEAP Hand fingertips are clamped to the range [0g, 3000g], and mapped linearly to the K_P gain of the Dynamixel servos. For haptic feedback, we use waveform ID 56 from the haptic engine library, corresponding to Pulsing Sharp 1–100%.

This combination strategy for haptic+force retargeting enables human operators to distinguish object shape, size and softness without visual input. It also improves performance in complex, contact-rich manipulation tasks. Further details are provided in Section VI.

VI. EXPERIMENTS

In this section, we use DOGlove to teleoperate the LEAP Hand [33] mounted on the Franka Robot Arm to evaluate its effectiveness through a series of challenging tasks across three key aspects:

- **Haptic+Force Perception:** Without visual input, how effectively can DOGlove assist human operators in perceiving object properties through haptic+force feedback?
- **Teleoperation Efficiency:** Does integrating haptic force feedback improve vision-based teleoperation success rates and reduce task completion time? Can DOGlove enable human operators to perform challenging, contact-rich manipulation tasks?
- **IL Compatibility:** Can the data collected via DOGlove be leveraged to train IL policies for dexterous manipulation?

Evaluation Setup: To evaluate the effectiveness of haptic force perception, we conduct a user study (Section VI-A) and a quantitative experiment (Section VI-B). Teleoperation efficiency is assessed in Experiment VI-C, while IL compatibility is evaluated in Experiment VI-D.

Comparisons: All experiments share the following comparison conditions, although a subset of these may be selected depending on the specific task setup:

- **Only Force:** Force feedback is enabled only when the force sensor readings exceed 10 g.
- **Only Haptic:** Haptic feedback is enabled only when the force sensor readings exceed 10 g.
- **Haptic+Force:** A combined feedback strategy is applied, as detailed in Section V-B.
- **No Haptic/Force:** DOGlove is used solely for MoCap, with no feedback provided.
- **Baseline:** AnyTeleop [27], a widely recognized vision-based hand retargeting method, is used as the MoCap baseline.

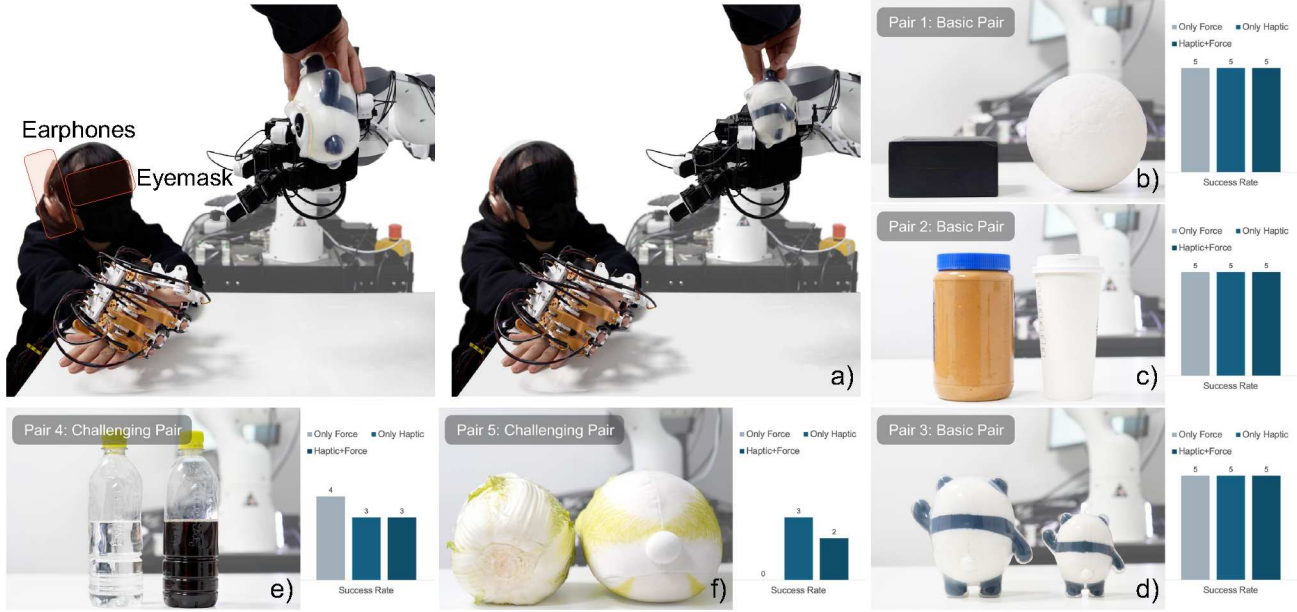


Fig. 8: **User Study.** (a) Experiment setup: Users wear an eyemask and headphones to eliminate visual and auditory input. (b)–(f) Object pairs tested in the study. All feedback modes perform well for the basic pairs. For the challenging pairs, force feedback is more sensitive to softness, while haptic feedback is more sensitive to shape.

A. User Study: Object Perception w/o Visual Input

Task: Five untrained human operators participate in this user study. During the experiment, they are required to distinguish between five pairs of objects solely through feedback from DOGlove, without any visual or auditory input (achieved by wearing an eyemask and headphones). In each trial, a pair of objects is randomly selected, and users provide their answers immediately after experiencing feedback from DOGlove for both objects. Fig. 8 illustrates the experiment setup and the five object pairs, selected based on factors such as shape, size, and softness.

Metrics: Users’ ability to distinguish object pairs is evaluated based on their success rate.

Challenges: The five object pairs are intentionally chosen based on the following considerations:

- Pair 1: Basic Pair, different shape. The ball and the box have distinctly different shapes (Fig. 8b).
- Pair 2: Basic Pair, similar shape, different size. The peanut bottle and the coffee paper cup share a similar cylindrical shape, but their diameters differ slightly (Fig. 8c).
- Pair 3: Basic Pair, similar softness, different size. The two toys have similar softness and shapes but vary in size (Fig. 8d).
- Pair 4: Challenging Pair, similar size and shape, different softness. Two identical bottles are used, one filled with pure water (soft) and the other filled with carbonated cola, shaken to increase its hardness (Fig. 8e).
- Pair 5: Challenging Pair, similar shape, different size and softness. A toy cabbage (softer, larger) and a real cabbage (Fig. 8f).

Performance: As shown in Fig. 8, even without visual and auditory input, all participants effortlessly distinguish basic pairs 1-3. For challenging pair 4, most participants can perceive softness using only force feedback. Some also discern

softness using only haptic feedback by evaluating the duration of contact during deformation.

For challenge pair 5, when the robotic hand grasps the softer toy cabbage, it deforms to resemble the size of the real cabbage. This deformation increases its perceived softness, making it difficult for participants to distinguish using force feedback alone.

For both challenge pairs, combining haptic and force feedback slightly reduces user sensitivity, leading to a marginally lower accuracy.

B. Bottle-Slipping

In both experimental setups, the human operator is required to perform a bottle-slipping action. Since the LEAP Hand is not equipped with shear sensors, slipping information is not directly transmitted to DOGlove. Instead, we infer slip indirectly through the relationship between friction and the applied normal force. A larger normal force increases the friction between the LEAP Hand and the bottle, resulting in a lower slipping velocity. In our experiments, we measure the slipping distance and task success rate (determined by whether the bottle falls) as indicators of the user’s ability to regulate normal force effectively.

1) Teleoperation w/o Visual Input

Task: In this experiment, the human operator must perform a bottle-slipping action relying solely on feedback from DOGlove. A 15-second countdown timer is set for each trial. If the bottle successfully slips without falling within the 15 seconds, the trial is denoted as successful.

Metrics: The success rate.

Challenges: Without any visual or auditory input (achieved by wearing an eyemask and headphones), the operator must determine if the bottle is slipping at the right speed or too quickly, risking a fall.




Fig. 9: **Teleoperation experiments and quantitative results.** a) Without visual input, force feedback significantly improves the task success rate. With visual input, it enhances precise control. b) In in-hand rotation, the challenge is to slightly release the fingers, allowing the carton to rotate without slipping out (as shown in the middle two images).

Performance: As shown in Fig. 9a, force feedback significantly improves the success rate of this task. Additionally, incorporating haptic feedback further enhances overall performance. However, since the fingers of the LEAP Hand maintain continuous contact with the bottle during the task, haptic feedback does not provide additional information beyond using the glove solely as a MoCap device, resulting in the same success rate for both conditions.

Due to differences in retargeting strategies, even small changes in human finger positions can significantly affect LEAP Hand movements. As a result, AnyTeleop [27] struggles to perform the slipping task effectively.

2) Teleoperation with Visual Input

Task: Unlike the previous blindfolded experiment, this setup allows operators visual input. To further evaluate the operator’s control ability, they are required to slip the bottle to a specified distance (9 cm). A trial is successful if the bottle slips without falling. Additionally, We measure the deviation between the actual slipping distance and the target distance (9 cm).

Metrics: Performance is evaluated using two metrics:

- **Success Rate:** A trial is denoted as successful if the bottle slips without falling.
- **Slipping Deviation:** This measures the difference between the target sliding distance (9 cm) and the actual slipping distance, with a smaller deviation indicating greater operational accuracy.

Challenges: Operators must precisely control the bottle to achieve the desired distance. While a greedy approach often causes the bottle to fall and results in failure, a conservative approach leads to an unsatisfactory distance deviation.

Performance: This task evaluates not only success rate but also teleoperation precision. To minimize slipping deviation, operators are instructed to control the LEAP Hand carefully and optimally. As shown in Fig. 9a, similar to previous results,

haptic+feedback does not provide additional information and may even interfere with task precision. However, force feedback enables operators to minimize slipping deviation more effectively. While using DOGlove solely as a MoCap device achieves the same success rate as with haptic+force feedback, it results in a larger average slipping deviation.

C. Rotating and Placing the Carton

Task: This is a long-horizon contact-rich task. As shown in Fig. 9b, the operator must first pick up the carton horizontally, then perform an in-hand rotation, orienting the carton vertically before placing it into a small bucket.

Metrics: Performance is evaluated using two metrics:

- **Success Rate:** A trial is denoted as successful if the carton rotates more than 45 degrees and is successfully placed into the bucket.
- **Completion Time:** The total time taken to complete the entire process.

Challenges:

- **Precise Manipulation:** The operator must accurately teleoperate to rotate the carton while preventing it from falling.
- **Visual Obstacle:** Grasping the carton is hindered by visual obstacles, as the operator cannot see the contact points between the robotic hand’s fingers and the carton.

	Success Rate	Average Completion Time (s)
Only Force	9/10	18.92
Only Haptic	9/10	21.16
Haptic+Force	10/10	19.89
No Haptic/Force	4/10	24.76
AnyTeleop	1/10	54.85

Table II: **Quantitative experiment results.** Haptic+force feedback enables operators to achieve a higher success rate and a faster average completion time, as haptic feedback provides contact information, while force feedback indicates the proper timing for in-hand rotation.

Performance: Table II shows that both haptic and force




Fig. 10: **The imitation learning experiment.** (a) The robot must first locate the correct position of the box and then apply adequate force to press it. Excessive force prevents movement, while insufficient force causes the fingers to slip. (b) The robot must first locate the bear, then open its hand to grasp it. Due to the bear’s size, precise grasping control is required. An inaccurate grasp deforms the bear and causes it to slip out of the fingertips.

feedback significantly improve the teleoperation success rate and reduce completion time. Force feedback alone achieves a similar average completion time, but combining haptic and force feedback yields a higher success rate. The vision-based MoCap method AnyTeleop [27] struggles with in-hand rotation in this task.

D. Imitation Learning

We show DOGlove is capable of collecting high-quality demonstrations. **3D Diffusion Policy (DP3)** [49] is selected as our imitation learning algorithm, and we use *Realsense L515* to acquire the point cloud inputs, which are then down-sampled to 1024 points using farthest point sampling [26]. The data collected by DOGlove is used to train policies for various downstream tasks. We evaluate imitation learning performance on 2 basic contact-rich tasks and 1 long-horizon task:

Press and Move Box: As shown in Fig. 10a, the robot must continuously press down on a box and move it to a specified target location. During data collection, the box is randomly placed within a 30×20 cm area, and DOGlove collects 40 demonstrations to train the policy. In evaluation, the box is also randomly placed in the same area. Across 20 trials, the success rate is 85% (17/20).

Pick and Place Teddy Bear: As shown in Fig. 10b, the robot must grasp a teddy bear and place it into a designated box. During data collection, the teddy bear’s initial position is randomized within a 30×20 cm area, and DOGlove collects 40 demonstrations to train the policy. In evaluation, the bear is again randomly placed in the same area. Across 20 trials, the success rate is 70% (14/20), with failures primarily due to the teddy bear slipping out of the robotic hand when not grasped firmly.

Rotating and Placing the Carton. This task follows the same setup as Section VI-C. For this contact-rich task, we use 3 human-collected demonstrations to train the policy. Across 10

trials, the success rate is 90% (9/10).

VII. LIMITATIONS AND FUTURE WORK

DOGlove is a powerful haptic+force feedback glove for dexterous manipulation, but several limitations remain. First, the weight of DOGlove is inevitably high, as it utilizes 5 commercial servos, bringing the total weight to 450g. Second, in agile teleoperation scenarios, performance is constrained by the servos’ maximum speed and torque output. Servos use gearboxes to balance torque and speed, resulting in an inverse relationship between these two factors. Higher servo speed allows DOGlove to capture rapid movements such as playing a piano, while still providing haptic and force feedback. However, increasing speed reduces torque, potentially limiting realistic force sensations at higher speeds. Thanks to DOGlove’s modular design, users prioritizing speed can easily replace the default XC330-M288 servo with lower gear-ratio options like XC330-M188 or XL330-M077, according to their needs. Third, although DOGlove is designed to accommodate most hand sizes, it may be uncomfortable for some users. To enhance adaptability and wearability, we develop CAD files for linkages in multiple sizes, enabling customization for various hand sizes.

VIII. CONCLUSION

In this paper, we present DOGlove, a low-cost, open-source haptic+force feedback glove designed for dexterous manipulation. DOGlove enables precise and efficient execution of long-horizon, contact-rich tasks. Experimental results show that DOGlove enhances the operator’s immersive teleoperation experience while also serving as an effective tool for training imitation learning policies. Moreover, the user study demonstrates that DOGlove provides precise perception of object properties through its integrated haptic+force feedback. To support further research and contribute to the community, all hardware designs and code are open-sourced.

ACKNOWLEDGMENT

We would like to thank Zhengrong Xue, Gu Zhang, Changyi Lin, Mengda Xu, and Yifan Hou for their invaluable advice and fruitful discussions on hardware design and learning policies. We also appreciate Wenhao Ding and Laixi Shi for their insightful discussions and feedback. Additionally, we thank Yichuan Gao, Xiaoyan Yang, Xinyao Qin, and Botian Xu for their assistance with the user study. Special thanks to Skyentific, Gennady Plyushchev, for their innovative contributions to the unconventional cable-driven joint design. We are also grateful to Tiansheng Sun and Guanghan Pan for their open-source repository of the HTC Vive Tracker Python API. We extend our appreciation to Yitong Wang for her help in creating elegant graphic renderings of the hardware design. Tsinghua University Dushi Program supports this project.

REFERENCES

- [1] Jorge Aldaco, Travis Armstrong, Robert Baruch, Jeff Bingham, Sanky Chan, Kenneth Draper, Debiddatta Dwibedi, Chelsea Finn, Pete Florence, Spencer Goodrich, et al. Aloha 2: An enhanced low-cost hardware for bimanual teleoperation. *arXiv preprint arXiv:2405.02292*, 2024.
- [2] Homanga Bharadhwaj, Abhinav Gupta, Vikash Kumar, and Shubham Tulsiani. Towards generalizable zero-shot manipulation via translating human interaction plans. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6904–6911. IEEE, 2024.
- [3] Ilaria Cerulo, Fanny Ficuciello, Vincenzo Lippiello, and Bruno Siciliano. Teleoperation of the schunk s5fh under-actuated anthropomorphic hand using human hand motion tracking. *Robotics and Autonomous Systems*, 89: 75–84, 2017.
- [4] Pietro Cerveri, Elena De Momi, N Lopomo, Gabriel Baud-Bovy, RML Barros, and Giancarlo Ferrigno. Finger kinematic modeling and real-time hand motion estimation. *Annals of biomedical engineering*, 35:1989–2002, 2007.
- [5] Claire Chen, Zhongchun Yu, Hojung Choi, Mark Cutkosky, and Jeannette Bohg. Dexforce: Extracting force-informed actions from kinesthetic demonstrations for dexterous manipulation. *arXiv preprint arXiv:2501.10356*, 2025.
- [6] Xuxin Cheng, Jialong Li, Shiqi Yang, Ge Yang, and Xiaolong Wang. Open-television: Teleoperation with immersive active visual feedback. In *8th Annual Conference on Robot Learning*, 2024.
- [7] 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*, page 02783649241273668, 2023.
- [8] 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. In

- Proceedings of Robotics: Science and Systems (RSS)*, 2024.
- [9] Dexta Robotics. Dexta robotics official website. <https://www.dextarobotics.com/>, 2025.
- [10] Runyu Ding, Yuzhe Qin, Jiyue Zhu, Chengzhe Jia, Shiqi Yang, Ruihan Yang, Xiaojuan Qi, and Xiaolong Wang. Bunny-visionpro: Real-time bimanual dexterous teleoperation for imitation learning. *arXiv preprint arXiv:2407.03162*, 2024.
- [11] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J Black, and Otmar Hilliges. Arctic: A dataset for dexterous bimanual hand-object manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12943–12954, 2023.
- [12] Hongjie Fang, Hao-Shu Fang, Yiming Wang, Jieji Ren, Jingjing Chen, Ruo Zhang, Weiming Wang, and Cewu Lu. Airexo: Low-cost exoskeletons for learning whole-arm manipulation in the wild. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 15031–15038. IEEE, 2024.
- [13] Zipeng Fu, Tony Z. Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. In *Conference on Robot Learning (CoRL)*, 2024.
- [14] Xiaochi Gu, Yifei Zhang, Weize Sun, Yuanzhe Bian, Dao Zhou, and Per Ola Kristensson. Dexmo: An inexpensive and lightweight mechanical exoskeleton for motion capture and force feedback in vr. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 1991–1995, 2016.
- [15] HaptX. Haptx official website. <https://haptx.com/>, 2025.
- [16] Aadithya Iyer, Zhuoran Peng, Yinlong Dai, Irmak Guzey, Siddhant Haldar, Soumith Chintala, and Lerrel Pinto. OPEN TEACH: A versatile teleoperation system for robotic manipulation. In *8th Annual Conference on Robot Learning*, 2024.
- [17] Zhanat Kappasov, Juan-Antonio Corrales, and Véronique Perdereau. Tactile sensing in dexterous robot hands. *Robotics and Autonomous Systems*, 74: 195–220, 2015.
- [18] Uikyum Kim, Heeyeon Jeong, Hyunmin Do, Jongwoo Park, and Chanhun Park. Six-axis force/torque fingertip sensor for an anthropomorphic robot hand. *IEEE Robotics and Automation Letters*, 5(4):5566–5572, 2020.
- [19] Changyi Lin, Han Zhang, Jikai Xu, Lei Wu, and Huazhe Xu. 9dtact: A compact vision-based tactile sensor for accurate 3d shape reconstruction and generalizable 6d force estimation. *IEEE Robotics and Automation Letters*, 2023.
- [20] Hangxin Liu, Xu Xie, Matt Millar, Mark Edmonds, Feng Gao, Yixin Zhu, Veronica J Santos, Brandon Rothrock, and Song-Chun Zhu. A glove-based system for studying hand-object manipulation via joint pose and force sensing. In *2017 IEEE/RSJ International Conference on*

- Intelligent Robots and Systems (IROS)*, pages 6617–6624. IEEE, 2017.
- [21] Hangxin Liu, Zhenliang Zhang, Xu Xie, Yixin Zhu, Yue Liu, Yongtian Wang, and Song-Chun Zhu. High-fidelity grasping in virtual reality using a glove-based system. In *2019 international conference on robotics and automation (icra)*, pages 5180–5186. IEEE, 2019.
 - [22] Manus Meta. Manus meta robotics official website. <https://www.manus-meta.com/robotics>, 2025.
 - [23] Malte Mosbach, Kara Moraw, and Sven Behnke. Accelerating interactive human-like manipulation learning with gpu-based simulation and high-quality demonstrations. In *2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*, pages 435–441. IEEE, 2022.
 - [24] OptiTrack. OptiTrack Motion Capture Systems. <https://www.optitrack.com/>, 2025.
 - [25] Younghyo Park, Jagdeep Singh Bhatia, Lars Ankile, and Pulkit Agrawal. Dexhub and dart: Towards internet scale robot data collection. *arXiv preprint arXiv:2411.02214*, 2024.
 - [26] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
 - [27] Yuzhe Qin, Wei Yang, Binghao Huang, Karl Van Wyk, Hao Su, Xiaolong Wang, Yu-Wei Chao, and Dieter Fox. Anyteleop: A general vision-based dexterous robot arm-hand teleoperation system. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
 - [28] Felipe Sanches, Geng Gao, Nathan Elangovan, Ricardo V Godoy, Jayden Chapman, Ke Wang, Patrick Jarvis, and Minas Liarokapis. Scalable, intuitive human to robot skill transfer with wearable human machine interfaces: On complex, dexterous tasks. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6318–6325. IEEE, 2023.
 - [29] Gaspare Santaera, Emanuele Luberto, Alessandro Serio, Marco Gabiccini, and Antonio Bicchi. Low-cost, fast and accurate reconstruction of robotic and human postures via imu measurements. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2728–2735. IEEE, 2015.
 - [30] Ioannis Sarakoglou, Anais Brygo, Dario Mazzanti, Nadia Garcia Hernandez, Darwin G Caldwell, and Nikos G Tsagarakis. Hexotrax: A highly under-actuated hand exoskeleton for finger tracking and force feedback. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1033–1040. IEEE, 2016.
 - [31] Max Schwarz, Christian Lenz, Andre Rochow, Michael Schreiber, and Sven Behnke. Nimbrow avatar: Interactive immersive telepresence with force-feedback telemanipulation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5312–5319. IEEE, 2021.
 - [32] SenseGlove. Senseglove official website. <https://www.senseglove.com/>, 2025.
 - [33] Kenneth Shaw, Ananye Agarwal, and Deepak Pathak. Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning. *Robotics: Science and Systems (RSS)*, 2023.
 - [34] Kenneth Shaw, Yulong Li, Jiahui Yang, Mohan Kumar Srirama, Ray Liu, Haoyu Xiong, Russell Mendonca, and Deepak Pathak. Bimanual dexterity for complex tasks. In *8th Annual Conference on Robot Learning*, 2024.
 - [35] Aravind Sivakumar, Kenneth Shaw, and Deepak Pathak. Robotic telekinesis: Learning a robotic hand imitator by watching humans on youtube. In *Robotics: Science and Systems*, 2022.
 - [36] Shuran Song, Andy Zeng, Johnny Lee, and Thomas Funkhouser. Grasping in the wild: Learning 6dof closed-loop grasping from low-cost demonstrations. *IEEE Robotics and Automation Letters*, 5(3):4978–4985, 2020.
 - [37] Yuan Su, Gaofeng Li, Yongsheng Deng, Ioannis Sarakoglou, Nikos G Tsagarakis, and Jiming Chen. The joint-space reconstruction of human fingers by using a highly under-actuated exoskeleton. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9645–9651. IEEE, 2024.
 - [38] Omid Taheri, Nima Ghorbani, Michael J Black, and Dimitrios Tzionas. Grab: A dataset of whole-body human grasping of objects. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 581–600. Springer, 2020.
 - [39] Ultraleap. Ultraleap official website. <https://www.ultraleap.com/>, 2025.
 - [40] Vicon Motion Systems. Vicon Motion Capture Systems. <https://www.vicon.com/>, 2025.
 - [41] Chen Wang, Linxi Fan, Jiankai Sun, Ruohan Zhang, Li Fei-Fei, Danfei Xu, Yuke Zhu, and Anima Anandkumar. Mimicplay: Long-horizon imitation learning by watching human play. In *7th Annual Conference on Robot Learning*, 2023.
 - [42] Chen Wang, Haochen Shi, Weizhuo Wang, Ruohan Zhang, Li Fei-Fei, and C Karen Liu. Dexcap: Scalable and portable mocap data collection system for dexterous manipulation. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024.
 - [43] Han Xu, Mingqi Chen, Gaofeng Li, Lei Wei, Shichi Peng, Haoliang Xu, and Qiang Li. An immersive virtual reality bimanual telerobotic system with haptic feedback. *arXiv preprint arXiv:2501.00822*, 2025.
 - [44] Mengda Xu, Zhenjia Xu, Cheng Chi, Manuela Veloso, and Shuran Song. XSkill: Cross embodiment skill discovery. In *7th Annual Conference on Robot Learning*, 2023.
 - [45] Mengda Xu, Zhenjia Xu, Yinghao Xu, Cheng Chi, Gordon Wetzstein, Manuela Veloso, and Shuran Song. Flow as the cross-domain manipulation interface. In *8th*

Annual Conference on Robot Learning, 2024.

- [46] Jingyun Yang, Junwu Zhang, Connor Settle, Akshara Rai, Rika Antonova, and Jeannette Bohg. Learning periodic tasks from human demonstrations. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8658–8665. IEEE, 2022.
- [47] Wenzhen Yuan, Siyuan Dong, and Edward H Adelson. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors*, 17(12):2762, 2017.
- [48] Kevin Zakka. Mink: Python inverse kinematics based on MuJoCo, July 2024. URL <https://github.com/kevinzakka/mink>. Software.
- [49] Yanjie Ze, Gu Zhang, Kangning Zhang, Chenyuan Hu, Muhan Wang, and Huazhe Xu. 3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024.
- [50] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- [51] Tony Z Zhao, Jonathan Tompson, Danny Driess, Pete Florence, Seyed Kamyar Seyed Ghasemipour, Chelsea Finn, and Ayzaan Wahid. Aloha unleashed: A simple recipe for robot dexterity. In *8th Annual Conference on Robot Learning*, 2024.
- [52] Wenping Zhao, Jinxiang Chai, and Ying-Qing Xu. Combining marker-based mocap and rgb-d camera for acquiring high-fidelity hand motion data. In *Proceedings of the ACM SIGGRAPH/eurographics symposium on computer animation*, pages 33–42, 2012.
- [53] Christian Zimmermann, Duygu Ceylan, Jimei Yang, Bryan Russell, Max Argus, and Thomas Brox. Freihand: A dataset for markerless capture of hand pose and shape from single rgb images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 813–822, 2019.