Robotics: Science and Systems 2024
Delft, Netherlands, July 15-July 19, 2024

Demonstrating Learning from Humans on Open-Source Dexterous Robot Hands

Kenneth Shaw, Ananye Agarwal, Shikhar Bahl, Mohan Kumar Srirama, Alexandre Kirchmeyer,
Aditya Kannan, Aravind Sivakumar, Deepak Pathak
Carnegie Mellon University

Fig. 1: Our demonstration will show our three different open-source robot hands that each cost around $2000 and are extremely easy to fabricate and 3D print. First, LEAP Hand (RSS 2023) is an easy-to-use robot hand that will be demonstrated doing sim2real in-hand reorientation. Second, DASH hand is a soft, compliant and durable robot hand that will be shown doing real-world learning and teleoperation from human video. Finally, LEAP Hand v2 is our newest robot hand that has a human-like size and incredible strength. It will be shown using our new MoCap teleoperation system. The demo will engage with the attendees using real robot hands and will serve to demystify dexterous manipulation hardware.

***Abstract*—Emulating human-like dexterity with robotic hands has been a long-standing challenge in robotics. In recent years, machine learning has demanded robot hands to be reliable, inexpensive and easy-to-reproduce. For the past few years we have been investigating how to address these demands. [1, 2, 3, 4, 5, 6] We will demonstrate our three robot hands that address this problem ranging from rigid easy-to-simulate hand to soft but strong dexterous robot hands performing three different machine learning tasks. Our first machine learning task will be teleoperation, where we will develop a new mobile arm and hand motion capture system that we will bring to RSS 2024. Second, we will demonstrate how to use human-video and human motion to teach robot hands. Finally, we will show how to continually improve these policies using reinforcement learning in both simulation and the real-world. This demo will be engaging, will serve to demystify dexterous manipulation and inspire researchers to bring robot hands into their own projects. Please see our website at https://leaphand.com/rss2024demo for more interactive information.**

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 sensory abilities. This extraordinary sensing and adaptability are orchestrated by the impressive capabilities of our brains. The development of our brains is often linked to the necessity of manipulating our surroundings with our hands [7].

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. The absence of robot humanoids with efficient 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 such as are too large, some have fewer degrees of freedom and other are extremely difficult to produce and maintain. We believe this isn't an inherent flaw in robot hands but rather a consequence of not designing them correctly. To break through this prevailing belief, we will show demonstrative proof with three of our robot hands. **We will have hands-on real robot demos with at least 6 robot hands and 2 robot arms:**

1) **LEAP Hand:** The most popular open-source robot hand for academia with over 50 being actively used only 6 months after release. At RSS 2023, we provided demos of both our Mocap and sim2real pipelines with great audience participation.
2) **DASH Hand:** A completely soft, small tendon-driven robot hand with 3 DOF per finger.
3) **LEAP Hand v2:** Our smallest and strongest open-source robot hand that is still easy to assemble. It has 3D printed soft fingers and an articulated palm both with hard plastic skeletal structures to unlock great performance.

While having strong hardware is an important foundation, it is only a part of dexterous manipulation. To perform robot learning, we must leverage video data and demonstrations from kinematically similar human hands. **The robot hands will be mixed together to demonstrate these three key robot learning paradigms:**

1) **Motion Capture Teleoperation:** Significant advances in motion-capture such as the Manus Meta Glove enable accurate teleoperation on robot hands and arms. **Attendees will be able to teleoperate all three robot hands using our newly developed open-source mocap system.**
2) **Learning from Human Video:** Easily collectible data can be obtained through human video, such as from the internet. We develop how to utilize this human hand data to train robot hands. **Real robot hands will be running autonomous policy rollouts and videos will play of how to convert human video to pseudo-robot experience.**
3) **Real-world and simulation-based RL:** We enable robot hands to continually learn and iteratively improve from their training data through real-world experience and simulation. **Similar to RSS 2023, but at larger scale, sim2real policies will be doing various behaviors such as in-hand reorientation.**

These are hands-on demos with almost all real robot hands. For methods that are difficult to demo, such as real-world learning or retargeting from human video, a few accompanying videos will be shown during the 5 min presentation and also during the demo duration.

This demo serves to guide robotics researchers and demystify the field of dexterous manipulation. The demo will be engaging to attendees with many strong hands-on examples of teleoperation for behavior cloning and sim2real. This is enabled by our novel class of robot hands, which are significantly more dexterous, open-source, stronger, cost-effective (each all around $2000 and easy-assembly) and more user-friendly than existing counterparts. **We hope this demo will inspire attendees to bring open-source dexterous robot hands into their own manipulation projects.** To further support this mission, we will bring 3D printed parts kits that some RSS 2024 attendees can bring home to help kick-start their research using these open-source robot hands.

Both the hardware instructions and the learning methods demonstrated are already open-sourced or will be by RSS 2024. Some of these efforts are based off of many papers that the authors have published in the past, LEAP Hand [4], DASH Hand [5] for hardware and Robotic Telekinesis, [8], Videodex [2] and DEFT [3] in learning from human videos and Dexterous Functional Grasping [6] for sim2real. However, this demo will also have elements that are unique and developed for RSS 2024. This includes our bimanual teleop system never seen before in our published work.

II. Related Work

Robot Hands Over the course of time, numerous robotic hands have been developed to mimic the capabilities of the human hand, with varying degrees of success and accessibility. Notably, the Shadow hand, as documented in [9, 10], has demonstrated remarkable achievements such as in-hand reorientation of a Rubik's cube [11]. Despite its impressive performance, the Shadow hand is widely acknowledged for its high cost (approximately $150k) and challenging usability.

Conversely, the Allegro Hand [12, 13], 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 [1, 14, 15, 2, 5], as well as in-hand reorientation [16]. 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) [17].

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 [4, 6] is easy to use and has been used by many research labs around the world. The Robel suite, exemplified by D'Manus, offers large yet durable hands employed in tasks such as reorientation [18] and grasping [19]. Other hands, such as Inmoov [20] and DexHand [21], cater to hobbyists but may be limited by inexpensive motors or fragile plastic components.

Some advanced robotic hands, while challenging to manufacture and acquire, showcase remarkable results in their respective laboratories. The MIT/Utah Hand has an early tendon-driven design [22]. [23, 24, 25, 26, 27] developed this area of tendon-driven hands. The dexterous all-soft hand, with palm articulations in a completely soft structure, is presented in [28].

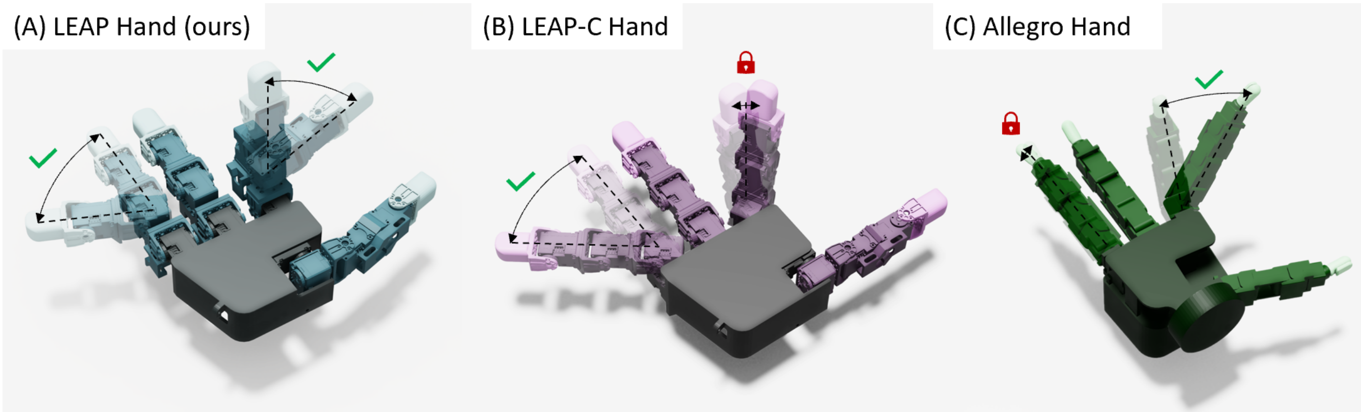

Fig. 3: In LEAP Hand, we introduce a MCP joint that allows for abduction and adduction in both flexed and extended positions. This allows for a large range of motion of the fingertip. In a conventional hand, LEAP-C Hand, the finger can move side to side in the open-palm, but in the flexed position it only spins in place. In Allegro, there is a large of motion at flexed but not in the extended position. [4]

Additionally, the Faive Hand [29] demonstrates noteworthy sim2real results in in-hand reorientation.

A resurgence of interest in humanoid robotics from industry players like Tesla Optimus [30], Figure [31], BD Atlas [32], 1x [33], Sanctuary AI [34], and Digit [35] 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.

Rapid Manufacturing The conventional approach to fabricating robust components involves machining, such as with aluminum which incurs high costs. The production of plastic parts, historically characterized by a sequence involving mold creation, casting, curing, and support removal [36], is suitable for large-scale production. Conversely, the advent of 3D printing has revolutionized the landscape of small-scale manufacturing [37], facilitating the autonomous, rapid printing of individual parts automatically.

The 3D printing field has seen significant material advancements, with materials like TPU/TPE by Ninjatek and Filaflex, offering new flexibility [38, 39]. Foaming materials such as Colorfabb Varioshore enable the adjustment of material properties through flow-rate modulation. Additionally, materials like Nylon and carbon-fiber, when utilized in 3D printing, provide noteworthy strength and durability. Consumer-friendly multimaterial 3D printers have become affordable and accessible.

Learning for Dexterous Manipulation In robot learning Andrychowicz et al. achieved in-hand rotation for various objects using a Shadow hand and Sim2real techniques. [40, 11]. Simulation-based training that scales to thousands of objects is explored in works such as [41, 42, 43, 6] which shows promise in robot learning. D'Hand is utilized by Nair et al. to reposition a valve [44]. Other notable instances of dexterous manipulation include Baoding Balls' in-hand rotation using the Shadow Hand trained exclusively in the real world [45].

Recent research emphasizes supervising policies for robot hands based on human actions such as from MANO [46] parameters of the human hand. Related work involves the teleoperation of robot hands from real-time video [14, 1] and can offer guidance for learning [2, 47, 48]. Hand poses extracted from online video data are leveraged for learning manipulation policies [47, 49]. Large-scale pre-training using internet videos proves beneficial for efficiently training robot hands for downstream tasks with a few task-specific demos [2, 3, 50] and extends to non-dexterous manipulation scenarios [51, 52].

III. Our Robot Hand Hardware

We will bring at least 6 examples of 3 different low-cost, fully open-sourced, dexterous anthropomorphic robot hands to RSS 2024 for our demonstration. In this section we describe all of our low-cost ($2000), strong, easy to assemble and repair, open-source anthropomorphic dexterous robot hands that attendees will be able to interact with.

A. LEAP Hand

LEAP Hand is an affordable, dexterous, and anthropomorphic hand designed for machine learning research which was presented at RSS 2023. Setting it apart from previous models, LEAP Hand incorporates a groundbreaking kinematic structure, ensuring maximum dexterity irrespective of finger pose. With a low assembly cost of 2000 USD and a four-hour assembly time using easily accessible components, LEAP Hand consistently performs over extended durations. Notably, LEAP Hand outperforms its closest competitor, the Allegro Hand, in all conducted experiments, while being only 1/8th of the cost. We have made detailed assembly instructions, the Sim2Real pipeline, and a development platform with valuable APIs accessible on our website at http://leaphand.com .

1) Kinematic Structure

Hands that are directly driven traditionally face limitations in kinematic structure due to the need to house motors within fingers, preventing precise imitation of the human hand. Hinge joints, like the PIP and DIP joints, can be easily modeled with a single actuator each. However, ball joints present a challenge and are typically approximated using two motors (MCP-1, MCP-2) positioned closely together, as noted in [53].

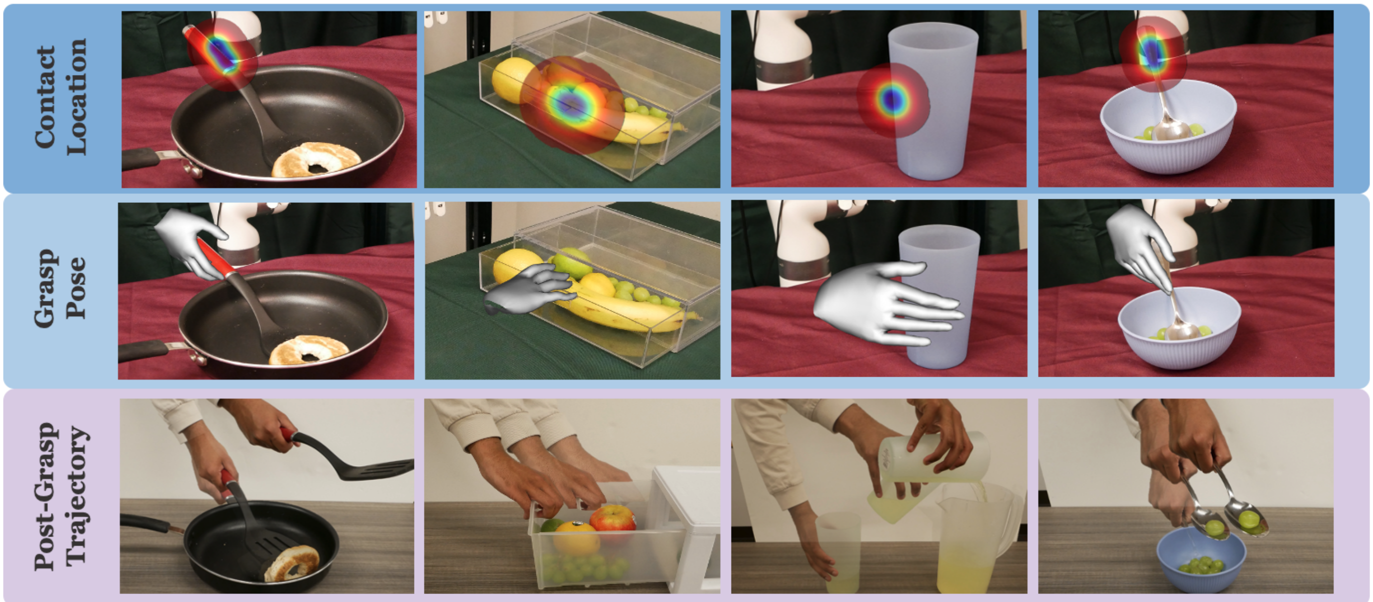

Fig. 4: We can extract different types of data from human motion. In our papers such as DEFT [3], we produce three priors from human videos: the contact location (**top row**) and grasp pose (**middle row**) from the affordance prior; the post-grasp trajectory (**bottom row**) from a human demonstration of the task.

Prior work have introduced two designs to address this challenge (see Fig.3). Both the Allegro and LEAP-C Hand designs, however, sacrifice one degree of freedom either in the extended or closed position. Allegro exhibits reduced dexterity when extended, while LEAP C-Hand (similar to C-Hand in[53]) is less dexterous when closed.

The common reason for this loss of dexterity in both LEAP C-Hand and Allegro is the fixed axis of the motor responsible for adduction-abduction (MCP-2) to the palm of the hand. In LEAP C-Hand, this axis is perpendicular to the palm's plane, while in Allegro, it lies in the plane of the palm. Consequently, when the finger aligns parallel to this axis, the degree of freedom becomes ineffective as seen in 3 In LEAP Hand, a novel ***universal abduction-adduction mechanism*** is proposed for the fingers, allowing them to maintain all degrees of freedom at all MCP positions. Instead of fixing the MCP-2 axis to the palm, the innovative approach brings the axis to the frame of reference of the first *finger* joint, ensuring it remains perpendicular to it at all times. This design enables adduction-abduction in all positions (Fig. 3), providing LEAP Hand with versatility in both the extended position (similar to LEAP C-Hand) and pronation/supination in the flexed position (similar to Allegro).

B. DASH Hand

Oftentimes, the rigid structure of a robot gripper undesirable. The rigid contact with objects can lead to poor grasping abilities and brittle fingers. Underactuated soft hands promise to adaptively conform to the surface of objects without the explicit need for complicated feedback systems. [54] Soft manipulators could tightly grasp around an object and provide good form closure without breaking the object within it, similar to how humans use their hands.

How can we simplify the design of a robot hand to achieve this easily? We adopt a straightforward method by 3D printing soft TPU fingers as a single piece using a standard 3D printer. We thread four tendons through the finger, connecting them to three motors with pulleys. Two tendons are linked to the ends of one pulley, regulating the MCP side-to-side movement of the finger. Another tendon is directed to control the MCP forward motion. The last tendon controls the curling action of the PIP and DIP in the final two flexure joints of the finger. This approach makes the actuation of a soft TPU finger easily attainable.

A key issue with soft robotics hand is how to control them. First, we tension the fishing line tendons which finds the limits of motion that the motor can reach. Second, we must know where the fingertip/end-effector is relative to the motor position. To do this, we collect data of motor angles paired with fingertip positions from an AR tag. We then train a small model which can predict this position open-loop accurately even under deformation.

An additional noteworthy observation is that the form factor of the DASH hand more closely resembles a human-sized hand. This configuration places motors in the wrist, with power transmitted forward using fishing-line tendons, mirroring the way humans utilize tendons to connect arm muscles to parts of their hand. See videos of DASH hand on our website.

C. LEAP Hand v2

LEAP Hand v2 seeks to most closely replicate the human-hand while still being inexpensive, reliable and easy to produce. It features three pivotal elements crucial for unlocking key capabilities in robotic manipulation. Firstly, our innovation involves the introduction of 3D-printed multi-material fingers that closely mimic the stiffness, softness, and durability of human fingers. This characteristic enables our robotic hand

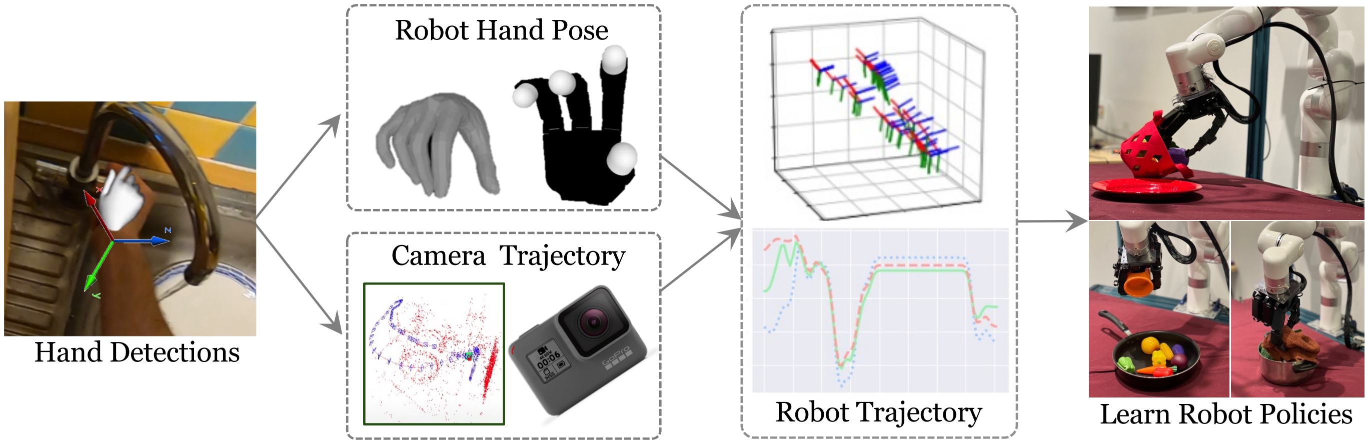

Fig. 5: To learn from human motion, we must retarget the human hand motion to the robot hand and arm. This is done in Robotic Telekinesis and Videodex by tracking the hand and wrist in the video, retargeting them using an internet-video learned retargeting method and then tracking the camera motion using SLAM. [1, 2]

to exert robust forces while maintaining compliance when required. Secondly, we introduce an agile palm with integrated joints meticulously designed to replicate the conformal qualities of the human palm which a feature often lacking in many robot hands but one we consider essential. Our robotic hand effectively utilizes its articulated soft palm for adept grasping, stabilizing, and supporting objects, providing critical opposition from the thumb to the rest of the hand. Thirdly, we design our robotic hand with finger kinematics that maximize dexterity and similarity to a human hand. The Metacarpophalangeal (MCP) joint is engineered for strength and enhanced flexibility compared to a human hand, while the Proximal Interphalangeal (PIP) and Distal Interphalangeal (DIP) joints are articulated with a single fishing line tendon, replicating the articulations of the human hand.

Our demonstration will showcase the hand's strength, dexterity, durability, and suitability for the demands of machine learning research. Significantly, our robotic hand can be fully 3D printed and assembled in a few hours by a novice roboticist using components costing less than $3000. We believe this hand will serve as a robust starting point for numerous labs embarking on their journey into dexterous manipulation. Demoing it with bimanual teleoperation will further help attendees get familiar with the hand.

IV. Our Robot Learning Methodologies

We will bring 6 physical robot hands to show these robot learning methodologies. First, teleoperation using Mocap gloves will be available for attendees to experience on our two physical robot arms and our open-source Mocap setup. Next, our demonstration will show how to use human experience to improve our policy's training distributions. This will be shown through teleoperation from human video and autonomous policy rollouts on the robot hands. Third, we will show how to continually improve these policies through real-world experience and sim2real training.

A. Learning from Human Mocap Glove Demonstrations

In conventional 2-finger gripper manipulation many teleoperation setups have worked successfully to collect demonstrations for behavior cloning. Kinesthetic Methods such as ALOHA [55], GELLO [56] or Da Vinci machines [57] can work accurately. With VR or camera-based hand tracking, methods such as [58] work reasonably well.

However, it is unclear how to scale these methodologies to robot hands. **In this part of the demo, we are developing a mobile system that we can bring to RSS. This will use two robot arms and an open-source Mocap glove teleoperation setup and our robot hands. Participants will be able to teleoperate LEAP v2 using the motion capture system to get hands on experience.** Attendees will have a variety of tabletop tasks and objects that they can try to manipulate and teach the robot how to do these tasks. We will provide information on our open-source pipeline and how to build this system.

Using this system, we can train behavior cloning policies to complete a variety of different tasks that during the event such as tool use or soft manipulation. We will show these policy rollouts, augmented with human-video pretraining on our robot hands.

B. Learning from Human Video

Collecting these demonstrations can be expensive and time-consuming and it is impossible to collect data in all the possible scenarios that the robot will see. This will lead to a distribution shift, where the robot's testing environments will be inevitably different from its training environments. While one could argue that robots could adapt or generalize to unseen scenarios, this is not a guarantee. How can we make our training set larger without collecting more teleoperated data?

Because the robot hand is in a similar embodiment and kinematic structure to the human hand, they will both naturally interact with the world and complete tasks in a very similar way. To leverage this similarity, we can learn from large-scale human motion in human video and motion capture datasets. We find that the kinematic similarity between the human and

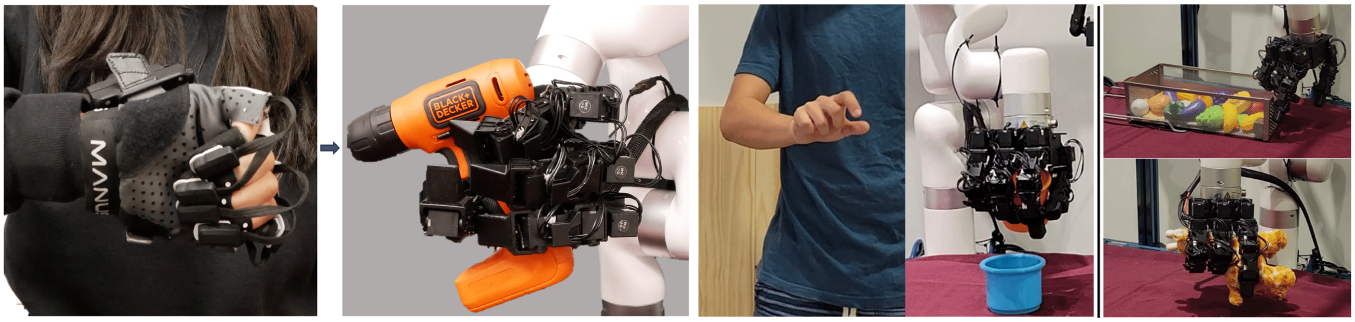

Fig. 6: We will demonstrate real robot hands doing teleoperation from VR glove and teleoperation from human video as developed in [4, 1, 6]. We will build a special version of the VR teleop system for the RSS 2024 demo that will be mobile and compatible with any arm that will be open-sourced if this demo proposal is accepted.

robot hand leads to this method being very effective. In this part of the demo, we will show videos of how to use human hands and motion to improve robot hand behavior.

Our first challenge is in teleoperation in real-time to a robotic hand. This conversion to the robot hand is particularly challenging due to the underconstrained nature of the problem, as the Allegro hand and the 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. In Robotic Telekinesis [1] from RSS 2022, we develop a video-to-robot hand retargeting system that is trained from a corpus of rich and diverse human hand 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. [59] 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. For the arm, we track the trajectory of the moving camera and convert that trajectory to the world frame. These elements are combined together to teleoperate the robot hand and arm. See Figure 5 for further details.

In VideoDex, we would like to use internet videos to directly aid in learning policies. To do this, we use a similar pipeline on EpicKitchens data [60] and retarget it from the human embodiment to the robot embodiment. This process makes human data in the same format as robot embodiment data. This unlocks the ability to easily teach robot embodiment from human videos.

In DEFT, we aim to achieve efficient and generalizable dexterous manipulation by learning from human videos and real-world fine-tuning using only a few samples. [3] The approach first trains a learned affordance model from human videos, extracting information such as grasp location, grasp pose, and task specifics for each task. This means that the model can output human-like hand poses from image inputs alone.

To demonstrate this line of work, **attendees will be able to teleoperate our robot hands from a single monocular camera similar to Robotic Telekinesis from RSS 2022. They will also be able to observe autonomous policy rollouts pre-trained from human video data on pre-recorded videos and on our real robot.**

C. Reinforcement learning

While these internet videos in DEFT can provide good policies on their own, how can robotic systems utilize this to continually learn and improve from experience? In prior work, it is found that learning is often sample-inefficient in the real-world and one must find ways to learn efficiently. [44, 45] The unique part of robot hands is being able to effectively use this human prior experience as a prior for robot hand behavior.

Because the policy does not work zero-shot from internet videos, a fine-tuning procedure is introduced. A residual policy is learned to improve upon the affordance model's predictions. It explores nearby behaviors in the real world, refining the robot's grasp behavior. This is a real-world reinforcement learning setup, where the policy is rewarded based on task

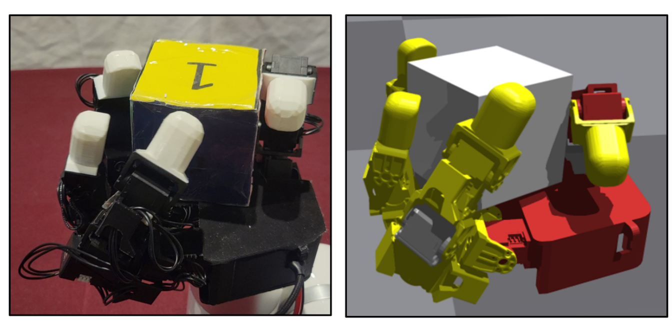

Fig. 7: **Sim2Real transfer.** *Left*: Simulated LEAP Hand in Isaac Gym [61] completing an in-Hand manipulation task. *Right:* LEAP Hand completing the same task in the real world. **This will be demoed at RSS 2024.** Please see our website https://leaphand.com for our open source pipeline and the paper [4] for further details.

completions. These rewards can come from a VLM or from hand labelled human rewards. The system is tested using a variety of kitchen tasks and the results improve the robot's manipulation capabilities through the fine-tuning process. [3]

In Dexterous Functional Grasping, we combine internet data and large-scale simulation training for real-world grasping that can generalize to a wide variety of objects. First, an affordance model is used to predict a functional pose for the hand, considering the intended use of the object. The pre-grasp pose is determined by matching DINO-ViT features to affordance masks obtained from annotated internet images. Sim2real training with an eigengrasp action space, derived from a mocap dataset, is used to restrict the action space to physically realistic hand poses. The reward function for training the policy incentivizes the pickup of objects and ensuring a firm grasp. [6] **This demo will have our fully open-source sim2real in-hand reorientation of a cube example on a real LEAP Hand on display. This is similar to our successful RSS 2023 demo. This demo has been recreated by 100s of people using the LEAP Hand since launch last year. We will also have videos of sim2real and real-world learning on our hand and arm based system doing dexterous grasping.**

V. Conclusion

In this submission, we propose an upcoming robotics demonstration at RSS 2024. It will be an engaging and insightful experience, featuring hands-on real robot demos with six different robot hands and two robot arms. The showcased robot hands include LEAP Hand, known for its popularity and open-source nature [4]; DASH Hand, a soft, tendon-driven robot hand with three degrees of freedom per finger[5], and the LEAP Hand v2, the smallest and strongest open-source robot hand with 3D printed fingers. The demonstration will show these hands doing three key robot learning paradigms: Motion Capture Teleoperation, Learning from Human Video, and Real-world and Simulation-based Reinforcement Learning.

Attendees will have the opportunity to teleoperate all three robot hands using a newly developed open-source mocap system. Additionally, the demo will illustrate the utilization of human hand data for training robot hands and showcase real-world and simulation-based reinforcement learning, including sim2real policies performing various behaviors. The goal is to demystify the field of dexterous manipulation and inspire robotics researchers to incorporate open-source dexterous robot hands into their projects. The demonstrated methods and hardware instructions are open-sourced or will be by RSS 2024. We will develop a unique aspect to this demo: a mobile bimanual teleoperation system and new sim2real results that will be open-sourced after the demonstration at RSS 2024. We hope that these demos inspire researchers to begin their journey in dexterous manipulation research. There are many areas of research, both task-driven and methodology-driven that can significantly benefit the field.

References

[1] A. Sivakumar, K. Shaw, and D. Pathak, "Robotic telekinesis: Learning a robotic hand imitator by watching humans on youtube," 2022. 1, 2, 3, 5, 6

[2] K. Shaw, S. Bahl, and D. Pathak, "VideoDex: Learning Dexterity from Internet Videos," in *Conference on Robot Learning (CoRL)*, 2022. 1, 2, 3, 5

[3] 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. 1, 2, 3, 4, 6, 7

[4] K. Shaw, A. Agarwal, and D. Pathak, "Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning," *arXiv preprint arXiv:2309.06440*, 2023. 1, 2, 3, 6, 7

[5] 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. 1, 2, 7

[6] A. Agarwal, S. Uppal, K. Shaw, and D. Pathak, "Dexterous functional grasping," in *Conference on Robot Learning*. PMLR, 2023, pp. 3453–3467. 1, 2, 3, 6, 7

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

[8] A. Sivakumar, K. Shaw, and D. Pathak, "Robotic telekinesis: learning a robotic hand imitator by watching humans on youtube," *RSS*, 2022. 2

[9] "Shadowhand," https://ninjatek.com/shop/edge/. 2

[10] 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

[11] 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

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

[13] 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

[14] 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

[15] 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

[16] A. Handa, A. Allshire, V. Makoviychuk, 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. 2

[17] "Psyonic ability hand," https://www.psyonic.io/ability-hand. 2

[18] 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. 2

[19] 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. 2

[20] "Inmoov hand," https://inmoov.fr/. 2

[21] "Dexhand," https://www.dexhand.org/. 2

[22] 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

[23] 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

[24] 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

[25] 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

[26] 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

[27] 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

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

[29] 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. 3

[30] "Tesla optimus," https://www.tesla.com/AI. 3

[31] "Figure.ai," https://www.figure.ai/. 3

[32] "Boston dyanmics atlas," https://bostondynamics.com/atlas/. 3

[33] "1x eve," https://www.1x.tech/. 3

[34] "Sanctuary ai," https://sanctuary.ai/. 3

[35] "Digit from agility robotics," https://agilityrobotics.com/. 3

[36] "Pneuflex tutorial," https://www.robotics.tu-berlin.de/menue/software_tutorials/pneuflex_tutorial/. 3

[37] "Reprap open-source 3d printer," https://www.reprap.org/wiki/RepRap. 3

[38] "Ninjatek ninjaflex edge," https://ninjatek.com/shop/edge/. 3

[39] "Recreus filaflex," https://recreus.com/gb/. 3

[40] 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

[41] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "Rma: Rapid motor adaptation for legged robots," *RSS*, 2021. 3

[42] T. Chen, J. Xu, and P. Agrawal, "A system for general in-hand object re-orientation," *Conference on Robot Learning*, 2021. 3

[43] 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

[44] A. Nair, A. Gupta, M. Dalal, and S. Levine, "Awac: Accelerating online reinforcement learning with offline datasets," *arXiv preprint arXiv:2006.09359*, 2020. 3, 6

[45] 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, 6

[46] 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

[47] 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

[48] 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

[49] P. Mandikal and K. Grauman, "Dexvip: Learning dexterous grasping with human hand pose priors from video," in *Conference on Robot Learning (CoRL)*, 2021. 3

[50] 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
[51] J. Pari, N. M. Shafiullah, S. P. Arunachalam, and L. Pinto, "The surprising effectiveness of representation learning for visual imitation," 2021. 3
[52] S. Bahl, A. Gupta, and D. Pathak, "Human-to-robot imitation in the wild," *RSS*, 2022. 3
[53] 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. 3, 4
[54] 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*. IEEE, 2013, pp. 2737–2743. 4
[55] 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. 5
[56] 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. 5
[57] 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. 5
[58] 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. 5
[59] 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. 6
[60] 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. 6
[61] 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. 6