

# 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](https://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](https://ruka-hand.github.io)

## I. INTRODUCTION

Achieving dexterity similar to human hands is essential for performing daily human tasks [30]. In recent years, significant progress has been made in robotics toward developing autonomous dexterous policies [18, 17, 35, 21]. Many of these advances have been achieved through learning-based approaches, including sim-to-real methods [49, 32] and imitation learning methods using teleoperated robot demonstrations [36, 25, 2] or human hand demonstrations [19, 47]. Robotic hands are now being applied in dexterous, multimodal [51, 5], and long-horizon [48, 7] tasks.

While these achievements are remarkable, progress has been limited by hardware. An ideal robotic hand must balance precision, compactness, strength, and affordability. Yet these requirements remain challenging to achieve simultaneously, so existing hand designs have needed to make trade-offs based on the available control methods and target applications. To prioritize precision, some hands integrate motors with encoders directly in their joints, but this significantly




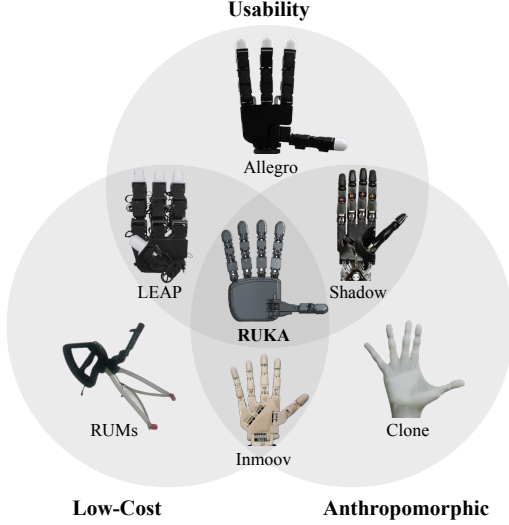
Fig. 1: RUKA is a tendon-driven humanoid hand that is simple, affordable, and capable. Its size and morphology closely match those of a human hand, enabling it to perform diverse human-like power, precision, and fine-grained grasps.

increases the hand’s size and weight [44, 1]. To prioritize compactness and strength, other hands adopt a tendon-driven design with actuators outside the hand [43, 11], but this introduces nonlinearities, elasticities, and uncertainties into the force transmission system, making it challenging to model how actuators affect joint angles and fingertip positions. This challenge can be addressed to an extent by incorporating small position encoders into the hand [43], but this is expensive, prone to mechanical failure, and difficult to repair.

Some degree of trade-offs is inevitable given current sensing

\*Equal contribution. Correspondence to [az61@nyu.edu](mailto:az61@nyu.edu), [irmakguzey@nyu.edu](mailto:irmakguzey@nyu.edu).

## A Design Principles



## B Size Comparisons

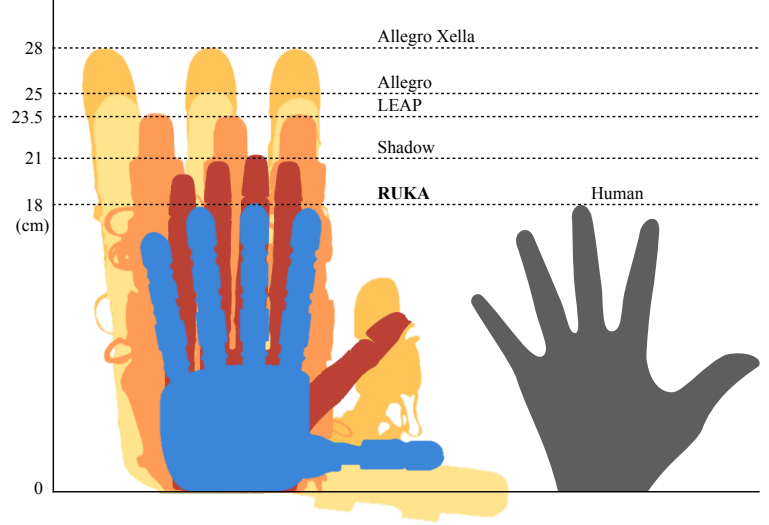


Fig. 2: (A) A Venn diagram of a variety of robotic hands [1, 44, 43, 13, 24, 11] demonstrates RUKA’s unique combination of low cost, anthropomorphism and usability. (B) An illustration of the sizes of different hands that are commonly used by the robotics community. RUKA is designed to closely match the average human hand.

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.

In this work, we introduce RUKA, a tendon-driven humanoid hand that is simple, affordable, and capable. Made from 3D-printed parts and off-the-shelf components, it takes approximately 7 hours to assemble and costs under \$1300 USD. RUKA has 5 fingers, including an opposable thumb, with dimensions matching the average human hand. These features enable smoother learning from human demonstrations and integrating into human environments. Its 15 degrees of freedom (DOFs) are driven by 11 actuators housed in the forearm and connected to finger linkages using flexible tendons. This tendon-driven actuation allows for diverse, human-like, powerful grasps in a compact, human-sized form factor. To tackle the associated control challenges, we use learning-based techniques to develop fingertip-to-actuator and joint-to-actuator models that predict the actuation commands required to produce target fingertip and joint positions. Training these models involves two key ideas: (1) We fit MANUS gloves [33]—off-the-shelf motion-capture gloves designed for human hands—onto our robot hands to collect fingertip and joint positions without using joint encoders. This crucially leverages our hand’s morphological similarity to human hands. (2) We autonomously collect labeled data points in the form of (fingertip/joint position, actuation command) pairs by procedurally sampling actuation commands within motor limits. This method captures the broad dataset needed for the models to learn these nonlinear relationships.

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 [44] and Allegro [1] across key metrics testing reachability, durability, and strength.

## II. RELATED WORK

### A. Robotic Hands

Direct-driven robotic hands like the LEAP [44] and Allegro [1] are popular in research due to their low cost and precise control using motors located directly in the joints. Despite its \$15,000 price, the Allegro hand has only four fingers, overheats easily, and is hard to repair. The LEAP hand improves endurance and repairability but, like other direct-drive designs, remains oversized and not human-like.

Tendon-driven hands with external actuators address these issues, but commercial options like the \$100,000 Shadow Hand [43] are costly and require complex repairs. Open-source alternatives, such as the Inmoov [24] and Dexhand [12], allow easier repairs but often lack infrastructure and precise control.

Soft robotic hands offer many advantages, including compliance and more natural grasps. The RBO3 hand is a soft, pneumatically actuated design that, while beneficial, is less precise and difficult to simulate due to its flexibility [39]. The ADAPT hand addresses this with a rigid link structure,

compliant joints, and a compliant skin [27]. Though both hands are robust, neither is open-source or available to buy.

Biomimetic hand designs, which mimic the exact tendon actuation of a human hand, are often extremely complex but robust. Hands such as the FLLEX hand [29] and the hand in Xu et al. [50] use complex pulley systems and custom parts. They are highly precise while remaining compliant. However, because of their complex nature, they are difficult to reproduce. In contrast, RUKA is a compact, tendon-driven, anthropomorphic hand with a simple, accessible design, making it an effective research tool.

TABLE I: Comparison of RUKA with LEAP, Allegro, Allegro Xella, Inmoov, Shadow robotic hands [44, 1, 43] and a human hand baseline. Evaluation across cost, degrees of freedom (DOF), degrees of actuation (DOA), actuation type (direct-drive or tendon-driven), and whether the design is open-source.

Robot Hand	Cost	DOF	DOA	Actuation	Open-Source
Human	—	22	—	Tendon	—
LEAP	\$2,000	16	16	Direct	✓
Allegro	\$15,000	16	16	Direct	✗
Allegro Xella	\$75,000	16	16	Direct	✗
Inmoov	\$100	14	5	Tendon	✓
Shadow	\$100,000	22	20	Tendon	✗
<b>RUKA</b>	<b>\$1,300</b>	<b>15</b>	<b>11</b>	<b>Tendon</b>	<b>✓</b>

### B. Controllers for Hands

Traditionally, hand controllers use well-defined kinematic models to control joint angles or end-effector positions. The LEAP [44] and Allegro [1] hands employ simple controllers where motor positions directly dictate joint angles. The HRI hand [37] uses a single motor to drive a rigid two-four-bar linkage, allowing geometric calculation of fingertip positions. The Shadow Hand [43] features joint encoders for closed-loop control, while the Faive Hand [45] estimates joint angles based on tendon displacement.

When kinematics and proprioception are ill-defined, as in soft and tendon-driven hands, prior work has turned to data-driven approaches for learning controllers. While these methods have achieved great success in policy learning for dexterous manipulation [18, 17, 19, 20, 8, 35, 9], applying them to controller learning remains challenging due to difficulties in collecting ground-truth supervised data.

Existing approaches [16, 41, 42, 44] rely on infrared sensors, Vicon motion capture, or AR tags to collect robot position data, but these methods are rigid and labor-intensive. Motion capture gloves have also been used to collect human motion data [42, 4], which is then retargeted to robots via AR tags or video-based pose estimation. However, these approaches are limited by their dependence on the wearer’s morphology and the need for carefully curated training poses.

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.

1) *Morphologically Accurate*: The RUKA hand mimics human morphology to enable tool use and direct application of human hand data. True anthropomorphism requires more than matching degrees of freedom—it demands accurate size, finger count, and overall form [46]. While hands like LEAP [44] and Allegro [1] match human degrees of freedom, they often fall short in form with fewer fingers and oversized designs. In contrast, RUKA replicates the human hand’s shape to minimize the need for retargeting human data, allowing for seamless interaction with tools and everyday objects.

2) *Low-Cost*: Existing morphologically accurate robotic hands, such as the Shadow Hand [43], are often prohibitively expensive. For RUKA hand to serve as an accessible research tool, keeping costs low is a priority. To do so, we prefer 3D-printed parts and off-the-shelf components. The total cost of the raw materials needed, excluding tools (a 3D-printer and soldering iron), is under \$1,300 USD. There is also a \$500 and \$900 version of RUKA with varying Dynamixel motors.

3) *Reliability*: For RUKA to serve as a reliable research tool, it must consistently reach commanded positions, operate for long durations without degradation, and be easily repairable. Rigid hinge joints are needed for repeatability, as compliance in actuation mechanisms can introduce uncertainty [6]. However, to balance rigidity with functional flexibility, we added soft pads to the fingers, improving task performance while maintaining durability. Additionally, the hand’s open-source design allows for quick, in-house repairs, ensuring minimal downtime. The Dynamixel motors [40] allow for accurate positioning, and the motor case has ventilation to prevent overheating during long runtimes. To ensure consistency across builds, we avoid construction methods that introduce variability, such as drilling and gluing.

4) *Open-Source*: RUKA is fully open-source, with its 3D design and software freely available. Hands like Inmoov [24] lack editable CAD files, making design modification difficult. RUKA is designed in OnShape, allowing easy sharing and editing of CAD files. The control code and a MuJoCo model of the hand is available on GitHub. Open-source hardware is cheap, adaptable, and fosters collaboration among users and developers. However, it is also important to ensure easy assembly for inexperienced users and maintain documentation [38]. We provide step-by-step assembly instructions and repair guides. Assembling RUKA takes approximately 7 hours, and most repairs take under 20 minutes.

### B. Kinematics

The RUKA hand is designed with 15 degrees of freedom and 11 actuators. The thumb is actuated by 3 motors, 1 for each joint, while the other 4 fingers each use 2 actuators, actuating the proximal interphalangeal (PIP) and distal interphalangeal (DIP) joints together (Fig. 3). Although the


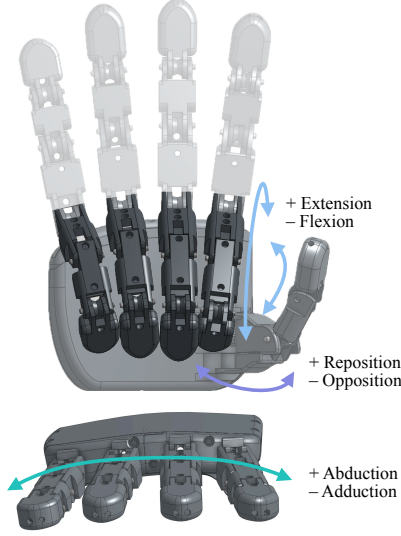
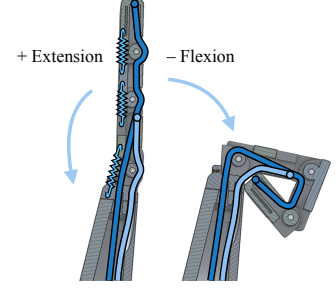
**A** Degrees of Freedom**B** Movements**C** Tendons and Springs

Fig. 3: (A) Joints enable 15 degrees of freedom of RUKA labeled with their corresponding joint names. (B) The splay of the fingers allow 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.

hand is underactuated, previous studies on robotic hands have demonstrated that underactuation can provide robust and reliable functionality while significantly reducing weight and mechanical complexity [26, 34, 39, 29, 45]. In human hands, DIP and PIP are rarely actuated independently [26], so RUKA actuates these joints with a single tendon per finger.

The metacarpophalangeal (MCP) or knuckle joint in human fingers is a ball joint, but for simplicity and rigidity, it is modeled as a revolute joint in RUKA fingers. To replicate the functions of the MCP joint, two rotations are applied to the knuckle where the joint attaches. The first rotation creates a splay, spreading the four fingers outward when the hand is fully open, allowing them to converge naturally as the hand closes, mimicking the adduction of the MCP joints (Fig. 3). The second rotation mimics the curve of human knuckles, which originates from the carpometacarpal (CMC) joints (Fig. 3). A concave curvature is applied to the palm and knuckles and is enhanced by slightly rotating the fingers inward toward each other. These adaptations improve the hand's ability to achieve grasps dependent on the MCP (Fig. 7).

The thumb's kinematics are also simplified while retaining critical functionality. A human thumb typically has

five degrees of freedom from three joints, enabling abduction/adduction, flexion/extension, and opposition motions. The RUKA thumb is modeled with three degrees of freedom from three joints to simplify control while optimizing grasping performance (Fig. 3), with the ranges of motion described in Table II. The carpometacarpal (CMC) joint handles opposition. The metacarpophalangeal (MCP) joint controls abduction and adduction and is oriented 90 degrees relative to the first joint. The interphalangeal (IP) joint controls flexion and extension, rotated 45 degrees toward the palm relative to the second joint. This configuration mimics the orientation of these joints in the natural position of the thumb.

The hand is made to be similar to a human hand size. Using the average hand length and width from a dataset of hand measurements [10], we compare RUKA to human hands and other robot hands (Fig. 2).

### C. Materials and Fabrication

1) *3D-Printed Parts*: The hand's parts are 3D-printed in 24 hours using the Bambu Lab X1C [3]. The finger joints, palm, wrist, and motor base are printed using PLA, chosen for its rigidity and ease of printing. The compliant pads for the fingers and palm are printed in FilaFlex Foamy TPU [15]. These pads are then attached using glue, and 3M friction tape is applied on top for additional grip.

2) *Off-The-Shelf Components*: Heat-set inserts are used to securely assemble the printed parts, allowing the components to be fastened with screws. These inserts are pressed into the printed parts using a soldering iron. Metal dowels are used for the pin joints, which are spring-loaded open with an extension spring. The tendons are made of braided fishing line rated for 200 lbs, selected for its combination of strength and flexibility. The line is secured to the printed components

TABLE II: The range of motion of each joint in RUKA in degrees compared to the human range of motion [23].

Joint Name	Range of Motion	Human
Distal Interphalangeal (DIP)	120°	85°
Proximal Interphalangeal (PIP)	120°	105°
Metacarpophalangeal (MCP, Finger)	140°	85°
Interphalangeal (IP)	120°	80°
Metacarpophalangeal (MCP, Thumb)	90°	56°
Carpometacarpal (CMC)	190°	—



using a slip knot, routed through the finger mechanism and a PTFE tube in the palm, which guides it to the motor. The low-friction properties of PTFE significantly reduce resistance during tendon movement, ensuring smooth operation. This routing system is highly adaptable, as the only fixed points are the start and end of the PTFE tubes, allowing for easy adjustments and reconfiguration.

3) *Actuators*: The actuators are Dynamixel XM430-W210-T motors for the thumb and Dynamixel XL330-M288-T motors [40] for the other fingers. Higher torque motors are used for the thumb. Communication with the motors is achieved through a USB-to-serial bridge. The motors are powered by a 12V and 5V dual-output power supply and connected via a custom wiring harness.

#### IV. HARDWARE EVALUATION

We run a variety of tests that assess RUKA hand’s capabilities and robustness. Specifically, we test its reachability, durability, and strength compared to other robot hands.

##### A. Reachability Tests

1) *Kapandji Test*: We evaluate the classical Kapandji test [28]. RUKA scores 10/10 for all poses, while the Allegro and LEAP hands each score 9/10, losing a point due to having fewer fingers.

2) *Range of Motion Test*: We evaluate the joint ranges of motion (Table II). We also evaluate the thumb’s opposition capabilities by randomly sampling 250,000 joint configurations for the thumb and each finger and recording instances where the fingertips touch (Fig. 4).

3) *Grasp Test*: We evaluate the 33 standard grasps from the GRASP Taxonomy [14]. RUKA successfully reproduces 29 out of 33 human hand grasps, as shown in Fig. 7. The hand is able to reach grasps that rely on MCP adduction and thumb degrees of freedom that are underactuated.




Fig. 4: The intersection space of the thumb fingertip and each of the fingertips overlaid on the hand. This demonstrates the large set of opposable grasps possible with RUKA.

##### B. Durability Tests

Durability is essential for a usable research robot. When robot hands run continuously for more than an hour, issues arise like motor overheating and declining hand precision. This means researchers cannot run continuous tasks for long periods, and the robot hand needs supervision. We were able to run the RUKA hand for 20 hours continuously, without a significant drop in motor precision. With 20 hours as a

lower bound, RUKA outperforms both Allegro [1] and LEAP [44] hands. We also recorded the motor temperatures over 90 minutes. Even run continuously, the temperatures stabilize below the maximum rating, allowing for long run times (Fig. 5).




Fig. 5: We run the hand continuously for 90 minutes, repeatedly doing a full range of motion. Here we show the temperature of each motor during. Note how the temperature stabilizes after some time.

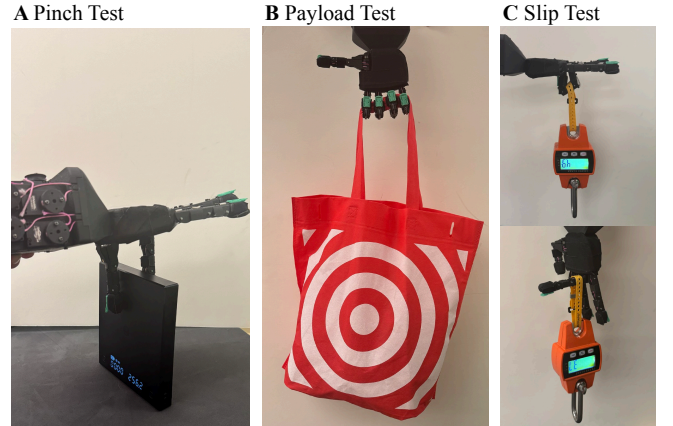


Fig. 6: The experimental setup for the strength tests conducted on robot hands.

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

Robot Hand	Pinch (N)	Payload (kg)	DIP/PIP (N)	MCP (N)
Allegro Xella	1.62	2.2	8.34	3.57
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	-
<b>RUKA</b>	<b>2.74</b>	<b>6.0</b>	<b>33.02</b>	<b>16.15</b>

##### C. Strength Tests

1) *Pinch Test*: We actuate the thumb and index fingers on each hand to pinch a vertical scale measuring force. Each robot

## RUKA Executing GRASP Taxonomy



Fig. 7: RUKA doing the 29 out of 33 Grasps from the GRASP Taxonomy [14], including a variety of power grasps, precision grasps, and intermediate grasps. Red grasps were not reached and yellow were partially reached or unstable to perturbations.

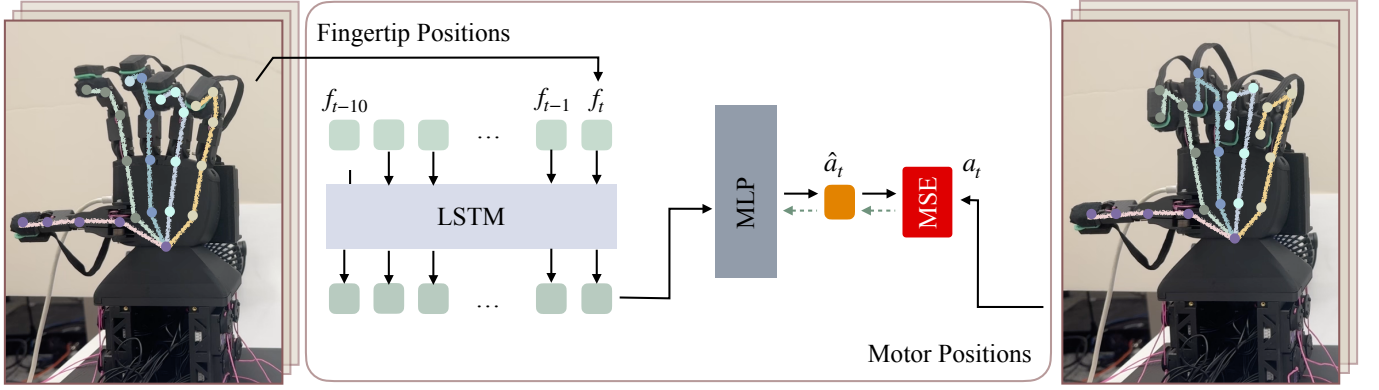


Fig. 8: Keyoints received from the MANUS Haptic Gloves (left and right) and the controller architecture (center). Fingertip positions are computed from the keyoints and passed as input to an LSTM, along with the previous 10 fingertip positions. The final sequential representation from the LSTM is fed into an MLP head to predict the motor positions for each finger.

hand is tested 3 times, with the highest force recorded. The average force across left/right hands is reported in Table III.

2) *Payload Test*: We actuate the PIP and DIP joints on the non-thumb fingers to curl around the handle of a cloth bag, to which we add weight until any joint angle error exceeds 15 degrees. The final weight is reported in Table III.

3) *Slip Test*: We actuate joints in combination to hold a hanging scale measuring force, to which we add weight until the joint angle error exceeds 15 degrees. We test the force supported by the combined DIP/PIP joints and the MCP joints of non-thumb fingers. The Inmoov hand [24] only has one actuator per finger, so DIP/PIP and MCP slip force cannot be measured; instead, the whole finger is measured jointly. The average force across the 4 fingers is reported in Table III.

#### D. Analysis

RUKA has the highest performance across all strength metrics. The LEAP Hand [44] uses the same Dynamixel motors [40] as RUKA, with three per finger instead of two for flexion. This suggests RUKA’s improved performance over LEAP is primarily due to its tendon-driven design, which removes actuator weight from the fingers. When compared to Inmoov [24], which is open-source and lower cost than RUKA, RUKA outperforms Inmoov with an increased payload capacity of 87.5% and a 64.2%<sup>1</sup> improvement in the slip test. This performance gain is primarily due to Inmoov’s single-actuator finger design and lower-torque motors. While Inmoov achieves a similar pinch force to RUKA, this is mainly due to Inmoov’s complaint silicone fingertips, as it is only able to exert 57.4% of force without the silicone tips. These results show that lightweight, tendon-driven designs offer significant advantages in performance.

#### V. CONTROLLER

The tendon-driven approach to hand design offers significant advantages, including durability and a compact form

<sup>1</sup>For comparison with Inmoov, which uses only one actuator, we use RUKA’s average slip force across the DIP/PIP and MCP joints.

factor, allowing robot hands to more closely mimic human hands. However, this introduces challenges in control. The absence of a direct mapping between motor positions and fingertip positions or joint angles makes it difficult to apply standard control techniques commonly used in robotics [2, 25].

In RUKA, we find the best of both worlds by employing a data-driven approach, achieving a balance between a compact, durable form and simplified control. We illustrate our framework involving the data collection process and controller learning architecture (Figure 8). This section provides a detailed explanation of our data collection process and the learning methodology used for the hand controls.

##### A. Autonomous Data Collection

To effectively utilize a learning-based approach, we first integrate the MANUS Haptic Gloves [33] with RUKA using a custom 3D-printed attachment that secures the gloves to the back of the hand, along with the fingertip attachments provided by MANUS. The gloves track fingertip motion using magnetic field sensors paired with small embedded magnets. As the fingers move, the gloves provide real-time fingertip positions, represented as  $f_t \in \mathbb{R}^{5 \times 3}$ , along with keypoint data estimating the full hand pose,  $k_t \in \mathbb{R}^{5 \times K \times 3}$ , where  $K = 5$  denotes the number of keypoints per finger.

Once the gloves are attached, we initiate autonomous data collection by performing a random walk over uniformly sampled motor positions within their respective limits. For each finger, we uniformly sample from its corresponding motor positions, move the finger to that position, and then, for 100 steps, randomly increase or decrease each motor position. As described in Section III, the thumb is actuated by three motors, whereas each of the other four fingers is actuated by two motors. To ensure coverage of the entire action space, we repeat this process 500 times for the thumb and 300 times for each of the other four fingers.

During data collection, we compute joint angles by taking the dot product between vectors connecting consecutive keypoints. Throughout data collection, we record fingertip






Fig. 9: RUKA’s controller testing framework. The top row shows how we collect our test data on the powered-off RUKA hand and bottom row shows corresponding controller performance on the desired poses.

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. Each controller predicts the corresponding motor positions based on the desired fingertip position for the thumb, and joint angles for the other fingers. This distinction is necessary because (a) the thumb joints differ between the human and the robot, and (b) MANUS outputs different fingertip representations for the human and the robot for the other four fingers.

We illustrate the controller learning architecture in Fig. 8. To incorporate temporal information, we use a simple Long Short-Term Memory (LSTM) [22] network as a sequential encoder, which processes the past 10 finger states to generate a sequential representation. The final output from the LSTM is then 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 [31] optimizer.

2) *Experiments*: We explore different architectures, input/output types, and learning parameters to optimize performance. In this section, we present our results on different data-driven controllers and the evaluations we conducted. Our experiments aim to answer the following questions: (a) How do architecture choices and learning parameters affect controller performance? (b) How do the trained controllers transfer to new hands?

We evaluate our approach on two separate test datasets:

- **Human Validation**: A human wears the MANUS glove, moves their hand into various poses, and the keypoints are recorded during this motion.

- **Robot Validation**: The MANUS glove is attached to the robot, and the robot’s fingertips are manually moved while recording data from the MANUS stream, as illustrated in the top row of Fig. 9.

During evaluation for both datasets, we attach the glove to the robot and replay the recorded keypoints by moving the robot motors to the positions predicted by the controllers. We then record the robot’s keypoints and measure accuracy by comparing the reproduced fingertip positions against the originally saved ones.

### How do architecture choices and learning parameters affect controller performance?

After we collected data as described in Section V-A, we experiment on four different ways to train controllers.

- **LSTM+MLP (RUKA)**: Controller training used in RUKA as described in Sec. V-B1.
- **MLP**: Instead of integrating the past observation we directly input the current fingertip position and predict the motor position.
- **k-Nearest Neighbor (k-NN)**: Instead of learning a parametric model, we use k-nearest neighbor retrieval on the desired fingertip positions and apply the mean of those neighbor’s motor positions.
- **Search-Based**: Instead of training an inverse model (i.e. predicting motor positions given finger state) we train a forward model, do a search on all possible motor positions, and apply the one that will give the closest predicted finger state.

Table IV presents the mean error of the thumb fingertip across different axes. We observe that while all methods perform relatively well on the Robot Validation set—reporting a maximum error of 5mm—on the Human Validation set, the



In-Hand Teleoperated Tasks



Teleoperated Tasks



Fig. 10: Teleoperated tasks made possible with RUKA. The top three rows showcase teleoperation of in-hand grasping tasks, while the remaining rows demonstrate arm-and-hand teleoperated tasks.

TABLE IV: Comparison of the performance of different architectures. The second column reports the mean error in Robot Validation, while the third column presents the mean error in Human Validation. Error values are computed over three positional axes for the Thumb finger.

Method	Error in Robot Val (cm)			Error in Human Val (cm)		
Search Based	0.55	0.59	0.48	2.03	2.00	2.01
MLP	0.53	0.56	0.43	1.80	1.90	2.00
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

MLP and search-based baselines perform significantly worse, with errors reaching up to 2cm on certain axes. Although the k-NN baseline outperforms the RUKA architecture on the Robot Validation set, it underperforms on the Human Validation set. We attribute this to the out-of-distribution nature of human fingertip positions, which the k-NN baseline fails to generalize to effectively.

#### How do the trained controllers transfer to new hands?

Since RUKA is a tendon-driven hand that can be assembled by different users, the way its tendons are tensioned may vary from user to user. To maintain the usability of RUKA across different builds, the learned controllers should be able to generalize to these variations.

To enable this, we implement an auto-calibration script that estimates the maximum range of each motor given the current tendon tensioning. The script performs a binary search over motor positions, slowly pulling the tendons to find the point at which the finger is as curled as possible for that actuation. To ensure consistency across runs, we executed this script multiple times on the same hand and find that the variation in the maximum motor ranges found across 10 runs as 0.5 degrees.

After calibration, to evaluate how well the learned controllers generalize to a differently built hand, we repeated the same experiments from the Robot Validation set using a newly assembled and calibrated hand. The average difference in fingertip positions between the original and new hands was 3mm across different axes. We hope that these experiments will help ensure consistency across different RUKA hands, making the system more accessible to a wider range of research labs.

## VI. APPLICATIONS OF RUKA

### A. Teleoperation

Using the controller described in Section V, we teleoperate RUKA to collect demonstrations for various dexterous tasks (Fig. 10). We use a MANUS glove for hand teleoperation and attach ArUco markers to it for arm control. By optimizing the execution time of the trained controller, ArUco detection, and RUKA’s control frequency, we achieve a teleoperation rate of 25Hz. If users prefer not to use the learned controllers and instead directly control the hand via motor positions, the system supports control at 40Hz.

Similar to the tests described in Section V, we input the fingertip position to the thumb controller and joint angles for the remaining fingers.




Fig. 11: Rollouts of tasks learned with HuDOR [19] on RUKA. For each task, the bottom row shows the open-loop rollout of the human trajectory, while the top row shows the final rollout after online finetuning, highlighting the importance of the training process.

### B. Policy Learning

To demonstrate how RUKA can be used for policy learning approaches, we deploy HuDOR [19], 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*. Rollouts of the policies trained using this method are shown in Fig. 11. Unlike the original HuDOR framework, for both tasks we learn a residual policy that adjusts motor positions instead of fingertip positions. On average, each policy is trained for 40 episodes, with training converging in approximately 45 minutes.

## VII. DISCUSSION

In this work, we introduced RUKA, a robotic hand designed to address key challenges in existing hand designs while maintaining cost efficiency.

**Limitations:** Currently, training the model and capturing hand poses relies on the Manus glove, which provides accurate data but introduces a cost barrier for full replication and scalability. Exploring alternative motion capture techniques could improve accessibility. The current design of RUKA lacks tactile sensing, which may limit its performance in complex dexterous tasks. Integrating touch sensors could provide crucial feedback for more precise and adaptive manipulation. To maintain the current design’s simplicity, it also lacks abduction at the MCP joints, which may limit the hand’s ability to perform dynamic tasks. In future versions, we aim to incorporate this functionality using a compliant ball joint.

**Hardware Failure Modes:** There are two primary failure modes of the hardware. First, since we use cheap Dynamixel motors with plastic gears, they wear out over time and may



require replacement. The hand’s mechanism is independent of the base actuators, so this can be avoided by using metal-gearred motors, albeit at a higher price point. Secondly, while the 3D-printed parts are robust to normal usage, collisions of the hand with obstacles tend to damage the MCP joint. For both failure modes, the replacement time is under 20 minutes given the simple design and easy access for hardware modifications. The design of RUKA is highly customizable, so if users require a more durable and higher-quality hand, they can adapt the design to use higher-cost components.

#### ACKNOWLEDGEMENTS

We thank Raunaq Bhirangi, Siddhant Halder and Venkatesh Pattabiraman for valuable feedback and discussions. This work was supported by grants from Honda, Hyundai, NSF award 2339096, and ONR award N00014-22-1-2773. LP is supported by the Sloan and Packard Fellowships. NXB is supported by the Fannie and John Hertz Foundation Fellowship.

#### REFERENCES

- [1] allegro. Allegro, 2010. URL <https://www.allegrohand.com/ah-v4>.
- [2] Sridhar Pandian Arunachalam, Irmak Güzey, Soumith Chintala, and Lerrel Pinto. Holo-dex: Teaching dexterity with immersive mixed reality, 2022.
- [3] bambux1c. Bambu 3d printer, 2010. URL <https://us.store.bambulab.com/products/x1-carbon>.
- [4] 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.
- [5] Raunaq Bhirangi, Abigail DeFranco, Jacob Adkins, Carmel Majidi, Abhinav Gupta, Tess Hellebrekers, and Vikash Kumar. All the feels: A dexterous hand with large area sensing. *arXiv preprint arXiv:2210.15658*, 2022.
- [6] Patricia Capsi Morales, Cristina Piazza, Manuel Catalano, Giorgio Grioli, lisa schiavon, elena fiaschi, and Antonio Bicchi. Comparison between rigid and soft poly-articulated prosthetic hands in non-expert myo-electric users shows advantages of soft robotics. *Scientific Reports*, 11, 12 2021. doi: 10.1038/s41598-021-02562-y.
- [7] 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.
- [8] Yuanpei Chen, Chen Wang, Li Fei-Fei, and C. Karen Liu. Sequential dexterity: Chaining dexterous policies for long-horizon manipulation, 2023. URL <https://arxiv.org/abs/2309.00987>.
- [9] Yuanpei Chen, Chen Wang, Yaodong Yang, and C Karen Liu. Object-centric dexterous manipulation from human motion data. *arXiv preprint arXiv:2411.04005*, 2024.
- [10] E. Churchill, L. L. Laubach, J. T. Mcconville, and I. Tebbetts. *Anthropometric source book. Volume 1: Anthropometry for designers*. NASA, 1978. URL <https://ntrs.nasa.gov/citations/19790003563>.
- [11] clone. Clone hand, 2010. URL <https://clonerobotics.com/hand>.
- [12] dexhand. Dexhand, 2010. URL <https://www.dexhand.org/>.
- [13] Haritheja Etukuru, Norihito Naka, Zijin Hu, Seungjae Lee, Julian Mehu, Aaron Edsinger, Chris Paxton, Soumith Chintala, Lerrel Pinto, and Nur Muhammad Mahi Shafiullah. Robot utility models: General policies for zero-shot deployment in new environments. *arXiv preprint arXiv:2409.05865*, 2024.
- [14] 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.
- [15] filaflex. Filaflex, 2010. URL <https://recreus.com/gb/filaments/24-filaflex-foamy.html>.
- [16] 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.
- [17] 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.
- [18] Irmak Guzey, Ben Evans, Soumith Chintala, and Lerrel Pinto. Dexterity from touch: Self-supervised pre-training of tactile representations with robotic play, 2023.
- [19] 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.
- [20] Siddhant Halder, Jyothish Pari, Anant Rai, and Lerrel Pinto. Teach a robot to fish: Versatile imitation from one minute of demonstrations, 2023. URL <https://arxiv.org/abs/2303.01497>.
- [21] 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.
- [22] 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>.
- [23] Mary C. Hume, Harris Gellman, Harry McKellop, and Robert H. Brumfield. Functional range of motion of the joints of the hand. *The Journal of Hand Surgery*, 15(2):240–243, 1990. ISSN 0363-5023. doi: [https://doi.org/10.1016/0363-5023\(90](https://doi.org/10.1016/0363-5023(90)

- 90102-W. URL <https://www.sciencedirect.com/science/article/pii/S036350239090102W>.
- [24] inmoov. Inmoov, 2010. URL <https://inmoov.fr/inmoov-hand/>.
- [25] Aadithya Iyer, Zhuoran Peng, Yinlong Dai, Irmak Guzey, Siddhant Halder, Soumith Chintala, and Lerrel Pinto. OPEN TEACH: A Versatile Teleoperation System for Robotic Manipulation. *arXiv e-prints arXiv:2403.07870*, March 2024.
- [26] Atzori M Jarque-Bou NJ, Scano A and Müller H. Kinematic synergies of hand grasps: a comprehensive study on a large publicly available dataset. *Journal of NeuroEngineering and Rehabilitation*, 2019.
- [27] Kai Junge and Josie Hughes. Robust anthropomorphic robotic manipulation through biomimetic distributed compliance, 2024. URL <https://arxiv.org/abs/2404.05262>.
- [28] A. Kapandji. Cotation clinique de l’opposition et de la contre-opposition du pouce. *Annales de Chirurgie de la Main*, 5(1):67–73, 1986. ISSN 0753-9053. doi: [https://doi.org/10.1016/S0753-9053\(86\)80053-9](https://doi.org/10.1016/S0753-9053(86)80053-9). URL <https://www.sciencedirect.com/science/article/pii/S0753905386800539>.
- [29] 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.
- [30] 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.
- [31] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [32] 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.
- [33] manus. Manus haptic gloves, 2010. URL <https://www.manus-meta.com/>.
- [34] 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>.
- [35] 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.
- [36] Jyothish Pari, Nur Muhammad Shafiullah, Sridhar Pandian Arunachalam, and Lerrel Pinto. The surprising effectiveness of representation learning for visual imitation, 2021.
- [37] 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>.
- [38] Vatsal V. Patel, Minas V. Liarokapis, and Aaron M. Dollar. Open robot hardware: Progress, benefits, challenges, and best practices. *IEEE Robotics & Automation Magazine*, 30(3):123–148, 2023. doi: 10.1109/MRA.2022.3225725.
- [39] 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.
- [40] robotis. Dynamixel motors, 2010. URL <https://www.robotis.us/dynamixel-xc330-m288-t/>.
- [41] 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.
- [42] Cornelia Schlegenhaut, 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.
- [43] shadow hand. Shadow hand, 2010. URL <https://www.shadowrobot.com/dexterous-hand-series/>. Shadow Hand.
- [44] Kenneth Shaw, Ananye Agarwal, and Deepak Pathak. Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning. *Robotics: Science and Systems (RSS)*, 2023.
- [45] Yasunori Tshimitsu, 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.
- [46] Anand Vazhapilli Sureshbabu, Giorgio Metta, and Alberto Parmiggiani. A systematic approach to evaluating and benchmarking robotic hands—the ffp index. *Robotics*, 8(1), 2019. ISSN 2218-6581. doi: 10.3390/robotics8010007. URL <https://www.mdpi.com/2218-6581/8/1/7>.
- [47] 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.



- [48] 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.
- [49] 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.
- [50] Zhe Xu and Emanuel Todorov. Design of a highly biomimetic anthropomorphic robotic hand towards artificial limb regeneration. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3485–3492, 2016. doi: 10.1109/ICRA.2016.7487528.
- [51] 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.