可泛化具身智能中的数据难题

许华哲

交叉信息研究院
Institute for Interdisciplinary Information Sciences

Mobile Manipulation

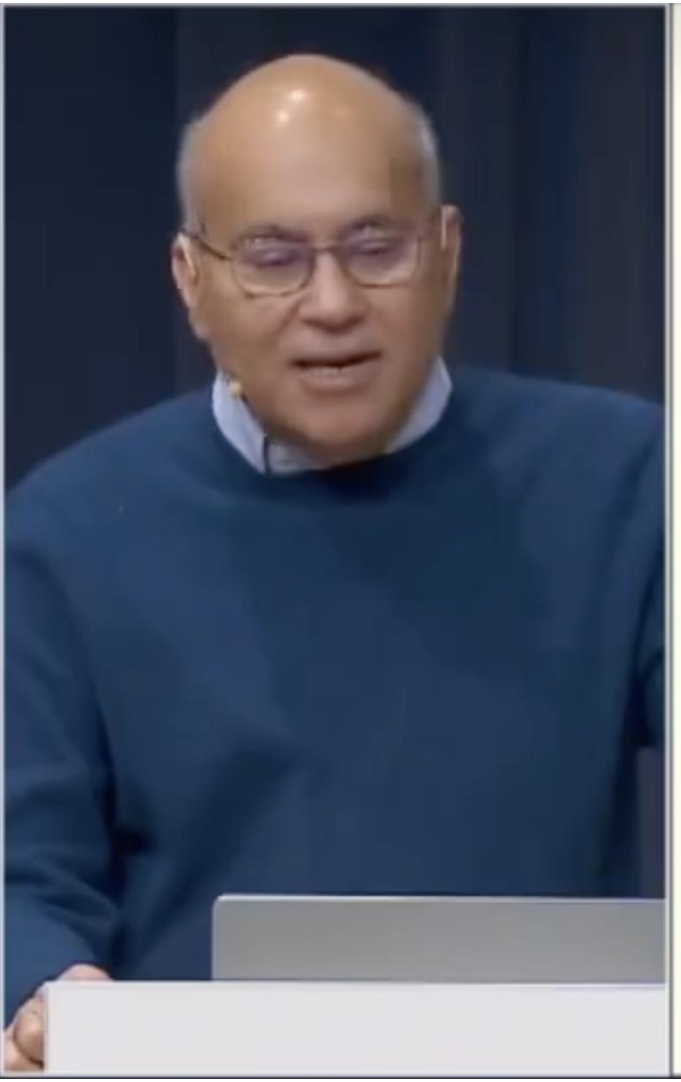

So how close to “solved” are we?

- Locomotion – *good progress* — Infinite simulation data!
- Navigation – *nearly done* — Autonomous cars + sim data!
- Manipulation – *long way to go* — Lack of modality! Sim and video? Inaccurate! Real? Expensive !

The big challenge is data. Much harder to obtain than for language and vision

LLMs are at least 10,000 times more data-hungry than human children

Tasks we need to do!

- More modalities, especially touch
- Better use of data


GL

DOGlove: Dexterous Manipulation with a Low-Cost Open-Source Haptic Force Feedback Glove

Han Zhang[1], Songbo Hu[1], Zhecheng Yuan[1,2,3], Huazhe Xu[1,2,3]

[1] Tsinghua University, [2] Shanghai Qi Zhi Institute, [3] Shanghai AI Lab

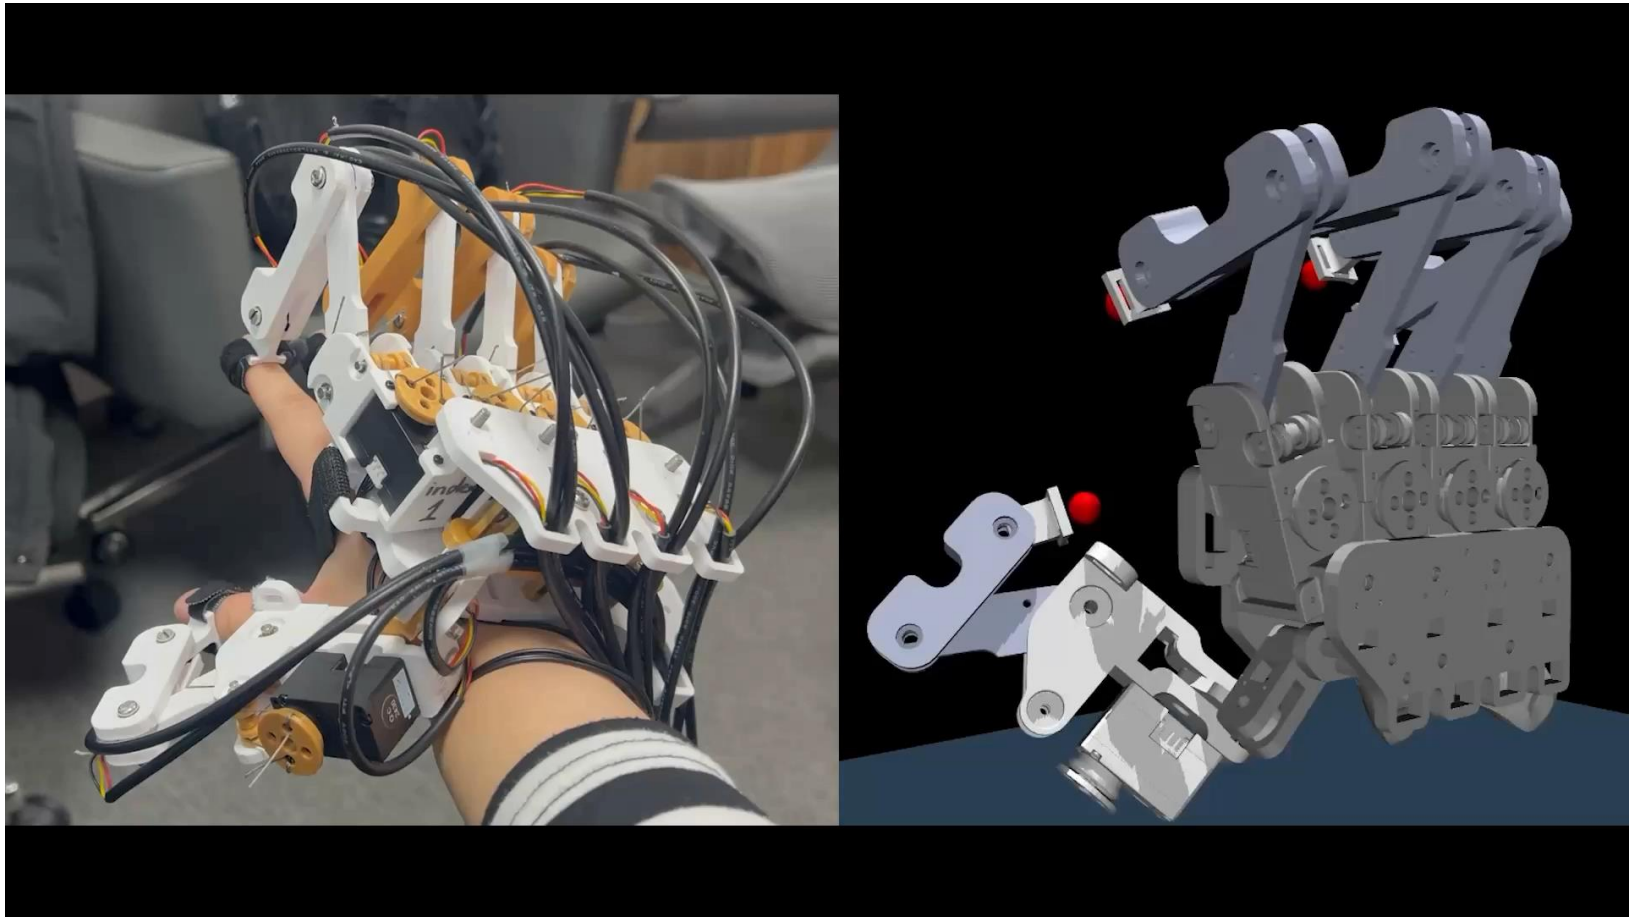

LEAP Hand

Shadow Hand

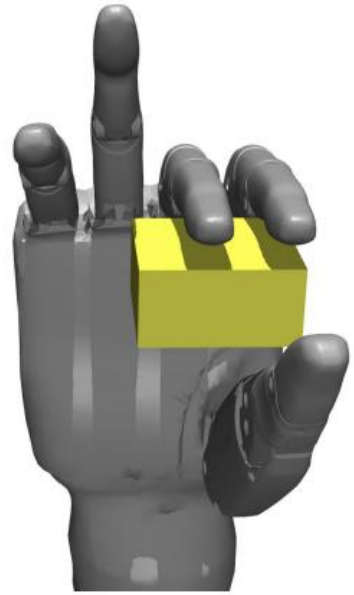

Inspire Hand

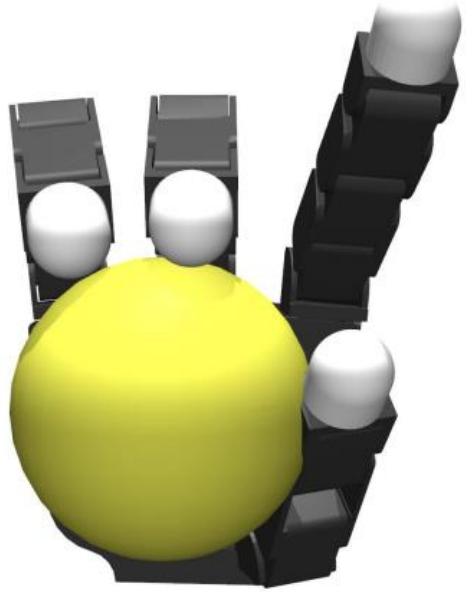

Allegro Hand

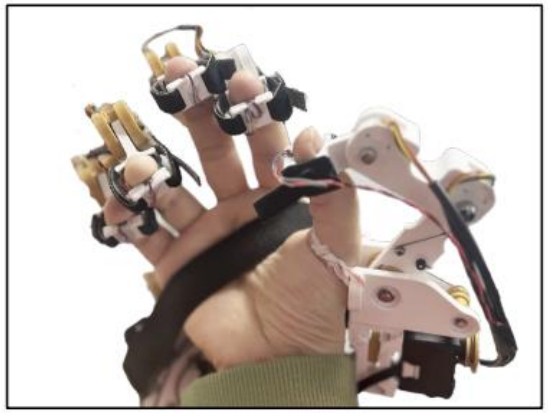

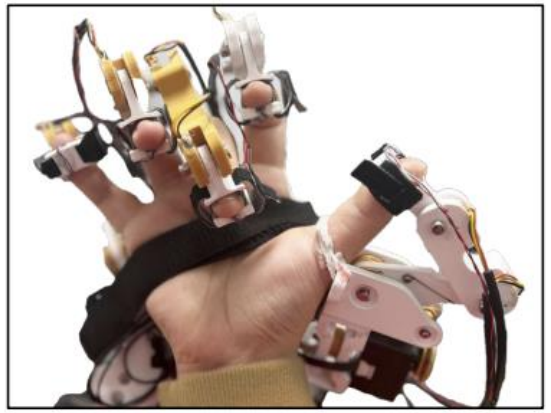

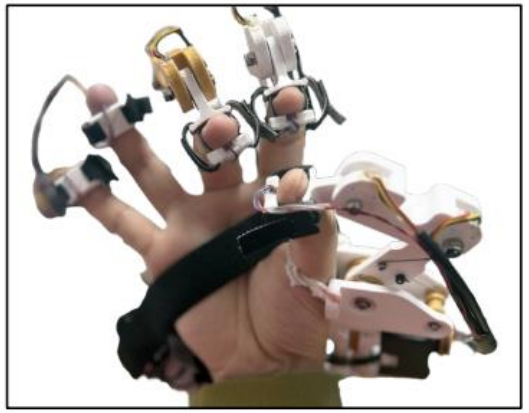

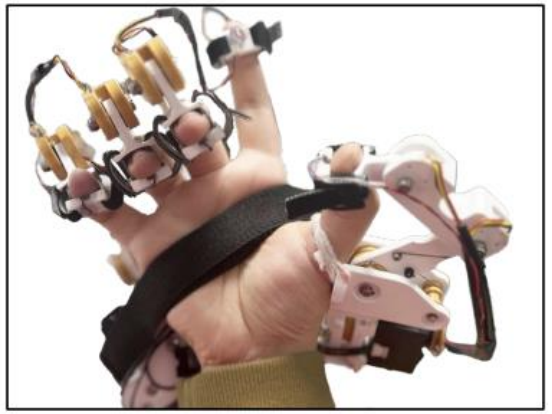


5x
Teleoperate the LEAP Hand to
control condensed milk flow


10x
Press and Move Box
(autonomous)

1. Teleoperation System: TactAR

Calibration & Functions

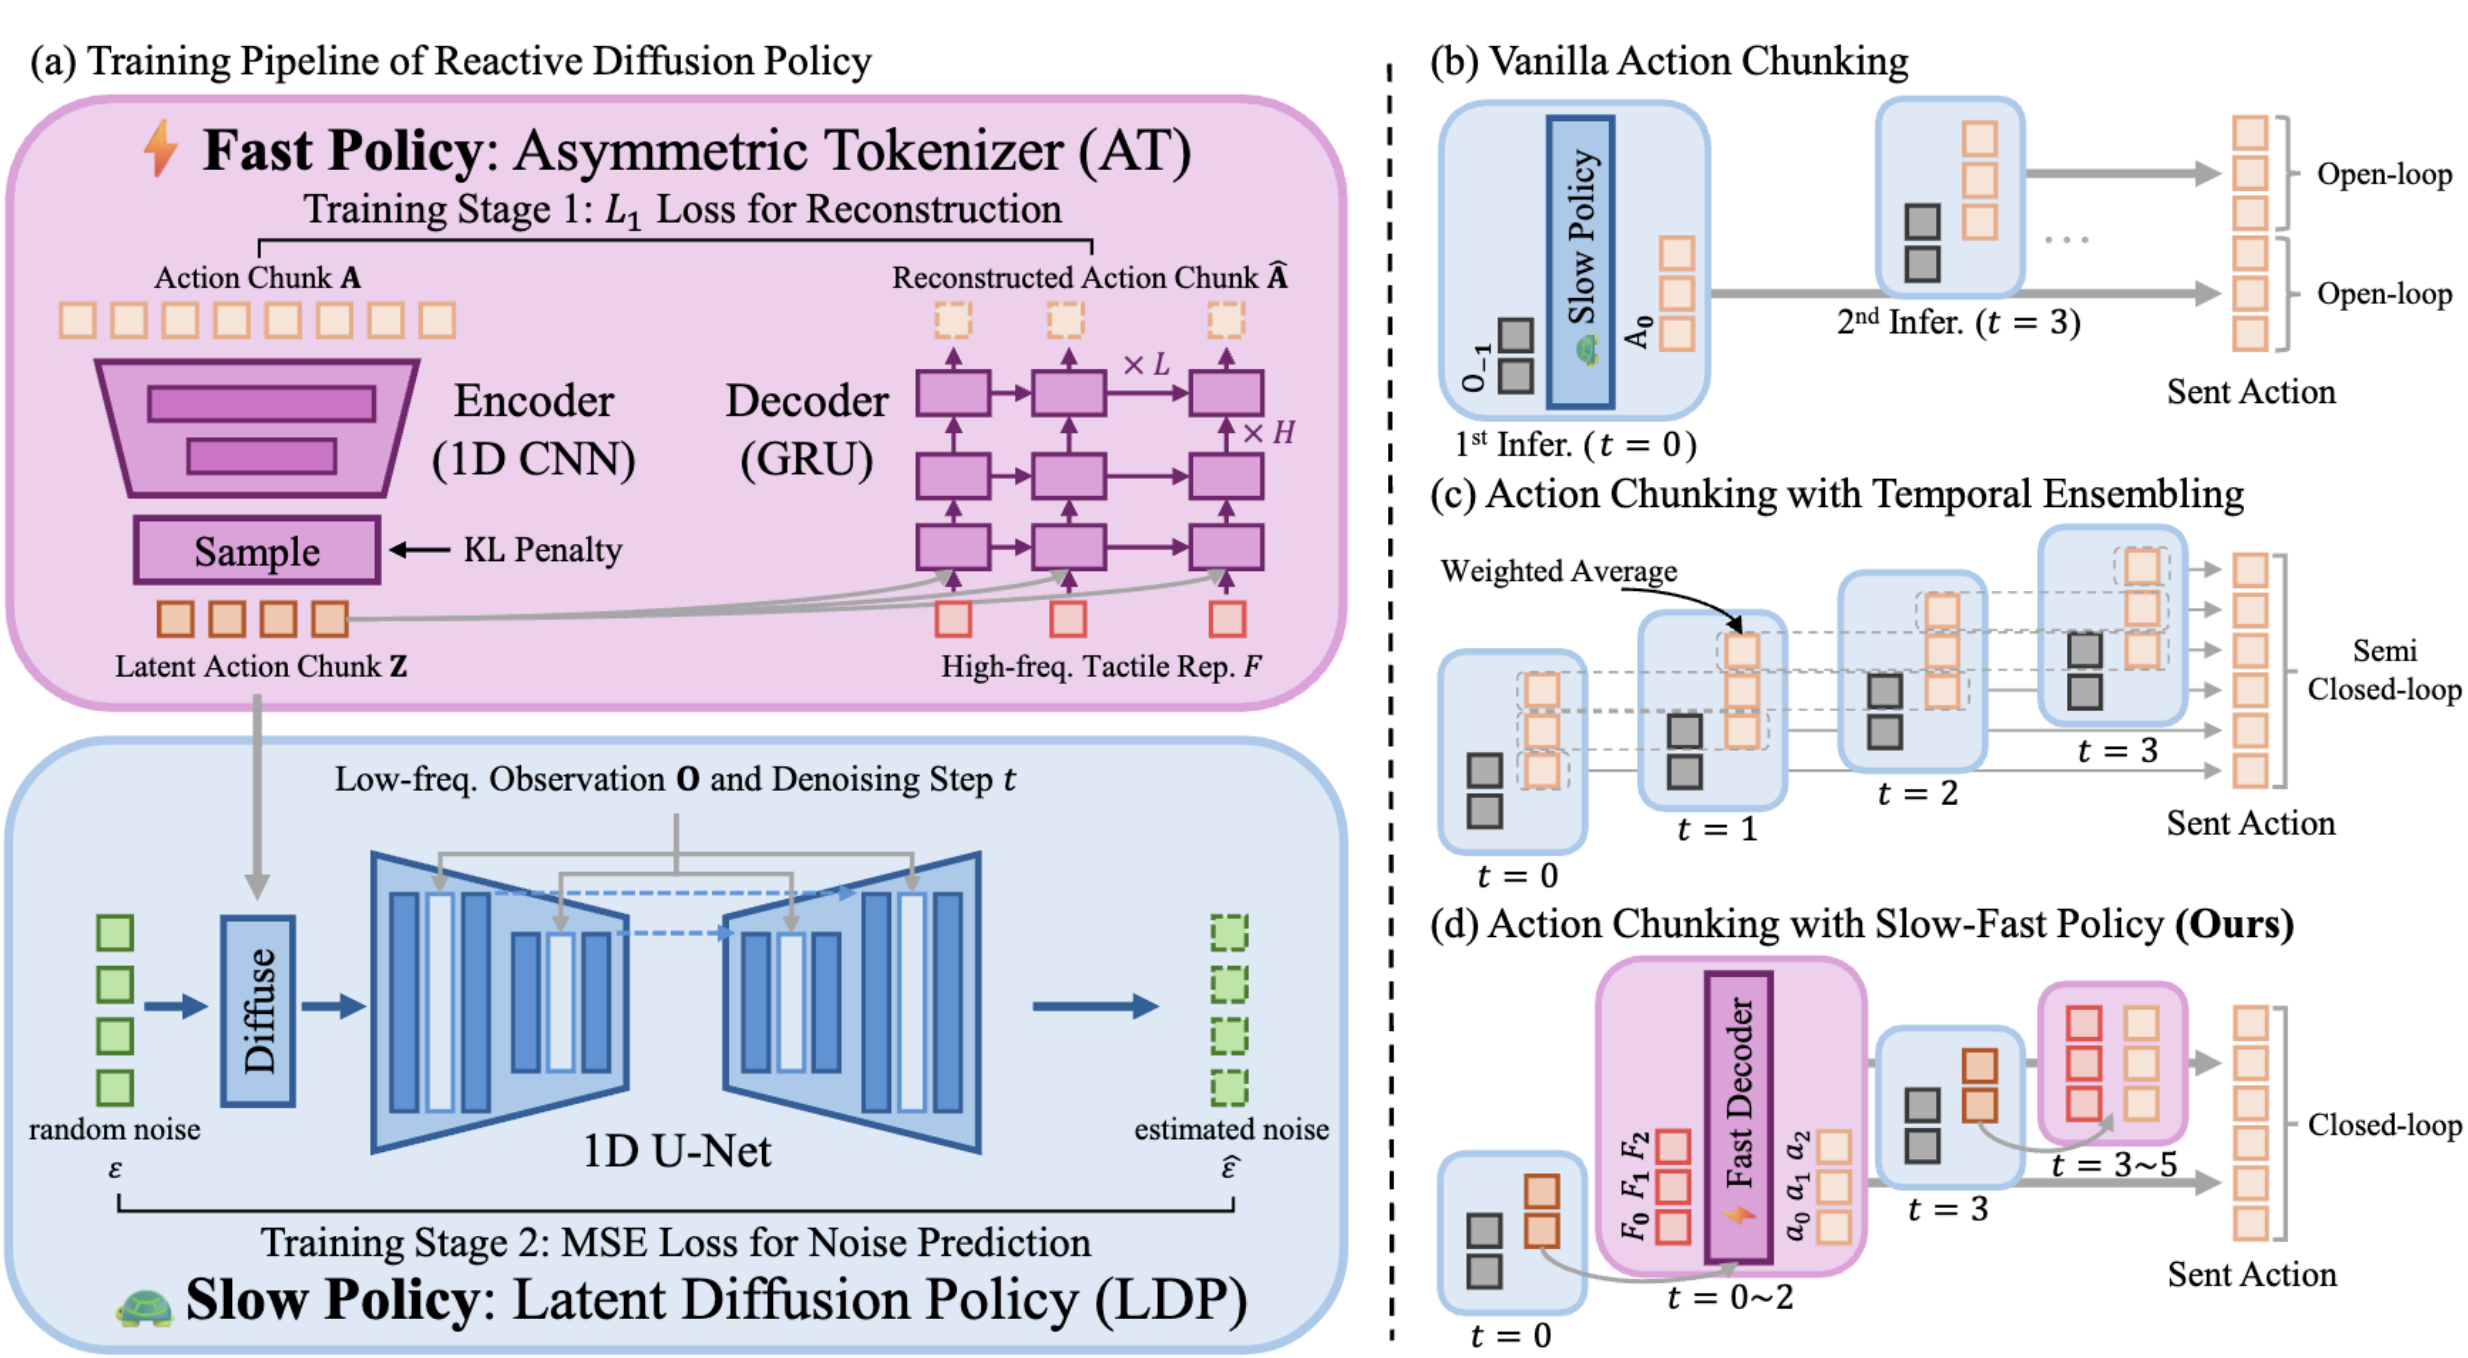
(a) Training Pipeline of Reactive Diffusion Policy
Fast Policy: Asymmetric Tokenizer (AT)
Training Stage 1: L_1 Loss for Reconstruction
Action Chunk A
Reconstructed Action Chunk $\hat{\mathbf{A}}$
Encoder (1D CNN)
Decoder (GRU)
× L
× H
Sample
KL Penalty
Latent Action Chunk Z
High-freq. Tactile Rep. F
Low-freq. Observation O and Denoising Step t
Diffuse
random noise
ε
1D U-Net
estimated noise
$\hat{\varepsilon}$
Training Stage 2: MSE Loss for Noise Prediction
Slow Policy: Latent Diffusion Policy (LDP)
(b) Vanilla Action Chunking
Slow Policy
O_{-1}
A_0
1st Infer. (t = 0)
2nd Infer. (t = 3)
Open-loop
Open-loop
Sent Action
(c) Action Chunking with Temporal Ensembling
Weighted Average
t = 0
t = 1
t = 2
t = 3
Semi Closed-loop
Sent Action
(d) Action Chunking with Slow-Fast Policy (Ours)
Fast Decoder
F_0 F_1 F_2
a_0 a_1 a_2
t = 0
t = 0~2
t = 3
t = 3~5
Closed-loop
Sent Action

2. Learning Algorithm: RDP

Demo with Baselines

How to better use the existing data?

- Generate them!

DemoGen: Synthetic Demonstration Generation for Data-Efficient Visuomotor Policy Learning

Zhengrong Xue*, Shuying Deng*, Zhenyang Chen,
Yixuan Wang, Zhecheng Yuan, Huazhe Xu


Human-Collected Demo

Protocols

Single-Arm Setups

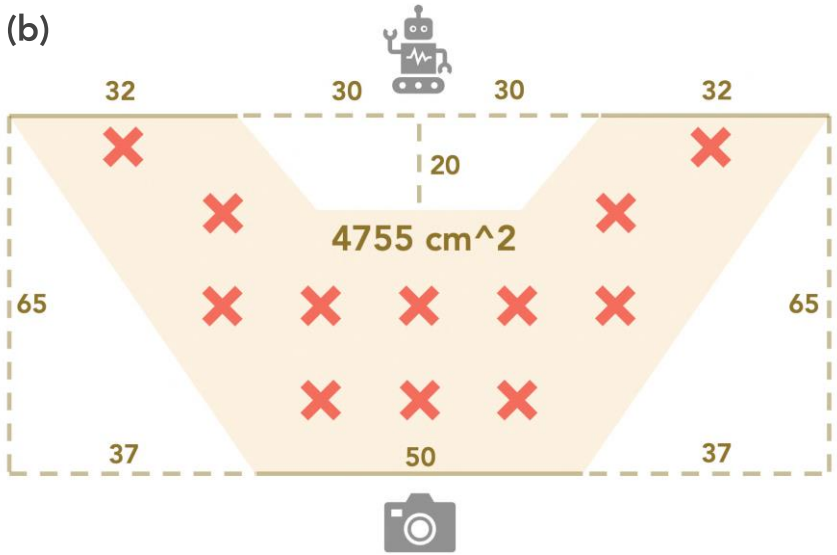

Full-size evaluation workspace

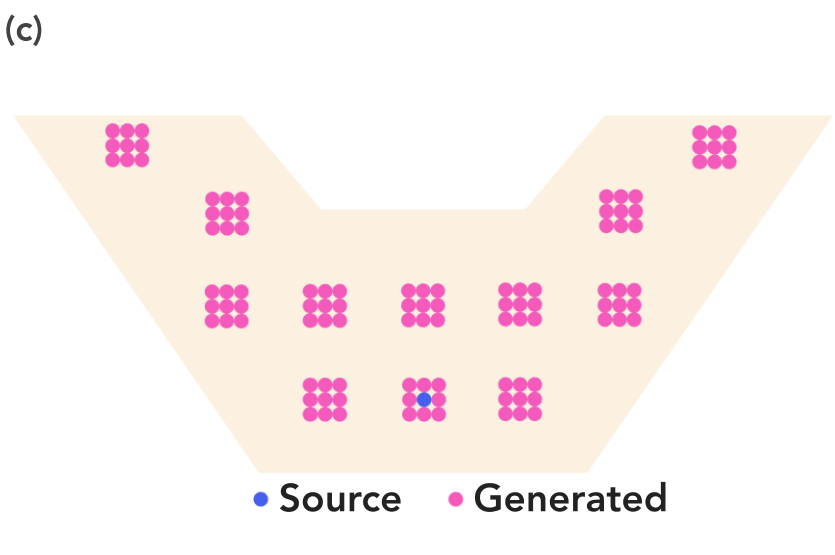

Generation strategy

Single-Arm Tasks

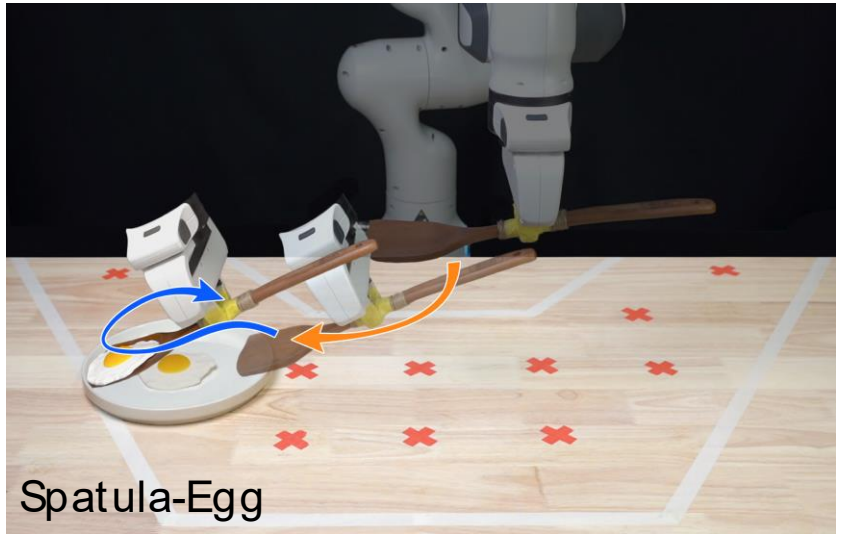

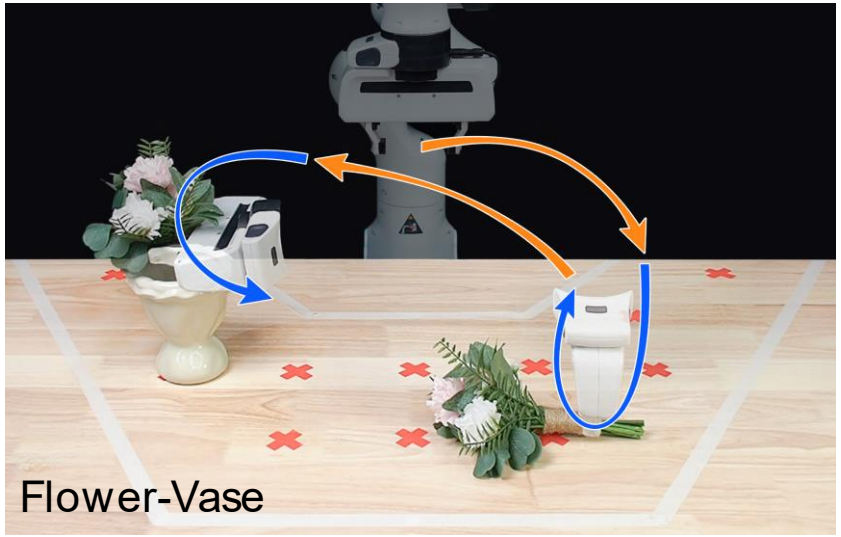

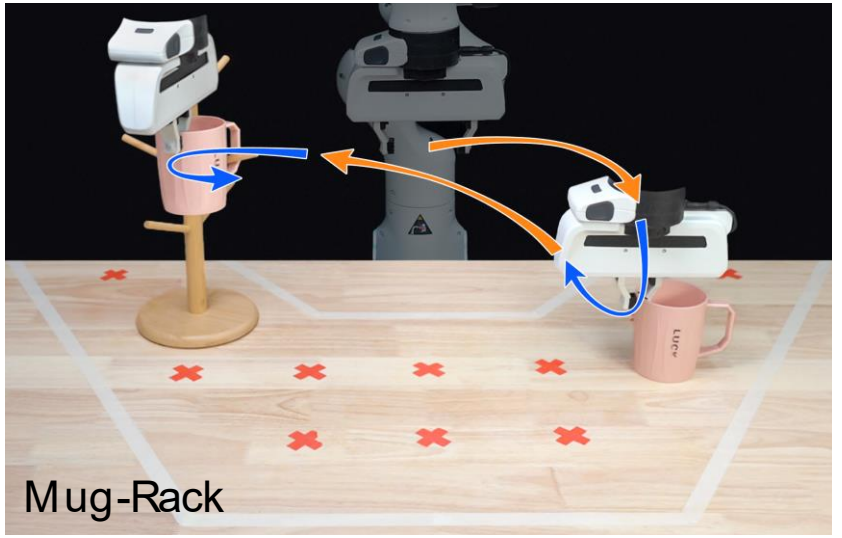

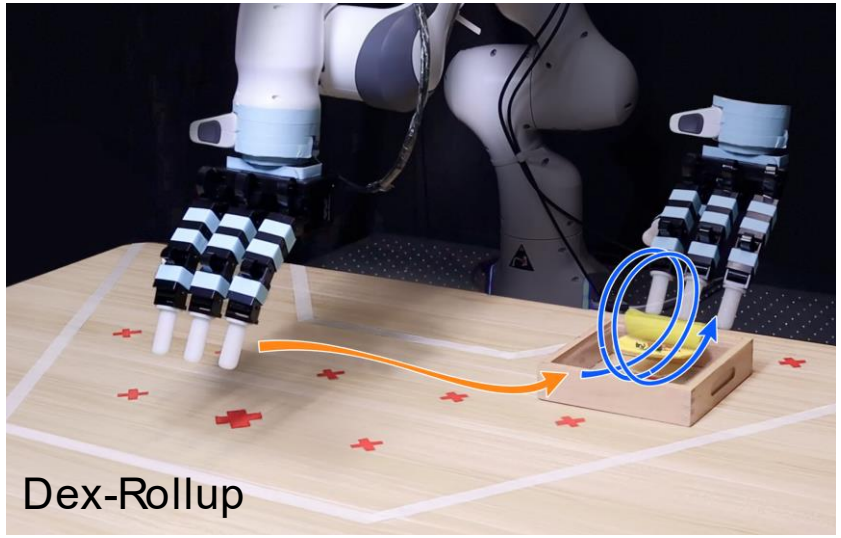

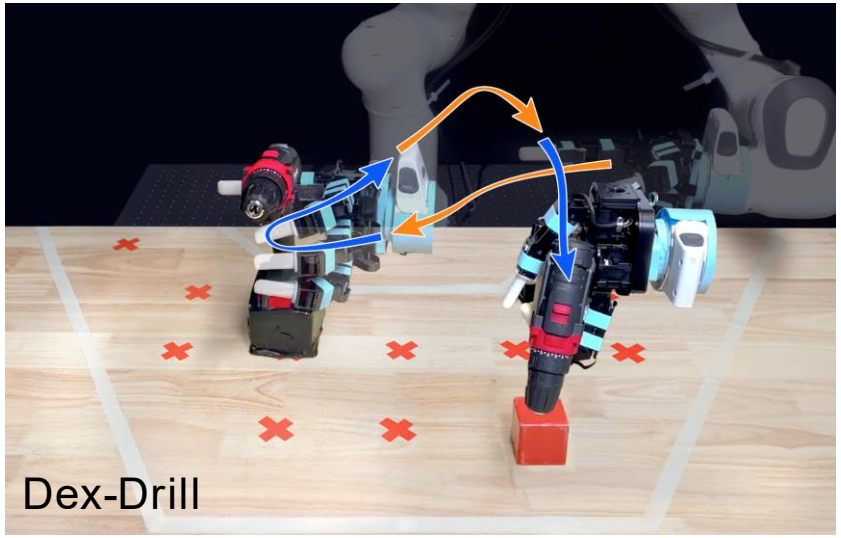

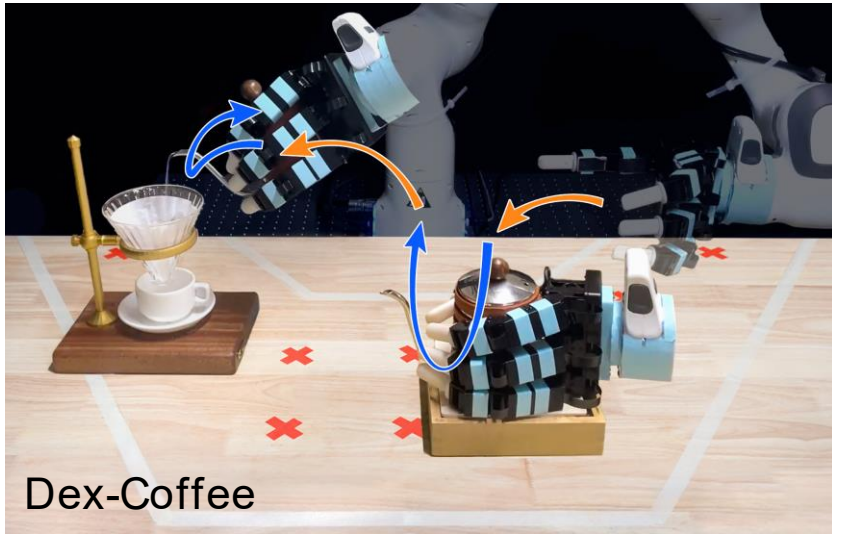

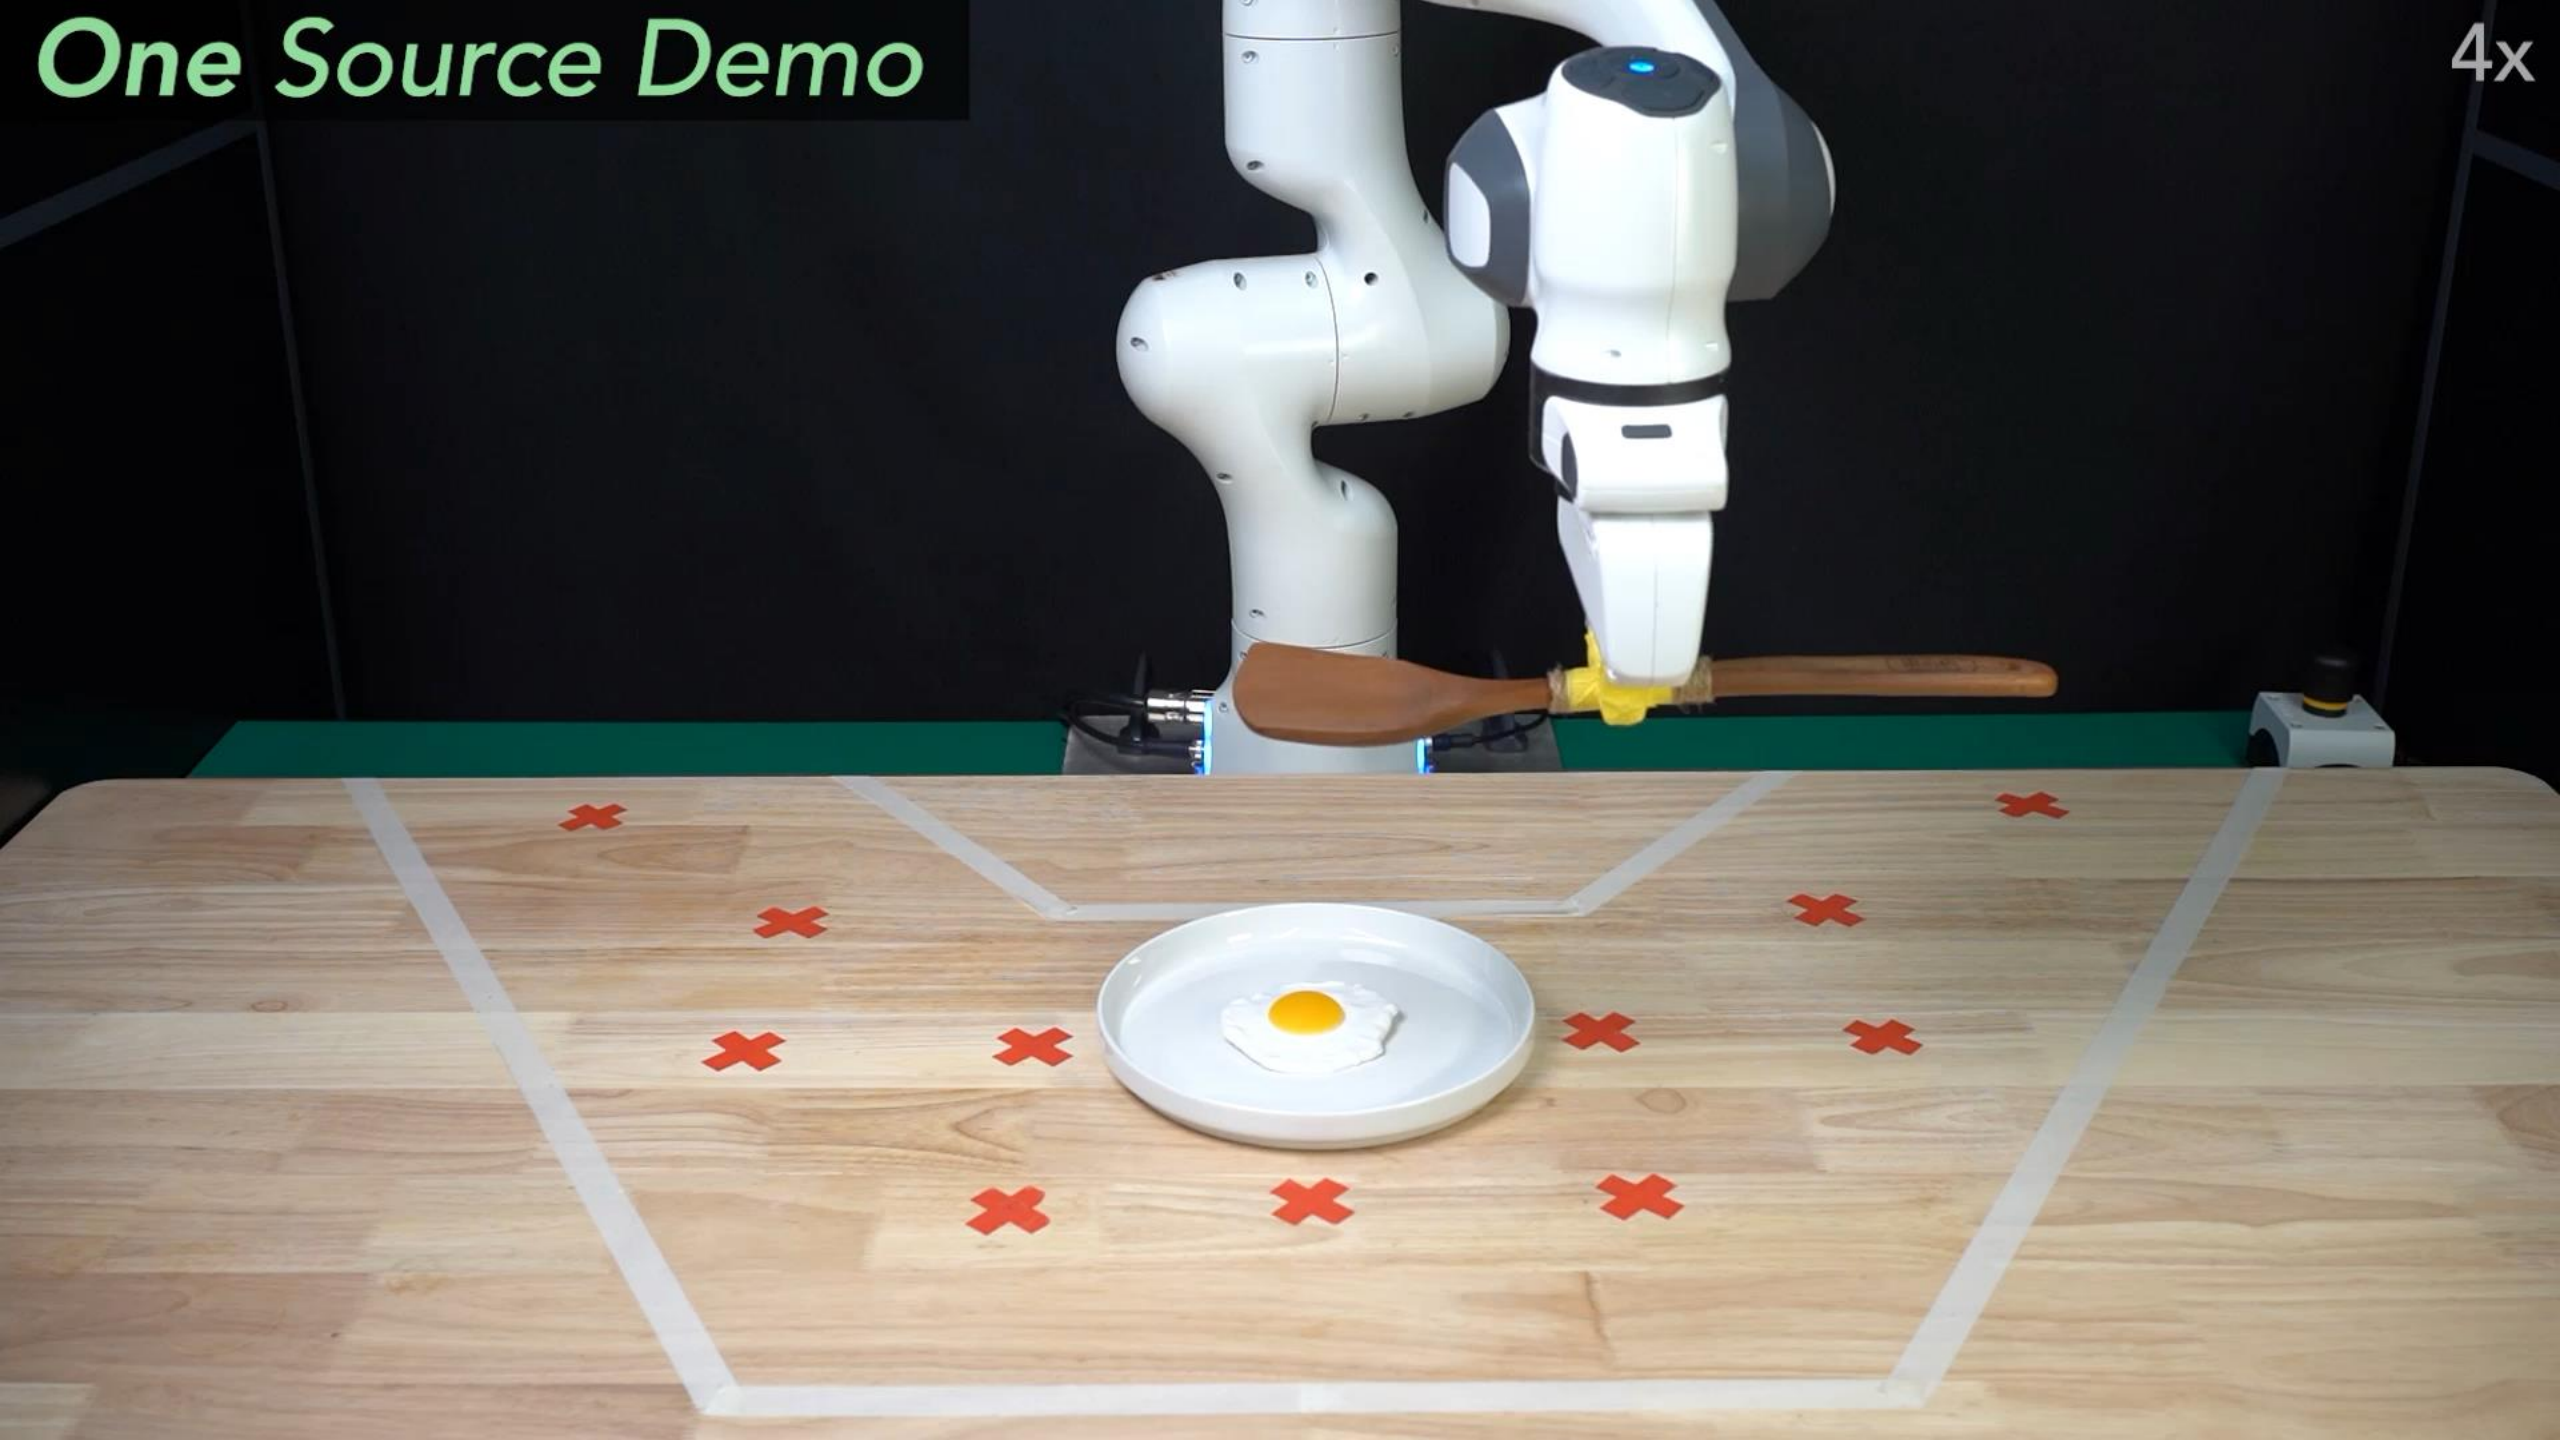
One Source Demo
4x

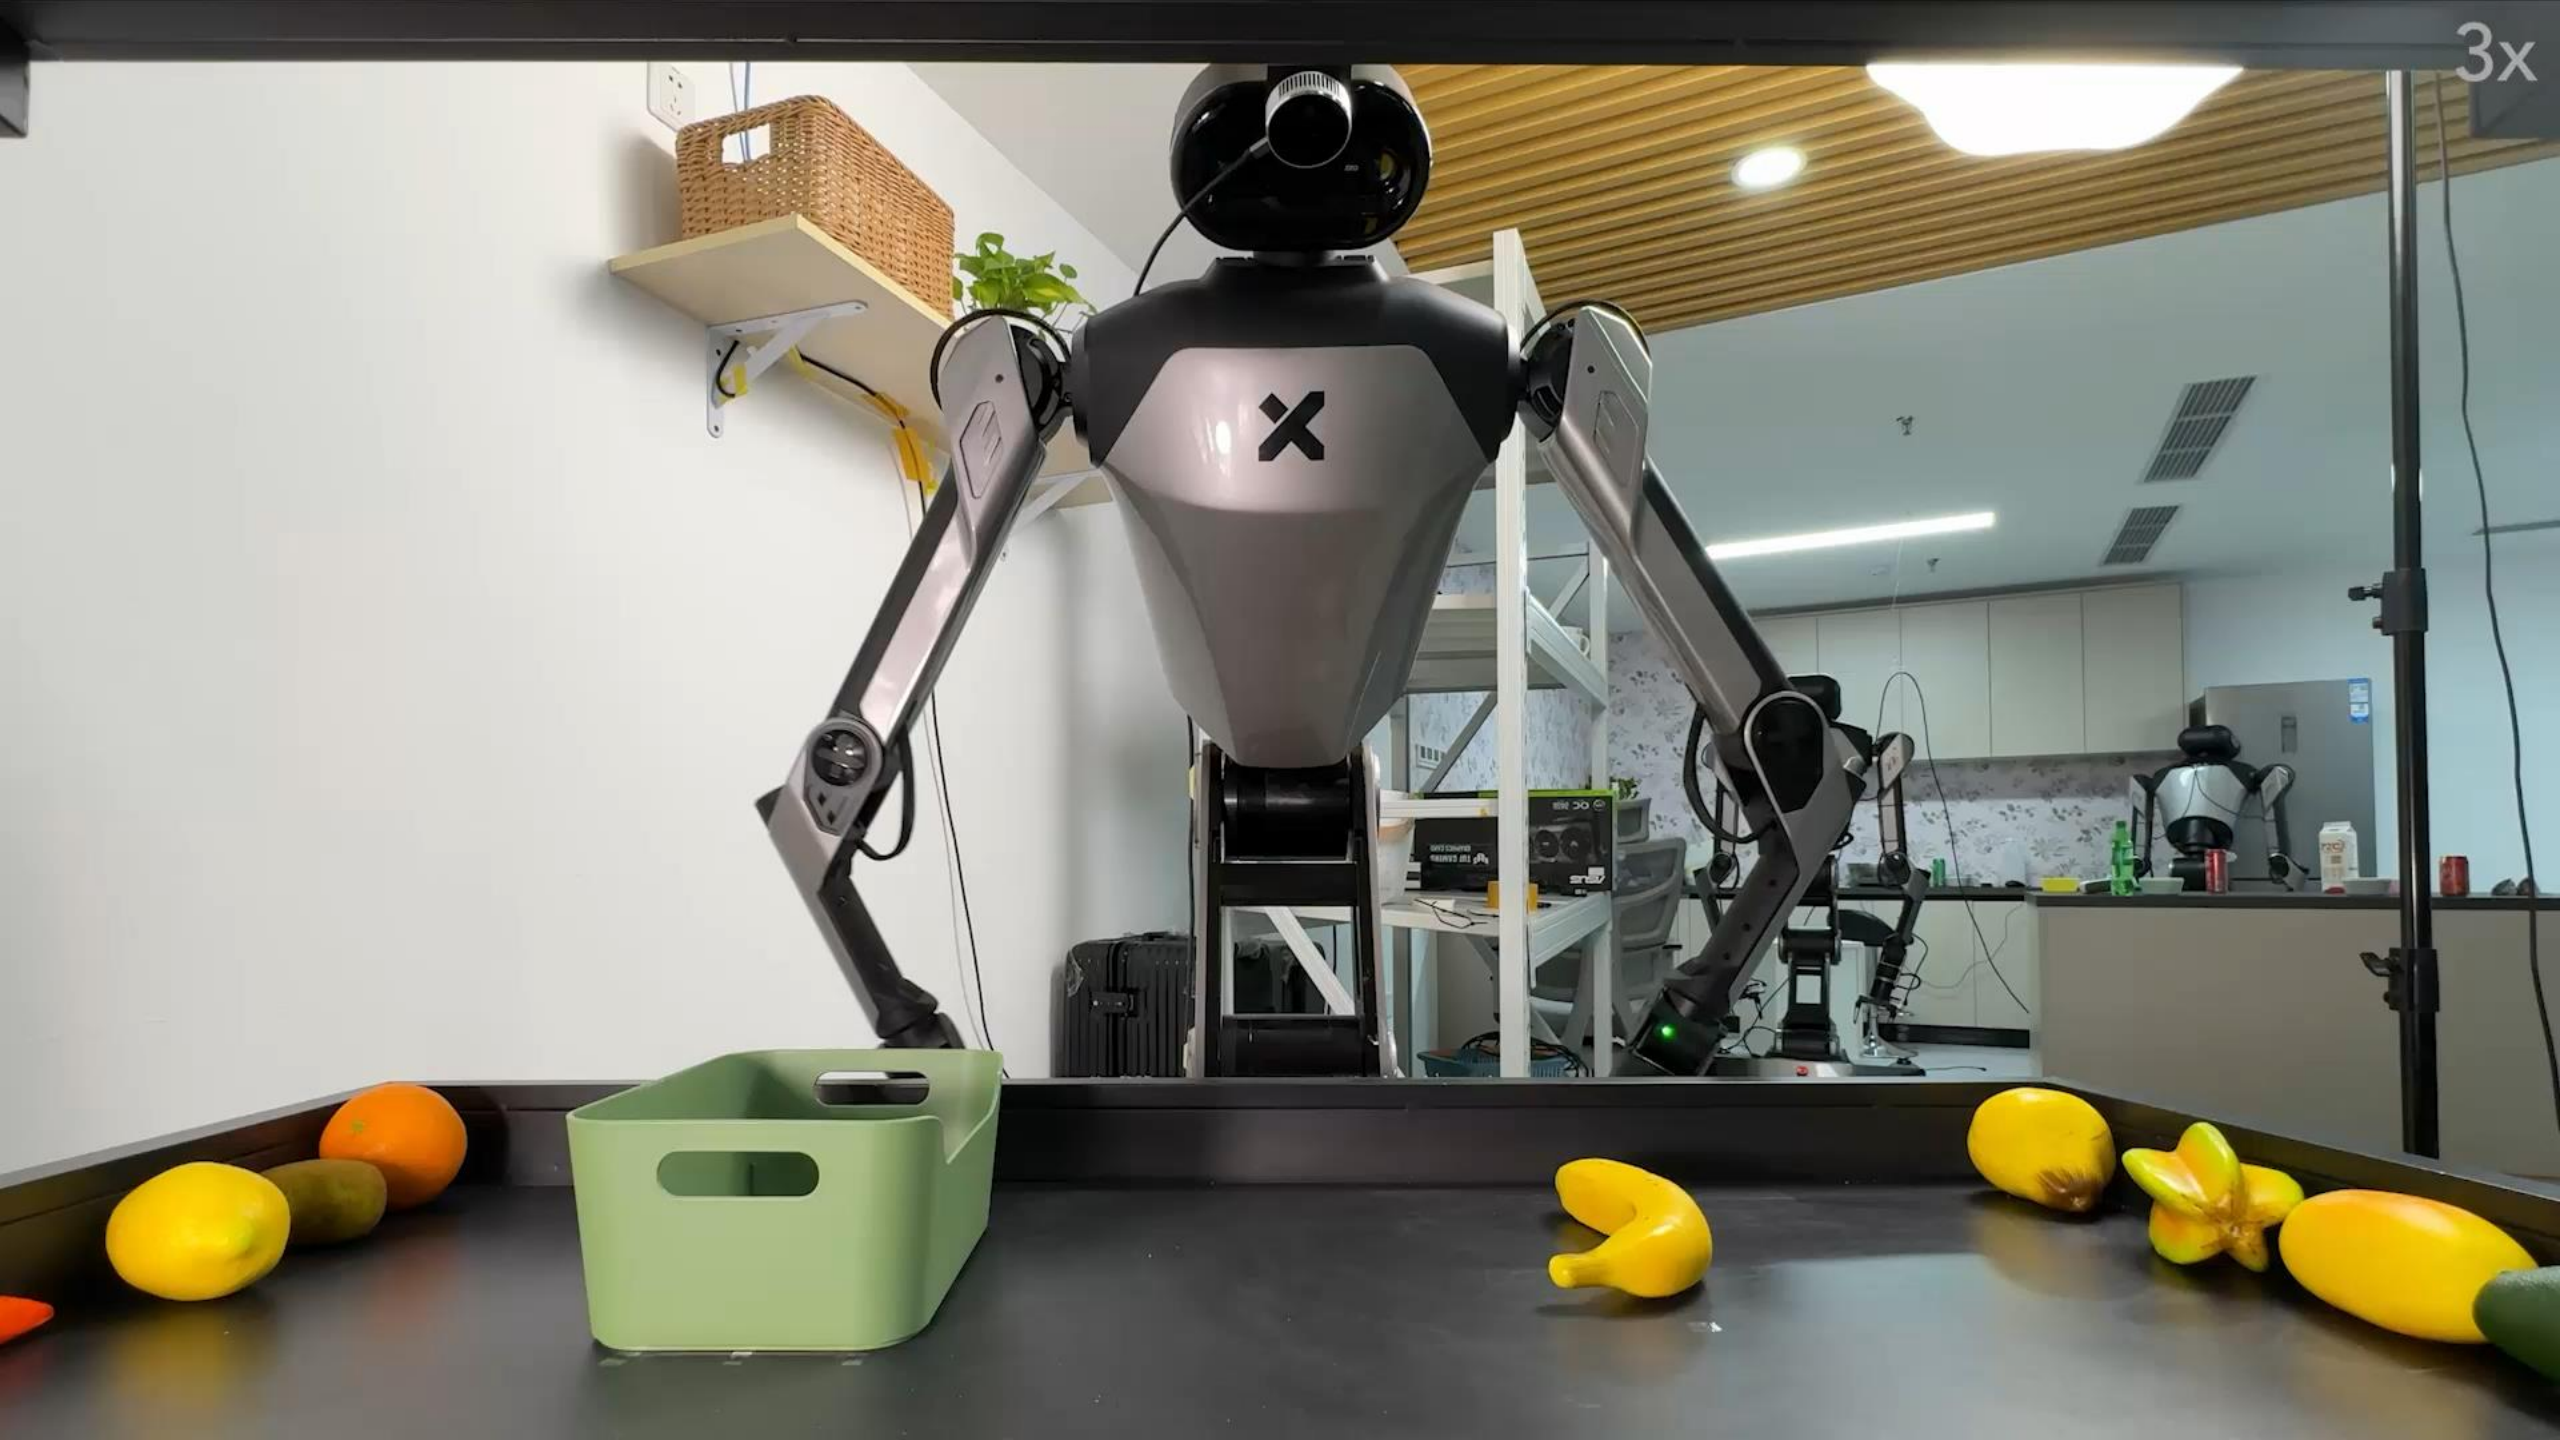
3x

Generation Cost

MimicGen
Expensive on-robot rollouts

DemoGen
Fully synthetic, cost-effective

	Single o-a Pair	A Trajectory	Whole Dataset
MimicGen	2.1 s	2.1 min	83.7 h
DemoGen	**0.00015** s	**0.010** s	**22.0** s

Time cost for generating real-world demonstrations

Ajay Mandlekar, Soroush Nasiriany, Bowen Wen, Iretiayo Akinola, Yashraj Narang, Linxi Fan, Yuke Zhu, and Dieter Fox. Mimicgen: A data generation system for scalable robot learning using human demonstrations. CoRL'23.

Thank you, but two more things!

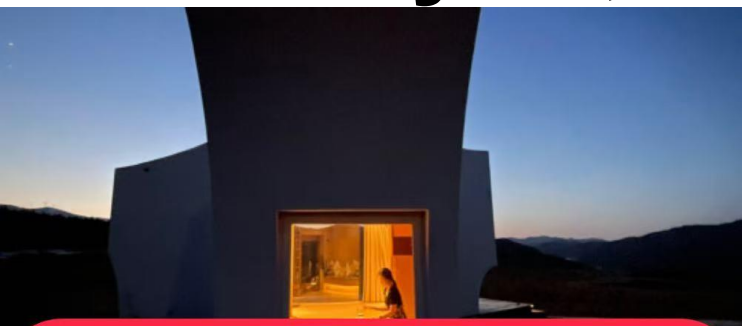

许华哲 Harry

小红书号：26214366243

Slow science. Cool robots. 清华姚班助理教授. 星海图联创

小红书

扫描二维码
在小红书找到我

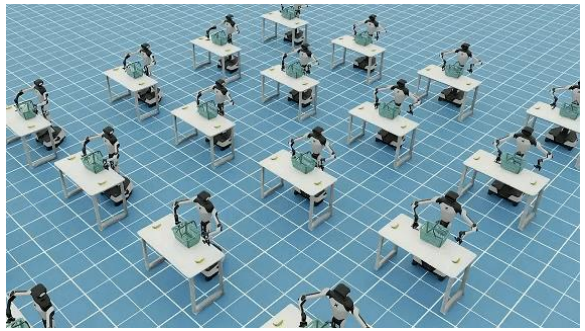

星辰大海，永不止步

欢迎关注星海图公众号

访问官网 https://galaxea.ai 或 http://galaxea.tech

了解更多关于星海图团队以及产品技术的详细信息

Hugging Face