

# LEAP Hand: Low-Cost, Efficient, and Anthropomorphic Hand for Robot Learning

Kenneth Shaw, Ananye Agarwal, Deepak Pathak  
Carnegie Mellon University




Fig. 1: (a) LEAP Hand is an anthropomorphic dexterous robot hand designed for robot learning research. It can be assembled in under 4 hours for 2000 USD, is composed of readily available parts, and is robust. (b) to-scale comparison of LEAP Hand and a human hand (c-h) LEAP Hand in different power and precision grasps holding common objects. The hand design and code will be open-sourced to democratize access to hardware for anthropomorphic dexterous manipulation. Videos at <https://leap-hand.github.io/>

**Abstract**—Dexterous manipulation has been a long-standing challenge in robotics. While machine learning techniques have shown some promise, results have largely been currently limited to simulation. This can be mostly attributed to the lack of suitable hardware. In this paper, we present LEAP Hand, a low-cost dexterous and anthropomorphic hand for machine learning research. In contrast to previous hands, LEAP Hand has a novel kinematic structure that allows maximal dexterity regardless of finger pose. LEAP Hand is low-cost and can be assembled in 4 hours at a cost of 2000 USD from readily available parts. It is capable of consistently exerting large torques over long durations of time. We show that LEAP Hand can be used to perform several manipulation tasks in the real world—from visual teleoperation to learning from passive video data and sim2real. LEAP Hand significantly outperforms its closest competitor Allegro Hand in all our experiments while being 1/8th of the cost. We release the URDF model, 3D CAD files, tuned simulation environment, and a development platform with useful APIs on our website at <https://leap-hand.github.io/>

## I. INTRODUCTION

Hand dexterity has been critically responsible for human cognition through active manipulation, tool use, and governing how humans learn from the world [1, 2, 3]. Replicating the dexterity of the human hand with a robot hand has been a long-standing challenge in robotics. Machine learning techniques have recently shown promise in areas such as learning from humans. However, unlike the learning successes in locomotion

[4, 5] across truly diverse terrains, robotic manipulation results in the real world have mostly been limited to one degree-of-freedom parallel jaw grippers [6, 7, 8]. In contrast, dexterous manipulation has largely been limited to simulation [9, 10] with comparatively fewer real-world results [11, 12, 13, 14, 15].

A major bottleneck in democratizing dexterous manipulation has been the hardware. Tendon-based hands like Shadow [16], while impressively capable [11], cost over 100K USD and often require significant maintenance [17] due their complicated nature. While Inmoov [18] is inexpensive and open source, it only has 5 actuators on weak tendons. Therefore, direct-driven hands have been the popular alternative for many applications [19, 20]. The Allegro Hand has been a popular direct-driven hand, but it is often unreliable, difficult to repair, and does not have an anthropomorphic kinematic structure, (see Fig 3) and is expensive at over \$16K. Please see Section II for further analysis.

As a result, only a few labs have access to hardware capable of complex dexterous tasks. This is in stark contrast to 2-finger grasping or locomotion [22, 23, 24, 25] where readily available hardware allows results to be easily reproduced and improved upon by the community. Following this analogy, good hand hardware for machine learning must be durable, repeatable, low-cost, versatile, and ideally anthropomorphic to enable easy transfer learning from humans.



Fig. 3: **Relative size of popular robot hands to scale.** *Left to right*, adult human hand, Allegro Hand [20], LEAP-C Hand, LEAP Hand, Inmoov [18], D’Manus [21]. LEAP Hand is similar in size to Allegro and  $\sim 30\%$  larger than a human hand. D’Manus is considerably larger than the rest. Because of the tendon-driven nature, Inmoov is the smallest robotic hand. The hands are accurate to scale.

We propose LEAP Hand— a dexterous, extremely *low-cost* and *robust* hand for robot learning, built from off-the-shelf or 3D-printed parts. Our hand can be assembled in under 4 hours at a cost of 2000 USD, which is 1/8th the cost of the Allegro Hand and 1/50th to that of ShadowHand. While we acknowledge this is still not affordable for all, we believe it is a step towards democratizing dexterous manipulation research. We show through a number of rigorous experiments that LEAP Hand is both robust, durable, and able to exert large torques over long periods of time. Additionally, our robot hand it is easily repairable in-house just using a standard \$250 3D printer and does not need to be sent out for repair.

Although robustness and low-cost is critical, they should not come at the cost of dexterity and anthropomorphism. We believe a good versatile hand is the one that is both *dexterous* as well as *anthropomorphic* because much of the world around us, for instance, doors, kitchens, tools, or instruments, are designed with human hands in mind, making it easier to learn by watching humans act. In LEAP Hand, we aim to maximize dexterity while being kinematically similar to a human hand. Since the ball joint at the human knuckle (Metacarpophalangeal aka MCP) cannot be replicated with direct-driven hands, it must instead be approximated using two separate motors. Prior work in direct driven hands has converged primarily to two designs, see Fig 4, one that allows abduction-adduction of fingers in open hand pose and the other that only allows with the finger flexed upwards. However, both of these lose one degree of freedom (DoF)—either in the flexed or extended position of the finger. In LEAP Hand, we propose a new kinematic mechanism to facilitate *universal abduction-adduction* for direct-driven hands that retain all degrees of freedom in all finger positions. We demonstrate that this leads to higher dexterity and for improved grasping and in-hand manipulation.

Finally, we show that LEAP Hand easily integrates with existing results in robot learning. For instance, YouTube video-based learned teleoperation and behavior cloning. In addition to the physical robot hardware, we also release an Isaac Gym-compatible simulator for the LEAP Hand and show sim2real transfer for a contact-rich task of blind in-hand rotation of a cube [26]. This shows that the hardware and simulation are

accurate and that complex tasks trained in simulation can be transferred to the real hand. We **open source** the URDF model, assembly instructions, ROS/Python API, mapping methods from human hands to LEAP Hand, and an Isaac Gym simulation environment at <https://leap-hand.github.io/>.

## II. RELATED WORK

**Robot Hands** Shadow [16] and ADROIT [27] hands paved the way to enable complex, contact-rich dexterous tasks with an anthropomorphic ball joint MCP. [11, 17]. However, they are costly (100k USD) and require constant maintenance. In contrast, the Inmoov hand is 3D printable, tendon-driven, and human-like[18]. It has only one DoF per finger and is reliant on tendon actuation which is difficult to calibrate and can be inaccurate. Bauer et. al. [28] present a soft tendon-driven hand that is very flexible for many configurations. Unfortunately, it is difficult to simulate due to its deformable nature [29, 30]. In contrast to tendon-driven hands, which have motors in the wrist, the Allegro Hand [20] has its motors in the finger joints. It is most popular in research labs [31, 32, 33, 34, 35] because it is relatively cheap (16k USD). However, users find the motors in the fingers to be weak for many everyday tasks. Moreover, the closed-source components are difficult to repair or replace. Additionally, its kinematic structure is not anthropomorphic or dexterous as we demonstrate. The ROBEL suite (which includes D’Manus) [21] is more robust, open-sourced, and easy to build. However, it has only two fingers and a thumb—making it significantly different from a human hand. Yuan et. al. [36] accomplish within-hand manipulation using rollers attached to the fingers. Humans lack this degree of freedom and manipulate objects in a very different manner. [37, 38, 39] have shown impressive results and hold a lot of promise by using fluids and linear actuators to move the fingers, but these hands are not readily available and too complicated to quickly produce, use, and maintain for robot learning.

**Rapid Manufacturing** Aluminum machining is traditionally used to create strong parts but is prohibitively difficult and expensive. Manufacturing plastic parts includes a cumbersome process of mold making, casting, curing, and support removal [40]. In contrast, additive manufacturing can be used

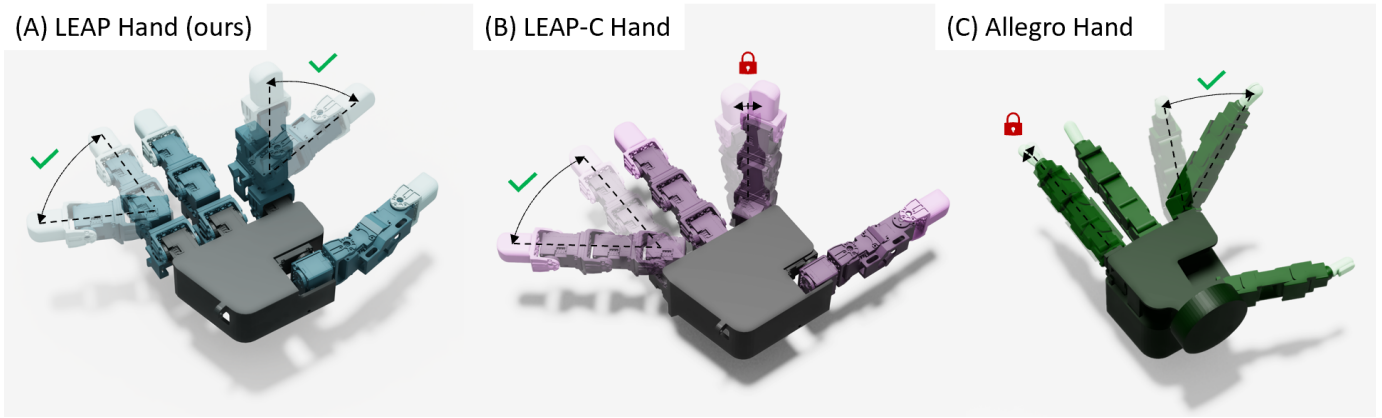


Fig. 4: Comparison of MCP joints in different robot hands and their dexterity and two different positions. (A) In LEAP-C Hand there is a large range of motion at extended but not flexed position (B) In LEAP Hand, at flexed and extended positions, the fingertip has a large range of motion. (C) In Allegro, there is a large of motion at flexed but not extended position.

to create parts very quickly for prototyping. In our paper, we leverage recent advancements in extruders, hot-ends, and motors made by the open-source Reprap Community [41]. This allows us to directly print soft flexible filaments like Ninjabflex [42]. We use this to create many parts for LEAP Hand like the hard palm and soft rubber fingertips.

**Learning Dexterity** Using a Shadow hand and Sim2real, Andrychowicz et. al. [11, 17, 43] accomplish in-hand rotation for a variety of objects. Policies that scale to thousands of objects can also be trained in simulation [9, 44]. [45] uses the D’Hand to reposition a valve. In-hand rotation of Baoding Balls using the Shadow Hand trained purely in the real-world [46], and pipe insertion using the D’Hand [47] are other notable examples of dexterous manipulation.

Several recent works focus on supervising policies of robot hands [48, 49, 50, 51]. from MANO [52] parameters which parameterize a human hand. Closely related is the teleoperation of robot hands from real-time video [31, 53], which can be used to guide learning and improve sample-efficiency [54, 55]. Hand poses can be extracted from video data available on the web to learn manipulation policies [54, 56]. Large-scale pre-training using internet videos is helpful for efficiently training robot hands for downstream tasks using a few task specific-demos [57] and also on non-dexterous manipulation [58, 59].

### III. KINEMATIC DESIGN AND ANALYSIS

The kinematic structure of a hand refers to the arrangement of its joints which determines the different poses and forces of motion it can apply. First, LEAP Hand should be as anthropomorphic as possible so that data from humans can be used to learn skills [58, 57, 54] with machine learning. This can be done using methods such as teleoperation with VR gloves or extracting keypoints from videos of human hands [53]. In addition, LEAP Hand should be dexterous for tasks such as in-hand manipulation from sim2real. In this section, we propose a robot hand design that is both anthropomorphic and dexterous.

In the human hand, there are four main degrees of freedom in each finger (Fig. 5). The knuckle or metacarpophalangeal (MCP) is a ball joint with two degrees of freedom that allows

abduction/adduction and flexion/extension. The joint closest to the knuckle is called the proximal interphalangeal (PIP) joint. The last joint, closest to the fingertip, is the distal interphalangeal (DIP). The PIP and DIP are hinge joints, each with one degree of freedom. The human hand features an opposable thumb which allows the application of force in opposition to other fingers. This enables a variety of power and precision grasps [61]. To easily map motions, a robot hand must have analogous joints to a human hand.

To replicate this structure, it is alluring to use tendons like robot hands such as the ShadowHand [16]. Such tendon-driven hands can store the large motors needed to drive them in the wrist, enabling greater flexibility in joint design and introduce ball joints. However, they are very expensive (100K USD), complicated or hard to maintain. As a result, cheaper direct-driven alternatives [20, 19] have been more popular.

#### A. Universal Abduction-Adduction Mechanism

Direct-driven hands must store the motors inside the fingers so they are limited in kinematic structure and cannot precisely imitate the human hand. Since the PIP and DIP joints are hinge joints, they are easily modeled, each with a single actuator. A ball joint cannot be modeled in this way and is typically approximated using two motors (MCP-1, MCP-2) arranged close together [19]. Prior seminal work has proposed two designs for this (Fig. 4). However, both of these designs, Allegro and LEAP-C Hand, lose one degree of freedom in either the extended or closed position. As a result, Allegro is less dexterous when extended whereas LEAP-C Hand (like C-Hand in [19]) 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 4 and the kinematic tree in the supplemental.

In LEAP Hand, we propose a new **universal abduction-adduction mechanism** for the fingers such that they can retain

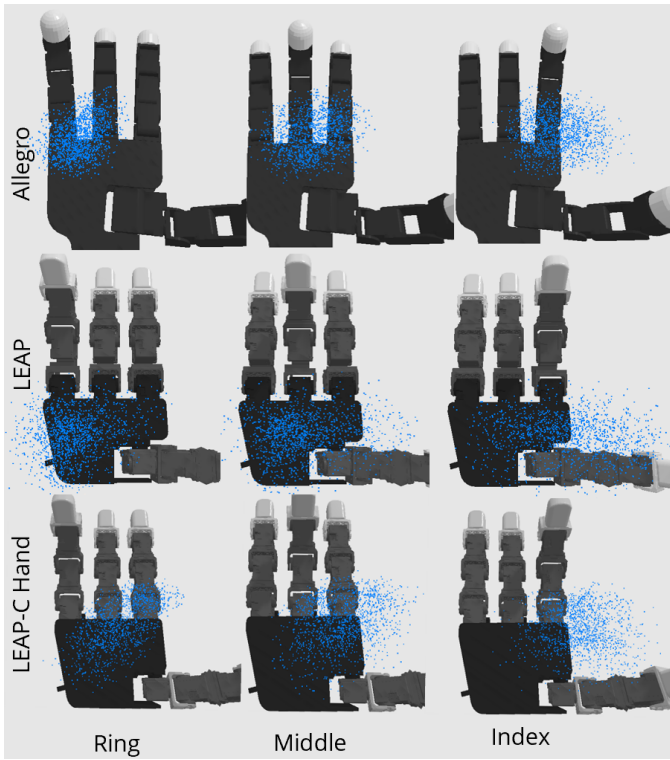


Fig. 6: We compare the possible positions of opposability of the thumb and each of the other fingers on each of the three hands. We find that LEAP Hand has the best even spread on top of the palm and a very large contact area.

all degrees of freedom at all MCP positions. Instead of the MCP-2 axis being fixed to the palm (i.e., of motor responsible for adduction-abduction), the key idea is to bring the axis to the frame of reference of the first *finger* joint and arrange it such that it is always perpendicular to it. This allows the finger to have adduction-abduction in all positions (Fig. 4). Thus, LEAP Hand has adduction-abduction in the extended position

(similar to LEAP-C Hand) as well as pronation/supination in the flexed position (similar to Allegro).

### B. Evaluating Manipulability via Thumb Opposability

Chalon et.al. [62] and Lee et.al. [19], have shown that what makes a hand more versatile is not merely the degree of abduction-adduction but also its thumb opposability volume. We test our design against Allegro and LEAP-C Hand, a baseline hand we manufacture with the same motors and parts as LEAP Hand. In Fig. 6, we plot the intersection of the thumb and finger workspaces for each hand and compute a thumb opposability metric [62]. In Table II, we show that LEAP Hand combined with the new MCP joints in the secondary fingers is better placed and is more dexterous compared to other available hands because of the increased opposable volume.

Next, manipulability measures the ease with which the fingertip can be moved in various directions at a particular joint pose. We use metrics introduced by Yoshikawa et. al. [63]. To evaluate this, many calculate the manipulability ellipsoid from the end-effector Jacobian which models the directions in which the end-effector can move. We compute the volume of this ellipsoid using the following equation:

$$w = \sqrt{\det(\mathbf{J}(\mathbf{q})\mathbf{J}(\mathbf{q})^T)}$$

where  $\mathbf{q}$  is the joint configuration and  $\mathbf{J}$  is the Jacobian of the end-effector. Note that because we are calculating volumes, a hand that can only move in one or two cartesian directions will have a volume close to zero at that pose. In three key poses, we show that LEAP Hand has consistently larger ellipsoids of greater volume for both the cartesian and angular components of the jacobian. This means that LEAP Hand has better movement at the fingertips in these few poses and a higher manipulability metric leading to more dexterity (Table II).

Finally, we show that increased dexterity leads to practical benefits as well. In the grasping test (Sec. VI-A), we find that

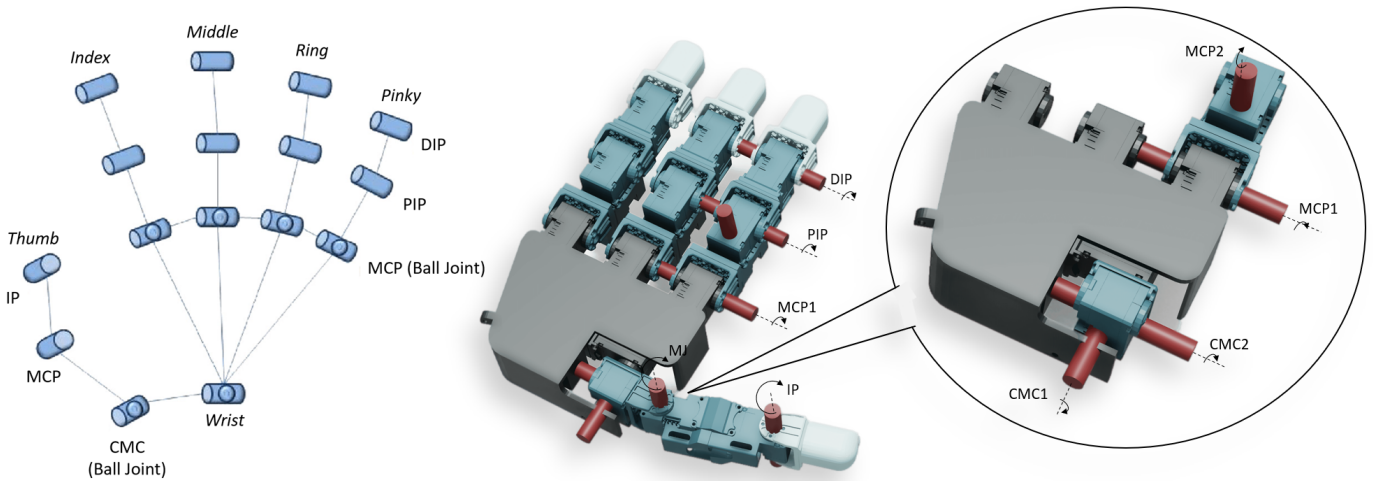


Fig. 5: The human hand kinematics above has ball joints at the MCP and CMC joints. These are difficult joints for low-cost hands to include. Left Figure from [60]. Comparison of MCP joints in different robot hands. (A) In LEAP-C Hand there is a large range of motion at extended but not flexed position (B) In LEAP Hand, at flexed and extended positions, the fingertip has a large range of motion. (C) In Allegro, there is a large range of motion at flexed but not extended position.

Robot/Position	Down (m <sup>3</sup> )	Up (m <sup>3</sup> )	Curled (m <sup>3</sup> )
<i>Allegro Hand</i>			
Linear	$8.11 \times 10^{-9}$	$3.98 \times 10^{-13}$	$2.39 \times 10^{-5}$
Angular	0	0	0
<i>LEAP-C Hand</i>			
Linear	$1.60 \times 10^{-12}$	$1.23 \times 10^{-10}$	$9.28 \times 10^{-5}$
Angular	$1.02 \times 10^{-13}$	$1.02 \times 10^{-9}$	$2.02 \times 10^{-13}$
<i>LEAP Hand (ours)</i>			
Linear	$2.02 \times 10^{-6}$	$2.42 \times 10^{-6}$	$4.51 \times 10^{-5}$
Angular	$1.20 \times 10^{-5}$	$1.20 \times 10^{-5}$	$1.20 \times 10^{-5}$

TABLE I: We show the manipulability ellipsoid volume for both the linear and angular component at three different finger positions, down, all the way up, and then halfway/curled. We find that LEAP Hand has a large manipulability ellipsoid at all three configurations.

LEAP Hand is able to grasp more objects tightly. In the blind in-hand cube rotation task (Sec. VI-D), we find that LEAP Hand is able to rotate the cube much faster than Allegro.

#### IV. HAND DESIGN PRINCIPLES

A good kinematic design must be realized effectively in hardware. In particular, the hardware should be low-cost, easy to repair, and robust.

##### A. Low-cost and Easy to Repair

In contrast to locomotion or manipulation with two-fingered grippers, real-world research in dexterous manipulation has been limited. This can be attributed in large part to the lack of suitable dexterous hand hardware. Commonly used dexterous hands such as ShadowHand [16] and AllegroHand [20] cost 100K and 16K USD, respectively, and must be sent back to manufacturers for repair in case of damage. This hardware is out-of-budget or impossible to maintain for many researchers allowing only a small fraction to work on real-world dexterous manipulation. On the other hand, due to the availability of cheap and reliable locomotion [24, 25] and manipulation [22, 23, 64] hardware, a large community of researchers is able to build off of each others’ work and drive progress.

A suitable hand should therefore be as accessible as possible. This implies that it should be low-cost and easy to repair. In LEAP Hand, we accomplish this by using as many off-the-shelf parts as possible and fabricating the rest using only a commodity 3D printer that costs around 200 USD. It can be assembled in under 4 hours.

LEAP Hand is designed to be modular. This allows key features of the robot hand to be changed, such as the length or number of fingers and the distances between each of the fingers in the palm for particular learning tasks or for analysis. Additionally, the modularity makes the hand easily repairable with only a few distinct parts.

##### B. Robustness

Learning, whether via teleoperation, behavior cloning, or reinforcement learning, on a robot hand can be notoriously harsh on hardware, especially when it is placed on a robot arm [31, 33]. Due to the movement of the arm, the hand may

Robot/Fingers	Index (mm <sup>3</sup> )	Middle (mm <sup>3</sup> )	Ring (mm <sup>3</sup> )
Allegro	409,135	348,809	204,281
LEAP-C Hand	834,516	743,764	638,605
LEAP Hand (ours)	<b>1,125,556</b>	<b>1,056,746</b>	<b>804,618</b>

TABLE II: We show the finger-to-thumb opposability volume in mm<sup>3</sup> by randomly sampling 25,000 joint configurations and finding the instances at which both fingers touch and recording that contact point. The volume of this area of contact is calculated and reported.

repeatedly collide with the table and objects it is trying to grasp. A robot hand should be robust to such treatment and continue to function reliably without breaking. In addition, a robot hand must be able to impart large torques. This is required for lifting heavy objects or using heavy tools like drills or hammers.

While 3D printing is fast and inexpensive, the resulting parts are often not strong enough. One alternative is custom metal machined parts. However, we avoid these as they add significant cost and require specialized skill and equipment to manufacture. We instead rely on inexpensive (\$10) off-the-shelf professionally extruded reinforced plastic brackets from Robotis [65] that are designed to withstand wear and tear. We only print the palm and smaller wire guide spacers using a commercial 3D printer.

The joints in LEAP Hand are designed to exceed the strength of the human hand. We choose motors geared to high torque output for their size while still being capable of a hand-like joint movement velocity of around 8 rad/sec. The amount of motor mass inside the hand is maximized compared to the size of the hand, and every other component is minimized. This enables the hand to be as strong as possible for its human-like form factor. Because these motors are so powerful, we support current- or torque-limiting them as in Section V to manipulate fragile objects and increase the durability of the hand.

##### 1) Endurance test

To test the strength of the hand over a long period of time, we hang a 2kg weight on one of the fingertips. This pose is similar to if a person was holding a half gallon of milk up with one finger without using their palm as support. We find that

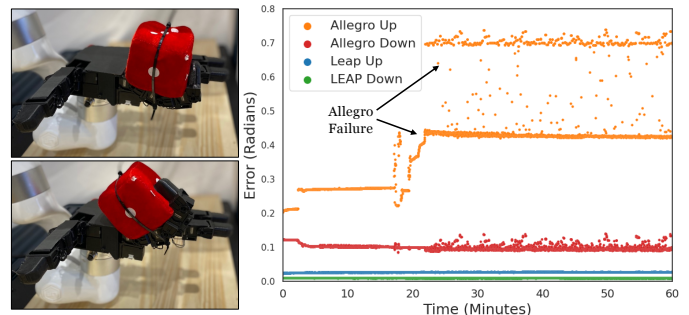


Fig. 7: **Repeatability test.** *Left:* An illustration of LEAP-C Hand performing repeating grasp-ungrasp on a small plush dice using one joint for an hour. *Right:* Comparison of LEAP-C Hand and Allegro [20]. After just 15 minutes, the Allegro Hand cannot maintain movement. In contrast, LEAP Hand continues to maintain movement with minimal joint error (< 0.05 rad). Videos at <https://leap-hand.github.io/>

Hand	Strength (N)	Power Density $N \times DOF / (cm^2)$
Bauer et. al [28]	37.4	0.677
Allegro Hand [20]	8.5	0.35
D'Manus [21]	27.8	0.313
Inmoov Hand [18]	5.8	0.116
Adult Human Hand	26.5	2.199
LEAP Hand	19.5	<b>1.045</b>
LEAP-C Hand	21.5	<b>1.15</b>

TABLE III: **Pullout Test.** A resistance comparison of each hand to pullout force which correlates to grasping strength. Power density is the total amount of motor force per square area of the hand.

LEAP Hand is able to continuously hold the grasp for an hour with only a small angle error. While the current usage gradually increases initially, it stabilizes along with the temperature of the motors, which remain cool. The current usage of the top motor reaches 250mA, which is still less than half of the maximum possible. The Allegro Hand is not powerful enough to complete this test. See Figure 8 for a plot of the results.

### 2) Repeatability test

We test the consistency and accuracy of LEAP Hand against Allegro Hand by running them continuously for 1 hour in a grasping scenario as in Figure 7. We continuously raise and lower a small (25g) plush dice strapped onto the finger by commanding one of the base finger joints up and down at 5Hz. The error of the desired joint angle compared to the actual joint angle is graphed through time.

LEAP Hand has a consistent error of 0.025 radians in the up position and 0.005 in the down position (Fig. 7), which is reasonable given the PID controller and the 750mA current limit. On the other hand, the Allegro hand starts at a much higher error. After 15 minutes, it begins to fail and then completely fails to move on one out of three grasps. This was not a failure of the position sensor in the motor. The strain of the continuous grasping on the motor caused overheating such that the motor was not able to apply required torques.

### 3) Pull-out force test

This test measures the amount of momentary outward force that can be resisted by a flexed finger from a hooked force

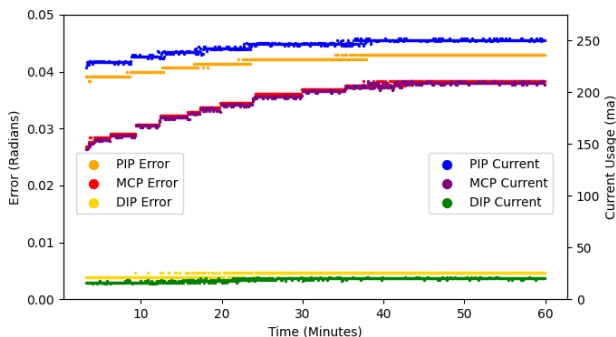


Fig. 8: **Endurance Test.** We balance a heavy 2kg weight on only one fingertip of LEAP Hand for one hour. On the left axis, we show that the angle error of commanded vs actual remains small. The right axis shows that the current use does initially increase with the temperature of the motor. However, it still withholds the weight and uses less than half of its maximum possible current of 600ma.

gauge before failure. Failure is defined as a motor or gear slipping or a finger deviating more than 15 degrees from its commanded position (Fig.9). The force returned correlates with the grip strength.

Tab. III compares force for robot and human hands. D'Manus [21] is the strongest due to its large motors. Of the anthropomorphic hands, LEAP Hand performs the best, exceeding the grip strength of a human. The Allegro Hand is weak because of smaller motors explaining why it struggles in many grasping tasks. The tendon-driven hands, Bauer et. al, and Inmoov do not perform that much better in this test even though they can store large motors in their wrists and arms. We find that these tendons often slip and cannot provide that much force at the end-effector.

## V. FABRICATION AND SOFTWARE

**Fabrication** First, each of the 3D printed components must be fabricated ( Fig. 1). A \$200 Ender 3 3D printer [66] was used with PLA plastic over a 2 day period, but any consumer-grade FDM printer will suffice. Each of the two palm pieces is printed along with fingertips and finger spacers. We collect the 3D printed parts, plastic extruded brackets, the Dynamixel motors [67], U2D2 controller, and assorted cabling. The fingers are assembled individually using brackets, 3D-printed finger spacers, and motors. The assembly process for LEAP Hand takes around 4 hours. Then each of the fingers is mounted onto the palm, and their firmware is flashed for control. The hand interfaces with the computer using a USB cable and ROS, Python, or C++. The 4-finger LEAP Hand weighs 595g and can be easily mounted to a variety of robot arms. Full video instructions of the assembly process is on our website.

**Software** A variety of control modes are supported on LEAP Hand: position control, current control, current-based position control, and velocity control. Position control enables the hand to create torques to match a desired position on the motors which is typical of many PID-based controllers. Current control mode enables a desired torque to be applied to the motors. Current-based position control enables the hand to follow position commands but also prevents large torques, which can be unsafe for the robot and the environment around it.

**Simulation** We construct a detailed 3D assembly of the hand as used in (Fig. 6) on Pybullet. This will enable anyone to 3D print and design their own version of LEAP Hand. In




Fig. 9: **Pullout Test.** A pullout force is applied and the maximum force is recorded before the hand has a 15° error or slipping.




Fig. 10: **Teleoperation and Behavior Cloning.** *Left:* We perform dexterous teleoperation using Telekinesis [53] with a single view color camera. *Right:* We perform behavior cloning from internet video and teleoperated demonstrations using Videodex [57].

addition to hardware, we release an Isaac Gym and Pybullet-based simulator for LEAP Hand. Its faithfulness to the real world is verified by performing sim2real in Sec. VI-D. We release the sim2real platform to jumpstart lab research with LEAP Hand.

## VI. LEAP HAND APPLICATIONS

First, we compare all of the hands in a grasping test with various everyday objects. Next, we compare the two most robust, human-like hands, the 4-finger LEAP Hand and the Allegro Hand [20] against each other in a variety of machine learning tasks. Teleoperation from human video demonstrates grasping capabilities and human-like form factor. Next, leveraging internet video shows the capability of learning from humans. Finally, we show LEAP Hand on in-hand manipulation via sim2real, which demonstrates that the simulation and hardware are precise. In this task, LEAP Hand is able to rotate the cube faster and is more robust to disturbances. Please see the supplemental and our [website](#) for videos of these results.

### A. Grasping Test using Teleoperation

We compare each of the hands and their ability to perform different types of grasp when holding objects. To quickly experiment and find these poses, we use the Manus Meta VR

#	Teleoperated Task	Success Rate		Completion Time (in s)	
		LEAP Hand	Allegro	LEAP Hand	Allegro
1	Pickup Dice Toy	<b>1.0</b>	0.9	<b>6.5 (1.7)</b>	8.6 (2.65)
2	Pickup Dino Doll	<b>1.0</b>	0.9	<b>6.0 (1.5)</b>	8.2 (3.49)
3	Box Rotation	<b>0.7</b>	0.6	<b>28.2 (15.7)</b>	37.2 (12.6)
4	Scissor Pickup	0.6	<b>0.7</b>	32.4 (7.8)	<b>28.6 (9.4)</b>
5	Cup Stack	<b>0.8</b>	0.6	<b>15.4 (7.0)</b>	21.5 (7.6)
6	Two Cup Stacking	<b>0.6</b>	0.3	<b>18.2 (9.2)</b>	27.3 (11.0)
7	Pour Cubes in Plate	<b>0.8</b>	0.7	<b>30.2 (15.2)</b>	36.8 (17.7)
8	Cup Into Plate	<b>0.8</b>	<b>0.8</b>	<b>6.2 (2.5)</b>	10.6 (4.4)
9	Open Drawer	<b>0.9</b>	<b>0.9</b>	<b>18.2 (11.2)</b>	23.6 (12.3)
10	Open Drawer & Pick	<b>0.7</b>	0.6	37.2 (10.2)	<b>33.7 (8.1)</b>
<b>Outperform rate</b>		<b>9/10</b>	3/10	<b>8/10</b>	2/10

TABLE IV: **Teleoperation—comparing LEAP Hand and Allegro.** Success rate and average completion time of a trained operator completing a variety of teleoperated tasks. LEAP Hand outperforms or matches the Allegro performance on 9/10 tasks.

Object	Grasp Type [61]	LEAP	LEAP-C	Allegro	D’Manus	Inmoov
<i>Power:</i>						
Mustard	Med. Palm+Pad	20	20	13	8	Y
Toy Kick Ball	Lrg. Palm+Pad	20	20	9	20	N
Golf Ball	Small Pad	16	20	7	0	Y
Softball	Large Pad	20	20	10	15	N
Drill	Trigger Press	20	20	15	0	N
Pringle Can	Power Palm	19	20	20	0	Y
Pan (from rim)	Disk Grasp	20	20	20	14	N
<i>Intermediate:</i>						
Chopsticks	Tripod Grasp	16	13	0	0	N
Wood Cylinder	Cigarette Grasp	4	5	0	0	N
<i>Precision:</i>						
1” Cube	2 Finger Precision	20	20	20	0	N
M&M	Tip Pinch Grasp	Y	Y	Y	N	N
Wine Glass	Flat Hand Cupping	20	20	4	0	N
Credit Card	Lateral Pinch	20	20	8	0	N

TABLE V: We test each robot hand on a variety of objects and grasps types and see how much perturbation force they can resist (in newtons). The dexterous morphology of LEAP Hand as well as its strong motors enables the cigarette and flat-hand cupping grasps.

glove [68] to accurately teleoperate the first three hands (see appendix for details). Since D’Manus is not anthropomorphic enough to teleoperate from human motion, we manually control keyframes for it. We show various types of grasps that each hand can perform, and the amount of perturbation force they can resist (up to 20N). Once the object is grasped we push on it with the force gauge until it slips or the force gauge crosses 20N. Because InMoov is too fragile to teleoperate and apply perturbation forces to, we only test if it can grasp the object securely and report these results in the table.

Table V shows LEAP Hand can grasp all objects and can perform both many power and precision grasps. While LEAP Hand and LEAP-C Hand perform similarly, the latter has weaker grasps because its MCP side motors cannot be used to adjust the grasp. Allegro’s motors are significantly weaker which leads to objects like the golf ball or the soccer ball to easily slip out. Additionally, because its kinematics lacks adduction/abduction in an extended position, it cannot perform the tripod grasp for the chopsticks, the cigarette grasp for the wooden cylinder, or the flat hand cupping grasp for the wine glass properly. D’Manus can complete extremely strong power grasps on larger objects, but its lack of opposability and inability to provide resistive force on all 4 sides of the objects hurts its performance. Due to its large size, it fails to grasp smaller objects.

### B. Teleoperation from Uncalibrated Human Video

Teleoperation enables control of high DOF robots in real-time via human feedback. This is also a useful method for collection demonstrations. Because the Allegro Hand’s morphology does not have a human-like MCP joint we must borrow the human-to-robot re-targeting method from Robotic Telekinesis [53] (Fig. 10 (left)), that manually defines key vectors between palms and fingertips on both robot  $v_i^r$  and human hand  $v_i^h$ . These vectors define an energy function  $E_\pi$  which minimizes the distance between human hand poses (parameterized by the

	Pick		Rotate		Open		Cover		Uncover		Place		Push		Overall
	train	test	train	test	train	test	train	test	train	test	train	test	train	test	
Allegro Hand	0.81	0.75	<b>0.89</b>	0.69	0.90	<b>0.80</b>	0.78	0.67	<b>1.00</b>	<b>0.90</b>	0.90	0.70	<b>1.00</b>	<b>1.00</b>	6/14
LEAP Hand	<b>0.92</b>	<b>0.84</b>	<b>0.89</b>	<b>0.72</b>	<b>0.94</b>	0.76	<b>0.80</b>	<b>0.75</b>	0.96	<b>0.90</b>	<b>0.94</b>	<b>0.75</b>	<b>1.00</b>	<b>1.00</b>	<b>12/14</b>

TABLE VI: **Learning from videos via VideoDex** [57]. Hand policies are pretrained on internet videos of humans and finetuned using minimal ( $\approx 100$ ) teleoperated demos. On this practical use case, LEAP Hand performs better on 12 of 14 {task} $\times$ {train, test} pairs.

tuple  $(\beta_h, \theta_h)$ ) and the robot hand poses  $q$  scaled by  $c_i$ :

$$E_\pi((\beta_h, \theta_h), q) = \sum_{i=1}^{10} \|v_i^h - (c_i \cdot v_i^r)\|_2^2 \quad (1)$$

[53] trains an MLP  $H_R(\cdot)$  to implicitly minimize the energy function described in Equation 1.

Because LEAP Hand includes similar joints to a human, we can directly map joint angles between the human and robot. In table IV, we observe that LEAP Hand performs better than Allegro Hand on 9/10 of these teleoperated tasks. LEAP Hand is easier to control due to its better morphology, accuracy, and responsiveness to hand input. As users in [53] mention, it is difficult to teleoperate a robot hand with an energy function. Additionally, the better opposability and strength of LEAP Hand allows the operator to reliably grasp objects that are difficult to grasp on the Allegro Hand. While Allegro Hand needed to take breaks to avoid overheating like in [14], LEAP Hand kept running without a degradation of performance.

### C. Behavior Cloning from Demonstrations

Behavior cloning enables agents to learn a policy for a particular task given demonstrations. However, collecting demonstrations for behavior cloning is expensive. We utilize video from Epic-Kitchens [69] as pre-training for our policy using VideoDex [57] along with NDP [70], see Fig. 10 (right). We also only use demonstrations collected from prior work [57] on Allegro Hand and map those to LEAP Hand. LEAP Hand still outperforms the Allegro Hand in task performance as in Table VI. This is because of its consistency and strength while completing these tasks.

### D. Sim2Real In-Hand Manipulation

We perform in-hand rotation of a cube along an axis perpendicular to the palm. LEAP Hand can return current




Fig. 11: **Sim2Real transfer**. *Left*: Simulated LEAP Hand in Isaac Gym [71] completing an in-Hand manipulation task. *Right*: LEAP Hand completing the same task in the real world. Please see our website <https://leap-hand.github.io/> for videos.

joint position, velocity, and torque and can be controlled from both torque or position commands. In this case, the robot infers the object pose through the history of observed joint angles alone. This is a challenging task since it is contact-rich, and the policy cannot directly observe the pose of the cube. The policy receives joint angles (16 values) from the motors and outputs the target joint angles (16) at 20 Hz which is passed as position commands to the motors.

We choose a GRU [72] architecture for our policy. We first generate a cache of stable grasps similar to [73]. The policy is then rewarded for turning the cube  $r_{\text{rot}} = \text{clip}(\omega_z, -0.25, 0.25)$ , where  $\omega_z$  is the angular velocity along the vertical axis. We add additional penalties for deviation from the stable grasp pose, mechanical work done, motor torques, and object linear velocity. The scale for the rotation reward is 1.25. The scales for the penalties are -0.1, -1, -0.1, -0.3 respectively. We train using PPO [74] with backpropagation through time [75] in IsaacGym [71].

In simulation, we compare the average angular velocity of a cube for different hands and find that LEAP Hand leads to faster rotations (Tab. VII) than Allegro. This is because the joint structure of LEAP Hand allows it to support the cube from the sides, whereas since Allegro does not have adduction/abduction it must let go of the cube periodically in order to re-orient it.

## VII. CONCLUSION AND FUTURE WORK

We introduce LEAP Hand and its core design principles. Following these principles, we demonstrate that LEAP Hand can perform exceedingly well compared to other hands on the market in strength, grasping, and durability. We show its usefulness in a variety of real-world tasks, including teleoperation, behavior cloning, and sim2real. We **open source** the URDF model, 3D CAD files, and a development platform with useful APIs. In future work, we plan to develop and integrate LEAP Hand with low-cost touch sensors. We believe that touch feedback will be increasingly important for many machine learning tasks.

Hand	Angular velocity (rad/s)
Allegro	0.0828
LEAP-C Hand	0.2205
LEAP Hand (ours)	<b>0.2288</b>

TABLE VII: Comparison of angular velocity for the blind in-hand rotation of a cube in simulation. Since the Allegro Hand lacks abduction/adduction at its extended position, it has low angular velocity. However, LEAP Hand and LEAP-C Hand have this ability and have better performance.



**Acknowledgement** We thank Shikhar Bahl, Russell Mendonca, Unnat Jain and Jianren Wang for fruitful discussions about the project. KS is supported by NSF Graduate Research Fellowship under Grant No. DGE2140739. This work is supported by ONR N00014-22-1-2096 and the DARPA Machine Common Sense grant.

#### REFERENCES

- [1] K. Libertus, A. S. Joh, and A. W. Needham, “Motor training at 3 months affects object exploration 12 months later,” *Developmental Science*, vol. 19, no. 6, pp. 1058–1066, 2016. 1
- [2] T. Bruce, *Learning through play, for babies, toddlers and young children*. Hachette UK, 2012. 1
- [3] E. J. Gibson, “Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge,” *Annual review of psychology*, vol. 39, no. 1, pp. 1–42, 1988. 1
- [4] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, “Learning robust perceptive locomotion for quadrupedal robots in the wild,” *Science Robotics*, vol. 7, no. 62, p. eabk2822, 2022. 1
- [5] A. Agarwal, A. Kumar, J. Malik, and D. Pathak, “Legged locomotion in challenging terrains using egocentric vision,” *CoRL*, 2022. 1
- [6] Bhardwaj, Mohak and Sundaralingam, Balakumar and Mousavian, Arsalan and Ratliff, Nathan and Fox, Dieter and Ramos, Fabio and Boots, Byron, “STORM: An Integrated Framework for Fast Joint-Space Model-Predictive Control for Reactive Manipulation,” in *Conference on Robot Learning (CoRL)*, 2021. 1
- [7] S. Dasari, J. Wang, J. Hong, S. Bahl, Y. Lin, A. Wang, A. Thankaraj, K. Chahal, B. Calli, S. Gupta, D. Held, L. Pinto, D. Pathak, V. Kumar, and A. Gupta, “Rb2: Robotic manipulation benchmarking with a twist,” in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, J. Vanschoren and S. Yeung, Eds., vol. 1, 2021. [Online]. Available: <https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/file/3988c7f88ebcb58c6ce932b957b6f332-Paper-round2.pdf> 1
- [8] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, “Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 438–13 444. 1
- [9] T. Chen, J. Xu, and P. Agrawal, “A system for general in-hand object re-orientation,” *Conference on Robot Learning*, 2021. 1, 3
- [10] 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. 1
- [11] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray *et al.*, “Learning dexterous in-hand manipulation,” *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020. 1, 2, 3
- [12] L. Sievers, J. Pitz, and B. Bäuml, “Learning purely tactile in-hand manipulation with a torque-controlled hand,” *arXiv preprint arXiv:2204.03698*, 2022. 1
- [13] A. S. Morgan, K. Hang, B. Wen, K. Bekris, and A. M. Dollar, “Complex in-hand manipulation via compliance-enabled finger gaing and multi-modal planning,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4821–4828, 2022. 1
- [14] A. Handa, A. Allshire, V. Makovychuk, A. Petrenko, R. Singh, J. Liu, D. Makoviichuk, K. Van Wyk, A. Zhurkevich, B. Sundaralingam *et al.*, “Dextreme: Transfer of agile in-hand manipulation from simulation to reality,” *arXiv preprint arXiv:2210.13702*, 2022. 1, 8
- [15] T. Chen, M. Tippur, S. Wu, V. Kumar, E. Adelson, and P. Agrawal, “Visual dexterity: In-hand dexterous manipulation from depth,” *arXiv preprint arXiv:2211.11744*, 2022. 1
- [16] “Shadowhand,” <https://ninjatek.com/shop/edge/>. 1, 2, 3, 5
- [17] 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. 1, 2, 3
- [18] “Inmoov hand,” <https://inmoov.fr/>. 1, 2, 6
- [19] 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. 1, 3, 4
- [20] “Allegro hand,” <https://www.wonikrobotics.com/research-robot-hand>. 1, 2, 3, 5, 6, 7
- [21] R. Bhirangi, A. DeFranco, J. Adkins, C. Majidi, A. Gupta, T. Hellebrekers, and V. Kumar, “All the feels: A dexterous hand with large area sensing,” *arXiv preprint arXiv:2210.15658*, 2022. 2, 6
- [22] “xarm6 by ufactory,” <https://www.ufactory.cc/xarm-collaborative-robot>. 1, 5
- [23] “Franka panda,” <https://www.franka.de/>. 1, 5
- [24] “Unitree a1,” <https://www.unitree.com/en/a1>. 1, 5
- [25] “Unitree go1,” <https://www.unitree.com/en/go1>. 1, 5
- [26] R. R. Ma and A. M. Dollar, “On dexterity and dexterous manipulation,” in *2011 15th International Conference on Advanced Robotics (ICAR)*. IEEE, 2011, pp. 1–7. 2
- [27] 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
- [28] D. Bauer, C. Bauer, A. Lakshminpathy, 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. 2, 6
- [29] F. Faure, C. Duriez, H. Delingette, J. Allard, B. Gilles, S. Marchesseau, H. Talbot, H. Courtecuisse, G. Bousquet, I. Peterlik, and S. Cotin, *SOFA: A Multi-Model Framework for Interactive Physical Simulation*. Berlin, Heidelberg:

- Springer Berlin Heidelberg, 2012, pp. 283–321. [Online]. Available: [https://doi.org/10.1007/8415\\_2012\\_125](https://doi.org/10.1007/8415_2012_125) 2
- [30] C. Duriez and T. Bieze, “Soft robot modeling, simulation and control in real-time,” in *Soft Robotics: Trends, Applications and Challenges: Proceedings of the Soft Robotics Week, April 25-30, 2016, Livorno, Italy*. Springer, 2017, pp. 103–109. 2
- [31] 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, 5
- [32] B. Sundaralingam and T. Hermans, “Relaxed-rigidity constraints: kinematic trajectory optimization and collision avoidance for in-grasp manipulation,” *Autonomous Robots*, vol. 43, no. 2, pp. 469–483, 2019. 2
- [33] A. Sivakumar, K. Shaw, and D. Pathak, “Robotic telekinesis: learning a robotic hand imitator by watching humans on youtube,” *RSS*, 2022. 2, 5
- [34] S. P. Arunachalam, S. Silwal, B. Evans, and L. Pinto, “Dexterous imitation made easy: A learning-based framework for efficient dexterous manipulation,” *arXiv preprint arXiv:2203.13251*, 2022. 2
- [35] F. O. H. to Multiple Hands: Imitation Learning for Dexterous Manipulation from Single-Camera Teleoperation, “Qin, yuzhe and su, hao and wang, xiaolong,” 2022. 2
- [36] S. Yuan, A. D. Epps, J. B. Nowak, and J. K. Salisbury, “Design of a roller-based dexterous hand for object grasping and within-hand manipulation,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 8870–8876. 2
- [37] “Clone robotics,” <https://www.clonerobotics.com/>. 2
- [38] 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. 2
- [39] U. Kim, D. Jung, H. Jeong, J. Park, H.-M. Jung, J. Cheong, H. R. Choi, H. Do, and C. Park, “Integrated linkage-driven dexterous anthropomorphic robotic hand,” *Nature communications*, vol. 12, no. 1, pp. 1–13, 2021. 2
- [40] “Pneuflex tutorial,” [https://www.robotics.tu-berlin.de/menue/software\\_tutorials/pneuflex\\_tutorial/](https://www.robotics.tu-berlin.de/menue/software_tutorials/pneuflex_tutorial/). 2
- [41] “Reprap open-source 3d printer,” <https://www.reprap.org/wiki/RepRap>. 3
- [42] “Ninjatek ninjaflex edge,” <https://ninjatek.com/shop/edge/>. 3
- [43] A. Kumar, Z. Fu, D. Pathak, and J. Malik, “Rma: Rapid motor adaptation for legged robots,” *RSS*, 2021. 3
- [44] 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
- [45] A. Nair, A. Gupta, M. Dalal, and S. Levine, “Awac: Accelerating online reinforcement learning with offline datasets,” *arXiv preprint arXiv:2006.09359*, 2020. 3
- [46] 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
- [47] A. Gupta, J. Yu, T. Z. Zhao, V. Kumar, A. Rovinsky, K. Xu, T. Devlin, and S. Levine, “Reset-free reinforcement learning via multi-task learning: Learning dexterous manipulation behaviors without human intervention,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6664–6671. 3
- [48] J. Wang, F. Mueller, F. Bernard, S. Sorli, O. Sotnychenko, N. Qian, M. A. Otaduy, D. Casas, and C. Theobalt, “Rgb2hands: real-time tracking of 3d hand interactions from monocular rgb video,” *ACM Transactions on Graphics (TOG)*, vol. 39, no. 6, pp. 1–16, 2020. 3
- [49] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik, “End-to-end recovery of human shape and pose,” *CoRR*, vol. abs/1712.06584, 2017. [Online]. Available: <http://arxiv.org/abs/1712.06584> 3
- [50] Y. Feng, V. Choutas, T. Bolkart, D. Tzionas, and M. J. Black, “Collaborative regression of expressive bodies using moderation,” *arXiv preprint arXiv:2105.05301*, 2021. 3
- [51] 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. 3
- [52] 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
- [53] A. Sivakumar, K. Shaw, and D. Pathak, “Robotic telekinesis: Learning a robotic hand imitator by watching humans on youtube,” 2022. 3, 7, 8
- [54] 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
- [55] 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
- [56] P. Mandikal and K. Grauman, “Dexvip: Learning dexterous grasping with human hand pose priors from video,” in *Conference on Robot Learning (CoRL)*, 2021. 3
- [57] K. Shaw, S. Bahl, and D. Pathak, “VideoDex: Learning Dexterity from Internet Videos,” in *Conference on Robot Learning (CoRL)*, 2022. 3, 7, 8
- [58] S. Bahl, A. Gupta, and D. Pathak, “Human-to-robot imitation in the wild,” *RSS*, 2022. 3
- [59] J. Pari, N. M. Shafiullah, S. P. Arunachalam, and L. Pinto, “The surprising effectiveness of representation learning for visual imitation,” 2021. 3
- [60] I. Cerulo, F. Ficuciello, V. Lippiello, and B. Siciliano,

- “Teleoperation of the schunk s5fh under-actuated anthropomorphic hand using human hand motion tracking,” *Robotics and Autonomous Systems*, vol. 89, pp. 75–84, 2017. 4
- [61] J. Liu, F. Feng, Y. C. Nakamura, and N. S. Pollard, “A taxonomy of everyday grasps in action,” in *2014 IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2014, pp. 573–580. 3, 7
- [62] M. Chalon, M. Grebenstein, T. Wimböck, and G. Hirzinger, “The thumb: Guidelines for a robotic design,” in *2010 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2010, pp. 5886–5893. 4
- [63] T. Yoshikawa, “Manipulability of robotic mechanisms,” *The international journal of Robotics Research*, vol. 4, no. 2, pp. 3–9, 1985. 4
- [64] “Ur5,” <https://www.universal-robots.com/products/ur5-robot/>. 5
- [65] “Robotis dynamixels,” <https://www.robotis.us/>. 5
- [66] “Crealitty ender 5 3d printer,” <https://www.crealitty.com/>. 6
- [67] “Robotis dynamixel,” <https://www.robotis.us/dynamixel-xc330-m288-t/>. 6
- [68] “Manus,” <https://www.manus-meta.com>, note=Accessed on 2022-11-28. 7
- [69] 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
- [70] S. Bahl, M. Mukadam, A. Gupta, and D. Pathak, “Neural dynamic policies for end-to-end sensorimotor learning,” in *NeurIPS*, 2020. 8
- [71] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa *et al.*, “Isaac gym: High performance gpu-based physics simulation for robot learning,” *arXiv preprint arXiv:2108.10470*, 2021. 8
- [72] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, “Empirical evaluation of gated recurrent neural networks on sequence modeling,” *arXiv preprint arXiv:1412.3555*, 2014. 8
- [73] H. Qi, A. Kumar, R. Calandra, Y. Ma, and J. Malik, “In-Hand Object Rotation via Rapid Motor Adaptation,” in *Conference on Robot Learning (CoRL)*, 2022. 8
- [74] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017. 8
- [75] P. J. Werbos, “Backpropagation through time: what it does and how to do it,” *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550–1560, 1990. 8