Robotics: Science and Systems 2025
Los Angeles, CA, USA, June 21–25, 2025

Demonstrating LEAP Hand v2: Low-Cost, Easy-to-Assemble, High-Performance Hand for Robot Learning

Kenneth Shaw, Deepak Pathak
Carnegie Mellon University

Fig. 1: Our demonstration will feature four different low-cost, open-source robotic hand designs, including a new model introduced in this paper, LEAP Hand V2. (left) These hands are highly dexterous, easy to build, durable, and affordable. We also provide a suite of open-source software tools, including motion capture teleoperation, human video-based learning, and reinforcement learning capabilities. Building on our successful demonstrations at RSS 2023 and 2024, we aim to further highlight the potential of low-cost, open-source robotic hands and strengthen our open-source robot hand community at RSS 2025. Please visit our website at https://leaphand.com for more details.

***Abstract*—Replicating human-like dexterity in robotic hands has been a long-standing challenge in robotics. Recently, with the rise of robot learning and humanoids, the demand for dexterous robot hands to be reliable, affordable, and easy to reproduce has grown significantly. To address these needs, we present LEAP Hand v2, a $200 8-DOF highly dexterous robotic hand designed for robot learning research. It is strong yet compliant, using a hybrid rigid-soft structure that is very durable. Its universal dexterous MCP joint provides exceptional finger mobility, enabling a variety of different grasps. The parts are all 3D printed and can be assembled very easily in under two hours using our instructions. Importantly, we offer a suite of advanced open-source software tools to support robot learning research. This includes human video retargeting code from MANO and Vision Pro, motion capture teleoperation code using the Manus Glove, and a URDF with simulation examples for various simulation engines. We will showcase LEAP Hand v2—designed specifically for this demonstration—alongside our previous robot hands with real robot interactive demos. Following our successful demos at RSS 2023 and 2024, we will again offer an engaging opportunity for attendees to get hands-on experience and information about the accessibility of low-cost, open-source robotic hands. Please visit our website at https://leaphand.com**

I. Introduction

Think about activities such as typing on your keyboard, hammering a nail, or using chopsticks, and you'll realize the pivotal role our hands play in manipulating the world. With remarkable strength at the fingertips, capable of over 70 different pinching and grasping motions, our hands possess unparalleled abilities to manipulate. This extraordinary sensing and adaptability are orchestrated by both our hand "hardware" itself as well as the impressive capabilities of our brains. The development of our brains is often linked to the necessity of manipulating our surroundings with our hands [1, 2].

In the realm of robotics, manipulation has predominantly relied on claw grippers or suction cups for pick-and-place tasks in factories. However, the collective aspiration is to witness humanoid robots coexisting with humans, undertaking similar tasks in similar environments. The absence of robot humanoids with efficient and low-cost robotic hands raises the question: Why haven't they become a reality?

One major bottleneck is that while there are a few robot hands are available today, the prevailing opinion is that they are challenging to use, expensive, and difficult to acquire. The belief has been that the human kinematic structure and strength is difficult to produce in robot hands. Some robot hands are too large, some have fewer degrees of freedom and other are extremely difficult to produce and maintain. We believe this isn't necessarily an inherent flaw in robot hands but rather a consequence of not designing them ideally for machine learning research.

To break through this prevailing belief, we introduce LEAP Hand V2: Low-cost Easy-to-Assemble high-Performance robot hand for robot learning. LEAP Hand V2 is our latest robot hand designed for this demonstration that has these characteristics:

1) **Hybrid Rigid-Soft Hand:** Combines flexibility and compliance with exceptional strength and durability in each finger.
2) **Universal MCP Joint:** This universal abduction-adduction mechanism enables extreme dexterity for many tasks.
3) **Extreme Low-cost and Reproducibility:** At $200 and with under an hour of assembly time, it makes the hand readily available for many robot learning researchers.

Robot hands designed for machine learning require both a physical design and advanced software tools to support research. By developing open-source tools, the goal is to enable the community to build upon and iterate these tools for a variety of applications and reshare them with the community. Our open-source tools serve as a solid foundation for advancing machine learning research:

1) **Motion Capture Teleoperation:** Motion-capture technologies, such as the Manus Meta Glove, enable highly precise teleoperation of the LEAP Hand v2 with any robot arms.
2) **Learning from Human Video:** A wide range of freely available human video sources can be used to teach robot hands like LEAP Hand v2 to mimic human-like behaviors.
3) **Simulation:** A kinematically accurate simulator allows us to perform both forward and inverse kinematics and geometric grasping analysis.

In summary, we present a new robot hand designed specifically for this demo, LEAP Hand v2, that is affordable at under $200 and incredibly easy to produce and assemble in under 1 hour. It is extremely dexterous, durable and with a human-like size. In this RSS demo, we show its ability to handle a variety of machine learning applications such as teleoperation and real-world reinforcement learning. These tools are released on our website for the community to use. These RSS demos have played a crucial role in the success of open-source robot hands and inspire attendees to incorporate open-source, dexterous robot hands into their own manipulation projects. They have captivated conference attendees, even those not interested in dexterous manipulation. We look forward to the opportunity of showcasing LEAP Hand V2 and our other robot hands at RSS 2025 and on our website at https://leaphand.com

II. Related Work

Dexterous Robot Hands Many robotic hands have been developed to mimic the capabilities of the human hand, with varying degrees of success and accessibility. The MIT/Utah Hand has an early tendon-driven design [4]. [5, 6, 7] used variable stiffness actuators and soft materials in their IIT hands. [8] used a novel compliant actuator and [9] developed this area of tendon-driven hands further. [5, 10] use underactuation similar to our hand to enable high dexterity without using too many actuators which makes the hand large and heavy. A dexterous all-soft hand, with palm articulations in a completely soft structure, is presented in [11].

While these robot hands are very advanced, they are difficult to reproduce and obtain for many research labs. There are now a few robot hands for purchase. The Shadow hand, as documented in [12, 13], has demonstrated remarkable achievements such as in-hand reorientation of a Rubik's cube [14]. Despite its impressive performance, the Shadow hand is widely acknowledged for its high cost (approximately $150k) and challenging usability. Conversely, the Allegro Hand [15, 16], has been historically recognized as a more affordable option priced at $20k. However, it is often criticized for its tendency to break down and the associated difficulties in repair. Nevertheless, the Allegro Hand has showcased commendable capabilities, including teleoperation from video [17, 18, 19, 20, 21], as well as in-hand reorientation [22]. The Psyonic Ability Hand, designed as a prosthetic with a robust internal hard skeleton and soft exterior but only possesses 6 degrees of freedom (6DOF) [23]. Inspire Hand is lower in costs but is not durable and often breaks[24]. The Faive Hand [25] demonstrates noteworthy sim2real results in in-hand reorientation but is not readily available yet.

A resurgence of interest in humanoid robotics from industry players like Tesla Optimus [26], Figure [27], BD Atlas [28], 1x [29], Sanctuary AI [30], and Digit [31] has been observed. These hands are designed for strength and mass production to handle daily tasks for humanoids. However, they often feature limited degrees of freedom and are not readily available for purchase, evaluation, or research purposes.

The emergence of rapid-prototyping technologies, such as 3D printers and CNCs, has led to the development of a plethora of low-cost, open-source hands tailored for academic research purposes. The LEAP Hand, detailed in our papers [3, 32] is easy to use and has been used by many research labs around the world. The Robel suite, exemplified by D'Manus, offers

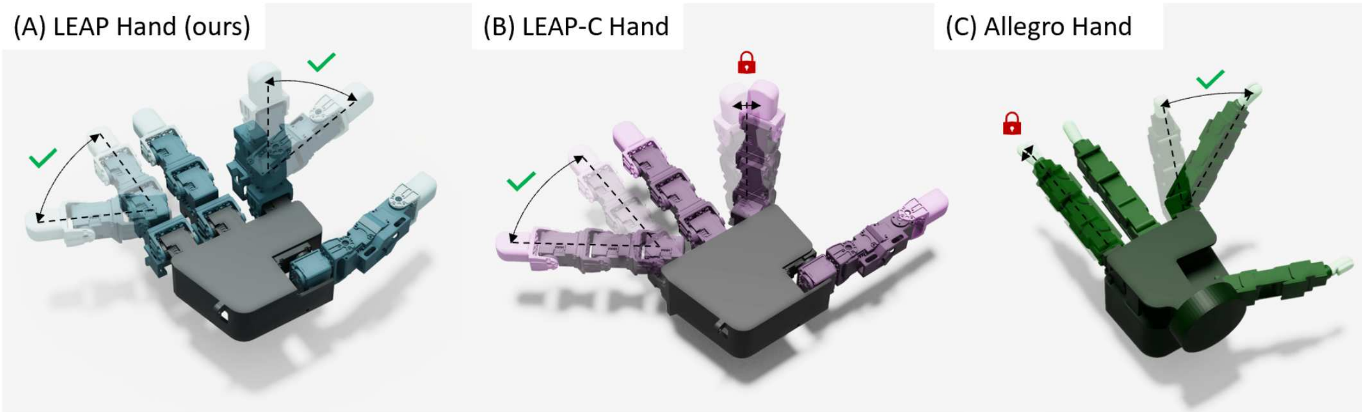

Fig. 3: In all of our LEAP Hands we introduce a MCP joint that allows for abduction and adduction in both flexed and extended positions. In this figure we show LEAP Hand v1 on the left, where this dexterous kinematic structure is most apparent. In a conventional robot hand, LEAP-C Hand, the finger can move side to side in the open-palm, but in the flexed position it only spins in place. In Allegro, there is a large of motion at flexed but not in the extended position. [3]

large yet durable hands employed in tasks such as reorientation [33] and grasping [34]. Other hands, such as Inmoov [35] and DexHand [36], cater to hobbyists but may be limited by inexpensive motors or fragile plastic components.

Rapid Manufacturing The traditional method for creating robust components typically involves machining, such as with aluminum, which can be costly. Plastic parts, on the other hand, are generally produced through a process that includes mold creation, casting, curing, and support removal, making it suitable for large-scale production [37]. In contrast, 3D printing has transformed small-scale manufacturing by enabling the rapid, autonomous printing of individual parts [38, 39].

In recent years, the 3D printing industry has made substantial progress in material innovation. Flexible materials like TPU/TPE from Ninjatek and Filaflex have introduced new possibilities for creating more adaptable components [40, 41]. Foaming materials such as Colorfabb Varioshore and Recreus Filaflex [42, 41] allow for the adjustment of material properties by modulating the flow rate. Additionally, the use of materials like Nylon and carbon fiber in 3D printing has resulted in components with enhanced strength and durability. As a result, consumer-friendly multimaterial 3D printers have become both affordable and widely accessible.

Learning for Dexterous Manipulation Because of the high dimensionality of dexterous manipulators, it is different to used traditional model-based controls or planning methods to achieve dexterous results. [43, 44] In robot learning Andrychowicz et al. achieved in-hand rotation for various objects using a Shadow hand and Sim2real techniques. [45, 14]. Simulation-based training that scales to thousands of objects is explored in works such as [46, 47, 48, 32, 22] which shows promise in robot learning. D'Hand is utilized by Nair et al. to reposition a valve [49]. Other notable instances of dexterous manipulation include Baoding Balls' in-hand rotation using the Shadow Hand trained exclusively in the real world [50].

Recent studies highlight the importance of supervising robot hand policies based on human actions, such as those derived from MANO [51] human hand parameters. Related work includes teleoperating robot hands through real-time video [18, 17], which can assist in learning [20, 52, 53]. Hand poses extracted from online video data are utilized for learning manipulation policies [52, 54]. Large-scale pre-training using internet videos has proven effective for training robot hands efficiently for downstream tasks with minimal task-specific demonstrations [20, 55, 56], and this approach extends to non-dexterous manipulation tasks [57, 58].

III. LEAP Hand v2

The LEAP Hand v2 is a low-cost, highly dexterous robotic hand designed specifically for robot learning. It combines soft and rigid materials through multi-material 3D printing to achieve human-like compliance and durability, while keeping the design simple and easy to reproduce. With underactuated joints, a novel MCP mechanism, integrated tactile sensing, and an intuitive assembly process, the hand is purpose-built for machine learning workflows. This section details the hardware innovations that enable LEAP Hand v2 to serve as a practical and accessible platform for robotic manipulation research.

A. 3D printed Hybrid Rigid-Soft Hand

The flexibility of human hands enables them to adapt to objects and their surroundings during interactions like grasping [59, 60]. For example, when reaching for a fragile object, the hand molds around it and applies gentle pressure. Similarly, when encountering an obstacle like a table, our fingers bend away without breaking. To replicate this, the soft robotics field often employs a casting method to create soft robots [61, 62]. However, this approach can result in robot hands being excessively compliant in under-actuated directions, which is not ideal. This makes the hand weak and objects can slip out. In contrast, some robot hands, like the LEAP Hand V1 or Shadow Hand, consist entirely of rigid joints. While these hands can impart a lot of force, they lack the ability to conform to their environment as a human hand does. This leads to brittleness and a tendency to snap upon contact. Finally, other designs, such as those in [25, 63], attempt to replicate the softness of human hand joints but are highly complex and difficult to reproduce.

Our goal is to fabricate conformal fingers with stiffness properties closely resembling those of a human finger, aiming to replicate both its softness and rigidity. We aim to ensure our robotic hand is durable and capable of withstanding the rigors of manipulation while also being easy to produce. The careful selection of materials and a suitable 3D printer is essential to achieve this hybrid rigid soft hand finger as seen in Figure 4.

For the soft outer skin, Recreus Filaflex Foamy is chosen due to its strong restitution and its ability to foam up buy taking in air from the environment when leaving the nozzle. This enables one to adjust the TPU rubber density based on the flow rate and temperature of the 3D printer nozzle. [42] Underneath the skin and within the fingers, we opt for PLA, known for its rigidity and smooth texture. The 3D printer used is an Independent Dual Extrusion (IDEX) printer, specifically the $1000 Snapmaker J1S [64]. This printer allows for the simultaneous use of two materials automatically in one seamless print. The realization of these materials together in one seamless print includes intricately intertwining the soft and rigid materials in CAD and the slicer software.

These material properties are used in a few key areas. The exterior layer of the palm's skin is crafted with denser material to resist cuts and bruises. The exterior of the fingers are created using denser high-flow Foamy material which makes them resistant to bruises. The interior of the finger between the flexure joints is printed with PLA bones which resist undesirable twisting and compression of the fingers as seen in Figure 4. This makes the fingers very strong. The internal flexure joints themselves are created with lower-density soft material to enable easy flexure actuation. The strength of each flexure joint in the fingers (MCP, PIP, DIP) are carefully tuned relative to each other by modulating the density as explained in the next section.

1) Underactuated Curling

A key factor in designing a robot hand is determining the appropriate number of underactuated and fully actuated degrees of freedom (DOF) the robot hand should have. Many works such as [65, 10] use underacted DOF to extend the

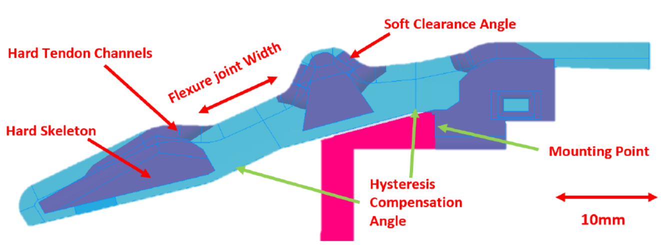

Fig. 4: **Finger Analysis** We show the key dimensions that create our soft-hard hybrid flexure joint as described in Section III-A. The flexure joint is soft material that lies between the hard skeleton. It is governed by the flexure joint width, the clearance angle, the hysteresis compensation angle and the size of the hard skeleton. The darker portions in the finger is the hard tendon channels and rigid structure. The blue material encompassing that is the soft material of the finger itself.

capabilities of these manipulators. Recent hands designed for humanoids keep the number of motors very small to make the hands light and durable [24]. LEAP Hand V2 is designed with 8 motors, two for each finger—one for curl and one for abduction—which contributes to the hand's light weight and simplicity in construction. This configuration results in substantial underactuation, with a single motor controlling three curling joints (MCP Forward, PIP, DIP) on each finger simultaneously as shown in Figure 5. However, after further analysis, we have identified methods to manage and adjust the underactuated behavior to improve its suitability for grasping.

Consider this example: when the curl motor activates the pulley, it shortens the tendon to a new length, applying a force F on the finger to achieve this movement. The flexure joints in the finger then deform to distribute this force from the motor. An important characteristic of these flexure joints is that they are initially easy to close, but as they approach their fully closed position, they increasingly resist additional force making it harder and harder to curl.

Suppose each joint in the finger is treated equally. In this scenario, each joint would experience the same amount of movement, with one-third of the tendon shortening force being distributed to each joint. As the finger curls, the motor increasingly exerts more effort to achieve the motion. However, this uniform motion is not ideal. In manipulation studies of the human hand, the MCP joint typically curls with a greater magnitude than the PIP and DIP joints to grasp over the palm and around objects

To replicate this behavior, our key insight is that each joint in the robot hand can be designed with different strengths. For instance, a weaker joint will generate less resistive force throughout its range of motion and will therefore actuate more. This means that a weaker joint will move further with the same tendon force compared to a stronger joint that is coupled with it. As the actuation force increases through the range of motion, the stronger joints will still actuate, but to a lesser extent relative to the weaker joint. The ratio of strength between the weaker and stronger joints determines the proportion of actuation across the different finger joints.

There are various ways to weaken a joint, one of the most straightforward being to print it thinner. However, thinning the joint makes it more vulnerable to twisting and compression, which are undesirable movements that reduce the finger's manipulation capabilities. Instead, we opt to adjust the flow rate of each hinge in the finger. By modifying the flow rate, we can easily control the relative density of the joints which affects strength of the joints while minimizing the risk of twisting and compression.

The MCP joint is designed using the lowest density material, followed by the PIP joint, and then the DIP joint. This design ensures that the MCP joint experiences the greatest actuation, with progressively smaller amounts of actuation for the PIP and DIP joints. By carefully selecting materials with varying densities for each joint, we can fine-tune the actuation behavior to better replicate natural human finger movement, where the MCP joint typically curls more significantly than the PIP and

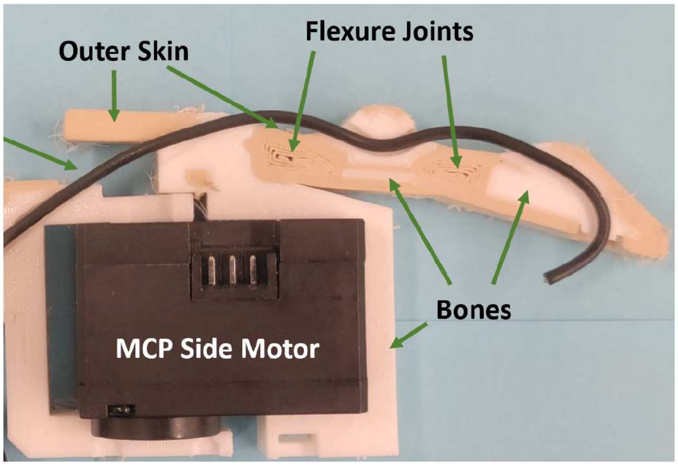

Fig. 5: A cross-section of the 3D printed finger showing the soft rubber joints, hard PLA bones, and resilient, dense outer skin. The MCP side rotates using an embedded motor, and the PIP and DIP joints are actuated together by a tendon connected to a pulley.

DIP joints.

A potential drawback is that a weaker joint is more compliant or less strong in the actuated direction, as an external disturbance allows the joint to move further with the same amount of force. This effect is particularly noticeable when the finger is open, as the flexure joint and tendon are not under tension. Wen the finger is curled, the tension still means that the DIP joint requires more force to actuate and becomes the most resistant to the applied force.

This effect is mitigated in two ways. First, since the tendon remains at approximately the same length from the motor when an external disturbance is applied, pushing on the MCP joint causes the PIP and DIP joints to curl more in compensation. This behavior is actually beneficial, as it enables the finger to curl and conform more effectively around the object being grasped without crushing it.

Secondly, the motor's controller is adjusted. A simple approach is to use current control, which enables the motor to apply a constant force to the robot's finger. However, this method has a drawback: it allows for greater compliance because the motor doesn't pull hard to maintain tendon length when external forces are applied. This weakens the grip. Position-based PID control offers a better solution by enabling the motor to resist external disturbances. This approach allows the motor to actively reduce compliance, applying more force when the tendon is pulled. As a result, the position controller stiffens the finger joints, improving the finger's ability to grip heavier and larger objects.

While it is possible to incorporate external position sensors and set position targets for the fingertips directly, this approach is costly and challenging to implement on a low-cost, easily produced 3D-printed hand. Furthermore, the inclusion of such sensors and their control policy could result in reduced compliance for the fingers, potentially counteracting the desired flexibility.

B. Dexterous MCP Joint

Off-the-shelf motors impose limitations on kinematic structure choices, making it challenging to accurately replicate the MCP ball joint of the human hand. As a result, it is typically approximated using two motors (MCP-1, MCP-2) positioned close together [66]. Previous work has proposed two designs for this configuration (Fig. 3). However, both the Allegro and LEAP C-Hand designs sacrifice one degree of freedom either in the extended or closed position. Consequently, the Allegro hand is less dexterous when extended, while the LEAP C-Hand (similar to the C-Hand in [66]) is less dexterous when closed.

The reason for the lost dexterity in both LEAP C-Hand and Allegro is that the axis of the motor responsible for adduction-abduction (MCP-2) is fixed to the palm of the hand. In LEAP C-Hand, the axis is perpendicular to the plane of the palm, whereas, in Allegro, it lies in the plane of the palm. Thus, when the finger becomes parallel to this axis, that DoF is ineffective. Please see Figure 3 for a visualization of this deficiency for these baseline hands.

In LEAP Hand v2, a **universal abduction-adduction mechanism** is used for the fingers, ensuring that dexterous motion is maintained at every MCP position. Rather than fixing the MCP-2 axis to the palm (i.e., the motor responsible for adduction-abduction), the key innovation is to align the axis after the first finger joint and orient it so that it remains perpendicular to the finger at all times. This design enables the finger to achieve adduction-abduction across all positions as shown more visibly on LEAP Hand V1 in Fig. 3. As a result, LEAP Hand V2 provides the best of both worlds and has both adduction-abduction in the extended position (similar to the LEAP C-Hand) and pronation/supination in the flexed position (similar to the Allegro).

C. Purpose built for Robot Learning

For machine learning research, a robot hand must not only be easy to assemble, maintain, and modify but also designed with flexibility in mind to accommodate rapid iterations and experimental changes. Given the iterative nature of machine learning development, having a hand that is low-cost and robust is crucial to minimize the barrier to entry and enable frequent updates without significant downtime or expense. Moreover, a highly dexterous design ensures the hand can perform a wide range of tasks, making it adaptable to various learning scenarios. Achieving these goals requires a thoughtful balance between functionality and simplicity, ensuring the system remains reliable while offering the versatility needed for research.

To start, we have reduced the total number of parts and ensured that all components used are easy to source. By utilizing a multi-material printer, we can create complex parts with multiple integrated components in a single print. For instance, the palm piece combines a soft top, a rigid undercarriage, tendon channels, motor mounts, the MCP forward joint, and an attachment for the MCP side servo horn—all in one simple, easy-to-print part. Similarly, the finger is printed as a single piece, incorporating the PIP and DIP

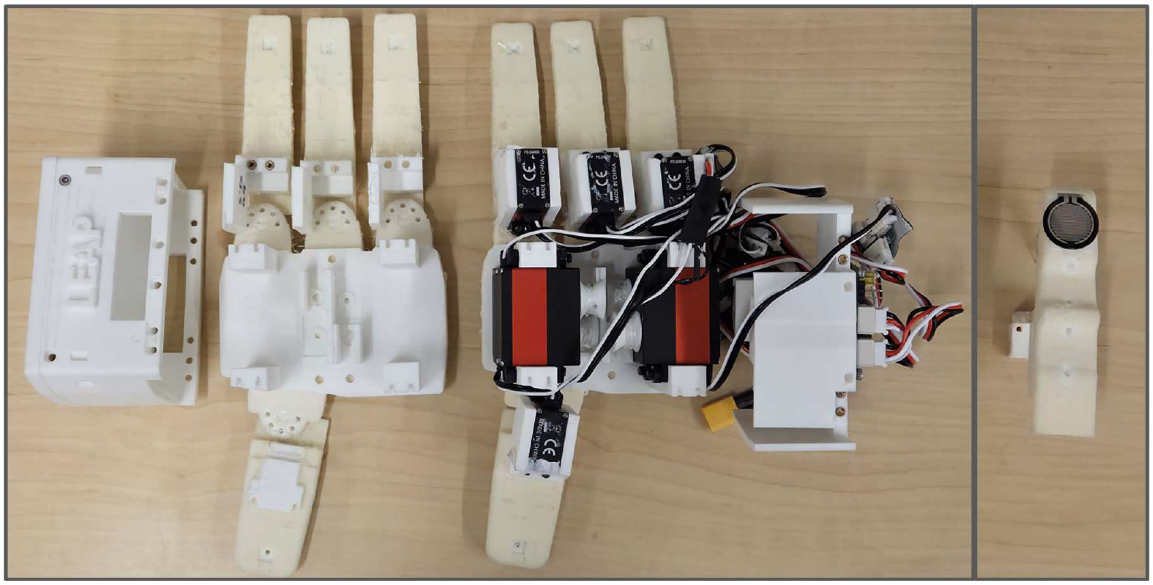

Fig. 6: The multi-material components, shown on the left, are fabricated using a dual-extrusion 3D printer and are designed for straightforward assembly in only 1 hour. The eight actuators are inserted into designated slots within the palm and fingers, followed by basic tendon routing and wiring to complete the setup. To facilitate widespread adoption, we release all print files, assembly instructions, and accompanying software as open-source materials at https://leaphand.com, enabling rapid replication by the research community. Additionally, off-the-shelf tactile sensors can be affixed to the fingertips to enhance sensing capabilities. (right)

joints, tendon channels, and motor mounts. This streamlined design reduces the overall part count (fewer than 10 3D-printed parts in total) and simplifies the assembly process, making it more manageable for the end user as seen in Figure 6.

For motor selection, we opted for Feetech Bus-Based Current-Limited control servos. These servos are cost-effective (under $30 each) and easily accessible. They are simple to program, and their internal PID loop parameters can be customized for both position-based current-limited control and current-control modes. Additionally, the built-in sensors provide useful data, such as current and position readings. This functionality is comparable to the more expensive Dynamixel servos used in previous versions of the LEAP Hand and other robotic hands, but at a significantly lower cost and in more suitable sizes for this robot hand.

A common challenge with tendon-based robot hands is the need for frequent retensioning and replacement of worn tendons. However, our design uses only one tendon per finger, and the pulleys are securely housed within channels in the palm, preventing the tendons from slipping off the motors. Unlike designs like the Shadow Hand [12] or DASH Hand [21], which require tensioning of tendon pairs, our system simplifies calibration, allowing it to be easily adjusted through automated software.

The assembly process is designed to be simple, with no complicated steps or hard-to-reach components. The motors are easily accessible and can be installed with just a screwdriver and a basic knot on the fishing line. Thanks to the intricate multi-material 3D-printed parts, the number of screws needed for assembly is minimal. While LEAP Hand v1 requires nearly 300 screws, LEAP Hand V2 uses fewer than 50 screws.

We make this comprehensive documentation available on our website at https://leaphand.com It includes an updated shopping list, detailed assembly instructions, and all the example code necessary to build LEAP Hand V2. This will enable anyone to easily replicate the robot hand and quickly begin working on their own projects.

D. Tactile Sensors

Since LEAP Hand v2 is user-friendly and affordable, we aim to ensure that the tactile sensors are equally easy to use, cost-effective, and capable. However, there is a significant challenge. Vision-based sensors like Gelsight or Digit, which rely on cameras and a gel-based touch surface, are too large for our needs. Similarly, sensors such as ReSkin, AnySkin, and Xela are designed for larger fingers. While they could theoretically be miniaturized, they often face interference issues when placed too closely together, as in a smaller hand.

For our hand, we have chosen resistive-based sensors similar to STAG [67] or 3D-ViTac [68]. These sensors are affordable ($2 each), and a LEAP Touch control board ($20) can be mounted on the palm of the hand. While they provide only a single continuous force reading per finger, this is adequate for many machine learning applications. Additionally, these sensors can be mounted on Manus glove [69] to collect data both hand tracking and touch data in a format that closely mirrors what the robot hand will encounter.

IV. Software for Robot Learning

LEAP Hand V2 is designed specifically for robot learning, both in terms of the aforementioned hardware and the software which we describe below. The software package we release on our website includes all the essential tools that researchers need to get started quickly. First, teleoperation using motion capture

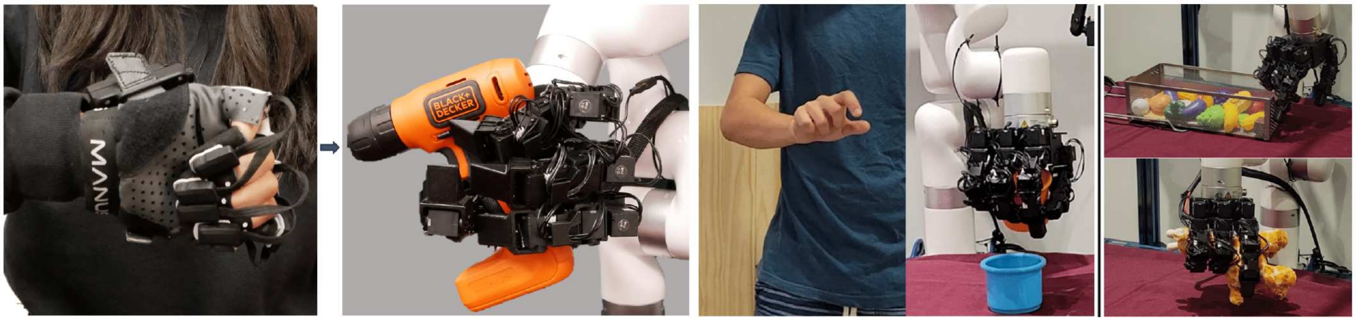

Fig. 7: We will demonstrate real robot hands doing teleoperation from VR glove and teleoperation from human video as developed in [3, 17, 32] Attendees will be able to teleoperate a variety of low-cost robot hands such as LEAP Hand V1 shown in this figure or LEAP Hand v2 as discussed in the paper.

gloves is a key feature for collecting high-quality data for behavior cloning and related approaches. Retargeting human video data, easily captured with web cameras or the Apple Vision Pro, is useful for control and enhancing autonomous policy training. Finally, we offer simulation tools for forward and inverse kinematics as well as kinematic geometric analysis. These software tools are crucial for LEAP Hand V2 which designed for machine learning. These tools are available on our website at https://leaphand.com.

A. Learning from Human Mocap Glove Demonstrations

In conventional 2-finger gripper manipulation many teleoperation setups have worked successfully to collect demonstrations for use in behavior cloning. Kinesthetic Methods such as ALOHA [70], GELLO [71] or Da Vinci machines [72] can work accurately. With VR or camera-based hand tracking, methods such as [73, 74] work reasonably accurately. However, it is unclear how to scale these methodologies to robot hands.

Our key insight is to take inspiration from the motion-capture community. The motion capture community has been using gloves for accurate tracking for movie production or game production. These motion capture gloves often rely on EMF sensors and are relatively not bulky to the user. However, the data returned is in the human hand morphology which is not the same as the robot hand kinematics.

If the robot hand's kinematic structure is roughly similar to that of a human, one approach is to map the joint angles from human fingers to the robot's. While gloves can calculate these human hand angles using an inverse kinematics solver, differences in finger size and proportions can lead to misalignment, especially with the complex human thumb and the differing kinematics that robot hands have. This can cause inaccurate pinch grasps which makes accurate fine-manipulation difficult.

Previous work has addressed this by optimizing joint positions for consistent pinch grasps between human and robot hands. [75, 18, 76] We employ Manus gloves [69] with an inverse kinematics approach to ensure accurate pinch grasps and proper thumb positioning, improving manipulation reliability.

We have collaborated with Manus Meta on our full-featured open-source Python/ROS2 repository that converts the MANUS data to any robot hand that has a URDF. We have examples for all LEAP Hands and are the main repository for these motion capture gloves for robotics.

Using this software to collect data, behavior cloning policies can be trained to complete a variety of different tasks that during the event such as tool use or deformable manipulation. We will show these policy rollouts, augmented with human-video pretraining on our robot hands. However, teleoperation data does not scale to novel environments so it must be collected in large quantities in many different environments. Our key insight is to leverage data from the web to enable further generalization which we will explain in the next section.

B. Learning from Human Video

Collecting teleoperated demonstrations is both costly and time-consuming, and it is impossible to capture all potential scenarios the robot may encounter. This leads to a distribution shift, where the robot's testing environments will inevitably differ from its training environments. While one might suggest that robots can adapt or generalize to unseen scenarios, this is not guaranteed. Therefore, a key question is how can we effectively expand our training set without resorting to additional laborious teleoperated data collection?

Given that robot hands share a similar embodiment and kinematic structure with the human hand, they both will naturally interact with the world and perform tasks in much the same way. By utilizing this similarity, we can learn from extensive human motion data, such as video and motion capture datasets. However, converting from human hand demonstrations to robot hand demonstrations is a non-trivial problem.

This conversion to the robot hand is particularly challenging due to the under-constrained nature of the problem, as robot hands and human hand possess numerous degrees of freedom (DOF) and exhibit substantial differences in shape, size, and joint structure. The retargeting process must cater to any human operator attempting to execute various tasks in diverse environments. Additionally, an essential criterion is the efficiency of the solution, demanding a real-time performance at a rate exceeding 30 Hz. Unlike the glove data, this data can be exceptionally noisy and difficult to use. To handle these issues, we develop a video-to-robot hand retargeting system that is trained from a corpus of rich and diverse human hand

Fingertip Strength (N)	Curled	Open	Side
Allegro Hand [15]	8.5	10.5	7.5
Inmoov Hand [35]	5.8	2.5	-
Bauer et. al [39]	37.4	-	-
DASH Hand [21]	34.5	8.4	6.2
LEAP Hand [3]	27.5	19.5	16.2
LEAP Hand V2	32.2	20.5	17.2

TABLE I: **Strength Test:** We apply force to both the curled and open parts of the finger, as well as the side of an open finger. The force is measured when the angle error exceeds 15 degrees. LEAP Hand V2 is highly powerful in all directions, thanks to its rigid MCP side joint, internal rigid bone structure, and efficient power transfer from robust servo motors.

videos. The system understands human hands and retargets the human video stream into a robot hand-arm trajectory that is smooth, swift, safe, and semantically similar to the guiding demonstration. This methodology has a few stages. First, we detect the human hand in the image by using a state of the art hand detector such as FrankMocap. [77] Then, we retarget this robot hand pose using a NN trained on an energy function and human data. This ensures that the retargeted robot hand poses are human like and semantically similar to the human hand demonstration. These elements are combined together to teleoperate the robot hand and arm in Robotic Telekinesis [17] from RSS 2022. See Figure 7 for an example.

We also show that these internet videos can directly support the learning of autonomous policies, rather than just assist with teleoperation. To achieve this, we can apply this Telekinesis retargeting pipeline to the EpicKitchens dataset [78] to extract actions performed by human hands and arms to the robot embodiment. This process transforms human data into a format compatible with robot embodiment data. Then, this enables us to efficiently teach robot behaviors from these human videos using behavior cloning and co-training with teleoperation data. Details of this system is available on our website at https://leaphand.com

C. Simulation Tools

In robot learning applications, a simulation model of these hands is useful in a variety of applications. The kinematic model of the hand enables forward/inverse kinematics calculations for retargeting approaches such as for teleoperation from human video. It also enables geometric methods such as for grasping.

V. Analysis

A common question when starting out in robot learning is, "Which robot hand should I use?" While there are quite a few choices, many robot hands are not easy to acquire [25, 2]. There are others which are easier to acquire but still very expensive [12, 15, 24]. There are even fewer that are open source as well as being easy to produce for machine learning research. Selecting a robot hand involves more than simply considering performance metrics. **In our live demo this year, we will showcase these robot hands so that attendees can have hands on experience with all of these robot hands. This will help demystify this choice and is a key enabler of driving the open-source robot hand community.**

Hand	Grasp Type	LEAP V2	LEAP
Pickup Dice	Power	1.0	1.0
M&M	Pinch	0.4	0
Hammer	Hook	1.0	0.8
Wooden Cylinder	Cigarette	1.0	0.8
Egg	Pinch	0.8	0.8
Pringles Can	Power	1.0	1.0
Pan Handle	Hook	1.0	0.6
Bin Picking	Pinch	1.0	0.8

TABLE II: We teleoperate LEAP Hand V2 and LEAP Hand V1 on various different tasks and evaluate their success rate over 5 trials. We find that our new hand can still complete these grasps successfully.

A. Strength Test

Many soft hands, while strong in the actuated direction, often have undesirable degrees of freedom in other directions, such as twisting or compression. This makes them less suitable for a human-like robotic hand. For effective manipulation of the environment, robot hand fingers need to exhibit strength at the fingertips across all actuation directions.

In this test, we apply a force gauge to the fingers to measure how much force they can withstand without failing to hold their desired position. Some hands lack certain articulations or are too delicate to undergo the test, which is why some data points are missing.

As shown in the results in Table I, LEAP Hand V2 stands out as one of the strongest hands despite its small size and compliant nature, making it highly versatile for a wide range of tasks. The powerful motors beneath the palm enable the hand to curl with substantial strength. Meanwhile, the internal bone structure in the fingers and palm ensures rigidity, and the MCP-side motors, positioned directly on the fingers, provide efficient, high-strength power transfer.

B. Teleoperation from Motion Capture

Oftentimes we would like teleoperation to be as accurate as possible for robot hands. To do this the user wears two motion capture gloves and moves naturally to complete everyday dexterous tasks. The gloves captures accurate finger tracking to map motions naturally to the robot hands. The GELLO [71] inspired arm tracking accurately tracks the human wrist position and joint angles for the robot arm. The data collected by the system can be used to train effective policies using imitation learning, after which the bimanual robot system can perform the task autonomously. In Table II we find that LEAP Hand V2 can consistently perform well due to its small size and high levels of dexterity. Its soft construction makes is very durable to bumps and scrapes with the environment. We release videos of our results and instructions to recreate the setup on our website at https://leaphand.com **Demo attendees will have the opportunity to teleoperate a live robotic hand system and ask questions about how to build a similar setup themselves.**

VI. Conclusion

In conclusion, we introduce LEAP Hand v2, an affordable, highly dexterous, and easy-to-assemble robotic hand designed specifically for robot learning research. Priced at just $200 and with an assembly time of under one hour, it combines a hybrid rigid-soft structure for durability and strength, a universal MCP joint for exceptional dexterity, and a human-like size. Alongside the hand, we offer a suite of open-source software tools, including motion capture teleoperation, learning from human video, and reinforcement learning capabilities, all aimed at advancing machine learning applications. Through previous demonstrations at RSS 2023 and 2024, we have shown the potential of low-cost, open-source robotic hands to inspire and support the robotics community, and we look forward to showcasing these tools again at RSS 2025 to further encourage the use of open-source dexterous robot hands in various manipulation tasks.

VII. Acknowledgements

This work is supported in part by DARPA Machine Commonsense grant, AFOSR FA9550-23-1-0747, ONR N00014-22-1-2096, AIST CMU grant and Google Research Award.

References

[1] S. Goldstein, D. Princiotta, and J. A. Naglieri, "Handbook of intelligence," *Evolutionary theory, historical perspective, and current concepts*, vol. 10, pp. 978–1, 2015. 2

[2] A. Bicchi, "Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity," *IEEE TRANSACTIONS ON ROBOTICS AND AUTOMATION*, vol. 16, no. 6, 2000. 2, 8

[3] K. Shaw, A. Agarwal, and D. Pathak, "Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning," *arXiv preprint arXiv:2309.06440*, 2023. 2, 3, 7, 8

[4] S. Jacobsen, J. Wood, D. Knutti, and K. Biggers, "The utah/m.i.t. dextrous hand: Work in progress," *The International Journal of Robotics Research*, vol. 3, no. 4, pp. 21–50, 1984. [Online]. Available: https://doi.org/10.1177/027836498400300402 2

[5] M. G. Catalano, G. Grioli, E. Farnioli, A. Serio, M. Bonilla, M. Garabini, C. Piazza, M. Gabiccini, and A. Bicchi, "From soft to adaptive synergies: The pisa/iit softhand," in *Human and Robot Hands*. Springer, 2016, pp. 101–125. 2

[6] M. G. Catalano, G. Grioli, E. Farnioli, A. Serio, C. Piazza, and A. Bicchi, "Adaptive synergies for the design and control of the pisa/iit softhand," *The International Journal of Robotics Research*, vol. 33, no. 5, pp. 768–782, 2014. 2

[7] C. Della Santina, C. Piazza, G. Grioli, M. G. Catalano, and A. Bicchi, "Toward dexterous manipulation with augmented adaptive synergies: The pisa/iit softhand 2," *IEEE Transactions on Robotics*, vol. 34, no. 5, pp. 1141–1156, 2018. 2

[8] R. Deimel and O. Brock, "A novel type of compliant and underactuated robotic hand for dexterous grasping," *The International Journal of Robotics Research*, vol. 35, no. 1-3, pp. 161–185, 2016. [Online]. Available: https://doi.org/10.1177/0278364915592961 2

[9] J. Butterfaß, M. Grebenstein, H. Liu, and G. Hirzinger, "Dlr-hand ii: Next generation of a dextrous robot hand," in *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No. 01CH37164)*, vol. 1. IEEE, 2001, pp. 109–114. 2

[10] R. R. Ma, L. U. Odhner, and A. M. Dollar, "A modular, open-source 3d printed underactuated hand," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 2737–2743. 2, 4

[11] Y. Liu, H. Xiao, T. Hao, D. Pang, F. Wang, and S. Liu, "Dexterous all-soft hand (dash) with active palm: multi-functional soft hand beyond grasping," *Smart Materials and Structures*, vol. 32, no. 12, p. 125012, 2023. 2

[12] "Shadowhand," https://ninjatek.com/shop/edge/. 2, 6, 8

[13] V. Kumar, Y. Tassa, T. Erez, and E. Todorov, "Real-time behaviour synthesis for dynamic hand-manipulation," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 6808–6815. 2

[14] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas *et al.*, "Solving rubik's cube with a robot hand," *arXiv preprint arXiv:1910.07113*, 2019. 2, 3

[15] "Allegro hand," https://www.wonikrobotics.com/research-robot-hand. 2, 8

[16] D.-H. Lee, J.-H. Park, S.-W. Park, M.-H. Baeg, and J.-H. Bae, "Kitech-hand: A highly dexterous and modularized robotic hand," *IEEE/ASME Transactions on Mechatronics*, vol. 22, no. 2, pp. 876–887, 2017. 2

[17] A. Sivakumar, K. Shaw, and D. Pathak, "Robotic telekinesis: Learning a robotic hand imitator by watching humans on youtube," 2022. 2, 3, 7, 8

[18] A. Handa, K. Van Wyk, W. Yang, J. Liang, Y.-W. Chao, Q. Wan, S. Birchfield, N. Ratliff, and D. Fox, "Dexpilot: Vision-based teleoperation of dexterous robotic hand-arm system," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 9164–9170. 2, 3, 7

[19] Y. Qin, W. Yang, B. Huang, K. Van Wyk, H. Su, X. Wang, Y.-W. Chao, and D. Fox, "Anyteleop: A general vision-based dexterous robot arm-hand teleoperation system," *arXiv preprint arXiv:2307.04577*, 2023. 2

[20] K. Shaw, S. Bahl, and D. Pathak, "VideoDex: Learning Dexterity from Internet Videos," in *Conference on Robot Learning (CoRL)*, 2022. 2, 3

[21] P. Mannam, K. Shaw, D. Bauer, J. Oh, D. Pathak, and N. Pollard, "Designing anthropomorphic soft hands through interaction," in *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*. IEEE, 2023, pp. 1–8. 2, 6, 8

[22] A. Handa, A. Allshire, V. Makoviychuk, A. Petrenko, R. Singh, J. Liu, D. Makoviichuk, K. Van Wyk, A. Zhurke-

vich, B. Sundaralingam *et al.*, "Dextreme: Transfer of agile in-hand manipulation from simulation to reality," *arXiv preprint arXiv:2210.13702*, 2022. 2, 3
[23] "Psyonic ability hand," https://www.psyonic.io/ability-hand. 2
[24] "Inspire hand," https://en.inspire-robots.com/. 2, 4, 8
[25] 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)*. IEEE, 2023, pp. 1–7. 2, 3, 8
[26] "Tesla optimus," https://www.tesla.com/AI. 2
[27] "Figure.ai," https://www.figure.ai/. 2
[28] "Boston dyanmics atlas," https://bostondynamics.com/atlas/. 2
[29] "1x eve," https://www.1x.tech/. 2
[30] "Sanctuary ai," https://sanctuary.ai/. 2
[31] "Digit from agility robotics," https://agilityrobotics.com/. 2
[32] A. Agarwal, S. Uppal, K. Shaw, and D. Pathak, "Dexterous functional grasping," in *Conference on Robot Learning*. PMLR, 2023, pp. 3453–3467. 2, 3, 7
[33] T. Chen, M. Tippur, S. Wu, V. Kumar, E. Adelson, and P. Agrawal, "Visual dexterity: In-hand reorientation of novel and complex object shapes," *Science Robotics*, vol. 8, no. 84, p. eadc9244, 2023. 3
[34] R. Bhirangi, A. DeFranco, J. Adkins, C. Majidi, A. Gupta, T. Hellebrekers, and V. Kumar, "All the feels: A dexterous hand with large-area tactile sensing," *IEEE Robotics and Automation Letters*, 2023. 3
[35] "Inmoov hand," https://inmoov.fr/. 3, 8
[36] "Dexhand," https://www.dexhand.org/. 3
[37] "Pneuflex tutorial," https://www.robotics.tu-berlin.de/menue/software_tutorials/pneuflex_tutorial/. 3
[38] "Reprap open-source 3d printer," https://www.reprap.org/wiki/RepRap. 3
[39] D. Bauer, C. Bauer, A. Lakshmipathy, R. Shu, and N. S. Pollard, "Towards very low-cost iterative prototyping for fully printable dexterous soft robotic hands," in *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*. IEEE, 2022, pp. 490–497. 3, 8
[40] "Ninjatek ninjaflex edge," https://ninjatek.com/shop/edge/. 3
[41] "Recreus filaflex," https://recreus.com/gb/. 3
[42] "Colorfabb varioshore," https://colorfabb.com/varioshore-tpu-medium-brown. 3, 4
[43] Z. Li and S. S. Sastry, "Task-oriented optimal grasping by multifingered robot hands," *IEEE Journal on Robotics and Automation*, vol. 4, no. 1, pp. 32–44, 1988. 3
[44] A. M. Okamura, N. Smaby, and M. R. 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. IEEE, 2000, pp. 255–262. 3
[45] 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. 3
[46] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "Rma: Rapid motor adaptation for legged robots," *RSS*, 2021. 3
[47] T. Chen, J. Xu, and P. Agrawal, "A system for general in-hand object re-orientation," *Conference on Robot Learning*, 2021. 3
[48] W. Huang, I. Mordatch, P. Abbeel, and D. Pathak, "Generalization in dexterous manipulation via geometry-aware multi-task learning," *arXiv preprint arXiv:2111.03062*, 2021. 3
[49] A. Nair, A. Gupta, M. Dalal, and S. Levine, "Awac: Accelerating online reinforcement learning with offline datasets," *arXiv preprint arXiv:2006.09359*, 2020. 3
[50] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar, "Deep dynamics models for learning dexterous manipulation," in *Conference on Robot Learning*. PMLR, 2020, pp. 1101–1112. 3
[51] J. Romero, D. Tzionas, and M. J. Black, "Embodied hands: Modeling and capturing hands and bodies together," *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, vol. 36, no. 6, Nov. 2017. 3
[52] 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," *arXiv preprint arXiv:2108.05877*, 2021. 3
[53] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," *arXiv preprint arXiv:1709.10087*, 2017. 3
[54] P. Mandikal and K. Grauman, "Dexvip: Learning dexterous grasping with human hand pose priors from video," in *Conference on Robot Learning (CoRL)*, 2021. 3
[55] A. Kannan, K. Shaw, S. Bahl, P. Mannam, and D. Pathak, "Deft: Dexterous fine-tuning for real-world hand policies," *arXiv preprint arXiv:2310.19797*, 2023. 3
[56] S. P. Arunachalam, I. Güzey, S. Chintala, and L. Pinto, "Holo-dex: Teaching dexterity with immersive mixed reality," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5962–5969. 3
[57] J. Pari, N. M. Shafiullah, S. P. Arunachalam, and L. Pinto, "The surprising effectiveness of representation learning for visual imitation," 2021. 3
[58] S. Bahl, A. Gupta, and D. Pathak, "Human-to-robot imitation in the wild," *RSS*, 2022. 3
[59] X. Zhou, C. Majidi, and O. M. O'Reilly, "Soft hands: An analysis of some gripping mechanisms in soft robot design," *International Journal of Solids and Structures*, vol. 64, pp. 155–165, 2015. 3
[60] L. U. Odhner, L. P. Jentoft, M. R. Claffee, N. Corson,

Y. Tenzer, R. R. Ma, M. Buehler, R. Kohout, R. D. Howe, and A. M. Dollar, "A compliant, underactuated hand for robust manipulation," *The International Journal of Robotics Research*, vol. 33, no. 5, pp. 736–752, 2014. 3

[61] F. Ilievski, A. D. Mazzeo, R. F. Shepherd, X. Chen, and G. M. Whitesides, "Soft robotics for chemists," *Angewandte Chemie International Editiion*, 2011. 3

[62] C. Majidi, "Soft-matter engineering for soft robotics," *Advanced Materials Technologies*, vol. 4, no. 2, p. 1800477, 2019. 3

[63] Y.-J. Kim, J. Yoon, and Y.-W. Sim, "Fluid lubricated dexterous finger mechanism for human-like impact absorbing capability," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3971–3978, 2019. 3

[64] "Snapmaker j1s," https://snapmaker.com/. 4

[65] M. Catalano, G. Grioli, E. Farnioli, A. Serio, C. Piazza, and A. Bicchi, "Adaptive synergies for the design and control of the pisa/iit softhand," *The International Journal of Robotics Research*, vol. 33, no. 5, pp. 768–782, 2014. [Online]. Available: https://doi.org/10.1177/0278364913518998 4

[66] D.-H. Lee, J.-H. Park, S.-W. Park, M.-H. Baeg, and J.-H. Bae, "Kitech-hand: A highly dexterous and modularized robotic hand," *IEEE/ASME Transactions on Mechatronics*, vol. 22, no. 2, pp. 876–887, 2016. 5

[67] S. Sundaram, P. Kellnhofer, Y. Li, J.-Y. Zhu, A. Torralba, and W. Matusik, "Learning the signatures of the human grasp using a scalable tactile glove," *Nature*, vol. 569, no. 7758, pp. 698–702, 2019. 6

[68] B. Huang, Y. Wang, X. Yang, Y. Luo, and Y. Li, "3d-vitac: Learning fine-grained manipulation with visuo-tactile sensing," *arXiv preprint arXiv:2410.24091*, 2024. 6

[69] "Manus," https://www.manus-meta.com, note=Accessed on 2022-11-28. 6, 7

[70] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," *arXiv preprint arXiv:2304.13705*, 2023. 7

[71] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel, "Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators," *arXiv preprint arXiv:2309.13037*, 2023. 7, 8

[72] F. Pugin, P. Bucher, and P. Morel, "History of robotic surgery: from aesop® and zeus® to da vinci®," *Journal of visceral surgery*, vol. 148, no. 5, pp. e3–e8, 2011. 7

[73] A. Mandlekar, Y. Zhu, A. Garg, J. Booher, M. Spero, A. Tung, J. Gao, J. Emmons, A. Gupta, E. Orbay *et al.*, "Roboturk: A crowdsourcing platform for robotic skill learning through imitation," in *Conference on Robot Learning*. PMLR, 2018, pp. 879–893. 7

[74] A. Iyer, Z. Peng, Y. Dai, I. Guzey, S. Haldar, S. Chintala, and L. Pinto, "Open teach: A versatile teleoperation system for robotic manipulation," *arXiv preprint arXiv:2403.07870*, 2024. 7

[75] C. Wang, H. Shi, W. Wang, R. Zhang, L. Fei-Fei, and C. K. Liu, "Dexcap: Scalable and portable mocap data collection system for dexterous manipulation," *arXiv preprint arXiv:2403.07788*, 2024. 7

[76] K. Shaw, Y. Li, J. Yang, M. K. Srirama, R. Liu, H. Xiong, R. Mendonca, and D. Pathak, "Bimanual dexterity for complex tasks," *arXiv preprint arXiv:2411.13677*, 2024. 7

[77] Y. Rong, T. Shiratori, and H. Joo, "Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, October 2021, pp. 1749–1759. 8

[78] D. Damen, H. Doughty, G. M. Farinella, S. Fidler, A. Furnari, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray, "Scaling egocentric vision: The epic-kitchens dataset," in *European Conference on Computer Vision (ECCV)*, 2018. 8