Low-Cost Exoskeletons for Learning Whole-Arm Manipulation in the Wild

Hongjie Fang*, Hao-Shu Fang*, Yiming Wang*, Jieji Ren, Jingjing Chen, Ruo Zhang, Weiming Wang, Cewu Lu
(* = equal contribution)

Overview

Humans can use parts of their arms other than the hands for manipulations like holding a baby and closing doors, which we refer to as whole-arm manipulations. This approach is more complex than typical manipulation using only end-effectors. In this paper, we use this challenging scenario as an illustration and introduce a novel toolkit to facilitate effortless robot demonstration collection for imitation learning, even enabling data collection in unstructured environments without the need for a physical robot.

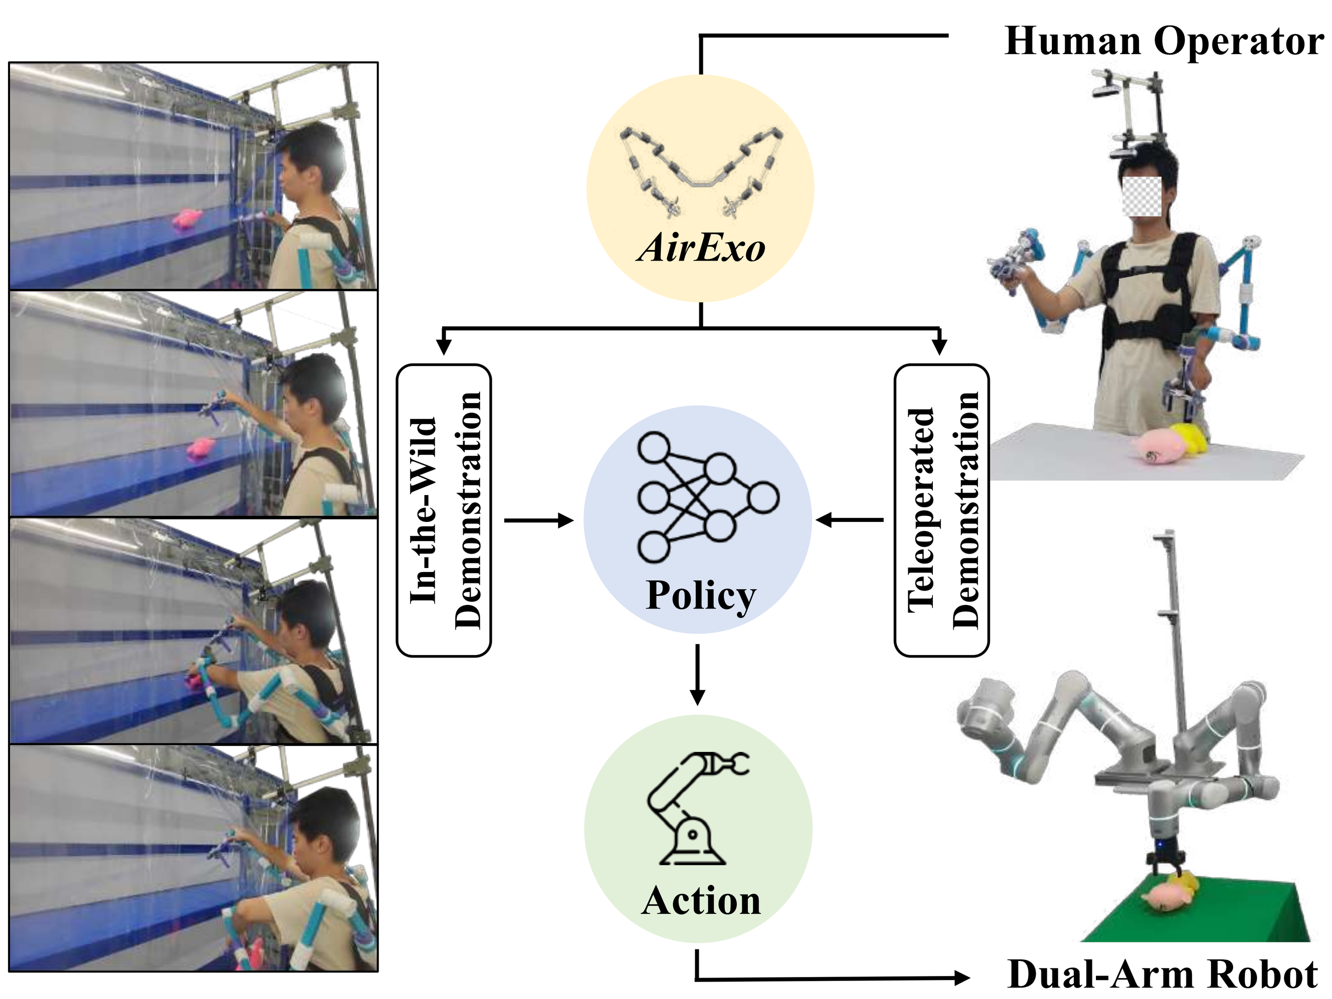

Fig. 1. The methodology of our in-the-wild learning framework with low-cost exoskeletons *AirExo*. It empowers the human operator to not only control the robots for collecting teleoperated demonstrations but also directly record human demonstrations in the wild. Our framework leverages both demonstrations in policy learning, resulting in a more general and robust policy compared to training with even more teleoperated demonstrations.

AirExo: An Open-Source, Portable, Adaptable, Inexpensive and Robust Exoskeleton

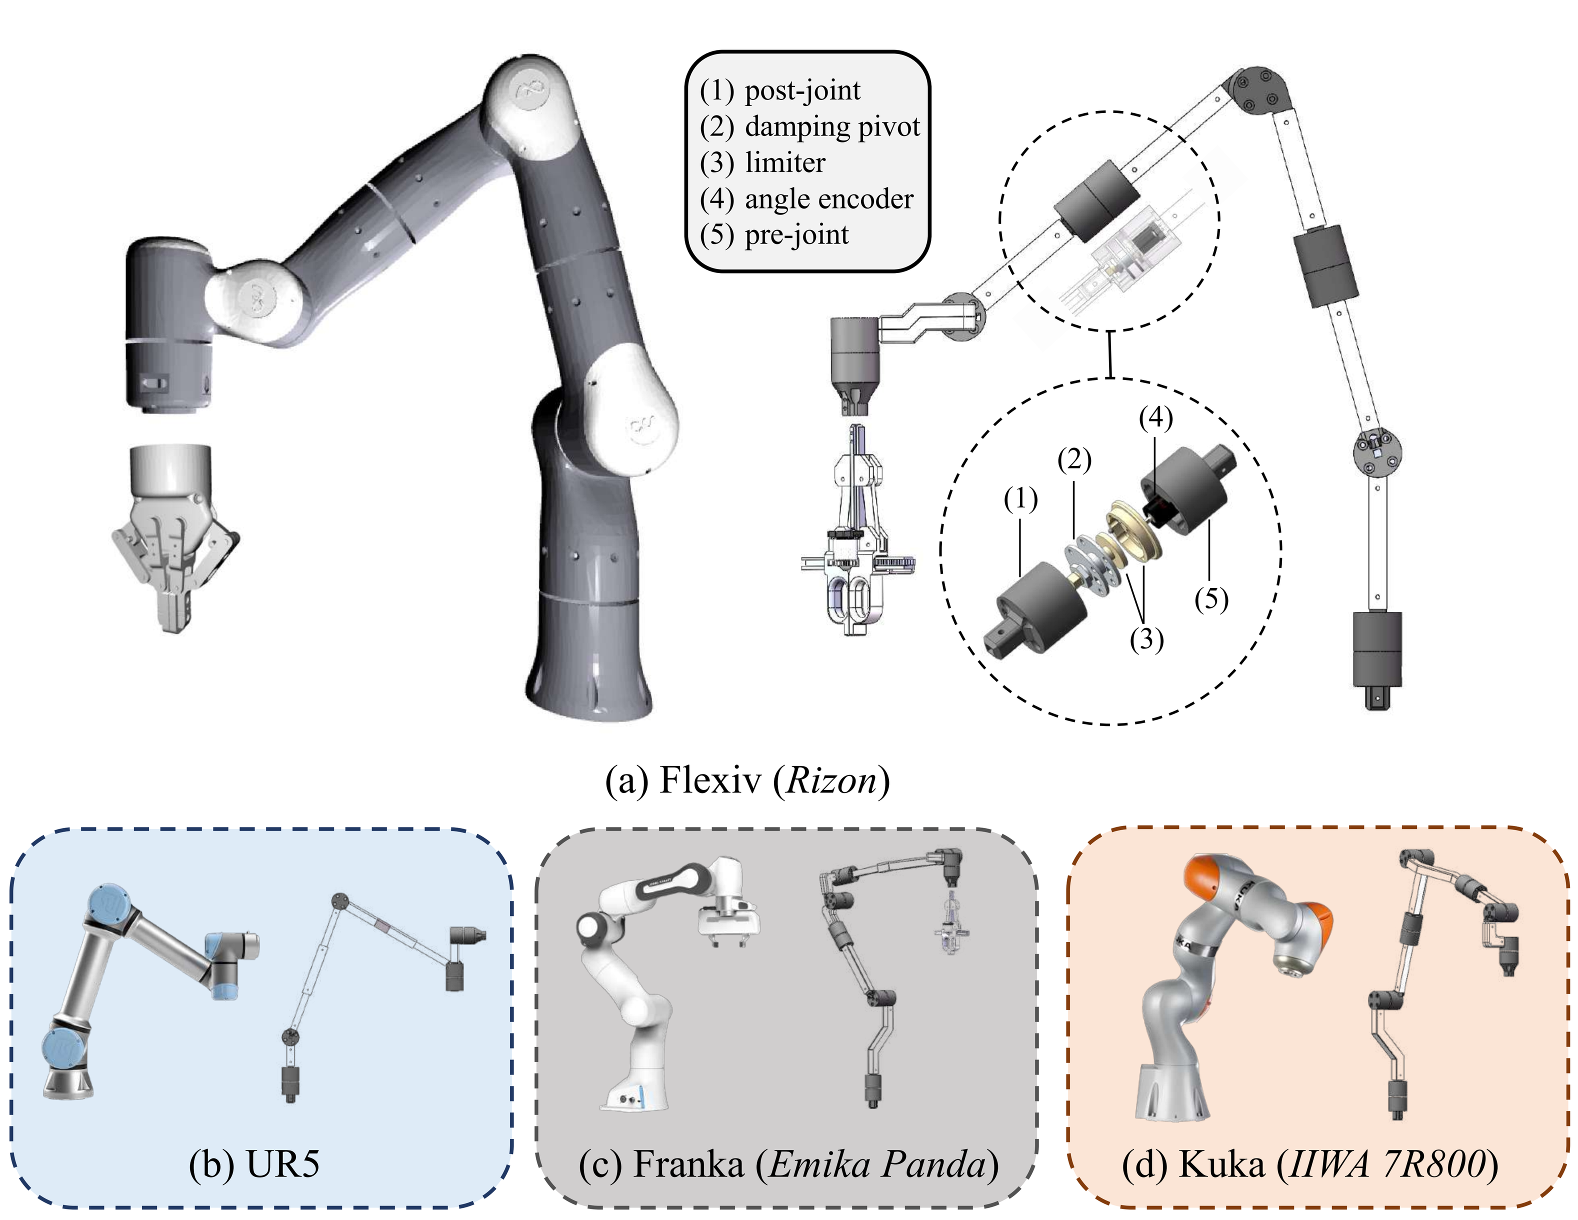

Fig. 2. *AirExo* models for different types of robots. Notice that the internal structure of the joints is standardized, only the linkages are altered to accommodate different robotic arm configurations.

With *AirExo*, the human operator can perform accurate teleoperation of the robot intuitively. Moreover, the portability and the one-to-one joint mapping property of *AirExo* allows human operator to collect demonstrations in the wild at scale, without needing a robot.

In-the-Wild Learning Framework

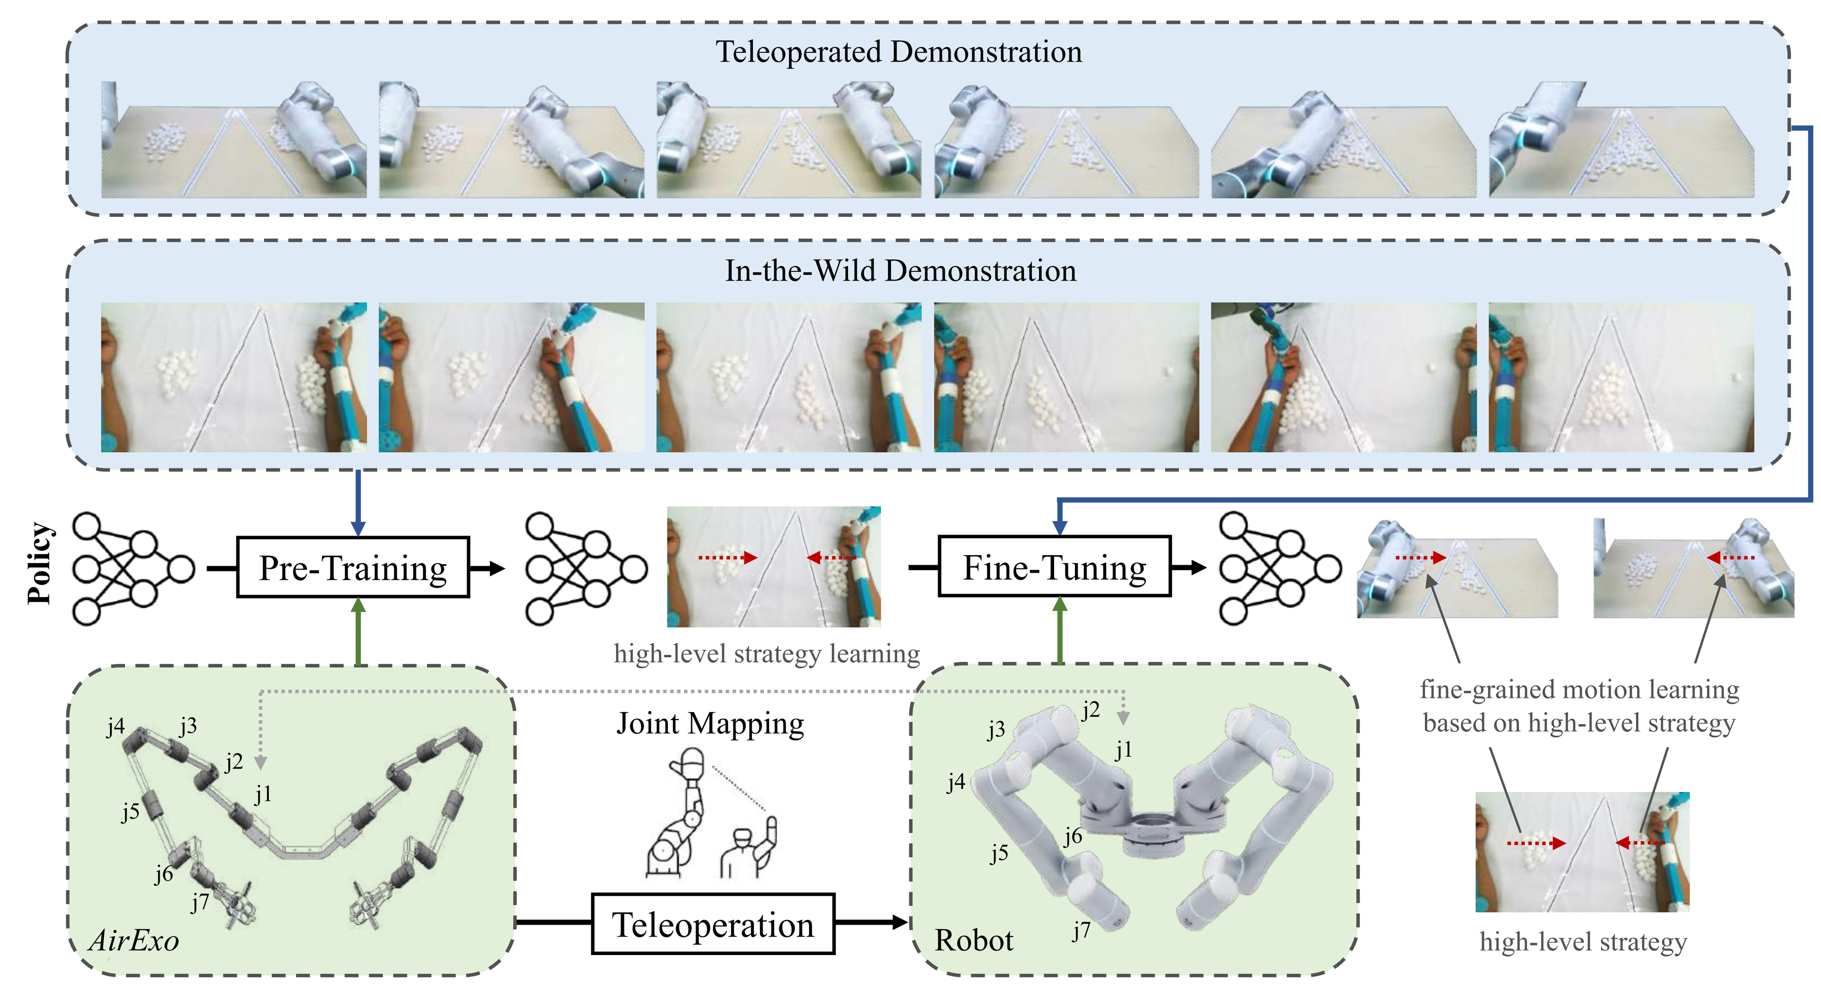

Fig. 3. Overview of learning whole-arm manipulations in the wild with *AirExo*. First, we use in-the-wild demonstrations and exoskeleton actions that are transformed into the robot's domain to pre-train the policy, which corresponds to learning the high-level strategy of task execution. Then, we use teleoperated demonstrations and robot actions to fine-tune the policy, which corresponds to learning fine-grained motion based on the learned high-level strategy.

Results

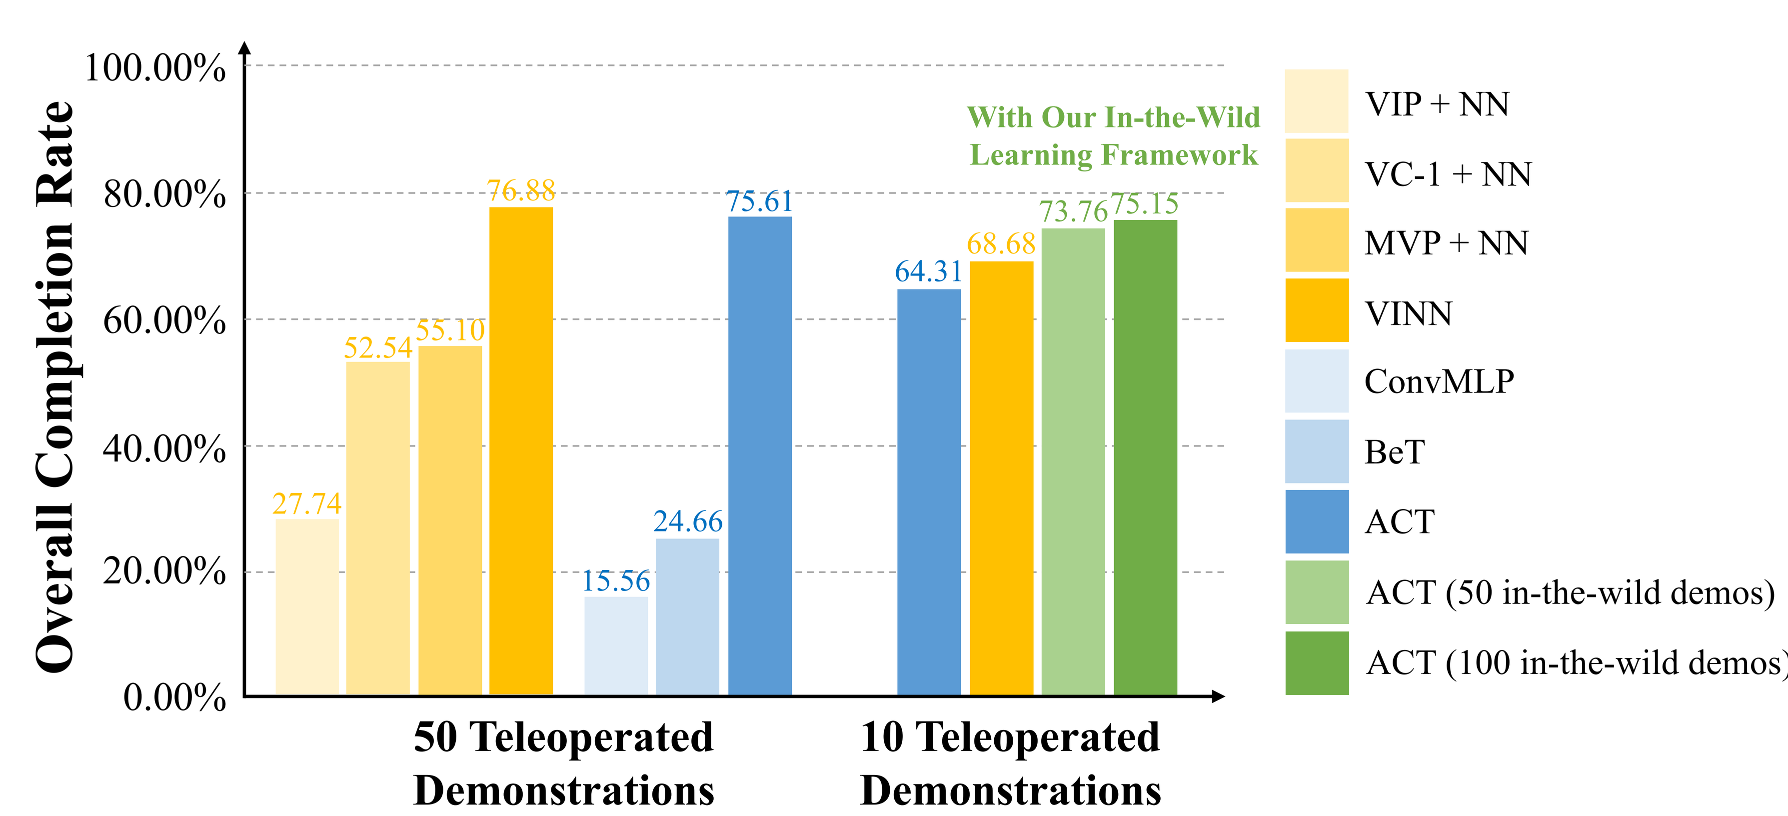

Fig. 4. Results of the "*Gather Balls*" task.

# Demos: Teleoperated	# Demos: In-the-Wild	Method	Success Rate (%) ↑: Reach in	Push aside	Approach	Grasp	Throw
50	-	VINN [25]	**100**	96	92	60	48
50	-	ACT [45]	**100**	**100**	**100**	**84**	**84**
10	-	VINN [25]	**100**	84	84	60	44
10	-	ACT [45]	**100**	**100**	96	72	44
10	50	ACT [45]	**100**	**100**	96	76	76
10	100	ACT [45]	**100**	**100**	**100**	**92**	**88**

Tab. I. Results of the "*Grasp from the Curtained Shelf*" task.

Disturbance	w/wo In-the-Wild Learning	Success Rate ↑ # Success / # Total
Novel Object	✘	4 / 8
	✔	**7** / 8
Different Background	✘	2 / 8
	✔	**6** / 8
Visual Distractors	✘	4 / 8
	✔	**8** / 8

Tab. II. Results of the robustness experiments on the "*Grasp from the Curtained Shelf*" task.

For more videos, see

airexo.github.io