

Being-H0.5: Scaling Human-Centric Robot Learning for Cross-Embodiment Generalization

BeingBeyond Team

<https://research.beingbeyond.com/being-h05>

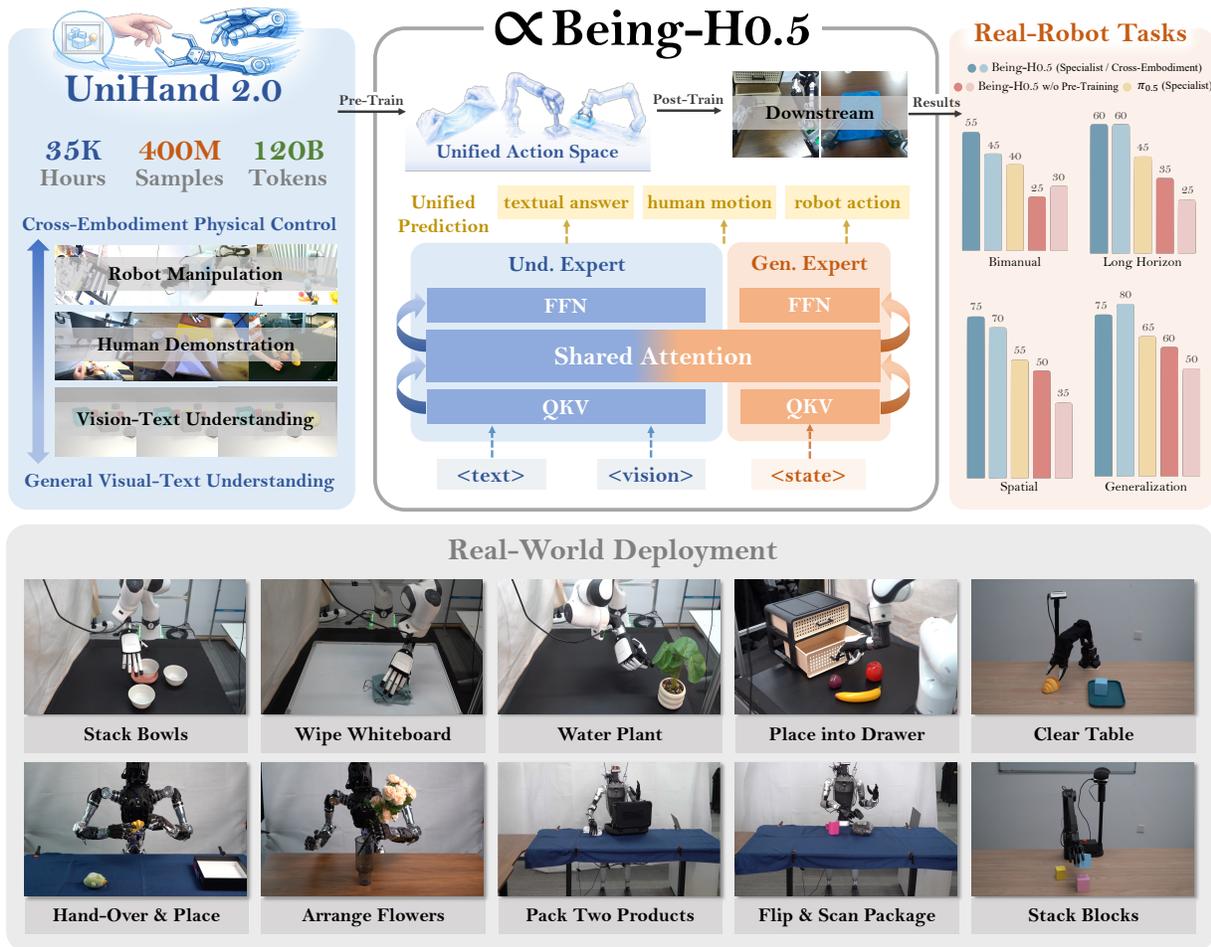


Figure 1: **Being-H0.5 at a Glance.** We scale **human-centric robot learning** with Being-H0.5 toward cross-embodiment generalization. We introduce UniHand-2.0, a large-scale corpus exceeding **35,000 hours** that spans both cross-Embodiment physical control and general visual-text understanding. Building on this data, we unify human hand motion and diverse robot embodiments with a **Unified Action Space**, and train all heterogeneous supervision through **unified sequence modeling** under a single framework. This yields a single foundation model that can *perceive, describe, and act* within one framework, enabling robust cross-embodiment generalization and real-world deployment across diverse robots and tasks. We empirically deploy a *single* checkpoint of Being-H0.5 to control PND Adam-U, Franka+Inspire Hand, Unitree G1+LinkerBot, BeingBeyond D1, and LeRobot SO-101 to accomplish diverse tasks.

Abstract

We introduce Being-H0.5, a foundational Vision-Language-Action (VLA) model designed for robust cross-embodiment generalization across diverse robotic platforms. While existing VLAs often struggle with morphological heterogeneity and data scarcity, we propose a human-centric learning paradigm that treats human interaction traces as a universal “mother tongue” for physical interaction. To support this, we present UniHand-2.0, the largest embodied pre-training recipe to date, comprising over 35,000 hours of multimodal data across 30 distinct robotic embodiments. Our approach introduces a Unified Action Space that maps heterogeneous robot controls into semantically aligned slots, enabling low-resource robots to bootstrap skills from human data and high-resource platforms. Built upon this human-centric foundation, we design a unified sequential modeling and multi-task pre-training paradigm to bridge human demonstrations and robotic execution. Architecturally, Being-H0.5 utilizes a Mixture-of-Transformers design featuring a novel Mixture-of-Flow (MoF) framework to decouple shared motor primitives from specialized embodiment-specific experts. Finally, to make cross-embodiment policies stable in the real world, we introduce Manifold-Preserving Gating for robustness under sensory shift and Universal Async Chunking to universalize chunked control across embodiments with different latency and control profiles. We empirically demonstrate that Being-H0.5 achieves state-of-the-art results on simulated benchmarks such as LIBERO (98.9%) and RoboCasa (53.9%), meanwhile performing strong cross-embodiment capabilities on five robotic platforms.

Date: Jan 20, 2026

Contents

1	Introduction	4
2	Related Work	6
3	UniHand-2.0: A Recipe for Large-Scale Human-Centric Learning	7
3.1	Human Demonstration Data	10
3.2	Robot Manipulation Data	11
3.3	Visual-Text Understanding Data	11
4	UniCraftor: A System for Portable, Extensible, and Affordable Data Collection	13
5	Being-H0.5: A Foundational VLA Unifying Cross-Embodiment Control	14
5.1	Model Architecture	14
5.1.1	Unified State-Action Space	15
5.1.2	Mixture of Flow	16
5.2	Pre-Training: Human-Centric Robot Learning	17
5.2.1	Unified Sequence Modeling	17
5.2.2	Human-Centric Multi-Task Objective	17
5.2.3	Hybrid Human Motion Representation	18
5.3	Post-Training: Towards Cross-Embodiment Adaptation	19
5.3.1	Embodiment-Specific Adaptation	19
5.3.2	Manifold-Preserving Gating	20

5.3.3	Universal Async Chunking	21
6	Infrastructure: Real-Time Cross-Embodiment Deployment	22
6.1	Manifold-Preserving Refinement	22
6.2	Universal Async Chunking Protocol	23
6.3	Dual-Thread Deployment Architecture	24
7	Experiment	24
7.1	Comparison across Real Robots	24
7.1.1	Experimental Setup	24
7.1.2	Result Analysis	26
7.2	Comparison on Simulation Benchmarks	28
7.2.1	Results on LIBERO	29
7.2.2	Results on RoboCasa	29
7.3	Ablation Study	30
7.3.1	How Do Human-Centric Learning Benefit Downstream Adaptation?	30
7.3.2	How Does Masked Motion Token Prediction Benefit Capturing Behavior Priors?	32
7.3.3	Real-Time Efficiency: MPG and UAC	33
8	Conclusion	34
A	Contributions	43
B	Full Rollouts	44

1 Introduction

True intelligence is characterized by the fluid expression of knowledge across diverse media. In Natural Language Processing (NLP) [1, 2], multilingual pre-training [3–5] has transcended linguistic boundaries by distilling a “universal semantics”, facilitating seamless knowledge transfer from high-resource to low-resource languages. Today, general-purpose robotics stands at a similar crossroads: the motor spaces of heterogeneous robots are, in essence, distinct physical languages. Consequently, a primary challenge in scaling robotic intelligence is cross-embodiment translation:

How can robotic intelligence master a new platform with limited data, much like a human learning a minor language?

While Vision-Language-Action (VLA) models show promise toward a general robotic policy, they are hindered by a profound “physical gap”. Current VLAs function as monolingual speakers that are highly proficient on specific hardware but effectively incompatible when deployed on different morphologies. This fragmentation is well-known exacerbated by the scarcity of robot-specific data [6, 7]. Unlike the trillions of tokens available in NLP, robotics lacks large-scale demonstration corpora for most individual platforms. For diffusion-based VLAs [8, 9], this is particularly acute: a model pre-trained on simpler hardware (e.g., parallel grippers) suffers from severe distribution shift especially when encountering the high-dimensional vector fields of complex entities (e.g., dexterous hands). This mismatch causes inference trajectories to “drift” from the valid motion manifold, resulting in unstable behaviors.

To overcome these hurdles, we propose a paradigm shift toward embodiment-agnostic VLAs via two core concepts:

- **Human-Centric Learning:** Just as different human languages share an underlying grammar (e.g., syntax and logic), different robots follow shared physical laws. We treat the vast library of human interaction acts as the “mother tongue”—the universal lingua franca of the physical world. Human movements provide dense semantic priors, capturing the logic of causal interaction and the physics of contact that remain invariant across all kinematic “dialects”. By grounding learning in human-centric data, we provide a source of physical common sense that enables VLAs to adapt rapidly to new hardware.
- **Unified Action Space:** By representing diverse robot actions as tokens in a shared physical vocabulary, we align disparate robot morphologies into a unified latent space, allowing low-resource, complex robots to bootstrap motor skills from data-rich platforms and human demonstrations.

Building upon these motivations, we introduce **Being-H0.5**, a foundational VLA towards cross-embodiment generalization. Our model is anchored by UniHand-2.0, an expansive corpus comprising **35,000+ hours of data and 120 billion tokens, totaling more than 400 million samples**. As an extension of [10], UniHand-2.0 integrates 16,000 hours of egocentric human video, 14,000 hours of robot manipulation, and 5,000 hours of general visual-language understanding data. Notably, UniHand-2.0 spans **30 distinct robotic platforms**, from tabletop arms to legged humanoids. Much like a multilingual LLM, Being-H0.5 is pre-trained on this “polyglot” dataset to internalize a robust, shared representation of physical manipulation. Furthermore, UniHand-2.0 includes large-scale vision-text corpora to ensure the model retains the high-level reasoning, strategic planning, and instruction-following capabilities of its VLM backbone. Compared to its predecessor, UniHand-2.0 represents a 200× increase in scale. To our knowledge, this represents the largest embodied pre-training recipe to date, featuring the most extensive use of human video and the most diverse collection of robotic embodiments. By centralizing pre-training on massive human data alongside dozens of robotic datasets, we endow Being-H0.5 with strong generalization across heterogeneous platforms, even when target-specific data is highly limited.

Furthermore, we developed a human-centric data collection system named **UniCraftor** to address the limitations of existing datasets and scale our data acquisition. Through meticulous design, our system integrates depth information, keyframe events, and precise camera extrinsics—modalities that are often absent but important in current benchmarks. The system is intentionally modular and extensible, specifically designed to accommodate diverse and flexible viewpoints, as well as future tactile information. Leveraging

this system, we have curated a comprehensive dataset exceeding 200 hours, covering 43 distinct tasks.

In addition to data scarcity, scaling generalist robotics is constrained by embodiment heterogeneity—the physical variance between human morphology and diverse robotic platforms with disparate kinematics, actuation limits, and control frequencies. Naively aggregating these datasets introduces action-space interference, where conflicting control signals generate significant noise. This often results in negative transfer, where cross-platform training degrades performance rather than fostering synergy. Some works [9] typically mitigate this by employing platform-specific action heads. However, this strategy is inherently limited as it bypasses the underlying structural misalignments. Without reconciling these configurations, models fail to develop proprioceptive-aware reasoning or distill transferable physical priors across domains [11]. Consequently, existing policies remain either embodiment-specific specialists or “shallow generalists” that struggle with high-DoF tasks, such as dexterous or bimanual manipulation. To bridge this gap, we propose a **Unified Action Space** that maps human trajectories and heterogeneous robot controls into semantically aligned slots. This acts as a “universal grammar” for disparate hardware, effectively decoupling functional intent (e.g., the semantic objective of a grasp) from the mechanical articulation. By doing so, our model internalizes the underlying physics of interaction rather than just embodiment-specific commands.

Building on this unified action space, we take a further step towards scalable pre-training by **casting all heterogeneous supervision into a unified sequence modeling problem**. Rather than maintaining separate pipelines for human demonstrations, robot trajectories, and vision-language corpora, we serialize them into a single multimodal token stream, where *vision* and *text* provide contextual grounding and unified *state/action* tokens carry physically meaningful interaction signals. This lets us optimize all data with various losses applied to the relevant supervised segments. Concretely, text-centric corpora (e.g., VQA and motion description) contribute a standard next-token prediction loss over **text**, while human and robot corpora supervise **action** prediction in the unified space, with human behavior supplying dense, transferable behavioral priors and robot trajectories anchoring high-fidelity kinematic control. By unifying both the representation and the optimization, pre-training becomes a coherent curriculum over one token stream, enabling the model to perceive, describe, and act within the same framework.

Architecturally, Being-H0.5 adopts a Mixture-of-Transformers (MoT) design that disentangles high-level multimodal reasoning from low-level execution experts. While a unified action space effectively represents diverse embodiments and prevents representational conflicts during pre-training, the limited capacity of conventional action experts remains a critical bottleneck. Inspired by the principle of Mixture-of-Experts (MoE) [12], we introduce a scalable architectural framework called **Mixture of Flow**, which decouples the action module into foundational experts with shared dynamic knowledge and specialized experts that utilize embodiment-aware task routing. To ensure practical scalability, we introduce two complementary ingredients that maintain flow-based action generation both *stable* and *transferable*. **1) Manifold-Preserving Gating** encourages the model to rely on reliable context and fall back to robust priors when perception is ambiguous, preventing unstable corrections from being amplified through iterative refinement. **2) Universal Async Chunking** extends real-time chunked control to the *cross-embodiment* setting. Rather than baking in a single robot’s timing and control assumptions, we train a unified policy to remain consistent across heterogeneous platforms with varying actuation frequencies and delay profiles, enabling a single checkpoint to operate fluently on all embodiments. Together, these innovations allow Being-H0.5 to remain fluent across varying sensory conditions and heterogeneous robot “dialects”.

We evaluate Being-H0.5 and the efficacy of human-centric learning through extensive benchmarking across real-world and simulated environments. Being-H0.5 achieves state-of-the-art (SoTA) results on LIBERO (98.9%) and RoboCasa (53.9%) using only low-resolution RGB input without auxiliary modalities. Furthermore, our model significantly outperforms existing VLAs, such as $\pi 0.5$, across five physically distinct embodiments, which demonstrates superior cross-embodiment generalization regardless of structural complexity. Notably, our real-robot experiments reveal an unexpected emergent embodiment-level zero-shot transfer signal: A single Being-H0.5 generalist checkpoint, trained jointly across embodiments under a unified action interface, achieves a non-zero success rate on unseen tasks —embodiment pairs without any data on the target robot. We believe this observation sheds light on a practical scaling direction for cross-embodiment VLAs and paves the way for improving emergent transfer through increasing diverse post-training datasets.

To support the community and ensure full reproducibility, we provide a comprehensive open-source release encompassing our model weights, training pipeline, and simulation scripts. Furthermore, we will include our real-world deployment infrastructure and a 1,000 GPU-hour pre-training recipe. Our primary contributions are summarized as follows:

- **Largest Training Recipe.** We introduce UniHand 2.0, the most extensive embodied VLA dataset to date, comprising 400M samples across 35,000 hours (16,000 human, 14,000 robot) and 30 embodiments. To facilitate future scaling, we also provide a portable, plug-and-play human data collection system.
- **Unified Training Paradigm.** For the first time, we unify human hand motion and diverse robotic controls into a single action space for cross-embodiment generalization. This is coupled with a unified sequence modeling paradigm that allows Being-H0.5 to *perceive, describe, and act* within a single sequence, enabling scalable human-centric pre-training on heterogeneous corpora.
- **Architectural Innovations.** To maximize model capacity and cross-domain transfer, we introduce several novel designs, including mixture-of-flow, manifold-preserving gating, and universal async chunking. These components address the inherent bottlenecks in scaling flow-based action generation.
- **Real-Time Infrastructure.** We develop an efficient inference infrastructure that enables low-latency, real-time control, providing the computational throughput necessary for high-DoF, complex platforms.
- **State-of-the-Art Results.** Evaluated on five distinct real-world platforms and major simulated benchmarks, Being-H0.5 sets new SoTA records, including 98.9% on LIBERO and 53.9% on RoboCasa, while demonstrating emergent zero-shot transfer to unseen robot morphologies.

2 Related Work

Vision-Language-Action Models. Recent advances in robotic manipulation [13–16] have shifted from narrow, single-task specialists toward generalist models trained on diverse, large-scale datasets spanning multiple scenes and embodiments. A prominent paradigm in this shift is the Vision–Language–Action models (VLAs) [6, 17–20], which bridges internet-scale perception and physical execution by fine-tuning pre-trained vision-language models (VLMs) [21–27] on robotic control data. While many VLAs share architectural backbones, they diverge significantly in action-head design. Early initiatives typically adopt autoregressive approaches, discretizing continuous actions into tokens to align with VLM training objective [6, 28]. While this facilitates seamless knowledge transfer, it often introduces high inference latency and limits precision in high-degree-of-freedom tasks. To mitigate these bottlenecks, recent frameworks [19, 29–31] such as π_0 [8] and GR00T-N1 [9] integrate diffusion-based architectures [32] (*e.g.*, flow matching [33]) to generate high-frequency action chunks, which has been widely adopted by follow-up works [34–36]. Despite these improvements, bridging the gap between high-level reasoning and granular execution remains a challenge. Some approaches [37–39] employ Chain-of-Thought (CoT) reasoning [40] to decompose long-horizon tasks such as OneTwoVLA’s adaptive “thinking” mode [38] or RAD’s language-based guides derived from human videos [39], yet textual planning often lacks the spatial precision required for physical interaction. Consequently, recent research has explored non-textual intermediate representations, such as bounding boxes [41], dense correspondence fields [42], or 3D points [43]. A notable example is MolmoAct [44], which transfers observations into depth-aware perception tokens to generate mid-level spatial plans as editable trajectory traces. Despite these advances, most VLAs remain bottlenecked by their reliance on lab-collected teleoperation data, which lacks the environmental variety of the real world. We address this by leveraging millions of egocentric videos to capture the behavioral richness necessary for robust generalization.

Heterogeneous Pre-Training Datasets. The efficacy of VLAs is fundamentally tied to the scale and diversity of pre-training data. Large-scale initiatives like Open X-Embodiment dataset [7], along with datasets such as Droid [45] and BridgeData [46], have aggregated thousands of hours of robot demonstrations across diverse platforms. Community-driven benchmarks, such as [47], have further expanded this landscape by leveraging accessible, portable hardware. To address the complexities of high-dimensional bimanual coordination and dexterous manipulation, recent datasets like AgiBot World [48], RoboMIND/2.0 [49, 50], and RoboCOIN [51] have emerged to fill critical gaps in fine-grained motor control, while Open Galaxea [52] targets the unique requirements of mobile manipulation. Despite these advances, most teleoperated data remains

confined to specific lab settings, creating a bottleneck for generalizable policy learning. While synthetic data [53, 54] offers scalability, it suffers from the persistent sim-to-real gap. Crucially, as these datasets span heterogeneous platforms with significant structural distinctions, existing VLAs lack a unified mechanism to integrate them effectively. Discrepancies in data formats, annotation granularity, and embodiment structures pose significant challenges for large-scale unified modeling. In this paper, we propose a unified state-action space to bridge these structural gaps.

Human-Centric Learning. The scarcity of real robot data, driven by the high cost of manual teleoperation, has prompted a shift toward human-centric learning. To lower the barrier for data collection, portable physical interfaces like UMI [55–57] allow for the direct transfer of “in-the-wild” human skills to deployable robot policies. Recent efforts have already scaled UMI-collected data to over 10,000 hours [58]. An even more scalable alternative is learning directly from abundant human videos. Compared to robot-collected data, human video corpora span a vastly broader distribution of environments, object configurations, and long-horizon task structures. This domain has evolved from traditional vision tasks such as recognition and grounding [59, 60], to large-scale egocentric benchmarks like Ego4D [61], Ego-Exo4D [62], EPIC-KITCHENS [63], and EgoDex [64]. While these provide massive interaction priors, the fundamental embodiment gap and the lack of task-aligned labels make their direct application to VLA modeling non-trivial.

Early efforts of human-centric learning attempted to exploit the visual diversity of human videos through representation learning [65, 66]. However, the resulting features are largely implicit and lack action-level supervision, which limits their utility for precise downstream control. Consequently, recent research has shifted toward more structured manners, which can be categorized into several trajectories. First, some works [67–69] explore latent action as intent abstraction by encoding temporal changes into discrete codes or continuous embeddings to represent intent. LAPA [70] employs VQ-VAE tokens derived from frame pairs to pre-train a VLM, which is subsequently fine-tuned with an action head. UniVLA [71] replaces pixel reconstruction with DINOv2 [72] feature prediction and introduces a two-stage codebook scheme to isolate task-centric latents, enabling autoregressive planning with lightweight latent-to-action decoding. In addition to latent representation, approaches like GR-1/2 [73, 74] and Gen2Act [75] model physical interactions by explicitly predicting future environment states or video frames, treating video generation as a proxy for understanding dynamics. A third line of work [76–80] leverages explicit 2D geometric cues, such as point trajectories, keypoints, and bounding boxes [81–83] to bridge the gap between video pixels and physical movement.

Moving beyond indirect proxies, some recent studies focus on recovering actual human actions from demonstrations to provide direct supervision for policy pre-training [84–89]. For instance, EgoVLA [90] constructs a shared action space using MANO hand-model parameters, retargeting 3D wrist and hand poses from egocentric videos to robot commands via inverse kinematics. Being-H0 [10] scales this paradigm by introducing a motion tokenizer that discretizes continuous MANO trajectories into tokens for large-scale instruction tuning. Despite these advancements, existing methods are often constrained by the difficulty of scaling high-quality action annotations and the lack of a framework that can seamlessly integrate heterogeneous human demonstrations with robot-collected data. In this work, we extend the paradigm established by Being-H0 to curate the largest human-centric pre-training corpus to date, supported by a unified paradigm for multi-source integration.

3 UniHand-2.0: A Recipe for Large-Scale Human-Centric Learning

Vision-Language-Action models (VLAs) require vast quantities of robotic interaction data to acquire actionable knowledge and world commonsense. However, the current landscape of robotic data is hindered by several fundamental limitations. Our work systematically addresses the following critical challenges:

- 1) **Limited Data Scale and Diversity.** Most existing VLAs [6, 91] are constrained by the insufficient scale and diversity of their pre-training corpora. Many rely on a narrow scope of data, such as the Open X-Embodiment dataset [7] which offers only six major subsets with restricted variety after rigorous filtering, while Agibot World [48] contains merely ~ 200 hours of tabletop manipulation data and lacks essential third-person camera views. This scarcity impedes generalization to novel tasks and dynamic environments. To overcome this, we curate a comprehensive robotic manipulation dataset that aggregates the vast majority of available robot data.



Figure 2: **Overview of UniHand 2.0.** UniHand 2.0 is our large-scale pre-training recipe for human-centric robot learning, comprising 35K hours of multimodal data from three complementary sources. 1) **Human Demonstration** with diverse scenes, multitask for motion alignment, and multi-granularity semantics. 2) **Robot Manipulation** spanning 30+ embodiments with multiple observation views and heterogeneous control signals. 3) **Vision-Text Understanding** covering general VQA, 2D spatial grounding & affordance, and task planning & reasoning.

- 2) **Restricted Embodiment Variety.** Beyond scale, existing datasets typically feature a limited suite of embodiments, often restricted to a single robot type (*e.g.*, Agibot World [48] and LET [92]). While some cross-embodiment datasets exist [51], unifying these sources for pre-training remains difficult due to significant structural discrepancies. Consequently, few VLAs successfully incorporate broad robot morphologies. For example, the π -series [8, 34] encompasses only 10 robot types, predominantly homogeneous bimanual platforms. In contrast, UniHand-2.0 incorporates data from **30 distinct embodiments**, spanning single/dual-arm, portable, half-humanoid, and legged humanoid robots. This represents the most extensive array of embodiments reported in VLA literature to date. We achieve this by projecting these heterogeneous sources into a unified state-action space to ensure training stability.
- 3) **Scarcity of Dexterous Hand Data.** Despite advances in general data collection, data involving dexterous hands remains exceptionally rare, comprising less than 5% of existing corpora. This inadequacy stems from the high cost of hardware and the low throughput of dexterous teleoperation. To mitigate this, we leverage large-scale human motion data similar to [10] as a scalable proxy, capitalizing on the relative ease of capturing human hand interactions in natural settings.
- 4) **Imbalance between Visual and Language Information.** Prior VLAs [8, 93] generally rely on robot-only data for pre-training, leading to a severe modality imbalance where the ratio of text to visual tokens can reach 1:3,000. This disparity causes models to lose substantial textual reasoning capabilities, which is critical for long-horizon task execution. While some works incorporate vision-language data [34, 54] or interleaved multimodal samples [94], we systematically incorporate massive

multimodal data during pre-training to ensure Being-H0.5 excels at both atomic action execution and long-horizon task planning & spatial reasoning.

Here, we introduce **UniHand-2.0**, a significantly expanded dataset built upon its predecessor, UniHand-1.0 [10]. The dataset (Figure 2) contains over **400 million** samples extracted from **35,000 hours** of multimodal data, totaling over 120B training tokens. These samples span three key domains: egocentric human motion, robot manipulation, and visual-language understanding. As illustrated in Figure 3, UniHand-2.0 is the largest human-centric VLA pre-training corpus to our knowledge. Note that the usage of the Open X-Embodimen dataset is ambiguous for different works. We thereby restrict our embodiment counting with >10 hours of data. UniHand-2.0 leverages low-cost human data as primary pre-training material, treating the human hand as a universal template for all end-effectors, imbuing models with foundational interaction knowledge and physics understanding. We advance this paradigm by extracting 134 million human data samples from **16,000 hours** of egocentric video, which is a 100× increase over UniHand-1.0. Furthermore, UniHand-2.0 incorporates over **14,000 hours** of diverse robotic data across **30 embodiment types** (*e.g.*, Franka, AgiBot-GR1, Unitree-G1, SO101), endowing the model with robust cross-embodiment generalization.

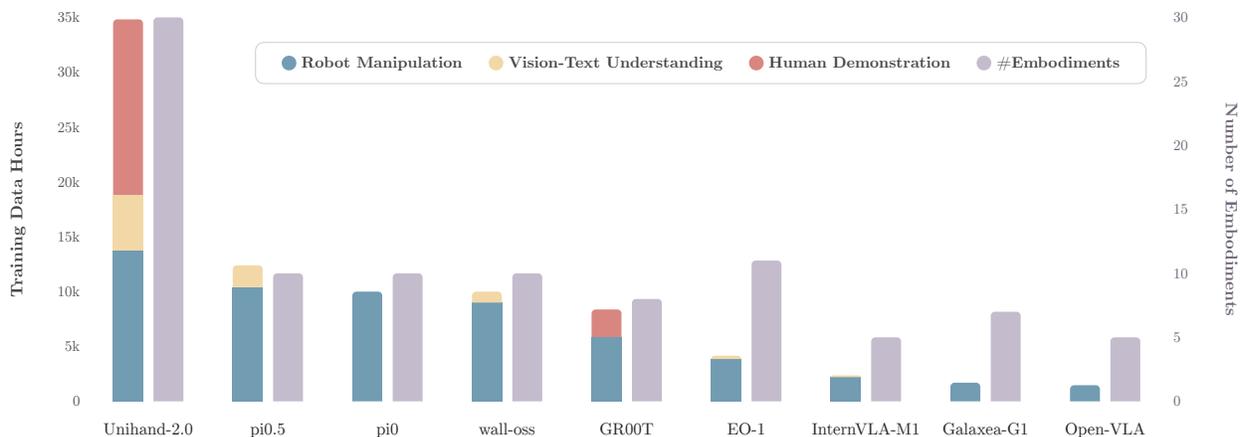


Figure 3: **Comparison of training scale and embodiment diversity.** Stacked bars represent hours of training data (left axis); hatched bars represent embodiment counts (right axis). UniHand-2.0 represents the largest and most diverse VLA pre-training recipe to date, totaling 35,000 hours of multimodal data. This includes 16,000 hours of human data, 14,000 hours of robot data across 30 embodiments, and 5,000 equivalent hours of VLM data.

We posit that such embodiment diversity is a prerequisite for existing diffusion-based VLAs. From a manifold learning perspective, simple embodiments (*e.g.*, parallel grippers) operate on a low-dimensional, smooth action manifold. In contrast, dexterous robots inhabit complex, high-dimensional spaces where the manifold structure is often non-linear and fragmented. The action distribution of these complex entities differs fundamentally from that of simple ones. For instance, the binary “open/close” command of a gripper is a simple scalar, whereas the “precision pinch” of a dexterous hand requires a high-dimensional, coordinated vector. This disparity leads to a severe target distribution shift during adaptation. Furthermore, for a diffusion framework, it must infer a continuous vector field to define the probabilistic evolution toward the next action state. When a model pretrained only on simple robots encounters the unseen state space of a complex entity, its vector field predictions suffer from accumulated errors. These errors cause generated trajectories to “drift” and rapidly deviate from the valid robot motion manifold, resulting in unstable or physically infeasible behaviors.

To maintain a critical balance between modalities, we integrate visual-language understanding data at a comparable scale, preserving the model’s reasoning and instruction-following abilities. UniHand-2.0 thus constitutes the largest pre-training dataset organized under the human-centric learning concept. The following sections detail each component, and statistics are visualized in Figure 4.

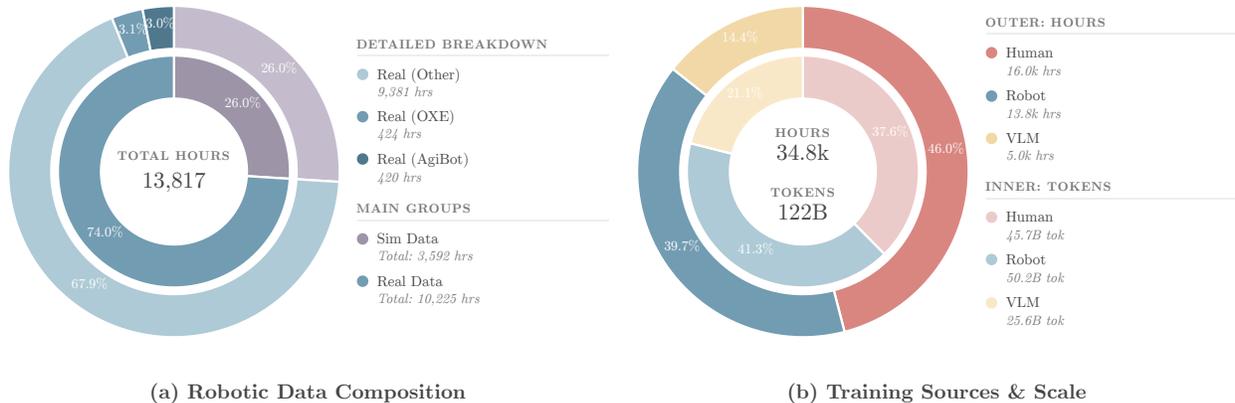


Figure 4: **Statistics of UniHand-2.0. (Left) Ratio of simulation vs. real-world data.** We maintain a balanced ratio with simulation data at 26%, while the widely used Open X-Embodiment (OXE) and AgiBot World datasets account for 3.1% and 3.0%, respectively. **(Right) Training sources and scale.** Human, robot, and visual-language (“VLM”) data consist of 16K hours (25.6B tokens), 14K hours (45.7B tokens), and 5K equivalent hours (50.2B tokens), respectively. These sources are curated to maintain a comparable scale for balanced pretraining.

3.1 Human Demonstration Data

Human demonstration data has emerged as a powerful and scalable foundation for pre-training VLAs, as evidenced by recent advancements [9, 10, 88, 90, 95]. Beyond serving as a substitute for scarce robotic interaction data, human-centric videos offer a direct window into the fundamental ways humans perceive, reason, and manipulate within the physical world. These corpora capture a vast diversity of everyday activities—ranging from fine-grained hand-object interactions to complex tool use and rich contextual scenes—that would be prohibitively expensive to replicate at scale using physical robots. This diversity establishes robust, generalizable behavioral priors and provides a broad coverage of motion patterns and affordance that are often absent in scripted robotic collection or simulation environments.

Following the paradigm established by Being-H0, we treat human hand motion as a generalized manipulator template. This paradigm physically grounds visual observations in real-world interactions, effectively transforming passive video frames into action-relevant supervision. With hand motion as a grounding anchor, our model leverages egocentric human videos for robotic policy learning through three primary mechanisms: **1) Action-Relevant Visual Context:** Hand motion functions as a mechanism for spatial attention, localizing where and what is relevant for a given action. This expands behavior-grounded visual context significantly beyond what is available in standard VLM corpora. **2) Rich Manipulation Structure:** Hand-centric interaction traces expose fine-grained contact patterns, dexterous control strategies, and diverse object affordance, capturing manipulation dynamics that are currently missing from most robot datasets. **3) Diverse Textual Context for Intent Grounding:** By aligning linguistic context with hand-centric interaction traces, the model can ground human language in physical outcomes and motion, enabling intent-level learning rather than superficial visual pattern matching.

Building upon Being-H0, we significantly scale and diversify our human-video corpus, extending its scope beyond lab-recorded, manipulation-specific settings to include both open-source repositories and self-curated data. To render these videos directly usable for robot learning, we extracted a unified hand-motion representation from diverse videos sources. This is further supplemented by a high-fidelity in-house dataset to provide the signal precision often lacking in open-web videos.

Curation of Heterogeneous Human Videos. In addition to the UniHand-1.0 dataset, we integrate an extensive collection of in-the-wild egocentric videos from large-scale public repositories, including Ego4D [61], EPIC-KITCHENS [96], Egocentric-10K [97], etc. This extension drastically increases environmental variety and embodied behavior, encompassing household chores, cooking, and industrial operations (see Figure 2). To

align these heterogeneous sources into a unified interface, we employ HaWoR [98] to estimate hand poses via MANO parameters [99] and camera extrinsics. Furthermore, we enrich the dataset with fine-grained semantic supervision using Gemini-2.5 [100], generating dual-level annotations: detailed per-second instructions and holistic, 10-second chunk-level intent descriptions.

We structure the training data into three interconnected task families: **motion generation** (the primary objective mapping vision/language to kinematic action), **motion description**, and **motion continuation**. To ensure high data fidelity and minimize spurious correlations, we implement a rigorous four-stage post-processing pipeline: **1) Language Augmentation:** We employ LLMs to paraphrase and diversify all template-formatted instructional texts, preventing the model from overfitting to rigid linguistic patterns. **2) Motion-Quality Filtering:** We ensure data reliability by filtering noisy HaWoR estimates based on detection confidence and DBA error, while removing high-frequency jitter or discontinuities in wrist-space. **3) Manipulation Relevance Filtering:** Using Gemini-assisted semantic screening, we exclude segments dominated by non-manipulative actions, such as pure locomotion. **4) Handedness Debiasing:** We apply left-right spatial mirroring to all samples to mitigate common right-handed bias, thereby promoting ambidextrous generalization in the learned policy.

3.2 Robot Manipulation Data

Robot control data serves as the cornerstone of policy learning, providing a critical pathway toward achieving both broad embodiment generalization and dexterous manipulation proficiency. To establish a comprehensive resource for this purpose, we curate a large-scale, heterogeneous collection of robot demonstrations. This dataset aggregates about 14,000 hours of interaction data (approximately 1.5 billion frames), making it one of the most extensive repositories of robotic behavior currently available. The collection systematically integrates a wide array of datasets, including OpenX-Embodiment [7], AgiBot-World [48], SO100-Community [47], InternData-M1 [54], RoboMIND [49], RoboCOIN [51], LET [92], etc. From this aggregate pool, we perform deduplication and downsample frames to 30% to maximize data diversity while minimizing redundancy.

As a result, our processed dataset encompasses **30 distinct robot embodiments**, such as Franka arms, Split ALOHA, and Agibot G1, representing a significant breadth of hardware morphologies and control interfaces. To further scale task coverage, we incorporate procedurally generated simulated samples [54] and inpainting-augmented trajectories [101]. While simulated data effectively boosts performance on standard benchmarks, excessive reliance can exacerbate the sim-to-real gap. To mitigate this, we strictly cap the proportion of simulated data in the pretraining mixture (26% of the total corpus, as shown in Figure 4), ensuring that real-world signals remain the dominant learning signal. Finally, we contribute novel demonstrations for hardware platforms such as the PND Adam-U and BeingBeyond D1 to improve data diversity and balance. Additional details regarding the robot data distribution are provided in Table 1.

3.3 Visual-Text Understanding Data

A pronounced imbalance often exists between the visual and textual components of standard robotic training corpora. While robotic interaction data provides dense visual signals, the accompanying linguistic supervision is frequently sparse and simplistic. This disparity creates a fundamental asymmetry in multimodal supervision, risking a degradation in the model’s capacity for complex textual reasoning, instruction following, and high-level planning. For instance, within our robotic interaction corpus alone, visual tokens number 45.7 billion, whereas textual tokens amount to merely 30 million—a staggering ratio of approximately 1,000:1. Left unaddressed, this discrepancy may cause the model to overfit to visual patterns at the expense of linguistic intelligence, potentially reducing a robotic agent to a visual-motor reflex system rather than a reasoning agent. To restore modality balance, we curate a comprehensive visual-language understanding corpus specifically designed to preserve and enhance the model’s semantic reasoning capabilities. We strategically categorize this collection into three pillars, ensuring coverage from general world understanding to precise, actionable robotic reasoning:

- 1) General Visual-Language QA.** To maintain robust general-purpose vision-language alignment, we incorporate established large-scale instruction-tuning datasets. This includes image-based benchmarks such as LLaVA-v1.5 [102], LLaVA-OneVision [103] and FineVision [104], as well as video-centric datasets

Table 1: **Statistical overview of robot manipulation data in UniHand-2.0.** The corpus aggregates demonstrations across **30 distinct embodiments**, totaling **14,000 hours** of interaction. The dataset is characterized by high diversity in camera views, kinematic structures, and operational environments, categorized by end-effector type—“Grp” (parallel gripper) vs. “Dex” (dexterous hand), and data source—“Real” (physical world) vs. “Sim” (simulation).

Robot Type	Camera Views	EEF Type	Source	Hours
<i>Single-Arm Robots</i>				
Franka	ego, 3rd×2, wrist	Grp/Dex	Real/Sim	2196.4
Google Robot	ego, 3rd	Grp	Real/Sim	1195.2
Kuka-iiwa	ego, 3rd	Grp	Real/Sim	338.2
WidowX	3rd×2, wrist	Grp	Real/Sim	305.2
Xarm7 Robot	top, 3rd	Grp	Real/Sim	594.4
UR5E	3rd, wrist	Grp	Real/Sim	208.3
Hello Stretch	ego, 3rd	Grp	Real	270.3
Kinova Jaco	3rd×3	Grp	Real/Sim	194.6
Kinova Gen3	3rd×3	Grp	Sim	197.4
Sawyer	3rd	Grp	Sim	196.1
Cobotta	3rd	Grp	Real	10.6
DLR Sara	3rd	Grp	Real	12.1
Willow Garage PR2	ego	Grp	Real	11.9
RMC RM65	3rd	Dex	Real	41.3
<i>Dual-Arm Robots</i>				
RMC Aida L	high, wrist×2	Grp	Real	325.7
Galaxea R1 Lite	high, wrist×2	Grp	Real	630.5
Agilex Split ALOHA	high, wrist×2	Grp	Real/Sim	1099.1
Agilex Cobot Magic	high, wrist×2	Grp	Real	399.2
Piper	ego, wrist×2	Grp	Real/Sim	904.5
<i>Portable Education Arm</i>				
BeingBeyond D1	ego	Dex	Real	100.0
SO101	ego, wrist×2	Grp	Real	194.4
<i>Half-Humanoid</i>				
PND AdamU	ego	Dex	Real	200.0
Agibot-G1	{ego, wrist}×2	Grp/Dex	Real/Sim	2391.7
Leju kuavo 4 LB	ego, wrist×2	Grp/Dex	Real	1198.2
AlphaBot 2	ego, wrist×2	Grp	Real	13.8
Airbot MMK2	3rd, wrist×2	Grp	Real	23.1
Galbot G1	wrist×2	Grp	Sim	195.8
Tianqing A2	high, wrist×2	Grp	Real	19.4
<i>Humanoid</i>				
Unitree G1edu-u3	ego, wrist×2	Grp/Dex	Real	135.7
Tiankung	top	Grp	Real	214.3
Total: 30 embodiments				13817.4

like LLaVA-Video [105]. These datasets serve to prevent the catastrophic forgetting of general world knowledge during robot-specific fine-tuning.

- 2) 2D Spatial Grounding & Affordance.** Robotic manipulation requires more than semantic description; it demands precise spatial localization. To bridge the gap between semantic understanding and physical actuation, we integrate datasets focused on grounded reasoning, object localization, and

affordance detection. Sources include RefCOCO [106] and RefSpatial [107] for referring expression comprehension, alongside robotics-specific grounding datasets such as RoboPoint [108], ShareRobot [109], RoboRefit [110], RoboVQA [111], MolmoAct [44], and A0-ManiSkill [112]. We also include point-level supervision benchmarks like PixMo-Points [113] and AsV2 [114] to enhance the model’s fine-grained spatial awareness within the 2D image plane.

- 3) **Task Planning & Reasoning.** In addition to immediate perception, a capable agent must reason over long horizons. We address this by incorporating high-level planning datasets, such as ShareRobot [109] and EO1.5M-QA [94]. These sources explicitly train the model to decompose complex, long-horizon commands into logical sub-task sequences, bridging the critical gap between abstract user intent and low-level motor execution.

4 UniCraftor: A System for Portable, Extensible, and Affordable Data Collection

While open-source corpora offer significant scale, they often lack critical supervision signals such as accurate depth, stable camera alignment, and temporally precise interaction events. For example, Ego4D [61] and its multi-view extension Ego-Exo4D [62] provide rich semantic descriptions but lack geometric depth. Egocentric-10K [115] features 10,000 hours of in-the-wild industrial footage but provides only raw RGB streams without annotation. Furthermore, benchmarks like HD-EPIC [116] and HOI4D [117] rely on offline calibration to approximate camera poses, and their interaction labels are typically aligned to clip boundaries or a handful of annotated frames. Such coarse labeling introduces temporal ambiguity regarding the exact moments of object contact and release, which is detrimental to fine-grained policy learning. To address these deficiencies, we develop a modular, plug-and-play data collection system to curate high-quality, multimodal recordings with synchronized supervision.

- 1) **Native Depth Acquisition:** We utilize a head-mounted Intel RealSense D435 for egocentric RGB-D capture. By leveraging active infrared stereo to provide raw physical depth, we bypass the artifacts and inconsistencies inherent in learning-based depth estimation, which often degrades under heavy occlusions, rapid ego-motion, or adverse lighting conditions.
- 2) **High-Precision Extrinsic:** To decouple hand-camera motion and establish a consistent world coordinate frame, we compute ground-truth camera poses via five tabletop AprilTags [118] using Perspective-n-Point (PnP) algorithms. This approach, rooted in classical robotic hand-eye calibration, offers superior stability against head-motion jitter and sensor noise compared to contemporary learning-based extrinsic predictors [98, 119]. Furthermore, it enables the seamless integration of exocentric views for flexible and synchronized view captures; additional D435 cameras can be freely repositioned without manual recalibration, provided at least one AprilTag remains within the field of view.
- 3) **Hardware-Synchronized Interaction Events:** To resolve the temporal ambiguity in automated labeling where frame downsampling often misses the precise moment of object contact or release, we incorporate a hardware-synchronized foot pedal to record explicit interaction events. This mechanism captures affordance-relevant keyframes with high temporal precision and minimal cognitive overhead for the demonstrator. These timestamps are synchronized across all sensor streams, serving as a ground truth for fine-grained semantic descriptions and verifying the temporal accuracy of manipulation cues.

Using this system (see Figure 5), we construct a dataset comprising **43 tabletop tasks** and **200+ hours** of multi-modal recordings. All sensors are synchronized via a central timestamping device, and the modular architecture permits rapid sensor interchange with planned extensions to tactile sensing and mobile scenarios. Our post-processing pipeline consists of three stages: first, AprilTag regions are tracked and inpainted using Grounded-SAM2 [120] and DiffuEraser [121] to ensure visual cleanliness; second, hand motions estimated via HaWor [98] are refined through multi-view depth integration to enforce cross-view spatial consistency; finally, fine-grained task descriptions are automatically generated by Qwen2.5-VL [122], conditioned on pedal-triggered event signals and undergo final human verification to ensure groundedness and correctness.

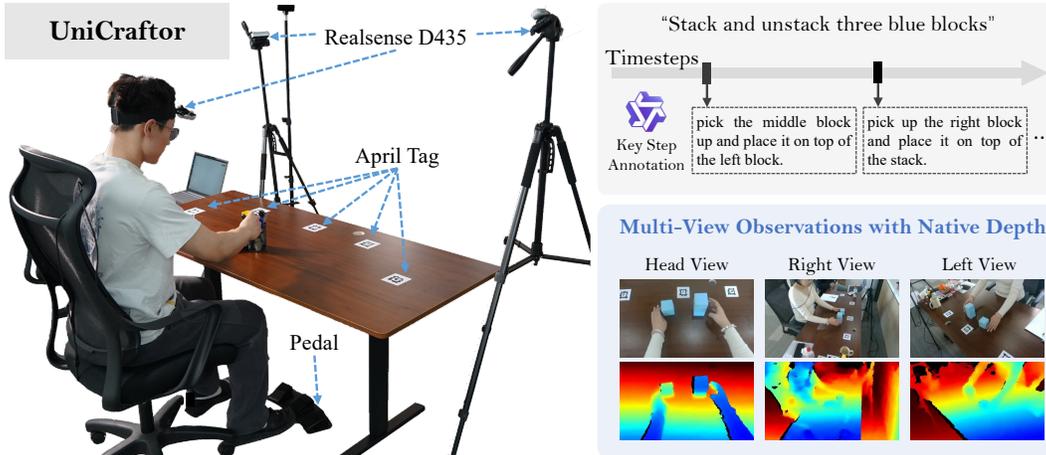


Figure 5: An overview of our data collection system UniCrafter.

5 Being-H0.5: A Foundational VLA Unifying Cross-Embodiment Control

5.1 Model Architecture

Following the architectural principles of BAGEL [123], Being-H0.5 adopts a specialized Mixture-of-Transformers (MoT) architecture [124] designed to disentangle high-level semantic reasoning from low-level motor control. Within this structure, Being-H0.5 integrates two distinct expert modules: **1) Multimodal Understanding Expert** is tasked with interpreting high-dimensional perceptual inputs. In the context of VLAs, this expert extends beyond traditional image captioning and is engineered for long-horizon planning, generating intermediate subgoals and providing the spatial reasoning necessary to ground the model in complex environments. **2) Action Generation Expert** serves as a dedicated policy network responsible for translating high-level plans into precise kinematic execution.

This dual-expert paradigm balances semantic comprehension with continuous control, aligning with the architectural consensus established by recent state-of-the-art models such as π_0 [8]. Crucially, while these experts maintain specialized functional roles, they operate within a unified MoT backbone. Both modules process a common token sequence and leverage shared self-attention mechanisms across each transformer layer. This design facilitates a seamless information flow, ensuring that action generation is deeply conditioned on visual-semantic context without the overhead of architectural bottlenecks.

Regarding the generation paradigm, we employ a hybrid approach tailored to the modalities of output. For textual outputs including high-level reasoning and motion description, we utilize the standard next-token-prediction paradigm which leverages the established effectiveness of VLM to produce coherent, logically structured instructional chains, while for discrete hand motion, a masked token prediction criterion is adopted. For action prediction, we adopt the Rectified Flow [33, 125] method, departing from discrete tokenization in favor of continuous actions. This aligns our policy head with cutting-edge diffusion-based VLA approaches, allowing for the generation of smooth, high-fidelity multimodal distributions over the action space.

The backbone of Being-H0.5 is initialized from InternVL-3.5 [126], a publicly available VLM featuring a robust decoder-only transformer architecture. We select InternVL-3.5 for its proven performance on complex visual reasoning benchmarks and accessibility. It is worth noting that the choice of the VLM backbone is critical, with empirical evidence suggests that the underlying visual features significantly dictate downstream VLA efficacy. Given that architectural biases can lead to substantial disparities in robotic performance, we intend to systematically benchmark alternative backbones in future exploration.

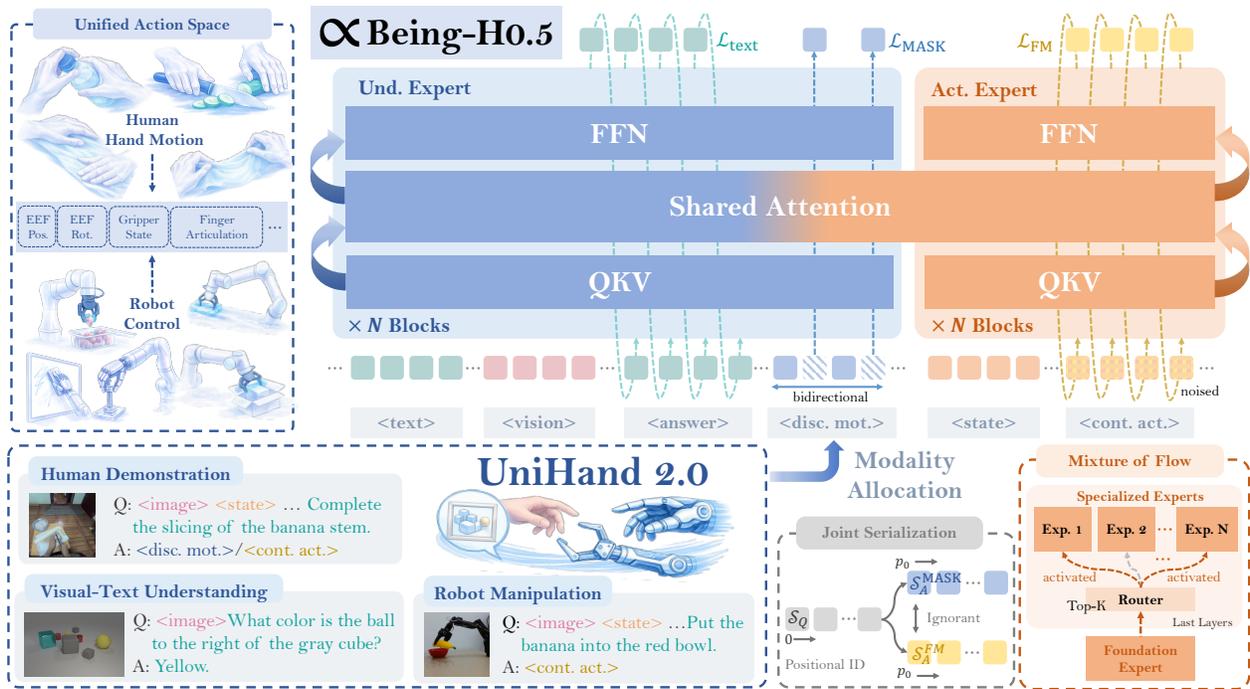


Figure 6: **Overview of Being-H0.5.** Being-H0.5 is a specialized MoT that disentangles multimodal understanding (**Und. Expert**) and action generation (**Act. Expert**) while maintaining coupling through shared attention mechanisms. A **unified state-action space** supports cross-embodiment pre-training by mapping human hand motion and diverse robot controls into semantically aligned slots. Our pre-training leverages **UniHand-2.0** by serializing multimodal data into a unified QA-style format, with each modality allocated to the relevant branch. Finally, a **Mixture-of-Flow** design scales action capacity by combining shared foundation layers with routed specialized experts for embodiment/task-specific dynamics.

5.1.1 Unified State-Action Space

A primary obstacle to scaling generalist robot learning is the extreme heterogeneity inherent in multi-source embodiment data. The “physical gap” between human hands and robotic end-effectors, combined with variances in kinematics, actuation limits, and control frequencies across diverse robotic platforms, results in a highly fragmented data landscape. Furthermore, existing datasets frequently utilize inconsistent nomenclature and disparate physical units for semantically identical actions. Prior approaches [9] often circumvent these discrepancies by assigning independent MLP projection layers (i.e., separate encoder/decoder heads) to each specific embodiment to handle dimensional variations. We posit that this strategy is suboptimal, as it unnecessarily consumes model capacity and fragments the learning of physical commonalities. In practice, diverse gripper configurations and dexterous hands exhibit high geometric consistency in End-Effector (EEF) trajectories, and many joint configuration possess underlying alignment potential. Isolating these through separate MLP heads prevents the model from leveraging shared physical priors, thereby degrading cross-embodiment generalization. Additionally, many existing methods tolerate mixed rotation representations (e.g., Euler angles, quaternions), forcing the network to squander computational resources on non-essential formatting discrepancies. This lack of structure introduces unnecessary noise, destabilizing the training process and severely limiting the model’s ability to transfer knowledge from human video data to robotic control.

To bridge these kinematic disparities, we introduce a **physically interpretable unified state-action space**. We formalize both state and action as fixed-length, high-dimensional vectors structured by the principle of **physical semantic alignment** proposed in [10]. Conceptually, we partition the vector space into semantically isolated subspaces where every dimension corresponds to a grounded physical quantity, such as bimanual EEF poses, joint positions, gripper width, finger articulations, or mobile base velocity and

heading commands. A critical innovation in our unified space is the treatment of **human hand motion as a generalized embodiment**. We map parameters from the MANO hand model directly into this unified space. Specifically, the global wrist pose of the human hand is aligned with the robotic EEF subspace, while finger articulations are mapped to reserved “fine-manipulation” slots. This architecture ensures that distinct degrees of freedom (DoF) do not conflict, allowing the model to learn a shared, embodiment-agnostic latent representation of manipulation logic. By utilizing this common “action language”, Being-H0.5 can be effectively pre-trained on vast libraries of human videos and episodes of different embodiment while remaining directly compatible with downstream robotic execution.

To guarantee transferability across heterogeneous platforms, we enforce strict standardization within each subspace. For Cartesian control, all end-effector actions are expressed as relative delta displacements in a unified world coordinate frame. Rotations are uniformly parameterized using Axis-Angle notation to prevent gimbal lock and ensure smooth interpolation on the SE(3) manifold. For joint-space control, positions are standardized as absolute radian values. Notably, we eschew traditional statistical normalization (e.g., scaling to [-1, 1]) in favor of preserving raw physical magnitudes. We argue that a movement of 1 radian or 10 centimeters carries intrinsic physical implications that normalization obscures. By applying only outlier filtering to mitigate sensor noise, we force the model to learn the true physical scale of actions, resulting in a policy that is both generalizable and physically grounded across diverse embodiments and environments.

5.1.2 Mixture of Flow

While a unified state-action space effectively represents diverse embodiments and prevents representational conflicts during pre-training, the limited capacity of conventional action experts remains a critical bottleneck. This capacity constraint often leads to performance degradation in VLAs when integrating robots with heterogeneous morphologies — particularly as action expert parameters are typically far fewer than those in flow-based experts for visual generation [123, 127]. Furthermore, such restricted capacity hampers the model’s ability to generalize across a broad spectrum of embodiments and complex downstream tasks.

To overcome these limitations, we introduce **Mixture of Flow (MoF)**: a scalable architectural framework designed to decouple distinct embodiments and skills while leveraging a shared foundational representation. Our approach is motivated by two key observations. First, many embodiments share partial control structures despite significant visual and kinematic differences. For instance, both Franka and Kuka robots are N-DoF manipulators defined by joint positions and end-effector states, while mobile platforms like wheeled half-humanoids share common velocity and orientation metrics for base control. Decomposing a unified expert into specialized sub-modules allows the model to master specific state-space intervals, thereby mitigating inter-embodiment interference. Second, human motor control is inherently modular and general motor primitives are dynamically adapted to specific tasks. Reflecting these principles, our flow-based action expert is structured into a two-tiered hierarchy: **1) Foundation Experts (Shared Dynamics)**: The initial layers of the action expert consist of standard transformer blocks shared across all inputs. These layers encode fundamental, transferable motor primitives (e.g., reaching, grasping dynamics, and collision avoidance) that remain invariant across disparate embodiments and tasks. **2) Specialized Experts (Embodiment & Task Routing)**: The upper layers utilize a suite of parallel, specialized experts managed by a learnable gating network, inspired by Mixture-of-Experts (MoE) architectures [12]. For a given input state and instruction, the router dynamically activates a sparse subset (e.g., Top-K) of experts, facilitating efficient specialization without a linear increase in computational overhead.

This architectural sparsity is fundamental to the efficiency and robustness of Being-H0.5. It allows the model to synthesize refined foundation primitives into complex, task-specific behaviors without cross-task interference. During training, gradients for a specific task update only the relevant expert pathway, thereby preserving the weights of other localized skills. Furthermore, this design decouples the total parameter count from the active parameter count. While the unified model hosts an extensive library of skills, only a fraction of its parameters are activated during inference. This makes Being-H0.5 highly deployable on resource-constrained edge hardware, such as the NVIDIA Orin-NX, allowing the robot to access broad, on-demand capabilities without exceeding memory bandwidth or latency constraints.

5.2 Pre-Training: Human-Centric Robot Learning

Our robot learning framework is fundamentally human-centric, utilizing the UniHand-2.0 dataset as its structural cornerstone. In this paradigm, human behavior is treated as a dense source of physical priors rather than a passive reference. UniHand-2.0 serves a dual role: its large-scale egocentric human data provides transferable behavioral intent to encode how complex interactions unfold in the open world, while its robot trajectories provide the high-fidelity kinematic supervision essential for low-level motor control.

To bridge the gap between human demonstration and robotic execution, we incorporate general Visual Question Answering (VQA) signals. This integration injects broad vision-language context and enhances scene understanding, addressing a critical dichotomy in embodied intelligence: the necessity for precise motor control coupled with rich contextual grounding. By synthesizing these heterogeneous streams, our multi-task objective trains the model to perceive with the nuance of a VLM while executing with the precise action.

5.2.1 Unified Sequence Modeling

We implement our learning framework by **casting all supervision into a unified, multimodal sequence modeling problem**. By treating the state-action space as an explicit modality, we define a state vector $\mathbf{s} \in \mathbb{R}^d$ and an action vector $\mathbf{a} \in \mathbb{R}^d$ shared across all embodiments. Each embodiment e (e.g., parallel grippers, dexterous hand, or mobile bases) is projected into this unified space via sparse slot assignments:

$$\mathbf{s} = \Phi_e(\mathbf{s}^{(e)}), \quad \mathbf{a} = \Phi_e(\mathbf{a}^{(e)}) \quad (1)$$

where $\mathbf{s}^{(e)}$ and $\mathbf{a}^{(e)}$ represent raw, embodiment-specific signals. The mapping function Φ_e routes these signals to relevant slots within the global vector, leaving unused slots as zeros. This allows human hand motion and robot trajectories to co-exist within a single, physically interpretable interface alongside vision and language. Formally, each training sample is serialized into a token stream \mathcal{S} composed of K modality segments:

$$\mathcal{S} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K] \quad (2)$$

where each segment $\mathbf{x}_k = \langle m_k, C_k \rangle$ consists of a modality tag $m_k \in \mathcal{M}$ and its corresponding content C_k . The modality set is defined as:

$$\mathcal{M} = \{\text{vision, text, state, action}\} \quad (3)$$

Each sample instantiates a subset $\mathcal{M}_i \subseteq \mathcal{M}$ based on available supervision, and each non-text modality is encapsulated within two distinct special tokens. During training, we adopt **Physical Instruction Tuning** [10] to organize data into a Query-Answer format $[\mathcal{S}_Q; \mathcal{S}_A]$. The model is conditioned on the context \mathcal{S}_Q and optimized via a generative loss exclusively on the response \mathcal{S}_A .

5.2.2 Human-Centric Multi-Task Objective

Within the unified sequence, we define a family of tasks distinguished by their modality organization. First, **motion generation** serves as the primary objective for manipulation learning. The model predicts an **action** chunk conditioned on **vision**, **text**, and **state**. This is the main pathway through which large-scale human interaction traces and robot trajectories provide policy-relevant supervision. To complement this, we introduce **motion description** and **continuation** as alignment objectives to strengthen vision-language-motion grounding and temporal coherence. For motion description, the model predicts **text** conditioned on **vision** and interaction traces (**state/action**), which enforces semantic grounding, aligning high-level linguistic intent with physically grounded outcomes. For motion continuation, the model predicts future **action** chunks conditioned on past observations and action history, encouraging the learning of temporally consistent interaction dynamics beyond single-step imitation.

While our pre-training is fundamentally human-centric, with egocentric human data forming the core of our data recipe, the QA-pair serialization seamlessly absorbs diverse auxiliary supervision. By instantiating only available modalities, the framework incorporates robot manipulation data for action generation ($\mathcal{S}_A = \{\text{action}\}$) and VQA datasets for standard text prediction ($\mathcal{S}_A = \{\text{text}\}$). All tasks share a common backbone and are optimized via a joint loss function:

$$\mathcal{L} = \lambda_{\text{text}} \mathcal{L}_{\text{text}} + \lambda_{\text{act}} \mathcal{L}_{\text{act}}. \quad (4)$$

To handle the multimodal nature of \mathcal{S}_A , we define token-level index sets Ω_{text} and Ω_{act} to identify segments of the unified sequence subject to each loss term. For VQA and motion-description tasks, we apply a standard cross-entropy loss over the target text tokens:

$$\mathcal{L}_{\text{text}} = - \sum_{i \in \Omega_{\text{text}}} \log p_{\theta}(y_i | \mathcal{S}_{<i}). \quad (5)$$

For motion generation and continuation (from both human and robot sources), we supervise the action tokens using a dedicated action loss \mathcal{L}_{act} , the formulation of which—dependent on the specific motion representation—will be detailed in the following section.

5.2.3 Hybrid Human Motion Representation

Human hand motion is inherently expressive; however, at scale, it inevitably exhibits subtle execution variations and observational noise. Relying exclusively on continuous supervision can result in brittle learned priors, while purely discrete representations may sacrifice the precision necessary for fine-grained control. To capture both the high-fidelity requirements of robot control and the stable behavioral priors of human-like movement, we supervise motion using two complementary representations within the same training instance: 1) continuous action chunks in the unified space and 2) discrete motion tokens derived via chunk quantization, as introduced in Being-H0 [10].

Specifically, given a motion chunk of length T , we represent it as a continuous sequence $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_T] \in \mathbb{R}^{T \times d}$, where each $\mathbf{a}_t \in \mathbb{R}^d$ resides in our unified, high-dimensional action space. Simultaneously, we leverage a pretrained tokenizer to quantize the motion into a discrete sequence $\mathbf{z} \in \{1, \dots, |\mathbb{C}|\}^{T_z}$. Here, \mathbb{C} denotes a codebook of size $|\mathbb{C}|$, and T_z represents the resulting tokenized length. This discrete channel serves as a robust, language-like abstraction that filters high-frequency execution noise, thereby stabilizing the motion prior across heterogeneous datasets. The joint objective is formulated as weighted combination of **Continuous Flow-Matching** and discrete **Masked Motion Prediction**:

$$\mathcal{L}_{\text{act}} = \lambda_1 \mathcal{L}_{\text{FM}} + \lambda_2 \mathcal{L}_{\text{MASK}}. \quad (6)$$

Continuous Flow-Matching. We model the continuous motion distribution using a time-conditioned velocity field $v_{\theta}(\mathbf{x}, t, c)$, where c represents the conditioning visual and textual context. This field is designed to transport samples from a standard Gaussian distribution $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ toward the target data distribution. For a target action \mathbf{a}_i at temporal index i , we define a probability path via linear interpolation: $\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{a}_i$ for $t \in [0, 1]$. The objective \mathcal{L}_{FM} minimizes the mean squared error between the predicted velocity and the ideal vector field $\mathbf{u}_t(\mathbf{x}_t) = \mathbf{a}_i - \mathbf{x}_0$. Let Ω_{FM} denote the set of indices for continuous action steps within the answer segment $\mathcal{S}_A^{\text{FM}}$. The objective can be denoted as:

$$\mathcal{L}_{\text{FM}} = \sum_{i \in \Omega_{\text{FM}}} \|v_{\theta}(\mathbf{x}_t, t, c) - (\mathbf{a}_i - \mathbf{x}_0)\|_2^2. \quad (7)$$

Masked Motion Token Prediction. In addition to continuous regression, we supervise the discrete channel to instill a stable, abstraction-level of motion primitives. Specifically, we augment the model backbone with a dedicated token embedding and a linear prediction head for quantized tokens. During training, we apply a masking strategy to the token sequence $\mathbf{z} \in \{1, \dots, |\mathbb{C}|\}^{T_z}$ by randomly sampling a subset of indices $\Omega_{\text{MASK}} \subseteq \{1, \dots, T_z\}$ according to a masking ratio ρ , replacing the original tokens with a specialized [MASK] token. The model is tasked to reconstruct the original codebook indices via a cross-entropy loss:

$$\mathcal{L}_{\text{MASK}} = - \sum_{i \in \Omega_{\text{MASK}}} \log p_{\theta}(z_i | c). \quad (8)$$

By predicting these discrete tokens, the model learns the underlying “grammar” of hand motion, providing a structural scaffold that supports the continuous flow-matching head.

Joint Serialization. To achieve seamless multi-modal integration, we serialize the context and hybrid targets into a unified sequence. For a given instance, we concatenate a shared context prefix \mathcal{S}_Q with two target segments: $[\mathcal{S}_Q; \mathcal{S}_A^{\text{FM}}; \mathcal{S}_A^{\text{MASK}}]$. Here, $\mathcal{S}_A^{\text{FM}}$ contains continuous action targets and $\mathcal{S}_A^{\text{MASK}}$ contains discrete motion-token targets, both conditioned on the same context \mathcal{S}_Q . To prevent trivial information leakage where the model might copy information between the continuous and discrete channels, we enforce both target segments $\mathcal{S}_A^{\text{FM}}$ and $\mathcal{S}_A^{\text{MASK}}$ can attend to the shared context \mathcal{S}_Q but remain mutually invisible. This is implemented via a modified attention mask with a gating matrix:

$$G = \begin{bmatrix} \mathbf{1} & \mathbf{0} & \mathbf{0} \\ \mathbf{1} & \mathbf{1} & \mathbf{0} \\ \mathbf{1} & \mathbf{0} & \mathbf{1} \end{bmatrix}, \quad \text{for } [\mathcal{S}_Q; \mathcal{S}_A^{\text{FM}}; \mathcal{S}_A^{\text{MASK}}]. \quad (9)$$

The effective attention mask is implemented by applying this gating onto the original mask. Furthermore, we align the positional indices to ensure that both targets are grounded from the same positional origin. Let $r(j)$ be the relative index within a segment, and $p_0 = \max_{j \in \mathcal{S}_Q} (j + 1)$. The positional encoding is defined as:

$$\text{PE}(j) = \begin{cases} j, & j \in \mathcal{S}_Q, \\ p_0 + r(j), & j \in \mathcal{S}_A^{\text{FM}}, \\ p_0 + r(j), & j \in \mathcal{S}_A^{\text{MASK}}. \end{cases} \quad (10)$$

This alignment allows the model to perceive the continuous and discrete representations as two complementary views of the same temporal event, grounded to an identical contextual origin.

5.3 Post-Training: Towards Cross-Embodiment Adaptation

While the UniHand-2.0 dataset provides Being-H0.5 extensive embodiment diversity, yielding stronger *structural priors* than existing VLAs, the **post-training** phase remains a non-trivial challenge. The primary objective is to adapt the pre-trained policy to specialized deployment robots, which often possess unique kinematics, constraints, and runtime characteristics, without eroding the general-purpose representations acquired during large-scale pretraining. In practice, this stage is dominated by a plasticity–stability dilemma: downstream fine-tuning can easily collapse into a narrow specialist, achieving high success within a limited trajectory distribution at the expense of reasoning depth and cross-embodiment transferability.

We identify three coupled sources of brittleness in cross-embodiment deployment: 1) morphological interference from embodiment-dependent action fields competing under shared parameters; 2) context unreliability under distribution shift, where degraded features induce action jitter in flow updates; and 3) real-time temporal mismatch between inference and execution under heterogeneous latency/frequency. To address them, we introduce two complementary post-training techniques (**Embodiment-Specific Adaptation** and **Manifold-Preserving Gating**), as well as a deployment-aware training protocol, **Universal Async Chunking**.

5.3.1 Embodiment-Specific Adaptation

Despite the powerful vision-language-action representations learned during pre-training, the model must still account for the unique physical dynamics of target hardware. Each embodiment e induces a different feasible action set governed by kinematic limits, self-collision, and contact constraints. consequently, the conditional distribution $p(\mathbf{a} \mid H, e)$ becomes increasingly multimodal and less smooth as embodiment complexity increases. From a flow-matching view, the action expert learns an embodiment-conditioned velocity field $v_\theta(\mathbf{a}_t; H, e)$ to match the target field $v^*(\mathbf{a}_t, \mathbf{a}_0)$:

$$\min_{\theta} \mathbb{E}_e \mathbb{E}_{(\mathbf{a}_0, H) \sim \mathcal{D}_e} \mathbb{E}_t \|v_\theta(\mathbf{a}_t; H, e) - v^*(\mathbf{a}_t, \mathbf{a}_0)\|_2^2. \quad (11)$$

When optimal fields differ substantially across embodiments, embodiment-specific gradients $\nabla_{\theta} \mathcal{L}_e$ become misaligned (morphological interference), slowing convergence and yielding unstable or “averaged” behaviors on complex hardware.

Key Design Goal. Our objective is to localize specialization to embodiment-relevant action components while maintaining the stability of the shared VLM backbone. Notably, we eschew the conventional strategy of

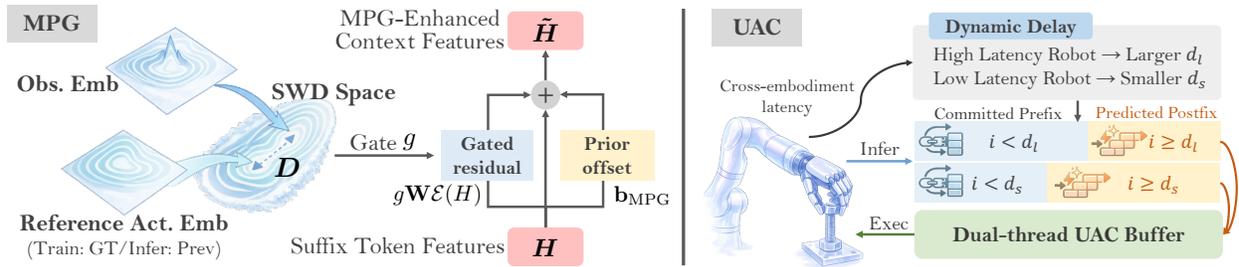


Figure 7: **MPG and UAC Overview.** **Left (MPG):** We compare observation embeddings with a reference action embedding (Train: ground truth; Inference: previous iterate) in the Sliced-Wasserstein Distance (SWD) space to obtain a discrepancy-guided gate g . The gate scales a *feature-conditioned* residual while an *ungated* learned prior offset provides a stable fallback, producing enhanced context features \tilde{H} for the action expert. **Right (UAC):** Based on embodiment-specific dynamic delay d , each predicted action chunk is split into a committed prefix $\mathbf{A}_{<d}$ (already queued/executing) and a predicted postfix $\mathbf{A}_{\geq d}$. A dual-thread buffer enables asynchronous inference/execution across robots with heterogeneous latency budgets.

assigning a separate projection head to each robot. Instead, ESA leverages our unified action space: different embodiments activate different index sets (with partial overlap for shared morphological components), and we only update parameters tied to these active indices.

Slot-Wise Adapters in a Unified Action Space. Let $\mathcal{S} = \{1, \dots, K\}$ denote the K semantic action slots (*e.g.*, arm joints, hand/gripper joints, base controls), and let $\mathcal{I}_e \subseteq \mathcal{S}$ denote the active slot indices for embodiment e . Embodiments sharing hardware components (*e.g.*, the same arm with different hands) feature overlapping index sets and therefore share a subset of adaptation parameters. We maintain a slot-wise bank of lightweight adapters $\mathbf{W}_{\text{ESA}} \in \mathbb{R}^{K \times d_{\text{out}} \times d_{\text{in}}}$, and during post-training, update only the subset indexed by \mathcal{I}_e on embodiment e :

$$\mathbf{W}_{\text{ESA}}^{(e)} \triangleq \{\mathbf{W}_{\text{ESA}}[k] : k \in \mathcal{I}_e\}, \quad \Delta \mathbf{W}_{\text{ESA}}[k] = \mathbf{0} \quad \forall k \notin \mathcal{I}_e. \quad (12)$$

Operationally, the ESA update is masked by the slot selector \mathcal{I}_e : only adapters tied to active slots contribute gradients, so parameters tied to shared slots can transfer knowledge across embodiments, while non-overlapping slots remain isolated. The result is a single checkpoint capable of supporting heterogeneous robots without parameter conflict.

5.3.2 Manifold-Preserving Gating

Flow matching transports noise into the action space conditioned on token-level context features H consumed by the action expert. In our implementation, H refers to the *suffix token features* formed by concatenating proprioceptive state embeddings and current action-token embeddings (projected to the VLM hidden size), while the language/vision prefix remains unchanged.

Motivation. A critical but often implicit assumption is that H faithfully captures task-relevant semantics. Under distribution shift (*e.g.*, lighting changes, viewpoint perturbations, partial occlusions), a policy may blindly regress actions from degraded features, producing unstable behavior. This issue is amplified in cross-embodiment adaptation: higher-DoF embodiments introduce additional unobserved state and stronger constraints, making the mapping from H to feasible actions less deterministic and more sensitive to small feature errors. In flow-matching, a denoising update with step size Δt is $\mathbf{a}_{t-\Delta t} = \mathbf{a}_t + \Delta t v_\theta(\mathbf{a}_t; H)$; thus perturbations in context features (*e.g.*, $H = H^* + \epsilon$) can induce action variance proportional to $\|\partial v_\theta / \partial H\|^2 \text{Var}(\epsilon)$, manifesting as action jitter.

Inspired by DiG-Flow [128], MPG treats H as a *noisy conditioning signal* and explicitly estimates its reliability. When H is out-of-distribution for a target embodiment, conditioning the action expert on H can inject spurious signals and force the model to fit inconsistent velocity fields, often resulting in jittery actions. MPG therefore utilizes a discrepancy-guided gate to *down-weight feature-dependent corrections* while retaining an *ungated* learned prior offset for graceful fallback (see Figure 7, left).

Gated Residual with an Ungated Prior Offset. MPG computes a gate $g \in (0, 1]$ (large discrepancy \Rightarrow small g) and modulates only the *feature-conditioned residual pathway*:

$$\tilde{H} = H + \lambda \cdot \mathcal{P}_{\text{MPG}}(g \cdot \mathcal{E}_{\text{obs}}(H)) = H + \lambda g \mathbf{W}_{\text{MPG}} \mathcal{E}_{\text{obs}}(H) + \lambda \mathbf{b}_{\text{MPG}}, \quad (13)$$

where \mathcal{E}_{obs} projects H to a common embedding space and $\mathcal{P}_{\text{MPG}}(\cdot) = \mathbf{W}_{\text{MPG}}(\cdot) + \mathbf{b}_{\text{MPG}}$ generates the enhancement residual. This yields an intuitive fallback: when $g \approx 1$, MPG trusts the current context and allows strong feature-conditioned adaptation; when $g \approx 0$, the feature-dependent term is suppressed, and the model falls back to a stable, learned prior offset $H + \lambda \mathbf{b}_{\text{MPG}}$. Because \mathbf{b}_{MPG} is *ungated*, it continues to influence \tilde{H} even in low-confidence regimes, effectively learning a robust default correction for uncertain contexts.

Variance Reduction vs. Output Gating. Unlike conventional output gating ($\tilde{H} = H + \lambda g \mathcal{R}(H)$ in DiG-Flow [128]), MPG applies the gate before the projection. Under conditions of gate fluctuations (*e.g.*, due to visual jitter), output gating scales both the projected term and the bias, inducing variance proportional to $\|\mathbf{W} \mathcal{E}_{\text{obs}}(H) + \mathbf{b}\|_2^2$. In contrast, MPG scales only the projected term, yielding variance proportional to $\|\mathbf{W} \mathcal{E}_{\text{obs}}(H)\|_2^2$. Since the ungated offset remains constant w.r.t. g , MPG produces smoother trajectories under contextual perturbation.

Discrepancy and Gate Computation. To assess context reliability, we construct an action prior anchor from noise-free actions. Let Z^{nf} denote the noise-free action-token embeddings encoded at zero noise level, summarize by mean-pooling $\bar{Z} = \text{MeanPool}(Z^{\text{nf}})$. Intuitively, \bar{Z} characterizes the manifold of *plausible* action embeddings under the model’s training distribution. We quantify the feature–action distributional discrepancy in a shared, scale-invariant space using the sliced Wasserstein distance (SWD) [129, 130]:

$$D(\mu_{\hat{H}}, \mu_{\hat{Z}}) \approx \frac{1}{M} \sum_{m=1}^M \left\| \text{sort}(\theta_m^\top \hat{H}) - \text{sort}(\theta_m^\top \hat{Z}) \right\|_2^2, \quad (14)$$

where θ_m represents random unit directions, $\hat{H} = \text{LN}(\mathcal{E}_{\text{obs}}(H))$, and $\hat{Z} = \text{LN}(\mathcal{E}_{\text{act}}(\bar{Z}))$. The \bar{Z} is broadcast along the observation sequence length when computing D . The discrepancy yields a temperature-scaled reliability gate:

$$g = \exp(-D/\tau) \in (0, 1]. \quad (15)$$

In practice, we apply a stop-gradient operation to the gate, $g^{\text{sg}} = \text{stopgrad}(g)$, to prevent degenerate solutions where the model might learn to manipulate g rather than improving underlying feature quality. With this operation, g behaves as a pure scaling factor on the enhancement pathway. Ignoring backpropagation through g , the Jacobian is approximately:

$$\frac{\partial \tilde{H}}{\partial H} \approx \mathbf{I} + \lambda g \mathbf{W}_{\text{MPG}} \frac{\partial \mathcal{E}_{\text{obs}}(H)}{\partial H}, \quad (16)$$

This formulation ensures that the feature-dependent correction becomes increasingly insensitive when the context is unreliable (small g), providing a robust safeguard against out-of-distribution observations.

5.3.3 Universal Async Chunking

Motivation. Deploying action-chunking policies on physical hardware introduces a fundamental temporal mismatch: while the model computes a subsequent action chunk, the robot must continue executing previously committed trajectories. In simulated environments, this latency gap is often negligible. However, on real-world robot, inference latency directly induces control discontinuities and catastrophic task failures.

Training-Time RTC [131] mitigates this by simulating inference delay during the training phase: prefix actions are assigned a clean timestep ($t=1$), while postfix actions receive stochastic noisy timesteps ($t<1$), with the loss computed exclusively on the postfix. Nevertheless, RTC is typically instantiated for a single platform. This remains a significant limitation for a cross-embodiment VLA where diverse robotic platforms operate under heterogeneous control periods and distinct latency budgets.

UAC extends RTC by rendering the delay *embodiment-aware*. For an embodiment e characterized by a control period $\Delta t^{(e)}$ and an expected inference latency budget $L^{(e)}$, the effective delay in control steps is scaled as

$\lceil L^{(e)}/\Delta t^{(e)} \rceil$. UAC utilizes an embodiment-specific delay distribution and maximum delay threshold, and is trained jointly with ESA, enabling a single checkpoint to accommodate diverse runtime characteristics (see Figure 7, right).

Embodiment-Aware Delay Sampling and Objective. We sample a delay

$$d \sim \pi^{(e)}(d), \quad d \in \{0, 1, \dots, d_{\max}^{(e)} - 1\}, \quad (17)$$

where both the distribution $\pi^{(e)}$ and maximum delay $d_{\max}^{(e)}$ are configured per embodiment based on its specific control frequency and expected inference latency. Following RTC, we partition each action chunk into a committed prefix $\mathbf{A}_{<d}$ and a predicted postfix $\mathbf{A}_{\geq d}$, assigning per-token timesteps as follows:

$$t_i = \mathbb{1}[i < d] + \mathbb{1}[i \geq d] \cdot t_{\text{base}}, \quad t_{\text{base}} \sim p(t). \quad (18)$$

The training objective is formulated to exclude prefix positions:

$$\mathcal{L}_{\text{UAC}} = \sum_{i \geq d} \|\hat{v}_i - v_i^*\|_2^2. \quad (19)$$

Combined with embodiment-specific adaptation, UAC trains a unified model to generate smooth continuations across disparate robots, each with its unique delay characteristics, while remaining compatible with real-time asynchronous deployment via a dual-thread buffer, as detailed in Section 6.3.

6 Infrastructure: Real-Time Cross-Embodiment Deployment

Bridging the gap between high-capacity VLA models and real-world robotic control requires a deployment pipeline that prioritizes *temporal consistency*, *robustness to distribution shift*, and *low-latency execution*. In this work, we design an inference infrastructure that translates Being-H0.5’s cross-embodiment and asynchronous capabilities into robust real-time control. Concretely, the infrastructure consists of three tightly coupled components: (i) **rectified-flow action inference** with cached prefix features for efficiency; (ii) **manifold-preserving refinement** powered by MPG to stabilize inference under shift; and (iii) a **universal async chunking protocol** with a dual-thread ring-buffer architecture for continuous control across heterogeneous platforms. We refer the reader to Figure 7 for recalling the MPG/UAC mechanisms.

6.1 Manifold-Preserving Refinement

Rectified-Flow Denoising. Action generation follows an iterative denoising process based on rectified flow. Starting from Gaussian noise $\mathbf{a}^{(0)} \sim \mathcal{N}(0, \mathbf{I})$, we integrate the learned velocity field v_θ using Euler steps:

$$\mathbf{a}^{(k+1)} = \mathbf{a}^{(k)} + \Delta t \cdot v_\theta(\mathbf{a}^{(k)}, t^{(k)} | H), \quad \Delta t = 1/K, \quad t^{(k)} = k/K, \quad (20)$$

where K is the number of denoising steps and H denotes the token-level context features consumed by the action expert. As described in Section 5.3.2, H consists of a static vision-language prefix and a dynamic suffix (state tokens concatenated with the current action-token embeddings). During denoising, only the *action-token embeddings* evolve with $\mathbf{a}^{(k)}$, so the vision-language prefix can be cached (kv-cache) for efficiency.

In practice, we keep $K < 10$ to balance action fidelity and latency. This low-step regime is precisely where MPG becomes valuable: by providing a stable fallback under low confidence, MPG can reduce sensitivity to noisy context and allow strong performance with fewer denoising steps (Figure 7, left).

Two-Stage Refinement with MPG. When MPG is enabled at inference time, we use a lightweight two-stage strategy: (1) a baseline pass produces a reference prediction, and (2) a small number of refinement rounds use the previous prediction as a proxy anchor to compute discrepancy and gate feature enhancement.

(1) Stage 1. Baseline prediction (w/o MPG). We first run rectified-flow denoising with the raw context features H to obtain an initial prediction $\mathbf{a}_{\text{pred}}^{(0)}$. This baseline may be suboptimal under a distribution shift, but it provides a concrete reference for constructing the MPG action anchor.

(2) Stage 2. Iterative refinement (w/ MPG). We then perform N_{ref} refinement rounds ($N_{\text{ref}} \in \{1, 2, 3\}$ in practice). At refinement round n , we form a noise-free reference action embedding from the previous prediction:

$$\bar{Z}^{(n-1)} = \text{MeanPool}\left(\text{Enc}(\mathbf{a}_{\text{pred}}^{(n-1)}, \sigma = 0)\right), \quad \mathbf{a}_{\text{pred}}^{(n)} = \text{FlowDenoise}\left(H; \bar{Z}^{(n-1)}\right), \quad (21)$$

where σ denotes the action noise level and $\sigma = 0$ is noise-free. Here, MPG is applied *inside the denoising loop* at each Euler step, operating on the current suffix features and using $\bar{Z}^{(n-1)}$ as the reference action anchor for discrepancy computation (Section 5.3.2). Intuitively, this creates a feedback loop: better actions \Rightarrow better anchors \Rightarrow more accurate gates \Rightarrow more robust features.

Why MPG Helps in the Low-Step Regime. MPG’s input-gating formulation (Section 5.3.2) is particularly compatible with small K . When the gate g is low (high discrepancy), MPG suppresses feature-dependent correction and falls back to the ungated prior offset, effectively shifting to $\tilde{H} \approx H + \lambda \mathbf{b}_{\text{MPG}}$. This stabilizes the velocity field prediction $v_\theta(\cdot | \tilde{H})$ under uncertainty, reducing the “distance” the flow needs to traverse from noise to a feasible action mode. Empirically, performance typically saturates within 2–3 refinement rounds even with $K < 10$.

6.2 Universal Async Chunking Protocol

Motivation. Real robots execute actions continuously while inference happens intermittently. With chunked action generation, this induces a temporal mismatch: while the model computes a new chunk, the robot continues executing previously committed actions. If the policy ignores this mismatch, inference latency causes trajectory discontinuities and control stuttering.

Training-Time RTC [131] addresses this mismatch in *training* by simulating inference delay and computing loss only on the postfix. Here we formalize a *deployment protocol* that makes this principle explicit and cross-embodiment: we commit a prefix based on latency, lock it throughout denoising, and stitch only the postfix back into the execution buffer (Figure 7, right).

Latency Commitment. Before each inference cycle, we declare a delay d that upper-bounds the number of control steps that will be executed before the new chunk is ready:

$$d \geq \left\lceil t_{\text{inference}}/t_{\text{control}}^{(e)} \right\rceil + \epsilon_{\text{safety}}, \quad (22)$$

where $t_{\text{control}}^{(e)}$ is the embodiment-specific control period and ϵ_{safety} is a small margin to absorb timing jitter. A key property is that *overestimation is safe while underestimation breaks continuity*. Overestimating d simply increases the locked prefix length; underestimating d means the robot executes actions that the model did not condition on, causing discontinuity.

For cross-embodiment deployment, we configure d per platform using: (1) **control period** $t_{\text{control}}^{(e)}$ (higher frequency \Rightarrow more steps per second); (2) **expected inference latency** $t_{\text{inference}}$ (model/hardware dependent); and (3) **safety margin** ϵ_{safety} (platform-specific tolerance).

Hard Prefix Locking (denoising-time constraint). During denoising, we enforce prefix consistency for all Euler steps:

$$\mathbf{a}_i^{(k)} \leftarrow \begin{cases} \mathcal{B}_i, & i < d, \\ \mathbf{a}_i^{(k)}, & i \geq d, \end{cases} \quad \forall k \in \{0, \dots, K\}, \quad (23)$$

where \mathcal{B} is the execution buffer containing the most recently committed actions. This prevents the flow from “editing” actions that are already being executed (or will be executed before inference finishes).

Hard Stitching (buffer update rule). After denoising, we discard the prefix $\mathbf{a}_{<d}$ and write only the postfix $\mathbf{a}_{\geq d}$ into the buffer:

$$\mathbf{a}_{\text{exec}}^{(n)} = \mathbf{a}_{<d}^{(n-1)} \oplus \mathbf{a}_{\geq d}^{(n)}, \quad (24)$$

which guarantees continuity provided the latency commitment (22) is honored.

Table 2: **Hardware specifications of heterogeneous real-robot embodiments.**

Embodiment			Structure		Vision	
Group	Name	DoF	Joints	Hand	Camera	Configuration
Upper-body Humanoid	PND Adam-U	31	Bimanual+Head+Waist	Dexterous (6DoF)	ZED Mini	Movable (ego), dual-camera
	Unitree G1 + LinkerBot O6	26	Bimanual	Dexterous (6DoF)	D435	Fixed (ego)
Single arm + Dexterous hand	FR3 + Inspire Hand	13	Single-arm	Dexterous (6DoF)	2×D435	Fixed (3rd-person)
	BeingBeyond D1	14	Single-arm	Dexterous (6DoF)	D435	Movable (ego)
Single arm + Gripper	LeRobot SO-101	6	Single-arm	Gripper	D435	Fixed (3rd-person)

6.3 Dual-Thread Deployment Architecture

To realize UAC in real time, we decouple inference and control into separate threads communicating through a shared ring buffer \mathcal{B} (Figure 7, right):

- **Control thread (consumer).** Runs at fixed frequency $1/t_{\text{control}}^{(e)}$, popping actions from \mathcal{B} and sending them to the robot. If the buffer runs low, it holds the last action or executes a platform-specific safe fallback to avoid discontinuities.
- **Inference thread (producer).** Runs asynchronously: it fetches observations, performs rectified-flow denoising (optionally with MPG refinement, Section 6.1), and writes *only* postfix actions into \mathcal{B} starting at offset d .

We size the ring buffer at least $2\times$ the chunk length to reduce underflow risk under typical latency jitter. A lightweight mutex protects concurrent access during brief read/write operations, ensuring thread safety without introducing meaningful blocking.

Benefits. This architecture provides: (1) **latency absorption:** inference jitter does not translate to control jitter; (2) **graceful degradation:** buffer underflow triggers safe fallback rather than undefined behavior; and (3) **cross-embodiment compatibility:** the same protocol supports robots with different control frequencies and compute budgets by adjusting (d, \mathcal{B}) parameters and latency margins.

Together, these infrastructure components enable a single generalist Being-H0.5 checkpoint to achieve real-time performance across diverse platforms (e.g., 10 Hz tabletop manipulators to 50 Hz humanoids) under both cloud computing (high network latency) and edge deployment (limited computing performance, e.g., NVIDIA Orin NX), while maintaining smooth and continuous control.

7 Experiment

7.1 Comparison across Real Robots

We evaluate Being-H0.5 across five real-world embodiments characterized by substantially different morphologies, sensing setups, and manipulation profiles. A total of 10 tasks are designed to rigorously assess four core objectives: 1) analyzing spatial, long-horizon, bimanual, and generalization capabilities; 2) validate that a **single checkpoint** can be deployed across heterogeneous embodiments (generalist); 3) quantify the performance gap relative to embodiment-specific optimization (specialist); and 4) isolating the architectural and data components critical for real-time robustness (ablation).

7.1.1 Experimental Setup

The hardware configurations for our experiments are illustrated in Figure 8, with evaluated embodiments and sensory setups summarized in Table 2. The comprehensive real-robot task suite is detailed in Table 3. For each task, we collect 30 to 60 minutes of real-robot demonstration data. We evaluate Being-H0.5 against the following methods under the same evaluation infrastructure:

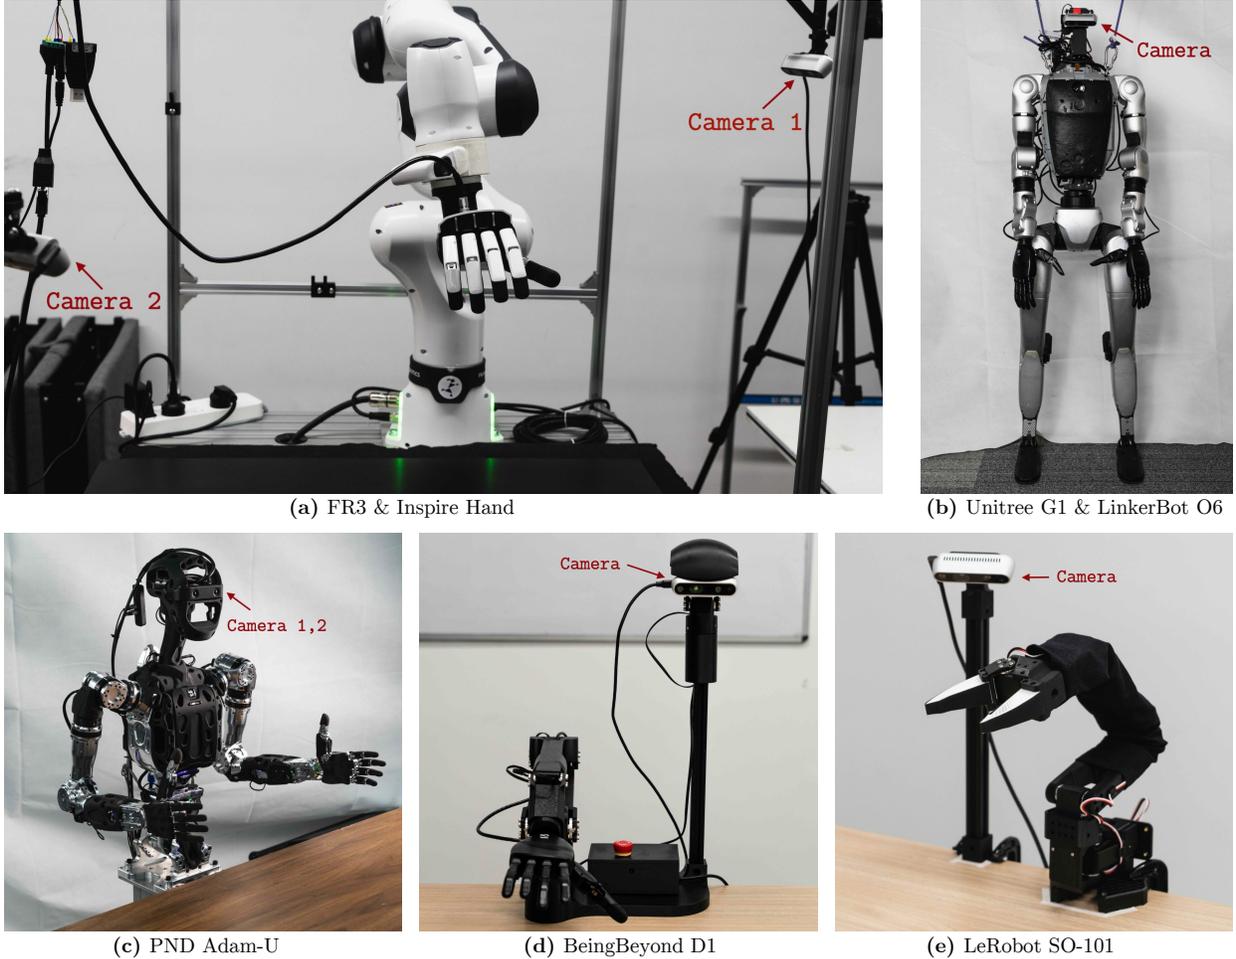


Figure 8: **Real-robot embodiments.** We evaluate on five robotic platforms spanning upper-body humanoids, single-arm dexterous manipulation, and lightweight parallel grippers.

- **Being-H0.5-specialist.** Starting from our pre-trained backbone, we conduct embodiment-specific post-training, where the model is fine-tuned independently for each robotic platform using the full task suite collected for that specific embodiment.
- **Being-H0.5-generalist.** A single checkpoint jointly post-trained on all embodiments and all tasks, enabling deployment across the five platforms.
- **Being-H0.5-scratch:** A baseline without UniHand-2.0 pre-training (specialist & generalist). Utilizing the same MoT architecture and post-training pipeline as Being-H0.5, but initialized without human-centric robot learning via UniHand-2.0. We report both specialist and generalist variants to isolate the contribution of large-scale human-robot pre-training to cross-embodiment transfer.
- $\pi_{0.5}$. A competitive open-source VLA baseline fine-tuned on the same per-embodiment data. Notably, due to its embodiment-specific action interface (requiring zero-padding to a fixed dim), $\pi_{0.5}$ does not natively support the cross-embodiment generalist setting; it’s therefore evaluated exclusively in the specialist regime.

Evaluation Protocol. A primary risk in real-robot evaluation is operator bias, which can significantly compromise the fairness of comparisons. To avoid such bias, we implement a **blind evaluation protocol** via a unified black-box inference server. The core of this protocol is the **encapsulation of all candidate**

Table 3: **Real-robot task suite and instructions.** Tasks are categorized by their requisite manipulation capabilities.

Category	Task	Embodiment	Instruction
Spatial	Arrange Flowers	Adam-U	“Arrange the flowers neatly into the vase.”
	Water Plant	FR3+Inspire	“Use your hand to hold the sprinkling can and water the plants”
	Stack Bowls	FR3+Inspire	“Hold the edge of the bowls and stack them on the plate one by one.”
	Stack Blocks	D1	“Stack the blocks following the color order: <ORDER>.”
Long-horizon	Drawer: Open → Place → Close	FR3+Inspire	“Open the drawer, put <OBJECTS> inside, then close the drawer.”
	Package: Flip → Scan	G1+O6	“Flip the package then scan the barcode with the scanner.”
Bimanual	Hand-over → Box	Adam-U	“Use both hands to pick and pass <OBJECTS> into the box.”
	Pack two products → Close lid	G1+O6	“Use both hands to put <OBJECTS> into the box, then close the lid.”
Generalization	Wipe Whiteboard	FR3+Inspire	“Use the wiper to clean any marked area of the whiteboard.”
	Clear Table	Adam-U	“Clear the table: put all objects into the box.”
	Clear Table	SO-101	“Clear the table: put all objects into the plate.”

policy endpoints into a unified black-box inference server. This server presents a single, standardized interface, ensuring that the evaluation infrastructure is identical for every model. Before each evaluation session, we pre-define a comprehensive set of N distinct layouts for each task, systematically varying object positions, orientations, and scene layouts to cover the expected operational distribution and challenge edge cases. During evaluation, each trials is executed through the following automated and blinded pipeline:

- *Configuration Sampling.* The system randomly samples one of the N pre-defined configurations and resets the scene accordingly.
- *Policy Sampling.* The system randomly selects a policy from the candidate pool. Critically, the policy’s identity completely hidden from the human operator.
- *Blinded Execution and Scoring.* The operator will be unaware of which model is being evaluated. Upon task completion or timeout, the operator records a binary success/failure outcome based solely on the pre-defined, objective success criteria.

The above process repeats until stopping condition is met: each policy must be evaluated for exactly K trials (default $K=20$) under each preset configuration. This protocol guarantees three key properties for statistical validity: 1) all policies encounter identical environmental conditions; 2) the human operator **cannot consciously or unconsciously favor any particular model**, and 3) the evaluation order is randomized to avoid temporal confounds (*e.g.*, robot wear, lighting drift).

7.1.2 Result Analysis

Figure 9 shows the example rollouts of the experiments. Figure 10 summarizes category-level success on real robots. We analyze the experimental results in detail in the following section.

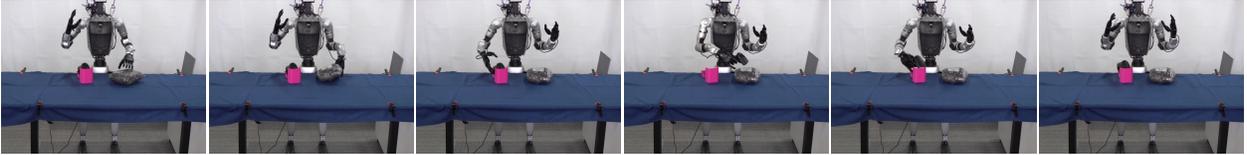
Specialist is strongest, but generalist stays surprisingly close. Being-H0.5-specialist performs best on most categories, as expected from embodiment-specific adaptation. Notably, Being-H0.5-generalist is only marginally behind on spatial, long-horizon, and bimanual categories. In several overlap-heavy settings, the generalist can even be *better* than the specialist. Empirically, joint training increases exposure to shared sub-skills (approach–grasp–place, lid/handle interactions, planar wiping/clearing), which improves robustness to nuisance factors (layout shifts, partial occlusions, camera drift) without overfitting to a single embodiment.

Clear margin over $\pi_{0.5}$, especially on long-horizon and bimanual tasks. Both Being-H0.5 variants significantly outperform $\pi_{0.5}$. The gap is most pronounced for long-horizon and bimanual categories, where small perception/control mismatches compound. A key difference is that Being-H0.5 is trained and deployed in a unified state–action space with an explicit cross-embodiment design, whereas $\pi_{0.5}$ in our setting remains bound to embodiment-specific action interfaces and thus cannot exploit joint training.

#1. Arrange Flower: Arrange the flowers neatly into the vase.



#2. Scan Package: Flip the package then scan the barcode with the scanner.



#3. Water Plant: Use your hand to hold the sprinkling can and water the plants.



#4. Stack Blocks: Stack the blocks following the color order: red, blue, and yellow.



#5. Clean Table: Clear the table: put all objects into the plate.



Figure 9: Example rollouts of representative real-robot task suites. (Only show one task for each embodiment. The full rollouts of all tasks are shown in Appendix B).

UniHand-2.0 pretraining is crucial for a strong generalist. The UniHand-2.0 ablation remains non-trivial in the *specialist* regime (still able to learn from real-robot finetuning), highlighting that our architecture and post-training pipeline provide a strong inductive bias. However, the *generalist* ablation degrades substantially and fails to match its specialist counterpart. This contrast indicates that large-scale cross-embodiment human-centric pretraining is not just “extra data”: it supplies broad manipulation priors that stabilize optimization under heterogeneous embodiment mixtures and enables positive transfer across platforms.

Embodiment-Level Zero-Shot Transfer. Beyond in-distribution evaluation, we observe a notable embodiment-level zero-shot transfer phenomenon. Specifically, after cross-embodiment post-training, a single **Being-H0.5-generalist** checkpoint can be deployed on **Adam-U** to solve *previously unseen task-embodiment pairs*—*i.e.*, tasks for which we collect **no Adam-U demonstrations at all**. Concretely, in environments resembling tasks demonstrated only on other embodiments, Adam-U achieves a *non-zero* success rate and consistently exhibits task-consistent multi-step execution. Examples include: (i) *flip-and-scan* behaviors reminiscent of the G1 package pipeline (reorienting the parcel to expose the label region and bringing the scanner toward it), (ii) *drawer interaction* behaviors aligned with the FR3 long-horizon task (reaching for handle-like affordances, pulling to open, and attempting to place objects before closing), and (iii) *stacking* behaviors resembling the D1 block task (sequential pick-and-place with an explicit stacking intent). While the zero-shot success rate remains lower than in-distribution performance and the motion precision is less reliable, the policy demonstrates a clear ability to *instantiate task structure learned on other morphologies* under a

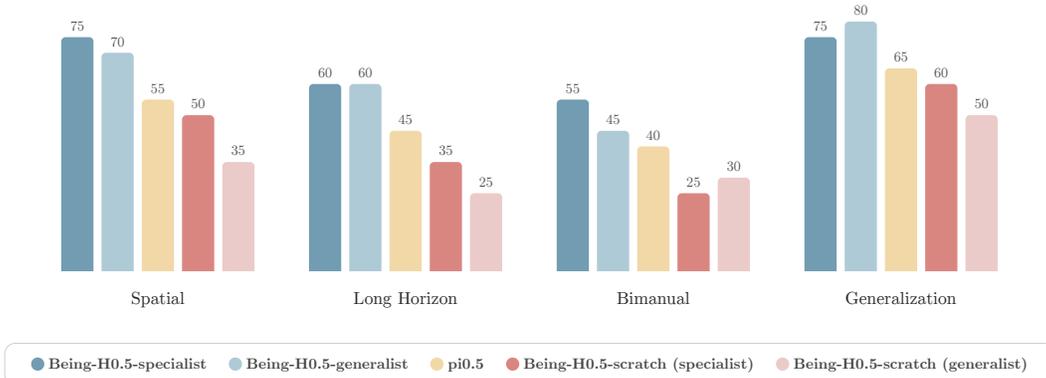


Figure 10: **Real-robot success rates (%) on four task suites.** We compare Being-H0.5 specialist/generalist, scratch ablations (both specialist and generalist), and $\pi_{0.5}$.

new kinematic chain.

We attribute this capability to two design choices emphasized throughout the paper: (1) a **unified action space** that encourages transferable motor primitives rather than embodiment-specific action conventions, and (2) **cross-embodiment joint post-training** that exposes the model to diverse realizations of the same underlying task abstractions (objects, goals, and subgoal compositions). To our knowledge, this is the first empirical evidence that a VLA policy can achieve embodiment-level *zero-shot* task completion (albeit at low success rates) in real-robot deployment. We believe this result sheds light on a promising scaling direction for VLAs. In particular, it suggests that increasing the diversity and coverage of post-training data, together with a unified action interface, can pave the way for cross-embodiment policies that generalize compositionally beyond the limited task-embodiment pairs and progressively acquire more transferable, emergent intelligence.

7.2 Comparison on Simulation Benchmarks

Table 4: **Success rates (%) on the LIBERO benchmark.** Being-H0.5 is evaluated against a wide range of state-of-the-art VLA models. We report the mean success rate over 50 evaluation episodes per task. Being-H0.5 achieves leading performance on the overall average, with a significant advantage in LIBERO-Long, which requires complex multi-step reasoning. Best results are **bolded**, and second-best are underlined.

Method	L-Spatial	L-Object	L-Goal	L-Long	Average
Diffusion Policy [20]	78.5	87.5	73.5	64.8	76.1
OpenVLA [6]	84.7	88.4	79.2	53.7	76.5
SpatialVLA [28]	88.2	89.9	78.6	55.5	78.1
CoT-VLA [132]	87.5	91.6	87.6	69.0	83.9
π_0 -Fast [133]	96.4	96.8	88.6	60.2	85.5
GR00T-N1 [9]	94.4	97.6	93.0	90.6	93.9
π_0 [8]	98.0	96.8	94.4	88.4	94.4
F1 [134]	98.2	97.8	95.4	91.3	95.7
InternVLA-M1 [54]	98.0	99.0	93.8	92.6	95.9
Discrete Diffusion VLA [31]	97.2	98.6	97.4	92.0	96.3
$\pi_{0.5}$ [34]	98.8	98.2	98.0	92.4	96.9
OpenVLA-OFT [91]	97.6	98.4	97.9	94.5	97.1
X-VLA [11]	98.2	98.6	97.8	97.6	98.1
EO1 [94]	99.7	99.8	<u>99.2</u>	94.8	<u>98.2</u>
Being-H0.5 (generalist)	97.0	98.2	99.0	96.2	97.6
Being-H0.5 (specialist)	<u>99.2</u>	<u>99.6</u>	99.4	<u>97.4</u>	98.9

We evaluate **Being-H0.5** on two widely-used simulation benchmarks, LIBERO [135] and RoboCasa [136],

and compare against recent VLA systems and specialized robotic policies. Unless otherwise specified, Being-H0.5 uses **RGB-only** observations at 224×224 resolution with a **2B** backbone. **Specialist vs. Generalist.** Throughout this section, we similarly report two training regimes: (i) **specialist**, which trains a benchmark-specific checkpoint using data from *only one* benchmark (LIBERO-only or RoboCasa-only); and (ii) **generalist**, which trains a *single* checkpoint jointly on LIBERO + RoboCasa. To ensure a fair comparison in terms of optimization, the generalist regime runs for $2 \times$ the training steps of a specialist model, so that the model sees a comparable number of update steps *per benchmark* while learning a unified policy across both testbeds. At evaluation time, the same generalist checkpoint is directly evaluated on each benchmark without any additional adaptation.

7.2.1 Results on LIBERO

Experimental setup. We follow the standard LIBERO evaluation protocol [135, 137]. The policy takes multi-view RGB observations (wrist-mounted and third-person cameras), each resized to 224×224 . We use action chunking with chunk size 8 and train with packed sequences (7,680 tokens per GPU) for an effective batch size of 128. Optimization runs for 45k steps on $4 \times$ A800 GPUs for the specialist model. For the generalist model, we train for approximately $2 \times$ steps under the joint LIBERO+RoboCasa mixture. For evaluation, we run 500 trials per suite.

Results and analysis. Table 4 shows that Being-H0.5(**specialist**) reaches a **98.9%** average success rate on LIBERO, with strong performance across all suites including LIBERO-Long (97.4%). Being-H0.5(**generalist**), trained jointly on LIBERO+RoboCasa with $2 \times$ steps, remains highly competitive at **97.6%** average. The modest gap indicates that a single checkpoint can absorb substantial cross-benchmark diversity without sacrificing much LIBERO performance, even though the generalist must allocate capacity to cover RoboCasa’s long-horizon kitchen interactions in addition to LIBERO’s task distribution. Crucially, both results are achieved using only 224×224 RGB inputs and a **2B** backbone, matching or outperforming many higher-resolution, larger, or multi-modal systems.

Table 5: **Success rates (%) on the RoboCasa 24-task benchmark.** Being-H0.5 is evaluated across three task categories, averaged over 50 trials per task. Despite using only 224×224 RGB inputs, Being-H0.5 outperforms state-of-the-art 3D-based methods and existing VLAs on the overall average. Best results are **bolded**, and second-best are underlined.

Modality	Method	Pick & Place	Doors/Drawers	Others	Total Avg.
<i>3D (Multi-modal)</i>					
3D	3DA [138]	0.0	2.3	13.1	5.5
3D	DP3 [139]	1.5	41.7	32.0	22.8
3D	GWM [140]	14.8	54.3	49.8	39.3
<i>RGB-only (Standard)</i>					
RGB (256×256)	BC [136]	4.3	47.0	42.2	28.9
RGB (256×256)	GR00T-N1 [9]	18.6	50.2	39.1	36.0
RGB (256×256)	$\pi_{0.5}$ [34]	21.5	57.8	44.9	41.4
RGB (256×256)	π_0 [8]	14.0	53.1	58.5	42.4
RGB (224×224)	Being-H0.5 (generalist)	40	73	52	53.3
RGB (224×224)	Being-H0.5 (specialist)	<u>36</u>	<u>71.7</u>	<u>57.6</u>	53.9

7.2.2 Results on RoboCasa

Experimental setup. We evaluate on RoboCasa [136], which contains 24 long-horizon household tasks in diverse kitchen environments. We adopt the challenging **Human-50** few-shot setting, using 50 human demonstrations per task. Evaluation is performed with 50 trials per task across five held-out scenes, featuring unseen object instances and unseen kitchen styles. We use **RGB-only** inputs at 224×224 and natural-language task commands, without depth or point clouds. The specialist model is trained for the same budget as above (RoboCasa-only), while the generalist is trained jointly on LIBERO+RoboCasa for $2 \times$ steps.

Table 6: **Impact of human-centric learning in single-task 5-shot adaptation.** We investigate this by selectively freezing key architectural components. **Native VLM:** Initialized from the original large multimodal model. **Human-Centric:** Initialized with our human-centric pretrained weights. The action expert remains trainable in all settings. Δ denotes the absolute performance gain.

Model Configuration		Single-Task LIBERO 5-Shot Success Rate (%)				
Frozen Components	Initialization	L-Spatial	L-Object	L-Goal	L-Long	Avg.
Und + Proj + ViT	Native VLM	67.8	76.8	57.2	29.8	57.9
	Human-Centric	78.6	78.9	72.6	46.2	69.0
	Gain (Δ)	+10.8	+2.1	+15.4	+16.4	+11.1
Und + ViT	Native VLM	70.6	50.0	67.0	17.4	51.3
	Human-Centric	85.0	84.2	80.0	59.0	77.1
	Gain (Δ)	+14.4	+34.4	+13.0	+41.6	+25.8
Und Only	Native VLM	78.4	84.8	62.0	49.2	68.6
	Human-Centric	83.2	79.6	74.0	63.6	75.1
	Gain (Δ)	+4.8	-5.2	+12.0	+14.4	+6.5
Proj + ViT	Native VLM	87.2	86.6	80.8	58.0	78.2
	Human-Centric	86.2	86.2	84.6	66.8	81.0
	Gain (Δ)	-1.0	-0.4	+3.8	+8.8	+2.8
None (Full FT)	Native VLM	84.2	83.6	75.0	65.8	77.2
	Human-Centric	87.2	91.4	81.6	66.8	81.8
	Gain (Δ)	+3.0	+7.8	+6.6	+1.0	+4.6

Results and analysis. As shown in Table 5, Being-H0.5 achieves a new state of the art on RoboCasa: **53.9%** success for the specialist checkpoint and **53.3%** for the generalist single-checkpoint model. Notably, the generalist preserves essentially the same performance while simultaneously maintaining strong LIBERO performance, suggesting that joint training does not introduce severe negative transfer between these benchmarks under our unified action interface. The largest gain appears in Pick & Place, where Being-H0.5 reaches 36–40% success, substantially exceeding prior VLAs (e.g., $\pi_{0.5}$ at 21.5%) and even outperforming multi-modal baselines that leverage explicit 3D inputs. This is particularly notable because RoboCasa pick-and-place demands reliable spatial reasoning and consistent grasp–transport–place coordination under appearance variation. Our results indicate that large-scale human-centric pretraining can distill transferable spatial priors that remain effective even under RGB-only inputs and reduced resolution.

Takeaway. Across both benchmarks, the specialist regime provides an upper bound when benchmark-specific adaptation is allowed, while the generalist regime shows that a **single** Being-H0.5 checkpoint can jointly cover LIBERO and RoboCasa with only a small degradation on LIBERO and near-identical performance on RoboCasa. This points to a practical path toward scalable deployment where maintaining many benchmark-specific checkpoints is undesirable.

7.3 Ablation Study

7.3.1 How Do Human-Centric Learning Benefit Downstream Adaptation?

Our human-centric learning provides foundational priors that facilitate efficient adaptation to downstream robotic tasks. In this section, we analyze the specific impact of these pretrained weights on the adaptation process. To mitigate the risk of overfitting, which is a prevalent challenge in robotic benchmarks, and to rigorously evaluate few-shot generalization, we adopt the LIBERO 5-shot benchmark with a restricted 10K-step training window. By utilizing a minimal set of training samples, we can effectively isolate performance gains directly attributable to pretrained knowledge versus task-specific rote learning. Our validation follows a dual-protocol approach: **1) Single-task-suite:** Training separate models for each LIBERO suite to observe pretrained effectiveness in singular task scenarios. **2) Multi-task-suite:** Training a unified model across all

Table 7: **Impact ablation of human-centric learning in multi-task 5-shot adaptation.** In contrast to the single-task setup, the model is jointly trained on data from all four task suites and subsequently evaluated on each individually. **Native VLM:** Initialized from the original large multimodal model. **Human-Centric:** Initialized with our human-centric pretrained weights. This multi-task protocol serves to rigorously assess the generalization capabilities and robustness of VLAs under diverse cross-task scenarios.

Model Configuration		Multi-task-suite LIBERO 5-Shot Success Rate (%)				
Frozen Components	Initialization	L-Spatial	L-Object	L-Goal	L-Long	Avg.
Und + Proj + ViT	Native VLM	61.6	77.0	68.0	36.0	60.7
	Human-Centric	80.8	81.8	74.4	52.4	72.4
	Gain (Δ)	+19.2	+4.8	+6.4	+16.4	+11.7
Und + ViT	Native VLM	78.4	84.8	71.4	38.6	68.3
	Human-Centric	83.6	82.6	80.8	59.2	76.6
	Gain (Δ)	+5.2	-2.2	+9.4	+20.6	+8.3
Und Only	Native VLM	75.8	84.2	71.2	59.8	72.8
	Human-Centric	85.2	90.6	84.2	68.4	82.1
	Gain (Δ)	+9.4	+6.4	+13.0	+8.6	+9.3
Proj + ViT	Native VLM	87.6	91.0	82.2	62.6	80.9
	Human-Centric	89.6	88.8	89.6	65.8	83.5
	Gain (Δ)	+2.0	-2.2	+7.4	+3.2	+2.6
None (Full FT)	Native VLM	87.8	94.6	82.6	71.4	84.1
	Human-Centric	88.8	90.2	85.4	76.0	85.1
	Gain (Δ)	+1.0	-4.4	+2.8	+4.6	+1.0

suites to assess the capability of the pretrained VLA to adapt to large-scale, diverse task distributions. The results are detailed in Table 6 and Table 7. In these tables, we compare the performance of models initialized with our human-centric weights against a baseline initialized directly from the native VLM. Note that the Action Expert remains trainable in all experimental configurations. Additionally, Figures 11 and 12 further explore the sensitivity of the Action Expert to parameter freezing.

Impact of Backbone Pretraining. The results reveal critical insights into how weight initializations influence VLA performance across different paradigms. First, the impact of pretraining is most pronounced under maximum parameter constraints. As shown in Table 6, when the core reasoning components (Und + ViT) are frozen while allowing the projector and action expert to train, the pretrained model achieves a remarkable **+25.8%** average gain over the baseline.

A consistent trend across both single- and multi-task setups is that LIBERO-Long tasks benefit disproportionately from pretraining. In the single-task setup, freezing the Und + ViT components yields a staggering +41.6% improvement in the Long suite. This indicates that while simple spatial or object-oriented tasks can be learned from scratch with 50 demonstrations, the temporal consistency and intent-grounding required for long-horizon manipulation are almost entirely derived from the foundational “world knowledge” embedded in the pretrained VLM. Interestingly, the marginal utility of pretraining (Δ) decreases as the model gains more trainable parameters (moving toward Full FT). In some cases, such as L-Object, we observe slight negative transfer (e.g., -4.4% in Multi-task Full FT). This suggests that for simpler, object-centric tasks, aggressive joint fine-tuning may lead the model to “overwrite” specialized pretrained features in favor of task-specific heuristics. Our ablation confirms that the Language (Und) and Visual (ViT) backbones act as the primary carriers of generalizable robotic proficiency, providing a stable “intent manifold” that prevents overfitting when downstream data is sparse.

Sensitivity of Action Expert Plasticity. While the backbone benefits from stability, the Action Expert requires plasticity. We investigate the impact of freezing layers within the Action Expert, ranging from the

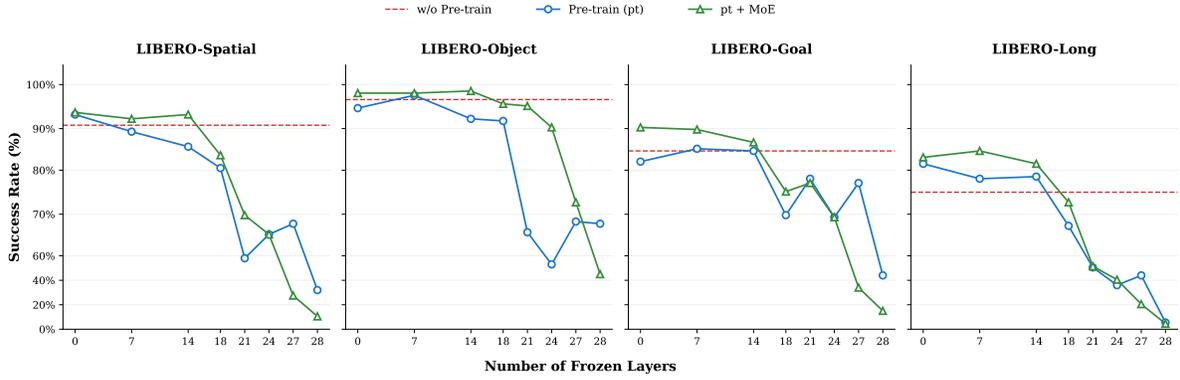


Figure 11: **Ablation study on the number of frozen layers in the action expert.** Results compare the performance of the model after human-centric pretraining (“pt”) and without it (“w/o pt”). “MoF” denotes our proposed Mixture of Flow architecture. All experiments freeze all parameters of the MLLM including the visual encoder and projector. Note that each task was evaluated with 20 rollouts due to time constraints.

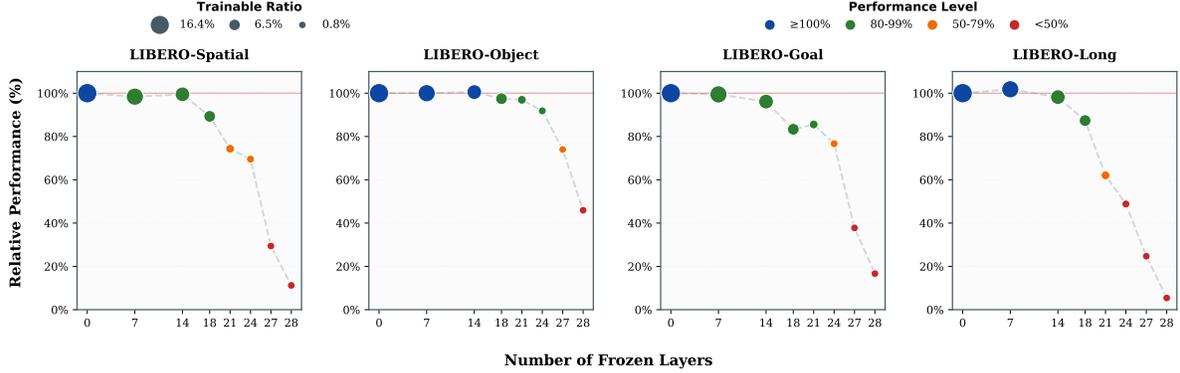


Figure 12: **Ablation Study on the Number of Frozen Layers in the Action Expert.** Results compare the performance of the model after human-centric pretraining (“pt”) and without it (“w/o pt”). “MoF” denotes our proposed Mixture of Flow architecture. All experiments freeze all parameters of the MLLM, including the visual encoder and projector. Note that each task was evaluated with 20 rollouts due to time constraints.

input to the output heads, as visualized in Figure 11. The results reveal that freezing the initial 0 to 7 layers results in negligible performance degradation, as the model sustains robust success rates exceeding 80% across the majority of task suites. However, as the number of frozen layers exceeds 14, performance degrades sharply, dropping below 20% when the entire expert is frozen. Figure 12 highlights this trade-off between trainable parameters and performance relative to the peak. This indicates that while the high-level semantic planning is transferable from human-centric pretraining, the low-level motor control—mapping features to precise joint velocities—must be actively learned from downstream demonstrations. Therefore, the optimal configuration for low-data adaptation is to freeze the pretrained semantic backbone (Und + ViT) while keeping the Action Expert fully trainable.

7.3.2 How Does Masked Motion Token Prediction Benefit Capturing Behavior Priors?

We further verify the benefit of introducing the discrete *Masked Motion Token Prediction* objective in our pre-training. Specifically, we conduct an ablation that removes $\mathcal{L}_{\text{MASK}}$ while keeping all other settings unchanged, and evaluate motion generation on a held-out subset of our human demonstration data. We split the evaluation set by data source into two domains: (i) **lab-curated** demonstrations collected under

controlled setups, and (ii) **in-the-wild** videos featuring diverse scenes and heterogeneous motion patterns.

Notably, since our unified action space adopts a *delta* representation, typical position error metrics such as Mean Per Joint Position Error (MPJPE) become less informative: the initial pose usually dominates such errors, while the initial state in our setting is provided as ground truth. Therefore, we instead measure whether the model captures the *intent* in language and grounds it in the image geometry using **Mean Wrist Displacement Similarity (MWDS)**. Concretely, MWDS averages the cosine similarity between the predicted and ground-truth wrist displacement vectors from the initial to the final state, $\cos(\Delta\hat{\mathbf{p}}_{\text{wrist}}, \Delta\mathbf{p}_{\text{wrist}})$.

Table 8: **Ablation on masked motion token prediction.** MWDS \uparrow on a held-out human demonstration test set with **Lab** (lab-curated) and **Wild** (in-the-wild) splits.

Method	Lab \uparrow	Wild \uparrow
Hybrid (Ours)	0.33	0.20
w/o $\mathcal{L}_{\text{MASK}}$	0.35	0.28

As shown in Table 8, removing masked prediction leads to a clear drop in MWDS on both the lab-curated and in-the-wild splits. This indicates that the discrete masked objective injects a stable, abstraction-level behavior prior that improves intention-aligned generation and enhances robustness across heterogeneous motion patterns. Notably, the improvement is more pronounced on the *Wild* split, where demonstrations contain substantially higher behaviour variability and observation noise. This suggests that the abstraction-level masked objective particularly matters for noisy, large-scale human-centric pre-training.

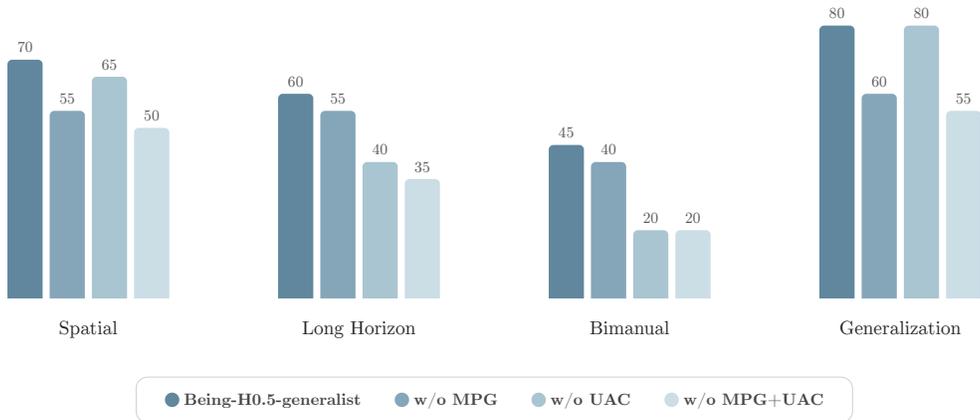


Figure 13: **Ablation of MPG+UAC on real robots.** Removing MPG+UAC hurts long-horizon and bimanual categories the most, where execution delay and unreliable context amplify compounding errors.

7.3.3 Real-Time Efficiency: MPG and UAC

We further ablate the deployment-time components by disabling both **MPG** (manifold-preserving gating) and **UAC** (universal async chunking), while keeping the backbone and training data fixed. Figure 13 shows consistent degradations.

Long-horizon drops sharply without UAC-style async consistency. Long-horizon tasks involve multi-stage behaviors with execution delays and asynchronous sensing/actuation. Without UAC, the policy becomes more sensitive to temporal mismatch between perceived state and the committed portion of the action chunk, leading to error accumulation (e.g., drifting grasps, premature contact, failed handle engagement) and a noticeable drop in success.

Bimanual coordination is brittle without stability mechanisms. Bimanual tasks require tight temporal coupling between two effectors under partial observability and contact. Disabling MPG+UAC increases oscillation and indecisive corrections when the observation context is noisy or changes rapidly, especially for platforms with movable cameras where viewpoint shifts can invalidate short-term visual cues.

Spatial/generalization are affected but less catastrophically. Spatial tasks still rely on precise placement, yet the execution horizon is shorter, and there are fewer compounding subgoals. Generalization tasks benefit from robustness mechanisms when the scene shifts, but typically allow more tolerance and recovery, hence the smaller (yet consistent) degradation.

8 Conclusion

In this work, we present Being-H0.5, a foundational Vision-Language-Action (VLA) model that establishes a scalable, human-centric paradigm for cross-embodiment robot learning. By unifying human and robotic motion within a physically grounded state-action space and leveraging the large-scale UniHand-2.0 dataset, we demonstrate that human interaction traces can serve as dense physical priors to bridge the gap between diverse robotic morphologies. Our Mixture-of-Transformers (MoT) architecture effectively disentangles high-level multimodal reasoning from low-level motor execution, while Manifold-Preserving Gating (MPG) and Universal Async Chunking (UAC) ensure robust deployment by mitigating contextual distribution shifts and hardware-specific inference latencies. Extensive evaluation across a heterogeneous fleet of five robotic platforms proves that a single Being-H0.5 checkpoint can internalize complex interaction dynamics, achieving high success rates in bimanual coordination and long-horizon tasks. Ultimately, Being-H0.5 signifies a transition toward versatile, general-purpose robotic intelligence capable of seamless operation across various environments and embodiments.

References

- [1] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186, 2019.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [3] Julian Eisenschlos, Sebastian Ruder, Piotr Czapla, Marcin Kadras, Sylvain Gugger, and Jeremy Howard. Multifit: Efficient multi-lingual language model fine-tuning. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 5702–5707, 2019.
- [4] Telmo Pires, Eva Schlinger, and Dan Garrette. How multilingual is multilingual bert? *arXiv preprint arXiv:1906.01502*, 2019.
- [5] Alexis Conneau and Guillaume Lample. Cross-lingual language model pretraining. *Advances in neural information processing systems*, 32, 2019.
- [6] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- [7] 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.
- [8] 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.
- [9] J Bjorck Nvidia, Fernando Castaneda, N Cherniadev, X Da, R Ding, L Fan, Y Fang, D Fox, F Hu, S Huang, et al. Gr00t n1: An open foundation model for generalist humanoid robots. *arXiv preprint arXiv:2503.14734*, 2025.
- [10] Hao Luo, Yicheng Feng, Wanpeng Zhang, Sipeng Zheng, Ye Wang, Haoqi Yuan, Jiazheng Liu, Chaoyi Xu, Qin Jin, and Zongqing Lu. Being-h0: vision-language-action pretraining from large-scale human videos. *arXiv preprint arXiv:2507.15597*, 2025.
- [11] Jinliang Zheng, Jianxiang 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.
- [12] N Shazeer, A Mirhoseini, K Maziarz, A Davis, Q Le, G Hinton, and J Dean. The sparsely-gated mixture-of-experts layer. *Outrageously large neural networks*, 2, 2017.
- [13] Lars Berscheid, Pascal Meißner, and Torsten Kröger. Robot learning of shifting objects for grasping in cluttered environments. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 612–618. IEEE, 2019.
- [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. *arXiv preprint arXiv:1910.11215*, 2019.
- [15] Hao-Shu Fang, Hongjie Fang, Zhenyu Tang, Jirong Liu, Chenxi Wang, Junbo Wang, Haoyi Zhu, and Cewu Lu. Rh20t: A comprehensive robotic dataset for learning diverse skills in one-shot. *arXiv preprint arXiv:2307.00595*, 2023.
- [16] Nur Muhammad Mahi Shafiullah, Anant Rai, Haritheja Etukuru, Yiqian Liu, Ishan Misra, Soumith Chintala, and Lerrel Pinto. On bringing robots home. *arXiv preprint arXiv:2311.16098*, 2023.

- [17] 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.
- [18] Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, pages 2165–2183. PMLR, 2023.
- [19] Michael Janner, Yilun Du, Joshua B Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022.
- [20] 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.
- [21] Andreas Steiner, André Susano Pinto, Michael Tschannen, Daniel Keysers, Xiao Wang, Yonatan Bitton, Alexey Gritsenko, Matthias Minderer, Anthony Sherbondy, Shangbang Long, et al. Paligemma 2: A family of versatile vlms for transfer. *arXiv preprint arXiv:2412.03555*, 2024.
- [22] Zhiqi Li, Guo Chen, Shilong Liu, Shihao Wang, Vibashan VS, Yishen Ji, Shiyi Lan, Hao Zhang, Yilin Zhao, Subhashree Radhakrishnan, et al. Eagle 2: Building post-training data strategies from scratch for frontier vision-language models. *arXiv preprint arXiv:2501.14818*, 2025.
- [23] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.
- [24] Wanpeng Zhang, Zilong Xie, Yicheng Feng, Yijiang Li, Xingrun Xing, Sipeng Zheng, and Zongqing Lu. From pixels to tokens: Byte-pair encoding on quantized visual modalities. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [25] Wanpeng Zhang, Yicheng Feng, Hao Luo, Yijiang Li, Zihao Yue, Sipeng Zheng, and Zongqing Lu. Unified multimodal understanding via byte-pair visual encoding. *arXiv preprint arXiv:2506.23639*, 2025.
- [26] Luo Hao, Yue Zihao, Zhang Wanpeng, Feng Yicheng, Zheng Sipeng, Ye Deheng, and Lu Zongqing. OpenMMego: Enhancing egocentric understanding for LMMs with open weights and data. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.
- [27] Yicheng Feng, Yijiang Li, Wanpeng Zhang, Sipeng Zheng, Hao Luo, Zihao Yue, and Zongqing Lu. Videorion: Tokenizing object dynamics in videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20401–20412, 2025.
- [28] Delin Qu, Haoming Song, Qizhi Chen, Yuanqi Yao, Xinyi Ye, Yan Ding, Zhigang Wang, JiaYuan Gu, Bin Zhao, Dong Wang, et al. Spatialvla: Exploring spatial representations for visual-language-action model. *arXiv preprint arXiv:2501.15830*, 2025.
- [29] Qixiu Li, Yaobo Liang, Zeyu Wang, Lin Luo, Xi Chen, Mozheng Liao, Fangyun Wei, Yu Deng, Sicheng Xu, Yizhong Zhang, et al. Cogact: A foundational vision-language-action model for synergizing cognition and action in robotic manipulation. *arXiv preprint arXiv:2411.19650*, 2024.
- [30] 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. *arXiv preprint arXiv:2410.07864*, 2024.
- [31] Zhixuan Liang, Yizhuo Li, Tianshuo Yang, Chengyue Wu, Sitong Mao, Tian Nian, Liua Pei, Shunbo Zhou, Xiaokang Yang, Jiangmiao Pang, et al. Discrete diffusion vla: Bringing discrete diffusion to action decoding in vision-language-action policies. *arXiv preprint arXiv:2508.20072*, 2025.
- [32] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [33] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- [34] 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.

- [35] Junjie Wen, Yichen Zhu, Jinming Li, Zhibin Tang, Chaomin Shen, and Feifei Feng. Dexvla: Vision-language model with plug-in diffusion expert for general robot control. *arXiv preprint arXiv:2502.05855*, 2025.
- [36] Yifan Zhong, Xuchuan Huang, Ruochong Li, Ceyao Zhang, Zhang Chen, Tianrui Guan, Fanlian Zeng, Ka Num Lui, Yuyao Ye, Yitao Liang, et al. Dexgraspvla: A vision-language-action framework towards general dexterous grasping. *arXiv preprint arXiv:2502.20900*, 2025.
- [37] Michał Zawalski, William Chen, Karl Pertsch, Oier Mees, Chelsea Finn, and Sergey Levine. Robotic control via embodied chain-of-thought reasoning. *arXiv preprint arXiv:2407.08693*, 2024.
- [38] Fanqi Lin, Ruiqian Nai, Yingdong Hu, Jiacheng You, Junming Zhao, and Yang Gao. Onetwovla: A unified vision-language-action model with adaptive reasoning. *arXiv preprint arXiv:2505.11917*, 2025.
- [39] Jaden Clark, Suvir Mirchandani, Dorsa Sadigh, and Suneel Belkhal. Action-free reasoning for policy generalization. *arXiv preprint arXiv:2502.03729*, 2025.
- [40] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- [41] Brent Griffin. Mobile robot manipulation using pure object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 561–571, 2023.
- [42] Michael Laskin, Aravind Srinivas, and Pieter Abbeel. Curl: Contrastive unsupervised representations for reinforcement learning. In *International conference on machine learning*, pages 5639–5650. PMLR, 2020.
- [43] Andreas Ten Pas and Robert Platt. Using geometry to detect grasp poses in 3d point clouds. In *Robotics Research: Volume 1*, pages 307–324. Springer, 2017.
- [44] Jason Lee, Jiafei Duan, Haoquan Fang, Yuquan Deng, Shuo Liu, Boyang Li, Bohan Fang, Jieyu Zhang, Yi Ru Wang, Sangho Lee, et al. Molmoact: Action reasoning models that can reason in space. *arXiv preprint arXiv:2508.07917*, 2025.
- [45] 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. *arXiv preprint arXiv:2403.12945*, 2024.
- [46] 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.
- [47] Mustafa Shukor, Dana Aubakirova, Francesco Capuano, Pepijn Kooijmans, Steven Palma, Adil Zouitine, Michel Aractingi, Caroline Pascal, Martino Russi, Andres Marafioti, et al. Smolvla: A vision-language-action model for affordable and efficient robotics. *arXiv preprint arXiv:2506.01844*, 2025.
- [48] Qingwen Bu, Jisong Cai, Li Chen, Xiuqi Cui, Yan Ding, Siyuan Feng, Shenyuan Gao, Xindong He, Xuan Hu, Xu Huang, et al. Agibot world colosse: A large-scale manipulation platform for scalable and intelligent embodied systems. *arXiv preprint arXiv:2503.06669*, 2025.
- [49] Kun Wu, Chengkai Hou, Jiaming Liu, Zhengping Che, Xiaozhu Ju, Zhuqin Yang, Meng Li, Yinyu Zhao, Zhiyuan Xu, Guang Yang, et al. Robomind: Benchmark on multi-embodiment intelligence normative data for robot manipulation. *arXiv preprint arXiv:2412.13877*, 2024.
- [50] Chengkai Hou, Kun Wu, Jiaming Liu, Zhengping Che, Di Wu, Fei Liao, Guangrun Li, Jingyang He, Qiuxuan Feng, Zhao Jin, et al. Robomind 2.0: A multimodal, bimanual mobile manipulation dataset for generalizable embodied intelligence. *arXiv preprint arXiv:2512.24653*, 2025.
- [51] Shihan Wu, Xuecheng Liu, Shaoxuan Xie, Pengwei Wang, Xinghang Li, Bowen Yang, Zhe Li, Kai Zhu, Hongyu Wu, Yiheng Liu, et al. Robocoin: An open-sourced bimanual robotic data collection for integrated manipulation. *arXiv preprint arXiv:2511.17441*, 2025.
- [52] Tao Jiang, Tianyuan Yuan, Yicheng Liu, Chenhao Lu, Jianning Cui, Xiao Liu, Shuiqi Cheng, Jiyang Gao, Huazhe Xu, and Hang Zhao. Galaxea open-world dataset and g0 dual-system vla model. *arXiv preprint arXiv:2509.00576*, 2025.

- [53] Jianglong Ye, Keyi Wang, Chengjing Yuan, Ruihan Yang, Yiquan Li, Jiyue Zhu, Yuzhe Qin, Xueyan Zou, and Xiaolong Wang. Dex1b: Learning with 1b demonstrations for dexterous manipulation. *arXiv preprint arXiv:2506.17198*, 2025.
- [54] Xinyi Chen, Yilun Chen, Yanwei Fu, Ning Gao, Jiaya Jia, Weiyang Jin, Hao Li, Yao Mu, Jiangmiao Pang, Yu Qiao, et al. Internvla-m1: A spatially guided vision-language-action framework for generalist robot policy. *arXiv preprint arXiv:2510.13778*, 2025.
- [55] 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. *arXiv preprint arXiv:2402.10329*, 2024.
- [56] 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.
- [57] Huy Ha, Yihuai Gao, Zipeng Fu, Jie Tan, and Shuran Song. Umi on legs: Making manipulation policies mobile with manipulation-centric whole-body controllers. *arXiv preprint arXiv:2407.10353*, 2024.
- [58] Gen Robot. 10kh-realomin-opendata, 2025.
- [59] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019.
- [60] Kevin Qinghong Lin, Pengchuan Zhang, Joya Chen, Shraman Pramanick, Difei Gao, Alex Jinpeng Wang, Rui Yan, and Mike Zheng Shou. Univtg: Towards unified video-language temporal grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2794–2804, 2023.
- [61] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18995–19012, 2022.
- [62] Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, et al. Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19383–19400, 2024.
- [63] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European conference on computer vision (ECCV)*, pages 720–736, 2018.
- [64] Ryan Hoque, Peide Huang, David J Yoon, Mouli Sivapurapu, and Jian Zhang. Egodex: Learning dexterous manipulation from large-scale egocentric video. *arXiv preprint arXiv:2505.11709*, 2025.
- [65] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3m: A universal visual representation for robot manipulation. In *Conference on Robot Learning (CoRL)*, 2022.
- [66] Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. Vip: Towards universal visual reward and representation via value-implicit pre-training. In *The Eleventh International Conference on Learning Representations*, 2022.
- [67] Jiange Yang, Yansong Shi, Haoyi Zhu, Mingyu Liu, Kaijing Ma, Yating Wang, Gangshan Wu, Tong He, and Limin Wang. Como: Learning continuous latent motion from internet videos for scalable robot learning. *arXiv preprint arXiv:2505.17006*, 2025.
- [68] Xiaoyu Chen, Hangxing Wei, Pushi Zhang, Chuheng Zhang, Kaixin Wang, Yanjiang Guo, Rushuai Yang, Yucen Wang, Xinquan Xiao, Li Zhao, et al. Villa-x: enhancing latent action modeling in vision-language-action models. *arXiv preprint arXiv:2507.23682*, 2025.
- [69] Hao Luo, Ye Wang, Wanpeng Zhang, Haoqi Yuan, Yicheng Feng, Haiweng Xu, Sipeng Zheng, and Zongqing Lu. Predictive embedding as latent action: Towards vla pretraining in the wild. 2025.
- [70] Seonghyeon Ye, Joel Jang, Byeongguk Jeon, SeJune Joo, Jianwei Yang, Baolin Peng, Ajay Mandlekar, Reuben Tan, Yu-Wei Chao, Bill Yuchen Lin, et al. Latent action pretraining from videos. *arXiv preprint arXiv:2410.11758*, 2024.

- [71] Qingwen Bu, Yanting Yang, Jisong Cai, Shenyuan Gao, Guanghui Ren, Maoqing Yao, Ping Luo, and Hongyang Li. Univla: Learning to act anywhere with task-centric latent actions. *arXiv preprint arXiv:2505.06111*, 2025.
- [72] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- [73] Hongtao Wu, Ya Jing, Chilam Cheang, Guangzeng Chen, Jiafeng Xu, Xinghang Li, Minghuan Liu, Hang Li, and Tao Kong. Unleashing large-scale video generative pre-training for visual robot manipulation. *arXiv preprint arXiv:2312.13139*, 2023.
- [74] Chi-Lam Cheang, Guangzeng Chen, Ya Jing, Tao Kong, Hang Li, Yifeng Li, Yuxiao Liu, Hongtao Wu, Jiafeng Xu, Yichu Yang, et al. Gr-2: A generative video-language-action model with web-scale knowledge for robot manipulation. *arXiv preprint arXiv:2410.06158*, 2024.
- [75] Homanga Bharadhwaj, Debidatta Dwibedi, Abhinav Gupta, Shubham Tulsiani, Carl Doersch, Ted Xiao, Dhruv Shah, Fei Xia, Dorsa Sadigh, and Sean Kirmani. Gen2act: Human video generation in novel scenarios enables generalizable robot manipulation. *arXiv preprint arXiv:2409.16283*, 2024.
- [76] Hanzhi Chen, Boyang Sun, Anran Zhang, Marc Pollefeys, and Stefan Leutenegger. Vidbot: Learning generalizable 3d actions from in-the-wild 2d human videos for zero-shot robotic manipulation. In *CVPR*, 2025.
- [77] Teli Ma, Jia Zheng, Zifan Wang, Ziyao Gao, Jiaming Zhou, and Junwei Liang. Glover++: Unleashing the potential of affordance learning from human behaviors for robotic manipulation. *arXiv preprint arXiv:2505.11865*, 2025.
- [78] Alexey Gavryushin, Xi Wang, Robert JS Malate, Chenyu Yang, Xiangyi Jia, Shubh Goel, Davide Liconti, René Zurbügg, Robert K Katzschmann, and Marc Pollefeys. Maple: Encoding dexterous robotic manipulation priors learned from egocentric videos. *arXiv preprint arXiv:2504.06084*, 2025.
- [79] Shikhar Bahl, Abhinav Gupta, and Deepak Pathak. Human-to-robot imitation in the wild. *arXiv preprint arXiv:2207.09450*, 2022.
- [80] Shikhar Bahl, Russell Mendonca, Lili Chen, Unnat Jain, and Deepak Pathak. Affordances from human videos as a versatile representation for robotics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [81] Chuan Wen, Xingyu Lin, John So, Kai Chen, Qi Dou, Yang Gao, and Pieter Abbeel. Any-point trajectory modeling for policy learning. *arXiv preprint arXiv:2401.00025*, 2023.
- [82] Jianwei Yang, Reuben Tan, Qianhui Wu, Ruijie Zheng, Baolin Peng, Yongyuan Liang, Yu Gu, Mu Cai, Seonghyeon Ye, Joel Jang, et al. Magma: A foundation model for multimodal ai agents. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 14203–14214, 2025.
- [83] 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.
- [84] Simar Kareer, Dhruv Patel, Ryan Punamiya, Pranav Mathur, Shuo Cheng, Chen Wang, Judy Hoffman, and Danfei Xu. Egomimic: Scaling imitation learning via egocentric video. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13226–13233. IEEE, 2025.
- [85] Yicheng Feng, Wanpeng Zhang, Ye Wang, Hao Luo, Haoqi Yuan, Sipeng Zheng, and Zongqing Lu. Spatial-aware vla pretraining through visual-physical alignment from human videos. *arXiv preprint arXiv:2512.13080*, 2025.
- [86] Ryan Punamiya, Dhruv Patel, Patcharapong Aphiwetsa, Pranav Kuppili, Lawrence Y Zhu, Simar Kareer, Judy Hoffman, and Danfei Xu. Egobridge: Domain adaptation for generalizable imitation from egocentric human data. In *Human to Robot: Workshop on Sensorizing, Modeling, and Learning from Humans*, 2025.
- [87] Lawrence Y Zhu, Pranav Kuppili, Ryan Punamiya, Patcharapong Aphiwetsa, Dhruv Patel, Simar Kareer, Sehoon Ha, and Danfei Xu. Emma: Scaling mobile manipulation via egocentric human data. *arXiv preprint arXiv:2509.04443*, 2025.
- [88] Qixiu Li, Yu Deng, Yaobo Liang, Lin Luo, Lei Zhou, Chengtang Yao, Lingqi Zeng, Zhiyuan Feng, Huizhi Liang, Sicheng Xu, et al. Scalable vision-language-action model pretraining for robotic manipulation with real-life human activity videos. *arXiv preprint arXiv:2510.21571*, 2025.

- [89] Hongzhe Bi, Lingxuan Wu, Tianwei Lin, Hengkai Tan, Zhizhong Su, Hang Su, and Jun Zhu. H-rdt: Human manipulation enhanced bimanual robotic manipulation. *arXiv preprint arXiv:2507.23523*, 2025.
- [90] Ruihan Yang, Qinxi Yu, Yecheng Wu, Rui Yan, Borui Li, An-Chieh Cheng, Xueyan Zou, Yunhao Fang, Xuxin Cheng, Ri-Zhao Qiu, et al. Egovla: Learning vision-language-action models from egocentric human videos. *arXiv preprint arXiv:2507.12440*, 2025.
- [91] Moo Jin Kim, Chelsea Finn, and Percy Liang. Fine-tuning vision-language-action models: Optimizing speed and success. *arXiv preprint arXiv:2502.19645*, 2025.
- [92] LEJU Company. Let dataset. https://www.modelscope.cn/datasets/lejurobot/let_dataset, 2025.
- [93] Andy Zhai, Brae Liu, Bruno Fang, Chalse Cai, Ellie Ma, Ethan Yin, Hao Wang, Hugo Zhou, James Wang, Lights Shi, et al. Igniting vlms toward the embodied space. *arXiv preprint arXiv:2509.11766*, 2025.
- [94] Delin Qu, Haoming Song, Qizhi Chen, Zhaoqing Chen, Xianqiang Gao, Xinyi Ye, Qi Lv, Modi Shi, Guanghui Ren, Cheng Ruan, et al. Embodiedonevision: Interleaved vision-text-action pretraining for general robot control. *arXiv e-prints*, pages arXiv-2508, 2025.
- [95] Yankai Fu, Ning Chen, Junkai Zhao, Shaozhe Shan, Guocai Yao, Pengwei Wang, Zhongyuan Wang, and Shanghang Zhang. Metis: Multi-source egocentric training for integrated dexterous vision-language-action model. *arXiv preprint arXiv:2511.17366*, 2025.
- [96] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. The epic-kitchens dataset: Collection, challenges and baselines. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):4125–4141, 2020.
- [97] Build AI. Egocentric-10k, 2025.
- [98] Jinglei Zhang, Jiankang Deng, Chao Ma, and Rolandos Alexandros Potamias. Hawor: World-space hand motion reconstruction from egocentric videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1805–1815, 2025.
- [99] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *arXiv preprint arXiv:2201.02610*, 2022.
- [100] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- [101] Guanhua Ji, Harsha Polavaram, Lawrence Yunliang Chen, Sandeep Bajamahal, Zehan Ma, Simeon Adebola, Chenfeng Xu, and Ken Goldberg. Oxe-aug: A large-scale robot augmentation of oxe for scaling cross-embodiment policy learning. *arXiv preprint arXiv:2512.13100*, 2025.
- [102] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 26296–26306, 2024.
- [103] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024.
- [104] Luis Wiedmann, Orr Zohar, Amir Mahla, Xiaohan Wang, Rui Li, Thibaud Frere, Leandro von Werra, Aritra Roy Gosthipaty, and Andrés Marafioti. Finevision: Open data is all you need, 2025.
- [105] Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*, 2024.
- [106] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *European conference on computer vision*, pages 69–85. Springer, 2016.
- [107] Enshen Zhou, Jingkun An, Cheng Chi, Yi Han, Shanyu Rong, Chi Zhang, Pengwei Wang, Zhongyuan Wang, Tiejun Huang, Lu Sheng, et al. Roborefer: Towards spatial referring with reasoning in vision-language models for robotics. *arXiv preprint arXiv:2506.04308*, 2025.
- [108] Wentao Yuan, Jiafei Duan, Valts Blukis, Wilbert Pumacay, Ranjay Krishna, Adithyavairavan Murali, Arsalan Mousavian, and Dieter Fox. Robopoint: A vision-language model for spatial affordance prediction for robotics. *arXiv preprint arXiv:2406.10721*, 2024.

- [109] Yuheng Ji, Huajie Tan, Jiayu Shi, Xiaoshuai Hao, Yuan Zhang, Hengyuan Zhang, Pengwei Wang, Mengdi Zhao, Yao Mu, Pengju An, et al. Robobrain: A unified brain model for robotic manipulation from abstract to concrete. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1724–1734, 2025.
- [110] Yuhao Lu, Yixuan Fan, Beixing Deng, Fangfu Liu, Yali Li, and Shengjin Wang. VI-grasp: a 6-dof interactive grasp policy for language-oriented objects in cluttered indoor scenes. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [111] Pierre Sermanet, Tianli Ding, Jeffrey Zhao, Fei Xia, Debidatta Dwibedi, Keerthana Gopalakrishnan, Christine Chan, Gabriel Dulac-Arnold, Sharath Maddineni, Nikhil J Joshi, et al. Robovqa: Multimodal long-horizon reasoning for robotics. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 645–652. IEEE, 2024.
- [112] Rongtao Xu, Jian Zhang, Minghao Guo, Youpeng Wen, Haoting Yang, Min Lin, Jianzheng Huang, Zhe Li, Kaidong Zhang, Liqiong Wang, et al. A0: An affordance-aware hierarchical model for general robotic manipulation. *arXiv preprint arXiv:2504.12636*, 2025.
- [113] 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.
- [114] Weiyun Wang, Yiming Ren, Haowen Luo, Tiantong Li, Chenxiang Yan, Zhe Chen, Wenhai Wang, Qingyun Li, Lewei Lu, Xizhou Zhu, et al. The all-seeing project v2: Towards general relation comprehension of the open world. *arXiv preprint arXiv:2402.19474*, 2024.
- [115] Build AI. Egocentric-10k, 2025.
- [116] Toby Perrett, Ahmad Darkhalil, Saptarshi Sinha, Omar Emara, Sam Pollard, Kranti Parida, Kaiting Liu, Prajwal Gatti, Siddhant Bansal, Kevin Flanagan, Jacob Chalk, Zhifan Zhu, Rhodri Guerrier, Fahd Abdelazim, Bin Zhu, Davide Moltisanti, Michael Wray, Hazel Doughty, and Dima Damen. Hd-epic: A highly-detailed egocentric video dataset, 2025.
- [117] Yunze Liu, Yun Liu, Che Jiang, Kangbo Lyu, Weikang Wan, Hao Shen, Boqiang Liang, Zhoujie Fu, He Wang, and Li Yi. Hoi4d: A 4d egocentric dataset for category-level human-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21013–21022, 2022.
- [118] Edwin Olson. Apriltag: A robust and flexible visual fiducial system. In *2011 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3400–3407. IEEE, 2011.
- [119] Zhengdi Yu, Stefanos Zafeiriou, and Tolga Birdal. Dyn-hamr: Recovering 4d interacting hand motion from a dynamic camera. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 27716–27726, 2025.
- [120] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*, 2024.
- [121] Xiaowen Li, Haolan Xue, Peiran Ren, and Liefeng Bo. Diffueraser: A diffusion model for video inpainting. *arXiv preprint arXiv:2501.10018*, 2025.
- [122] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- [123] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, et al. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025.
- [124] Weixin Liang, Lili Yu, Liang Luo, Srinivasan Iyer, Ning Dong, Chunting Zhou, Gargi Ghosh, Mike Lewis, Wen-tau Yih, Luke Zettlemoyer, et al. Mixture-of-transformers: A sparse and scalable architecture for multi-modal foundation models. *arXiv preprint arXiv:2411.04996*, 2024.
- [125] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

- [126] Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, et al. Internvl3. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*, 2025.
- [127] Team Seedream, Yunpeng Chen, Yu Gao, Lixue Gong, Meng Guo, Qiushan Guo, Zhiyao Guo, Xiaoxia Hou, Weilin Huang, Yixuan Huang, et al. Seedream 4.0: Toward next-generation multimodal image generation. *arXiv preprint arXiv:2509.20427*, 2025.
- [128] Wanpeng Zhang, Ye Wang, Hao Luo, Haoqi Yuan, Yicheng Feng, Sipeng Zheng, Qin Jin, and Zongqing Lu. Dig-flow: Discrepancy-guided flow matching for robust vla models. *arXiv preprint arXiv:2512.01715*, 2025.
- [129] Nicolas Bonneel, Julien Rabin, Gabriel Peyré, and Hanspeter Pfister. Sliced and radon wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51(1):22–45, 2015.
- [130] Soheil Kolouri, Kimia Nadjahi, Umut Simsekli, Roland Badeau, and Gustavo Rohde. Generalized sliced wasserstein distances. *Advances in neural information processing systems*, 32, 2019.
- [131] Kevin Black, Allen Z Ren, Michael Equi, and Sergey Levine. Training-time action conditioning for efficient real-time chunking. *arXiv preprint arXiv:2512.05964*, 2025.
- [132] Qingqing Zhao, Yao Lu, Moo Jin Kim, Zipeng Fu, Zhuoyang Zhang, Yecheng Wu, Zhaoshuo Li, Qianli Ma, Song Han, Chelsea Finn, et al. Cot-vla: Visual chain-of-thought reasoning for vision-language-action models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1702–1713, 2025.
- [133] Karl Pertsch, Kyle Stachowicz, Brian Ichter, Danny Driess, Suraj Nair, Quan Vuong, Oier Mees, Chelsea Finn, and Sergey Levine. Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747*, 2025.
- [134] Qi Lv, Weijie Kong, Hao Li, Jia Zeng, Zherui Qiu, Delin Qu, Haoming Song, Qizhi Chen, Xiang Deng, and Jiangmiao Pang. F1: A vision-language-action model bridging understanding and generation to actions. *arXiv preprint arXiv:2509.06951*, 2025.
- [135] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems*, 36:44776–44791, 2023.
- [136] Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi, Ajay Mandlekar, and Yuke Zhu. Robocasa: Large-scale simulation of everyday tasks for generalist robots. *arXiv preprint arXiv:2406.02523*, 2024.
- [137] Moo Jin Kim, Chelsea Finn, and Percy Liang. Fine-tuning vision-language-action models: Optimizing speed and success. *arXiv preprint arXiv:2502.19645*, 2025.
- [138] Tsung-Wei Ke, Nikolaos Gkanatsios, and Katerina Fragkiadaki. 3d diffuser actor: Policy diffusion with 3d scene representations. *Arxiv*, 2024.
- [139] Yanjie Ze, Gu Zhang, Kangning Zhang, Chenyuan Hu, Muhan Wang, and Huazhe Xu. 3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024.
- [140] Guanxing Lu, Baoxiong Jia, Puhao Li, Yixin Chen, Ziwei Wang, Yansong Tang, and Siyuan Huang. Gwm: Towards scalable gaussian world models for robotic manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9263–9274, 2025.

Appendix

A Contributions

† denotes leading the part

Core Contributors (alphabetical order)

- HAO LUO
- YE WANG
- WANPENG ZHANG
- SIPENG ZHENG

Model Training

- HAO LUO
- YE WANG
- WANPENG ZHANG
- SIPENG ZHENG

Infrastructure

- WANPENG ZHANG[†]
- YE WANG
- CHAOYI XU
- CHI ZHANG
- HAOQI YUAN

Human Data

- HAO LUO[†]
- ZIHENG XI
- YIQING WANG
- YICHENG FENG

Robot Data

- YE WANG[†]
- SIPENG ZHENG[†]
- HAIWENG XU

Experiments

- WANPENG ZHANG
- YE WANG
- HAO LUO
- SIPENG ZHENG

Project Lead

- SIPENG ZHENG
- ZONGQING LU

B Full Rollouts

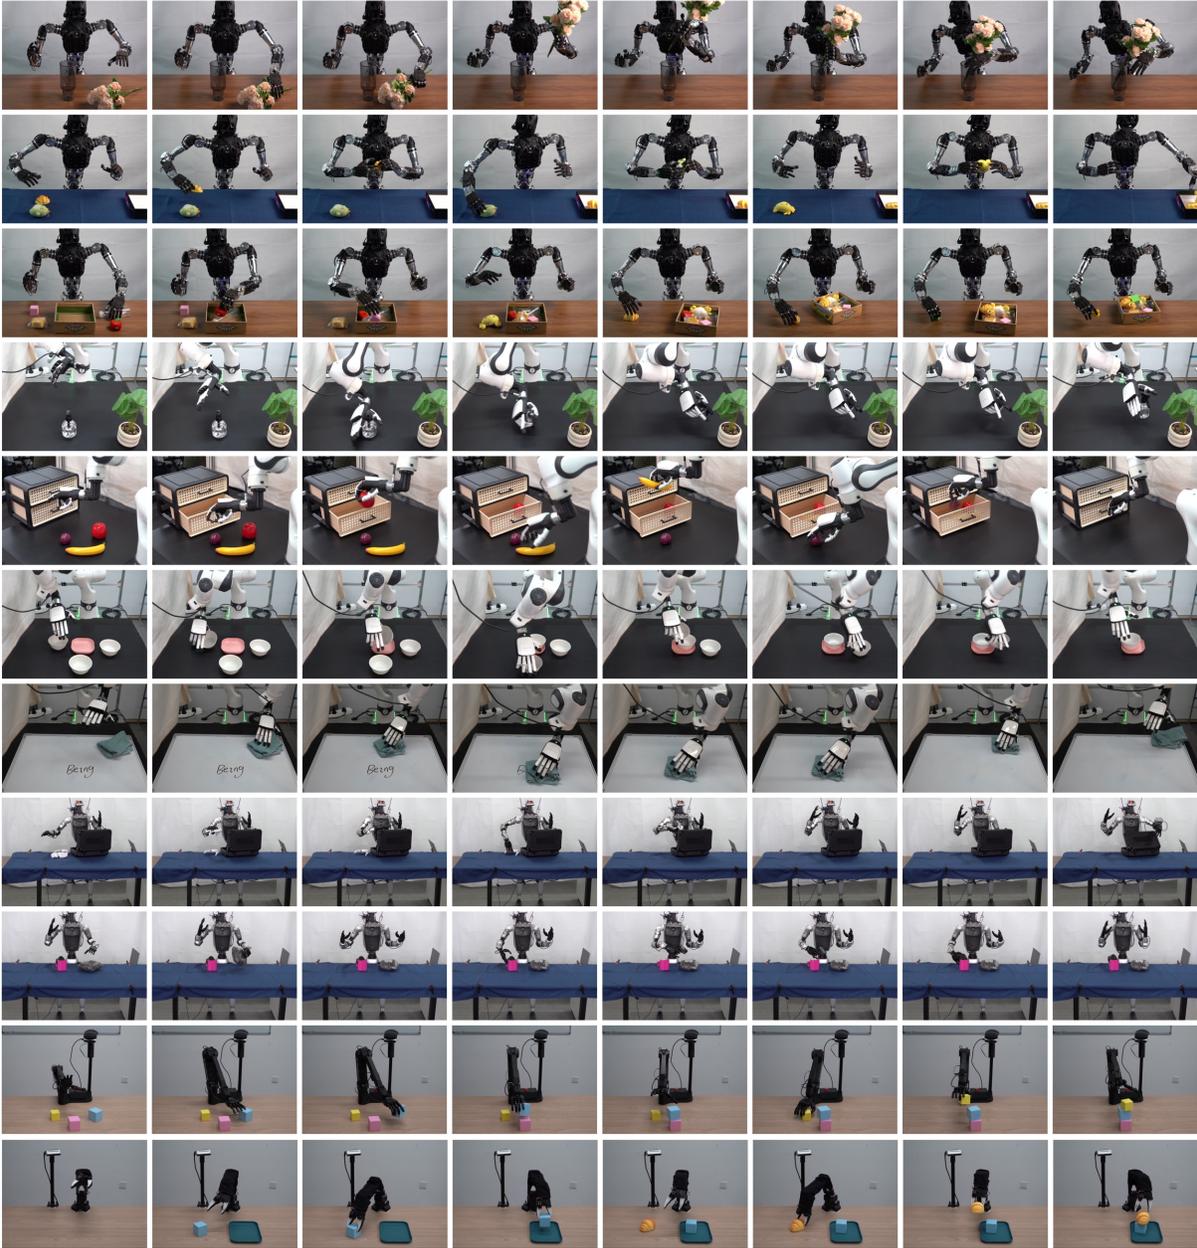


Figure 14: Example rollouts of all real-robot task suites.