

Tactile Sensing in Robotics

Binghao Huang @ Columbia

Today's Roadmap

1

Why Touch?

What does touch tell us that vision can't, for humans and for robots?

2

Today’s Tactile Sensing

How do tactile sensors work, how do you choose one, and what does it take to build one?

3

Learning with Touch

How do we fuse touch with vision in a policy, then scale it up with data and simulation?

Please interrupt with questions at any time!

Touch Matters in Everyday Human Tasks

Assemble Lego

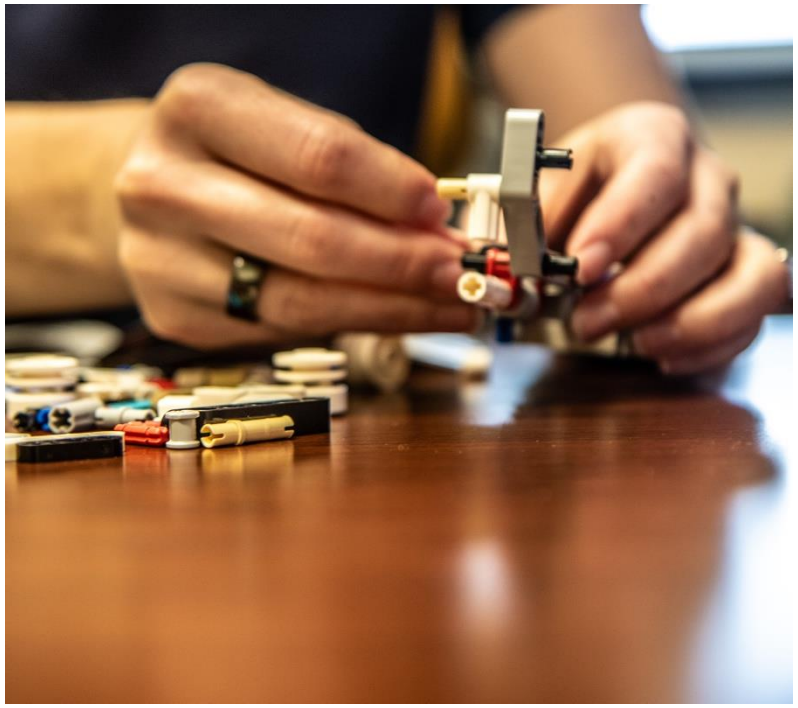

Pick Strawberries



Use Lab Dropper



Fine alignment

Touch Matters in Everyday Human Tasks

Assemble Lego

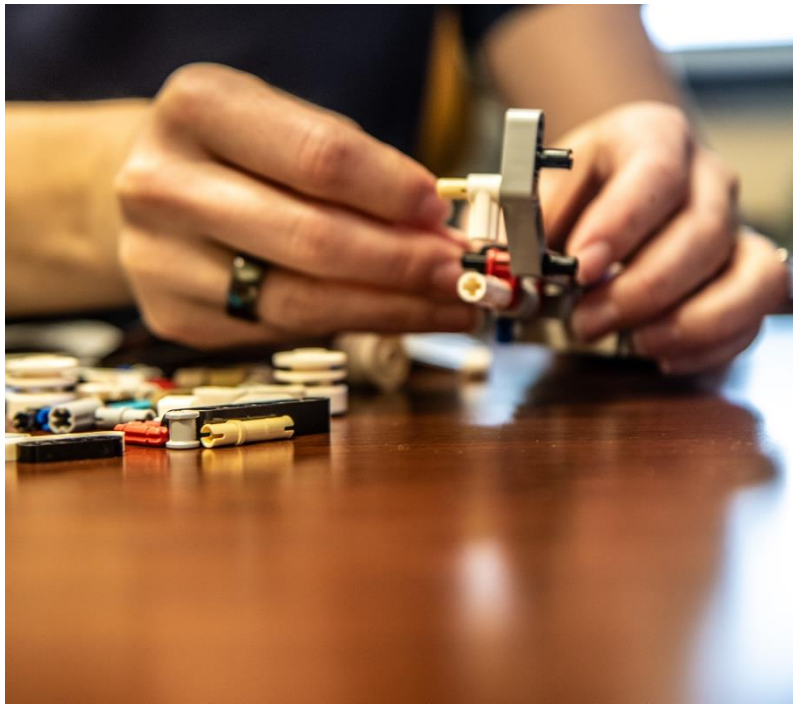

Fine alignment

Pick Strawberries



Handling fragile objects

Use Lab Dropper



Touch Matters in Everyday Human Tasks

Assemble Lego

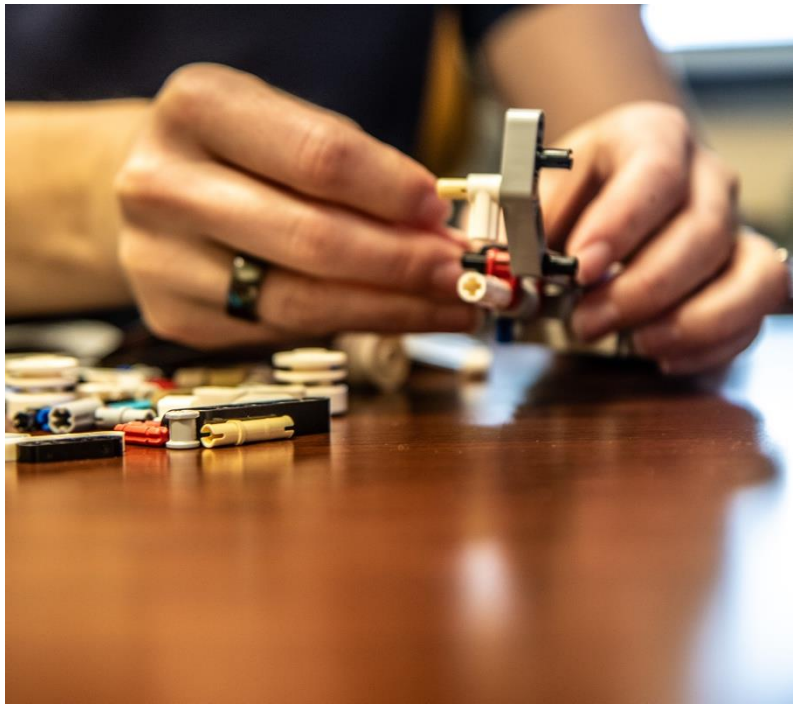

Fine alignment

Pick Strawberries



Handling fragile objects

Use Lab Dropper



Precise force control

How Do Humans Combine Vision and Touch?

TASK

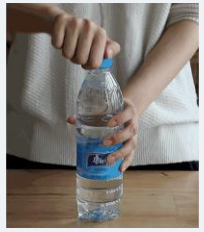

Open a cap



Zip a jacket

① SENSE

VISION
colors, edges, shapes
global; before contact

TOUCH
pressure, texture, slip
local; only at contact

② EXTRACT WHAT MATTERS

	Pose	Geometry	Force	Friction
Vision	●	●	○	○
Touch	○	○	●	●

● main source ○ can also estimate

③ ACT

ACTION
π picks the next move

act → make new contact → sense again: a closed loop

What Happens When Humans Lose Touch?



A CLASSIC EXPERIMENT

Same person, same task: pick up and strike a match.

With anesthetized fingertips, **vision is intact**, yet the movement becomes slow and clumsy.

Vision alone is not enough for dexterous manipulation.

Video: Roland Johansson lab, Umeå University

Do Robots Also Need Touch?



Does Touch Help Robots, Too?

Same robot, same task, same policy architecture:

trained **with** vs. **without** tactile input.

> **While you watch:** when exactly does the vision-only policy fail, and why?

Let's look at some videos

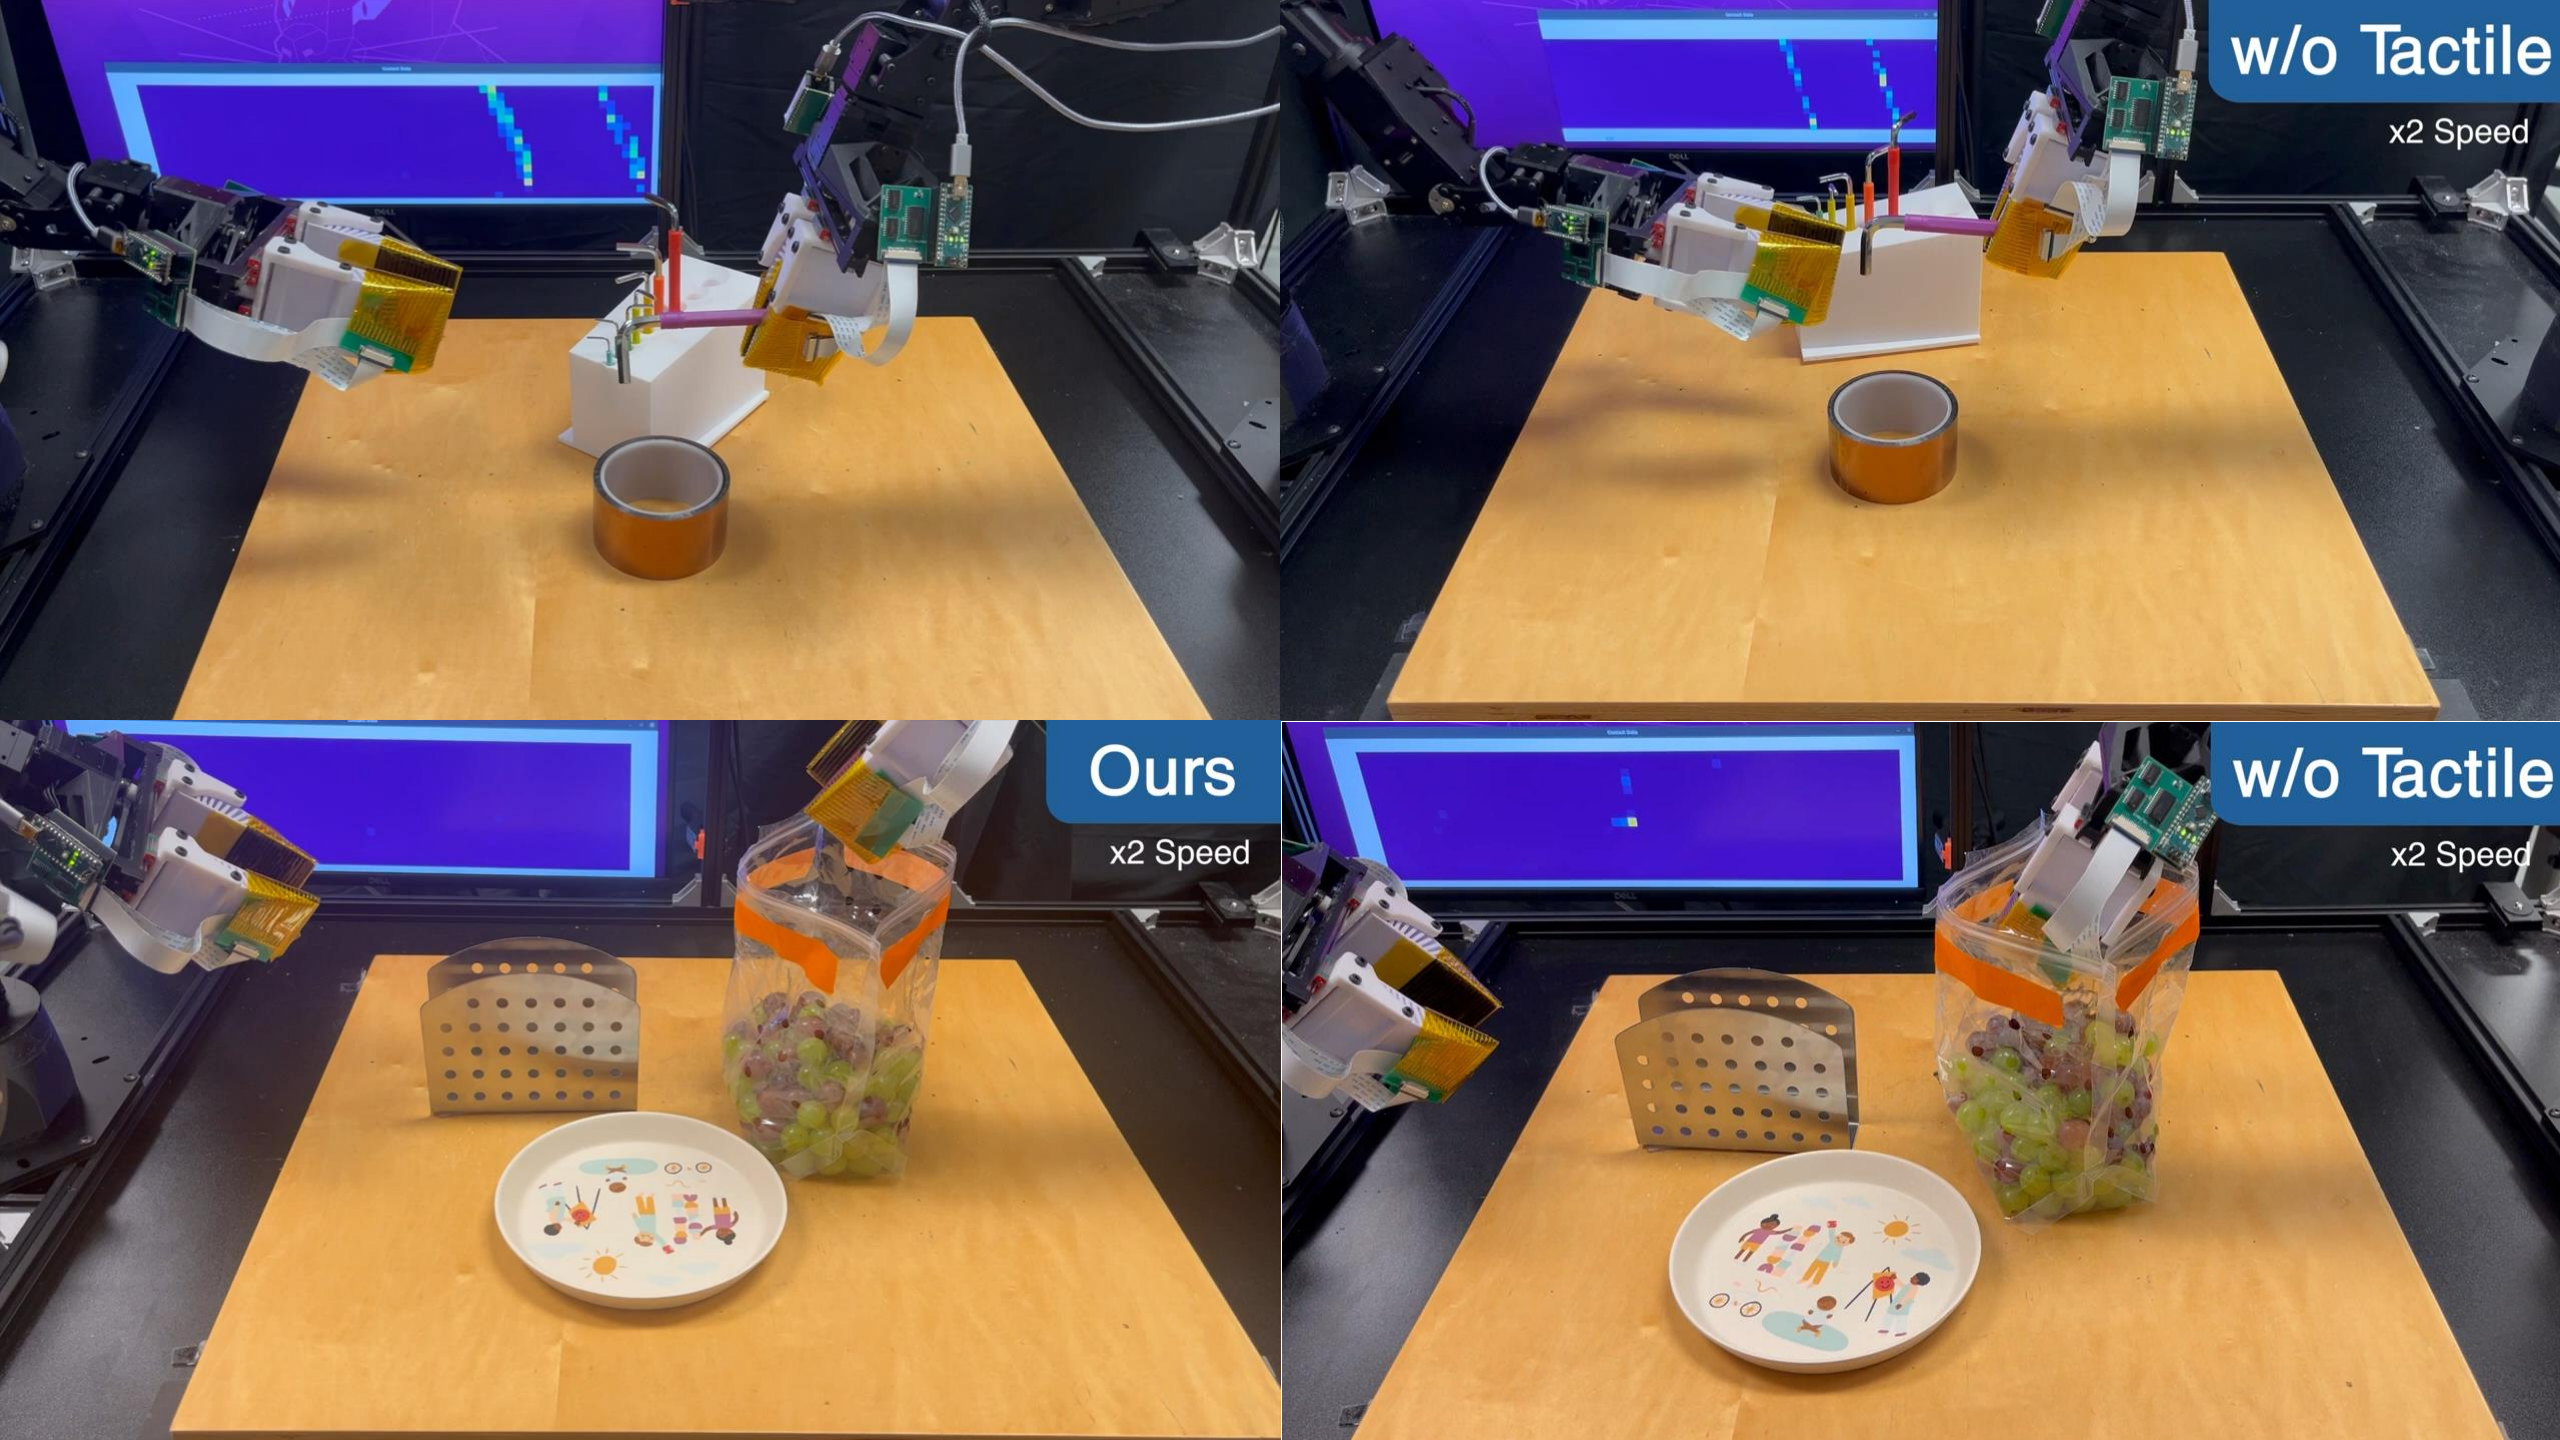
w/o Tactile
x2 Speed
Ours
x2 Speed
w/o Tactile
x2 Speed

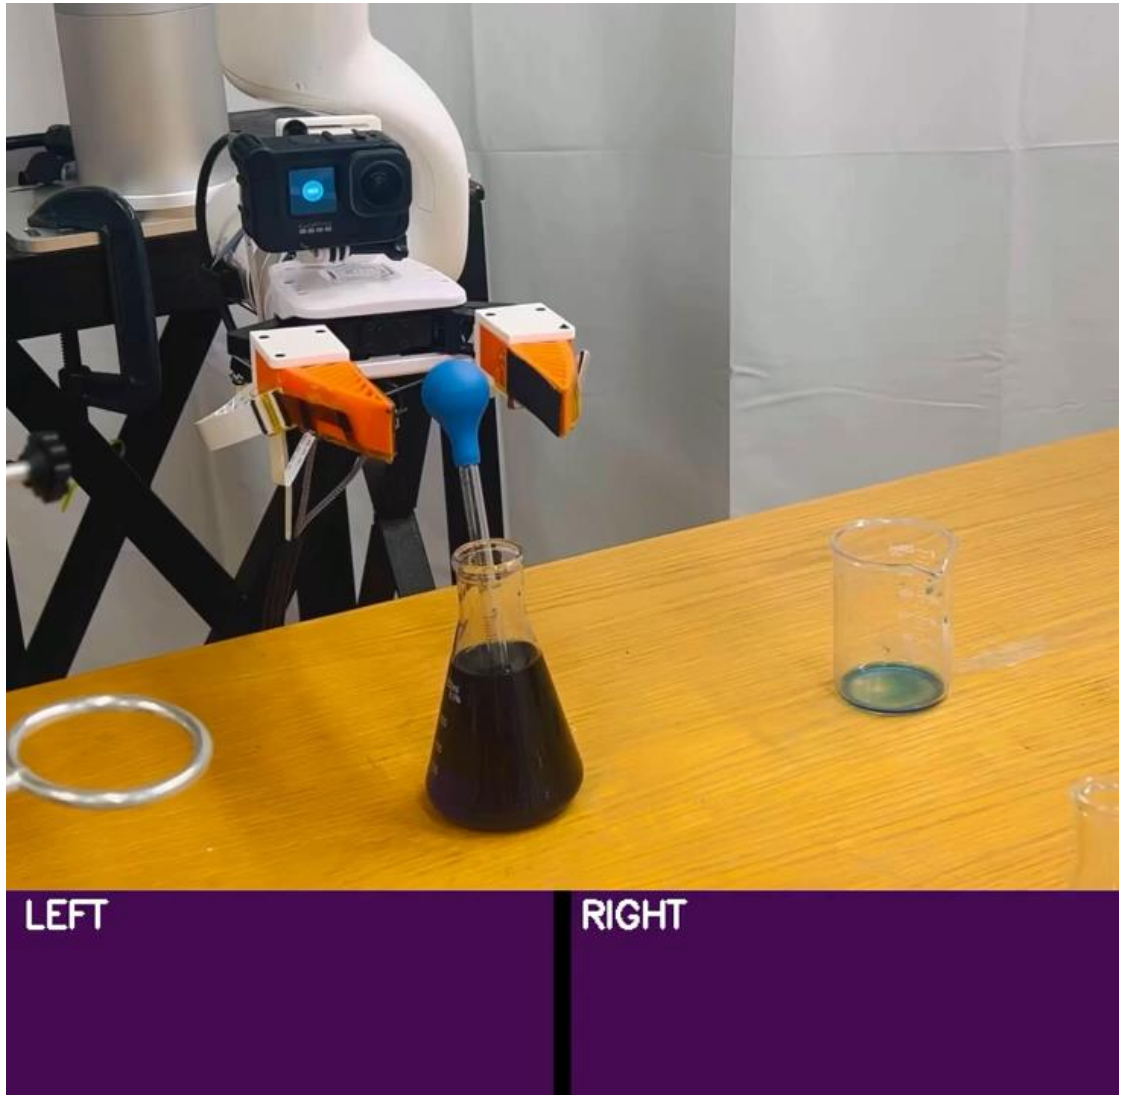

w/ Touch

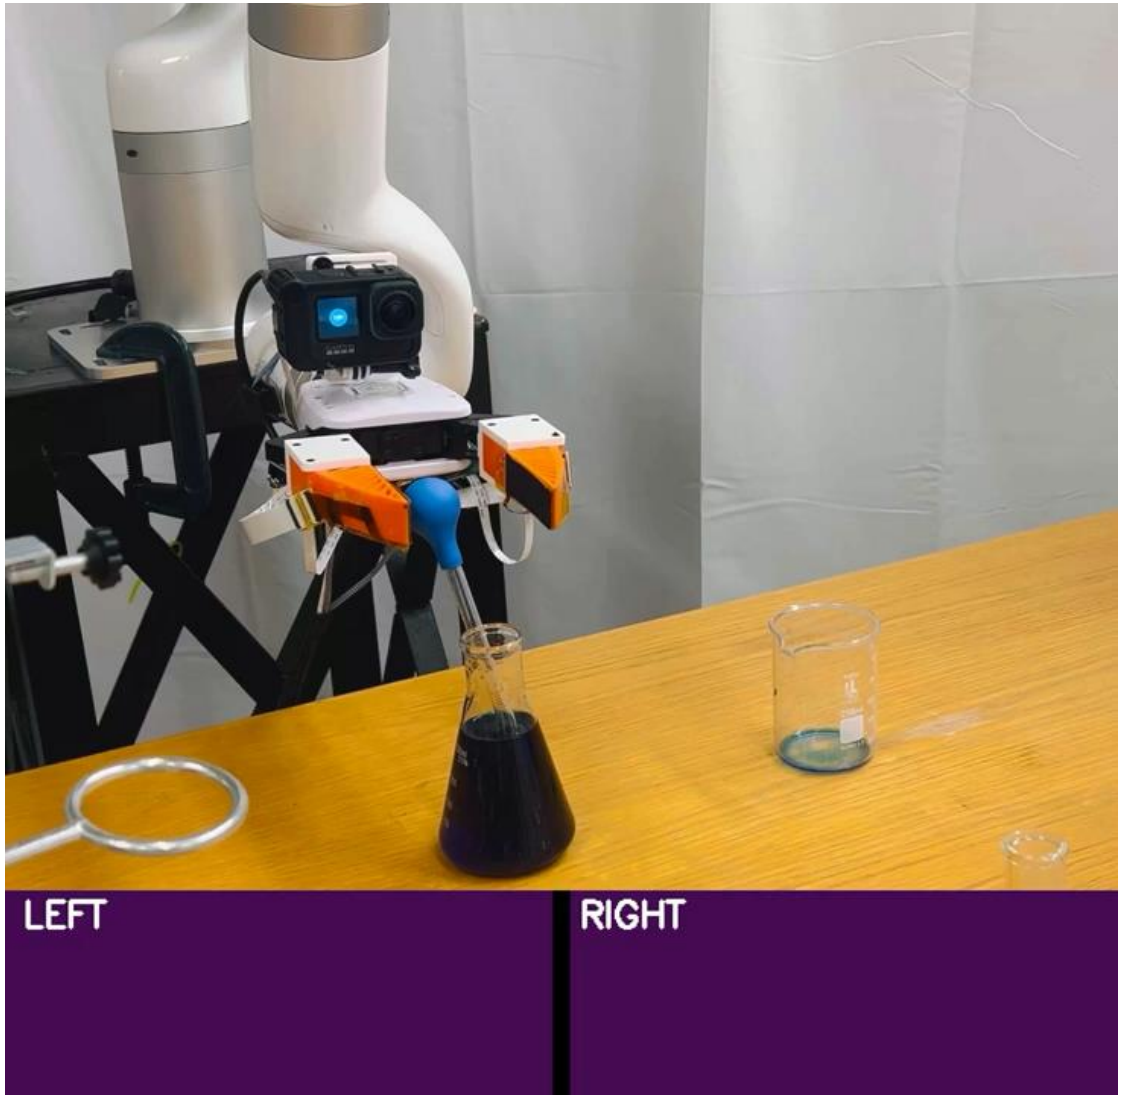

w/o Touch

So Why Isn't Touch on Every Robot?

If touch is this useful, what is stopping us?



Take a guess!

what makes touch hard to put on a robot, and hard to use once it is there?

So Why Isn't Touch on Every Robot?

HARDWARE BARRIERS → Part 2

No standard sensor

Many principles, no clear winner yet

Cost & accessibility

Expensive, hard to buy or customize

Reliability & robustness

Skins wear out; signals drift and get noisy

So Why Isn't Touch on Every Robot?

HARDWARE BARRIERS → Part 2

No standard sensor

Many principles, no clear winner yet

Cost & accessibility

Expensive, hard to buy or customize

Reliability & robustness

Skins wear out; signals drift and get noisy

LEARNING BARRIERS → Part 3

Hard to fuse with vision

Low-dimensional, local, only active at contact

Scarce data

No internet-scale tactile data to learn from

Hard to simulate

Contact physics → large sim-to-real gap

Today we tackle both: first the hardware, then the learning.

Roadmap

✓

Why Touch?

What does touch tell us that vision can't, for humans and for robots?

2

Tactile Sensing Hardware

NOW

How do tactile sensors work, how do you choose one, and what does it take to build one?

3

Learning with Touch

How do we fuse touch with vision in a policy, then scale it up with data and simulation?

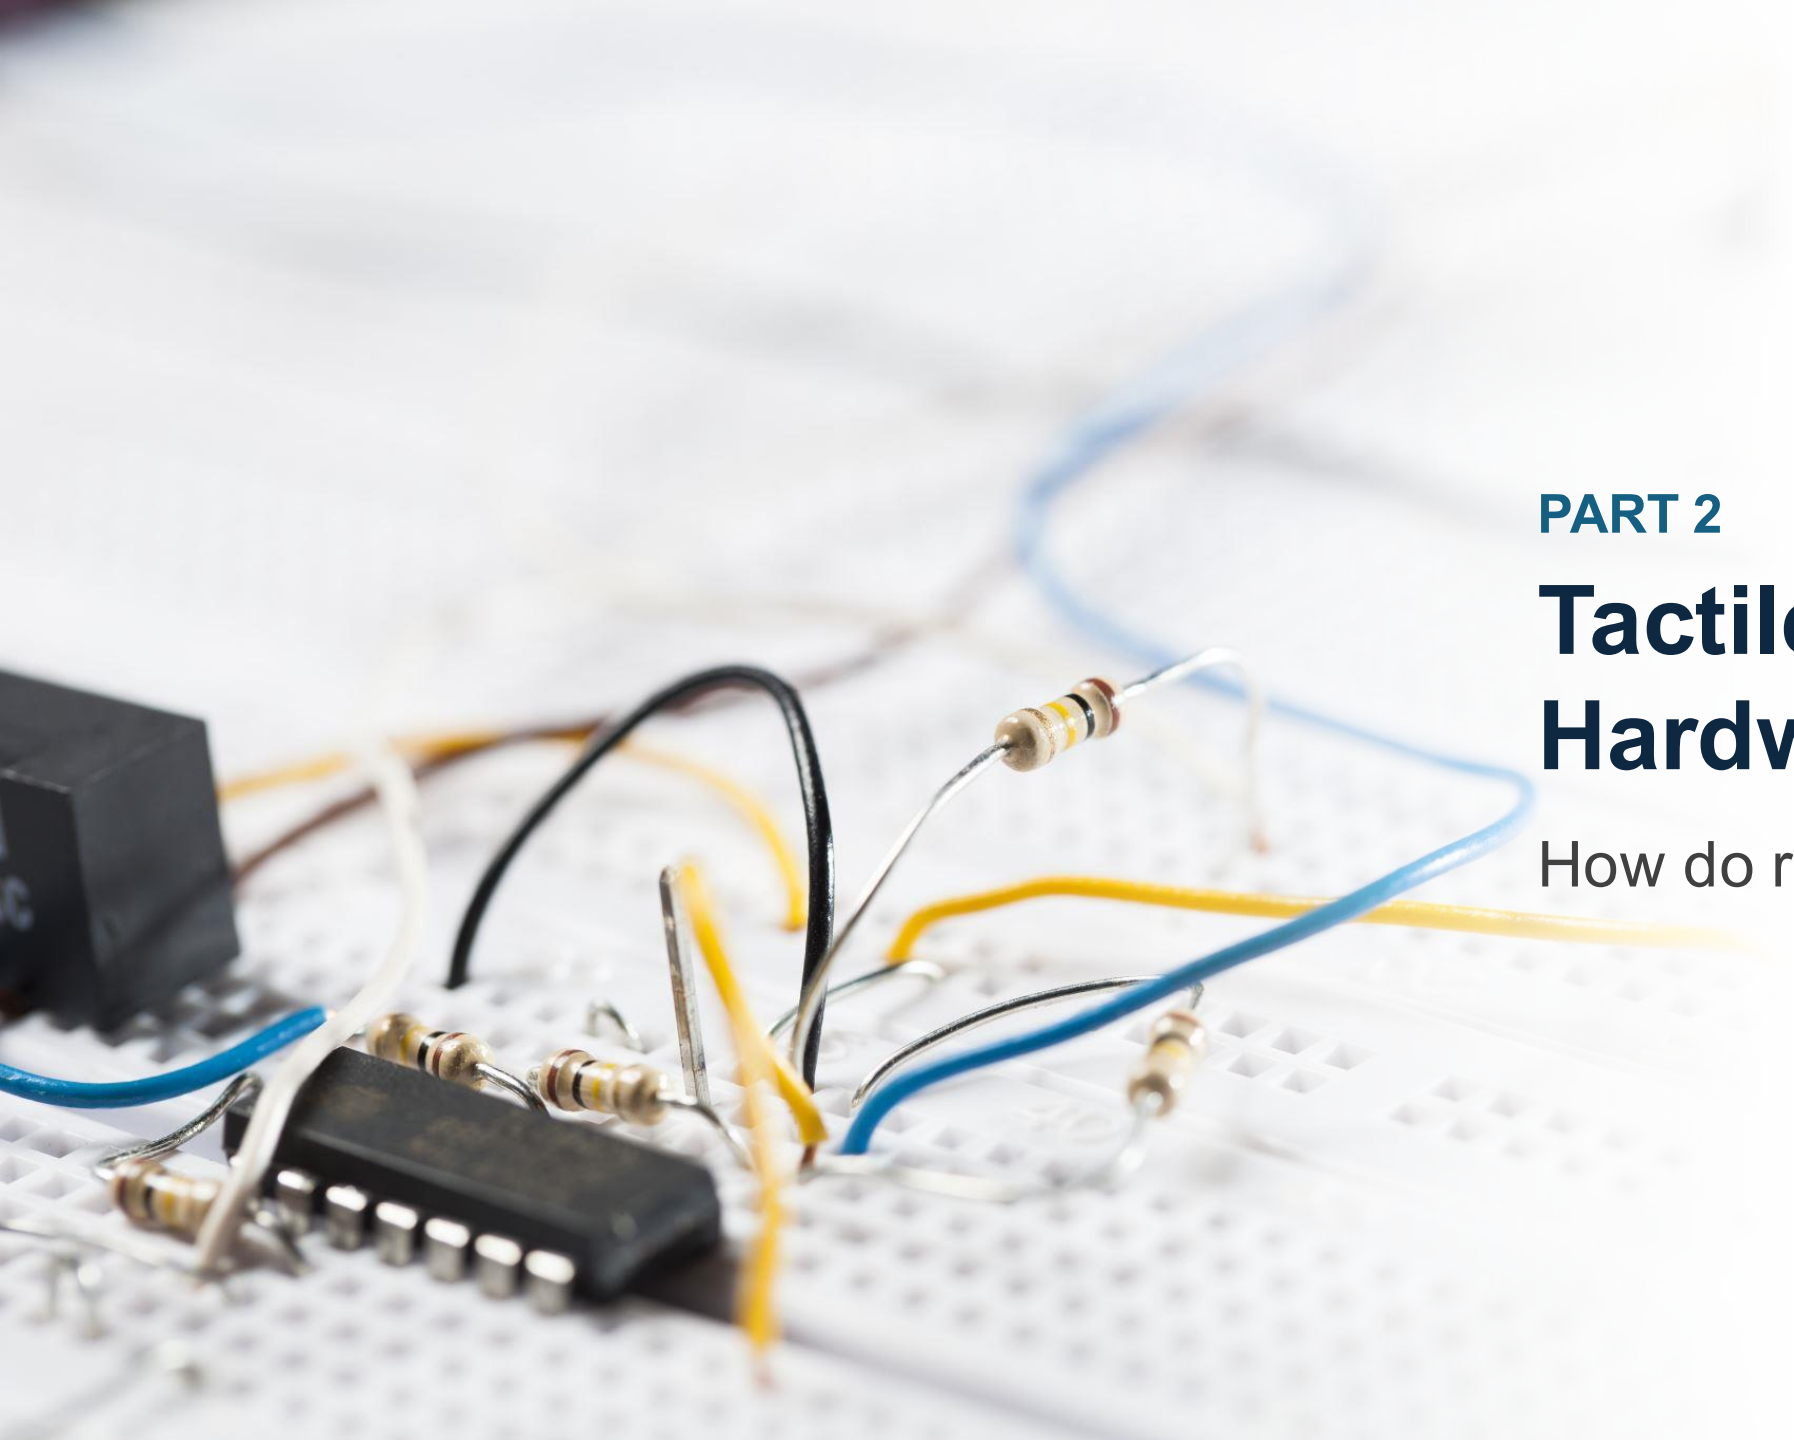

PART 2

Tactile Sensing Hardware

How do robots feel?

Tactile Sensors in Industry Today

Dexterous hands and grippers now ship with touch, and every sensing principle shows up.

OPTICAL

Camera under a soft, reflective skin



Sharpa
Wave W02 hand
Vision-based fingertips, 1 mm res.

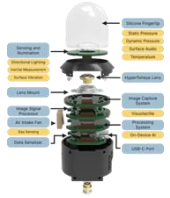

Meta × GelSight
Digit 360
~8.3 M taxels, ~1 mN



Daimon (戴盟)
DM-Tac W2
640×480 camera, 120 Hz

MAGNETIC

Magnetic skin over magnetometer chips



PaXini (帕西尼)
PX-6AX GEN3
0.1 mm res., 1 kHz



XELA Robotics
uSkin on Allegro Hand
368 three-axis taxels

PIEZORESISTIVE

Pressure-sensitive film arrays

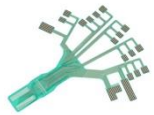

Tekscan
Grip System
18 regions, up to 750 Hz



Inspire (因时)
RH56E2-T1 hand
17 resistive tactile sensors



Wuji (舞肌)
Wuji Glove*
24×31 pressure matrix

CAPACITIVE

Electrode gap changes under load



Pressure Profile Systems
RoboTact
12-element array



Tashan (他山)
TS-F fingertip
0.01 N force res.

Images and specs: company product pages and papers (Oct 2026). *Wuji describes a "pressure sensor matrix" but does not publish its sensing principle.

Tactile Sensors Used in Robotics Research

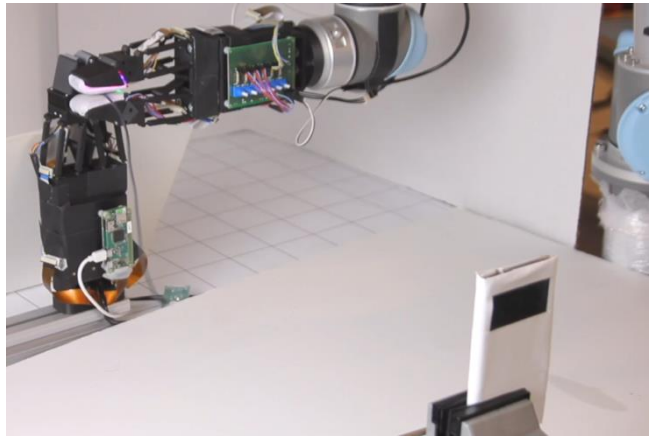

GelSight

Yuan, Dong & Adelson, Sensors 2017

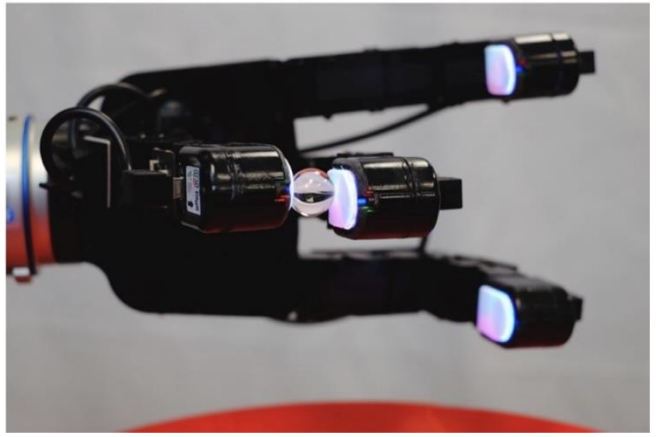

DIGIT

Lambeta et al., RA-L 2020

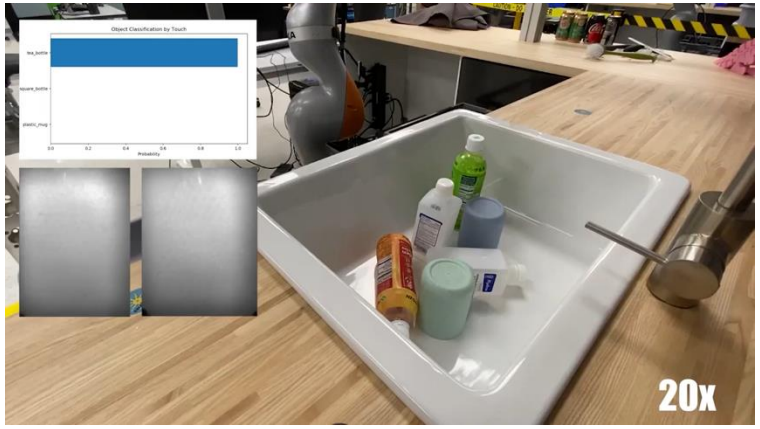

Soft Bubble

Kuppuswamy et al., IROS 2020

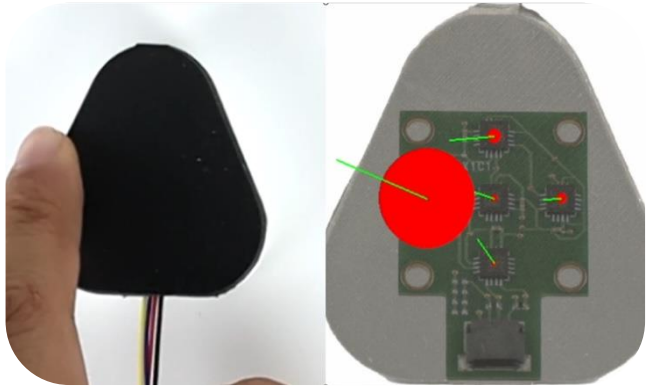

AnySkin

Bhirangi et al. [5]

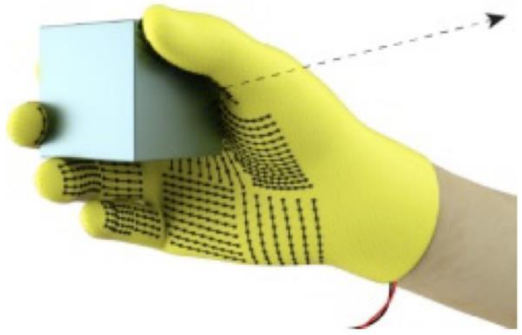

Scalable Tactile Glove (STAG)

Sundaram et al., Nature 2019

Tactile Sensors Used in Robotics Research

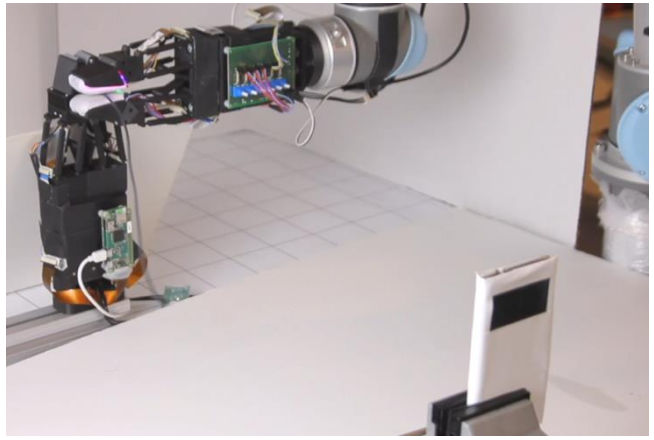

GelSight

Yuan, Dong & Adelson, Sensors 2017

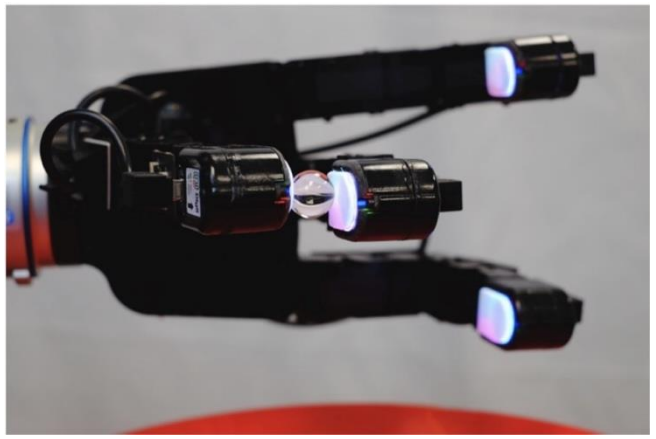

DIGIT

Lambeta et al., RA-L 2020

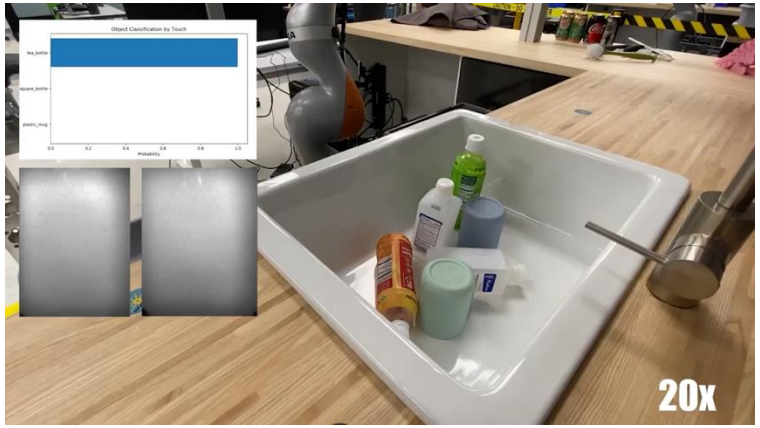

Soft Bubble

Kuppuswamy et al., IROS 2020

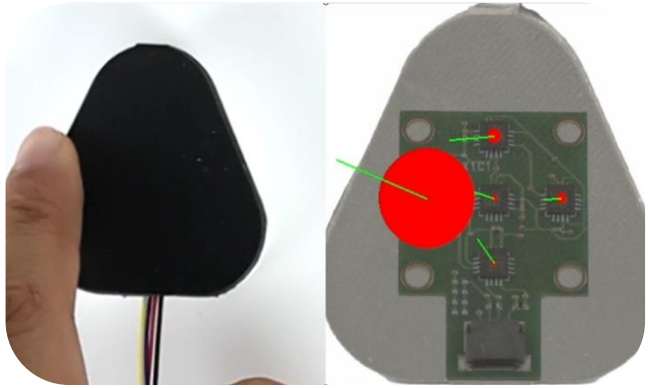

AnySkin

Bhirangi et al. [5]

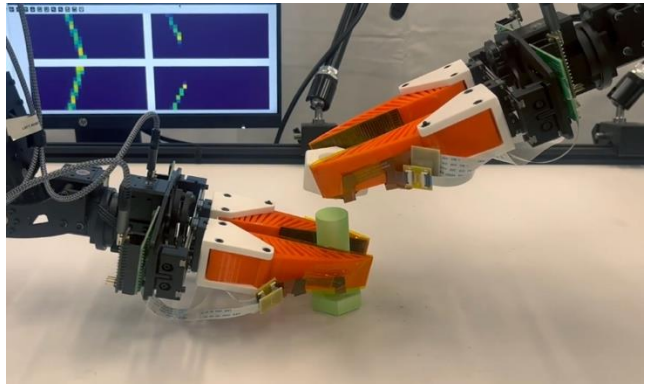

FlexiTac

Huang et al. 2025

Next: how each principle works →

Background: Optical-Based Tactile Sensors

SENSING PRINCIPLE

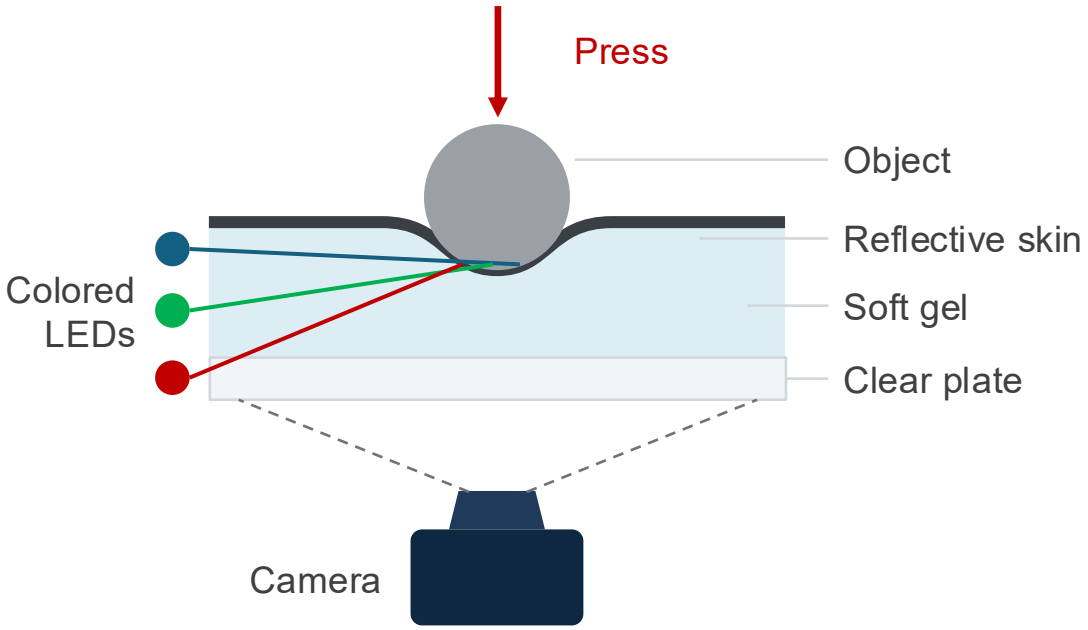

HOW IT WORKS

1. Contact deforms a soft gel coated with a thin reflective skin
2. Colored LEDs light the skin from different angles
3. A camera behind the gel images the deformed skin
4. Photometric stereo → surface normals → 3D depth map; printed markers track shear & slip

> **Output:** a high-resolution tactile image of contact geometry, texture and shear

A camera looking at a soft, reflective skin

Examples: GelSight · DIGIT · GelSlim · Soft Bubble (depth camera inside an air-filled membrane)

Example: GelSight, an Optical Tactile Sensor

It measures geometry directly; force and slip are inferred from how the skin deforms.

① GEOMETRY photometric stereo

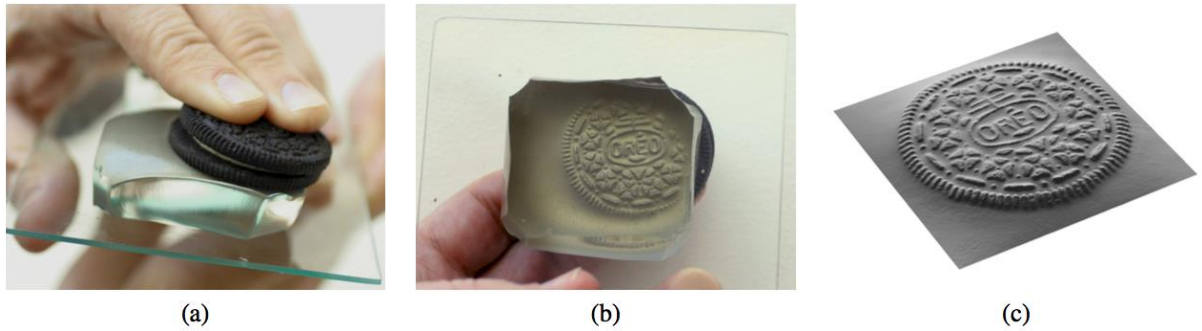

(a) press a cookie (b) the skin takes its exact shape (c) reconstructed 3D surface

THE HARDWARE

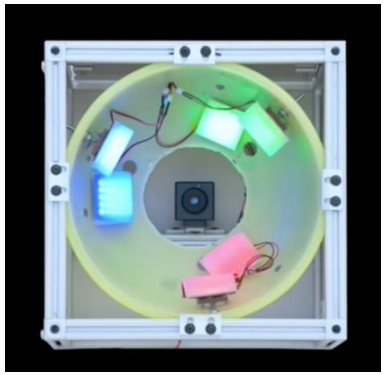

1–2 µm resolution (desktop)
30–100 µm on robot fingertips
RGB LEDs + webcam (30 Hz)

② FORCE & TORQUE printed markers

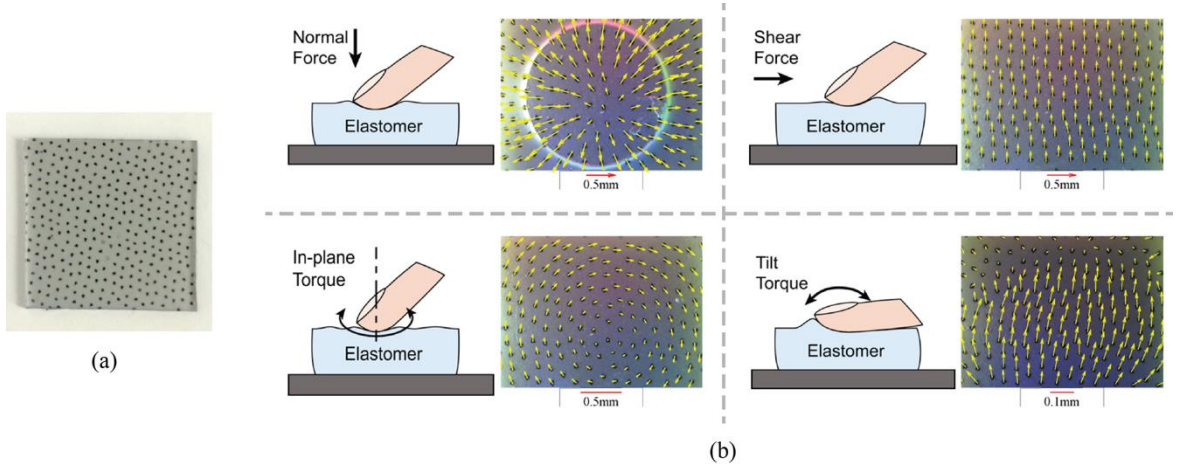

W. Yuan, S. Dong, E. H. Adelson, “GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force,” Sensors 17(12):2762, 2017 (CC BY 4.0). Figures 1, 2c, 3, 8 from the paper.

Background: Magnetic-Based Tactile Sensors

SENSING PRINCIPLE

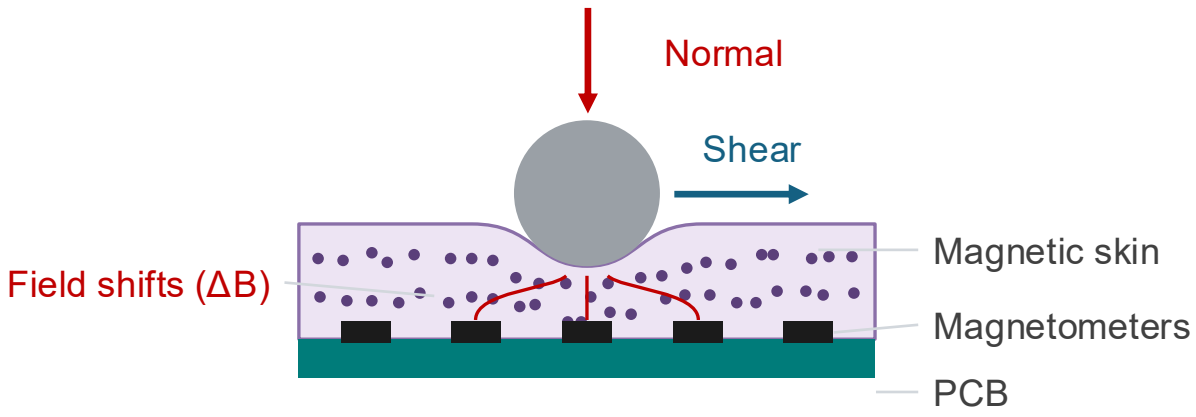

HOW IT WORKS

1. A soft elastomer skin is filled with magnetized particles (or tiny magnets)
2. Contact deforms the skin and shifts its magnetic field
3. Magnetometers on the PCB below measure the field in x, y and z
4. A calibrated or learned model maps field changes → contact location + 3-axis force

Output: a low-dim signal (ReSkin: 5 chips × 3 axes = 15 values) with normal + shear force

A soft magnetic skin/magnets over magnetometers

Examples: ReSkin · AnySkin · eFlesh · uSkin (XELA Robotics)

Background: Magnetic-Based Tactile Sensors

SENSING PRINCIPLE

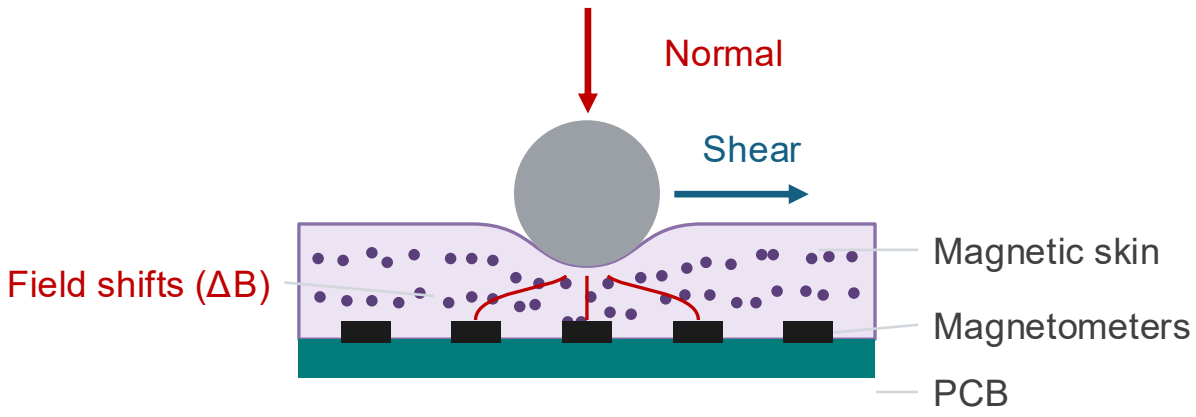

HOW IT WORKS

1. A soft elastomer skin is filled with magnetized particles (or tiny magnets)
2. Contact deforms the skin and shifts its magnetic field
3. Magnetometers on the PCB below measure the field in x, y and z
4. A calibrated or learned model maps field changes → contact location + 3-axis force

> **Output:** a low-dim signal (ReSkin: 5 chips × 3 axes = 15 values) with normal + shear force

A soft magnetic skin/magnets over magnetometers

Examples: ReSkin · AnySkin · eFlesh · uSkin (XELA Robotics)

What if we use magnetic sensor to manipulate your phone?

Example: eFlesh, a 3D-Printed Magnetic Sensor

No magnetized elastomer needed: print a lattice and drop in off-the-shelf magnets.

① FABRICATION



Send design to 3D printer



Pause print at pouches



Insert magnets into pouches



Resume print until complete

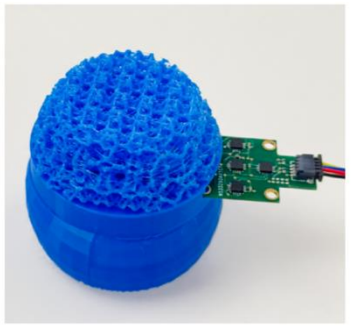

Insert magnetometer PCB into slot

② IN ACTION

Press → live contact map

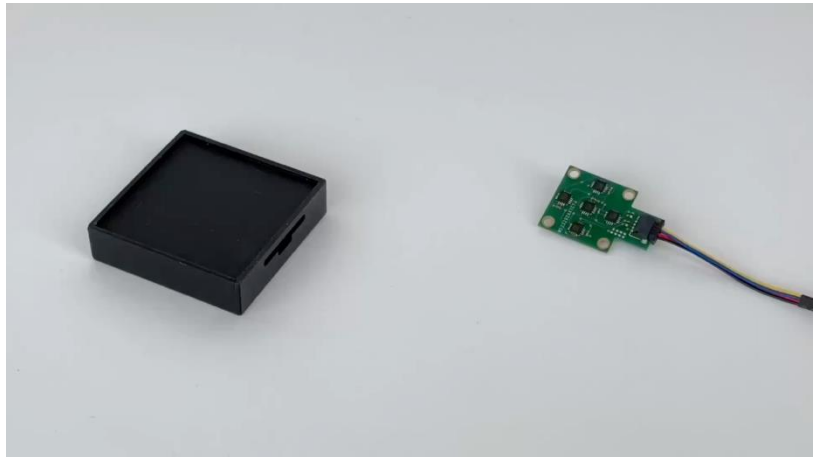

Gentle grasp, precise localization

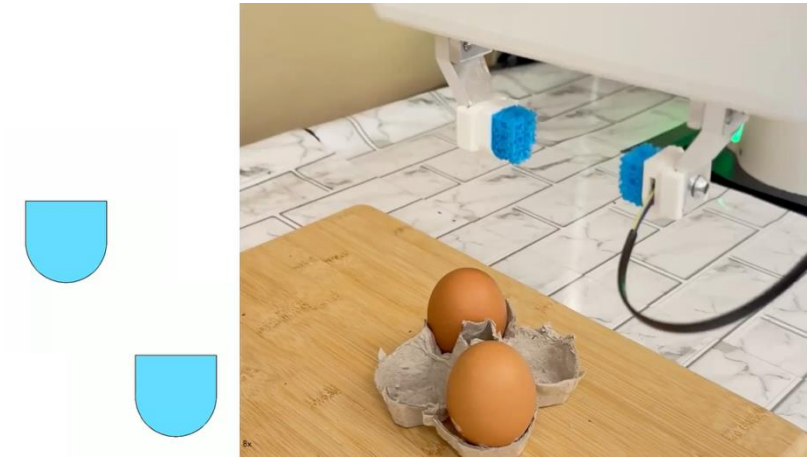

V. Pattabiraman, Z. Huang, D. Panozzo, D. Zorin, L. Pinto, R. Bhirangi, “eFlesh: Highly Customizable Magnetic Touch Sensing using Cut-Cell Microstructures,” arXiv:2506.09994, 2025.

Background: Resistive Tactile Sensors

SENSING PRINCIPLE

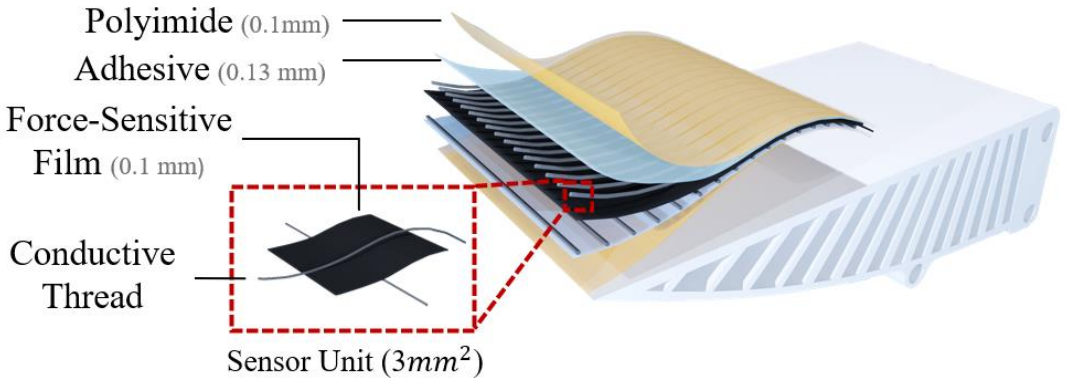

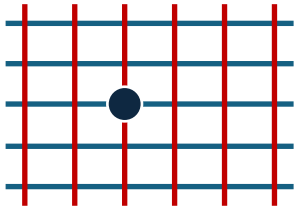

Rows (top) × **columns** (bottom)
Each crossing = one taxel; scan one row at a time to read the whole map

HOW IT WORKS

1. A force-sensitive (piezoresistive) film sits between two electrode layers
2. Pressure compresses the film → resistance drops
3. Top electrodes run in rows, bottom ones in columns: each crossing is one taxel
4. Electronics scan one row at a time, so N + M wires read N × M taxels

Output: a 2D map of normal force

A simple resistive "sandwich"

Examples: Force-sensing resistors (FSR) · Tekscan · MIT tactile glove (STAG) · FlexiTac

Interactive Course for FlexiTac

LEARN FLEXITAC

TOUCH · SCAN · SENSE

COURSE PROGRESS 2/6

12 × 32 taxels
100 frames / s

HOME
LESSONS
GLOSSARY
ABOUT
THEME

AN INTERACTIVE COURSE · 6 LESSONS

How does a robot feel a touch?

FlexiTac is an open-source touch sensor thinner than 1 mm. A grid of 12 × 32 tiny pressure spots, read 100 times a second, gives a robot fingertip a picture of what it is holding. Take it apart, press it and scan it, one idea at a time.

▶ START LESSON 1 | ALL LESSONS

2 done · next up: 3. Pick one crossing

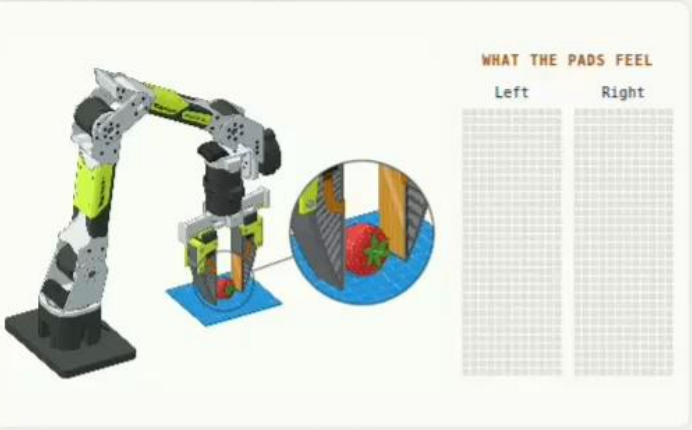

LESSON 1 · ON A ROBOT · ILLUSTRATIVE

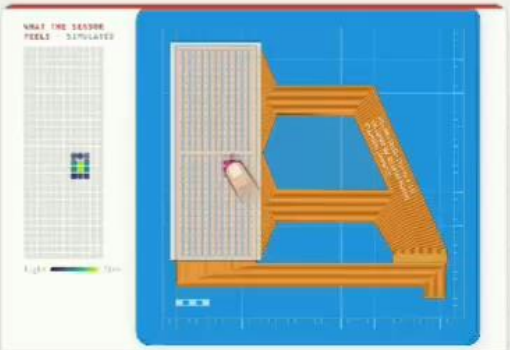

THE PAD · LESSONS 1–2

Press the sensor and watch its touch map light up, then take it apart layer by layer.

1 MEET FLEXITAC ›



THE BOARD · LESSONS 3–4

Watch the reading board pick one crossing at a time, then plug in the pad and see it read a whole frame, row by row.

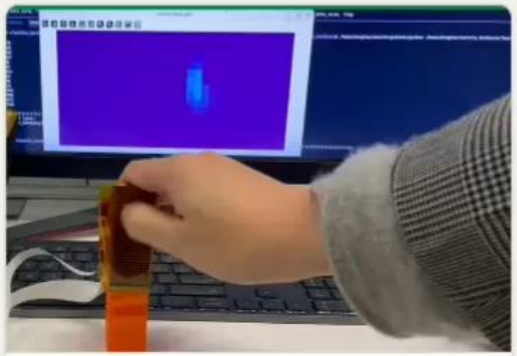

USE IT · LESSONS 5–6

Watch a live heatmap follow your finger, build your own pad in about 3 minutes, and fix it when the picture looks wrong.

localhost:5173/#/lesson/pick-one-crossing

Tactile Hardware Has Not Yet Converged



Optical	gel + reflective skin	**camera**	high-res image
Magnetic	magnetized elastomer	**magnetometers**	3-axis field, few taxels
Piezoresistive	force-sensitive film	**resistance change**	pressure map

WHY THERE'S NO WINNER YET

Every design trades off the same four things:

Resolution **Form factor** **Cost** **Durability**

no standard sensor.
Next: compare them,

Comparison

Optical-based:
High spatial resolution
Shear/texture cues
Bulky form factor
Hard to Customize

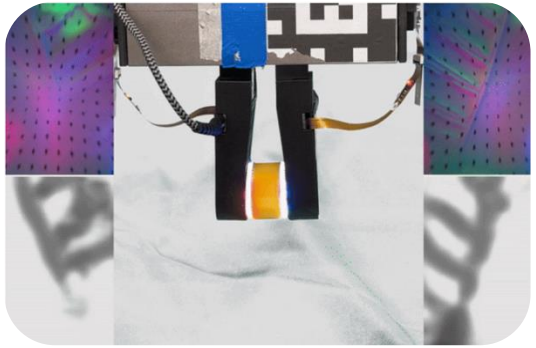

GelSlim — Sipos *et al.* [1]

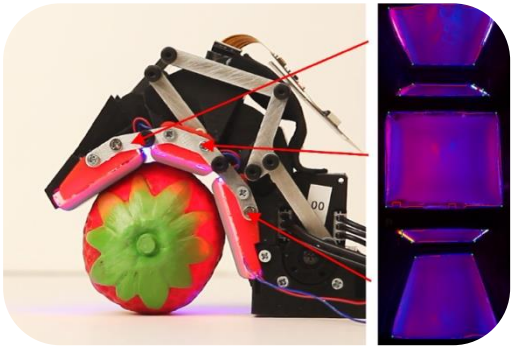

GelLink — Ma *et al.* [2]

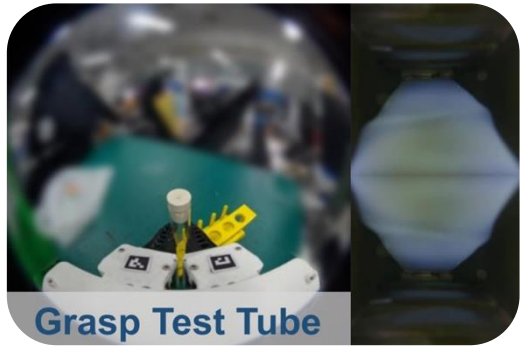

ViTaMIn — Liu et al. [3]

Magnetic-based:
Shear Sensing
Easy to shape
Noisy signals
Low-Resolution



eFlesh — Pattabiraman et al. [4]

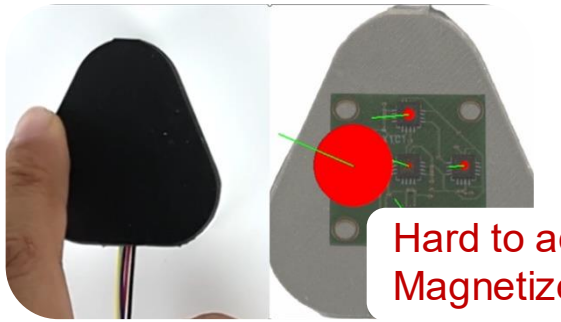

AnySkin — Bhirangi *et al.* [5]



ReSkin — Bhirangi *et al.* [6]

Hard to access:
Magnetized Elastomer required

Resistive-based:
Flexible and Compliant
"Good" Resolution (2 mm)
Ultra-thin (~0.5mm)
Easy to customize
No Explicit Shear

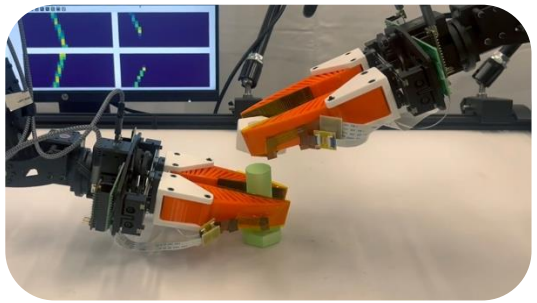

FlexiTac — Huang et al.



OpenTac — Slepyanet al.

[1] Sipos *et al.,* "GelSlim 4.0: Focusing on Touch and Reproducibility," 2024.
[2] Ma *et al.,* "GelLink: A Compact Multi-Phalanx Finger with Vision-Based Tactile Sensing and Proprioception," 2024.
[3] Liu et al., "ViTaMIn: Learning Contact-Rich Tasks Through Robot-Free Visuo-Tactile Manipulation Interface," 2025.
[4] Pattabiraman *et al.,* "eFlesh: Highly customizable Magnetic Touch Sensing using Cut-Cell Microstructures," 2025.
[5] Bhirangi *et al.,* "AnySkin: Plug-and-play Skin Sensing for Robotic Touch," 2024.
[6] Bhirangi *et al.,* "ReSkin: versatile, replaceable, lasting tactile skins," 2021.

Quick Check: Which Sensor Would You Pick?

Options: optical · magnetic · resistive. Be ready to say why.

A

Inspect tiny scratches on a phone screen at a factory station

Your pick: ?

B

Cover a humanoid's whole palm and fingers with touch

Your pick: ?

C

Detect an object starting to slip sideways in a gripper

Your pick: ?

Think for 30 seconds, then we vote.

Quick Check: Which Sensor Would You Pick?

Options: optical · magnetic · resistive. Be ready to say why.

A

Inspect tiny scratches on a phone screen at a factory station

Optical

Micron-level geometry; size and cost matter less on a fixed station

B

Cover a humanoid's whole palm and fingers with touch

Your pick: ?

C

Detect an object starting to slip sideways in a gripper

Your pick: ?

There is no universal best sensor: match the sensor to the task.

Quick Check: Which Sensor Would You Pick?

Options: optical · magnetic · resistive. Be ready to say why.

A

Inspect tiny scratches on a phone screen at a factory station

Optical

Micron-level geometry; size and cost matter less on a fixed station

B

Cover a humanoid's whole palm and fingers with touch

Resistive

Thin, flexible, cheap per area; normal force is often enough

C

Detect an object starting to slip sideways in a gripper

Your pick: ?

There is no universal best sensor: match the sensor to the task.

Quick Check: Which Sensor Would You Pick?

Options: optical · magnetic · resistive. Be ready to say why.

A

Inspect tiny scratches on a phone screen at a factory station

Optical

Micron-level geometry; size and cost matter less on a fixed station

B

Cover a humanoid's whole palm and fingers with touch

Resistive

Thin, flexible, cheap per area; normal force is often enough

C

Detect an object starting to slip sideways in a gripper

Magnetic

Slip shows up as shear, which resistive films cannot sense (optical + markers also works)

There is no universal best sensor: match the sensor to the task.

Optical-based for Precision, Resistive-Based for Scale

Optical Sensors:
High-precision system for industrial settings

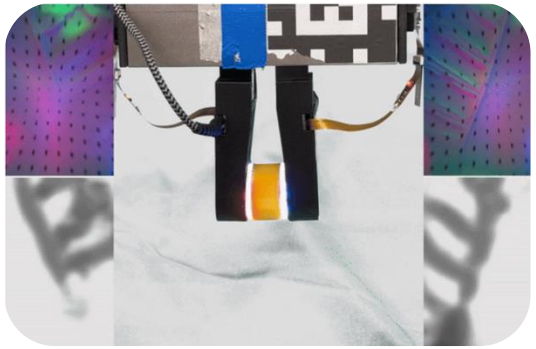

GelSlim — Sipos *et al.* [1]



Precision Robots in Factory



GelSight for Defect Inspection

Resistive:
Scalable systems for everyday robotics

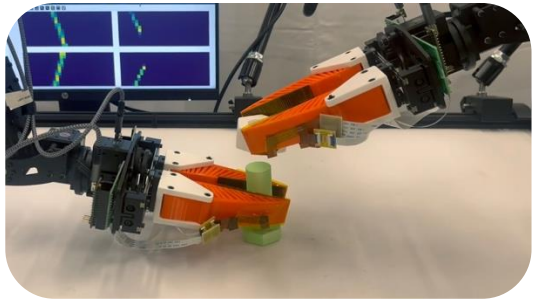

Scalable Tactile Sensing



Scalable Robot Platform

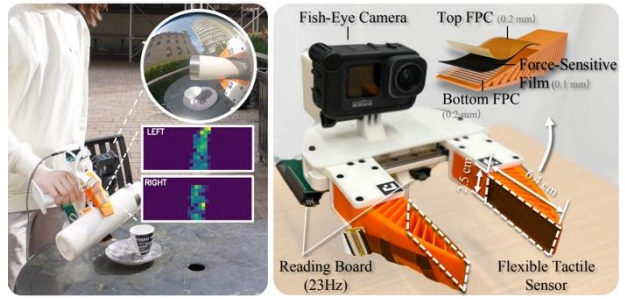

Portable Data Collection Device

Case Study: Building a Tactile Sensor from Scratch

My own hands-on experience: what it takes to turn a sensing principle into a tool other labs use

1

Prototype

Hand-built sensors (2023): slow to make, inconsistent, noisy

2

Iterate

Redesign the sensor and readout electronics for fabrication

3

Share

Open-source FlexiTac, integrate with LeRobot, grow a community

Hardware → System → Learning

As we go, ask yourself: what turned a fragile lab prototype into something others could reproduce?

My first prototype — built by hand

BACK TO 2023

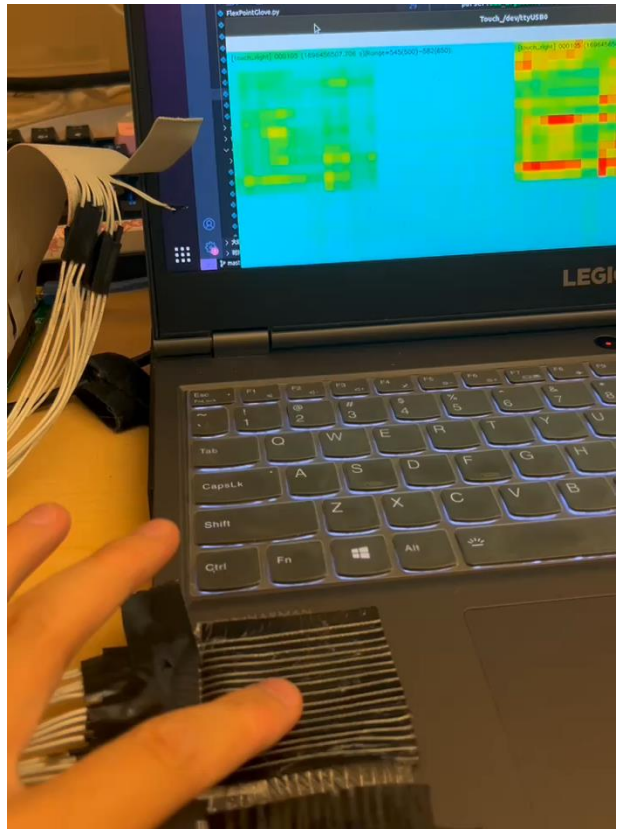

Hand-built prototype, 2023

Where it started — four real problems:

6 hours

lead time per sensor

≠ Inconsistent

piece-to-piece variation

~ Noisy

unstable raw signals

? Unknown

robustness / lifetime

Couldn't Reproduce It Reliably, So I Tried to Buy One

BACK TO 2023

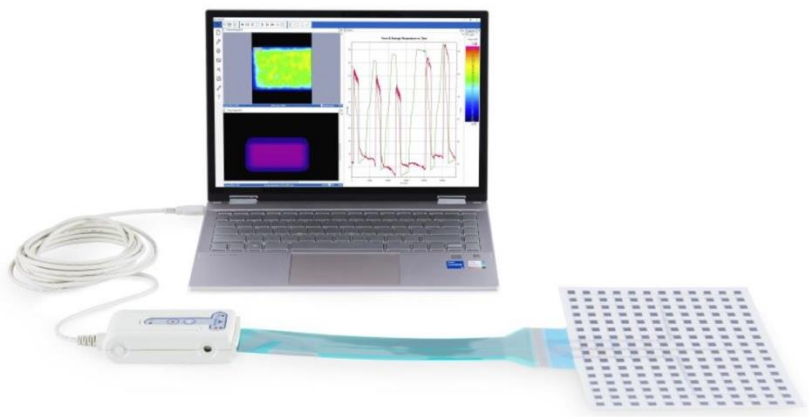

Commercial Pressure Sensor (Tekscan)

COMMERCIAL
Tekscan and similar
✗ ~$1,000 for system, and no customization.

FACTORY
OEM
✗ Cheap off-the-shelf, but ~$6K to customize.

ACADEMIC
Lab's prototype
✗ No response, and a steep licensing fee.

Why not build from scratch, from principle?

SENSING PRINCIPLE

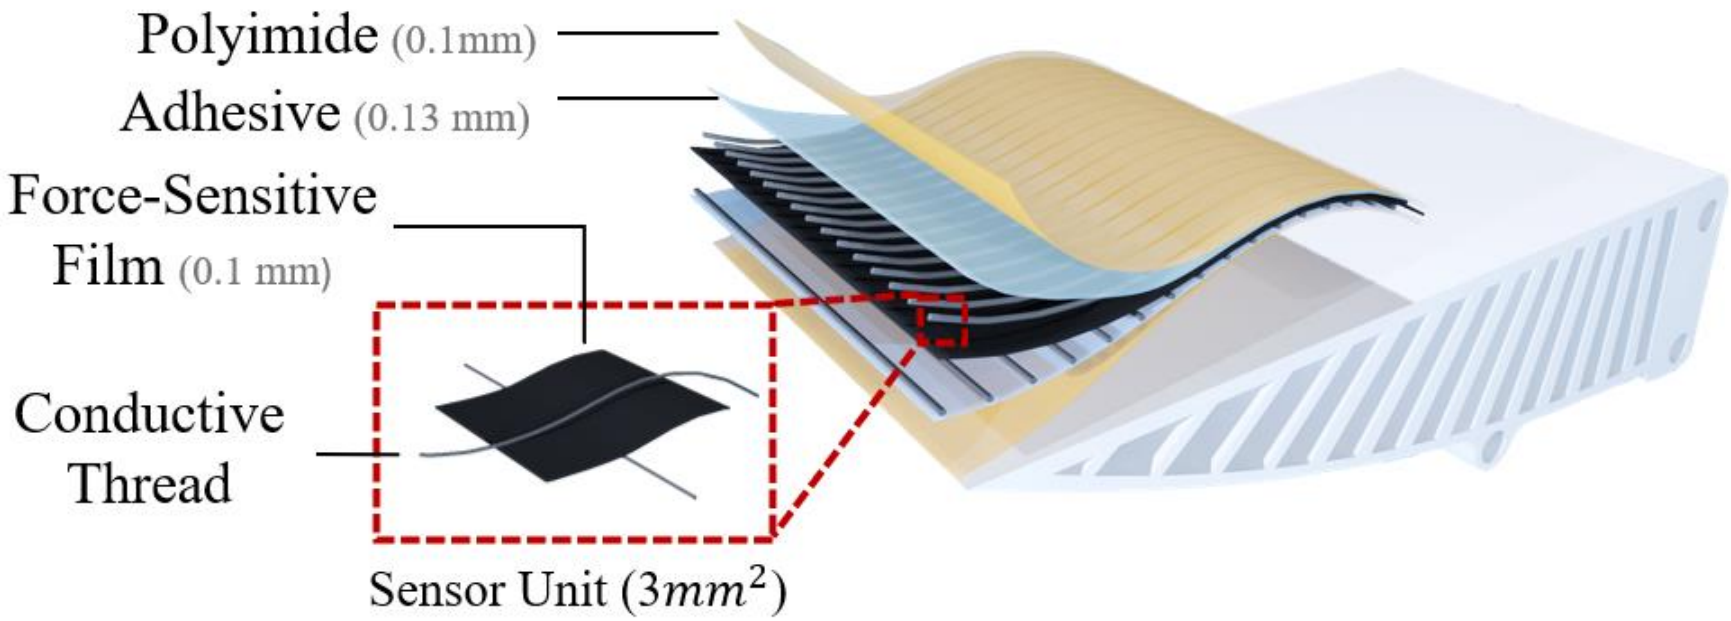

A simple resistive "sandwich"

Why not build from scratch, from principle?

THE SENSOR · faster to fabricate each version

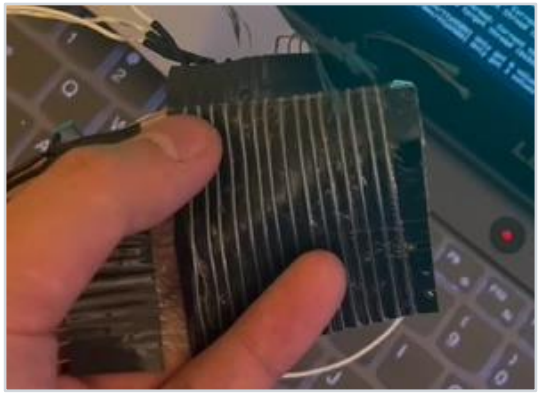

6 Hours

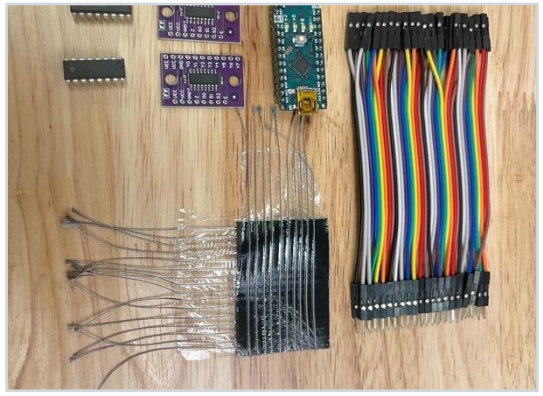

2 Hours

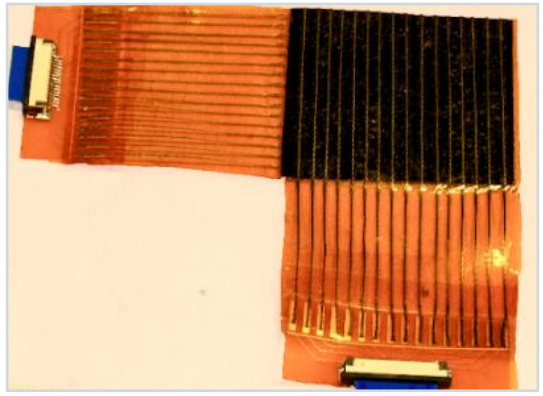

30 Minutes

THE READOUT ELECTRONICS · from design to production

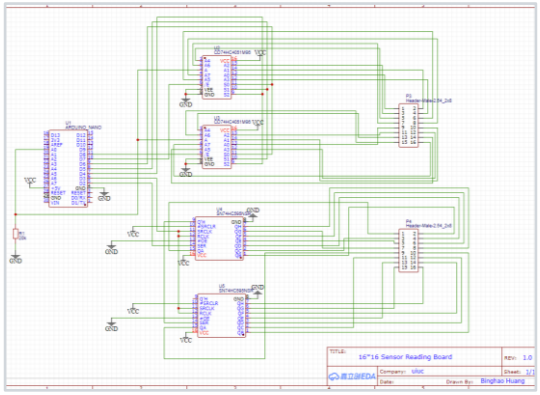

Design

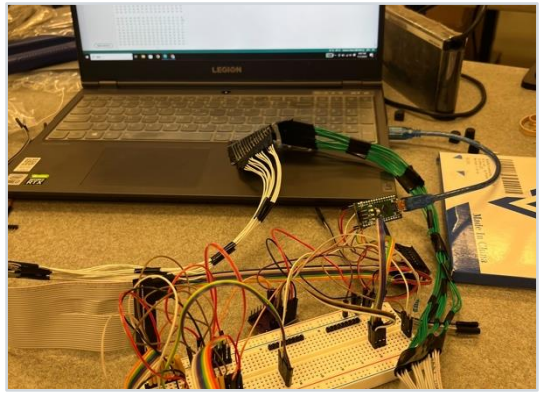

Testing

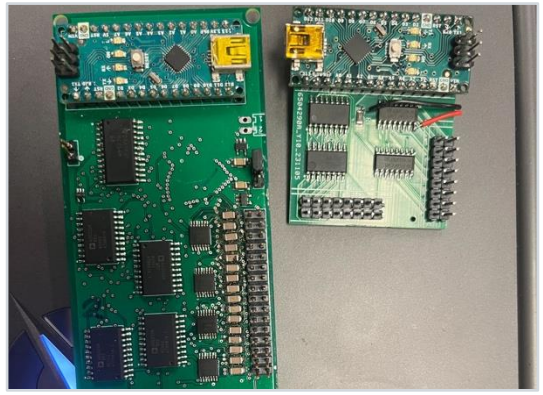

Production

After months of testing and iteration


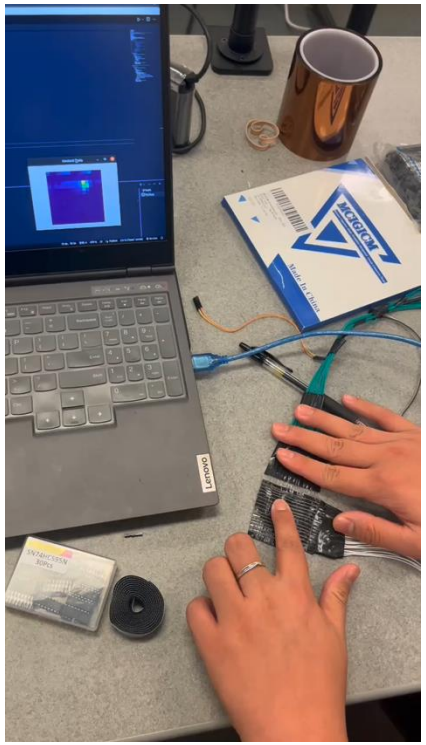
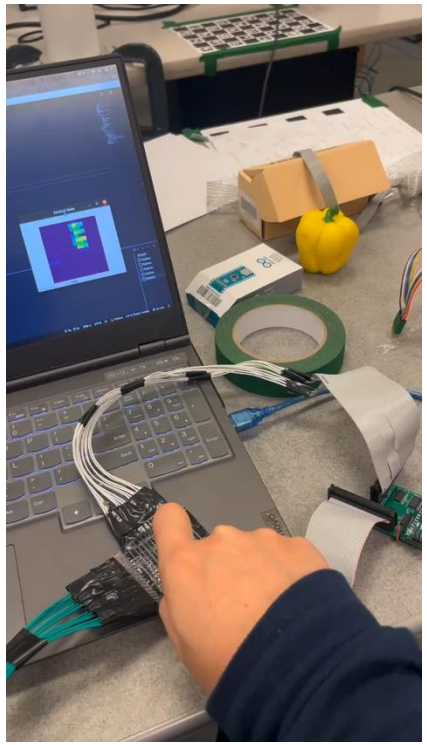
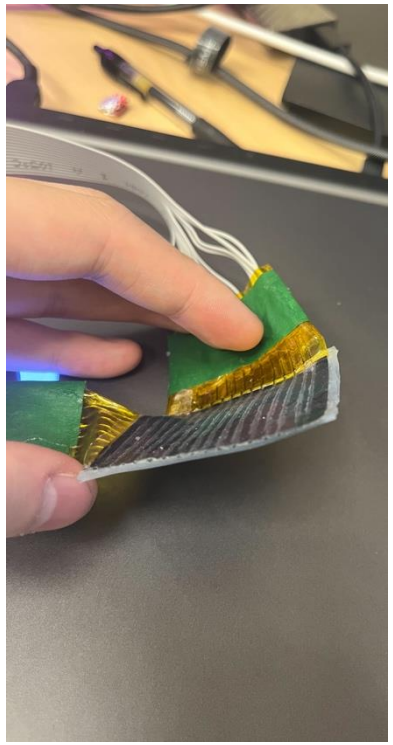
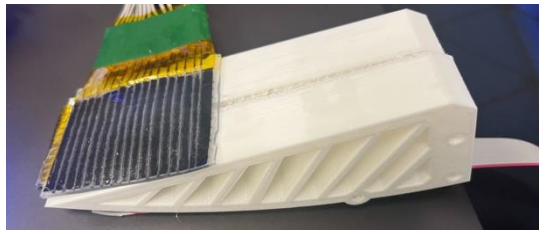
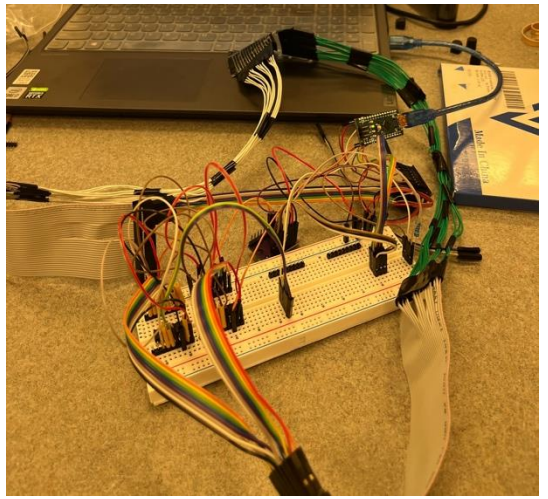
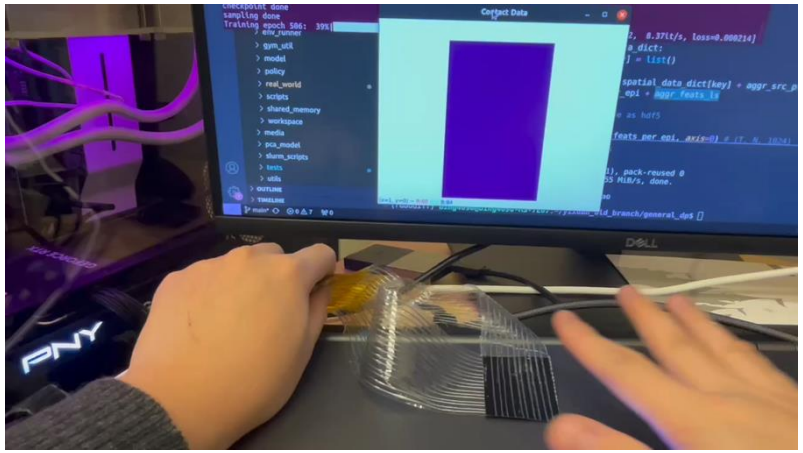
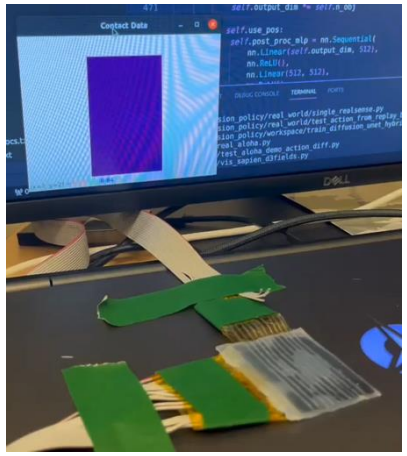
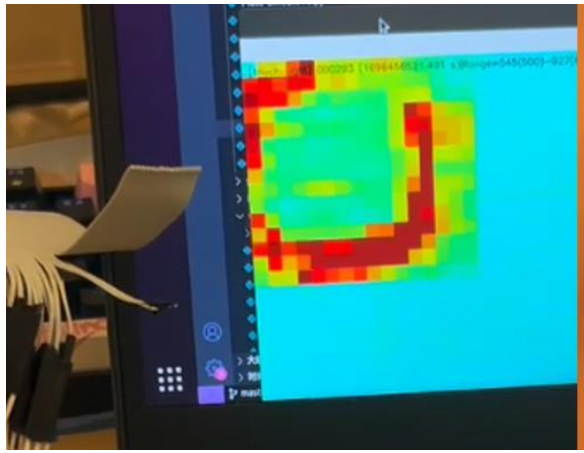


Finally, we got a stable version

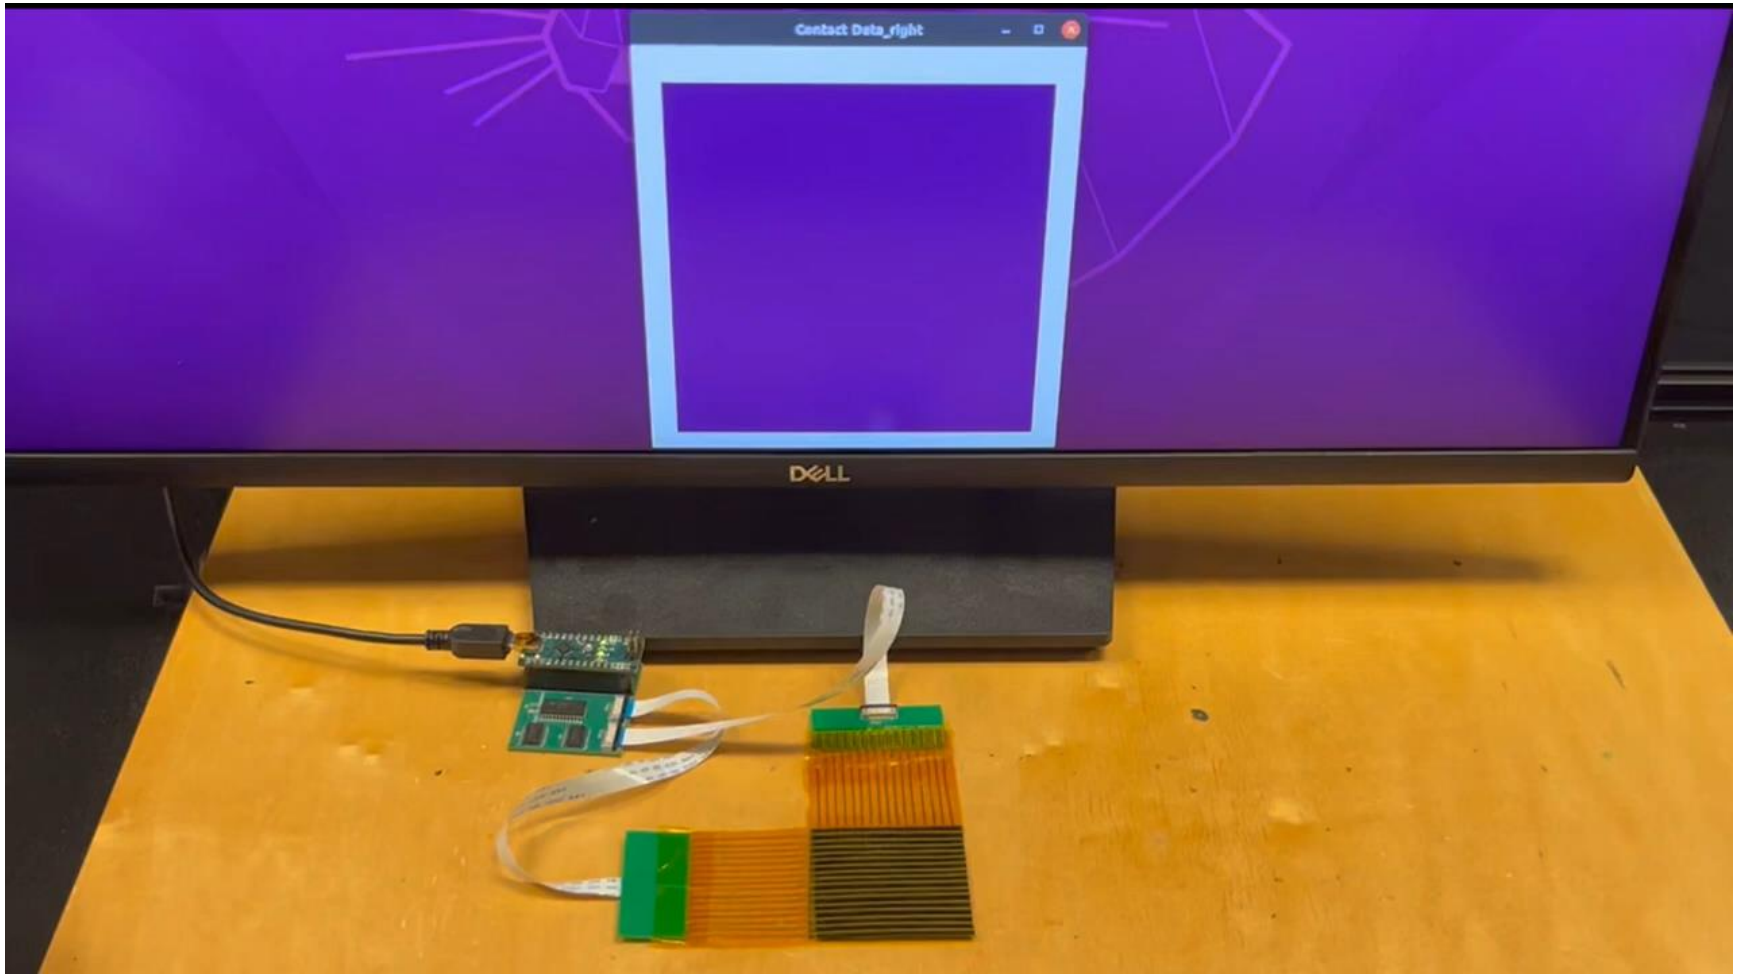

Accessible, Low-cost, Robust!

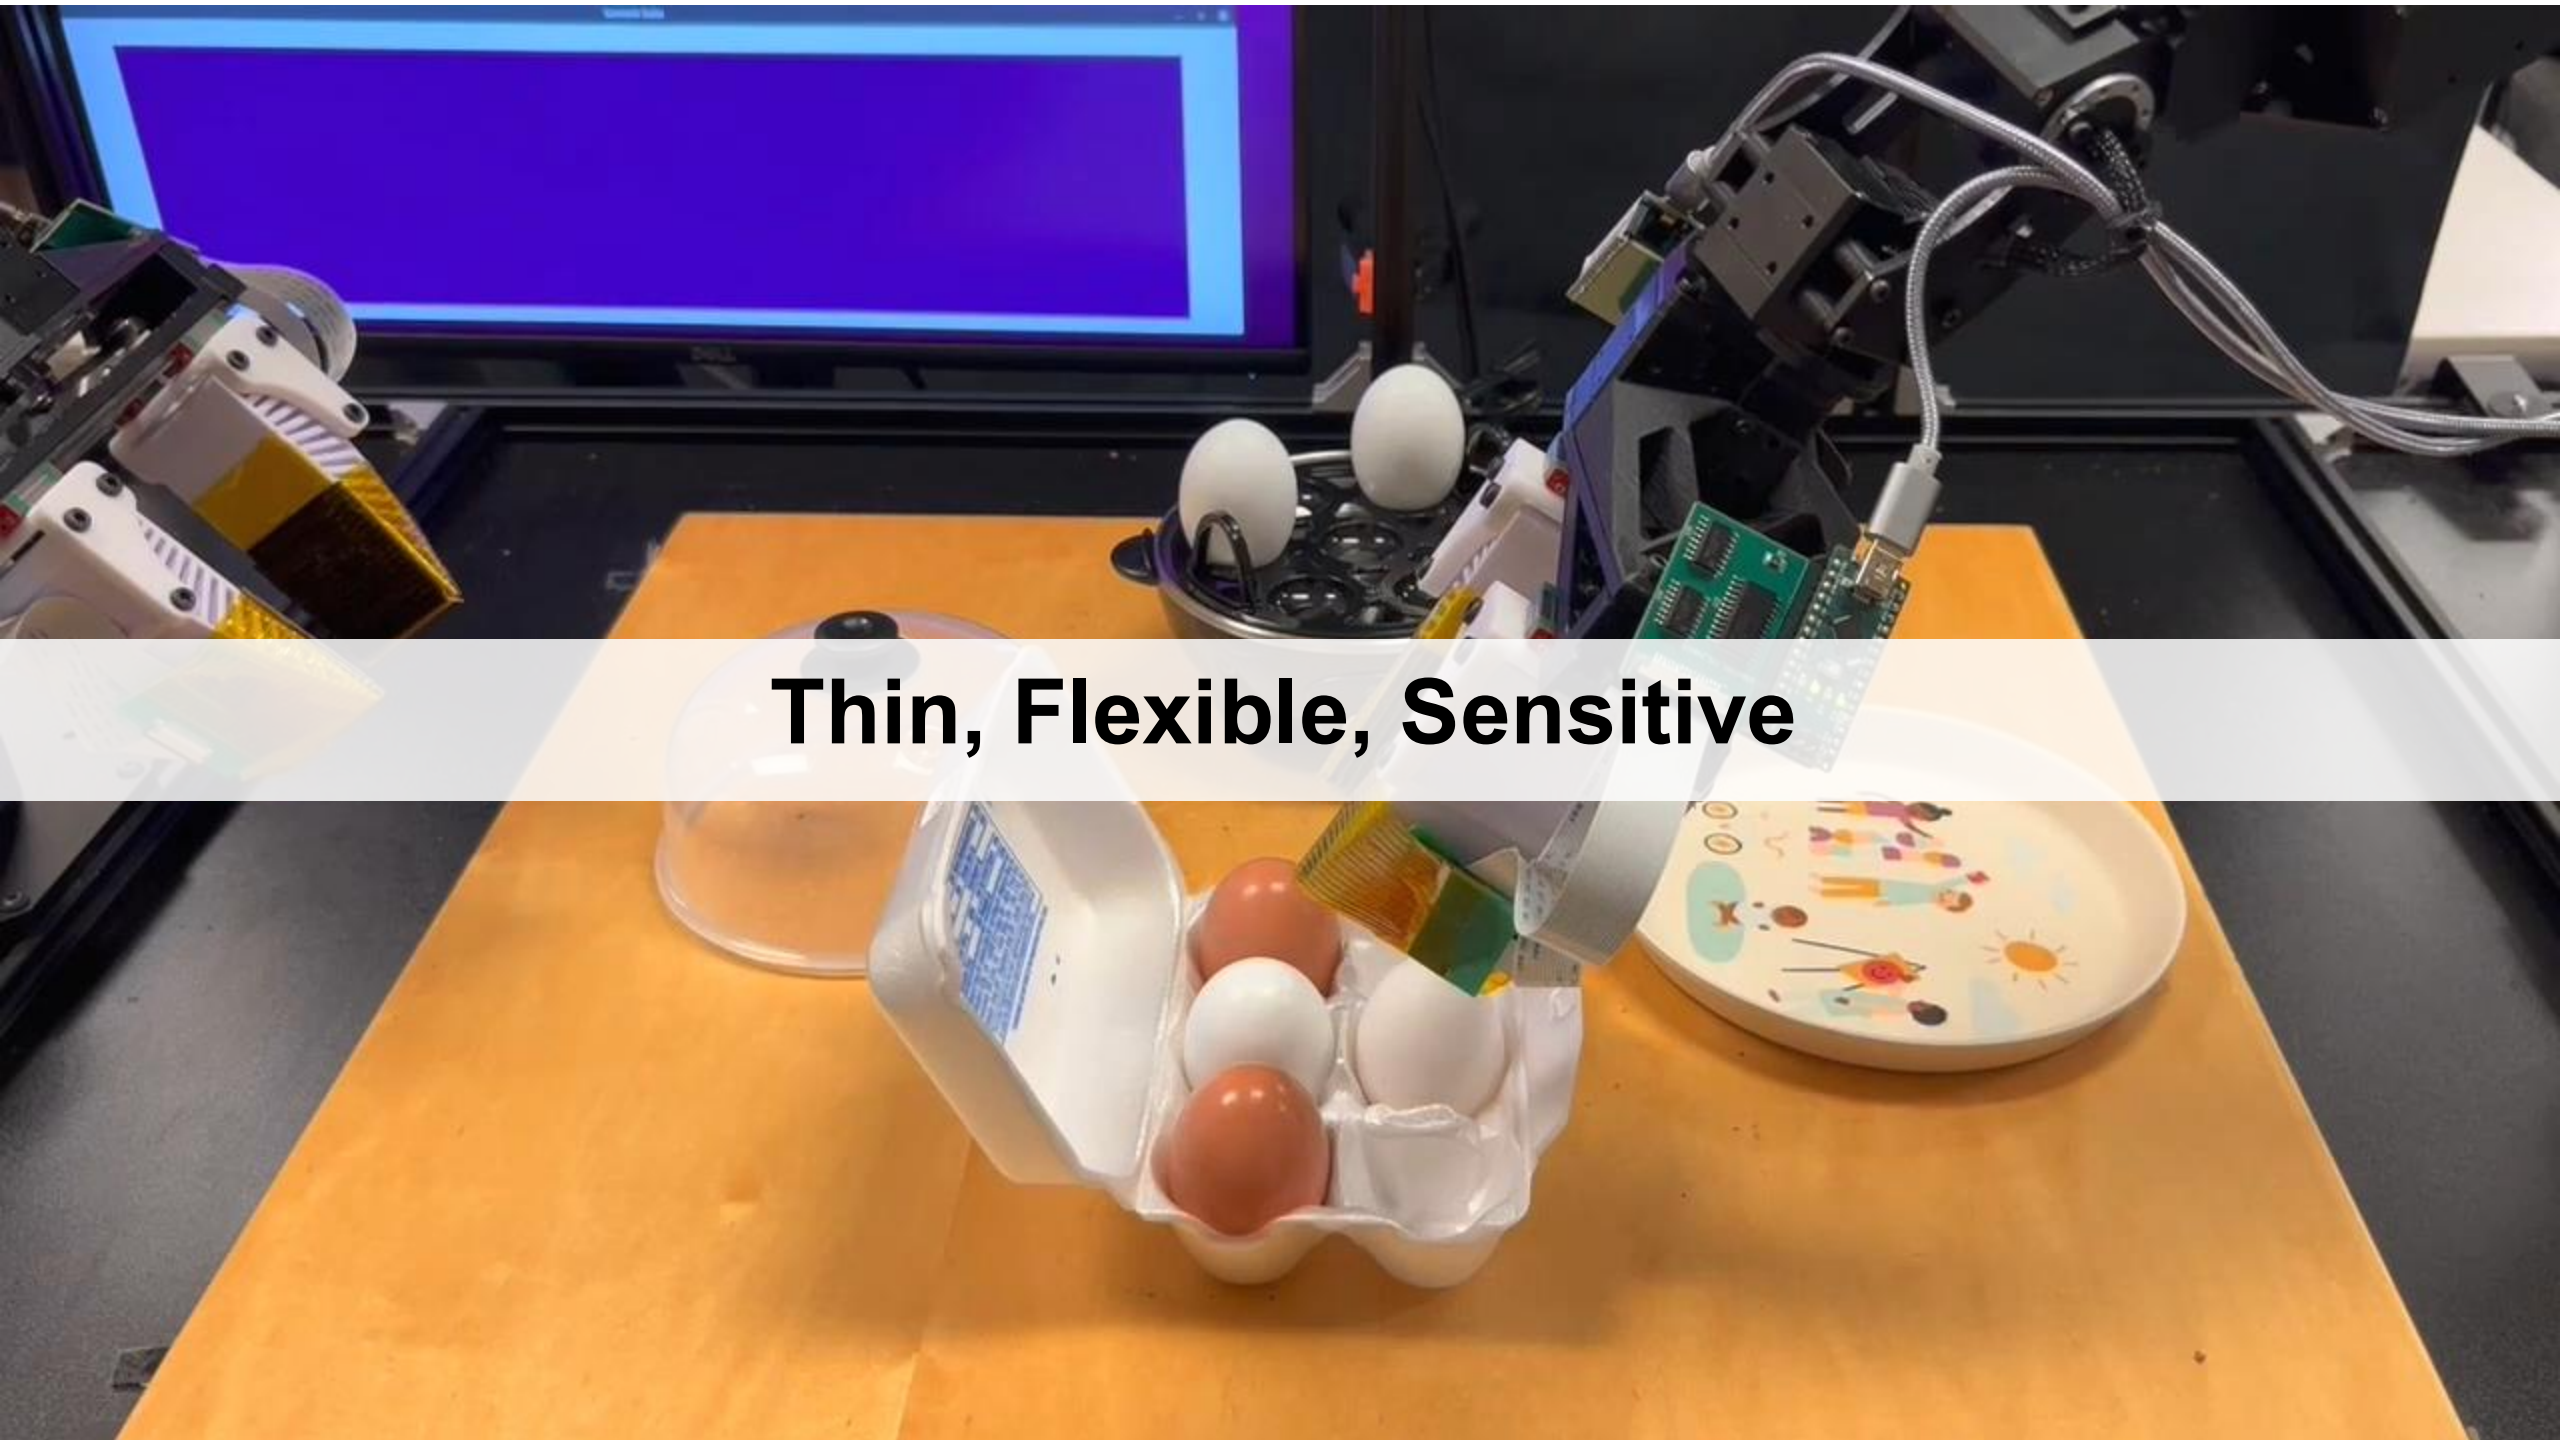

Thin, Flexible, Sensitive



Our Platform Today

Making Tactile Sensors Accessible to Everyone

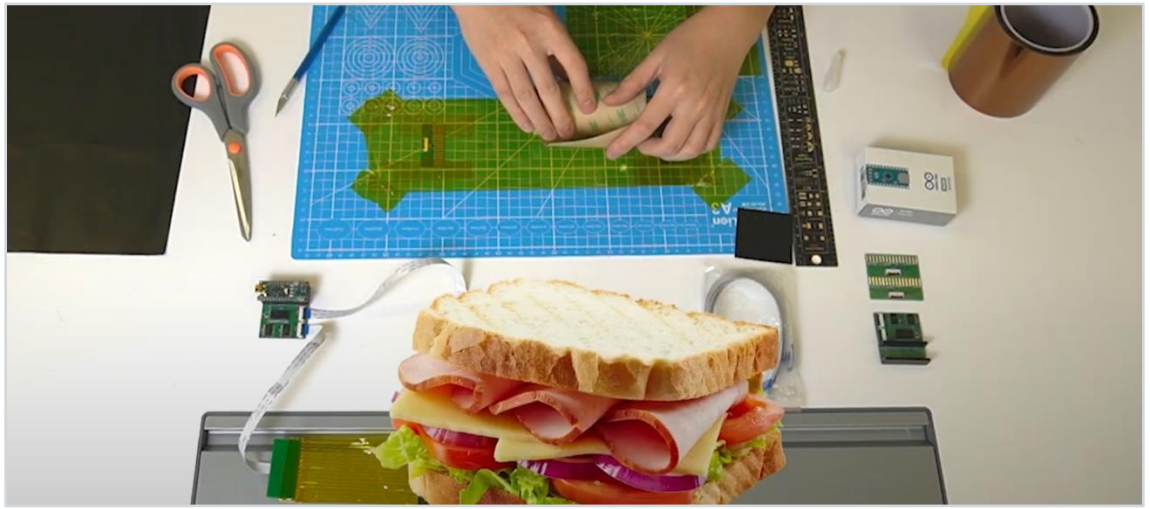

FlexiTac-V1

Fabrication	~30 min
Price	~$5/piece
Customize	hand-made
Lead time	in-lab ~30 min

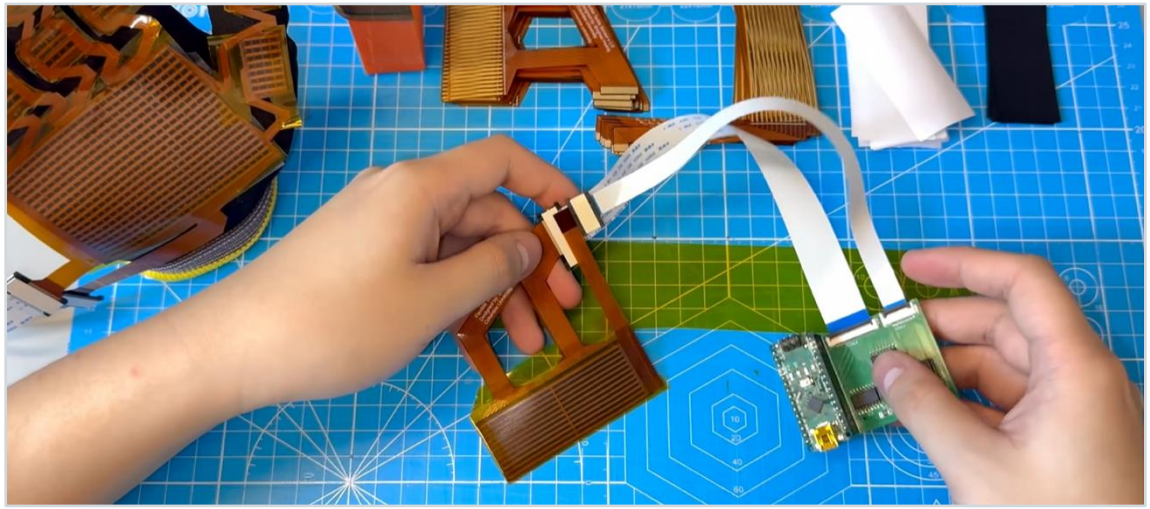

FlexiTac-V2

Fabrication	~3 min
Price	$4.19/piece (20 unit) $1.36/piece (1000 unit)
Customize	from PCB factory
Lead time	manufacture + ship < 1 week

Why Adoption Matters in Robotic Touch



2019

Only MIT used it

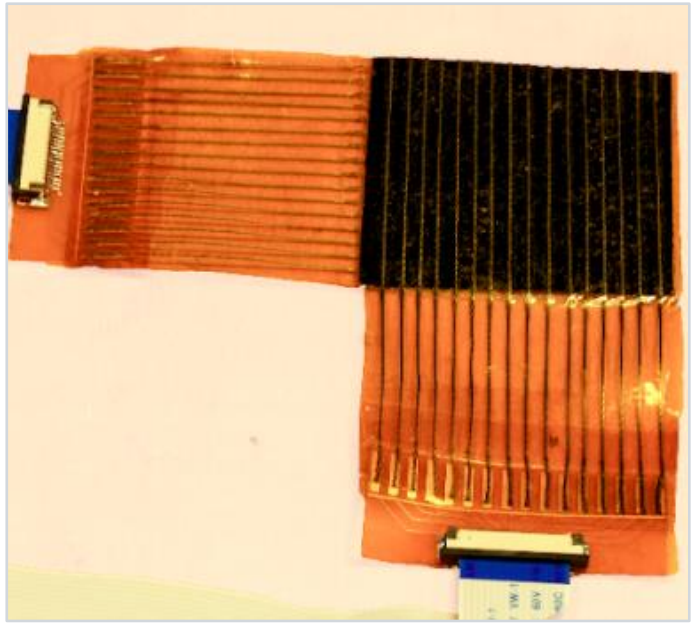

2024

More research labs adopt it

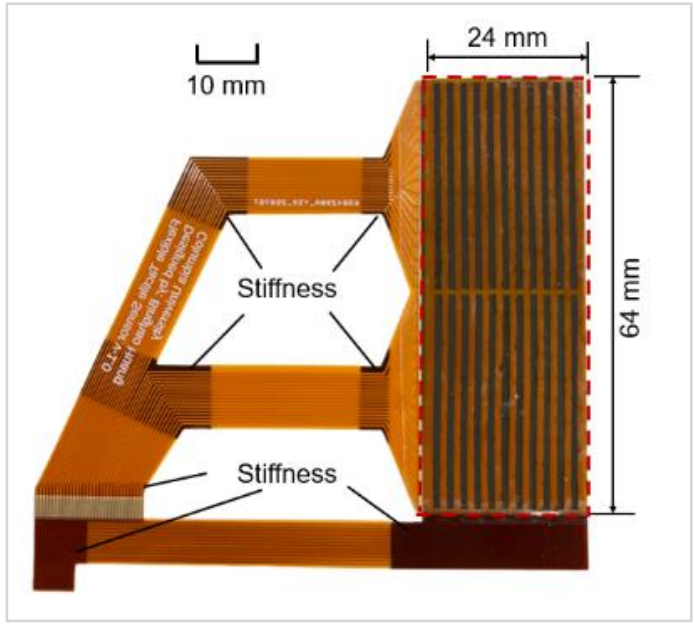

2026 ✓

Broad academia + industry use

Our claim: make sensors accessible, and the field builds on them.

We Open-Source FlexiTac!

FlexiTac Hardware Fabrication Guide

Binghao Huang[1], Yunzhu Li[1]
[1]Columbia University
March 2026

Simple but Works!

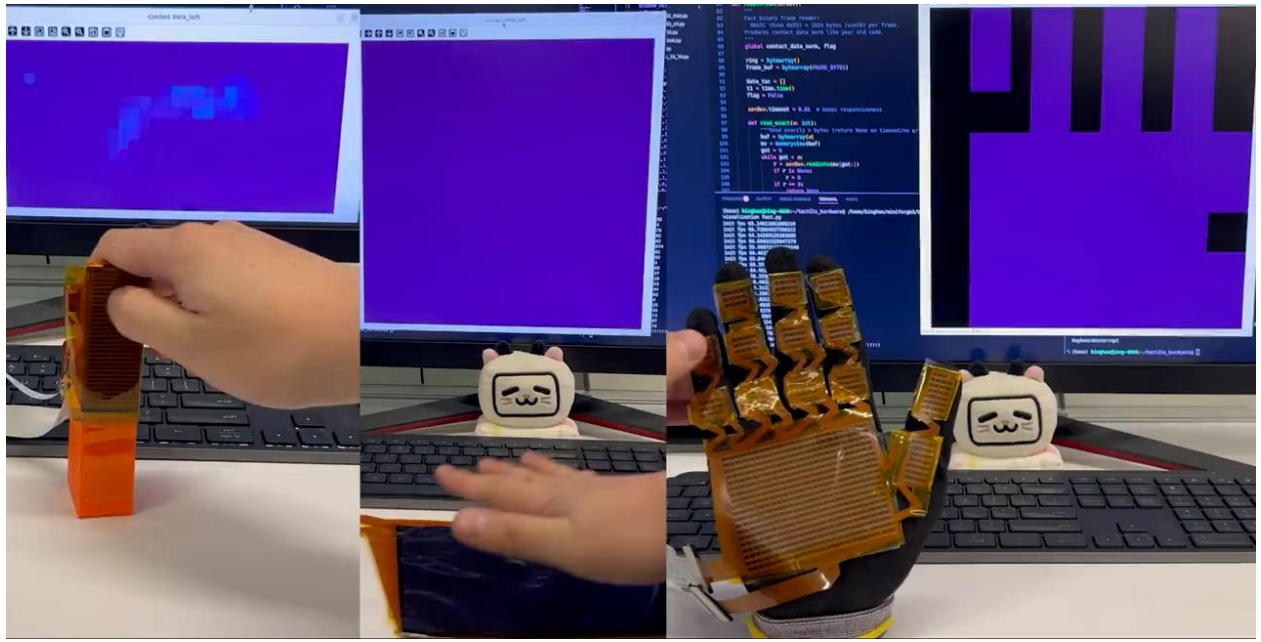

Get Your Tactile Sensor in 3 mins!

FlexiTac Hardware Fabrication Guide

Binghao Huang[1], Yunzhu Li[1]
[1]Columbia University
March 2026

Simple but Works!

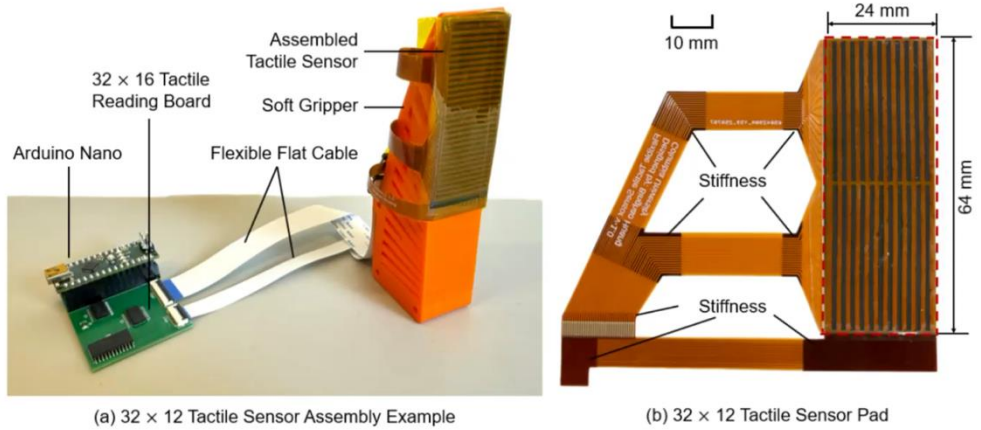

(a) 32 × 12 Tactile Sensor Assembly Example (b) 32 × 12 Tactile Sensor Pad

Get Your Tactile Sensor in 3 mins!

Resources

Contact

Please email binghao.huang@columbia.edu for any questions.

Agent-Driven Flexitac Customization

02 / 13

Robotiq 2F-85

Robotiq 2F-85 CAD, opened (asset/2F85_Opened.stl)

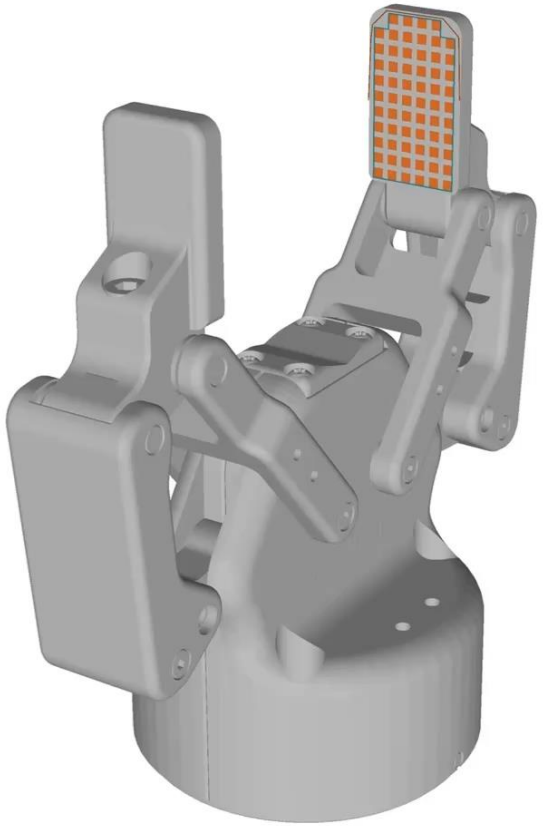

Taxel layout I plan

64 taxels 11 x 6 grid

pitch 3.40 mm electrode 2.21 mm

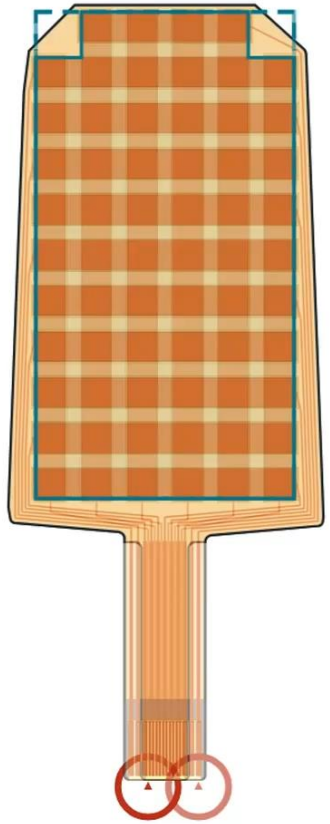

10 mm

Sensor I exploded 3D

64 taxels 11 x 6 grid

max rows 29 trace 0.127 mm tail pins 11 + 6

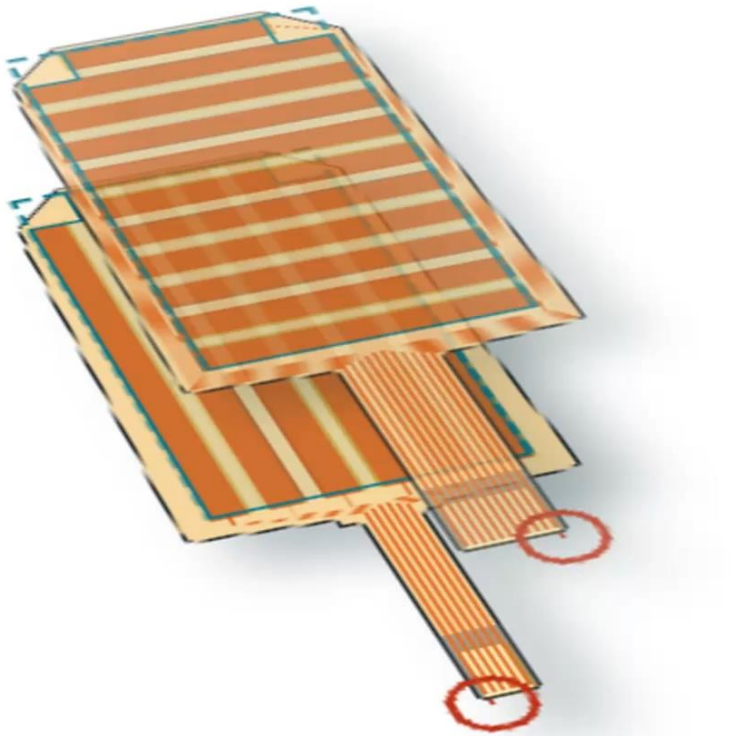

surface-11x6-c018478c9cfc Row FPC + Column FPC

… and Integrate into the LeRobot Community

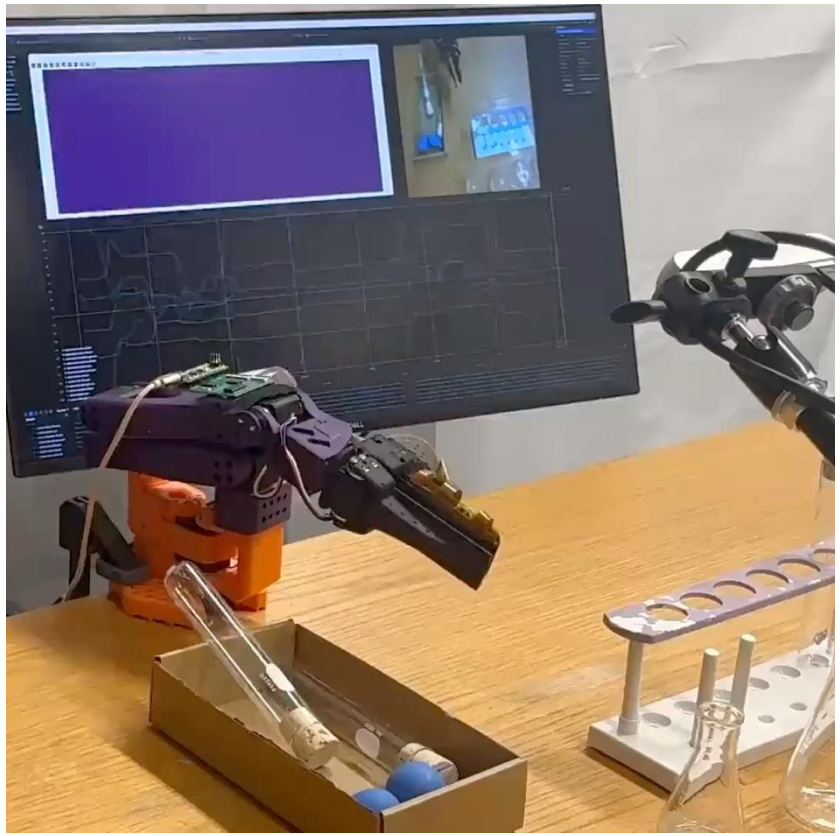

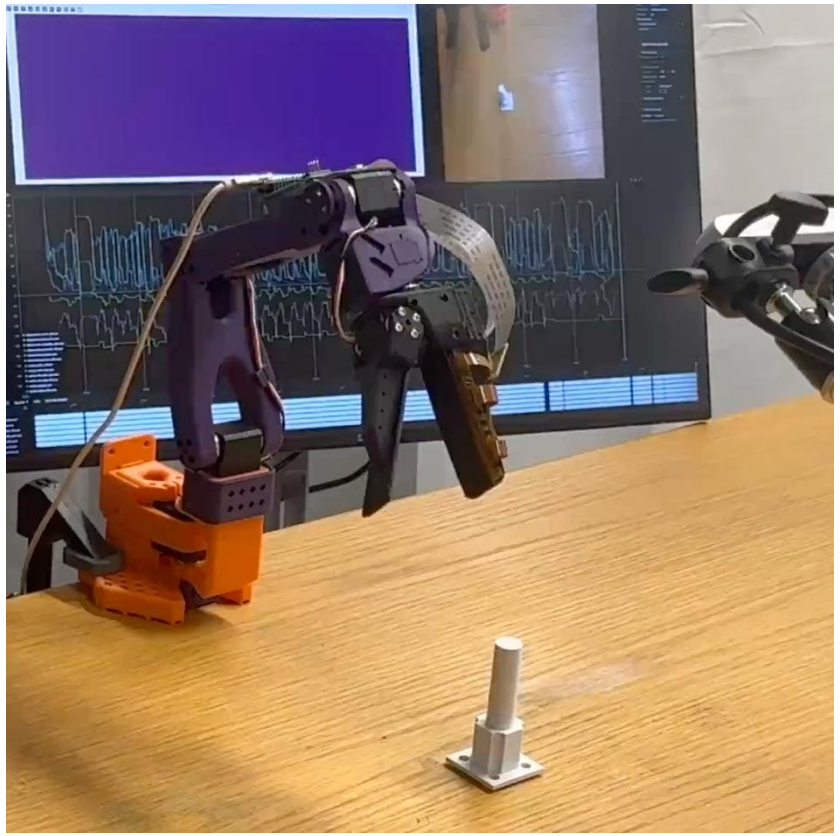



Naian Tao, Yifan He, Wesley Maa, Binghao Huang, and Yunzhu Li, “LeFlexiTac: Giving Robots a Sense of Touch,” Columbia University RoboPIL Blog, 2026. [Link]

One Year Later: A Growing FlexiTac Ecosystem

Awesome FlexiTac FlexiTac Homepage Join Discord

Categories
Filter Research

- All Works (25)
- Hardware Design (12)
- Glove/Hand (4)
- Multi-Modal Learning (11)
- Simulation (2)
- Locomanipulation (3)
- Industry (2)
- Dataset (2)

AWESOME LIST

Awesome FlexiTac

The growing ecosystem of research building on FlexiTac, together with the prior works.

Maintained by FlexiTac Team

Have a relevant work? Open a PR or issue to add it.

All Works | Hardware Design | Glove/Hand | Simulation | Multi-Modal Learning | Locomanipulation | Industry | Dataset

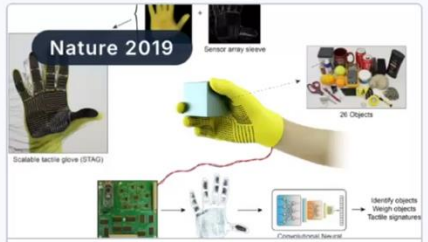

Learning the Signatures of the Human Grasp Using a Scalable Tactile Glove

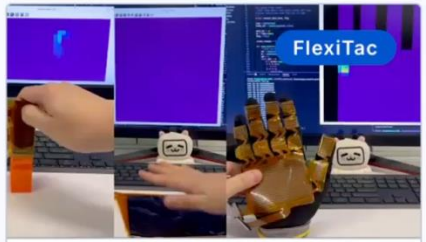

FlexiTac: An Open-Source, Scalable Tactile Solution for Robotic Systems

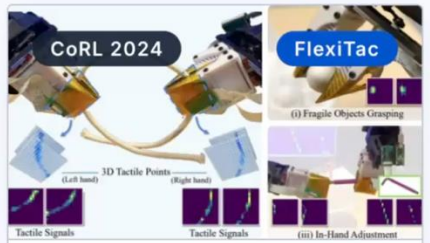

3D-ViTac: Learning Fine-Grained Manipulation with Visuo-Tactile Sensing

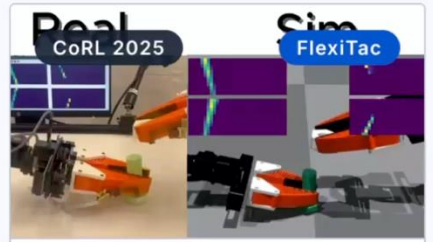

VT-Refine: Learning Bimanual Assembly with Visuo-Tactile Feedback via Simulation Fine-

One Year Later: A Growing FlexiTac Ecosystem



Zhou et al., “Tactile-Reactive Gripper with an Active Palm for Dexterous Manipulation,” Nature Partner Journals Robotics 2026.



Chen et al., “Multi-Modal Manipulation via Multi-Modal Policy Consensus,” ICRA 2026.

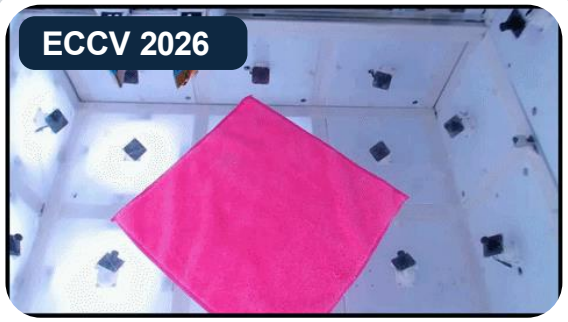

Li et al., “Deform360: A Massive Multi-view Visuotactile Dataset for Deformable World Models,” ECCV 2026.



Zhou et al., “Learning Tactile-Aware Quadrupedal Loco-Manipulation Policies,” IROS 2026.

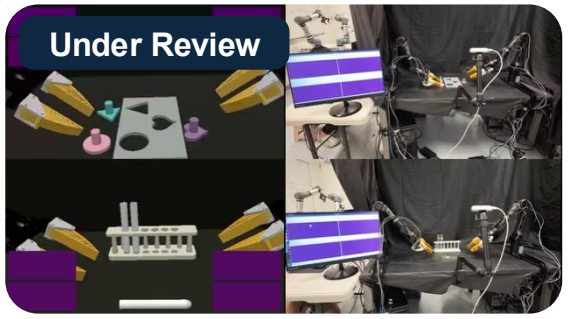

Ma et al., “TacCoRL: Integrating Tactile Feedback into VLA via Simulation,” 2026.



Lin & Zhao, “ART-Glove: Articulated Tactile Glove for Contact-Grounded Dexterous Interaction Capture,” 2026.

Hardware Perception Data

Acknowledgement

Massachusetts Institute of Technology
PURDUE UNIVERSITY
Carnegie Mellon University
Penn
UC San Diego
Georgia Tech
UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN
USC
COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK
ANALOG DEVICES
SAMSUNG RESEARCH AMERICA
NVIDIA

Roadmap

✓

Why Touch?

What does touch tell us that vision can't, for humans and for robots?

✓

Tactile Sensing Hardware

How do tactile sensors work, how do you choose one, and what does it take to build one?

3

Learning with Touch

NOW

How do we fuse touch with vision in a policy, then scale it up with data and simulation?

Making Touch Work in a Policy



Perception

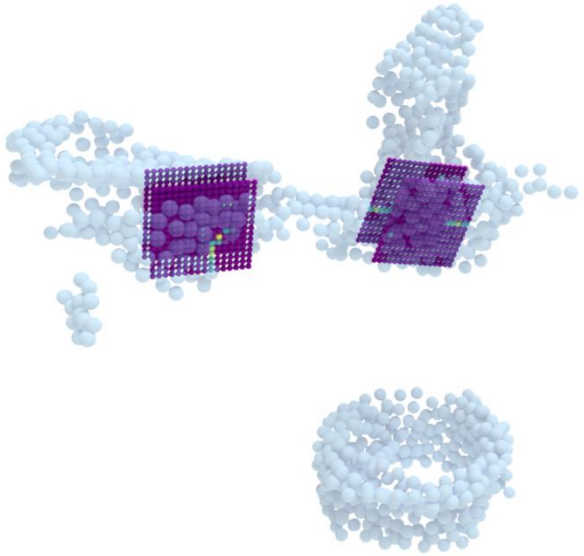

Adding touch should help.

So **why** does performance often **drop**?

Making Touch Work in a Policy



Perception

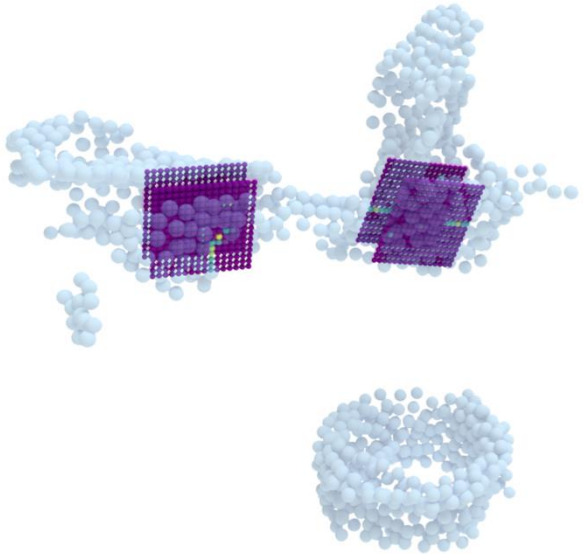

Adding touch should help.
So **why** does performance often **drop**?

THE REASON: the two signals are distinct

VISION

High-dimensional
Global scene
Always informative

TOUCH

Low-dimensional
Local contact
Only during contact

Making Touch Work in a Policy



Perception

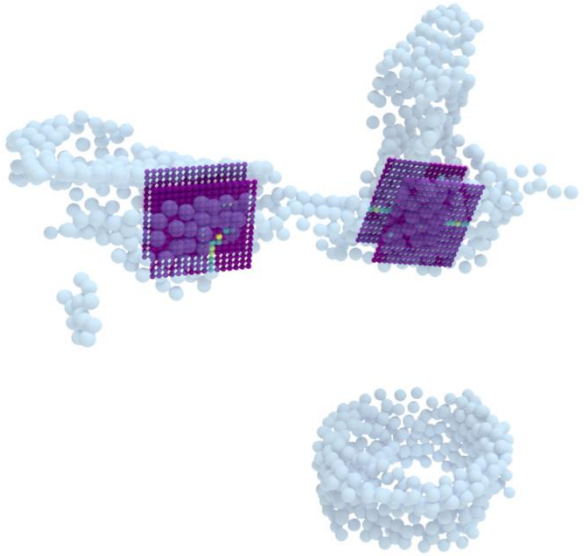

Adding touch should help.
So **why** does performance often **drop**?

THE REASON: the two signals are distinct

VISION

High-dimensional
Global scene
Always informative

TOUCH

Low-dimensional
Local contact
Only during contact

The key is: How we fuse it?

Visuo-Tactile Fusion Is Crucial to the Policy

Two Ways to Fuse Vision and Touch

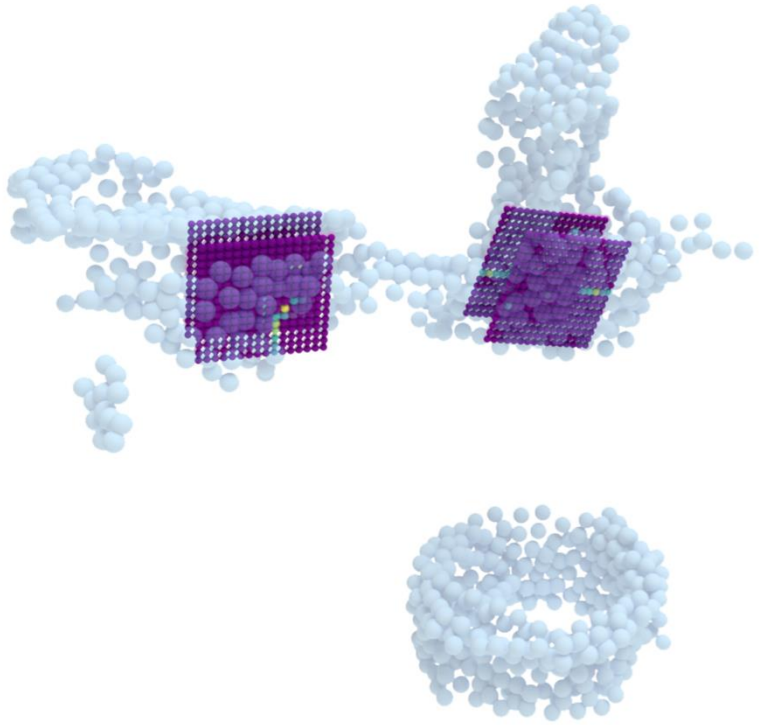

3D Tactile Fusion[1] (before policy)

[1] Huang et al., “3D-ViTac: Learning Fine-Grained Manipulation with Visuo-Tactile Sensing,” CoRL 2024.

[2] Zhu et al., “Touch in the Wild: Learning Fine-Grained Manipulation with a Portable Visuo-Tactile Gripper,” NeurIPS 2025.

Learning with Visuo-Tactile Policy

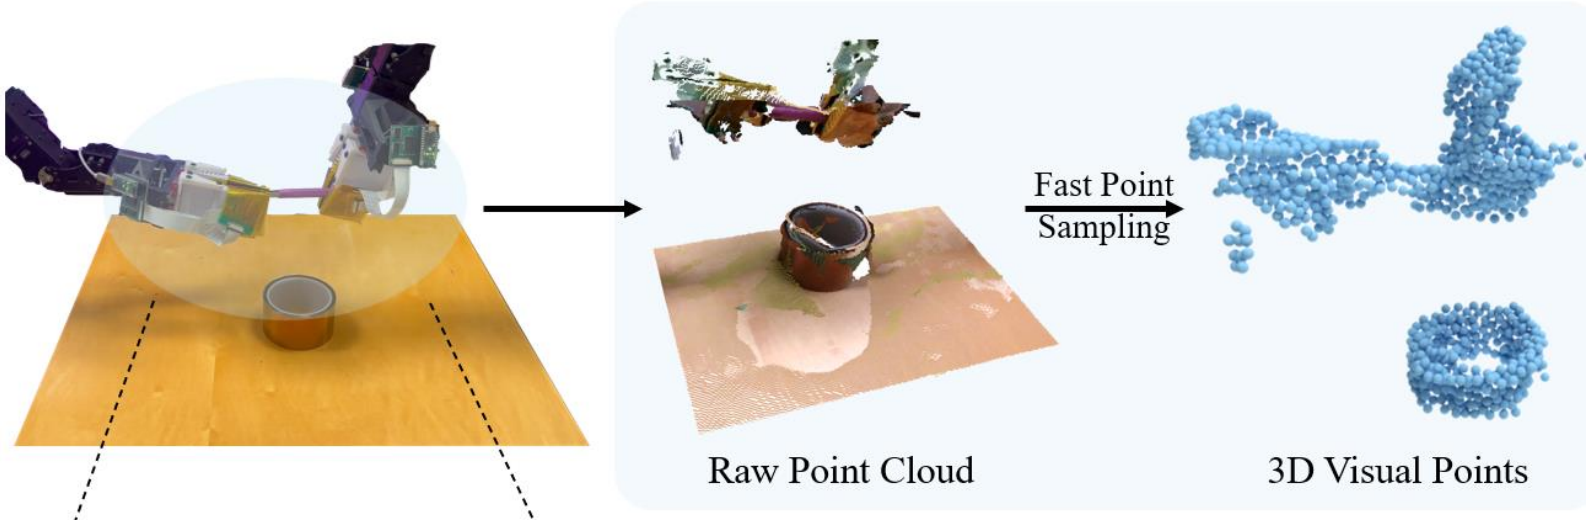

Learning with Visuo-Tactile Policy

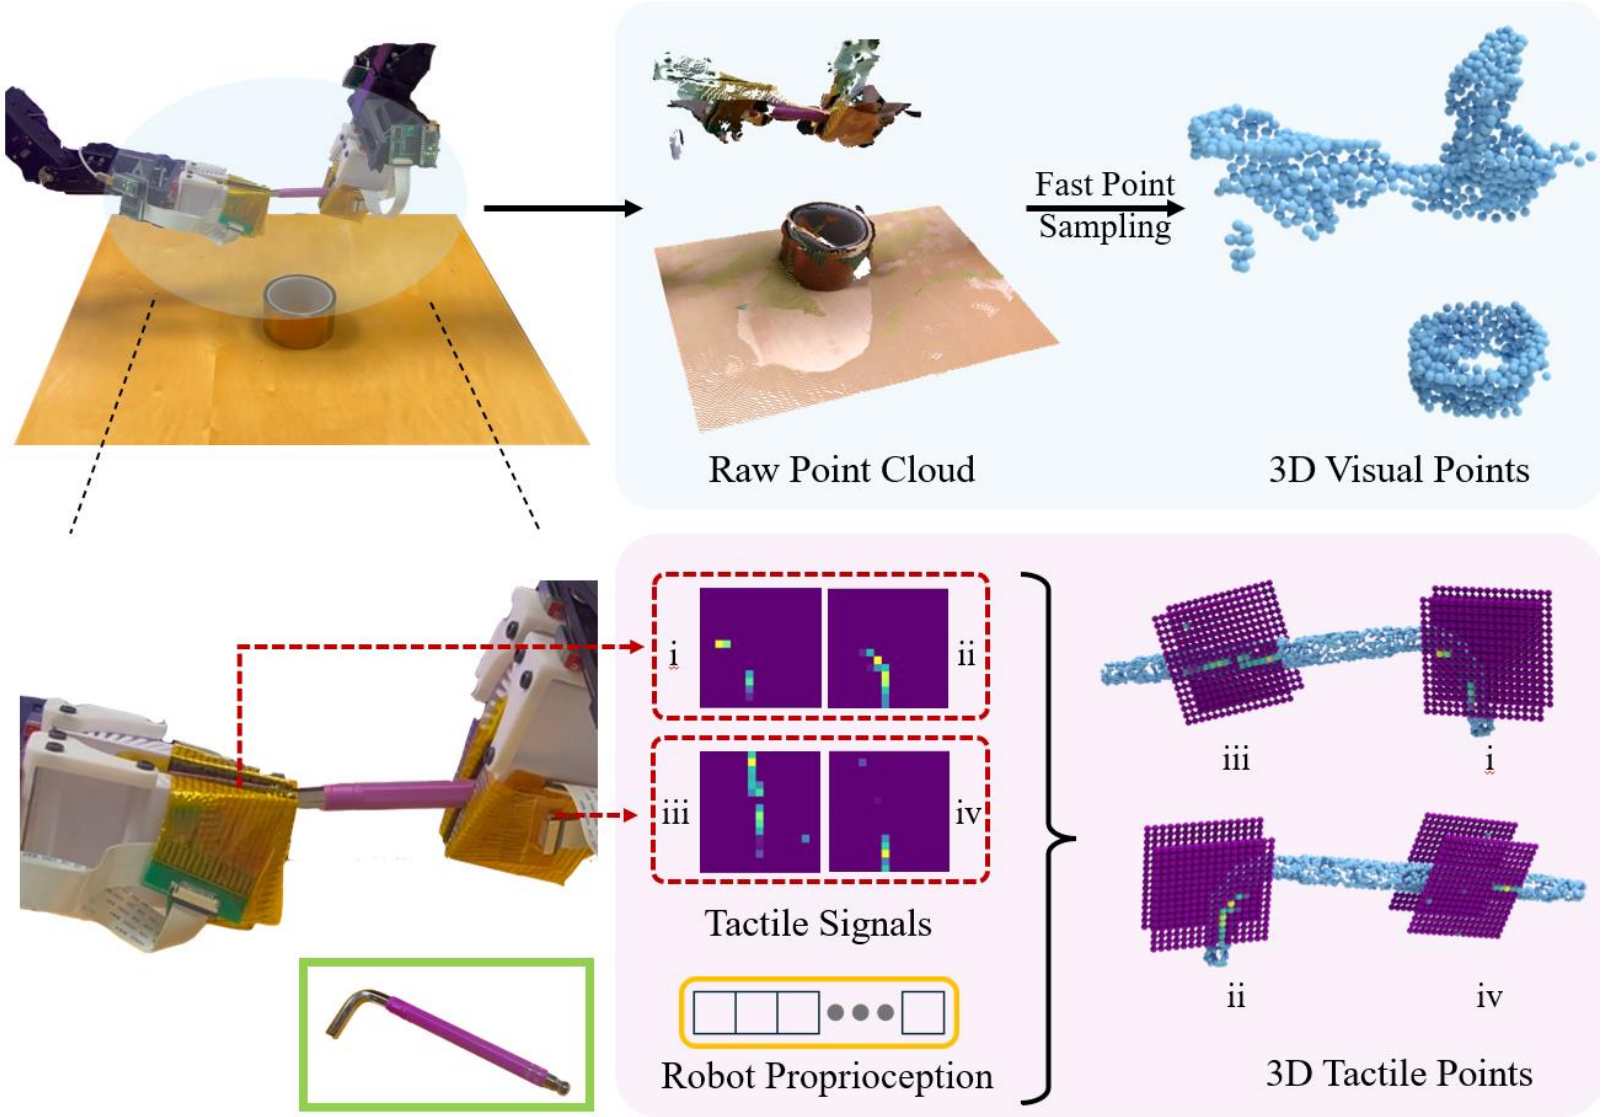

Learning with Visuo-Tactile Policy

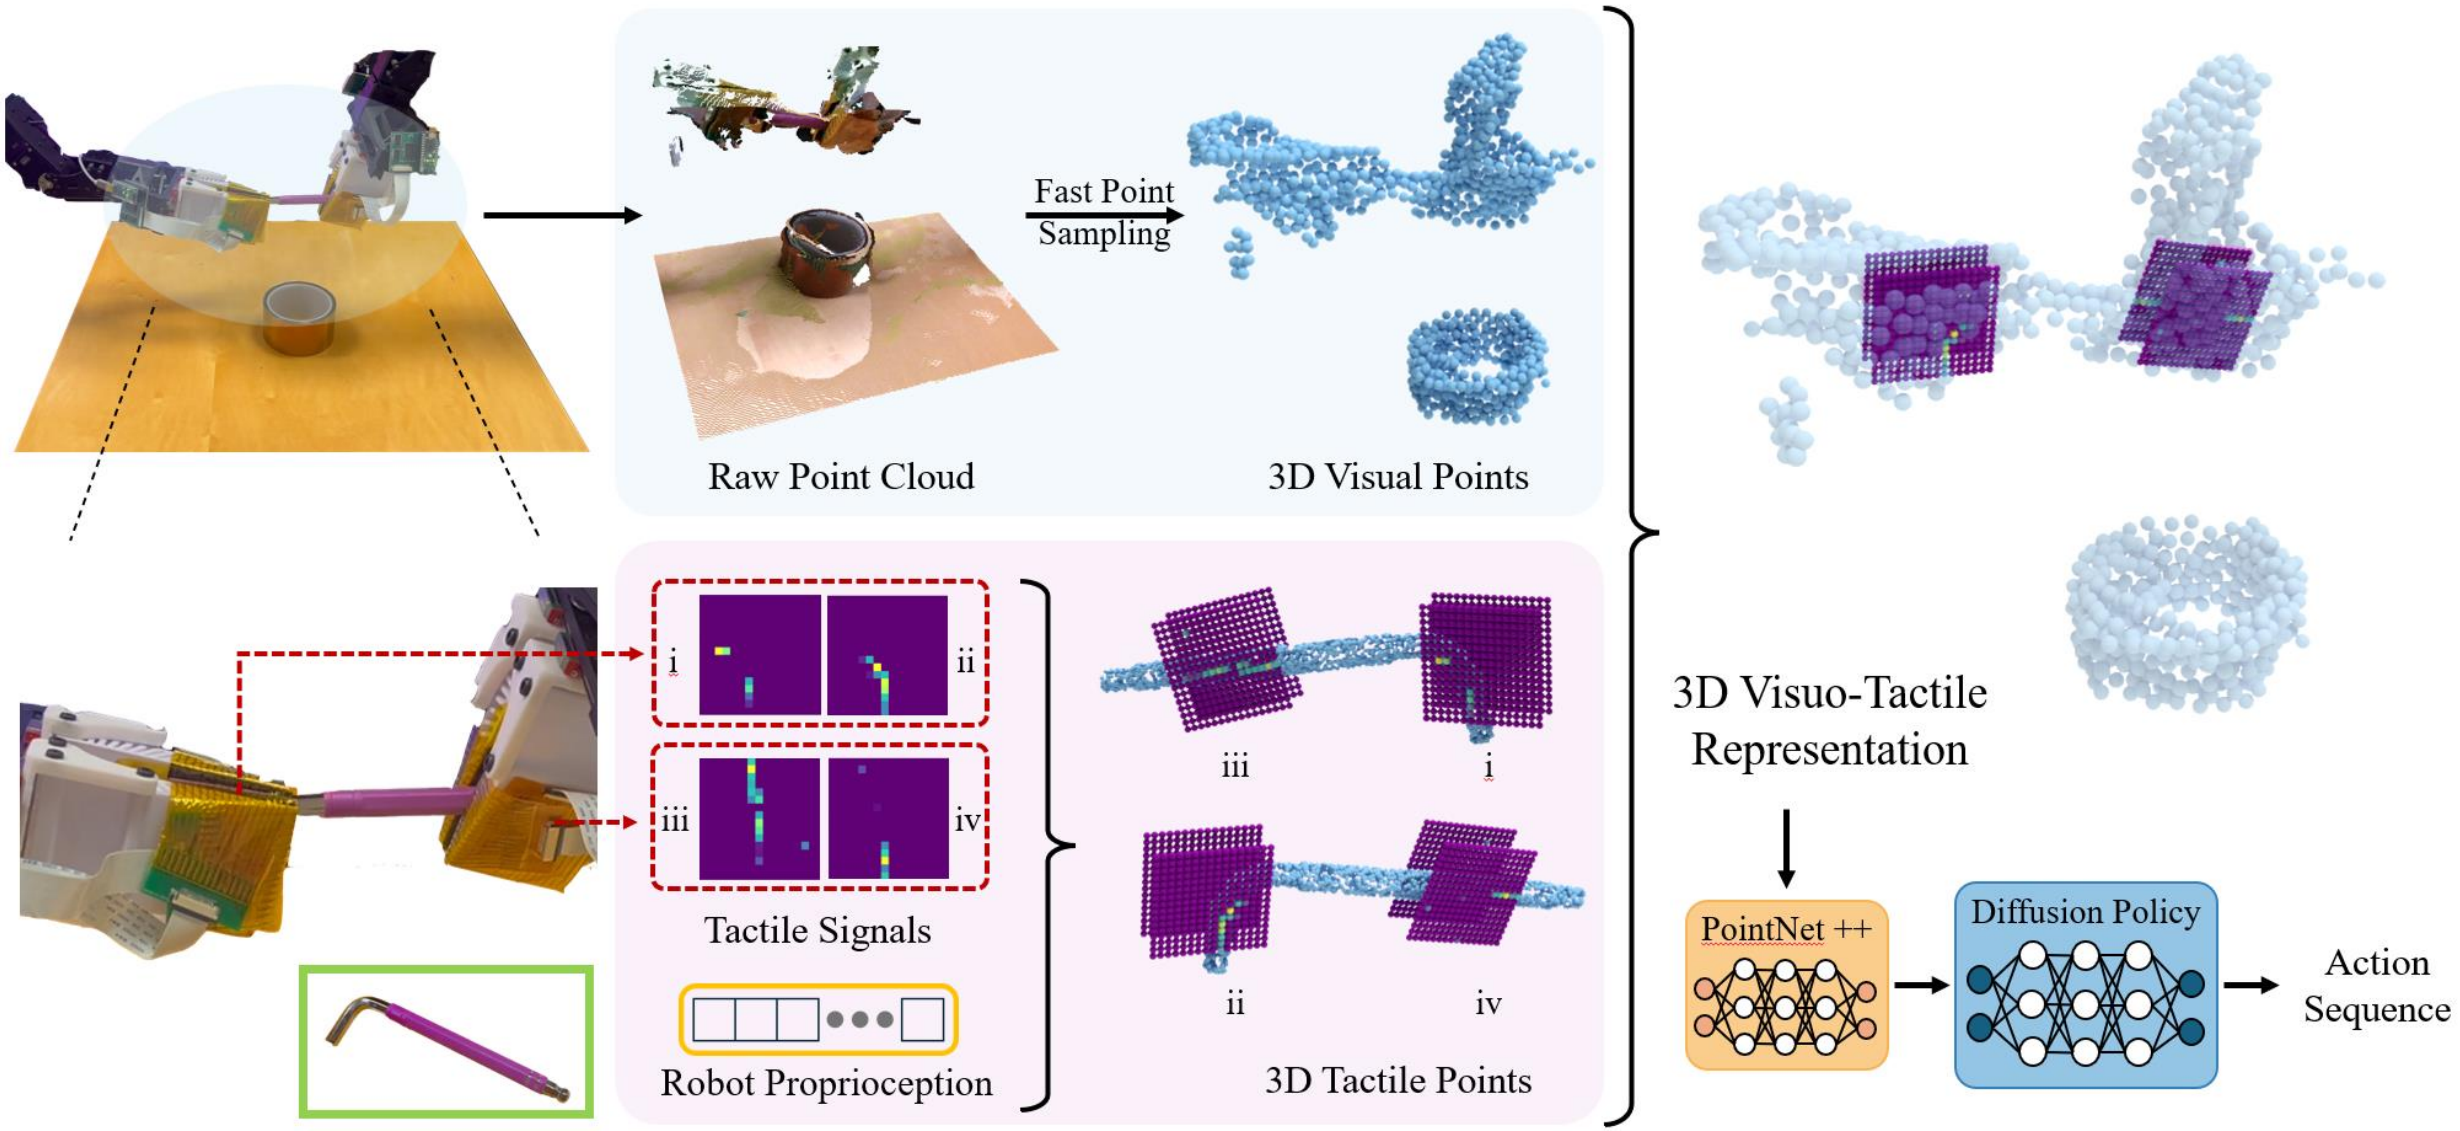


x2 Speed
Occluded Environments
with Fragile Objects


Ours
x2 Speed
In-Hand State Information

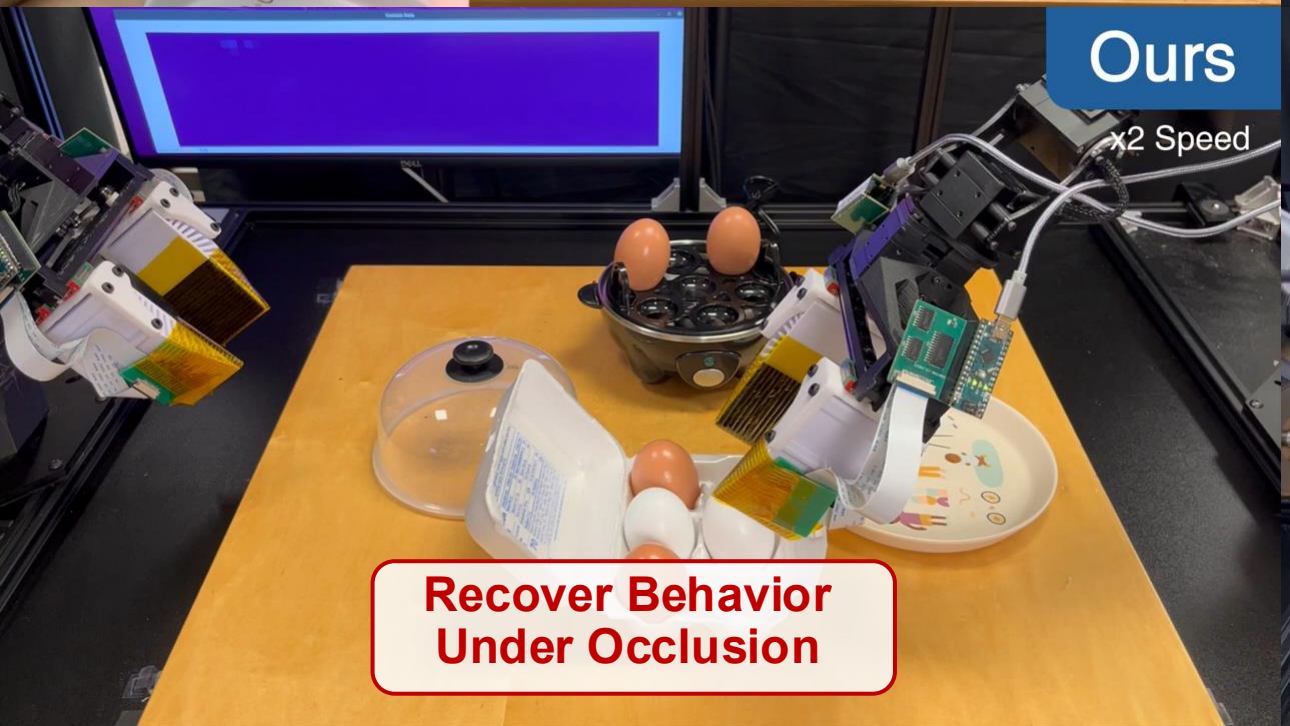
Ours
x2 Speed
Recover Behavior
Under Occlusion

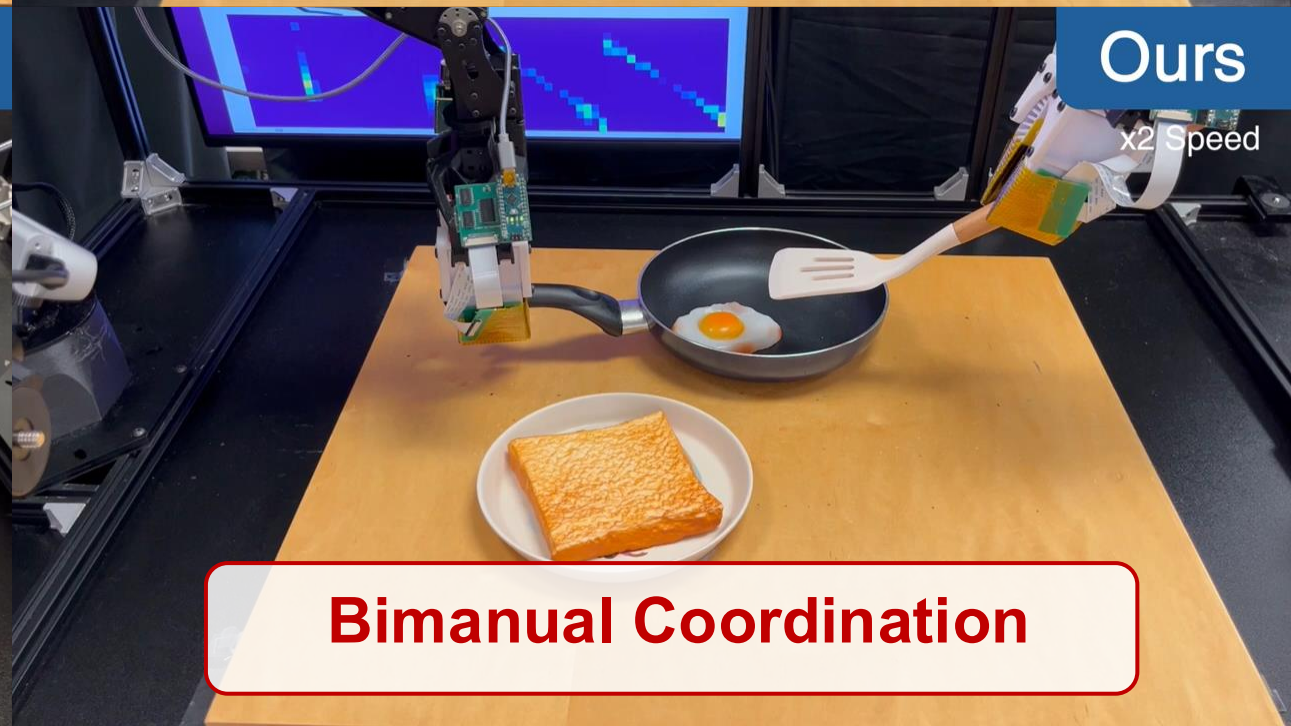
Ours
x2 Speed
Bimanual Coordination

Visuo-Tactile Fusion Is Crucial to the Policy

Two Ways to Fuse Vision and Touch

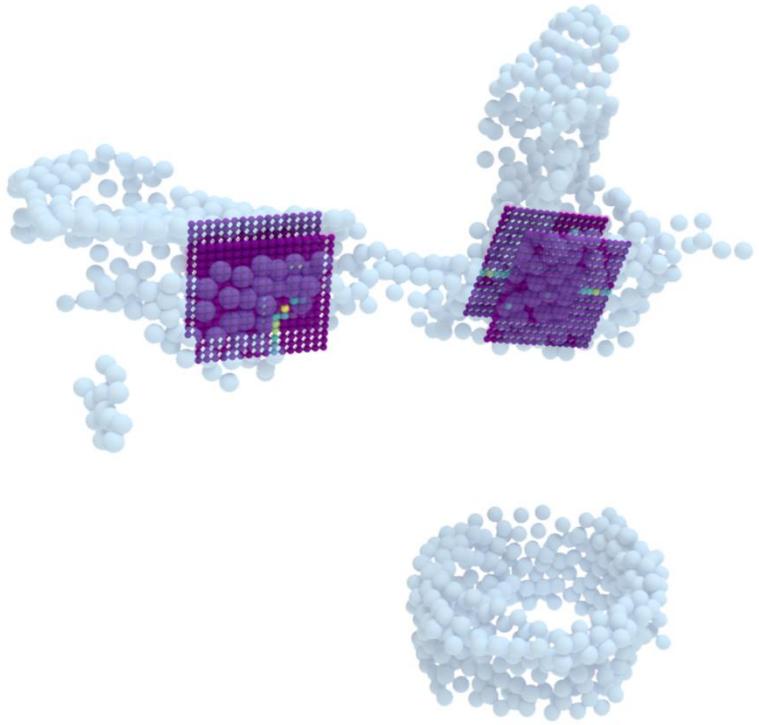

3D Tactile Fusion[1] (before policy)

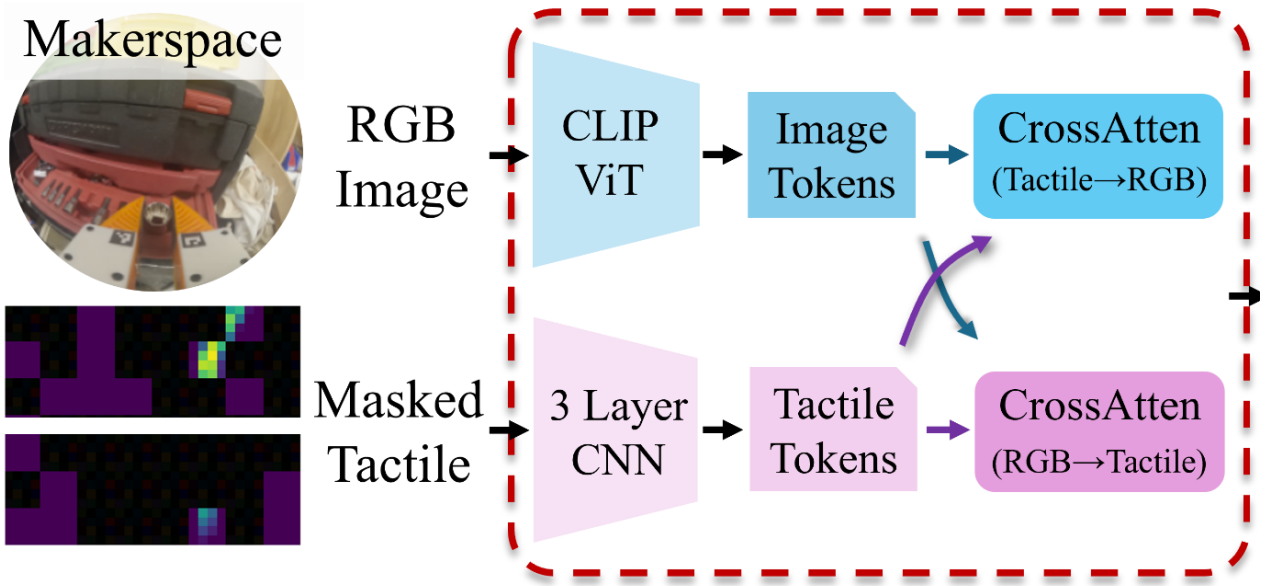

Tactile Fusion Encoder[2] (cross-attention)

[1] Huang et al., “3D-ViTac: Learning Fine-Grained Manipulation with Visuo-Tactile Sensing,” CoRL 2024.

[2] Zhu et al., “Touch in the Wild: Learning Fine-Grained Manipulation with a Portable Visuo-Tactile Gripper,” NeurIPS 2025.

Visuo-Tactile Pre-Training & Downstream Learning

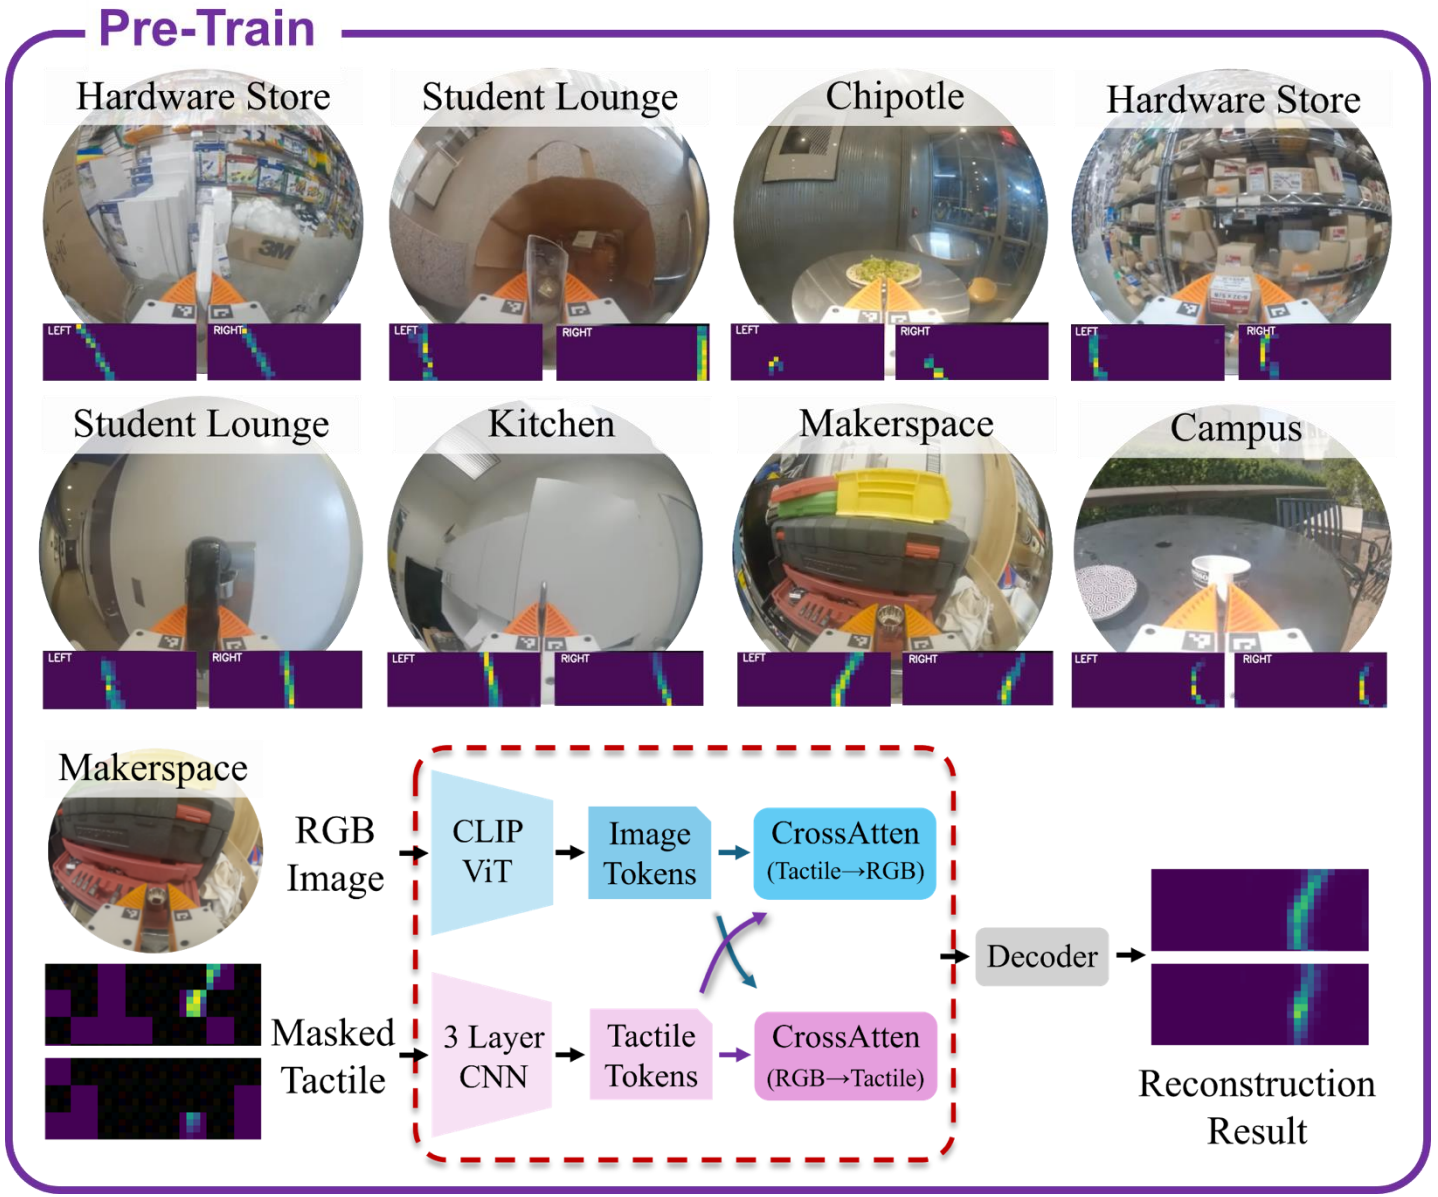

Visuo-Tactile Pre-Training & Downstream Learning

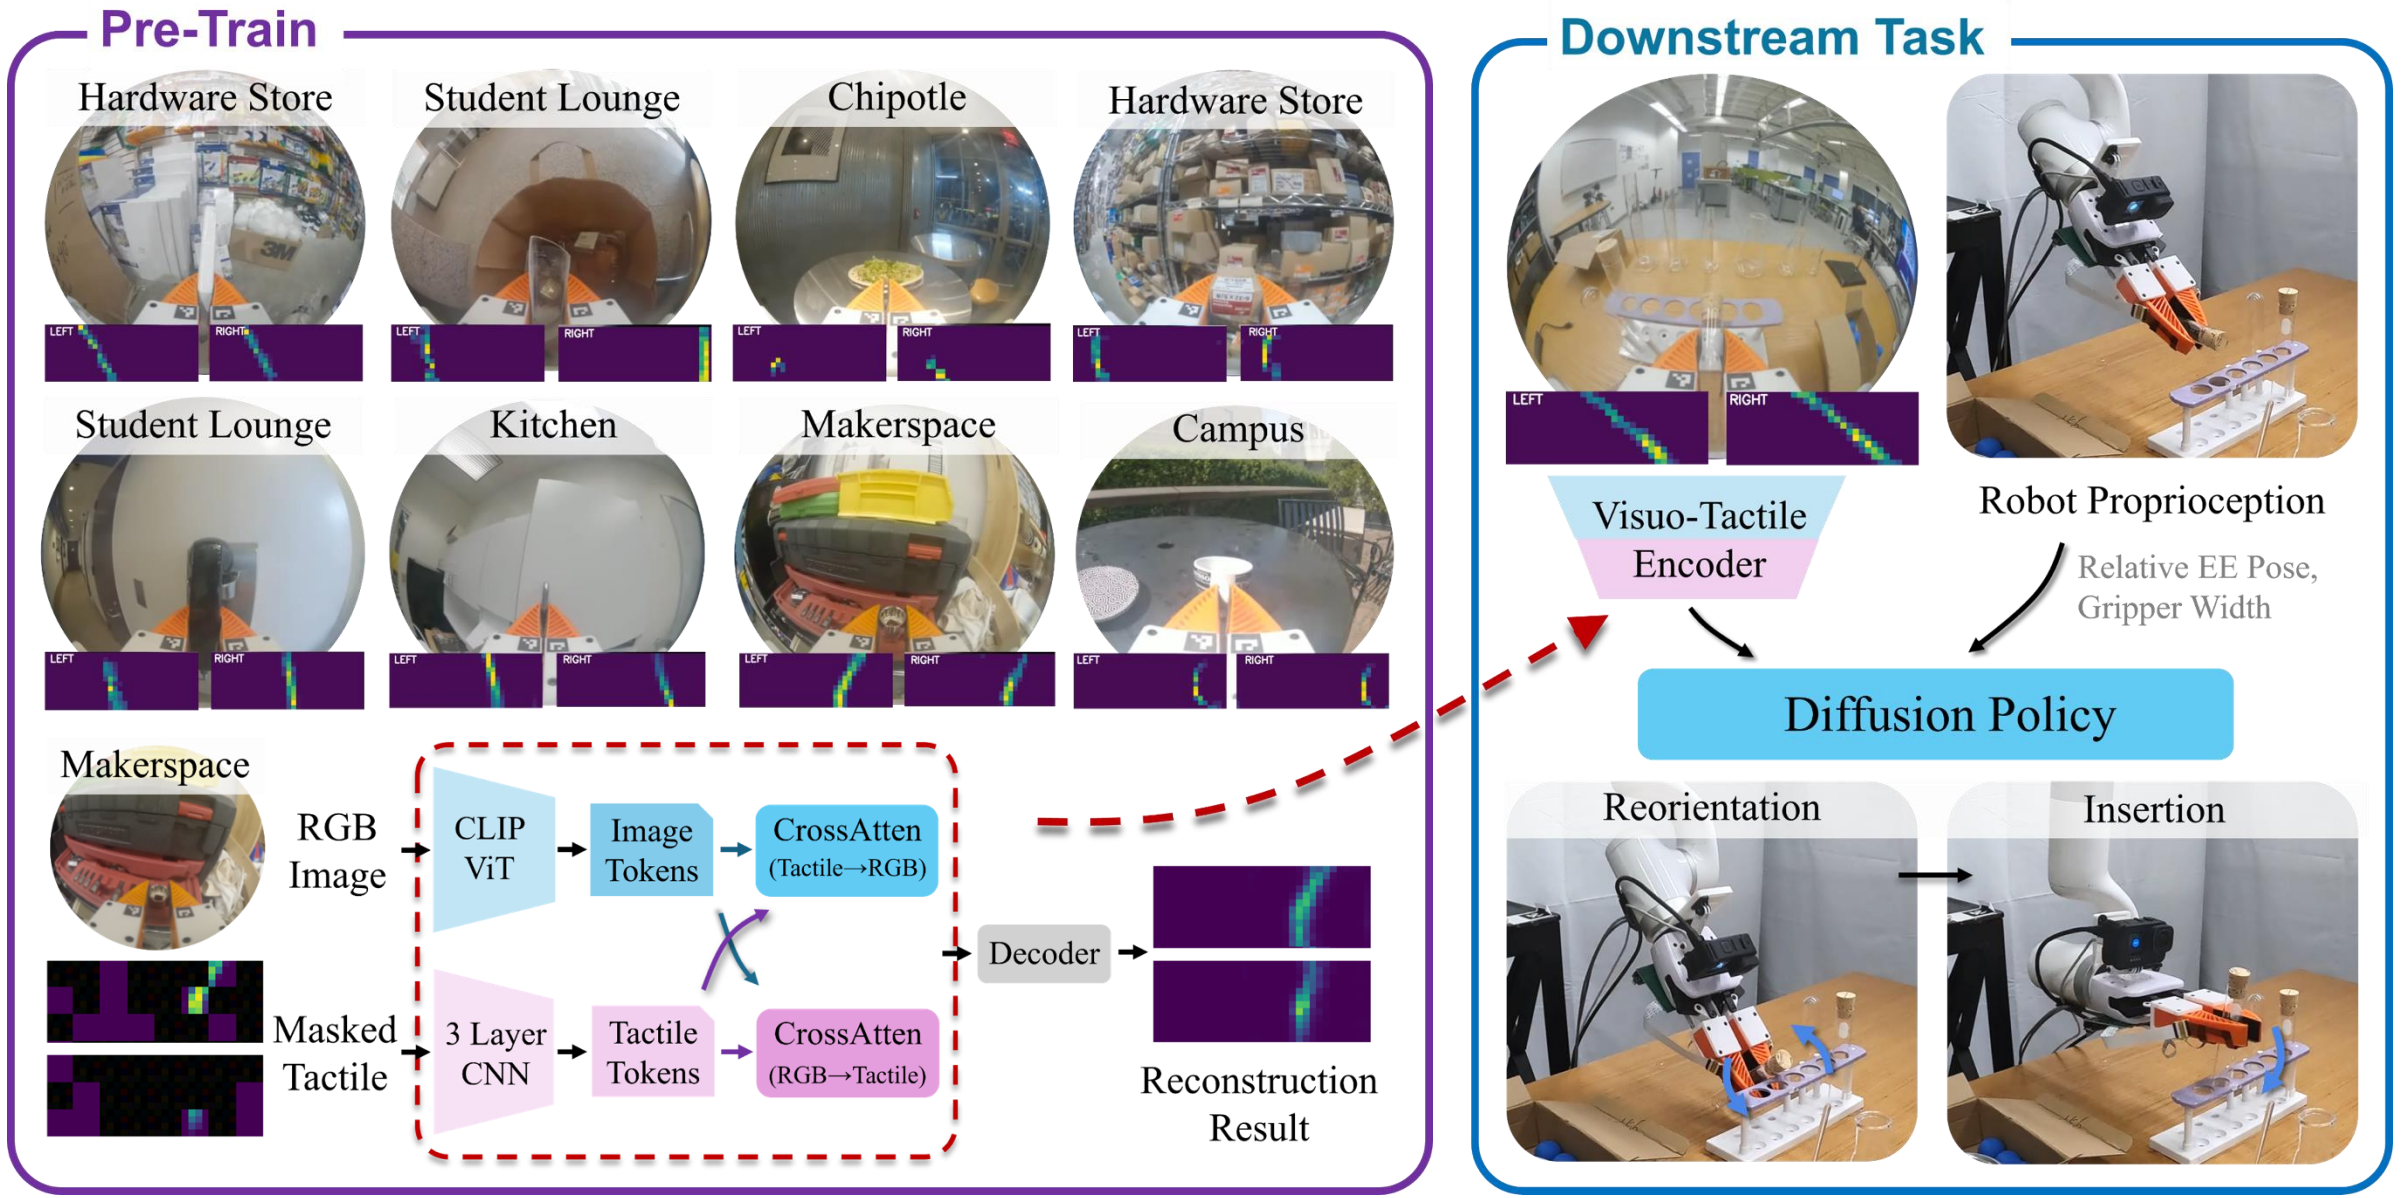

How well does the downstream policy work?

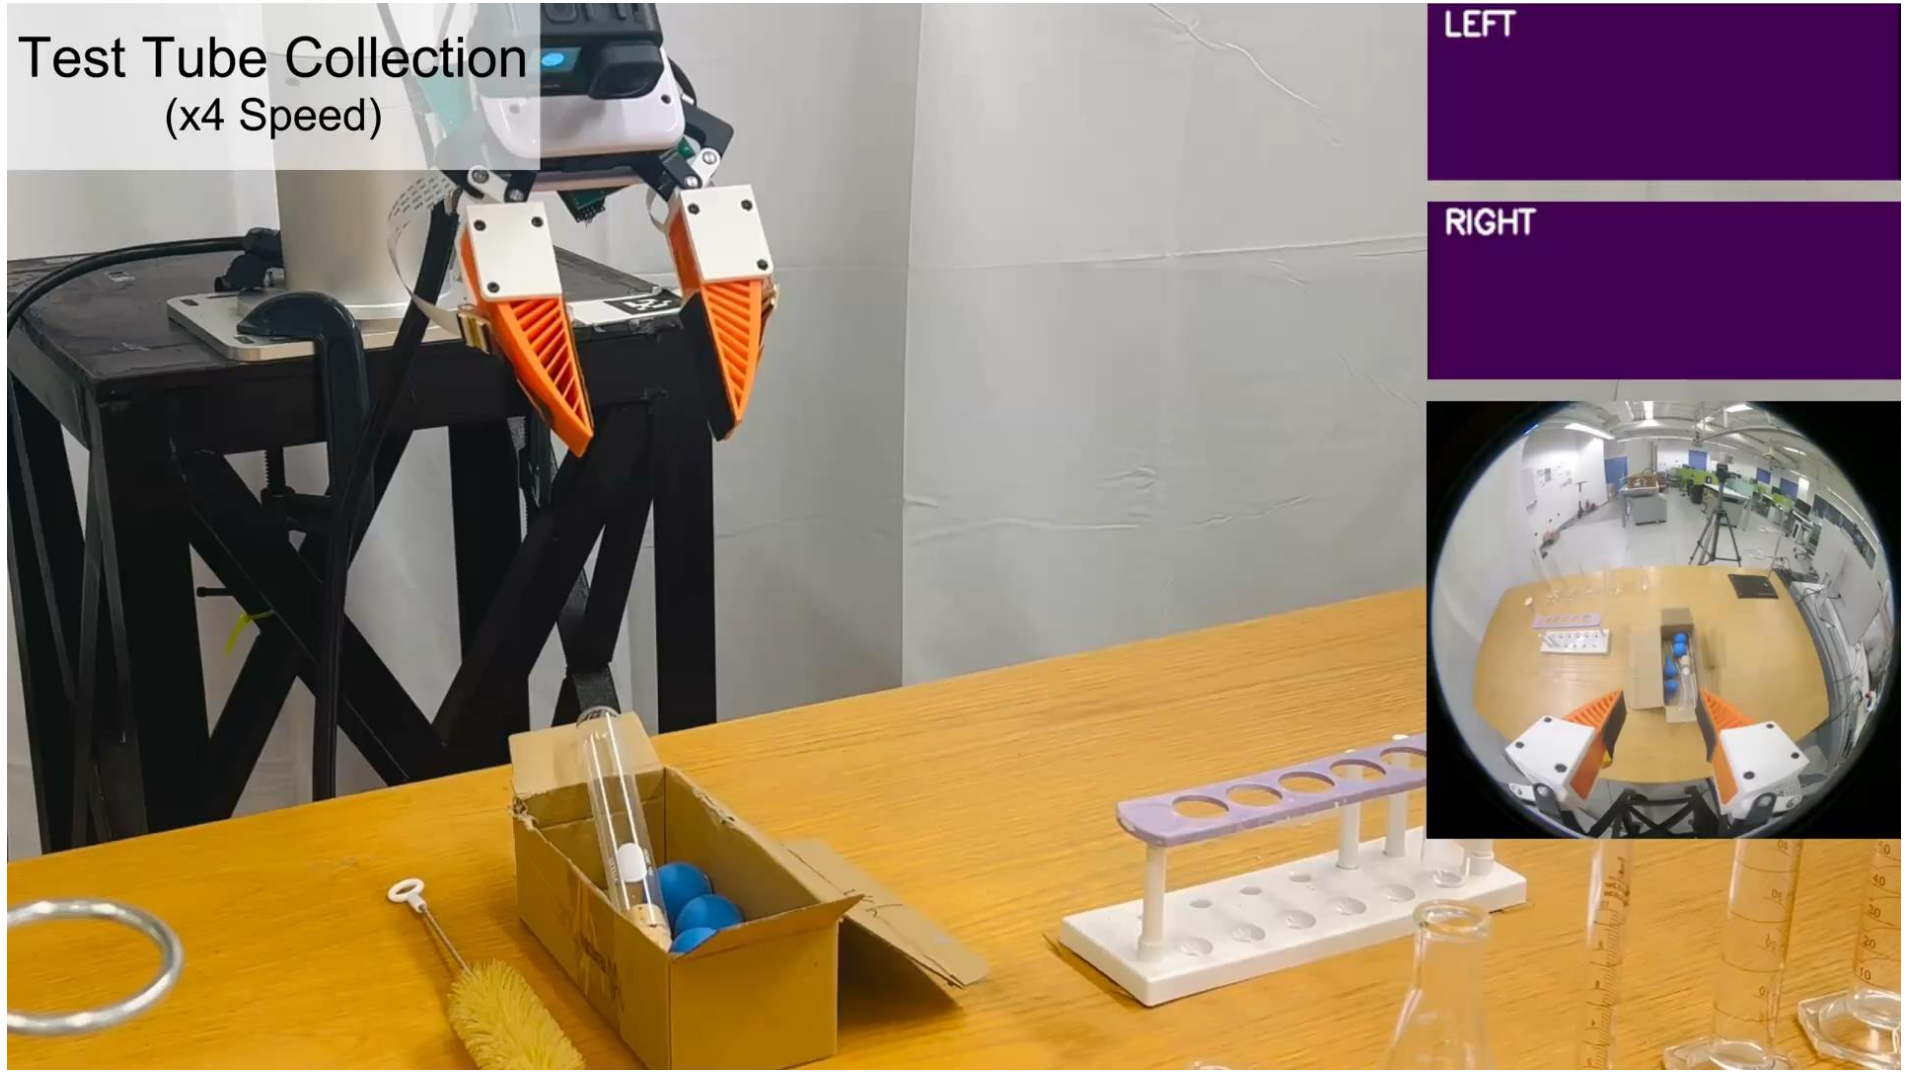

Xinyue Zhu, Binghao Huang, and Yunzhu Li, “Touch in the Wild: Learning Fine-Grained Manipulation with a Portable Visuo-Tactile Gripper,” NeurIPS 2025. **Best Demo Award**, RSS 2025 Workshop on Robot Hardware-Aware Intelligence

Visuo-Tactile Fusion Is Crucial to the Policy

Two Ways to Fuse Vision and Touch

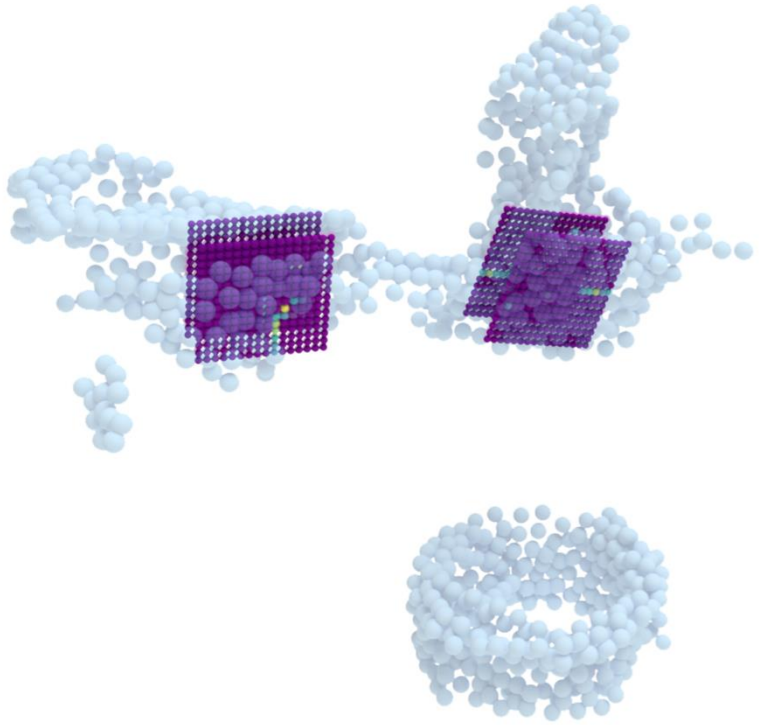

3D Tactile Fusion[1] (before policy)

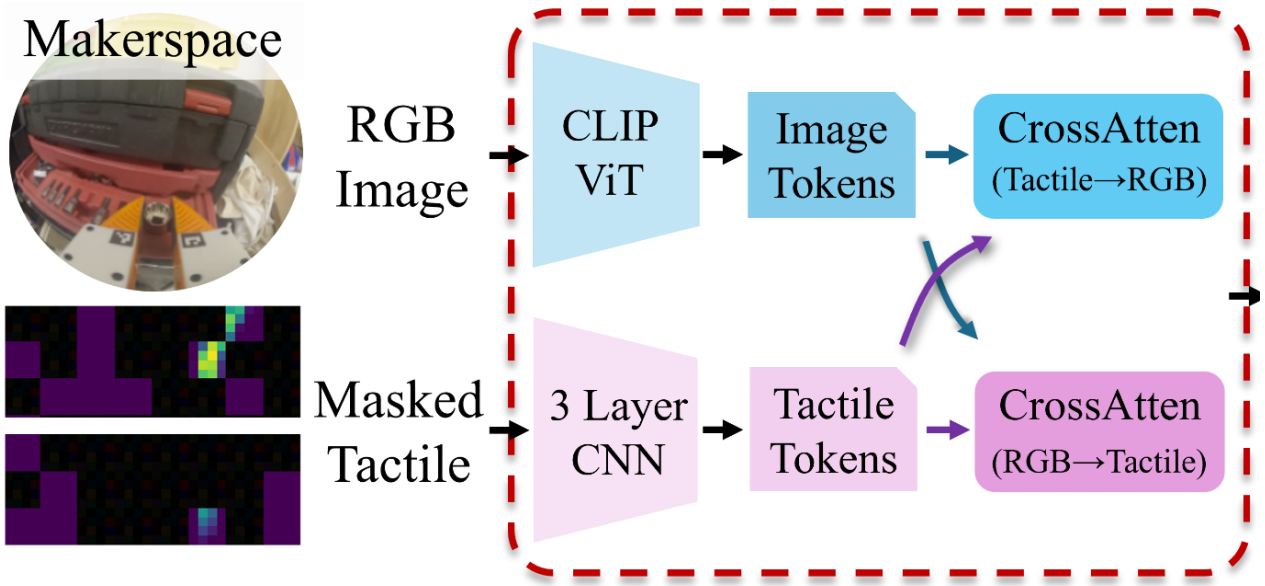

Tactile Fusion Encoder[2] (cross-attention)

[1] Huang et al., “3D-ViTac: Learning Fine-Grained Manipulation with Visuo-Tactile Sensing,” CoRL 2024.

[2] Zhu et al., “Touch in the Wild: Learning Fine-Grained Manipulation with a Portable Visuo-Tactile Gripper,” NeurIPS 2025.

Zooming Out: Where Should Fusion Happen?

EARLY

input level

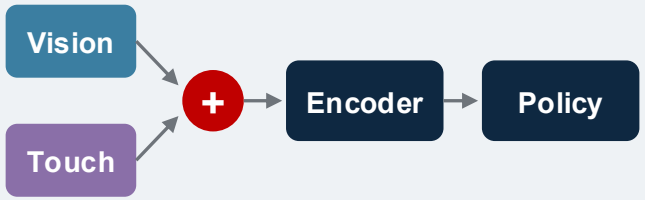

3D-ViTac: tactile → 3D points in the camera's frame

\+ explicit spatial alignment

– needs sensor poses + 3D camera

INTERMEDIATE

feature level

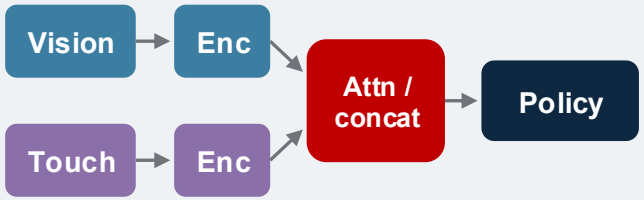

Touch in the Wild: ViT + CNN tokens, cross-attention

\+ flexible; can pre-train

– policy may learn to ignore touch

LATE

decision level

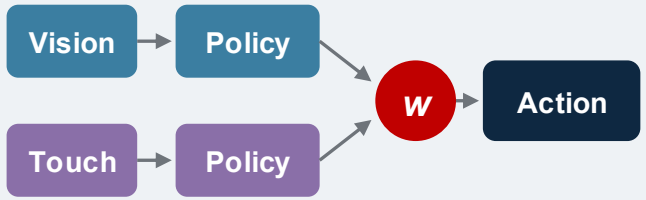

Policy Consensus: one policy per modality + router

\+ modular; easy to add sensors

– weaker cross-modal interaction

No single right answer: it depends on your sensors, your data, and your task.

[1] Huang et al., 3D-ViTac, CoRL 2024. [2] Zhu et al., Touch in the Wild, NeurIPS 2025. [3] Chen et al., "Multi-Modal Manipulation via Multi-Modal Policy Consensus," ICRA 2026.

Scaling Up Tactile Data

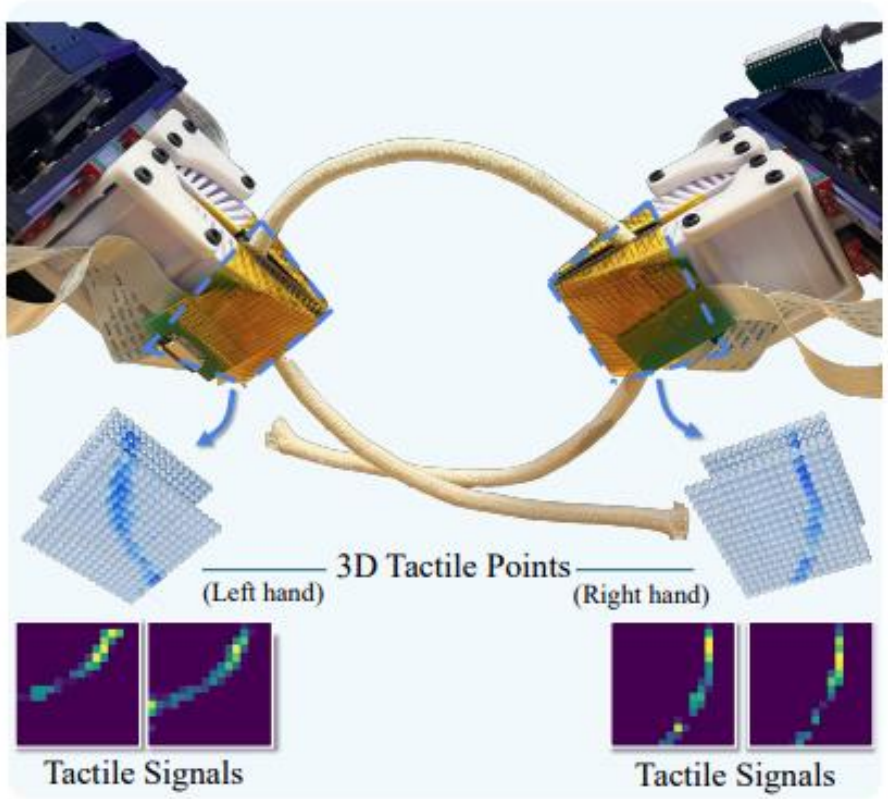

[1] Huang et al., “3D-ViTac: Learning Fine-Grained Manipulation with Visuo-Tactile Sensing,” CoRL 2024.

Scale in Real



Scale in Sim



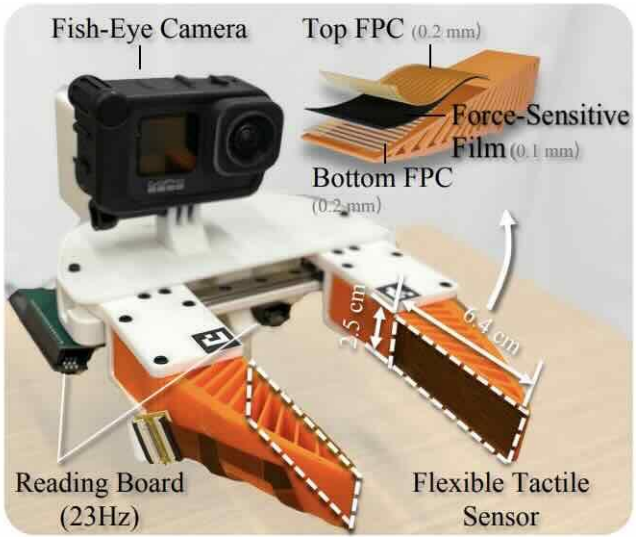

[2] Zhu et al., “Touch in the Wild: Learning Fine-Grained Manipulation with a Portable Visuo-Tactile Gripper,” NeurIPS 2025.

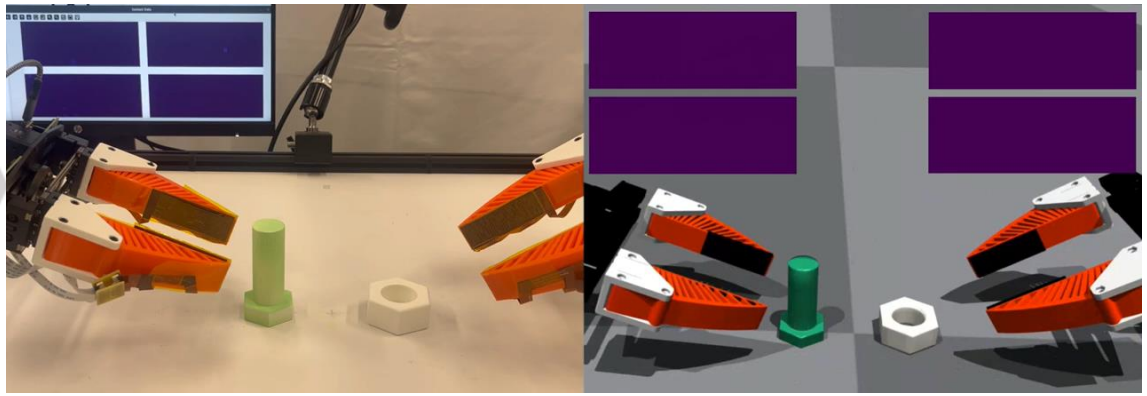

[3] Huang et al., “VT-Refine: Learning Bimanual Assembly with Visuo-Tactile Feedback via Simulation Fine-Tuning,” CoRL 2025.

Touch Is Taking Off

Tactile research is growing fast, and the community is now racing to scale up tactile data.

TACTILE PAPERS PER YEAR (arXiv, cs.RO)

+67% vs. all of 2025, with two months still to go

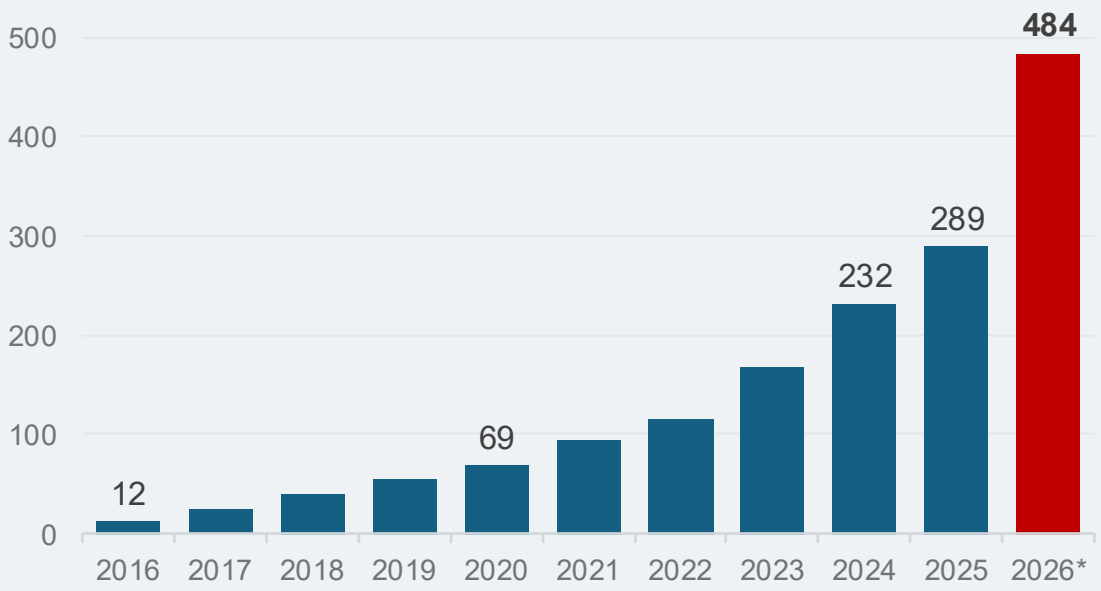

*2026 = Jan 1 to Oct 27. Papers in cs.RO with “tactile” in the title or abstract, counted with the arXiv API.

DATASETS GREW FROM TOUCHES TO HOURS

2019 — **STAG glove** Nature
135k tactile frames, 26 objects

2022 — **Touch and Go** NeurIPS
~13.9k touches, collected in the wild

2024 — **Sparsh** CoRL
460k+ tactile images for pre-training

2025 — **FreeTacMan** ICRA 2026
3M+ visuo-tactile pairs, 10k trajectories

2026 — **TouchScale** arXiv
500 h of human vision + touch

2026 — **FTP-1** arXiv
~3,000 h from 26 sources, 21 sensors

2026 — **TacClaw-Infinity** arXiv
10,000+ h, 1M episodes

2025: make collection cheap. 2026: pool tactile data across sensors, people and sites.

How the Community Is Scaling Up Tactile Data

1 Big datasets & cross-sensor models

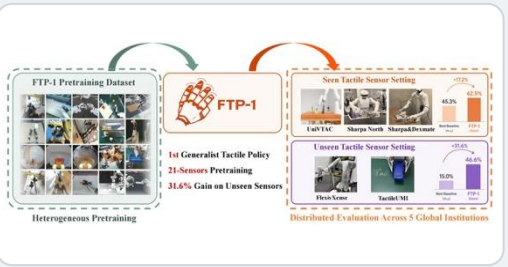

Sparsh CoRL 2024
460k+ tactile images, self-supervised

AnyTouch ICLR 2025
one representation across 4 sensors

FTP-1 arXiv 2026
~3,000 h, 26 sources, 21 sensors

2 Egocentric human touch

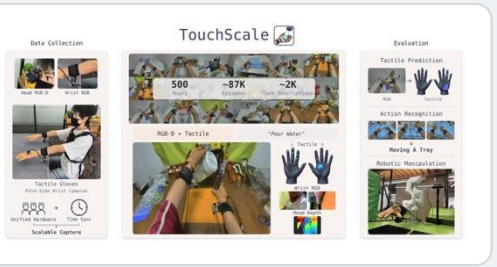

OpenTouch arXiv 2025
full-hand touch + egocentric video

EgoTactile ICML 2026
grasp pressure from egocentric video

EgoTac NeurIPS 2026
5.7M+ image–tactile pairs

TouchScale arXiv 2026
500 h of human vision + touch

3 Tactile gloves → robot skills



STAG Nature 2019
548-sensor knitted glove

Feel the Force arXiv 2025
robots learn from glove demos

OSMO arXiv 2025
12 three-axis sensors, no robot data

ART-Glove arXiv 2026
2,048 taxels, 120 Hz

4 Portable, robot-free collection



ViTaMIn arXiv 2025
handheld gripper with touch

exUMI CoRL 2025
1M+ tactile frames

FreeTacMan ICRA 2026
3M+ pairs, 50 tasks

TacClaw-Infinity arXiv 2026
10,000+ h, 1,000+ task types

The next jump in tactile learning will come from data, not just new sensors.

Numbers as reported by the authors; many 2026 papers are preprints. Figures: FTP-1, TouchScale, OSMO, TacClaw-Infinity papers. arXiv IDs in the speaker notes.

Key Takeaways

1

Why touch?

Touch senses what vision cannot see at contact: force, slip, in-hand state, occluded geometry.

2

Hardware

No sensor has won yet: optical = resolution, magnetic = shear, piezoresistive = thin and scalable. Accessible, reproducible hardware is what lets a field grow.

3

Learning

Touch does not help for free. How you fuse it matters (early / intermediate / late), and pre-training plus tactile simulation help scale beyond a few demos.

Thank You! Questions?

GET STARTED WITH TOUCH

FlexiTac fabrication guide
Open-source hardware: build a tactile sensor in 3 minutes

LeFlexiTac
FlexiTac in the LeRobot ecosystem (Columbia RoboPIL blog, 2026)

Awesome FlexiTac
Community list of hardware, learning, simulation and dataset work

PAPERS FROM TODAY

3D-ViTac

Gelsight

AnySkin

Background: GelSlim, ReSkin, STAG glove

Binghao Huang · binghao.huang@columbia.edu

Even with Touch, Imitation Learning Struggles…

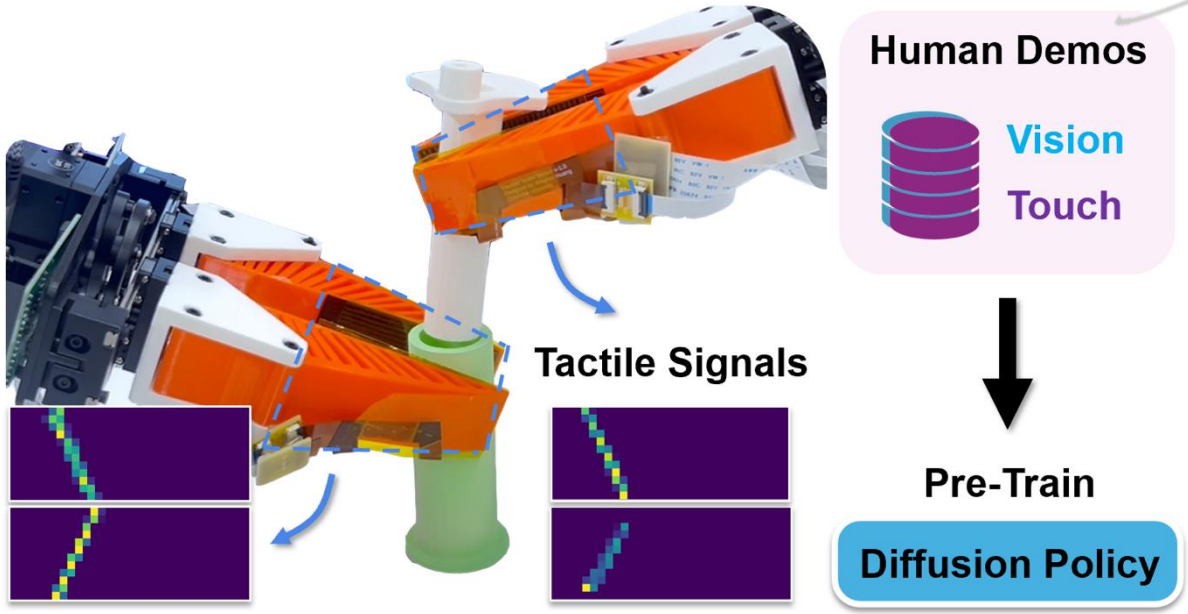

30 Demos

Human Demos

Vision

Touch

Pre-Train

Diffusion Policy

Even with Touch, Imitation Learning Struggles…

30 Demos

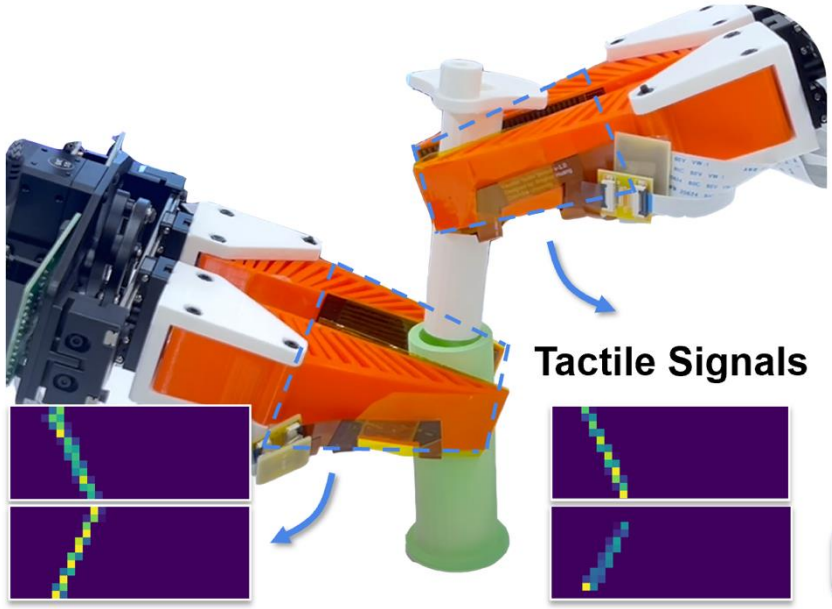

Tactile Signals

Human Demos

Vision

Touch

Pre-Train

Diffusion Policy

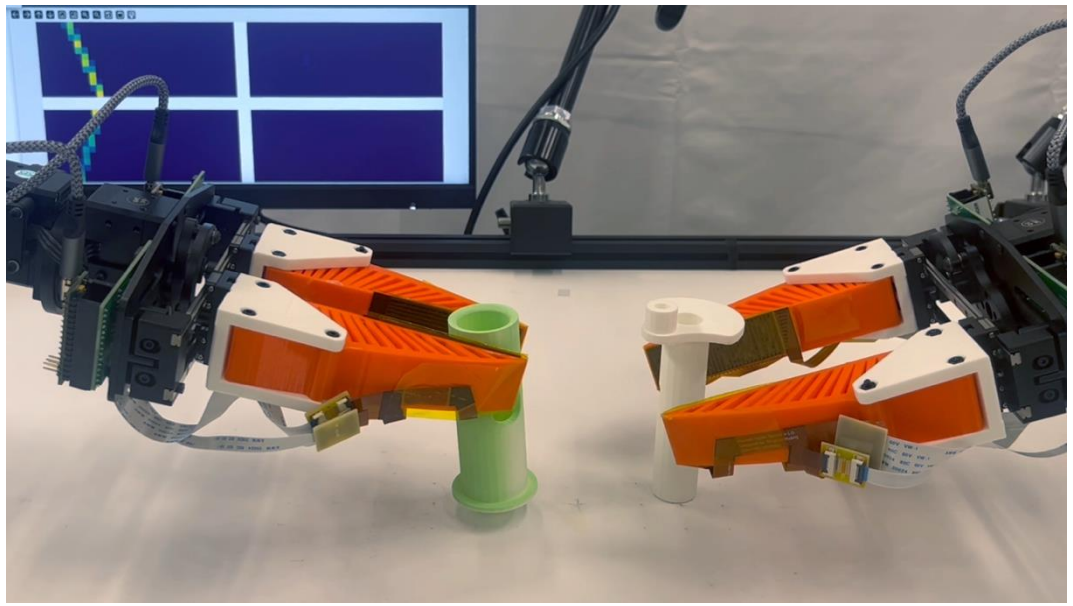

Even with Touch, Imitation Learning Struggles…

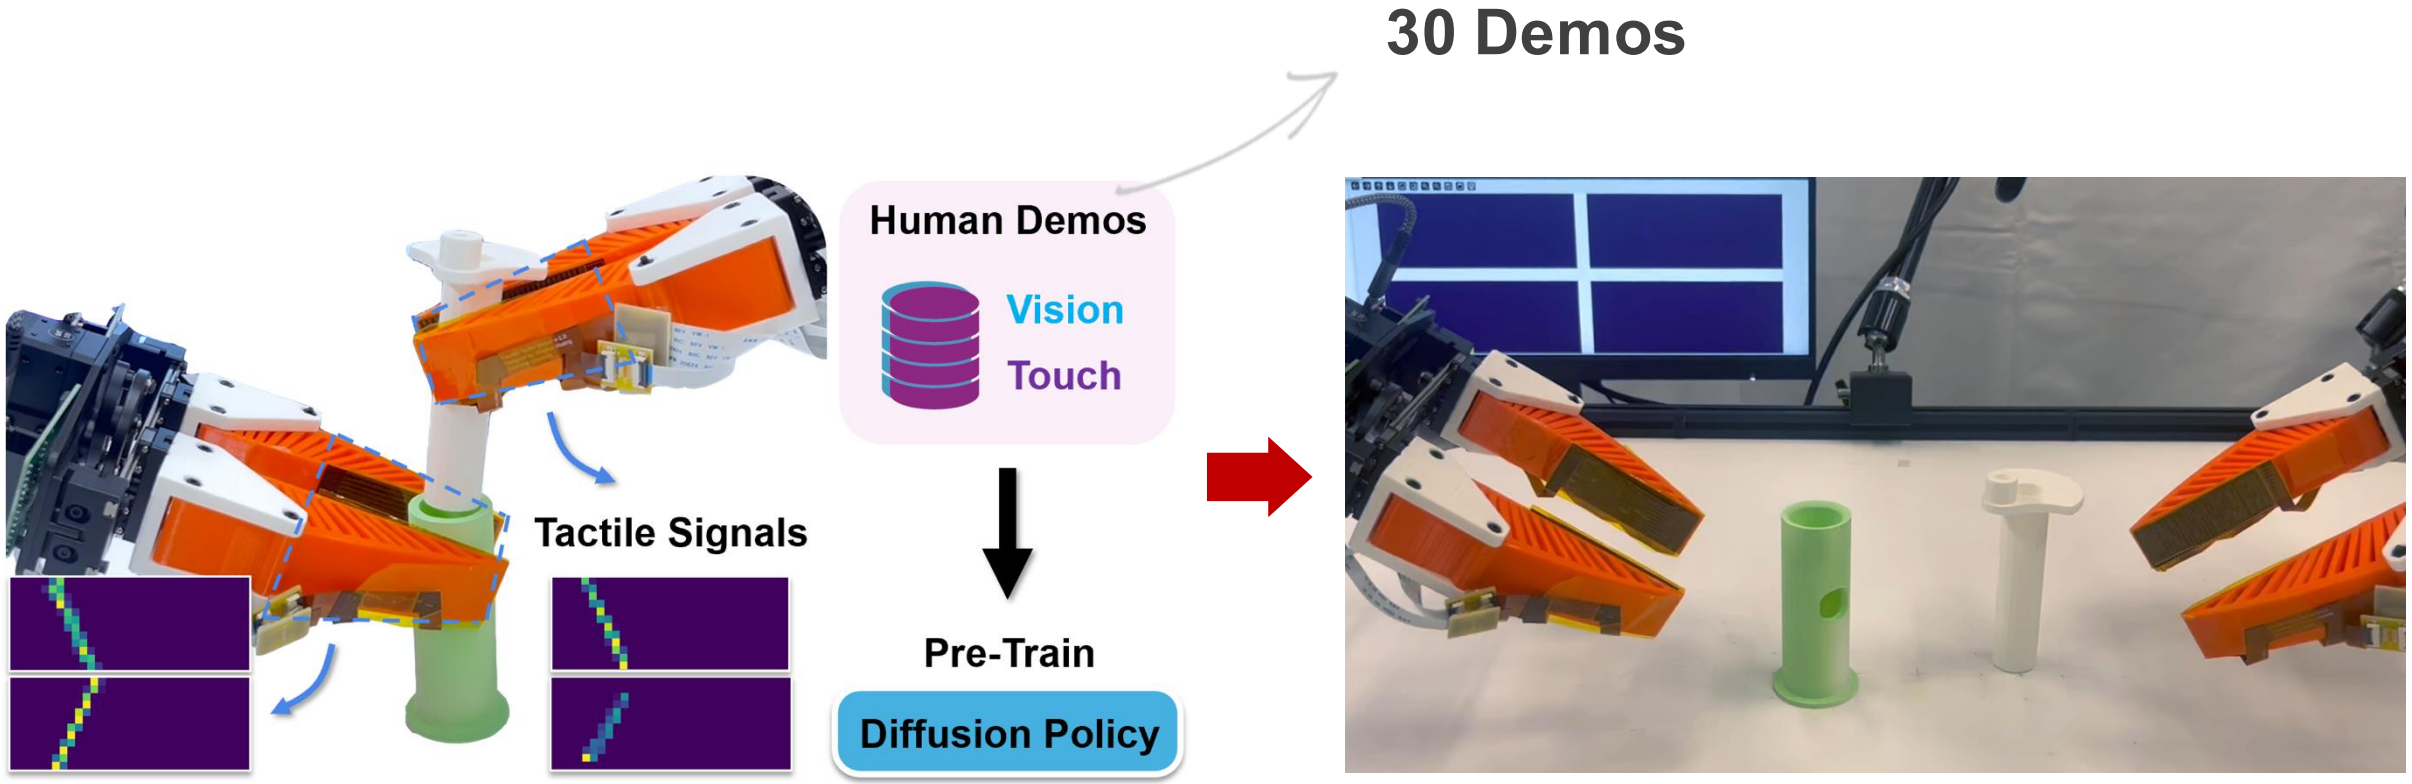

Train on a small amount of real data? Not enough.

Real World

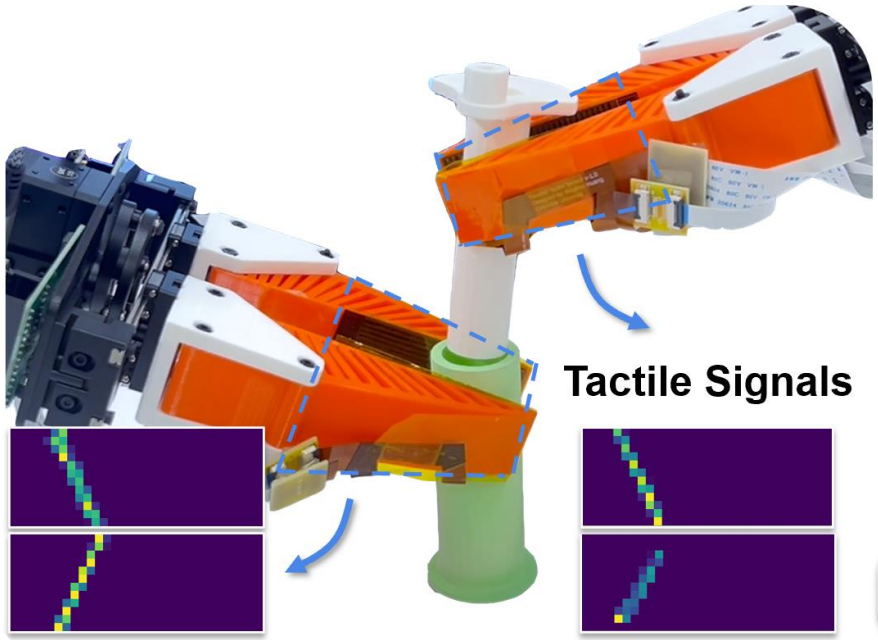

Tactile Signals

Human Demos

Vision

Touch

Pre-Train

Diffusion Policy

Simulation

GPU-Parallelized Tactile Simulation

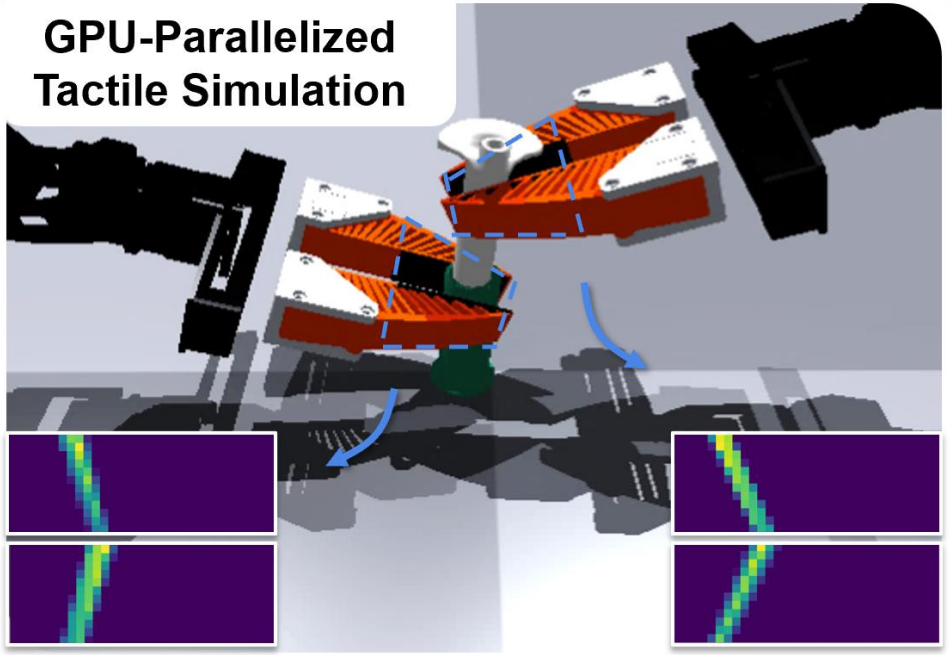

Large-Scale RL Fine-Tuning



Real World

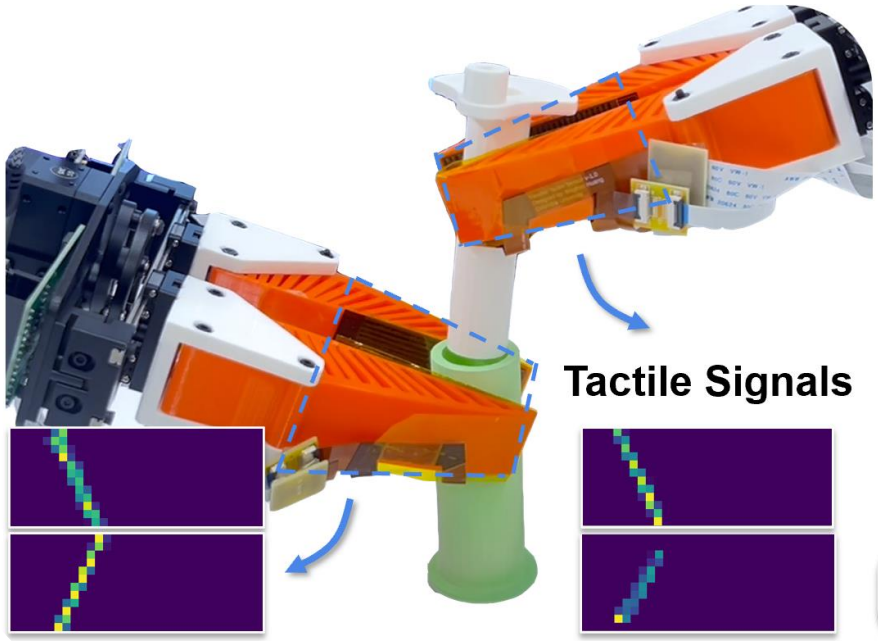

Human Demos

Vision

Touch

Tactile Signals

Pre-Train

Diffusion Policy

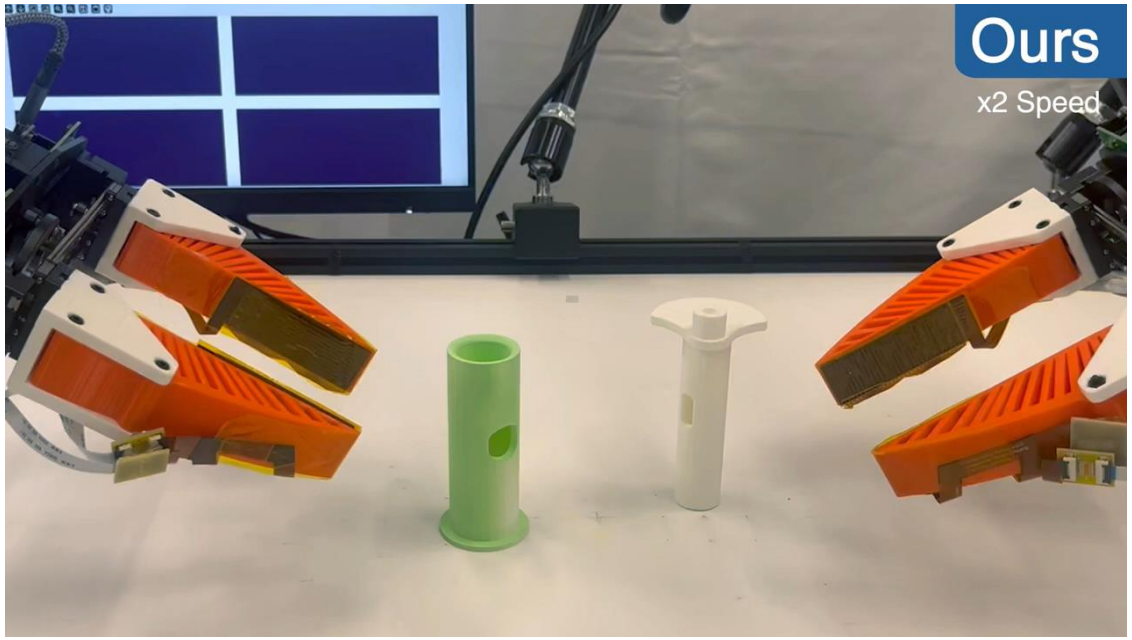

Simulation

GPU-Parallelized Tactile Simulation

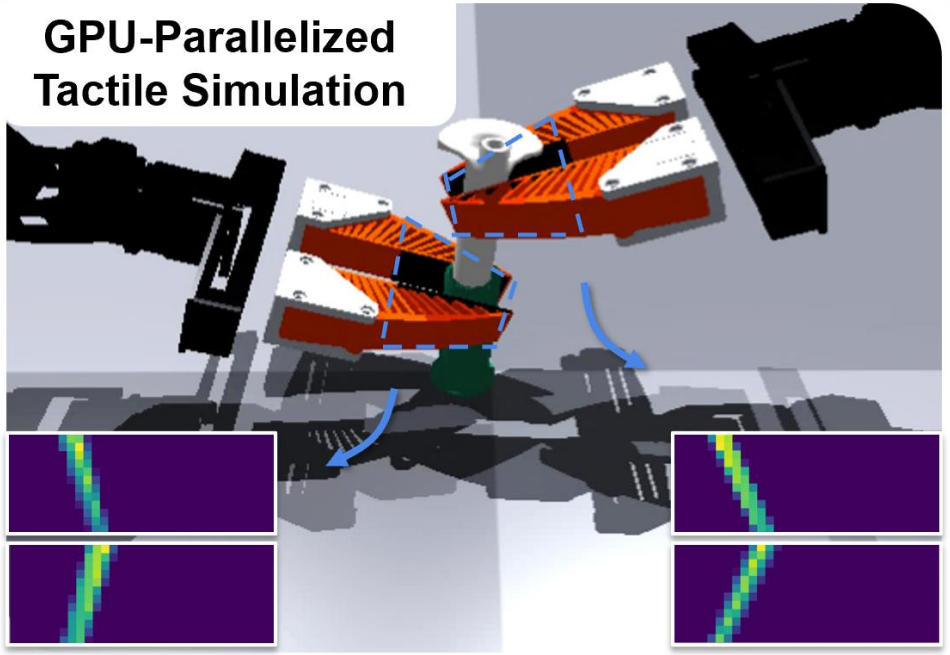

Large-Scale RL Fine-Tuning



Simulating Touch: Real vs. Sim

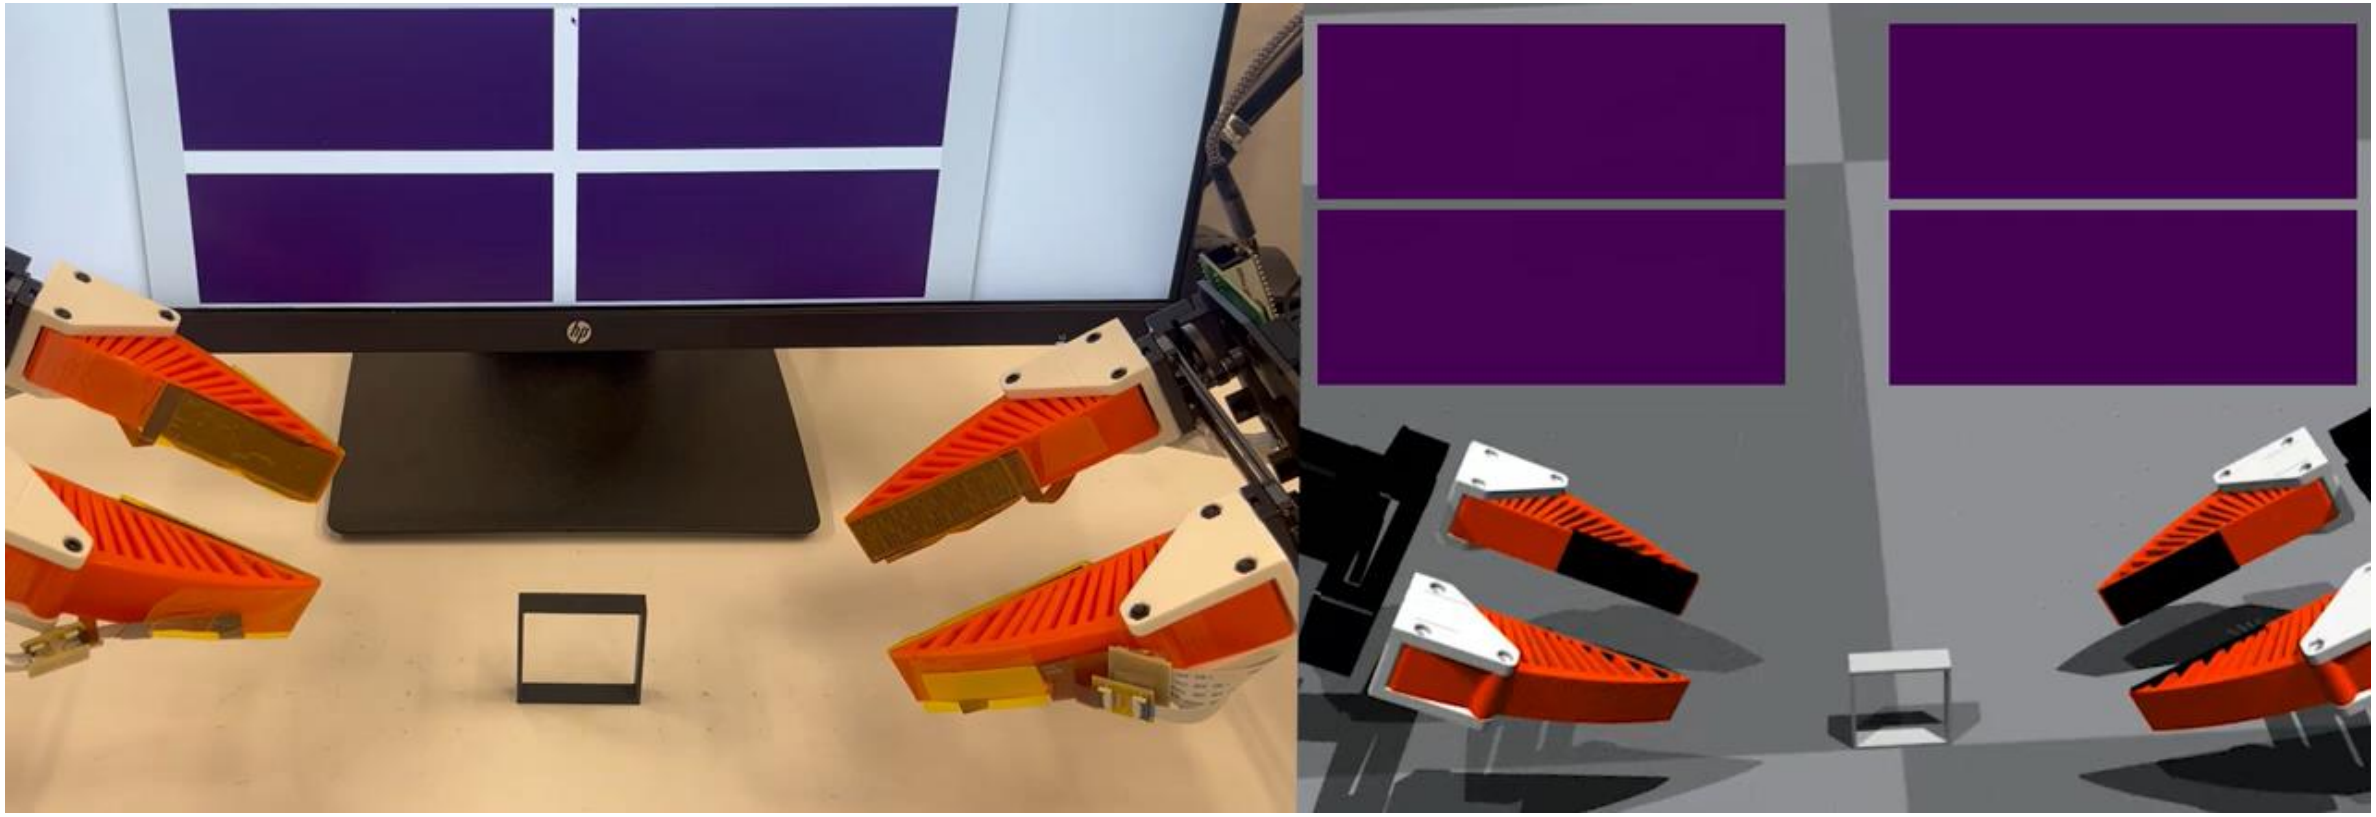

Real

Sim

Binghao Huang, Jie Xu, Iretiayo Akinola, Wei Yang, Balakumar Sundaralingam, Rowland O'Flaherty, Dieter Fox, Xiaolong Wang, Arsalan Mousavian, Yu-Wei Chao, and Yunzhu Li, "VT-Refine: Learning Bimanual Assembly with Visuo-Tactile Feedback via Simulation Fine-Tuning," CoRL 2025.
Best Paper Award, IROS 2025 AHFR Workshop.

Simulating Touch: Each Taxel Is a Spring–Damper

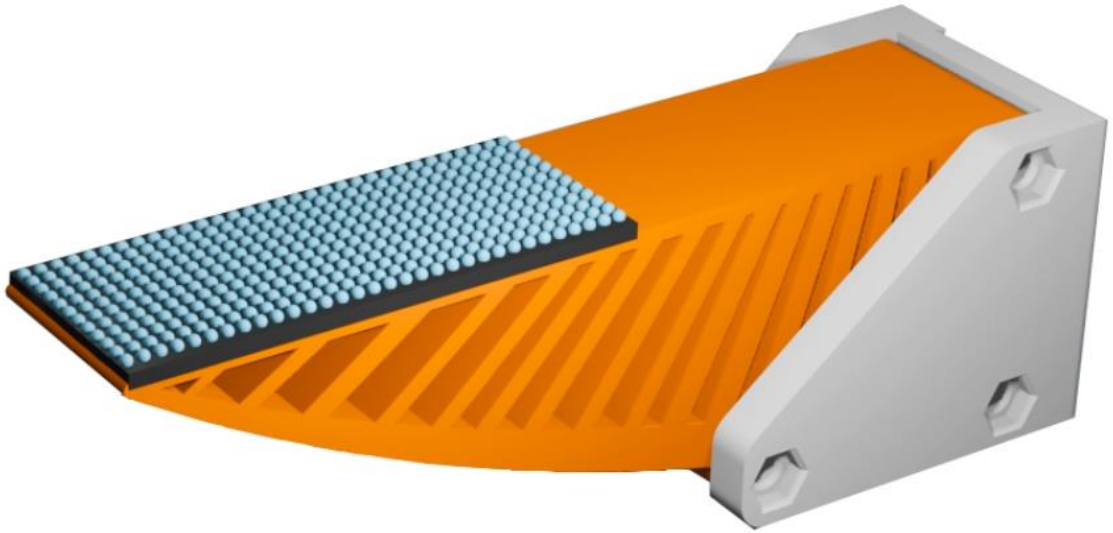

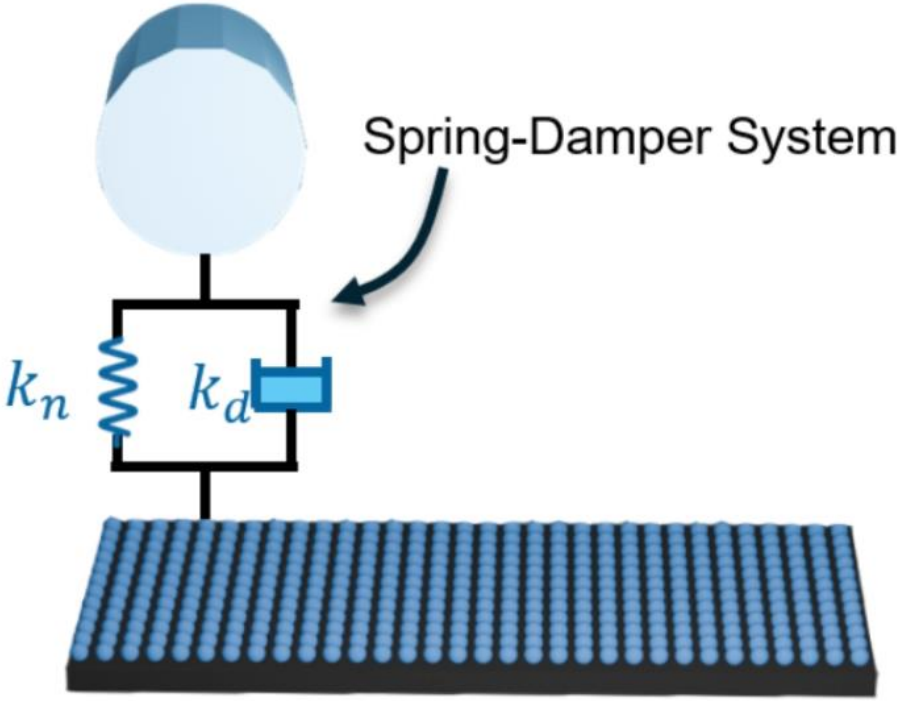

Kelvin–Voigt model:

$$f_n = -\left(k_n d + k_d \dot{d}\right)\boldsymbol{n}$$

- k_n, k_d= stiffness and damping coefficients
- $d, \dot{d}$= penetration depth and velocity

Binghao Huang, Jie Xu, Iretiayo Akinola, Wei Yang, Balakumar Sundaralingam, Rowland O'Flaherty, Dieter Fox, Xiaolong Wang, Arsalan Mousavian, Yu-Wei Chao, and Yunzhu Li, "VT-Refine: Learning Bimanual Assembly with Visuo-Tactile Feedback via Simulation Fine-Tuning," CoRL 2025.
Best Paper Award, IROS 2025 AHFR Workshop.

Real World

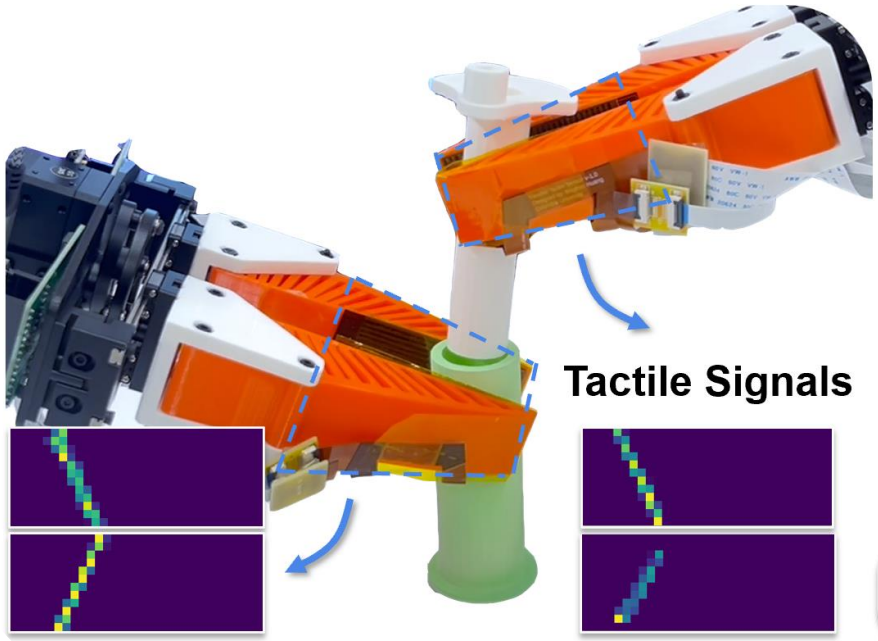

Human Demos

Vision

Touch

Tactile Signals

Pre-Train

Diffusion Policy

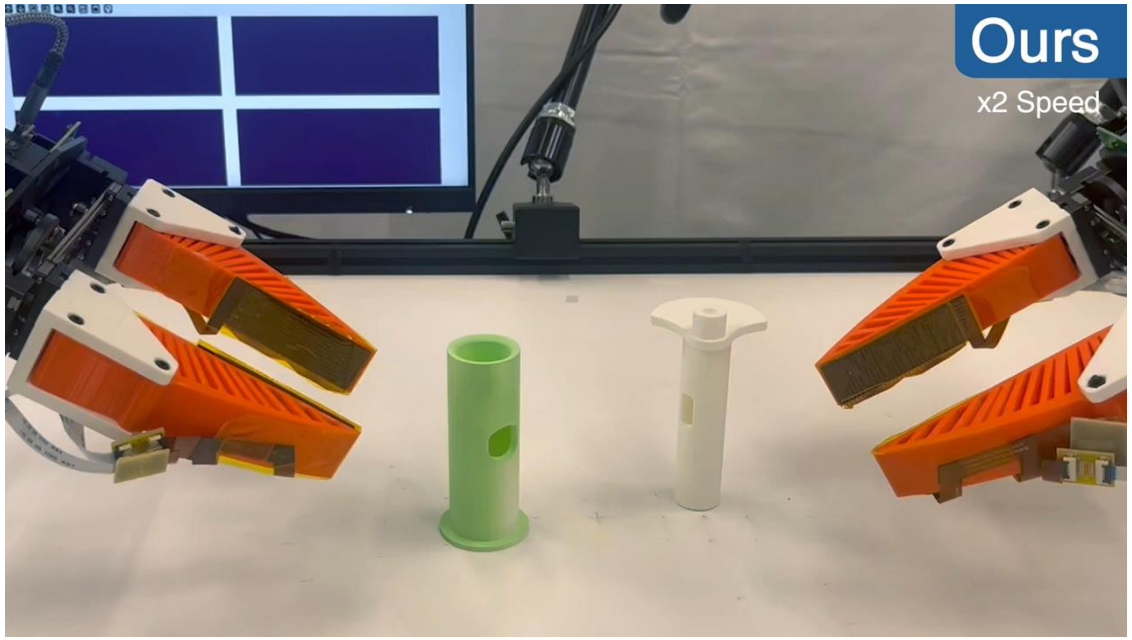

Simulation

GPU-Parallelized Tactile Simulation

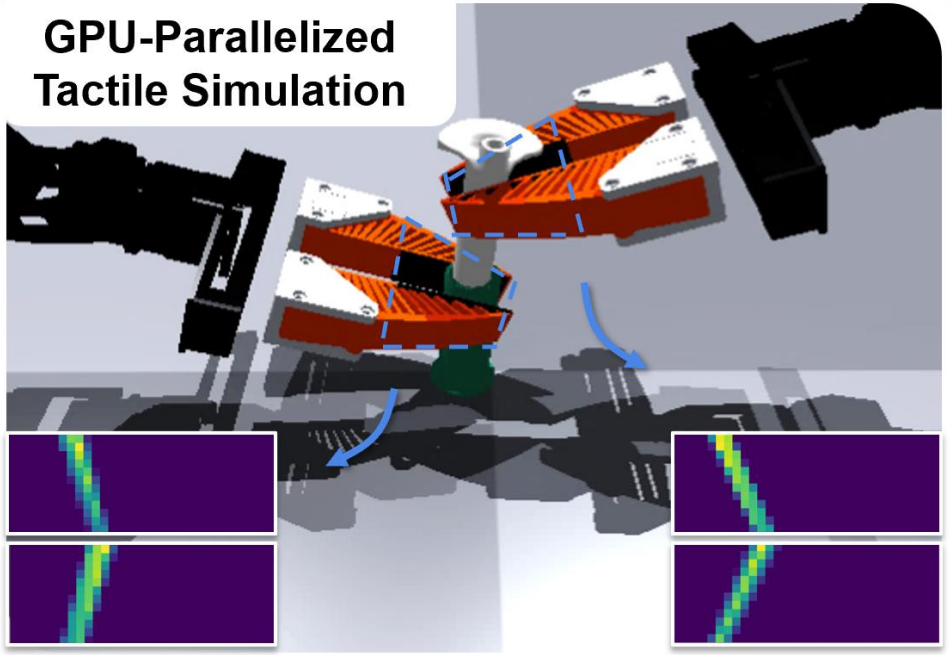

Large-Scale RL Fine-Tuning



Real World



30 Demos in Real

Collected within 30 mins

Simulation

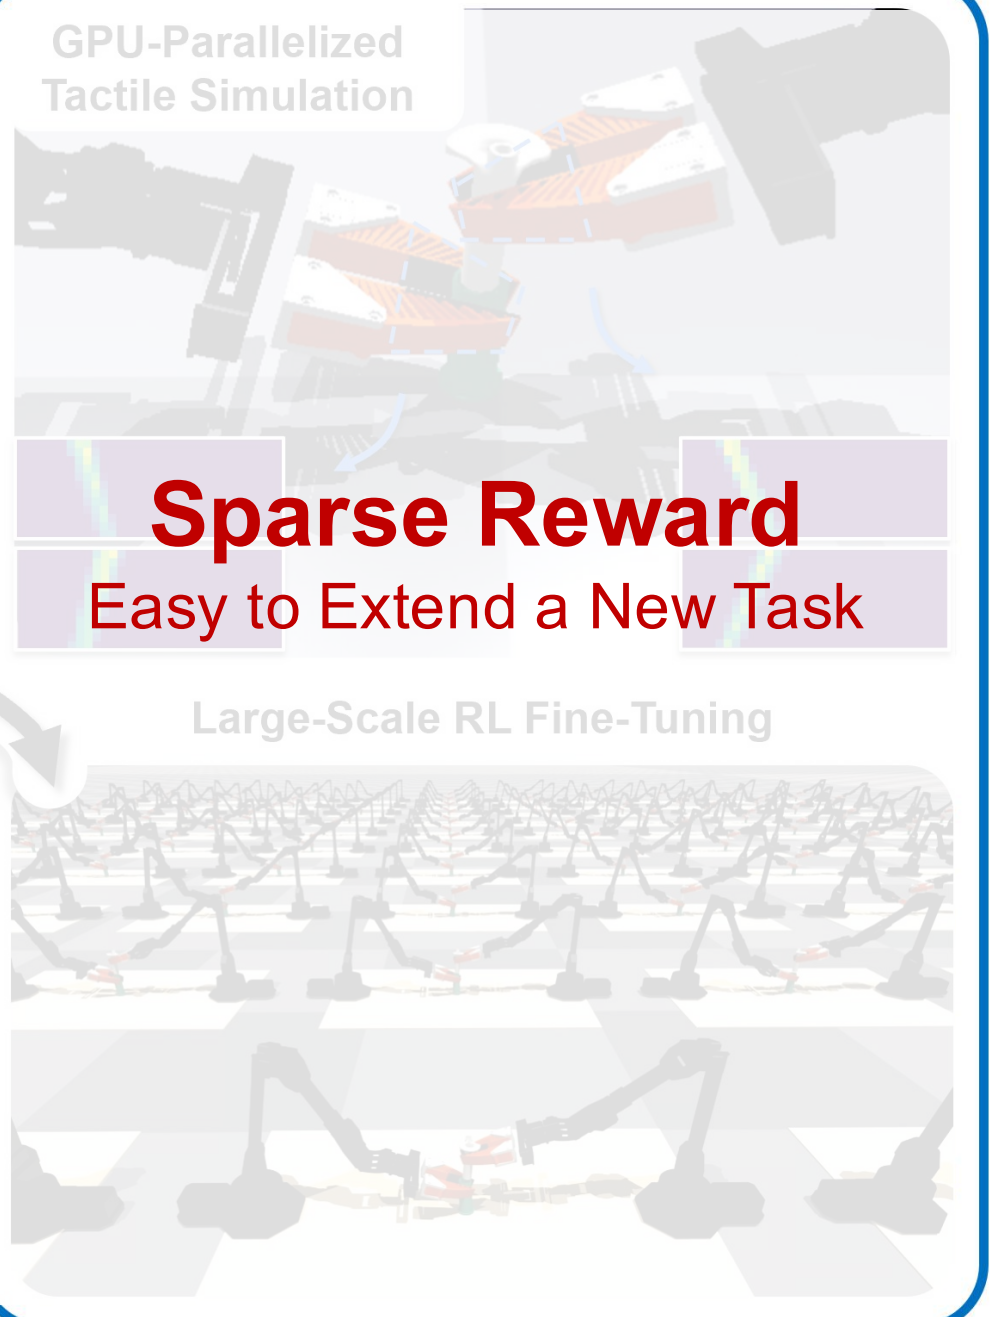

Sparse Reward

Easy to Extend a New Task

Real → Sim → Real: Scalable RL Training

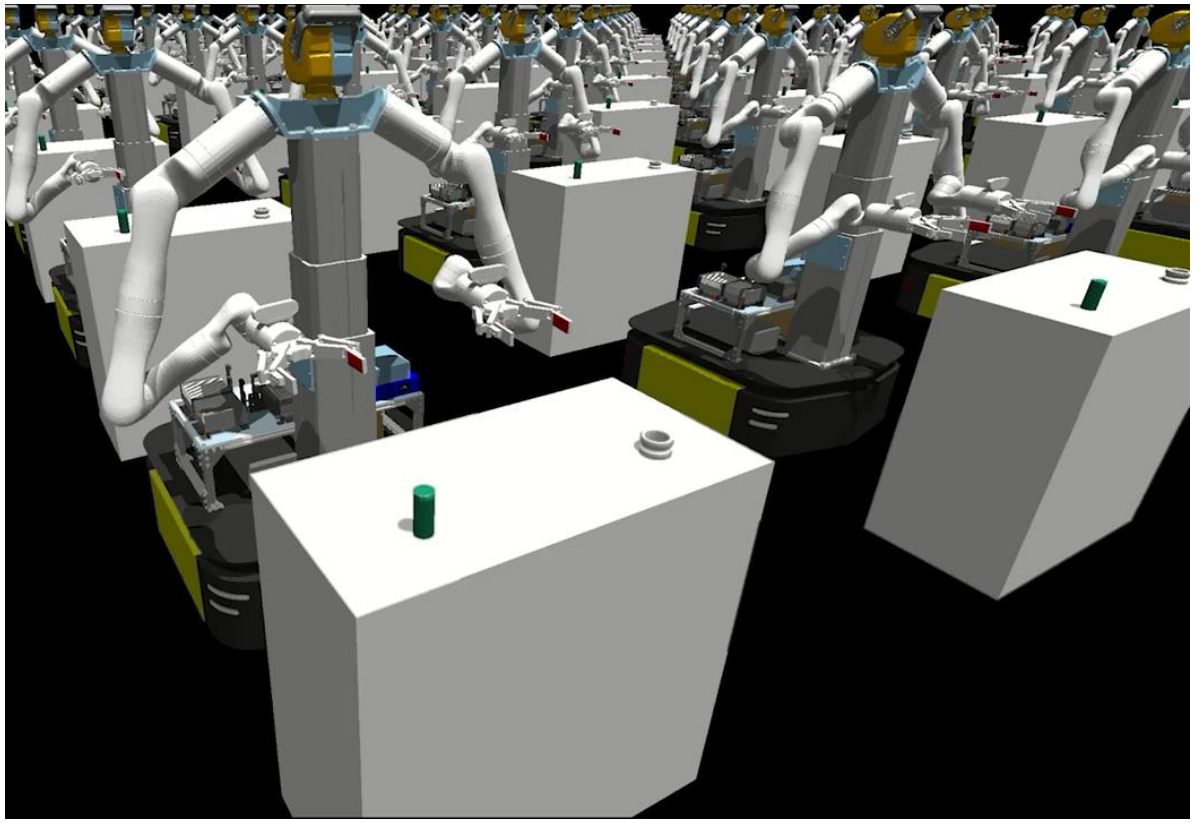

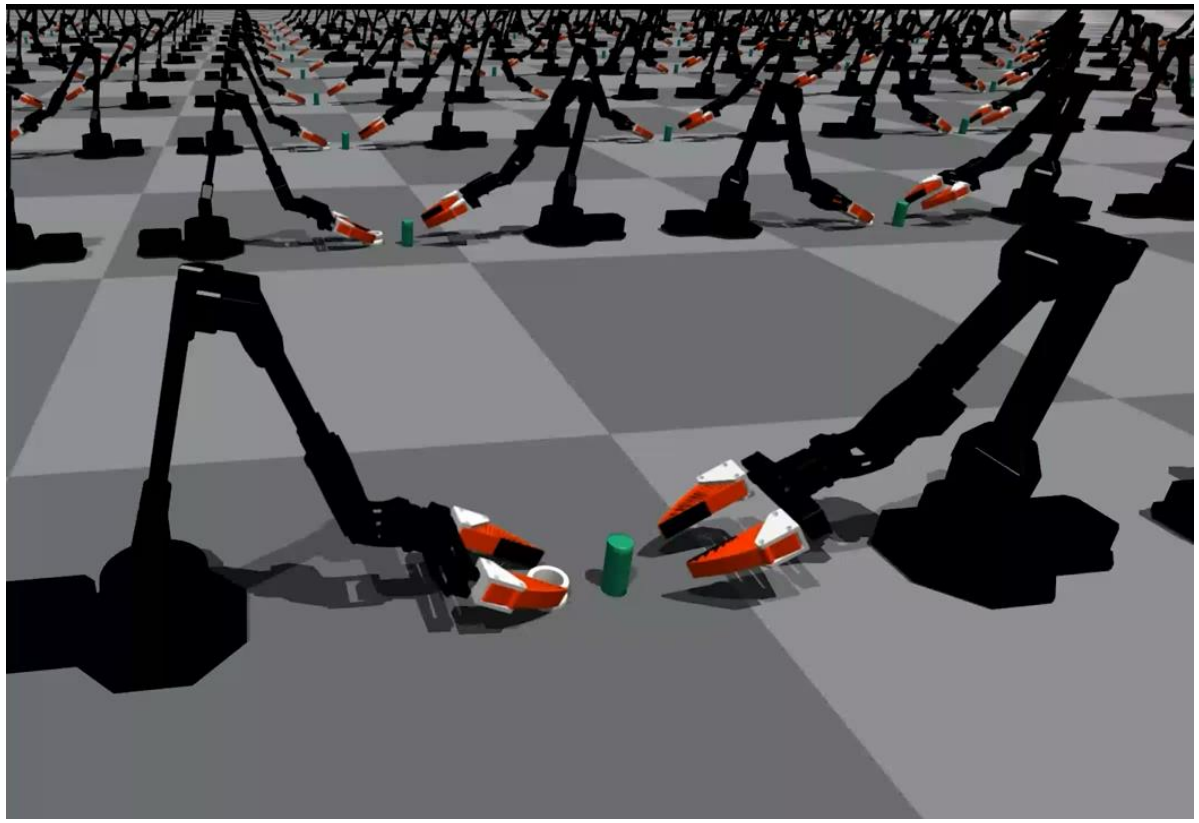

Binghao Huang, Jie Xu, Iretiayo Akinola, Wei Yang, Balakumar Sundaralingam, Rowland O'Flaherty, Dieter Fox, Xiaolong Wang, Arsalan Mousavian, Yu-Wei Chao, and Yunzhu Li, "VT-Refine: Learning Bimanual Assembly with Visuo-Tactile Feedback via Simulation Fine-Tuning," CoRL 2025.
Best Paper Award, IROS 2025 AHFR Workshop.













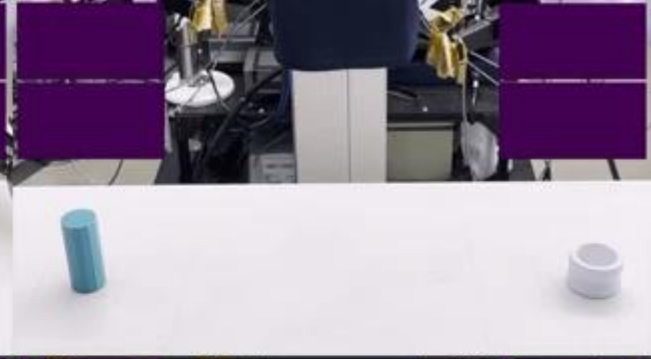

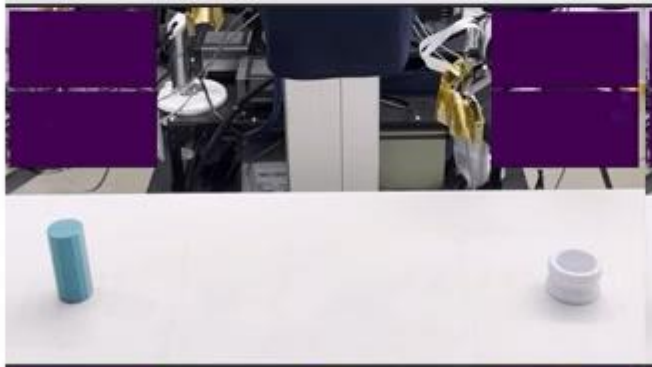

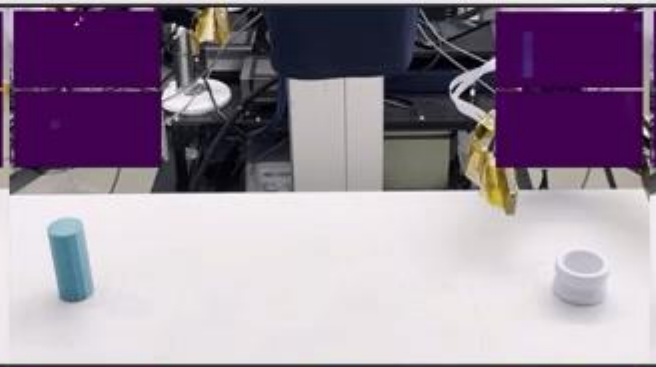

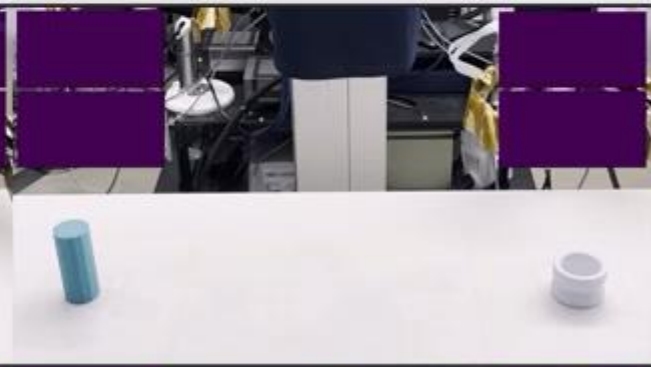







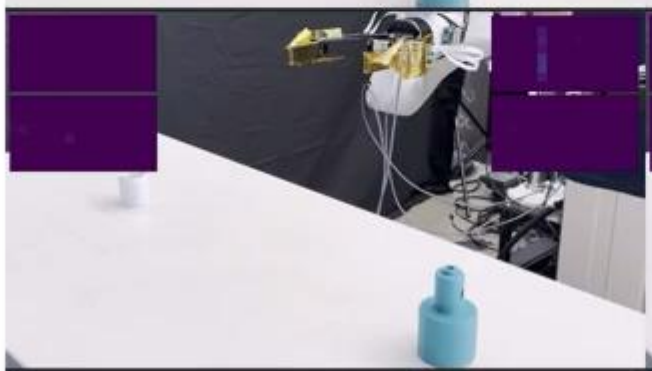

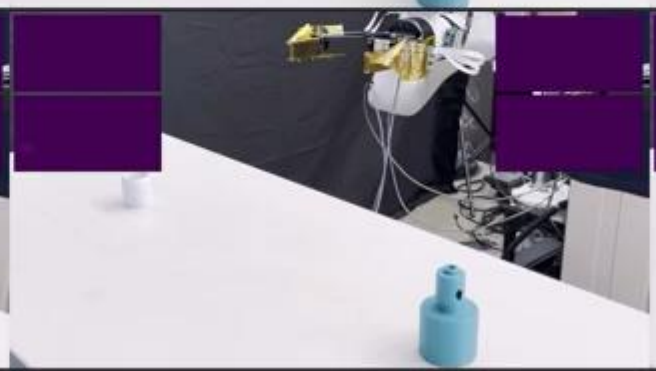



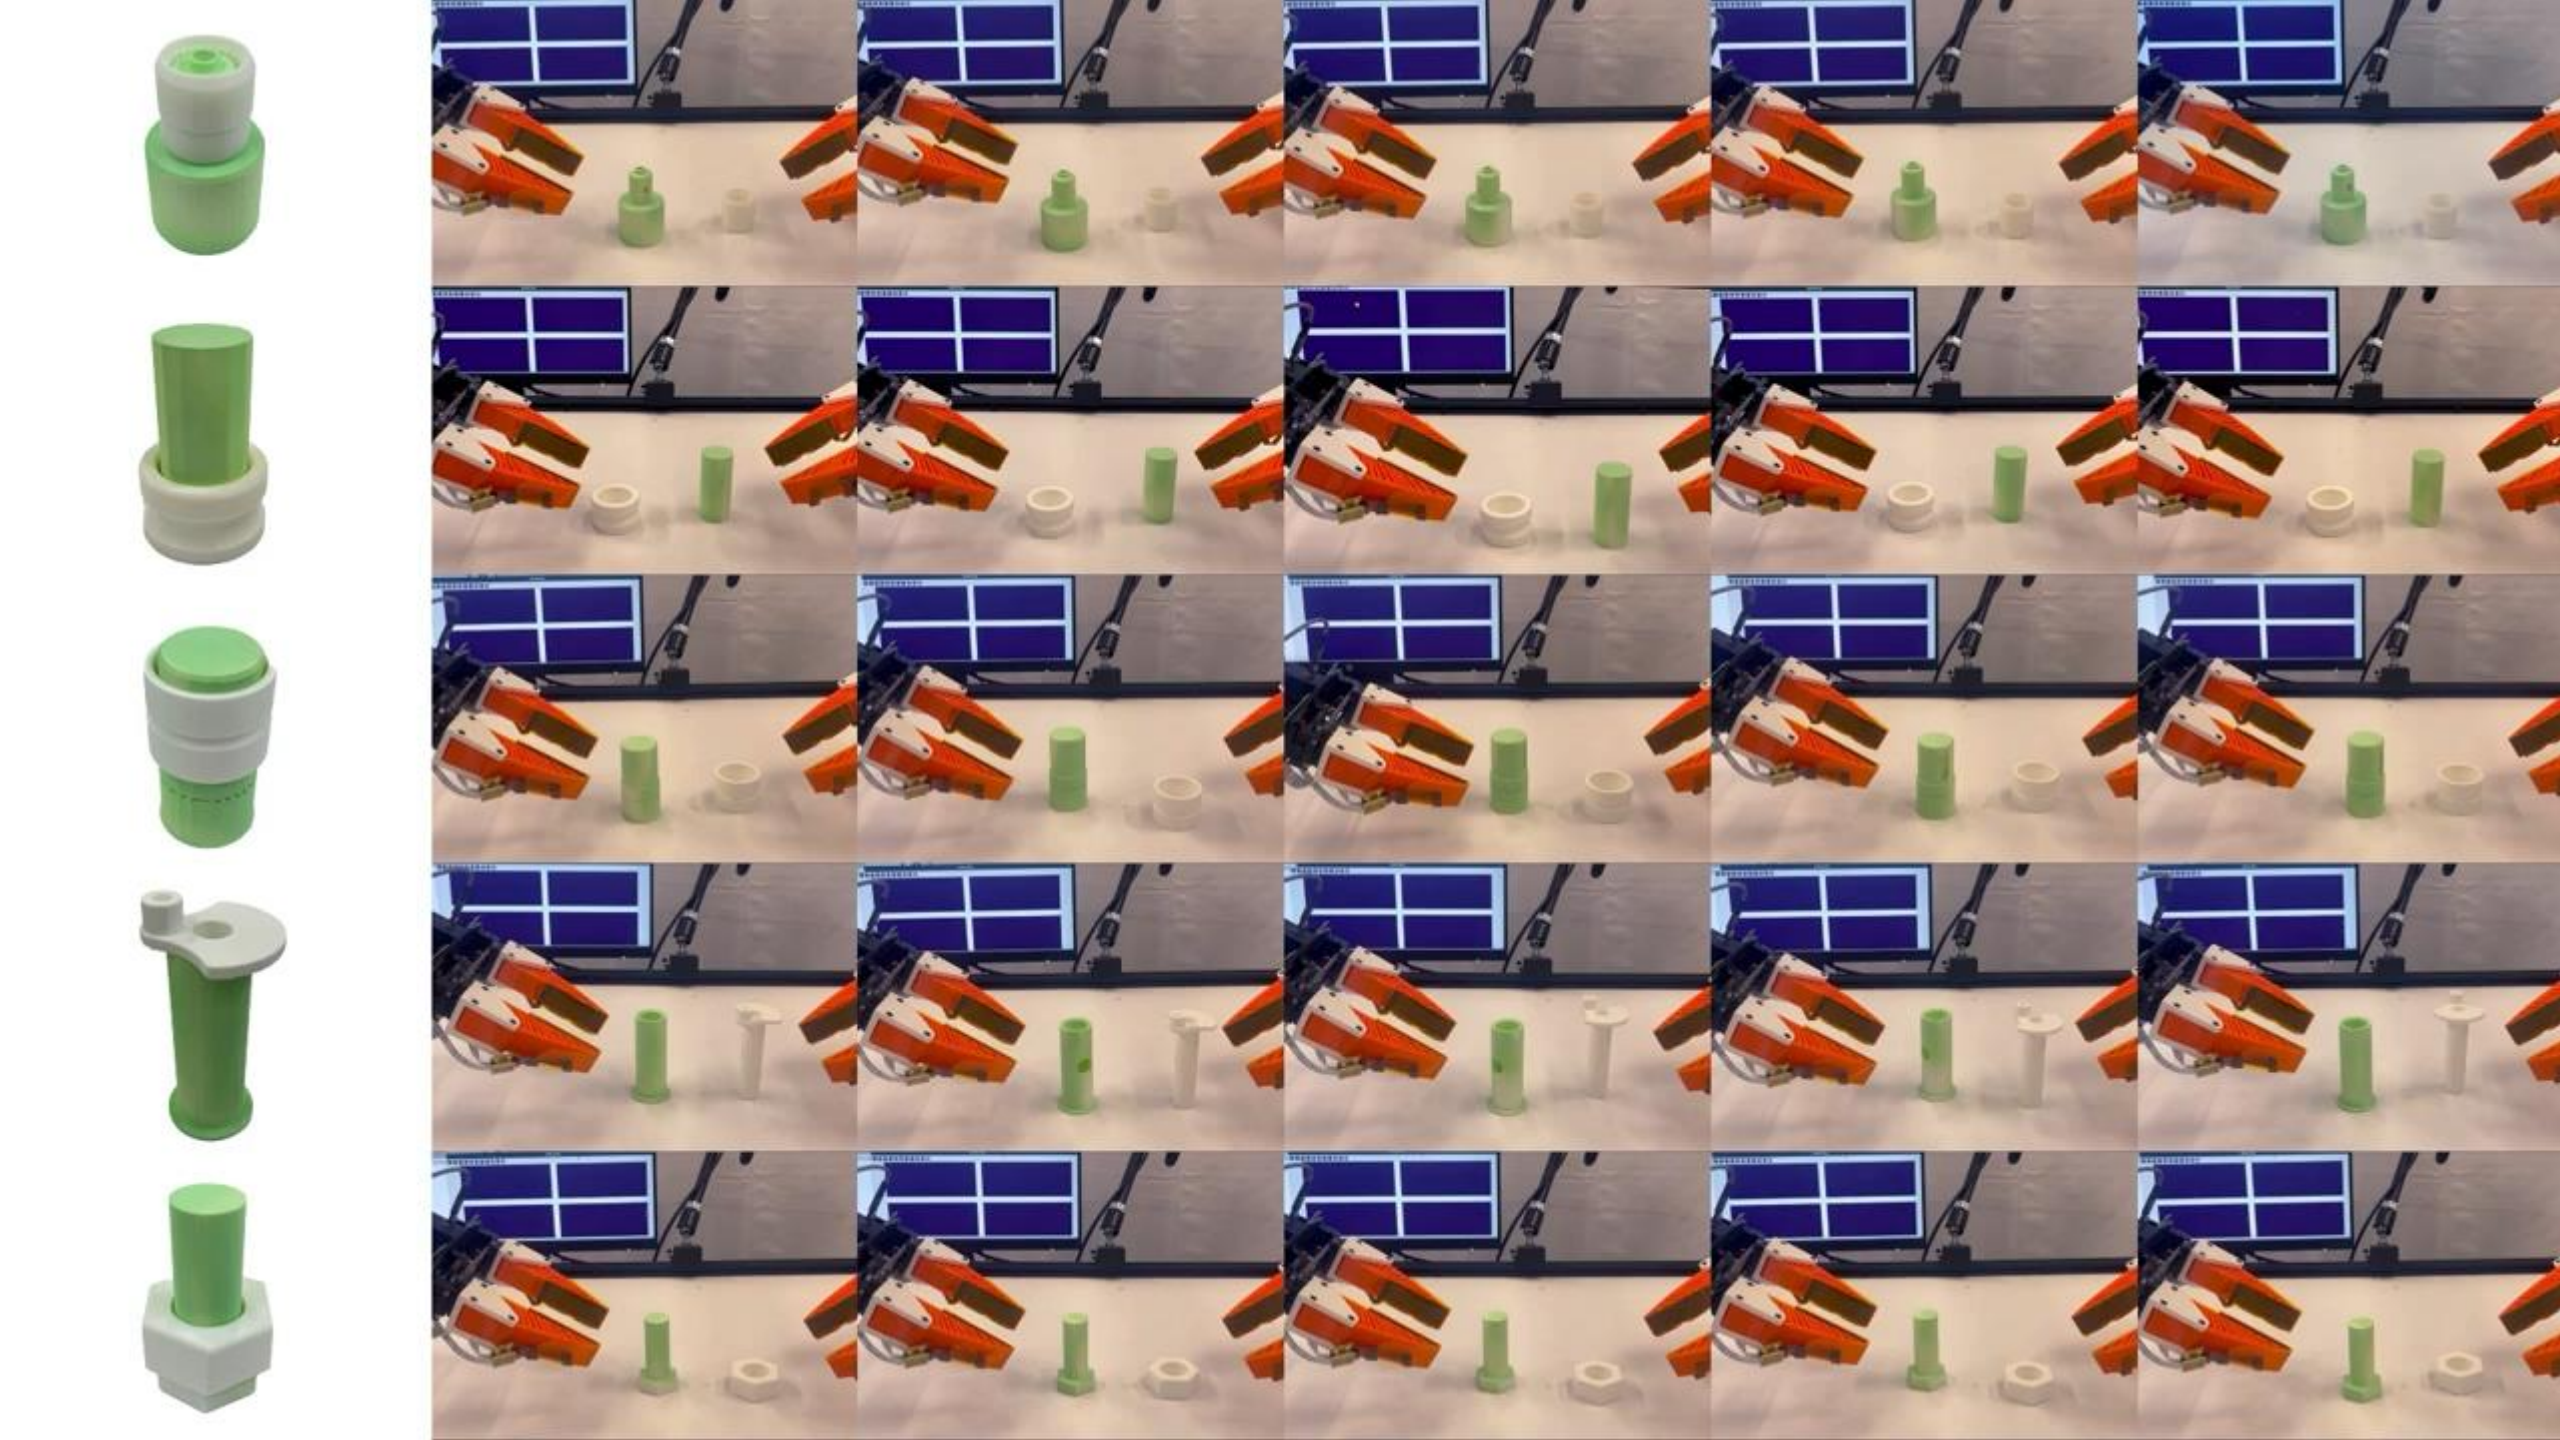

Next: 3D Multi-Modal World Model

