RUKA: Rethinking the Design of Humanoid Hands with Learning

Anya Zorin* Irmak Guzey* Billy Yan Aadhithya Iyer Lisa Kondrich Nikhil X. Bhattasali Lerrel Pinto

New York University
ruka-hand.github.io

***Abstract*—Dexterous manipulation is a fundamental capability for robotic systems, yet progress has been limited by hardware trade-offs between precision, compactness, strength, and affordability. Existing control methods impose compromises on hand designs and applications. However, learning-based approaches present opportunities to rethink these trade-offs, particularly to address challenges with tendon-driven actuation and low-cost materials. This work presents RUKA, a tendon-driven humanoid hand that is compact, affordable, and capable. Made from 3D-printed parts and off-the-shelf components, RUKA has 5 fingers with 15 underactuated degrees of freedom enabling diverse human-like grasps. Its tendon-driven actuation allows powerful grasping in a compact, human-sized form factor. To address control challenges, we learn joint-to-actuator and fingertip-to-actuator models from motion-capture data collected by the MANUS glove, leveraging the hand's morphological accuracy. Extensive evaluations demonstrate RUKA's superior reachability, durability, and strength compared to other robotic hands. Teleoperation tasks further showcase RUKA's dexterous movements. The open-source design and assembly instructions of RUKA, code, and data are available at ruka-hand.github.io.**

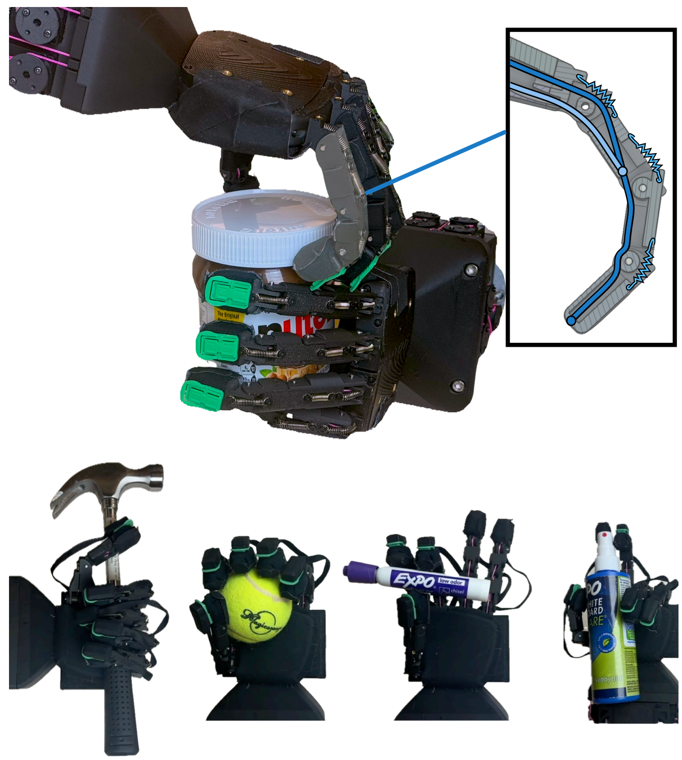

Fig. 1: RUKA is a tendon-driven humanoid hand that is simple, affordable, and capable. Its size and morphology closely match those of a human hand, enabling it to perform diverse human-like power, precision and fine-grained grasps.

I. INTRODUCTION

Achieving dexterity similar to human hands is essential for performing daily tasks [23]. Recent advances in robotics have enabled autonomous dexterous policies [16, 15, 27, 18], driven largely by learning-based methods such as sim-to-real [37, 25] and imitation learning from teleoperated robot or human hand demonstrations [28, 20, 9, 17, 35]. These approaches have been applied to dexterous, multimodal, and long-horizon manipulation tasks [38, 36, 11].

Despite this progress, hardware remains a key bottleneck. An ideal robotic hand must balance precision, compactness, strength, and affordability—goals that are difficult to achieve simultaneously. Designs that prioritize precision often integrate joint-level actuators and encoders, increasing size and weight [33, 1], while tendon-driven systems with external motors [8, 3] offer compactness and strength but introduce control challenges. Position encoders can help, but are costly [8].

Some degree of trade-offs is inevitable given current sensing and actuation technologies. However, we argue that learning-based approaches present an opportunity to rethink some of these trade-offs, particularly to tackle the challenges associated with tendon-driven actuation using low-cost materials.

We introduce RUKA, a simple, affordable, and capable tendon-driven humanoid hand. It is assembled from 3D-printed and off-the-shelf components in 7 hours for under $1300. With five fingers and human-like proportions, it supports smoother learning from demonstrations and integration into human environments. RUKA has 15 degrees of freedom driven by 11 actuators in the forearm and uses tendon actuation to enable diverse, powerful grasps. To address control challenges, we learn fingertip-to-actuator and joint-to-actuator models trained using two key ideas: (1) a MANUS motion-capture glove [6] retrofitted onto the robot for data collection without joint encoders, leveraging morphological similarity; and (2) self-supervised data collection through random actuation sampling to build a broad dataset.

We extensively evaluate RUKA against popular robotic hands and demonstrate its superior reachability, durability, and strength. We further apply RUKA in teleoperation tasks and show that it can perform dexterous movements. The key contributions of this work are:

1) RUKA provides an open-source design for a tendon-driven robotic hand that can be built for under $1,300.
2) RUKA introduces a data-driven control approach that leverages MANUS motion-capture gloves for data collection and learned controllers for fingertip positions and joint angles to support applications like teleoperation.
3) RUKA outperforms popular robotic hands such as LEAP [33] and Allegro [1] across key metrics testing reachability, durability, and strength.

II. Related Work

A. Robotic Hands

Direct-driven hands like the LEAP [33] and Allegro [1] are popular for their low cost and precise joint control, but the Allegro overheats, is hard to repair, and has only four fingers despite its $15,000 price. The LEAP improves durability but remains oversized. Tendon-driven hands with external actuators solve some issues but are expensive, like the $100,000 Shadow Hand [8], or lack precision and support, like the open-source Inmoov [5]. RUKA, in contrast, is a compact, anthropomorphic, tendon-driven hand with a simple, accessible design suited for research.

TABLE I: Comparison of RUKA with LEAP, Allegro, Inmoov, Shadow robotic hands [33, 1, 8] and a human hand baseline.

Robot Hand	Cost	DOF	DOA	Actuation	Open-Source
Human	–	22	–	Tendon	–
LEAP	$2,000	16	16	Direct	✓
Allegro	$15,000	16	16	Direct	✗
Inmoov	$100	14	5	Tendon	✓
Shadow	$100,000	22	20	Tendon	✗
RUKA	**$1,300**	**15**	**11**	**Tendon**	✓

B. Controllers for Hands

Traditional controllers use kinematic models to control joints or end-effector positions. Hands like LEAP [33], Allegro [1], and HRI [29] use direct or geometric mappings, while the Shadow hand [8] uses joint encoders and the Faive Hand [34] estimates joint angles via tendon displacement.

Prior approaches to data-driven control [14, 31, 32, 33] use Vicon motion capture or AR tags, which are rigid and labor-intensive. Human motion data [32, 10] requires pose retargeting and is limited by user morphology.

Inspired by these methods, RUKA also follows a data-driven approach for learning controllers. However, unlike prior work, we enable large-scale autonomous data collection by fitting a motion-capture glove directly to the robotic hand, simplifying the process of gathering supervised data

III. Hardware Design

A. Design Principles

The RUKA hand is designed for functionality and accessibility while balancing anthropomorphism, cost, and reliability.

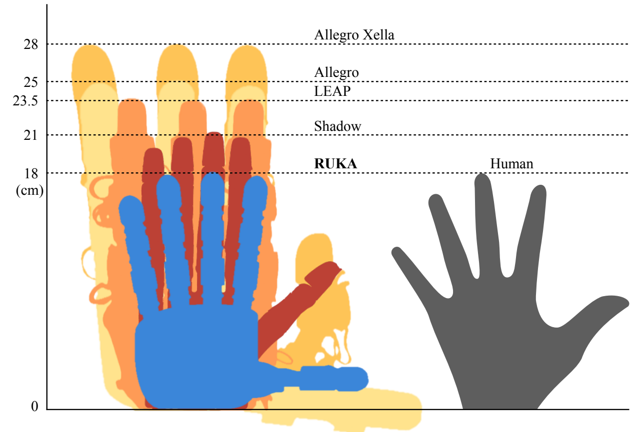

Fig. 2: An illustration of the sizes of different hands that are commonly used by the robotics community next to a human hand.

1) Morphologically Accurate: The RUKA hand mimics human morphology to enable tool use and direct application of human hand data. While hands like LEAP [33] and Allegro [1] match human degrees of freedom, they often fall short in form with fewer fingers and oversized designs.

2) Low-Cost: Morphologically accurate hands like the Shadow Hand [8] are often prohibitively expensive. To make RUKA accessible, we prioritize low cost by using 3D-printed parts and off-the-shelf components. The total material cost, excluding tools, is under $1,300 USD, with $500 and $900 versions available using different Dynamixel motors.

3) Reliability: For RUKA to be a reliable research tool, it must consistently reach commanded positions, operate for long durations without degrading, and be easily repairable. RUKA is designed with this in mind, and the open-source design also enables quick, in-house repairs with minimal downtime.

4) Open-Source: RUKA is fully open-source, with its 3D design and software freely available. RUKA is designed in OnShape for easy sharing and editing. We provide detailed, step-by-step instructions and repair guides. To ensure consistency, we avoid variable methods like drilling and gluing.

B. Kinematics

The RUKA hand features 15 degrees of freedom and 11 actuators. The thumb uses three motors, one per joint, while the four fingers each use two actuators. Although underactuated, this approach reduces weight and complexity while maintaining reliable function, as shown in prior work [21, 26, 30, 22, 34]. Since human DIP and PIP joints are rarely actuated independently [21], RUKA uses a single tendon per finger for both.

While human MCP joints are ball joints, RUKA uses revolute joints for simplicity and rigidity. Two rotations are applied to mimic MCP function: one creates finger splay for natural convergence during grasping, and the other curves the knuckles to support MCP-dependent grasps (Fig. 3). The thumb is also simplified but preserves key functions. It has three joints and degrees of freedom—opposition at the CMC joint, abduction/adduction at the MCP (oriented 90° to the first

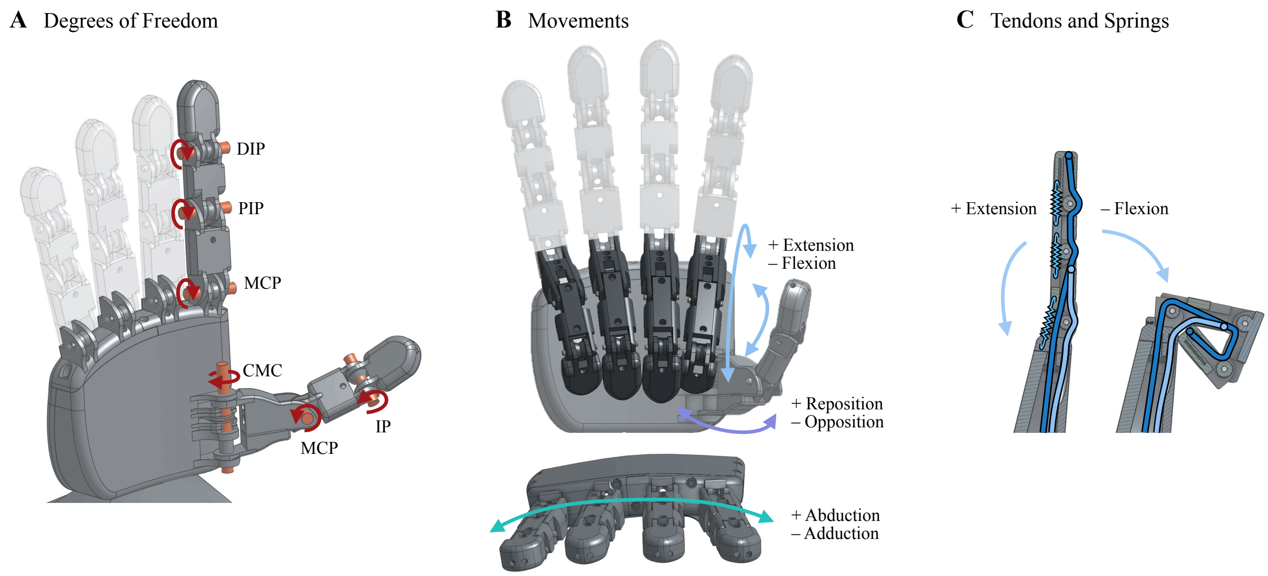

Fig. 3: (**A**) Joints enable 15 degrees of freedom of RUKA labeled with their corresponding joint names. (**B**) The splay of the fingers allows for natural abduction-adduction movement without an active degree of freedom. (**C**) The MCP and PIP / DIP coupled tendons (light blue and dark blue, respectively) are responsible for flexion, while the springs are responsible for extension.

joint), and flexion/extension at the IP (oriented 45° toward the palm). These orientations reflect the natural thumb position.

Using the average hand length and width from a dataset of hand measurements [12], we compare RUKA to human hands and other robot hands in Fig. 2.

C. Materials and Fabrication

The hand's parts are 3D-printed with PLA in 24 hours using the Bambu Lab X1C [2]. The compliant pads for the fingers and palm are printed in FilaFlex Foamy TPU [4]. Printed parts are assembled with heat-set inserts and other off-the-shelf hardware like springs and dowels. Tendons are 200 lb braided fishing line, secured with a slip knot and routed through the fingers into PTFE tubes in the palm, guiding them to the motors. The actuators used are Dynamixel XM430-W210-T motors for the thumb and the lower torque Dynamixel XL330-M288-T motors [7] for the other fingers.

IV. HARDWARE EVALUATION

To evaluate the hardware we run a variety of tests intended to access RUKA hand's robustness. We test the reachability, strength, payload, and ability to reach a variety of grasps, and compare its capabilities to those of other available robot hands.

Reachability Tests. RUKA is able to reach 29 out of 33 standard grasps from [13], showing its ability to achieve most grasps a human hand can. We evaluate the joints' range of motion in degrees and the thumb's opposition capabilities. We randomly sampled 250,000 joint configurations for the thumb and each finger, recording instances where the fingertips touch.

Durability Tests. We ran the RUKA hand for 20 hours without a significant drop in motor precision or overheating issues, outperforming Allegro [1] and LEAP [33].

Strength Tests We ran three strength tests—pinch, payload, and finger slip for the DIP/PIP and MCP joints—to evaluate the performance of each hand. Full test methodology is in the complete paper, and the results are summarized in Table II.

TABLE II: Results of strength tests across different robot hands.

Robot Hand	Pinch (N)	Payload (kg)	DIP/PIP (N)	MCP (N)
Allegro	1.60	3.6	17.8	12.2
LEAP	2.45	4.0	25.17	11.63
Inmoov	2.72	3.2	15.08	-
RUKA	**2.74**	**6.0**	**33.02**	**16.15**

RUKA outperforms other hands in strength metrics. Like RUKA, the LEAP Hand [33] uses XL-330 motors [7], but with three per finger for flexion, suggesting RUKA's performance gains come from its tendon-driven design, which removes actuator weight from the fingers. These results highlight the advantages of lightweight, tendon-driven designs.

V. CONTROLLER

Tendon-driven designs are compact and durable, but their uncertain kinematics make control challenging [9, 20].

In RUKA, we address this by using a data-driven approach for simplified control. Our framework, including data collection and controller learning, is illustrated in Figure 4. This section details our methodology for the hand's controls.

A. Self-Supervised Data Collection

The MANUS gloves [6] track fingertip motion using magnetic field sensors paired with small embedded magnets. As the fingers move, the gloves provide real-time fingertip positions (Fig. 4), represented as $f_t \in \mathbb{R}^{5\times3}$, along with keypoint data estimating the full hand pose, $k_t \in \mathbb{R}^{5\times K\times3}$, where $K = 5$ denotes the number of keypoints per finger.

We collect data autonomously by performing a random walk over uniformly sampled motor positions within their respective limits. To ensure coverage of the entire action space, we

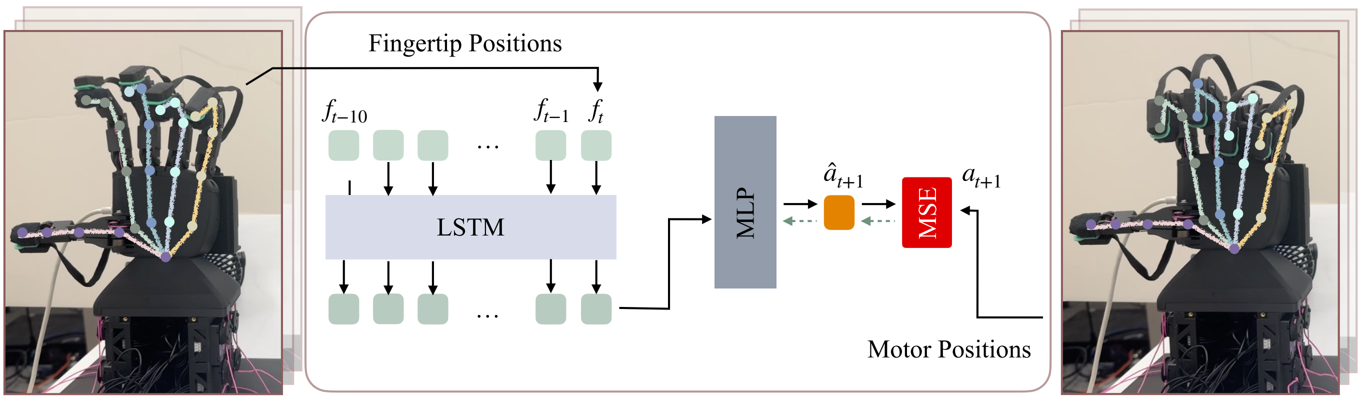

Fig. 4: An illustration of the keypoints received from the MANUS Haptic Gloves (left and right) and the controller architecture used (center) is shown. Fingertip positions are computed from the keypoints and passed as input to an LSTM, along with the previous 10 fingertip positions. The final sequential representation from the LSTM is fed into an MLP head to predict the motor positions for each finger.

repeat this process 500 times for the thumb and 300 times for each of the other four fingers. Throughout data collection, we record fingertip positions, keypoints, joint angles, and both commanded and actual motor positions at 15 Hz.

B. Controller Learning

1) Implementation Details: After data collection, we train a separate controller for each finger, which predicts the corresponding motor positions based on the desired fingertip position for the thumb, and joint angles for the other fingers. To incorporate temporal information, we use a simple Long Short-Term Memory (LSTM) [19] network as a sequential encoder, which processes the past 10 finger states to generate a sequential representation. The final output from the LSTM is passed to an MLP head to predict the next motor position. The loss between the predicted and ground truth motor positions is computed using mean squared error (MSE) and optimized with the AdamW [24] optimizer.

2) Experiments: We explore different architectures, input/output types, and learning parameters to optimize performance. Our experiments aim to answer the following questions: (a) How do architecture choices and learning parameters affect controller performance? (b) How do controllers transfer to different hands? Evaluations are performed on two datasets: a **Human Validation set**, where a user wearing the MANUS glove moves their hand through various poses, and a **Robot Validation set**, where the glove is mounted to the robot and the fingertips are manually manipulated. In both cases, RUKA replays recorded keypoints using the trained controllers, and accuracy is measured by comparing the reproduced fingertip positions to the originals. We compare the four architectures in Table III. A full analysis is included in the complete paper.

TABLE III: Comparison of the performance of different architectures. Error values are computed over three positional axes for the thumb.

Method	Error in Robot Val (cm)	Error in Human Val (cm)
Search Based	[0.55 0.59 0.48]	[2.03 2.0 2.01]
MLP	[0.53 0.56 0.43]	[1.8 1.90 2.0]
k-NN	[0.13 0.14 0.15]	[0.95 0.42 0.83]
RUKA	[0.20 0.27 0.22]	[0.83 0.60 0.54]

To maintain the usability of RUKA across different builds, we implement an auto-calibration procedure that estimates motor ranges based on tendon tension, enabling generalization across builds. After calibration, the average fingertip position difference between two builds is only 3mm. We hope that these experiments will help ensure consistency across different RUKA hands and make the system more accessible.

VI. Applications of RUKA

Teleoperation. Using the controller (Sec. V), we teleoperate RUKA with a MANUS glove at 25 Hz to collect demonstrations for various dexterous tasks. RUKA also supports OpenTeach [20] teleoperation.

Policy Learning. We use RUKA for policy learning by deploying HuDOR [17], an imitation learning framework that uses in-scene human videos, converts them into robot replays, and learns a residual policy to finetune the open-loop replay trajectory from the demonstrations. Rewards are computed by matching the trajectories of object centroids between the robot episodes and the human demonstrations. We apply this approach to two tasks: *Cube Flipping* and *Bread Pick and Drop*. On average, each policy is trained for 40 episodes, with training converging in approximately 45 minutes.

VII. Discussion

Limitations: Training and pose capture currently rely on the Manus glove, which, while accurate, adds a cost barrier to replication. RUKA lacks tactile sensing, which may limit performance in complex tasks—integrating touch sensors could improve precision and adaptability. To keep the design simple, abduction at the MCP joints is not supported, potentially limiting dynamic movements.

Hardware failure modes: The two main failure modes are motor wear and joint damage. The plastic-geared Dynamixel motors can degrade over time but are easily replaced or upgraded to metal-geared alternatives. Collisions may damage the MCP joint and other 3D printed parts, but repairs typically take under half an hour. RUKA is also highly customizable, allowing users to upgrade to more durable components as needed.

References

[1] Allegro. URL https://www.allegrohand.com/ah-v4.
[2] Bambu 3d printer. URL https://us.store.bambulab.com/products/x1-carbon.
[3] Clone hand. URL https://clonerobotics.com/hand.
[4] Filaflex. URL https://recreus.com/gb/filaments/24-filaflex-foamy.html.
[5] Inmoov. URL https://inmoov.fr/inmoov-hand/.
[6] Manus haptic gloves. URL https://www.manus-meta.com/.
[7] Dynamixel motors. URL https://www.robotis.us/dynamixel-xc330-m288-t/.
[8] Shadow hand. URL https://www.shadowrobot.com/dexterous-hand-series/. Shadow Hand.
[9] Sridhar Pandian Arunachalam, Irmak Güzey, Soumith Chintala, and Lerrel Pinto. Holo-dex: Teaching dexterity with immersive mixed reality, 2022.
[10] Dominik Bauer, Cornelia Bauer, Jonathan P. King, Daniele Moro, Kai-Hung Chang, Stelian Coros, and Nancy Pollard. Design and control of foam hands for dexterous manipulation. *International Journal of Humanoid Robotics*, 17(1), February 2020.
[11] Yuanpei Chen, Chen Wang, Li Fei-Fei, and C Karen Liu. Sequential dexterity: Chaining dexterous policies for long-horizon manipulation. *arXiv preprint arXiv:2309.00987*, 2023.
[12] E. Churchill, L. L. Laubach, J. T. Mcconville, and I. Tebbetts. *Anthropometric source book. Volume 1: Anthropometry for designers*. Number NASA-RP-1024. NASA, 1978. URL https://ntrs.nasa.gov/citations/19790003563.
[13] Thomas Feix, Javier Romero, Heinz-Bodo Schmiedmayer, Aaron M. Dollar, and Danica Kragic. The grasp taxonomy of human grasp types. *IEEE Transactions on Human-Machine Systems*, 46(1):66–77, 2016. doi: 10.1109/THMS.2015.2470657.
[14] Michele Giorelli, Federico Renda, Marcello Calisti, Andrea Arienti, Gabriele Ferri, and Cecilia Laschi. Neural network and jacobian method for solving the inverse statics of a cable-driven soft arm with nonconstant curvature. *IEEE Transactions on Robotics*, 31(4):823–834, 2015. doi: 10.1109/TRO.2015.2428511.
[15] Irmak Guzey, Yinlong Dai, Ben Evans, Soumith Chintala, and Lerrel Pinto. See to Touch: Learning Tactile Dexterity through Visual Incentives. *arXiv e-prints arXiv:2309.12300*, September 2023.
[16] Irmak Guzey, Ben Evans, Soumith Chintala, and Lerrel Pinto. Dexterity from touch: Self-supervised pre-training of tactile representations with robotic play, 2023.
[17] Irmak Guzey, Yinlong Dai, Georgy Savva, Raunaq Bhirangi, and Lerrel Pinto. Bridging the human to robot dexterity gap through object-oriented rewards. *arXiv preprint arXiv:2410.23289*, 2024.
[18] Ankur Handa, Arthur Allshire, Viktor Makoviychuk, Aleksei Petrenko, Ritvik Singh, Jingzhou Liu, Denys Makoviichuk, Karl Van Wyk, Alexander Zhurkevich, Balakumar Sundaralingam, Yashraj Narang, Jean-Francois Lafleche, Dieter Fox, and Gavriel State. Dextreme: Transfer of agile in-hand manipulation from simulation to reality. *arXiv*, 2022.
[19] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL https://doi.org/10.1162/neco.1997.9.8.1735.
[20] Aadhithya Iyer, Zhuoran Peng, Yinlong Dai, Irmak Guzey, Siddhant Haldar, Soumith Chintala, and Lerrel Pinto. OPEN TEACH: A Versatile Teleoperation System for Robotic Manipulation. *arXiv e-prints arXiv:2403.07870*, March 2024.
[21] Atzori M Müller H. Jarque-Bou NJ, Scano A. Kinematic synergies of hand grasps: a comprehensive study on a large publicly available dataset. *Journal of NeuroEngineering and Rehabilitation*, 2019.
[22] Yong-Jae Kim, Junsuk Yoon, and Young-Woo Sim. Fluid lubricated dexterous finger mechanism for human-like impact absorbing capability. *IEEE Robotics and Automation Letters*, 4(4):3971–3978, 2019. doi: 10.1109/LRA.2019.2929988.
[23] Vikash Kumar, Emanuel Todorov, and Sergey Levine. Optimal control with learned local models: Application to dexterous manipulation. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 378–383, 2016. doi: 10.1109/ICRA.2016.7487156.
[24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
[25] Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-Level Reward Design via Coding Large Language Models. *arXiv e-prints arXiv:2310.12931*, October 2023.
[26] Giulia Matrone, Christian Cipriani, Emanuele Lindo Secco, Giovanni Magenes, and Maria Chiara Carrozza. Principal components analysis based control of a multi-dof underactuated prosthetic hand. *Journal of NeuroEngineering and Rehabilitation*, 7:16 – 16, 2010. URL https://api.semanticscholar.org/CorpusID:7324783.
[27] OpenAI, Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, Jonas Schneider, Szymon Sidor, Josh Tobin, Peter Welinder, Lilian Weng, and Wojciech Zaremba. Learning dexterous in-hand manipulation, 2019.
[28] Jyothish Pari, Nur Muhammad Shafiullah, Sridhar Pandian Arunachalam, and Lerrel Pinto. The surprising effectiveness of representation learning for visual imitation, 2021.
[29] Hyeonjun Park and Donghan Kim. An open-source anthropomorphic robot hand system: Hri hand. *HardwareX*, 7:e00100, 2020. ISSN 2468-0672. doi: https://doi.org/10.1016/j.ohx.2020.e00100. URL https://www.sciencedirect.com/science/article/pii/

S2468067220300092.

[30] Steffen Puhlmann, Jason Harris, and Oliver Brock. Rbo hand 3: A platform for soft dexterous manipulation. *IEEE Transactions on Robotics*, 38(6):3434–3449, 2022. doi: 10.1109/TRO.2022.3156806.

[31] Matthias Rolf and Jochen J. Steil. Efficient exploratory learning of inverse kinematics on a bionic elephant trunk. *IEEE Transactions on Neural Networks and Learning Systems*, 25(6):1147–1160, 2014. doi: 10.1109/TNNLS.2013.2287890.

[32] Cornelia Schlagenhauf, Dominik Bauer, Kai-Hung Chang, Jonathan P. King, Daniele Moro, Stelian Coros, and Nancy Pollard. Control of tendon-driven soft foam robot hands. In *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*, pages 1–7, 2018. doi: 10.1109/HUMANOIDS.2018.8624937.

[33] Kenneth Shaw, Ananye Agarwal, and Deepak Pathak. Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning. *Robotics: Science and Systems (RSS)*, 2023.

[34] Yasunori Toshimitsu, Benedek Forrai, Barnabas Gavin Cangan, Ulrich Steger, Manuel Knecht, Stefan Weirich, and Robert 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)*, pages 1–7, 2023. doi: 10.1109/Humanoids57100.2023.10375231.

[35] Chen Wang, Haochen Shi, Weizhuo Wang, Ruohan Zhang, Li Fei-Fei, and C. Karen Liu. DexCap: Scalable and Portable Mocap Data Collection System for Dexterous Manipulation. *arXiv e-prints arXiv:2403.07788*, March 2024.

[36] Hecheng Wang, Lizhe Qi, Bin Fang, and Yunquan Sun. Hierarchical visual policy learning for long-horizon robot manipulation in densely cluttered scenes. *arXiv preprint arXiv:2312.02697*, 2023.

[37] Jun Wang, Ying Yuan, Haichuan Che, Haozhi Qi, Yi Ma, Jitendra Malik, and Xiaolong Wang. Lessons from learning to spin" pens". *arXiv preprint arXiv:2407.18902*, 2024.

[38] Ying Yuan, Haichuan Che, Yuzhe Qin, Binghao Huang, Zhao-Heng Yin, Kang-Won Lee, Yi Wu, Soo-Chul Lim, and Xiaolong Wang. Robot synesthesia: In-hand manipulation with visuotactile sensing. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6558–6565. IEEE, 2024.