Robotics: Science and Systems 2026
Sydney, Australia, July 13–17, 2026

RIO: Flexible Real-Time Robot I/O for Cross-Embodiment Robot Learning

Pablo Ortega-Kral*[1], Eliot Xing*[1], Arthur Fender Coelho Bucker[1], Vernon Luk[1], Junseo Kim[1,2], Owen Kwon[1], Angchen Xie[1], Nikhil Sobanbabu[1], Yifu Yuan[1], Megan Lee[1,4], Deepam Ameria[1], Bhaswanth Ayapilla[1], Jaycie Bussell[3], Guanya Shi[1], Jonathan Francis[1,3], and Jean Oh†[1,4]

[1]Carnegie Mellon University [2]TU Delft [3]Bosch Center for AI [4]Lavoro AI

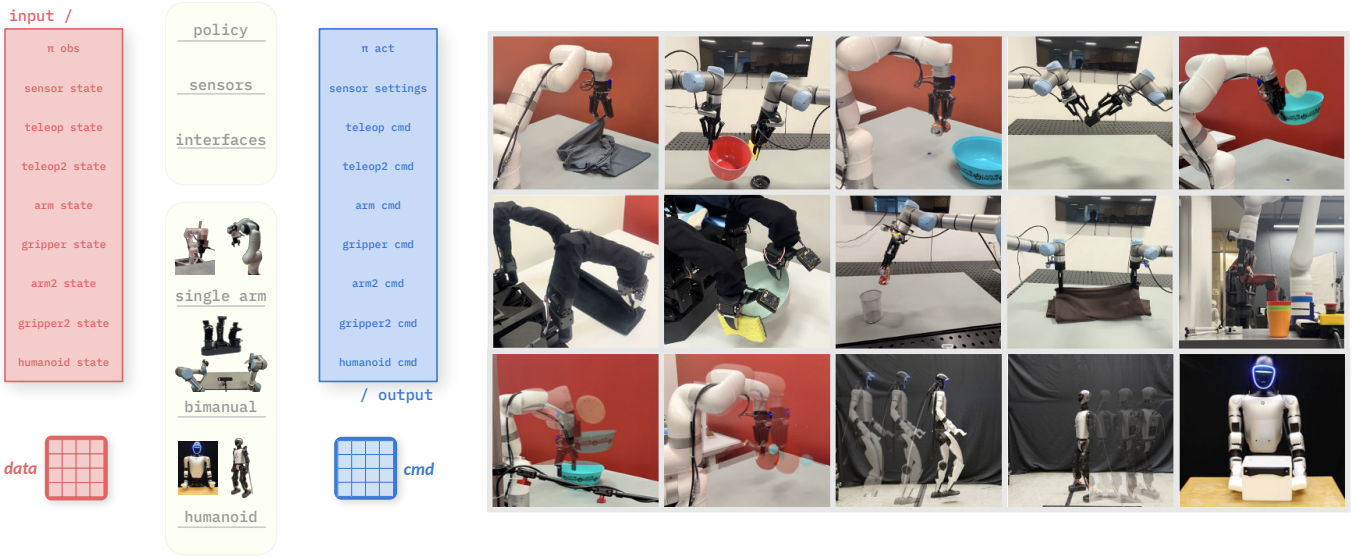

Fig. 1: **System overview.** We introduce RIO, flexible real-time Robot I/O for cross-embodiment robot learning, a lightweight Python-based framework to coordinate diverse robot morphologies, sensors, teleoperation interfaces, and policies.

***Abstract*—Despite recent efforts to collect multi-task, multi-embodiment datasets, to design recipes for training Vision-Language-Action models (VLAs), and to showcase these models on different robot platforms, generalist cross-embodiment robot capabilities remains a largely elusive ideal. Progress is limited by fragmented infrastructure: most robot code is highly specific to the exact setup the user decided on, which adds major overhead when attempting to reuse, recycle, or share artifacts between users. We present RIO (Robot I/O), an open source Python framework that provides flexible, lightweight components for robot control, teleoperation, data formatting, sensor configuration, and policy deployment across diverse hardware platforms and morphologies. RIO provides abstractions that enable users to make any choice and to switch between them, with minimal reconfiguration effort. We validate RIO on VLA deployment workflows across three morphologies (single-arm, bimanual, humanoid) and four hardware platforms with varying grippers and cameras. Using teleoperated data collected with RIO, we fine-tune state-of-the-art VLAs including $\pi_{0.5}$ and GR00T on household tasks such as pick-and-place, folding, and bowl scrubbing. By open sourcing all our efforts, we hope the community can accelerate their pace of robot learning on real-world robot hardware. Additional details at: `robot-i-o.github.io`.**

*Equal contribution, †Corresponding author: `jeanoh@cmu.edu`

I. Introduction

Vision-language-action models (VLAs) have recently emerged as a promising approach for training generalist robot policies, leveraging large-scale cross-embodiment datasets to learn broadly capable robot behaviors. Despite their potential, successfully deploying VLAs on new robot morphologies and platforms remains challenging in practice. VLAs are difficult to run out-of-the-box on new embodiments, and getting them to work still demands substantial engineering effort. This difficulty, however, extends well beyond VLAs.

Robotics practitioners have long contended with the fragmentation inherent in the field. Varying morphologies, diverse sensor configurations, heterogeneous hardware platforms, and manufacturer-specific driver code collectively result in robot infrastructure that is highly specific to a user's particular setup. This leads to significant overhead when attempting to reuse code, share datasets, or build on each other's work. Existing cross-embodiment datasets like Open X-Embodiment [36] are, in practice, aggregations of many individual collection efforts conducted across disparate infrastructure.

The cost of this fragmentation is growing. As robot hardware becomes increasingly affordable, more platforms are entering research labs and deployment settings. Yet, the specialized nature of most robotics infrastructure means that each new platform carries substantial integration overhead. Consider a common scenario: a research team wishes to reproduce real-world results released by another group. To use the original control code, they would need to replicate the exact hardware setup one-to-one, as in efforts like DROID [27]. If they instead have a different robot arm, then they face the burden of rewriting the entire control stack from scratch, before even trying to adapt any learned policies. This makes most robot learning hardware code *difficult to reuse*, and switching between platforms far harder than it should be.

What is the most important infrastructure for robot learning to advance? Beyond large datasets, we believe that a lack of flexible, reusable, accessible, and performant full-stack robot infrastructure has been a critical barrier to cumulative progress and collaboration within the field. Robot learning is missing reusable building blocks for hardware with flexible abstractions that have become standard in other areas of machine learning. Just as specialized high-performance GPU kernels within high-level auto-differentiation frameworks have enabled the rapid development and iteration of neural networks, robotics requires analogous foundational components for hardware and control that can be reliably shared, extended, and built upon across the community.

In this paper, we present the following contributions:

- **i)** We introduce RIO, a flexible real-time Robot I/O framework that provides reusable infrastructure for data collection and policy deployment across diverse robot embodiments. RIO is not intended as a comprehensive robot learning solution, but rather a lightweight set of reusable building blocks that can be quickly combined to deploy policies on real robots, tailored to each user's desired configuration. RIO is designed to be flexible, reusable, accessible, performant, and consistent, with abstractions such that the user is free to make any choice at each layer of the stack and switch between them with minimal effort.
- **ii)** We validate RIO for the VLA deployment workflow spanning diverse embodiments across single arm, bimanual, and humanoid robots with different grippers and sensors. This includes different robots, cameras, teleoperation interfaces, middlewares, data formats, and policies.
- **iii)** We demonstrate real-world deployment by collecting teleoperated data to fine-tune state-of-the-art VLAs such as $\pi_{0.5}$ and GR00T, on household tasks such as pick-and-place, folding, and bowl scrubbing, including a cross-embodiment policy trained on mixed data from two different robot morphologies and platforms.

II. Related Works

A. Generalist Robot Policies

Recent advances in vision-language-action models (VLAs) [7, 19, 5, 24, 40, 4] aim to leverage the robust image-to-language alignment learned by internet-scale pre-trained vision-language models (VLMs) [50, 18, 3, 2, 15] to train generalist robot policies. VLAs adapt VLMs to predict actions through imitation learning on robot datasets collected via human teleoperation of robots, scaling foundational work on imitation learning for visuomotor policy learning, such as ALOHA [51] and Diffusion Policy [13]. Due to the computational resources and data scale required, state-of-the-art VLAs are predominantly trained by industry labs with substantial infrastructure and engineering personnel. Open source efforts have sought to reproduce and democratize these results [44, 31, 39, 15, 28], providing fully open source implementations and model weights. However, a significant limitation remains: current VLAs must in practice be fine-tuned for each robot setup. Released VLA model checkpoints are typically fine-tuned for specific embodiments, such as the Franka arm from DROID [27] or the WidowX arm from BridgeData V2 [42]. Consequently, end-users must either reproduce the exact hardware setup used during training [43], or undertake substantial engineering effort to implement their own robot control stack, before even attempting to adapt learned policies to their own platforms. In this work, we lower this barrier by introducing a flexible cross-embodiment robot control stack, validated on the VLA adaptation workflow and deployed across diverse robot configurations.

B. Cross-Embodiment Robot Data

The effectiveness of scaling VLAs depends on access to large-scale robot demonstration data. Prior work has established that scaling robot data across both task diversity and robot embodiments [14, 45, 11, 17, 41, 47] shows promise at training better generalist robot policies, and cross-embodiment robot data may also enable learning directly from humans [26]. Training truly general robot policies requires diversity in both tasks and embodiments. Open X-Embodiment [36] aggregates 60 datasets spanning over 1 million robot trajectories across 22 embodiments. However, the heterogeneous collection techniques and sensor configurations across these datasets necessitate substantial curation for effective policy training, such as through filtering [19, 27] or data mixture re-weighting [23]. In this work, we aim to facilitate the collection of high-quality cross-embodiment robot data by developing flexible and reusable robot infrastructure.

C. Robot Control Stacks

Over the years, many robot control stacks have emerged [34, 29, 16, 9, 53, 30, 35, 25, 20, 33, 12] that are capable of cross-embodiment robot control. ROS [33] was developed to facilitate system compartmentalization and distributed communication. While this modular, distributed approach offers benefits for complex robotic systems, it requires compounding systems-level engineering to coordinate all modules together. Furthermore, ROS presents a high barrier to entry for researchers and practitioners new to robotics, as it requires wrangling a complex configuration and build system.

Table I: **Comparison of cross-embodiment robot stacks.** We compare various cross-embodiment robot stacks on the basis of native platform support, at different layers of the robot learning pipeline, e.g., data-collection system support, robot hardware support, middleware, data formats, and policy architecture support. For example, some stacks combine robot arm and robot gripper drivers, making it difficult to use other end effectors on arms.

Framework	Humanoids	Bimanual	Single arm	Robot grippers	Teleop	Cameras	Middleware(s)	Data format(s)	Policies
Ark [16]	✓	✓	✓	✗	✓	✓	✗: LCM	✗: Pickle	✓
LeRobot [9]	✓	✓	✓	✗	✓	✓	✗: Threads/gRPC[1]	✗: LeRobotDataset	✓
ManiUniCon [53]	✗	✗	✓	✗	✓	✓	✗: Shm	✗: Zarr	✓
PAPRLE [30]	✓	✓	✓	✓	✓	✓	✗: ROS	✗: Pickle	n/a
PyRobot [35]	✗	✗	✓	✓	✓	✓	✗: ROS	✗: Pickle	n/a
RCS [25]	✗	✗	✓	✓	✓	✓	✗: RPC	✗: Parquet	✓
RoBits [20]	✗	✓	✓	✓	✓	✓	✗: ZMQ	✗: NPZ/JSON	n/a
UMI, DP [12, 13]	✗	✓	✓	✗	✓	✓	✗: Shm	✗: Zarr	✗: DP
RIO (ours)	✓	✓	✓	✓	✓	✓	✓: **any**	✓: **any**	✓

[1] *LeRobot uses Threads for hardware drivers and gRPC for asynchronous policy inference.*

Despite this proliferation of frameworks, robot code remains highly platform-specific. We attribute this to two factors: first, most roboticists work with a single hardware platform, which incentivizes writing vendor-specific code quickly rather than abstractable solutions; second, existing frameworks lack flexible abstractions at every layer of the stack to ensure cross-embodiment robot code is easy to write in the first place. LeRobot [9] has also found mainstream adoption among the broader robotics community. Its popularity stems in part from its support for low-cost robot hardware along with Python-only implementations, which eliminates the need for complex build systems such as in ROS and multi-language dependencies. By streamlining the development workflow, LeRobot has lowered the barrier to deploying real-world robot systems and enabled a wider range of practitioners to experiment with robot learning. RIO shares the motivation to build an accessible open source community for cross-embodiment robot learning.

In Table I, we compare the flexibility of RIO to a variety of existing cross-embodiment robot infrastructure, across every layer of the robot learning stack. RIO offers a reusable Python-based set of real-time robot infrastructure in a similar style to DP/UMI and LeRobot, while supporting reconfigurability at each layer to facilitate VLA deployment workflows across diverse robot morphologies. RIO enables combinatorial configuration of robots, teleoperation devices, cameras, middlewares, data formats, and policies, for maximum flexibility.

III. RIO (Robot I/O)

RIO is a Python-based framework for flexible real-time Robot I/O, with reusable components for robot control, teleoperation, data collection, and policy deployment across diverse robot embodiments. Users are free to make any choices at every layer of the stack (humanoid robots, robot arms, robot grippers, teleoperation interfaces, cameras, middlewares, data formats, policies) and to switch between them with minimal effort for reconfigurability. See Table II for a detailed list of currently supported hardware.

Figure 2 illustrates the overall system architecture of RIO. Section III-A describes the *Design Philosophy*, Section III-B introduces the client-server *Nodes* abstraction and *Middlewares* implementation for message passing, Section III-C outlines the reconfigurable instantiation of *Robot Stations*, Section III-D describes *Teleoperation* and *Data Collection* used to collect real-world robot trajectories across multiple embodiments, and Section III-E describes the implementation of asynchronous *Policy Inference* to obtain smooth real-world rollouts with observation/action chunking.

Table II: **Current hardware support** (*more incoming*). RIO provides flexibility across robot hardware (humanoid robots, robot arms, robot grippers), teleop interfaces, cameras, and middlewares for multi-process distributed communication. These can be combined in any configuration depending on the user's hardware requirements.

Humanoid Robots	Unitree G1, Booster T1
Robot Arms	UFactory (xArm5/6/7, 850, Lite6), UR (UR5e, UR7e), Franka (FR3, Panda), Kinova (Gen3), SO-100/SO-101
Robot Grippers	UFactory Gripper, Franka Gripper, Robotiq Gripper (2F-85/2F-140), DH-Robotics Gripper (AG-105-145)
Teleop Interfaces	Spacemouse, Gamepad, Keyboard, VR (Apple Vision Pro, Meta Quest, Pico 4 Ultra), Leader-Follower (GELLO), Phone
Cameras	RealSense, ZED, UVC (Webcams, USB cameras), iPhone (Record3D)
Middlewares	Shared Memory, Thread, Portal, Zenoh, ZeroRpc

A. Design Philosophy

We describe the five design tenets of RIO: *flexible*, *reusable*, *accessible*, *performant*, and *consistent*.

- *Flexible.* RIO is agnostic to each component and does not make any locked-in choices. The user is free to choose between options at every layer, and switch between them with minimal effort.
- *Reusable.* RIO is composed of a lightweight set of reusable building blocks, that can be quickly combined and modified to support a user's desired configuration.
- *Accessible.* RIO is quick to install and uses Python-based hardware modules, with a command-line interface over a single configuration file.
- *Performant.* RIO is fully capable of high-frequency real-time robot control, and uses asynchronous policy inference to yield smooth robot trajectories.

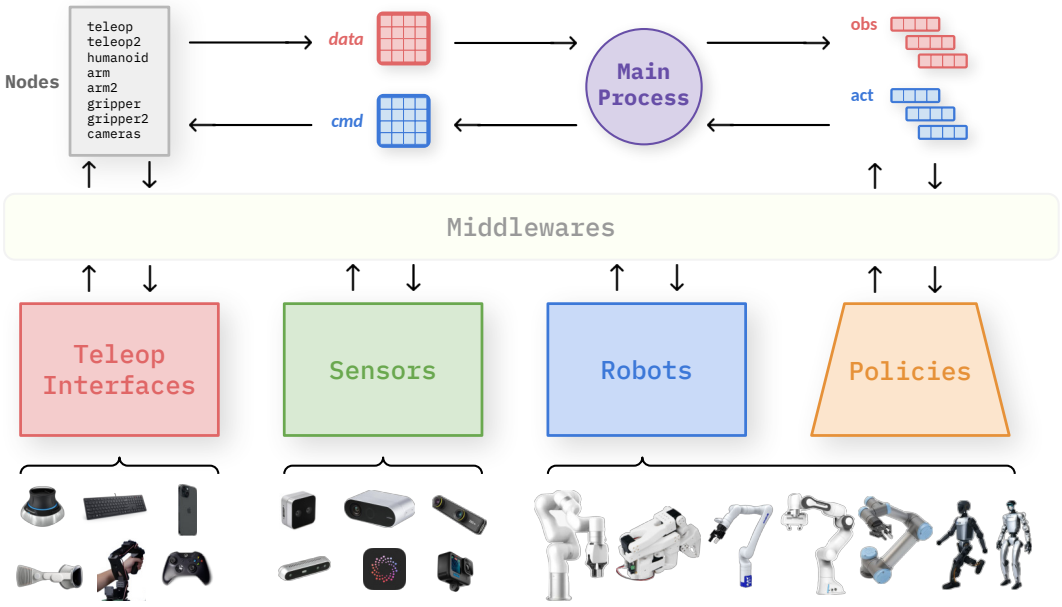

Fig. 2: **Architecture.** High-level overview of the architecture of RIO. Every component of the stack is flexible, meaning that the user is free to choose between different options (robots, sensors, teleoperation interfaces, middlewares, data formats, policies) and switch between them, with minimal effort.

- *Consistent.* RIO is designed to ensure consistent, scalable, reproducible data collection and robot learning.

Taken together, our design values serve a single goal: abstract user logic should be written once and should be reusable on any robot station. Although a wide range of robot infrastructure exists today (Table I), most robot code is highly platform-specific, with systems rarely reusable outside their original hardware setup. For a given task, such as teleoperation or policy deployment, RIO exposes abstract scripts that operate on arbitrary embodiments without modification. Swapping robot platforms, cameras, teleoperation devices, or middlewares is a change to the configuration rather than to the control logic. With this design, RIO aims to facilitate the writing of reusable control stack components rather than platform-specific code.

B. Nodes and Middlewares

Nodes for teleoperation interfaces, sensors, robots, and policies are implemented from the same template Node, requiring minimal boilerplate to enable flexible, real-time I/O across diverse hardware and deployment configurations.

Nodes. A Node dynamically inherits from a given Middleware that automatically handles message passing. Factory functions produce matched server-client Node pairs, for which its dynamic parent class implements the specified Middleware. Nodes support three execution patterns for publishing data and handling requests between server-client pairs:

1) Publish-only: `pub()` publishes data in the run loop.
2) Request-only: `req()` handles requests in the run loop.
3) Combined: `pubreq()` publishes data and handles requests in a single loop.

For patterns (1) and (2), the complementary operation can be optionally run in a separate worker loop thread. For example, `pub()` in the run loop with `req()` in a worker loop thread, or vice versa. To implement a Node, the user defines the relevant methods: a `pub()` implementation calls `ring_buffer.put(..)` to publish data, while a `req()` implementation calls `request_queue.get()` to receive and process requests. Published data flows through a ring buffer that continuously streams state at a fixed frequency, providing time-synchronized access to sensor readings, robot poses, and other data. Requests flow through a queue that enables asynchronous command communication, allowing multiple clients to send timestamped commands independently and at arbitrary rates. For each, a server Node executes "publish/request", while a client Node automatically resolves "subscribe/reply". The user specifies `example_data` and `example_request` in each Node definition, which are used to infer buffer shapes and data types when initializing ring buffers and request queues. Each Node exposes a public API (defined via `__api__`), whose methods are automatically wrapped for serialization on the server side and deserialization on the client side, enabling transparent bidirectional communication. Because Nodes are middleware-agnostic, they can be paired with different middleware backends depending on deployment requirements. Process synchronization is managed through ready and exit events that signal when "publish/request" loops are ready or exited, ensuring that user logic in the main process blocks until every Node has fully initialized and that the process cleans up when it completes.

Middlewares. Nodes interact with Middleware through a common interface, hiding transport-level details. For network-based communication, middlewares such as Zenoh or ZeroRpc

handles serialization and transport over TCP or IPC. For high-throughput local communication, shared memory middleware uses a SharedMemoryManager to allocate ring buffers and request queues in shared memory, allowing zero-copy data exchange between processes. The shared memory implementation runs the Node's loops in a separate process and communicates buffer handles back to the parent through pipes, enabling multiple processes to access the same underlying memory regions without serialization overhead. Middlewares can be interchanged depending on the user's requirements. For instance, switching to the Thread middleware can help with debugging when orchestrating many Nodes on the same computer, since running everything multi-threaded within one process simplifies stack traces, breakpoints, and exception handling compared to multi-process or networked deployments. Alternatively, for embodiments such as mobile manipulator robots that may not have a powerful onboard computer, the user may use network-based middleware to communicate between multiple machines.

C. Robot Stations

We aggregate the instantiation of all nodes that make up an environment into a single station configuration file. A robot station is instantiated through a composable dataclass configuration that specifies the hardware topology of the deployment, defining the set of robots, end effectors, and sensors, e.g., {`arm, gripper, arm2, gripper2, wrist_camera, wrist_camera2`} for a bimanual robot station. The effect is to simplify the main routine's logic in robot control loops: a context manager initializes all server Nodes with the specified middleware backend, while client Nodes are started within nested context managers that yield proxy objects for transparent communication. While the nodes internally use the specified middleware, queues, and ring buffers, the main routine interacts only with the APIs, resulting in simple, Python-based code. This pattern enables the same application logic to operate over arbitrary station configurations without modification.

D. Teleoperation and Data Collection

Teleoperation. We design three teleoperation scripts to support data collection. The first one controls relative end-effector poses (Keyboard, Phone, Gamepad, Spacemouse), and the second maps absolute joint positions using a leader-follower setup (ALOHA [51], GELLO [46]), for either single-arm or bimanual tabletop robot arms. The third script uses wrist pose retargeting from VR headsets such as Apple Vision Pro [37] to control the upper-body of a humanoid robot, similar to [21, 22]. These scripts do not assume any particular hardware platform, so teleoperation devices and robots can be swapped based on each user's needs. To ensure smooth teleoperation across devices and control frequencies, we include interpolators and signal-processing filters, e.g., low-pass filters. These can also be used during policy inference to mitigate the effects of noisy actions.

Data Collection. We define a recorder for logging robot demonstrations. To ensure consistent data collection across different hardware, we enforce standardized units: meters for world coordinates and radians for angular measurements. We adapt the RLDS-style format [38] to aggregate multiple data streams into a unified state representation, as shown in Figure 9. To support different robot platforms, we introduce the concept of morphologies: abstract descriptions of a robot's structure, each defining its own set of state keys. Each morphology overloads the observation field of the RLDS step with its specific keys. This design standardizes state reporting across platforms, regardless of the underlying hardware.

One challenge with scaling robotic data collection is the sheer storage required to manage it, given that a typical robot demo consists of the internal state of the robot, as well as multiple camera streams (often including depth) and other relevant sensors. Additionally, depending on the specific policy learning setup, the data may need to be exported and processed differently. To this end, we use RoboDM [10], a toolkit that employs flexible compression schemes and streamlined file structures, to efficiently record demos in a highly compressed, lightweight format that is quick to save and load. For training or fine-tuning, we encourage users to write minimal converters so that the exported demos can directly integrate with their intended dataloader format.

E. Policy Inference

Aside from providing lightweight building blocks for hardware components, we also want to reduce the overhead needed to swap between different policies. To reduce the challenge of integrating robot policies into our control stack, we design a high-level API for policies. For each policy, we only require a lightweight interface to instantiate the policy, convert observations from a standardized format to the policy-specific observation format, and run inference. We design a policy wrapper node that uses this API to directly instantiate policies and asynchronously handle inference requests. By handling inference requests directly through our middleware, we avoid the additional overhead of a dedicated policy server. Additionally, we design a configurable, policy-agnostic, hardware-agnostic inference script that queries the policy wrapper, handles logic for continuously obtaining observations from hardware, post-processes actions for smoothness, and sends commands to hardware. This allows for seamless switching between different policies and hardware.

IV. Evaluation

We evaluate whether RIO can support the robot learning workflow, from teleoperated data collection to fine-tuning and deployment, across diverse embodiments and policy classes (VLAs, Diffusion Policy, RL). Our experiments aim to answer the following questions regarding RIO's core functionalities:

- *Is RIO performant for real-time policy deployment?*
- *Can RIO support data collection across diverse morphologies that require different teleoperation interfaces?*

Fig. 3: **VLA manipulation trajectories.** We showcase rollouts of $\pi_{0.5}$ across 3 morphologies on 5 diverse tasks, shown at 0%, 20%, 40%, 60%, 80%, and 100% of task completion.

Table III: **Policy deployment.** We deploy state-of-the-art policies ($\pi_{0.5}$, GR00T N1.5, Diffusion Policy) across 3 morphologies (single arm, bimanual, humanoid) and two task regimes (quasi-static and dynamic), achieving $\geq$60% success across 20 trials on all tasks with finetuning on 50 teleoperated demonstrations. Policy rollout times closely match human demonstrations, and some tasks complete faster than demonstrations. Asynchronous inference maintains high GPU utilization throughout execution, showing that RIO efficiently saturates available compute. For the T-shirt folding task, half points were awarded for each fold.

Robot	Policy	Task	Success Rate	Task Completion Time (s)	Demo Time (s)	RAM (GB)	CPU (%)	GPU Util (%)	GPU Mem (%)
xArm7	BC $\pi_{0.5}$	Fold Shirt	92.5	41.96 $\pm$ 14.58	41.57 $\pm$ 9.25	22.5 $\pm$ 2.7	13.0 $\pm$ 1.4	56.7 $\pm$ 1.7	79.1 $\pm$ 0.0
xArm7	BC $\pi_{0.5}$	Place Can	95.0	16.08 $\pm$ 3.41	14.46 $\pm$ 2.00	24.8 $\pm$ 1.5	13.2 $\pm$ 1.4	54.6 $\pm$ 3.1	79.1 $\pm$ 0.1
SO-100	BC $\pi_{0.5}$	Fold Cloth	60.0	27.50 $\pm$ 5.51	22.43 $\pm$ 3.30	19.6 $\pm$ 0.5	14.1 $\pm$ 1.5	46.3 $\pm$ 10.0	78.6 $\pm$ 0.0
SO-100	BC $\pi_{0.5}$	Scrub Bowl	64.0	40.33 $\pm$ 13.68	27.66 $\pm$ 5.22	19.7 $\pm$ 0.7	15.0 $\pm$ 2.2	52.0 $\pm$ 4.8	78.6 $\pm$ 0.1
Unitree G1	BC GR00T N1.5	Pick Box	95.0	9.07 $\pm$ 6.10	10.38 $\pm$ 4.04	25.2 $\pm$ 0.1	26.6 $\pm$ 3.3	61.7 $\pm$ 4.7	33.8 $\pm$ 0.0
xArm7	BC DP	Flip Tortilla	66.7	12.36 $\pm$ 2.57	7.59 $\pm$ 0.79	17.2 $\pm$ 0.1	8.8 $\pm$ 1.9	8.6 $\pm$ 3.7	9.3 $\pm$ 0.4
xArm7	BC DP	Throw Ball	100.0	14.73 $\pm$ 1.28	13.26 $\pm$ 2.23	15.8 $\pm$ 0.8	6.8 $\pm$ 0.1	0.6 $\pm$ 0.8	9.1 $\pm$ 0.6
Unitree G1	RL PPO	Navigate	100.0	31.27 $\pm$ 6.56	N/A	23.0 $\pm$ 0.1	10.3 $\pm$ 0.4	5.1 $\pm$ 0.1	10.3 $\pm$ 0.1
Booster T1	RL PPO	Navigate	100.0	29.73 $\pm$ 4.49	N/A	22.6 $\pm$ 0.1	10.4 $\pm$ 0.4	5.3 $\pm$ 0.2	10.4 $\pm$ 0.3

- *Does RIO enable effective deployment of different policy classes (VLAs, Diffusion Policy, RL) across embodiments?*

Beyond this, we also show in Appendix V-B that RIO readily supports cross-embodiment robot learning workflows, though effectively leveraging such multi-embodiment data to improve policy performance remains an open research problem for future exploration.

Experimental Setup. All evaluations are performed on a desktop equipped with NVIDIA GeForce RTX 4090 GPU, AMD Ryzen 7 5700X CPU, and 64 GB of RAM.

A. Policy Fine-tuning and Deployment

RIO targets two stages of the robot learning pipeline: data collection for policy fine-tuning, and policy deployment. We adopt a bring-your-own-training-stack approach. Researchers collect demonstrations with RIO, export them to their preferred training format, and import the resulting policy weights back into RIO for deployment. This ensures consistency between data collection and policy rollout, while remaining agnostic to model architecture and training setup.

We validate this workflow across three policy families: VLAs ($\pi_{0.5}$, GR00T N1.5), Diffusion Policy (DP), and RL (PPO); spanning three morphologies (single-arm, bimanual, humanoid) and four robot platforms. The same RIO stack drives data collection and policy rollout for every case; only the policy and station configuration changes.

Training Setup. For $\pi_{0.5}$, we fine-tune[1] from the DROID checkpoint for single-arm tasks, and ALOHA checkpoint for bimanual tasks (20K steps each); for GR00T N1.5, we use 150 demonstrations on the humanoid manipulation task; for DP, we collect 50 demonstrations per task; locomotion policies

[1]We use one checkpoint per-task, per embodiment.

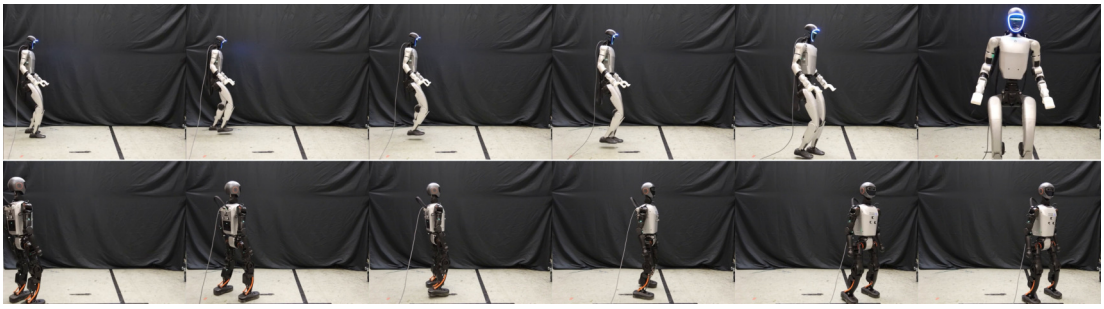

Fig. 4: **Humanoid locomotion trajectories.** RL policies on Unitree G1 (top) and Booster T1 (bottom), two humanoid robots from different manufacturers with different hardware drivers. RIO is capable of real-time control for humanoid locomotion

are trained with PPO in simulation. All demonstrations are collected at 50–80 Hz and stored in a compressed format exportable to each target training pipeline, e.g., 150 three-camera episodes occupy only 1.31 GB. To validate policies, we report success-rate metrics and compare the average task completion time with the average demonstration time.

Single-Arm Tabletop Manipulation. We use a UFactory xArm7 with a Robotiq 2F-140 gripper and three RealSense D400 cameras, on two task regimes with two different policies.

Tabletop tasks with $\pi_{0.5}$. For pick-and-place (*place can*, *fold shirt*), we collect demonstrations with both a Spacemouse and a GELLO leader-follower interface. Empirically, end-effector space devices such as the Spacemouse yield cleaner demonstrations for tasks where rotation is not a major factor. We also find that training VLAs with binary gripper actions yields poor gripper control, so we apply trajectory interpolation and low-pass filtering to smooth collected actions. As shown in Table III, both tasks achieve success rates above 90%, with completion times within 2 s of demonstration time.

Dynamic tasks with Diffusion Policy. Tortilla flipping and ball throwing are both dynamic tasks that demand precise, high-frequency control, which makes them well-matched for Diffusion Policy. We collect 50 GELLO demonstrations at 80 Hz and deploy DP through the same RIO pipeline. As reported in Table III, DP reaches 66.7% success on *flip tortilla* and 100% on *throw ball* across 20 trials, with execution times within a few seconds of the demonstration average.

Bimanual Tabletop Manipulation. We use SO-100 arms to evaluate RIO's support for bimanual coordination with $\pi_{0.5}$. RIO natively supports ALOHA-style robot-to-robot teleoperation; configuring the leader arms as teleoperation devices requires only a change to the station configuration file. We perform bimanual cloth folding and bowl scrubbing. All bimanual tasks achieve at least a 60% success rate. Bimanual tasks exhibit longer policy completion times relative to their demonstrations than the single-arm tasks do (average difference of -8.87 s), with *scrub bowl* showing the largest gap. We attribute this to rollout-time retry behavior rather than system overhead: the policy often does not complete the task on the first attempt and instead reattempts until succeeding.

Humanoids. Beyond tabletop arms, we validate RIO on humanoids from two different vendors, Unitree G1 and Booster T1 (Figure 4), covering both upper-body manipulation and locomotion with an RL controller. On the Unitree G1, GR00T N1.5 reaches 95% success on *pick box* (Table III), with rollouts faster on average than the human demonstrations. For locomotion, RL-PPO policies trained in simulation control robots through the same RIO deployment path used for the manipulation policies above, achieving successful locomotion on each platform. Together, these runs show that RIO accommodates RL and imitation policies, vendor-specific hardware drivers, and high-frequency control required for humanoids.

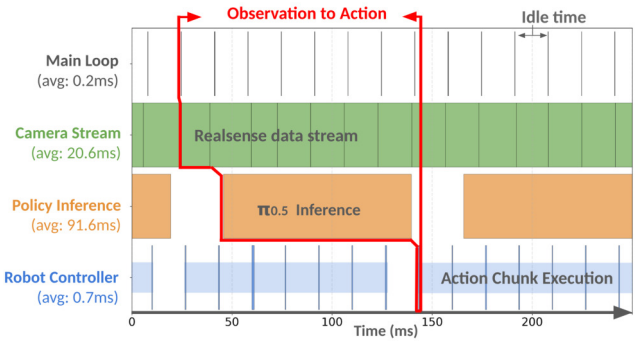

Fig. 5: **Node profiling during policy deployment.** RIO distributes blocking operations (camera streaming, policy inference, robot control) across separate nodes, keeping the main loop free for precise timekeeping.

B. Performance Analysis

We evaluate RIO at two levels: middleware latency and end-to-end latency under realistic workload conditions.

Middleware Latency. To establish an approximate lower bound on RIO's communication overhead, we measure latency across all supported middlewares. We follow the Open Messaging Benchmark by defining latency as half the median round-trip time, removing the 1st and 99th percentiles to improve robustness to outliers. We use a synthetic 2048-byte payload to reflect typical observation sizes. This benchmark isolates middleware performance because the main loop performs only ring-buffer reads and writes. Table IV shows results across five of RIO's supported backends. Zenoh and Shared Memory achieve sub-millisecond round-trip latency, making them the preferred choice for high-frequency control; Thread and ZeroRpc backends remain close at $\sim$1 ms, trading a small latency penalty for easier debugging or network flexibility,

while Portal adds further overhead (∼2 ms) in exchange for its distributed-deployment features.

Table IV: **Middleware latency.** We report latency across RIO's supported middlewares, mean ± stddev, averaged across 1,000 passes with 2048 bytes payload. Network-based backends (Zenoh, ZeroRpc, Portal) enable distributed multi-machine deployments, while local options (Shared Memory, Threads) minimize overhead for single-machine setups. This flexibility allows balancing performance, distribution, and system complexity depending on requirements.

Middleware	Latency (ms)
Zenoh	0.4287 ± 0.1344
Shared Memory	0.5413 ± 0.6166
Thread	0.9877 ± 0.3001
ZeroRpc	1.0476 ± 0.1687
Portal	1.9699 ± 0.3353

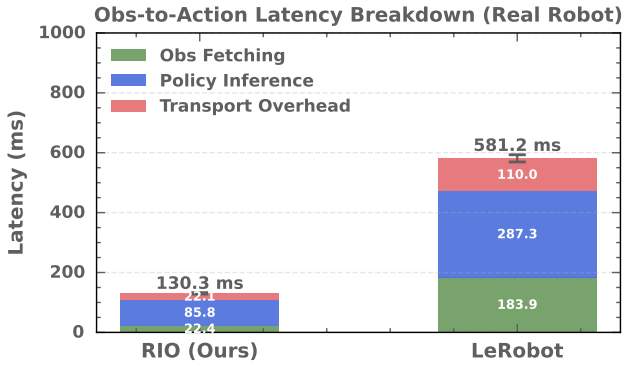

Fig. 6: **Real-world observation-action latency.** End-to-end latency during $\pi_{0.5}$ deployment with an SO-100 in the loop: RIO reaches 130.3 ms versus 581.2 ms for LeRobot.

End-to-end Profiling. To quantify performance under realistic conditions, we profile RIO during $\pi_{0.5}$ rollouts, receiving inputs from three Intel RealSense cameras (two D415s and one D405) at 640 × 480 resolution. Figure 5 shows the execution timeline across nodes: the main loop remains non-blocking, allowing for precise time-keeping, with blocking operations distributed to dedicated nodes. Asynchronous inference allows the system to preemptively request actions, maintaining continuous control despite the ∼85.8 ms model forward pass. To compare observation-to-action latency against LeRobot, a widely adopted Python-based framework, we run action chunks from the fine-tuned $\pi_{0.5}$ policy with an SO-100 in the loop under both stacks, keeping the camera setup identical to previous tasks; the robot's own read/write overhead is negligible relative to observation capture, inference, and action execution. As shown in Figure 6, RIO reaches 130.3 ms end-to-end observation-to-action latency versus 581.2 ms for LeRobot. This efficiency stems from its streamlined architecture: while LeRobot threads observations before network transmission to an asynchronous policy server, RIO directly leverages the middleware for asynchronous inference. Lower pipeline latency directly raises the effective control frequency, which is what makes the same stack viable across policy classes from the slower VLA rollouts to the high-frequency DP and RL policies reported in Table III.

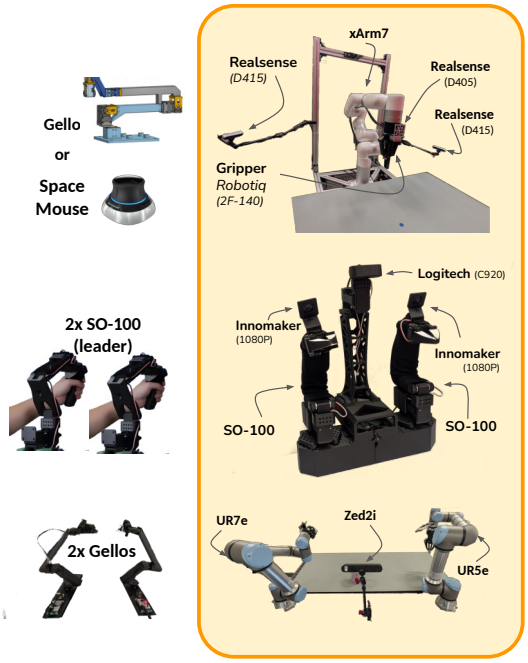

Fig. 7: **Example robot stations.** We illustrate single arm and bimanual robot stations with different cameras, controlled with different teleoperation interfaces using RIO.

V. Conclusion

The lack of flexible, reusable, accessible, performant, and consistent full-stack robot infrastructure has proven to be a critical barrier to cumulative progress and collaboration in robotics. In this work, we present RIO, a flexible real-time Robot I/O framework for cross-embodiment robot learning. RIO provides lightweight, middleware-agnostic building blocks for robot control, teleoperation, data collection, and policy deployment that can be freely combined and reconfigured across diverse hardware platforms. We validate RIO across the complete robot learning workflow, demonstrating its effectiveness on three morphologies (single-arm, bimanual, humanoid) and four robot platforms. We open source RIO, and hope that it will lower the barrier for robotics practitioners to deploy, benchmark, and iterate on modern robot learning approaches across diverse hardware configurations.

Limitations and Future Directions. In this work, we focus on single-embodiment fine-tuning. Since cross-embodiment generalization remains an active area of research, we primarily leave cross-embodiment fine-tuning of VLAs to future work. Additionally, while we demonstrate dynamic tasks through teleoperation, reliable dynamic task performance by VLAs remains an open research question that warrants further investigation. Future directions include systematic benchmarking of distribution shift across embodiments or extending support to include other robot hardware such as mobile manipulators and multi-fingered dexterous robot hands.

Acknowledgments

Guanya Shi holds concurrent appointments as an Assistant Professor at Carnegie Mellon University and as an Amazon Scholar. This paper describes work performed at Carnegie Mellon University and is not associated with Amazon.

References

[1] Amazon FAR, Pieter Abbeel, Juyue Chen, Rocky Duan, Alejandro Escontrela, Manan Gandhi, Samuel Gundry, Xiaoyu Huang, Angjoo Kanazawa, Tomasz Lewicki, Jiaman Li, Karen Liu, Clay Rosenthal, Younggyo Seo, Carlo Sferrazza, Guanya Shi, Linda Shih, Jonathan Tseng, Zhen Wu, Lujie Yang, Brent Yi, and Yuanhang Zhang. Holosoma. URL https://github.com/amazon-far/holosoma.

[2] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.

[3] Lucas Beyer, Andreas Steiner, André Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschannen, Emanuele Bugliarello, et al. Paligemma: A versatile 3b vlm for transfer. *arXiv preprint arXiv:2407.07726*, 2024.

[4] Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, et al. Gr00t n1: An open foundation model for generalist humanoid robots. *arXiv preprint arXiv:2503.14734*, 2025.

[5] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0: A Vision-Language-Action Flow Model for General Robot Control. *arXiv preprint arXiv:2410.24164*, 2024.

[6] Kevin Black, Manuel Y Galliker, and Sergey Levine. Real-Time Execution of Action Chunking Flow Policies. *arXiv preprint arXiv:2506.07339*, 2025.

[7] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.

[8] Herman Bruyninckx. Open robot control software: the OROCOS project. In *Proceedings 2001 ICRA. IEEE international conference on robotics and automation (Cat. No. 01CH37164)*, volume 3, pages 2523–2528. IEEE, 2001.

[9] Remi Cadene, Simon Alibert, Alexander Soare, Quentin Gallouedec, Adil Zouitine, Steven Palma, Pepijn Kooijmans, Michel Aractingi, Mustafa Shukor, Dana Aubakirova, Martino Russi, Francesco Capuano, Caroline Pascal, Jade Choghari, Jess Moss, and Thomas Wolf. Lerobot: State-of-the-art machine learning for real-world robotics in pytorch, 2024. URL https://github.com/huggingface/lerobot.

[10] Kaiyuan Chen, Letian Fu, David Huang, Yanxiang Zhang, Lawrence Yunliang Chen, Huang Huang, Kush Hari, Ashwin Balakrishna, Ted Xiao, Pannag R Sanketi, et al. Robo-DM: Data Management For Large Robot Datasets. *arXiv preprint arXiv:2505.15558*, 2025.

[11] Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Zixuan Li, Qiwei Liang, Xianliang Lin, Yiheng Ge, Zhenyu Gu, et al. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. *arXiv preprint arXiv:2506.18088*, 2025.

[12] Cheng Chi, Zhenjia Xu, Chuer Pan, Eric Cousineau, Benjamin Burchfiel, Siyuan Feng, Russ Tedrake, and Shuran Song. Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots. In *Robotics: Science and Systems*, 2024.

[13] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 44(10-11):1684–1704, 2025.

[14] Sudeep Dasari, Frederik Ebert, Stephen Tian, Suraj Nair, Bernadette Bucher, Karl Schmeckpeper, Siddharth Singh, Sergey Levine, and Chelsea Finn. RoboNet: Large-Scale Multi-Robot Learning. In *Conference on Robot Learning*, pages 885–897. PMLR, 2020.

[15] Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, et al. Molmo and pixmo: Open weights and open data for state-of-the-art vision-language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 91–104, 2025.

[16] Magnus Dierking, Christopher E Mower, Sarthak Das, Huang Helong, Jiacheng Qiu, Cody Reading, Wei Chen, Huidong Liang, Huang Guowei, Jan Peters, et al. Ark: An Open-source Python-based Framework for Robot Learning. *arXiv preprint arXiv:2506.21628*, 2025.

[17] Ria Doshi, Homer Rich Walke, Oier Mees, Sudeep Dasari, and Sergey Levine. Scaling Cross-Embodied Learning: One Policy for Manipulation, Navigation, Locomotion and Aviation. In *Conference on Robot Learning*, pages 496–512. PMLR, 2025.

[18] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. PaLM-E: an embodied multimodal language model. In *Proceedings of the 40th International Conference on Machine Learning*, pages 8469–8488, 2023.

[19] Dibya Ghosh, Homer Rich Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, Jianlan Luo, et al. Octo: An Open-Source Generalist Robot Policy. In *Robotics: Science and Systems*, 2024.

[20] Markus Grotz, Mohit Shridhar, Tamim Asfour, and Dieter Fox. PerAct2: Benchmarking and Learning for

Robotic Bimanual Manipulation Tasks. *arXiv preprint arXiv:2407.00278*, 2024.

[21] Tairan He, Zhengyi Luo, Wenli Xiao, Chong Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Learning human-to-humanoid real-time whole-body teleoperation. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8944–8951. IEEE, 2024.

[22] Tairan He, Zhengyi Luo, Xialin He, Wenli Xiao, Chong Zhang, Weinan Zhang, Kris M Kitani, Changliu Liu, and Guanya Shi. OmniH2O: Universal and Dexterous Human-to-Humanoid Whole-Body Teleoperation and Learning. In *Conference on Robot Learning*, pages 1516–1540. PMLR, 2025.

[23] Joey Hejna, Chethan Anand Bhateja, Yichen Jiang, Karl Pertsch, and Dorsa Sadigh. ReMix: Optimizing Data Mixtures for Large Scale Imitation Learning. In *Conference on Robot Learning*, pages 145–164. PMLR, 2025.

[24] Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, et al. $\pi_{0.5}$:a Vision-Language-Action Model with Open-World Generalization. *arXiv preprint arXiv:2504.16054*, 2025.

[25] Tobias Jülg, Pierre Krack, Seongjin Bien, Yannik Blei, Khaled Gamal, Ken Nakahara, Johannes Hechtl, Roberto Calandra, Wolfram Burgard, and Florian Walter. Robot Control Stack: A Lean Ecosystem for Robot Learning at Scale. *arXiv preprint arXiv:2509.14932*, 2025.

[26] Simar Kareer, Karl Pertsch, James Darpinian, Judy Hoffman, Danfei Xu, Sergey Levine, Chelsea Finn, and Suraj Nair. Emergence of Human to Robot Transfer in Vision-Language-Action Models. *arXiv preprint arXiv:2512.22414*, 2025.

[27] Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. DROID: A large-scale in-the-wild robot manipulation dataset. In *Robotics: Science and Systems*, 2024.

[28] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan P Foster, Pannag R Sanketi, Quan Vuong, et al. OpenVLA: An Open-Source Vision-Language-Action Model. In *Conference on Robot Learning*, pages 2679–2713. PMLR, 2025.

[29] Vikash Kumar, Rutav Shah, Gaoyue Zhou, Vincent Moens, Vittorio Caggiano, Abhishek Gupta, and Aravind Rajeswaran. Robohive: A unified framework for robot learning. *Advances in Neural Information Processing Systems*, 36:44323–44340, 2023.

[30] Obin Kwon, Sankalp Yamsani, Noboru Myers, Sean Taylor, Jooyoung Hong, Kyungseo Park, Alex Alspach, and Joohyung Kim. PAPRLE (Plug-And-Play Robotic Limb Environment): A Modular Ecosystem for Robotic Limbs. *arXiv preprint arXiv:2507.05555*, 2025.

[31] Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang Su, and Jun Zhu. RDT-1B: a Diffusion Foundation Model for Bimanual Manipulation. In *The Thirteenth International Conference on Learning Representations*, 2025.

[32] Yunchao Ma, Yizhuang Zhou, Yunhuan Yang, Tiancai Wang, and Haoqiang Fan. Running vlas at real-time speed. *arXiv preprint arXiv:2510.26742*, 2025.

[33] Steven Macenski, Tully Foote, Brian Gerkey, Chris Lalancette, and William Woodall. Robot operating system 2: Design, architecture, and uses in the wild. *Science robotics*, 7(66):eabm6074, 2022.

[34] Giorgio Metta, Paul Fitzpatrick, and Lorenzo Natale. YARP: yet another robot platform. *International Journal of Advanced Robotic Systems*, 3(1):8, 2006.

[35] Adithyavairavan Murali, Tao Chen, Kalyan Vasudev Alwala, Dhiraj Gandhi, Lerrel Pinto, Saurabh Gupta, and Abhinav Gupta. Pyrobot: An open-source robotics framework for research and benchmarking. *arXiv preprint arXiv:1906.08236*, 2019.

[36] Abby O'Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, et al. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6892–6903. IEEE, 2024.

[37] Younghyo Park and Pulkit Agrawal. Using apple vision pro to train and control robots, 2024. URL https://github.com/Improbable-AI/VisionProTeleop.

[38] Sabela Ramos, Sertan Girgin, Léonard Hussenot, Damien Vincent, Hanna Yakubovich, Daniel Toyama, Anita Gergely, Piotr Stanczyk, Raphael Marinier, Jeremiah Harmsen, et al. Rlds: an ecosystem to generate, share and use datasets in reinforcement learning. *arXiv preprint arXiv:2111.02767*, 2021.

[39] starVLA Community. StarVLA: A Lego-like Codebase for Vision-Language-Action Model Developing, January 2026. URL https://github.com/starVLA/starVLA.

[40] Gemini Robotics Team, Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montserrat Gonzalez Arenas, Travis Armstrong, Ashwin Balakrishna, Robert Baruch, Maria Bauza, Michiel Blokzijl, et al. Gemini robotics: Bringing ai into the physical world. *arXiv preprint arXiv:2503.20020*, 2025.

[41] RDT Team. Rdt2: Enabling zero-shot cross-embodiment generalization by scaling up umi data, September 2025. URL https://github.com/thu-ml/RDT2.

[42] Homer Rich Walke, Kevin Black, Tony Z Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, et al. Bridgedata v2: A dataset for robot learning at scale. In *Conference on Robot Learning*, pages 1723–1736. PMLR, 2023.

[43] J Wang, M Leonard, K Daniilidis, D Jayaraman, and ES Hu. Evaluating pi0 in the wild: Strengths, problems, and the future of generalist robot policies, 2025. URL https://penn-pal-lab.github.io/

Pi0-Experiment-in-the-Wild/.

[44] Junjie Wen, Yichen Zhu, Jinming Li, Minjie Zhu, Zhibin Tang, Kun Wu, Zhiyuan Xu, Ning Liu, Ran Cheng, Chaomin Shen, et al. Tinyvla: Towards fast, data-efficient vision-language-action models for robotic manipulation. *IEEE Robotics and Automation Letters*, 2025.

[45] Kun Wu, Chengkai Hou, Jiaming Liu, Zhengping Che, Xiaozhu Ju, Zhuqin Yang, Meng Li, Yinuo Zhao, Zhiyuan Xu, Guang Yang, et al. Robomind: Benchmark on multi-embodiment intelligence normative data for robot manipulation. *arXiv preprint arXiv:2412.13877*, 2024.

[46] Philipp Wu, Yide Shentu, Zhongke Yi, Xingyu Lin, and Pieter Abbeel. Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12156–12163. IEEE, 2024.

[47] Wei Wu, Fan Lu, Yunnan Wang, Shuai Yang, Shi Liu, Fangjing Wang, Qian Zhu, He Sun, Yong Wang, Shuailei Ma, et al. A Pragmatic VLA Foundation Model. *arXiv preprint arXiv:2601.18692*, 2026.

[48] Bin Xie, Erjin Zhou, Fan Jia, Hao Shi, Haoqiang Fan, Haowei Zhang, Hebei Li, Jianjian Sun, Jie Bin, Junwen Huang, Kai Liu, Kaixin Liu, Kefan Gu, Lin Sun, Meng Zhang, Peilong Han, Ruitao Hao, Ruitao Zhang, Saike Huang, Songhan Xie, Tiancai Wang, Tianle Liu, Wenbin Tang, Wenqi Zhu, Yang Chen, Yingfei Liu, Yizhuang Zhou, Yu Liu, Yucheng Zhao, Yunchao Ma, Yunfei Wei, Yuxiang Chen, Ze Chen, Zeming Li, Zhao Wu, Ziheng Zhang, Ziming Liu, Ziwei Yan, and Ziyu Zhang. Dexbotic: Open-Source Vision-Language-Action Toolbox. *arXiv preprint arXiv:2510.23511*, 2025.

[49] Mengda Xu, Han Zhang, Yifan Hou, Zhenjia Xu, Linxi Fan, Manuela Veloso, and Shuran Song. DexUMI: Using Human Hand as the Universal Manipulation Interface for Dexterous Manipulation. *arXiv preprint arXiv:2505.21864*, 2025.

[50] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11975–11986, 2023.

[51] Tony Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. In *Robotics: Science and Systems XIX*, 2023.

[52] Jinliang Zheng, Jianxiong Li, Zhihao Wang, Dongxiu Liu, Xirui Kang, Yuchun Feng, Yinan Zheng, Jiayin Zou, Yilun Chen, Jia Zeng, et al. X-vla: Soft-prompted transformer as scalable cross-embodiment vision-language-action model. *arXiv preprint arXiv:2510.10274*, 2025.

[53] Zhengbang Zhu, Minghuan Liu, Xiaoshen Han, and Zhengshen Zhang. Maniunicon: A unified control interface for robotic manipulation, 2025. URL https://github.com/Universal-Control/ManiUniCon.