Cross-Embodiment Robot Manipulation via a Unified Hand Action Space

Luis Felipe Casas[1], Robert Teal[1], Keval Shah[1], Abhijit Tadepalli[2], Wanxin Jin[2], Yu Xiang[1]
[1]Intelligent Robotics and Vision Lab, University of Texas at Dallas
[2]Intelligent Robotics and Interactive Systems Lab, Arizona State University

❑Can we find a unified action space for different robot hands?

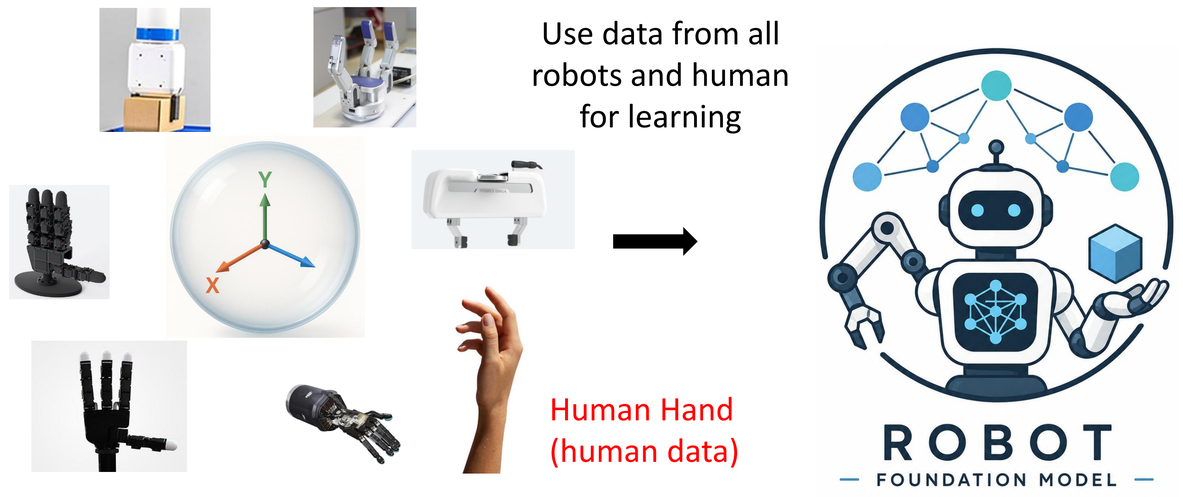

❑Previous Work: manual alignment of hands (retargeting)

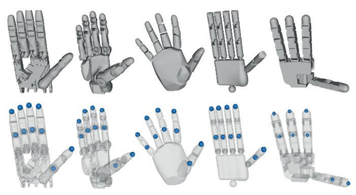

Learning Cross-hand Policies for High-DOF Reaching and Grasping (She et al., 2024)

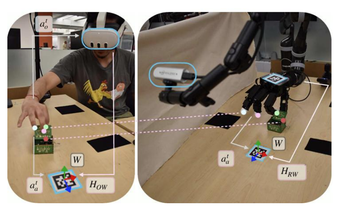

HuDOR, Guzey et al. NYU 2025
https://object-rewards.github.io/

- Manual mapping or hand-designed correspondence
- Hard to deal with different number of fingers
- Cannot handle unseen hands

❑Our idea: Let's use a sphere to align robotic hands

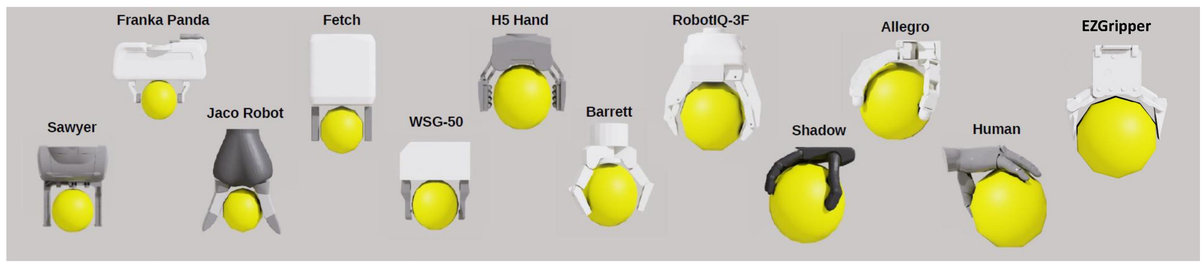

❑Automatic Sphere Creation for a Robotic Hand

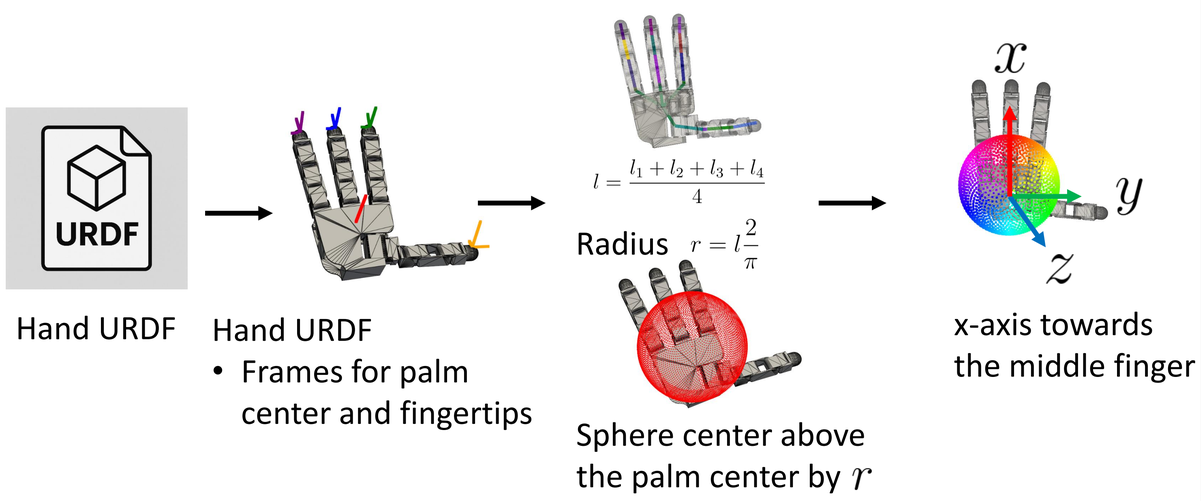

Hand URDF

Hand URDF
- Frames for palm center and fingertips

Sphere center above the palm center by r

x-axis towards the middle finger

❑Unified Hand Surface Correspondence

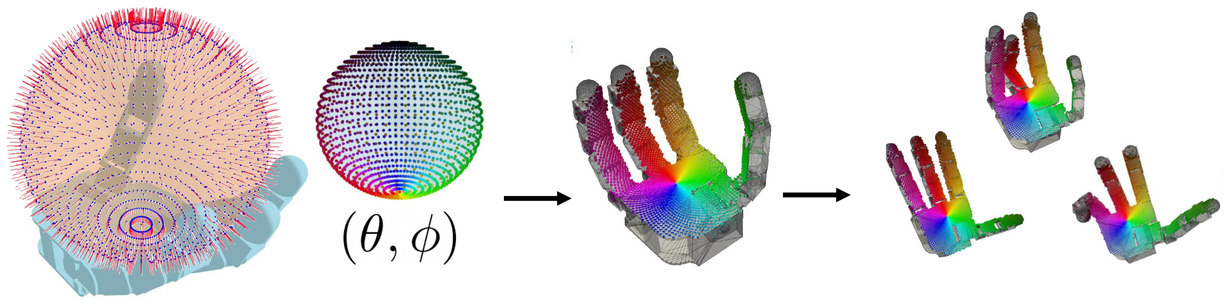

(a) Sample points on the sphere surface, compute spherical coordinates and normals of the sampled points

(b) Transfer the spherical coordinates to the robotic hand

(c) The spherical coordinates remain the same in different hand configurations

❑Sphere Deformation Action Space

Define several "driving planes"

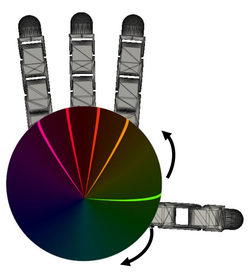

Define several "driving vectors" on each driving plane

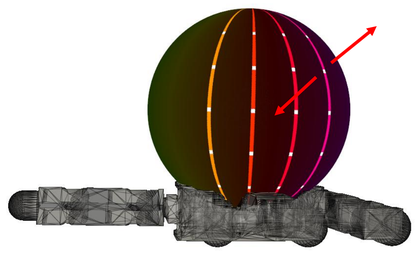

❑Cascade Inverse Kinematics

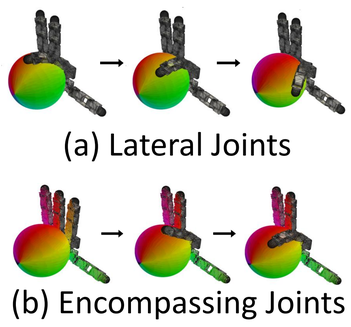

(a) Lateral Joints

(b) Encompassing Joints

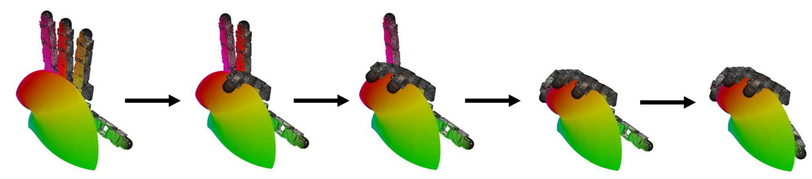

Input: a deformed sphere

Output: hand joint configuration

(c) Cascade inverse kinematics

❑Experiments: in-hand manipulation

- Task: repose a cube to a target orientation
- 10 consecutive reposing within 30 seconds
- RL training in Isaac Lab with PPO and our sphere controller

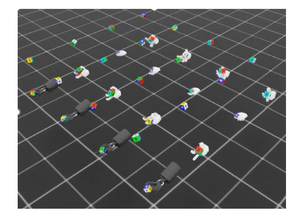

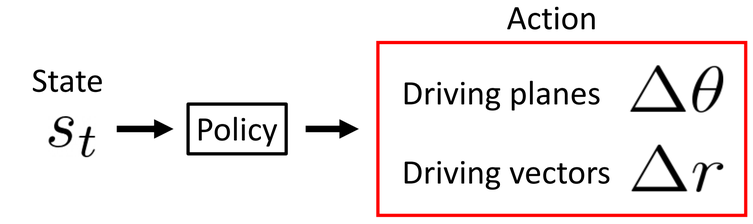

❑Simulation Results

Test Hand	Single-Hand	Joint Control	Multi-Hand	Zero-shot
Allegro	**99.1 / 9.6 ± 1.7**	98.5 / 9.2 ± 2.2	99.2 / 9.5 ± 1.9	95.3 / 7.7 ± 3.4
LEAP	**99.7 / 9.8 ± 1.1**	98.6 / 9.3 ± 1.2	99.1 / 9.5 ± 1.9	95.5 / 7.7 ± 3.5
Shadow	**99.3 / 9.6 ± 1.6**	98.0 / 9.1 ± 1.9	98.7 / 9.2 ± 2.3	85.7 / 4.4 ± 3.7
MANO	**99.8 / 9.9 ± 1.0**	99.6 / 9.8 ± 1.4	99.5 / 9.8 ± 1.2	98.1 / 8.9 ± 2.6

- Single-Hand: trained only on the target hand
- Joint Control: baseline joint-space controller trained on the target hand
- Multi-Hand: trained on all four hands
- Zero-shot: trained on all hands except the target hand

Test Hand	Train: Shadow + MANO (5-finger)	Train: Allegro + LEAP (4-finger)
Allegro (4F)	66.2 / 1.9 ± 1.9	**99.7 / 9.8 ± 1.2**
LEAP (4F)	80.8 / 3.7 ± 3.6	**99.8 / 9.9 ± 0.98**
Shadow (5F)	**98.6 / 9.3 ± 2.1**	83.2 / 4.0 ± 4.0
MANO (5F)	**99.7 / 9.8 ± 1.3**	95.0 / 7.6 ± 3.5

Cross-morphology generalization across robotic hands with different numbers of fingers.

Target Hand	Zero-shot	+500 Iter
Allegro	95.3 / 7.7 ± 3.4	**96.3 / 8.1 ± 3.3**
LEAP	95.5 / 7.7 ± 3.5	**96.2 / 8.0 ± 3.3**
Shadow	85.7 / 4.4 ± 3.7	**95.8 / 7.8 ± 3.5**

Fast adaptation from MANO to unseen robotic hands. The policy is first trained on MANO and finetuned on the target hand for only 500 iterations.

❑Real-World Results with Sim-to-Real Transfer

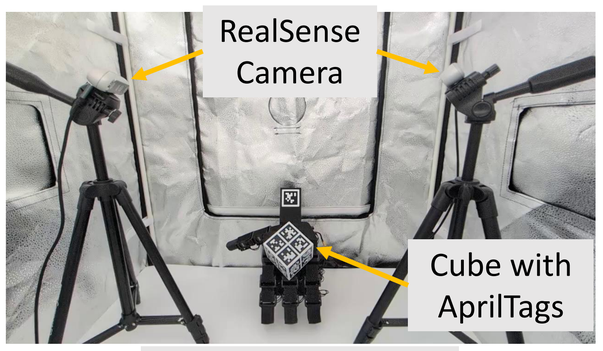

LEAP hand (4-finger)

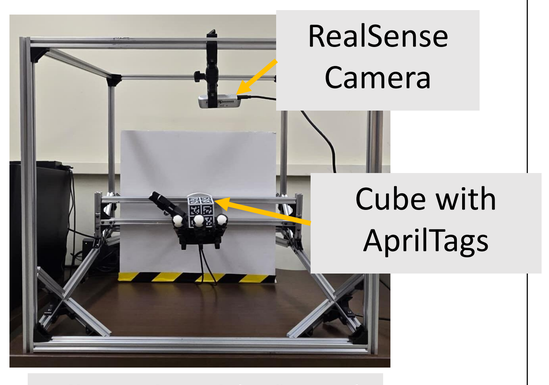

Allegro hand (4-finger)

LEAP Hand

Method	1	2	3	4	5	6	7	8	9	10	MEAN
Baseline (Joint Control)	0	0	1	2	0	0	0	1	2	0	0.6
UHAS (Zero-Shot)	2	4	2	0	1	0	0	0	0	0	0.9
UHAS (Trained on Multi-Hand)	3	0	1	0	0	5	0	0	0	2	1.1
UHAS (Trained on LEAP Hand)	0	2	1	0	1	6	2	1	2	5	**2.0**

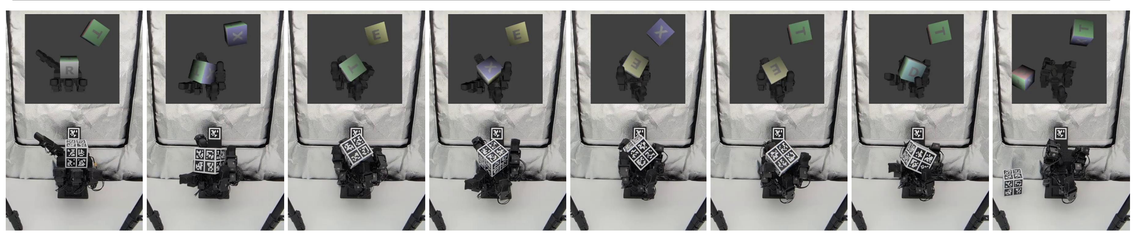

Allegro Hand

Method	1	2	3	4	5	6	7	8	9	10	MEAN
UHAS (Zero-Shot)	0	1	1	1	0	1	0	0	2	2	0.8
UHAS (Trained on Multi-Hand)	3	0	8	1	2	1	1	2	1	2	**2.1**
UHAS (Trained on Allegro Hand)	4	0	2	2	2	3	4	1	0	3	**2.1**

❑Acknowledgement

NSF
PENG
NVIDIA.