






The missing ingredients in scaling dexterous manipulation data

Akash Sharma

Research Scientist, Amazon FAR



Humans are incredibly **dexterous**!



Where is robot dexterity?




Object pick and place
[Figure AI '24]



Shoe pick and place
[Zhang et al., Alpha Unleashed '24]



Cloth folding
[Physical Intelligence '24]




Socks manipulation
[Sunday Robotics '25]


Agents can train dexterity




Pen spinning



Cube, single hand



Hammer spin






Drawing

PPO in 4,096–8,192 parallel simulations · Isaac Sim and MuJoCo · Sharpa and Wuji hands

Clips: @jianglong_ye (hands), @dimentary (drawing) on X

Still in (game-able) simulation, real world requires contact reasoning.

How is robot dexterity learned today?



Multi-fingered, contact-rich, in-hand manipulation — the fingers move the object.

**Learning in
Simulation**

**Learning from
Demonstrations**


Learning in simulation

Dexterity from experience


via reinforcement learning in simulation



Rotatelt [Qi et al. '23]




AnyRotate [Yang et al. '24]




In-hand translation [Yin et al. '24]

but only simple observations

Tactile sim is slow and needs approximations




Point-of-contact approximation



Binary contact approximation

The missing ingredient: **relax zero-shot**

Learn tactile dexterous policies
without ever simulating tactile sensors



PTLD · CoRL 2026

Rosy Chen, Mustafa Mukadam, Michael Kaess, Tingfan Wu, Francois R. Hogan, Jitendra Malik,
Akash Sharma

Relax the zero-shot requirement

Train in simulation, deploy


Zero-shot sim-to-real

Deploy with sensors you can simulate

Proprioception



RMA - Proprioception




HIM - Scan points

Train in simulation, use real robot experience

+ real data

Deploy with difficult to simulate sensors

Touch



Sparsh - Visuo-tactile data

A little real data lets the policy use **much richer sensors**


Privileged sensors

Information you can measure in an instrumented setup,
but not on the deployed robot

We use **object pose** for rotation and reorientation


In simulation

Free: read it from the simulator



In the real world

An instrumented cell: 4 cameras + markers




Step 1 · Teacher in simulation

Reinforcement learning (PPO)
in simulation

Policy sees privileged sensor: **object pose**,
plus friction, mass and scale

A **stronger teacher**
than a proprioception-only policy




Teacher reorienting objects to goal poses

Step 2 · Collect real data

Run the teacher on the real hand


No teleoperation, no kinesthetic teaching



Object pose comes from the instrumented cell


Record touch $\leftrightarrow$ latent pairs

Record real tactile states + record teacher predicted state



Tactile readings and the teacher's latent

Step 3 · Distill to touch



Only the tactile encoder is trained, on the step 2 data. At deploy time the cameras are not required.



Deployed on touch + proprioception: robust to weight, slip and wrist orientation

PTLD in the real world

+182%

Time to fall In-hand rotation vs. proprioception



Proprioception
drops it in ~13 s




PTLD
rotates for 2+ min

+57%




Reorientation goals reached, vs. proprio only



Touch + proprioception, reorienting to goal poses



How is robot dexterity learned today?



Multi-fingered, contact-rich, in-hand manipulation — the fingers move the object.

**Learning from
Simulation**

**Learning from
Demonstrations**

Learning from Demonstrations


Via teleoperation



DexterityGen [Yin et al., '25]

No haptic feedback to match forces

Via kinesthetic teaching



DexForce [Chen et al., '25]

Humans only have two hands · demonstrations need inpainting

Via handheld grippers




UMI [Chi et al., 24]

Embodiment gap · demonstrations need inpainting

Every approach requires the human to act like a robot.

The missing ingredient: **leverage human play**


Learn tactile dexterous policies
without ever simulating tactile sensors



PTLD · CoRL 2026

Rosy Chen, Mustafa Mukadam, Michael Kaess, Tingfan Wu, Francois R. Hogan, Jitendra Malik,
Akash Sharma





Learn dexterous policies
purely from human play




DSP · In submission

Botao He, Yiming Dou, Jitendra Malik, **Akash Sharma**, Ruoshi Liu

Dexterity from Supervised Play




From minutes of
goal-agnostic
Human Play



To goal-conditioned
Dexterity

Going from play, to dexterity



1 Minutes of Supervised play -> Play Policy


2 Unlimited Autonomous play -> Play Dataset

3 World modeling from play data

4 Model Predictive Control (MPC)

5 Policy distillation -> Fast goal-conditioned policy

Minutes of Human Play



15 min
of goal-free human play



Behavior
Cloning

Play Policy


1

2

3

4

Unlimited Autonomous Play



3h
of auto play



Play Dataset


1

2



3

4

World Modeling from Play Data



— Ground truth — WM prediction




1

2

3

4

Model Predictive Control



1 Propose

1,024 chunks from the play policy

2 Predict

outcomes in the world model

3 Score

distance to the goal

4 Execute

best prefix, then replan



1

2

3

4

Real-World Reorientation




68% SR@23°


Matches the sim performance

Emergent Dexterous Skills


Yaw Rot, Thumb Tip




Edge Flip




Fine Adjustment




Yaw Rot, Thumb Root





Corner Flip



Composite Actions



Simulated Reorientation vs. Baselines



91.4%

SR@0.6 rad vs. 53.9% hindsight BC


490x

less interaction data than PPO

15 Hz

runs in real time on single GPU

Going from play, to dexterity



1 Minutes of Supervised play -> Play Policy

2 Unlimited Autonomous play -> Play Dataset

3 World modeling from play data

4 Model Predictive Control (MPC)

5 Policy distillation -> Fast goal-conditioned policy

The missing ingredients

Real-world distillation · PTLD

Sim RL + rollouts in the
instrumented cell

PTLD closes the **sensing gap**

Play · DSP

3 h autonomous play,
1.3M self-play labels


DSP makes **data easier** for ergodic world models

Robot dexterity is hard


What is still missing?

A corpus of **dexterous human-object interaction**


Collected with **high-accuracy states**, including **tactile**




(1) Functional use




(2) Grasp type transition



(3) In-hand Rotation




(4) In-hand Translation




What is still missing?



An approach for **high fidelity dexterous**
****haptic retargeting****



From various sources:
mocap, tactile gloves, egocentric video




Thank you!




Sparsh

C. Higuera*, A. Sharma* et al., [CoRL '24]




Sparsh-skin

A. Sharma, et al., Sparsh-skin [CoRL '25]




Sparsh-X

C. Higuera*, A. Sharma*, T. Fan* et al., [CoRL '25, Oral]




PTLD

R Chen* et al., [CoRL '26]



DSP



Questions?