

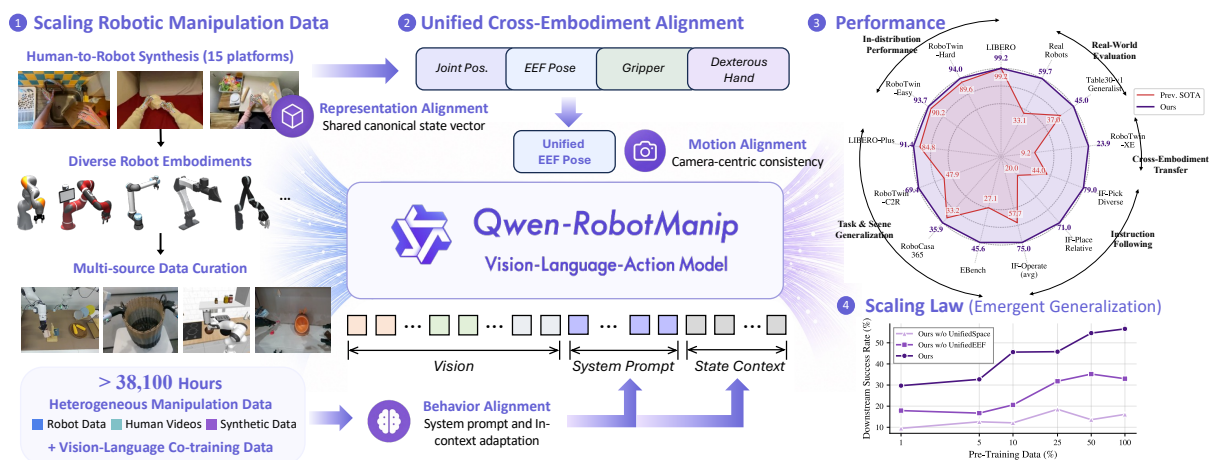
Qwen-RobotManip Technical Report: Alignment Unlocks Scale for Robotic Manipulation Foundation Models

Qwen Team

<https://qwen.ai/blog?id=qwen-robotmanip>
<https://github.com/QwenLM/Qwen-RobotManip>

Abstract

Foundation models in language and multimodality achieve strong generalization because heterogeneous data sources can be aligned under a unified formulation, and abundant low-cost data from the internet allows diverse training signals to reinforce one another at scale. In this report, **we investigate whether this scaling recipe can be applied to robotic manipulation to achieve genuine generalization.** This is challenging because, unlike text, manipulation data is heterogeneous by nature, expensive to collect, and narrow in diversity, making alignment and scale simultaneously difficult to achieve. We present QWEN-ROBOTMANIP, a generalizable Vision-Language-Action foundation model built upon Qwen-VL. QWEN-ROBOTMANIP introduces a unified alignment framework across the representation, motion, and behavioral dimensions of manipulation, making large-scale multi-source training coherent rather than conflicting. This alignment capability in turn enables QWEN-ROBOTMANIP to absorb manipulation data at a scale that prior training regimes could not sustain. To provide a scaling engine for manipulation data, a human-to-robot synthesis pipeline converts egocentric hand demonstrations into robot trajectories across 15 platforms, and a rigorous curation pipeline harmonizes heterogeneous real-robot and synthetic datasets. To our surprise, by leveraging **only** open-source robotic manipulation datasets and human demonstration videos without any proprietary data collection, QWEN-ROBOTMANIP constructs a ~38,100-hour pretraining corpus and already exhibits emergent generalization capabilities, including zero-shot instruction following, robustness to perturbations, reactive error recovery, and cross-embodiment knowledge transfer. In experiments, we further find that most standard benchmarks systematically fail to capture the quality of pretraining. Thus, we instead adopt OOD evaluation settings, including RoboCasa365, LIBERO-Plus, EBench, RoboTwin-Clean2Rand, RoboTwin-IF (our new instruction-following benchmark), and RoboTwin-XE (our new cross-embodiment transfer benchmark), as our north star for measuring genuine generalization. QWEN-ROBOTMANIP achieves substantially better performance than prior state-of-the-art models, including $\pi_{0.5}$, across all OOD settings, ranks 1st in RoboChallenge with a 20% relative improvement, and is validated on real-robot platforms including AgileX ALOHA, Franka, UR, and ARX.



1 Introduction

Foundation models in language and multimodality (Brown et al., 2020; OpenAI, 2023; Dubey et al., 2024; Team et al., 2023; Yang et al., 2025; Bai et al., 2025; Team, 2026) achieve strong generalization because heterogeneous data sources can be aligned under a unified formulation, and abundant low-cost data from the internet allows diverse training signals to reinforce one another at scale. In this report, we investigate whether this scaling recipe can be applied to robotic manipulation to achieve genuine generalization. This is challenging because, unlike text, manipulation data is heterogeneous by nature, expensive to collect, and narrow in diversity, making alignment and scale simultaneously difficult to achieve.

The current state of Vision-Language-Action (VLA) models (Kim et al., 2024; Bjorck et al., 2025; Black et al., 2024; 2025; Community, 2026; Liang et al., 2026; Kim et al., 2025) illustrates how far this gap remains. Despite a rapid succession of models reporting competitive numbers on standard benchmarks (Liu et al., 2023; Nasiriany et al., 2024; Mu et al., 2025), the generalization being demonstrated is largely superficial (Zhang et al., 2026c). OOD evaluations in most works involve only minor visual perturbations while preserving the same embodiment, task structure, and workspace layout as data collection, and performance degrades sharply when models are tested beyond these narrow settings.

The reason these pretrained priors fail to transfer is twofold. First, existing robotic demonstration corpora (Padalkar et al., 2024; Khazatsky et al., 2024; Fang et al., 2024) are concentrated in narrow teleoperation setups, far too limited in embodiment and task diversity for a scaling recipe to take hold. Second, and more fundamentally, data diversity alone is insufficient without alignment (Luo et al., 2026a; Wang et al., 2026). When demonstrations from different embodiments arrive with incompatible observation and action representations, scaling data volume produces interference rather than synergy. Prior cross-embodiment efforts (Bjorck et al., 2025; Zheng et al., 2025; Black et al., 2025) have adopted shared architectures, embodiment tokens, or unified action tokenization, but without a formulation that makes the same physical motion numerically consistent across embodiments, additional data cannot be converted into improved capability. Alignment is therefore not an independent engineering choice but a prerequisite for data scaling itself.

We present **QWEN-ROBOTMANIP**, a Vision-Language-Action foundation model built upon Qwen-VL (Yang et al., 2025; Bai et al., 2025; Team, 2026), designed around this principle: *alignment first, then scale*. QWEN-ROBOTMANIP introduces a unified alignment framework across the representation, motion, and behavioral dimensions of manipulation, making large-scale multi-source training coherent rather than conflicting: a canonical state-action representation with per-dimension binary masking accommodates diverse robot morphologies within a single template, a camera-frame delta pose parameterization makes visually similar motions numerically proximate across coordinate frames, and an in-context policy adaptation mechanism reads intra-episode execution history as an implicit embodiment identifier for kinematic-aware behavioral adjustment. Training is conducted under a dual-stream co-training strategy that jointly optimizes over manipulation data and a curated vision-language stream, preventing the VLM backbone’s perceptual and reasoning capabilities from eroding under action prediction pressure.

This alignment capability in turn enables QWEN-ROBOTMANIP to absorb manipulation data at a scale that prior training regimes could not sustain. A human-to-robot synthesis pipeline converts egocentric hand demonstrations into robot trajectories across 15 platforms, and a rigorous curation pipeline harmonizes heterogeneous real-robot and synthetic datasets, together accumulating $\sim 38,100$ hours of manipulation data. Notably, this entire corpus is constructed from **only** open-source robotic manipulation datasets and egocentric human videos without any proprietary data collection, yet QWEN-ROBOTMANIP already exhibits emergent generalization capabilities, including robustness to perturbations, zero-shot instruction following, reactive error recovery, and cross-embodiment transfer.

We further argue that the field’s evaluation methodology must evolve alongside its models. Standard in-domain benchmarks, where models without large-scale robot pretraining match or exceed pretrained ones (Yan et al., 2025; Community, 2026), systematically fail to distinguish genuine generalization from in-distribution pattern memorization. QWEN-ROBOTMANIP therefore introduces OOD evaluation settings including LIBERO-Plus (Fei et al., 2025), RoboTwin-Clean2Rand (Mu et al., 2025), RoboCasa365 (Nasiriany et al., 2026), and EBench (Laboratory, 2026), alongside two new benchmarks: **RoboTwin-IF**, an instruction-following benchmark that tests whether policies condition on language as a genuine control signal rather than exploit visual shortcuts, and **RoboTwin-XE**, which evaluates zero-shot transfer to morphologically distinct robots. QWEN-ROBOTMANIP achieves state-of-the-art on standard benchmarks and substantially outperforms existing VLA models including GR00T-N1.7 and $\pi_{0.5}$ across all OOD evaluation axes, ranks 1st on the RoboChallenge Table30-v1 generalist track with a 20% relative improvement, and is validated on various real-robot platforms and tasks. We hope these results and benchmarks together raise the standard for how VLA generalization is measured across the field.

Our contributions are as follows.

A unified alignment framework for cross-embodiment training. We address representational heterogeneity through three complementary mechanisms. A canonical state-action representation with per-dimension binary masking, a camera-frame delta pose parameterization that grounds end-effector actions in the visual domain, and an in-context policy adaptation mechanism that treats intra-episode execution history as an implicit embodiment identifier together enable consistent signal extraction across diverse embodiments. A dual-stream co-training strategy jointly optimizes manipulation and vision-language objectives to preserve the perceptual and reasoning capabilities that underpin generalization.

A scalable multi-source data corpus. We consolidate $\sim 38,100$ hours of manipulation data from open-source robot datasets and egocentric human demonstrations. A human-to-robot synthesis pipeline converts any egocentric demonstration into trajectories across 15 robot platforms, providing a scalable and embodiment-rich data source. A multi-stage curation pipeline ensures signal quality across all heterogeneous sources. Beyond manipulation data, a curated vision-language mixture including novel embodied chain-of-thought and egocentric video understanding data preserves the VLM backbone’s perceptual and reasoning capabilities during VLA training.

A new standard for evaluating VLA generalization. We introduce OOD evaluation settings including LIBERO-Plus, RoboTwin-Clean2Rand, RoboCasa365, and EBench, alongside RoboTwin-IF, a benchmark that diagnoses genuine language conditioning, and RoboTwin-XE, a benchmark for zero-shot cross-embodiment transfer. We argue that in-domain metrics are insufficient proxies for foundation model capability and that OOD transfer is the correct measure. QWEN-ROBOTMANIP substantially outperforms existing VLA models across all OOD settings while achieving state-of-the-art on standard benchmarks.

Real-robot validation across diverse deployment scenarios. We validate QWEN-ROBOTMANIP on four physical platforms (AgileX ALOHA, Franka, UR, and ARX) across in-domain, out-of-domain, few-shot adaptation, and zero-shot cross-embodiment transfer settings. On the RoboChallenge Table30 v1 generalist track, QWEN-ROBOTMANIP ranks 1st. Results confirm that the generalization capabilities of QWEN-ROBOTMANIP hold under real-world deployment conditions.

2 Data Sources for Robotic Manipulation

The quality and diversity of training data are foundational to the generalization ability we seek. To build such a VLA model with strong generalization across embodiments, tasks, and environments, we curate a large-scale heterogeneous training corpus comprising three complementary data modalities: robotic manipulation demonstrations across diverse hardware platforms, egocentric human manipulation videos, and synthetic robot data generated by our human-to-robot pipeline. A unified curation and pre-processing pipeline processes all sources to ensure high-quality and consistent state, action, video, and language annotations. Table 1 summarizes the composition of the full corpus.

Table 1: Overview of the training data corpus.

Data Type	Embodiment Type	Data Sources	Task Setting	Total Time
Robot	Single-arm	{OXE, RoboMIND, DROID, RH20T,	Tabletop	3,808 h
	Dual-arm	AgibotWorld-Beta, RoboCOIN, RDT,	Tabletop	6,744 h
	Mobile & humanoid	InternData-A1, Galaxea Open-World}	Tabletop & indoor	868 h
Human	Human hands	EgoDex, VITRA, EgoVerse	Tabletop & in-the-wild	1,933 h
Human-to-Robot	15 dual-arm platforms	Synthesized from human data.	Tabletop & in-the-wild	24,808 h

2.1 Robotic Datasets

Robot manipulation demonstrations constitute the core of our pretraining corpus, spanning single-arm and bimanual tabletop manipulation, dexterous manipulation, mobile manipulation, and humanoid loco-manipulation in both simulation and real world. We incorporate nine open-source datasets, totaling over 11,000 hours of demonstrations.

Open X-Embodiment (OXE) (Padalkar et al., 2024) aggregates real-world robotic datasets from diverse research institutions. We retain three subsets (*Fractal*, *Bridge*, and *BC-Z*) of high-quality single-arm tabletop manipulation data across the Google Robot and WidowX platforms, contributing about 600 hours.

AgibotWorld-Beta (Agibot-World-Contributors, 2025) is a large-scale real-world humanoid manipulation

dataset collected on the AgiBot G1 bimanual platform. We use datasets collected with grippers, covering about 200 task types and 2,400 hours.

RoboMIND (Wu et al., 2025a) and **RoboMIND 2.0** (Wu et al., 2025b) provide large-scale real-world datasets covering single-arm, dual-arm, ALOHA (Zhao et al., 2023), and humanoid robots across diverse tabletop manipulation tasks. RoboMIND spans four embodiments including Franka Emika Panda, UR5e, AgileX Cobot Magic V2.0, and Tien Kung humanoid; RoboMIND 2.0 extends coverage to six platforms including Franka, UR5, AgileX, ARX, Tien Kung, and Tian Yi. Together they contribute about 1,400 hours of demonstrations.

Galaxea Open-World Dataset (Galaxea AI, 2025) provides \sim 500 hours of real-world bimanual mobile manipulation demonstrations collected on Galaxea dual-arm robots across diverse household tasks.

RoboCOIN (Wu et al., 2025c) is a large-scale multi-embodiment real-world dataset covering a wide range of bimanual and humanoid platforms. We retain 10 embodiment types including AgiBot G1, Airbot MMK2, Alpha Bot 2, AgileX Cobot Magic, Unitree G1edu, Leju, Realman R1 Lite, Realman RMC-AIDA-L, ALOHA, and Tianqin A2, resulting in about 430 hours of demonstrations.

DROID (Khazatsky et al., 2024) is an in-the-wild single-arm dataset collected with Franka Panda robots across 86 diverse real-world environments, contributing about 95,000 trajectories totaling 500 hours.

RH20T (Fang et al., 2024) is a large-scale contact-rich real-world dataset spanning 4 embodiments (Flexiv, UR5, Franka, and Kuka) with multi-modal sensing including visual, force-torque, audio, and proprioception. It covers 140+ tasks across 42 skill categories, resulting in about 1,100 hours of demonstrations.

RDT-1B (Liu et al., 2025) provides 29 hours of bimanual manipulation demonstrations collected on the ALOHA platform.

InternData-A1 (Tian et al., 2025) is a large-scale dataset generated in high-fidelity simulation environments, covering various single-arm and dual-arm embodiments across pick-and-place, articulated manipulation, and long-horizon tasks, totaling over 3,600 hours.

2.2 Egocentric Human Datasets

Egocentric human hand manipulation data is naturally aligned with the perspective of robot-mounted cameras, serving as an efficient source for expanding manipulation data (Kareer et al., 2025; Qiu et al., 2025; Zheng et al., 2026; Luo et al., 2025; 2026b; Zhang et al., 2026a). We collect egocentric data from three sources with hand pose annotations.

EgoDex (Hoque et al., 2025), collected using Apple Vision Pro, contains 338K demonstrations across 194 tabletop manipulation tasks totaling 829 hours of 30 Hz egocentric video. It provides SE(3) annotations for 25 joints of both hands per frame, tracked on-device using multiple calibrated cameras and visual-inertial SLAM. We utilize 732 hours for training.

VITRA (Li et al., 2025a) performs fully automated 3D hand reconstruction, camera trajectory estimation, and atomic action segmentation on unstructured egocentric videos. It draws from five video sources: Ego4D (Grauman et al., 2022) (cooking-and-cleaning and general activity subsets), EPIC-KITCHENS (Damen et al., 2018), EgoExo4D (Grauman et al., 2024), and Something-Something v2 (Goyal et al., 2017), totalling approximately 1M trajectories. We utilize the Ego4D and EPIC-KITCHENS subsets, contributing about 247 hours of video.

EgoVerse (Punamiya et al., 2026) is a large-scale collaborative egocentric manipulation dataset spanning 1,362 hours across 1,965 tasks, 240 scenes, and 2,087 demonstrators. Hand poses (21 keypoints per hand) and 6-DoF head poses are recovered via vision-based methods including visual-inertial SLAM and model-based pose estimation. We utilize the industry-contributed portion, contributing approximately 954 hours of video.

The three sources collectively amount to approximately **1,933** hours of video. All hand poses are converted into a unified representation of MANO (Romero et al., 2022) parameters and 21 keypoints; for sources lacking native MANO annotations, parameters are recovered via optimization-based fitting.

2.3 Human-to-Robot Data Synthesis

A significant gap exists between egocentric human data and robot data in both morphology and visual domains. To bridge this gap, we map human hand trajectories to the robot action space, and replace human hands in videos with robot models. Inspired by prior work (Lepert et al., 2025b;a), we design an end-to-end synthesis pipeline that explicitly separates the process into action alignment and visual




Figure 1: **Human-to-robot data synthesis pipeline.** (Top) Given egocentric video, the pipeline performs action retargeting and smoothing, SAM3-based hand segmentation, ProPainter inpainting, base pose search with MuJoCo IK, and depth-guided compositing. (Bottom) $\sim 1,933$ h of egocentric data from 3 sources is rendered across 15 robot morphologies, yielding $\sim 24,808$ h of synthesized demonstrations.

alignment (Figure 1).

Action Alignment. This stage focuses on trajectory retargeting and smoothing to bridge the morphology gap between human hands and parallel-jaw grippers. We define the robot action at frame t as $\mathbf{a}_t = (\mathbf{p}_t, \mathbf{R}_t, w_t)$, where $\mathbf{p}_t \in \mathbb{R}^3$ is the end-effector position, $\mathbf{R}_t \in SO(3)$ is the gripper orientation, and $w_t \in \mathbb{R}_{\geq 0}$ is the gripper width. Using the 3D hand keypoints \mathbf{k}_i from MANO, we define a virtual finger \mathbf{k}_{vf} as a weighted combination of the index and middle fingertips. The end-effector position \mathbf{p}_t is retargeted as the midpoint between the thumb tip and the virtual finger, and the gripper width w_t is their Euclidean distance:

$$\mathbf{k}_{vf} = 0.7 \mathbf{k}_{\text{index}} + 0.3 \mathbf{k}_{\text{middle}}, \quad \mathbf{p} = \frac{1}{2} (\mathbf{k}_{\text{thumb}} + \mathbf{k}_{vf}), \quad w = \|\mathbf{k}_{\text{thumb}} - \mathbf{k}_{vf}\|_2. \quad (1)$$

The gripper orientation is constructed as a right-handed orthonormal frame $\mathbf{R} = [\mathbf{x} \ \mathbf{y} \ \mathbf{z}]$. We first establish the grasp axis \mathbf{z} along the jaw-line direction (the line connecting the thumb tip and the virtual finger). Together with the wrist-to-fingertip direction $\mathbf{d} = \mathbf{k}_{vf} - \mathbf{k}_{\text{wrist}}$, these two vectors define the jaw plane; the gripper-normal axis \mathbf{y} is the normal of this plane, and the approach axis \mathbf{x} completes the right-handed frame:

$$\mathbf{z} = \frac{s \cdot (\mathbf{k}_{\text{thumb}} - \mathbf{k}_{vf})}{w}, \quad \mathbf{y} = \frac{\mathbf{z} \times \mathbf{d}}{\|\mathbf{z} \times \mathbf{d}\|}, \quad \mathbf{x} = \mathbf{y} \times \mathbf{z} \quad (2)$$

where $s = +1$ for the right hand and $s = -1$ for the left hand. This sign flip ensures that \mathbf{z} points in a consistent direction regardless of handedness, so that both hands map to the same gripper frame. The three axes correspond to: \mathbf{x} – approach direction, \mathbf{y} – gripper normal (perpendicular to the jaw plane), \mathbf{z} – grasp axis (along the jaw line).

Per-frame hand detection introduces high-frequency noise. We apply Savitzky–Golay (Savitzky & Golay, 1964) filtering to positions and widths, and Gaussian-weighted SLERP to orientations, producing smooth trajectories while preserving motion structure.

Visual Alignment. This stage replaces the human appearance with a robot model through a sequence of masking, inpainting, and rendering steps to bridge the visual domain gap. First, SAM3 (Carion et al., 2025) generates a binary mask $M_t \in \{0, 1\}^{H \times W}$ for the human arm using text prompts. Next, ProPainter (Zhou et al., 2023) inpaints the masked regions guided by optical flow, creating a clean background sequence $\{\hat{I}_t\}$ without human hands.

A fundamental challenge in converting ego videos to robot data is determining the robot base placement. Unlike robot-to-robot transfer where a source base position is available, egocentric hand trajectories are embodiment-free: there is no physical robot base to reference. We formulate this as an optimization over base placements: given N target end-effector poses $\{\mathbf{T}_i^{\text{ee}}\}_{i=1}^N$ and a robot with maximum reach r_{max} , we seek:

$$\mathbf{T}_{\text{base}}^* = \arg \max_{\mathbf{T}_{\text{base}}} \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \mathbb{1} \left[\text{IK}(\mathbf{T}_{\text{base}}^{-1} \mathbf{T}_k^{\text{ee}}) \text{ is feasible} \right] \quad (3)$$

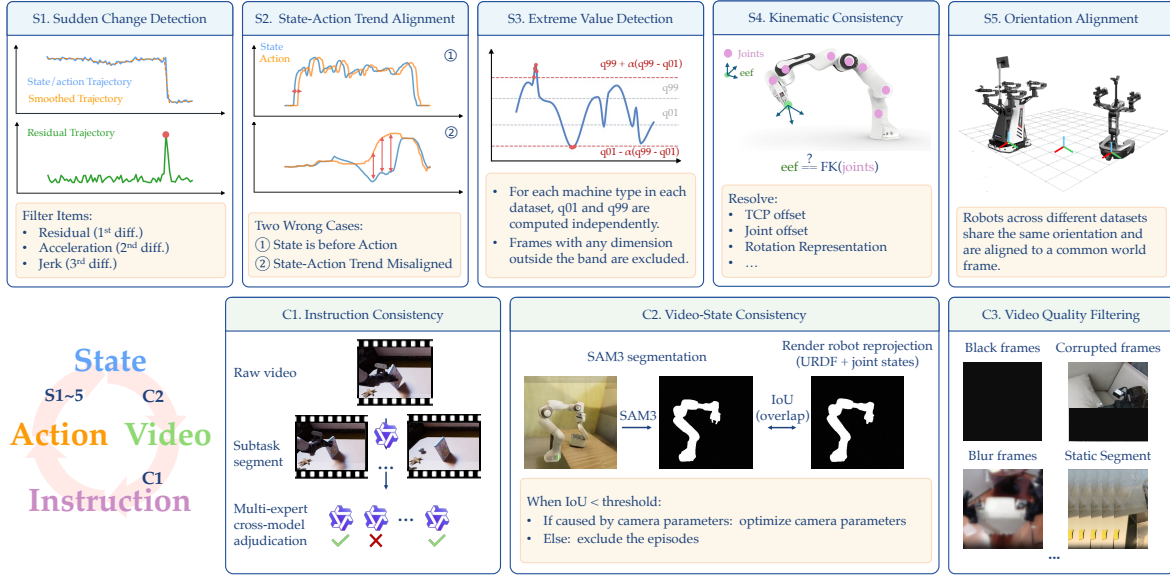


Figure 2: **Multi-stage data curation pipeline.** Five-stage state-action signal filtering (sudden change, trend alignment, extreme value removal, FK consistency, and base-frame alignment) followed by three cross-modal quality checks (instruction consistency, video-state consistency, and video quality filtering).

where $\mathcal{K} \subset \{1, \dots, N\}$ is a set of representative keyframes covering the spatial extremes of the trajectory. Candidate base placements are generated via grid search around the trajectory centroid, constrained by the per-morphology kinematic reach r_{\max} . This search is performed independently for each of the 15 robot morphologies, as different arm lengths and joint configurations require different base placements for the same trajectory.

Given the optimized base pose, we run inverse kinematics in a MuJoCo (Todorov et al., 2012; Zakka, 2026) virtual environment to track the smoothed action trajectory, rendering the robot image I_t^{robot} and its depth map D_t^{robot} . Depth Anything v3 (Lin et al., 2025) estimates a metric depth map D_t for the scene. We compute an occlusion mask $M_t^{\text{occ}} = \mathbb{1}[D_t^{\text{robot}} \leq D_t]$ to naturally composite the robot onto the clean background:

$$I_t^{\text{syn}} = M_t^{\text{occ}} \odot I_t^{\text{robot}} + (1 - M_t^{\text{occ}}) \odot \hat{I}_t. \quad (4)$$

Each human demonstration is rendered into 15 bimanual robot configurations (each composed of two identical arms from: Panda, UR5e, ARX-L5, xArm7, Sawyer, Kinova Gen3, IIWA, Jaco, FR3, UR10e, ViperX, WidowX, Piper, YAM, AgileX ALOHA), yielding approximately **24,808** hours of synthesized demonstrations in total.

Action Speed Alignment. Egocentric hand manipulation exhibits significantly higher action speeds than robot teleoperation data. To align the action speed distributions, we apply per-source frame subsampling during training to match the robot data speed. Specifically, EgoDex is downsampled to 60% of its original frame rate ($\sim 1.7\times$ slower), EgoVerse to 45% ($\sim 2.2\times$ slower), and ViTRA to 25% ($\sim 4\times$ slower).

2.4 Data Curation and Pre-Processing

Aggregating manipulation data across diverse embodiments, simulators, and collection pipelines introduces heterogeneous noise in recorded state and action signals, including discrete outliers from physical collisions, temporal misalignment between state and action streams due to unsynchronized clocks or packet loss, extreme values that destabilize optimization, and inconsistent end-effector conventions across datasets sharing the same robot embodiment. We address these through a five-stage filtering pipeline applied to all datasets prior to training.

Stage 1: Sudden Change Detection. For each signal dimension, we extract a smooth trend via cascaded median filtering and Savitzky–Golay smoothing (Savitzky & Golay, 1964), then compute three complementary deviation signals: the absolute residual between the raw and smoothed trajectory, the second-order finite difference (acceleration), and the third-order finite difference (jerk). A frame is flagged when the residual exceeds a scaled threshold *and* either acceleration or jerk also exceeds its threshold, reducing false positives from slow drift while preserving sensitivity to abrupt transients. Thresholds are

set per dataset according to embodiment type, rotation representation, data source (real vs. simulation), and base mobility. Exclusion strategies range from frame-level removal to full episode discard. For instance, in InternData-A1 (Tian et al., 2025) where sudden changes arise exclusively from physical collisions (e.g., a gripper contacting a rigid object), the corrupted episode is discarded entirely.

Stage 2: State-Action Trend Alignment. In a correctly recorded episode, action commands should temporally lead or coincide with resulting state changes, which is a causal invariant violated when timestamps are unsynchronized or there is packet loss. For each shared joint dimension, we smooth both the state and action trajectories, then estimate the optimal temporal lag via cross-correlation, then compute a *directional agreement* (DA) metric on lag-aligned first-order differences. Dimensions with DA below a dataset-specific threshold (typically 0.6-0.7) are flagged and their episodes excluded. For datasets using delta actions, we first integrate the action sequence to recover absolute values before comparison. This stage revealed severe quality issues in certain subsets: 81% of episodes in the RoboMIND UR-type data failed this check and were excluded.

Stage 3: Extreme Value Filtering. Frames with state or action values outside the expected range are removed to prevent distortion of the quantile-based normalization ($[q_{01}, q_{99}] \rightarrow [-1, 1]$) used during training. Per-dimension q_1 and q_{99} percentiles are computed per embodiment type, and frames outside the band $[q_1 - \alpha(q_{99} - q_1), q_{99} + \alpha(q_{99} - q_1)]$ are excluded. Gripper dimensions are exempt due to their bimodal distributions.

Stage 4: Joint-End-Effector Forward Kinematics Consistency. We compute forward kinematics (FK) via Pinocchio (Carpentier et al., 2019) from each robot’s URDF and compare against logged end-effector poses. The discrepancies can arise from differing joint-angle sign conventions, differing end-effector frame definitions, incorrect rotation representations, incorrect base-frame assumptions, and erroneous end-effector logging. Rather than aggressively filtering, this stage primarily performs *data correction*: constant positional offsets are resolved by adjusting the tool-center-point (TCP) definition, and shoulder-relative bimanual poses are transformed into the world frame. This process revealed that the same robot model can carry different joint-angle conventions across datasets, further motivating the unified state-action representation of Sec. 3.2.

Stage 5: Base Frame and End-Effector Orientation Alignment. We apply per-dataset rotation corrections to align world-frame orientation conventions, ensuring the positive x -axis consistently corresponds to the robot’s forward-facing direction and that the unified state-action representation is geometrically consistent across embodiments.

Beyond this five-stage state and action signal quality filtering, we apply three additional checks to ensure cross-modal consistency across video, language, and proprioceptive observations.

Check 1: Instruction Consistency. We verify semantic consistency between each demonstration and its language annotation via a three-stage VLM-based pipeline. First, long episodes are decomposed into subtask-level segments (Lei et al., 2026) so that each clip corresponds to a temporally localized action unit, keeping the visual evidence focused and tractable for automated assessment (Temporal Normalization for Evaluation Units). Second, each segment is evaluated through structured reasoning-guided prompting: rather than requesting an immediate binary label, the VLM is directed to attend to manipulated objects, action semantics, temporal ordering, and agent-environment interaction, producing an intermediate analytical judgment before issuing a final consistency decision. This structured prompting reduces superficial or heuristic responses and improves label interpretability (Structured Reasoning-Guided Annotation). Third, clips flagged as non-aligned or ambiguous by the initial model are adjudicated by multiple VLMs as independent evaluators, with the final label determined by cross-model voting, reducing single-model bias and improving label robustness (Multi-Expert Cross-Model Adjudication). Inconsistent samples are excluded from training.

Check 2: Video-State Consistency. We perform rigorous data cleaning to remove low-quality or misaligned samples. To verify video-state consistency, we render the robot projection into the image plane using the URDF and recorded joint states, segment the actual robot mask with a fine-tuned SAM3 (Carion et al., 2025) model, and measure their overlap. Samples with low overlap are filtered out.

Check 3: Video Quality Filtering We apply video-level data cleaning to remove frames that may degrade policy learning. We remove visually invalid frames including black, corrupted, blurred, and prolonged static segments, using image processing checks applied jointly with state and action signals to detect redundant static periods typically at episode boundaries. Task-critical key frames such as gripper

closure events are explicitly preserved to avoid discarding visually subtle but semantically important transitions.

2.5 Vision-Language Co-training Datasets

Prior works have demonstrated that co-training VLAs with vision-language (VL) data mixtures can significantly improve their generalization ability (Black et al., 2025; Driess et al., 2025; Fang et al., 2026). By incorporating VL data during VLA training, the model retains the rich visual and semantic knowledge acquired from web-scale multimodal pretraining and transfers this knowledge to action generation through the action expert. For example, this enables the model to follow novel language instructions, operate in unfamiliar scene backgrounds, or manipulate previously unseen objects.

To this end, we curate a comprehensive VL dataset from multiple sources, including proprietary data, open-source datasets (e.g., RoboPoint (Yuan et al., 2024), RefSpatial (Zhou et al., 2025), PixMo (Deitke et al., 2025), and CapsFusion (Yu et al., 2024)), and carefully synthesized embodied-centric data. The resulting VL mixture comprises approximately 28M data points, spanning the following categories:

- (1) **General Visual Understanding**, including visual question answering, multi-image reasoning, and image captioning at varying granularities (from single-sentence summaries to paragraph-level detailed descriptions), which preserves the model’s broad visual perception and commonsense reasoning capabilities;
- (2) **Spatial Perception and Reasoning**, covering 2D/3D visual grounding, point localization, counting, spatial relationship reasoning (depth comparison, distance estimation, camera viewpoint inference), and manipulation feasibility reasoning, which are directly transferable to robotic spatial understanding;
- (3) **OCR and Document Understanding**, which helps maintain the VLM’s ability to recognize text, numbers, and symbols, a capability that is also required in robotic tasks involving labeled objects (e.g., identifying a block marked with a specific number);
- (4) **Multimodal Specialized Knowledge**, covering domain-specific visual reasoning tasks such as STEM problem solving, chart interpretation, and visual puzzle reasoning, which helps prevent catastrophic forgetting of the VLM’s general multimodal reasoning capabilities during VLA fine-tuning;
- (5) **Instruction Following, Multilingual, and Pure Text data**, which strengthens the model’s ability to follow diverse natural language instructions, a capability that is critical for generalizing to novel robot tasks, while also enabling multilingual robot control and preserving text generation quality;
- (6) **Embodied-Centric VL Data**, which we specifically curate to bridge the gap between web-scale VL knowledge and robotic manipulation. This subset includes: (a) embodied chain-of-thought (ECoT) reasoning data derived from robot manipulation trajectories, where the model performs structured reasoning in three stages: first describing the current scene state from multi-view observations (including gripper status, object positions, and spatial layout), then assessing task progress by comparing the current state against the overall goal, and finally predicting the next atomic manipulation action; (b) egocentric video understanding data, where the model describes fine-grained hand/arm movements, hand-object interactions, and object state changes from short clips of first-person human manipulation videos; and (c) 2D trajectory prediction data, where the model predicts future movement trajectories of human hands or robot end-effectors as sequences of normalized 2D coordinates, conditioned on visual observations and task instructions.

Among these, categories (1)–(5) serve a dual purpose: they prevent catastrophic forgetting of the pre-trained VLM’s general capabilities, while certain subsets, such as spatial reasoning, visual grounding, and OCR, directly benefit robotic manipulation by strengthening the model’s spatial understanding, object recognition, and scene generalization. Category (6) is specifically curated to bridge VL understanding and action generation: ECoT data teaches the model to perceive embodiment-specific scene states, track task progress, and reason about next actions in language. This encourages the VLM backbone to build richer embodied representations that are more directly useful for downstream continuous action generation (Zawalski et al., 2024; Chen et al., 2025c); egocentric video data exposes the model to fine-grained human manipulation patterns and object state transitions from a first-person perspective, grounding its understanding of how physical interactions unfold, including how objects deform, slide, or topple under contact and how hand-object configurations evolve during grasping, placing, and tool use. This knowledge transfers to robotic manipulation despite the embodiment gap. In addition, 2D trajectory prediction data directly connects visual observations to spatial motion reasoning in image coordinates, and together these data sources establish a shared representational foundation that facilitates knowledge transfer to low-level action prediction through the action expert.

Table 2: Taxonomy of atomic action types.

Category	Action Type	Example
Movement	Reach (and grasp)	“Reach toward the red cup on the left side of the table and grasp it.”
	Move (and release)	“Move the held cup onto the wooden tray and release it.”
Manipulation	Flip	“Flip the golden pancake 180 degrees in the frying pan.”
	Rotate	“Rotate the black knob on the oven door clockwise.”
	Toggle	“Toggle the red power switch to on.”
	Open	“Open the wooden drawer on the left side of the cabinet.”
	Close	“Close the lid of the black laptop.”
	Push	“Push the yellow block forward along the table.”
	Pull	“Pull the silver drawer handle away from the cabinet.”
	Insert	“Insert the red peg into the circular hole in the board.”
	Press	“Press the sponge against the table surface.”
	Click	“Click the red button on the control panel.”
Strike	“Strike the silver nail with the wooden hammer.”	
Special	Handover	“Move the held cable toward the right arm and release it.”
	Return to home	“Return the arm to its home position.”
	Other	“Pour the water from the red cup into the glass.”

We describe the synthesis procedure for each type of embodied-centric VL data below.

Embodied Chain-of-Thought (ECoT) Data. Inspired by prior work (Zawalski et al., 2024; Feng et al., 2026; Li et al., 2026), we construct ECoT supervision that trains the VLM to jointly perform three forms of embodied reasoning: describing the current scene, assessing task progress, and predicting the next atomic manipulation action. For a sampled timestamp t in a manipulation trajectory with task instruction, we synthesize one ECoT annotation using both the current multi-view observation and additional trajectory context available only at annotation time.

Specifically, at timestamp t , we first extract synchronized images from all available camera views (e.g., front, wrist, and side views). We then construct three forms of additional context. First, we build a *memory summary* by uniformly subsampling frames from the beginning of the episode up to t and prompting a strong VLM to summarize completed actions and visible state changes. This memory mainly supports task progress assessment during annotation, since many intermediate goals or prior manipulations are no longer directly observable in the current frame. Second, we construct a *future action preview* from 6 frames sampled at 1-second intervals starting from t and ask a strong VLM to summarize the robot’s immediate future behavior in this short clip. This future preview provides direct annotation-time evidence for predicting the next atomic action. Third, we compute a coarse *temporal progress* estimate from the relative position of t within the full trajectory, which serves as a weak auxiliary cue for judging whether the task is likely close to completion.

Given the multi-view images at t , the task instruction, and the auxiliary annotation-time context described above, we prompt a strong VLM¹ to generate a structured three-part ECoT response: (1) a Scene Description summarizing observable objects, spatial relations, robot arm positions, and gripper states; (2) a Task Progress Assessment evaluating completed subgoals and ending with an explicit completion judgment (*Task complete.* or *Task not yet complete.*); and (3) a Next Action predicting a single atomic manipulation step from Table 2. Although the annotating VLM has access to privileged trajectory context, the prompt requires the generated ECoT text to be expressed using only evidence from the current observation and task instruction. During training, each ECoT data sample is converted into a standard VL input-output pair: the model receives only the multi-view images and task instruction, and is trained to generate the full three-part ECoT text. The privileged signals are used only during data synthesis to improve annotation quality and are excluded from training inputs.

Egocentric Video Understanding Data. We construct egocentric video understanding data from the human manipulation videos. For each episode, we split the main-camera video into non-overlapping clips of random duration (1.5–3 seconds) and extract 4 uniformly spaced frames per clip (at 0%, 33%, 67%, and 100% of the clip duration). A strong VLM is prompted to describe the fine-grained manipulation actions observed across the 4-frame sequence, including hand and arm movement directions, hand-object interactions with spatial relationships, and any object state changes. Clips with very little visual change, as determined by the VLM, are filtered out to avoid training on uninformative static segments. The resulting annotations teach the model to perceive and describe the dynamics of physical manipulation

¹In practice, we use Qwen3.6-Plus with thinking mode throughout the ECoT data synthesis process.

from an egocentric perspective, complementing the ECoT data, which operates on robot observations.

2D Trajectory Prediction Data. To further facilitate the VLM in learning motion planning from both robotic and human manipulation demonstrations, while alleviating ambiguity in depth perception, we project the trajectories of the robot end-effector (EEF) and the human hand in egocentric data onto the image using the estimated camera parameters. In addition, we filter out samples with little motion using a bounding-box-based criterion, thereby removing uninformative data points in which the EEF or human hand barely moves.

3 Qwen-RobotManip: The Generalizable Vision-Language-Action Model Design

3.1 Main Architecture

QWEN-ROBOTMANIP follows a decoupled architecture consisting of a vision-language backbone for multimodal perception and semantic reasoning, and a flow-matching action expert for continuous action generation. This decoupling allows the action expert to specialize in high-frequency, fine-grained motor control while the backbone retains and extends its pretrained perceptual and reasoning capabilities through joint end-to-end training.

Vision-language backbone. We adopt Qwen3.5-4B (Team, 2026) as the vision-language backbone. Qwen3.5 is a natively multimodal model trained with early vision-language fusion: visual tokens from a Vision Transformer with dynamic-resolution spatial merging are interleaved directly into the text token stream and processed uniformly across images and language instructions within a single transformer. Given one or more camera views together with a natural language task instruction, the backbone encodes them jointly into contextual representations (e.g., last-layer hidden states $D_{\text{vlm}}=2560$) that capture both fine-grained visual features and task-level semantics, which are then consumed by the action expert via cross-attention.

Action expert. We attach a Diffusion Transformer (DiT) (Peebles & Xie, 2023) as a flow-matching action expert (Chi et al., 2023; Black et al., 2024; Liang et al., 2023) for learning precise continuous actions from both robot trajectory data and egocentric human demonstrations. The expert consists of $N=10$ transformer blocks with hidden dimension $D_{\text{act}}=768$ and 12 attention heads. Each block performs self-attention over the concatenated state-and-action token sequence, followed by cross-attention to VLM hidden states and a SwiGLU feed-forward network. Cross-attention layers alternate between attending to *visual* tokens (even-indexed blocks) and *language* tokens (odd-indexed blocks), both extracted from the final layer of the VLM, letting the expert separately ground action predictions in spatial observations and linguistic instructions at each processing stage. The robot’s proprioceptive state is encoded by a two-layer MLP and prepended to the noisy action token sequence before entering the DiT blocks. The expert is further conditioned on denoising timestep embeddings and additional learned camera embeddings detailed in §3.3.

The expert is trained with a flow-matching objective (Lipman et al., 2023; Esser et al., 2024). Given a ground-truth action chunk \mathbf{a} , a timestep $t \sim \text{Beta}(1, 1.5)$ is sampled and an interpolant $\mathbf{x}_t = (1-t)\boldsymbol{\epsilon} + t\mathbf{a}$ is constructed from Gaussian noise $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The model is then trained to minimize mean squared error on the predicted velocity field $\mathbf{x}_1 - \mathbf{x}_0$. At inference, action sequences are produced via 4 Euler integration steps, enabling low-latency real-time control.

3.2 Cross-Embodiment State and Action Representation

Heterogeneous proprioceptive states and action spaces across embodiments make scalable training on mixed multi-embodiment datasets a key challenge. We address this by introducing an **80-dimensional canonical vector representation** for states and actions. The representation is structured as two 29-dimensional per-arm blocks followed by 22 reserved dimensions. Each per-arm block is organized into the following semantic groups:

- **Joint positions** (7 dims): joint positions for the robot arm;
- **End-effector pose** (9 dims): Cartesian position (3) and orientation in a 6D continuous rotation representation (Zhou et al., 2019) (6);
- **Gripper state** (1 dim): joint position for the parallel gripper;
- **Dexterous hand joints** (12 dims): active hand joint positions for embodiments equipped with multi-fingered dexterous hands.

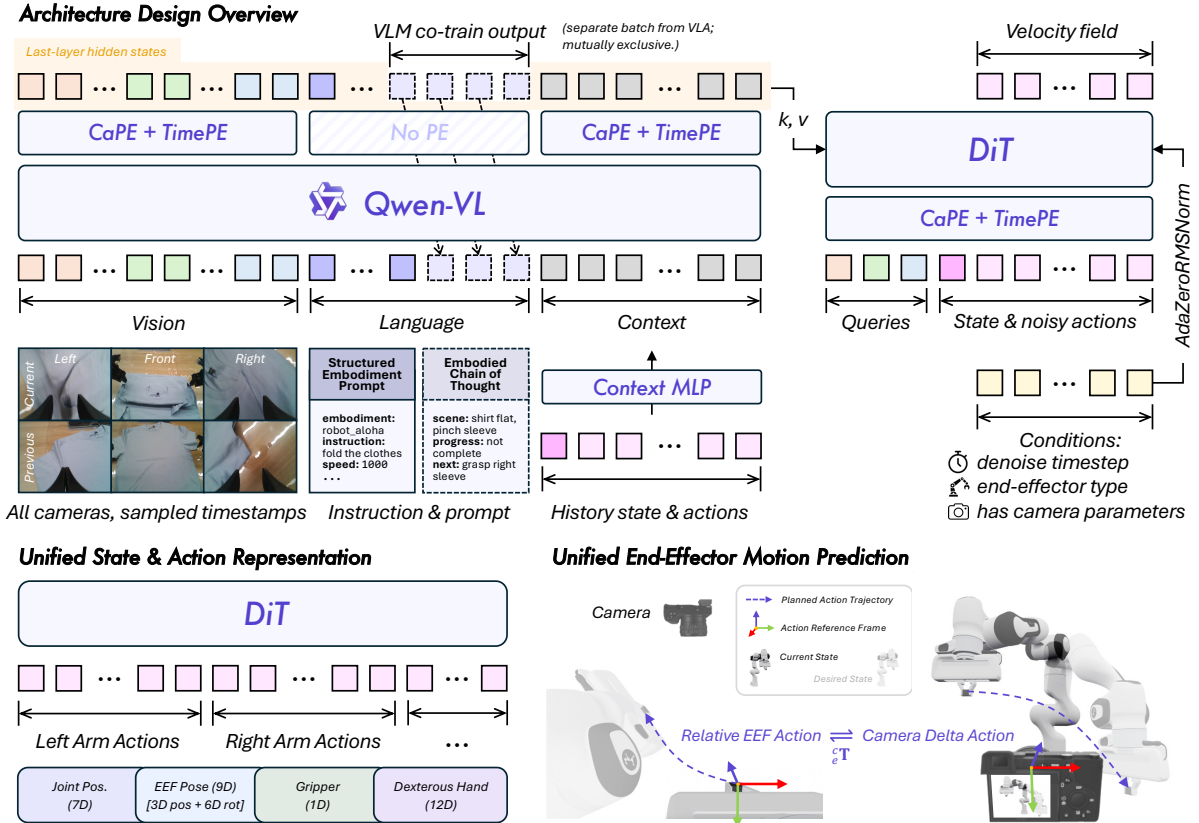


Figure 3: **Overview of QWEN-ROBOTMANIP.** The model couples a Qwen-VL backbone with a flow-matching Diffusion Transformer (DiT) action head. The backbone jointly encodes multi-view visual tokens, structured embodiment prompts, and historical context tokens, with last-layer hidden states injected into the DiT via alternating cross-attention. States and actions share a unified 80-dimensional canonical representation, with end-effector actions expressed as camera-frame delta poses to align the action space with visual observations across embodiments, conditioned on camera and end-effector type embeddings for embodiment-aware denoising. VLM co-training and VLA training use separate batches, where each batch contains either (Vision, QA) or (Vision, Language, Context, Action) data.

The trailing 22 reserved dimensions are shared across both arms and are available for additional degrees of freedom such as mobile-base velocity. For the *state* vector, all values are expressed in absolute coordinates. For the *action* vector, joint actions are expressed as absolute values and end-effector actions are expressed as relative deltas from the current state. In particular, end-effector orientation deltas are parameterized as 3D rotation vectors rather than the 6D representations used for states. § 3.3 further details the camera-frame delta representation for end-effector actions.

Different robot embodiments populate different subsets of this canonical template. For example, a 7-DOF single-arm gripper (*e.g.*, Franka Panda) fills the joint, end-effector, and gripper fields of one arm, leaving the remaining dimensions as zero. A dual-arm system (*e.g.*, ALOHA) fills both per-arm blocks. A robot with dexterous hands additionally populates the hand-joint dimensions. Zero-padded dimensions are excluded from the training loss via a per-dimension binary mask, ensuring that gradients flow only through semantically populated entries and preventing spurious supervision on structurally absent degrees of freedom.

3.3 Unified End-Effector Motion Prediction

The canonical state-action representation of § 3.2 unifies the structural layout of states and actions across embodiments, but does not yet address a subtler source of fragmentation: end-effector poses recorded in different coordinate frames across datasets (Lyu et al., 2026). When the same motion is expressed relative to different base frames or camera frames depending on the data source, the model must learn to reconcile these geometric inconsistencies rather than focusing on the underlying manipulation skill. We address this by grounding all end-effector actions in a shared camera-frame delta pose representation and injecting camera geometry directly into the action expert via positional encodings. Together with embodiment-

aware conditioning, this ensures that actions which appear visually similar are numerically proximate across embodiments, enabling the model to extract cross-embodiment synergies from heterogeneous data.

Concretely, we extract $N_{ee} \in \{1, 2\}$ independent **40-dimensional per-end-effector tokens** from the 80-dimensional state-action vector. The 29 active dimensions of each arm are packed into a 40-dimensional slot with 11 reserved for future extension, and the DiT processes these tokens jointly via self-attention.

Camera-frame delta pose action representation. Rather than representing end-effector motion as an absolute pose in the robot base frame, a relative pose in the end-effector local frame, or a world-frame delta, we adopt a *camera-frame delta pose* representation (Chen et al., 2025a). Its key property is that actions appearing visually similar in the image are also numerically proximate in the action space, directly aligning the action representation with the visual observation space and facilitating cross-embodiment transfer. This requires calibrated camera intrinsics and extrinsics at both training and inference.

Formally, let c denote the reference camera frame, e the current end-effector frame, and e^* the desired end-effector frame at a future step. The pose component of the predicted action is:

$$\mathbf{a}_p = \begin{bmatrix} {}^c\mathbf{R}_e {}^e\mathbf{R}_{e^*} & {}^c\mathbf{R}_e {}^e\mathbf{t}_{e^*} \\ \mathbf{0} & 1 \end{bmatrix} \quad (5)$$

The rotational block ${}^c\mathbf{R}_e {}^e\mathbf{R}_{e^*}$ expresses the relative end-effector rotation ${}^e\mathbf{R}_{e^*}$ in the camera frame by conjugating with the camera-to-end-effector extrinsics, while the translational block ${}^c\mathbf{R}_e {}^e\mathbf{t}_{e^*}$ projects the desired end-effector displacement into camera coordinates. The full expression is thus geometrically equivalent to projecting the relative end-effector action into the camera frame via the extrinsics. A more compact alternative is (Zhang et al., 2026b):

$$\mathbf{a}_p = {}^c\mathbf{T}_{e^*}^e \quad (6)$$

While eq. (6) eliminates end-effector definition inconsistencies entirely, its translational component is coupled with the relative end-effector rotation ${}^e\mathbf{R}_{e^*}$ and the camera-to-end-effector offset ${}^e\mathbf{t}_c$, making it more susceptible to long-tail distributions and more sensitive to calibration errors. We therefore adopt eq. (5) in implementation.

Camera-aware positional encoding. To enable the action expert to reason about camera geometry, we inject camera parameters into the DiT’s cross-attention layers via Camera Positional Encoding (CaPE) (Kong et al., 2024). Camera pose is encoded via CaPE, occupying 32 of each 64-dimensional attention head’s dimensions, with the remaining 32 used by RoPE (Heo et al., 2024) for temporal indexing. Each image token’s positional encoding is derived from the extrinsics of its corresponding camera, while each state/action token uses the extrinsics of its selected reference camera. Because CaPE is a rotational positional encoding, the global world-frame origin cancels algebraically in the dot-product attention, leaving only the relative pose between each visual token and the querying state/action tokens.

Following the practice of GTA (Miyato et al., 2024) and PRoPE (Li et al., 2025b), we apply CaPE not only to keys and queries but also to values and attention outputs, strengthening the geometric consistency of the cross-attention. Camera intrinsics are incorporated by projecting the normalized image-plane coordinates of each visual patch through a learned linear layer and adding them to the corresponding image token, providing per-token field-of-view awareness.

End-effector-aware conditioning. Beyond the denoising timestep, the DiT is further conditioned on two additional signals, both applied via additive embeddings through adaptive layer normalization (Peebles & Xie, 2023).

1. *End-effector type embedding*: a learned codebook entry per end-effector category (single-arm, dual-arm left, dual-arm right, egocentric head, or mobile base) associated with each state/action token, allowing the model to apply embodiment-specific action priors.
2. *Auxiliary flag embedding*: a binary embedding indicating whether calibrated camera parameters are available for the current sample, switching the predicted pose action space between camera-frame delta mode and robot-base relative mode.

Multi-view reference camera selection. In multi-view settings, the end-effector action is expressed relative to a chosen reference camera frame. During training, for single-arm datasets we randomly select any available external or wrist-mounted view as the reference. For dual-arm datasets, we randomly apply one of two strategies: (1) both arms share a head-mounted camera or any available third-person view as the common reference frame; (2) the left arm uses the left wrist camera and the right arm uses the right wrist camera as their respective reference frames.

Within the DiT’s cross-attention, each image token uses the pose of its corresponding camera for CaPE, while each state/action token uses the pose of its selected reference camera for CaPE, guiding the DiT to denoise the camera delta action expressed in that reference frame. Because CaPE is a rotational encoding, inter-view relative poses are encoded implicitly, with the world-frame superscript canceling algebraically.

3.4 Embodiment Prompt

We adopt a structured prompt to condition the policy on both task semantics and execution context. Each prompt consists of the following fields:

- **Embodiment:** the robot platform (*e.g.*, `robot_aloha`), enabling the model to account for morphological and control differences across embodiments.
- **Instruction:** the high-level task description, defining the overall objective of the episode.
- **Speed:** the episode length in timesteps, discretized into bins of 500 steps.
- **FPS:** the temporal sampling rate of the input sequence.
- **Camera View Direction:** the camera’s position relative to the robot arm, either `arm side` or `opposite side`.

```

Structured Embodiment Prompt Example

embodiment: robot_aloha
instruction: Take the toy off the table and put it on the mat.
speed: 1000
fps: 30
camera view direction: arm side
  
```

Together, these fields allow the model to capture not only what task should be performed, but also which robot is acting and how the behavior is temporally structured, reducing ambiguity in policy learning, improving adaptability across embodiments, and increasing robustness to variations in execution speed and frame rate. To further improve robustness to incomplete inputs, we randomly drop the embodiment, speed, and fps fields with probability 15% during training, encouraging the model to generalize when prompt information is partially unavailable at test time.

3.5 In-Context Policy Adaptation

Despite the strong generalization enabled by cross-embodiment pretraining, deploying a VLA policy to a new robot or environment often requires rapid behavioral adaptation without parameter updates. Inspired by in-context learning in large language models, we equip QWEN-ROBOTMANIP with an in-context policy adaptation mechanism that conditions current action prediction on a structured window of recent execution history (observation-action pairs) from the same episode, enabling the policy to adapt its behavior at deployment time without any parameter update or task-specific fine-tuning.

Execution context representation. A key design question is what information constitutes a useful policy context. We draw a direct analogy from the model’s own inference procedure. At each decision step, QWEN-ROBOTMANIP observes the current visual observations and proprioceptive state and predicts a complete action chunk of K steps. We therefore define one context chunk as exactly this triplet $(\mathbf{o}_h, \mathbf{s}_h, \mathbf{a}_h)$, consisting of the visual observation, proprioceptive state, and the K -step action sequence executed during chunk h . This records what the robot saw, was in, and did. A context of H such chunks thus provides the policy with a structured window of recent behavior it can directly reason about.

The two modalities within each context chunk are processed through complementary pathways, owing to their fundamentally different representational structures. Historical frames \mathbf{o}_h are prepended to the current frame and processed jointly by the VLM visual encoder within a single forward pass, with an image-count annotation appended to the language instruction to help the VLM attribute each visual token to its correct temporal position. Proprioceptive states and action chunks, which cannot be processed by the visual pathway, are projected into the VLM hidden space by two lightweight MLP encoders. The state encoder MLP_s and action encoder MLP_a produce per-chunk token representations with learned temporal position embeddings $\mathbf{e}_h^{\text{temp}}$ to distinguish chunks and slot embeddings $\mathbf{e}_{0:K}^{\text{slot}}$ to distinguish action tokens

within each chunk:

$$\mathbf{t}_h^s = \text{MLP}_s(\mathbf{s}_h) + \mathbf{e}_h^{\text{temp}} \in \mathbb{R}^{D_{\text{vlm}}}, \quad (7)$$

$$[\mathbf{t}_h^{a,0}, \dots, \mathbf{t}_h^{a,K'-1}] = \text{reshape}(\text{MLP}_a(\text{flatten}(\mathbf{a}_h))) + \mathbf{e}_h^{\text{temp}} + \mathbf{e}_{0:K'}^{\text{slot}} \in \mathbb{R}^{K' \times D_{\text{vlm}}}. \quad (8)$$

All H chunks are serialized chronologically into a single context token sequence:

$$\mathbf{C} = \underbrace{[\mathbf{t}_0^s, \mathbf{t}_0^{a,0:K'}]}_{\text{chunk 0}}, \underbrace{[\mathbf{t}_1^s, \mathbf{t}_1^{a,0:K'}]}_{\text{chunk 1}}, \dots, \underbrace{[\mathbf{t}_{H-1}^s, \mathbf{t}_{H-1}^{a,0:K'}]}_{\text{chunk } H-1} \in \mathbb{R}^{H(1+K') \times D_{\text{vlm}}}. \quad (9)$$

The current state \mathbf{s}_t is not encoded here and continues to flow through the action head’s dedicated state encoder unchanged, preserving full backward compatibility with the base QWEN-ROBOTMANIP model.

History integration. We study two strategies for injecting context token sequence into the policy. In the *unified* mode, context tokens \mathbf{C} are appended to the end of the VLM input sequence and processed jointly with visual and language tokens under causal self-attention, allowing the VLM to reason over history, task description, and visual observations together. The resulting history-fused last-layer hidden states are passed to the DiT action head via cross-attention. In the *dual* mode, the state-action context is injected directly into the DiT action head rather than the VLM, keeping the VLM context length unchanged at the cost of shallower history integration. Unified injection allows the VLM’s full self-attention to jointly reason over behavioral history, task description, and visual observations, enabling richer cross-modal context integration than is possible when history is confined to the action head alone. We therefore adopt unified injection as the default configuration.

Stochastic context sampling. A naive strategy of always providing the H most recent chunks leads to a degenerate shortcut. Because the last context chunk is temporally closest to the current step, the model can achieve low training loss by simply copying the most recent action chunk rather than genuinely reasoning about the episode’s behavioral dynamics. This collapses the context mechanism into a trivial action-copy heuristic, which breaks down whenever the immediate history is ambiguous, atypical, or does not reflect the robot’s broader execution style. What we want the model to learn is the behavioral profile of the current episode (Huang et al., 2025), its velocity patterns, grasping strategies, and interaction signatures, not a shortcut based on temporal proximity. To prevent this, we introduce *stochastic context sampling* during training. Rather than always supplying the H chunks immediately preceding the current step, the context window is drawn from a random position within the episode. The sampled chunks may therefore be temporally distant from the current step, forcing the model to reason about the robot’s behavioral profile across the full episode rather than exploiting recency as a shortcut.

This randomization serves as a form of curriculum diversification. The model must learn to extract consistent behavioral style from *any* subset of the episode history, making it robust to missing, partial, or temporally displaced context at inference time. At deployment, we supply a rolling window of the most recent H chunks as context, allowing the model to leverage the full available history. Empirically, stochastic context sampling proves critical for preventing this collapse. Without it, the policy achieves low training loss but poor task success, a clear sign that the model has learned to copy recent actions rather than reason about execution context. With it, the model genuinely exhibits in-context adaptation, adjusting its behavior based on the broader behavioral dynamics of the current episode.

4 Training

4.1 Pre-training Recipe

4.1.1 Dual-Stream Co-Training

We train QWEN-ROBOTMANIP on two complementary data streams simultaneously. The VLA stream is built from the full multi-source manipulation corpus described in Section 2, comprising real-robot demonstrations, egocentric human-hand manipulation videos, and human-to-robot synthesized trajectories.

The VLM stream consists of large-scale vision-language supervision data (§§ 2.5), co-trained alongside the VLA stream to prevent the pre-trained perceptual and language capabilities from degrading under action prediction optimization, which would directly weaken the model’s ability to interpret novel instructions and generalize to unseen visual contexts (Driess et al., 2025). In practice, we adopt a 9:1 ratio of robot data to VL data.

4.1.2 Training Objectives

Flow matching loss. For each VLA sample, we are given a ground-truth action chunk $\mathbf{a} \in \mathbb{R}^{T \times D}$. Following the flow-matching formulation, we construct a noisy interpolant

$$\mathbf{x}_t = (1 - t) \boldsymbol{\epsilon} + t \mathbf{a},$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $t \sim \text{Beta}(1, 1.5)$. The action expert is trained to predict the corresponding velocity field $\mathbf{v} = \mathbf{a} - \boldsymbol{\epsilon}$, minimizing:

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{\mathbf{a}, \boldsymbol{\epsilon}, t} \|f_{\theta}(\mathbf{x}_t, t, \mathbf{s}, \mathbf{o}) - (\mathbf{a} - \boldsymbol{\epsilon})\|_2^2, \quad (10)$$

where f_{θ} denotes the full model conditioned on the proprioceptive state \mathbf{s} and the visual-language observation \mathbf{o} . Gradients from \mathcal{L}_{FM} are applied to both the VLM backbone and the action expert.

Because different robot embodiments populate different subsets of the 80-dimensional canonical action space (§§ 3.2), we apply a composed binary mask $\mathbf{m} \in \{0, 1\}^{T \times D}$ that restricts the objective to only the dimensions and time steps carrying valid supervision. The mask is constructed from three complementary sources. The *per-dimension slot mask* identifies actively populated dimensions for the current embodiment. For instance, a single-arm gripper populates the joint, end-effector, and gripper fields of one arm while leaving the opposite arm and hand slots as zeros. The *step validity mask* excludes time steps flagged as anomalous by the data curation pipeline (§§ 2.4) or outside episode boundaries, with all subsequent steps also masked once any step is deemed invalid to preserve causal consistency. For egocentric human data, a *per-hand validity mask* zeros out an entire arm slot from the moment the corresponding hand exits the camera view, preventing the model from being trained on occluded hand trajectories. The three masks are AND-combined, and the masked flow matching loss becomes a per-sample average over valid entries only:

$$\mathcal{L}_{\text{FM}} = \frac{1}{B} \sum_{i=1}^B \frac{\sum_{t,j} m_{i,t,j} (f_{\theta}(\mathbf{x}_{i,t}, t_i, \mathbf{s}_i, \mathbf{o}_i)_j - v_{i,t,j})^2}{\sum_{t,j} m_{i,t,j}}, \quad (11)$$

where B is the batch size and subscripts t, j index the time step and dimension. This formulation ensures that every sample in the batch contributes equally to the gradient regardless of how many dimensions are active, preventing embodiments with more populated slots from dominating optimization.

VLM next-token prediction loss. For each VLM sample, the backbone is trained with the standard autoregressive next-token prediction objective:

$$\mathcal{L}_{\text{VLM}} = -\mathbb{E} \sum_i \log p_{\phi}(y_i | y_{<i}, \mathbf{c}), \quad (12)$$

where y_i is the target response token at position i , and \mathbf{c} is the input context, which may be text-only or an interleaved sequence of images and text. The overall objective is

$$\mathcal{L} = \mathcal{L}_{\text{FM}} + \lambda \mathcal{L}_{\text{VLM}}, \quad (13)$$

where λ controls the relative weight of the two losses. We set $\lambda = 0.1$ so that VLM supervision provides stabilizing regularization without overshadowing action learning. Separate learning rates are used for the backbone and the action expert to account for their different initialization scales. To amortize the cost of the VLM forward pass, the action expert performs $K_{\text{repeat}}=8$ repeated diffusion steps per training sample, drawing 8 independent noise samples and timesteps for the same action chunk, and substantially improving training efficiency without increasing data consumption.

4.2 Post-Training Recipe

4.2.1 Domain-specific Supervised Fine-tuning

Once the foundation model has been pre-trained on the full heterogeneous corpus, we adapt it to specific deployment scenarios through supervised fine-tuning (SFT). Rather than training specialist policies for individual tasks, we adopt a *generalist SFT* paradigm. For each benchmark or real-world deployment scenario, all available demonstration data is combined into a single training set, producing one unified fine-tuned model that can execute every task within the target domain.

Compared with pre-training, SFT differs in several aspects. First, the SFT procedure only optimizes the flow matching objective \mathcal{L}_{FM} of eq. (11) without the VLM next-token prediction loss. We disable the multi-stage data curation filtering of §§ 2.4 and train on the complete unfiltered data, preserving every valid demonstration for post-training. We apply color jitter augmentation to the input images. In addition, with the model initialized from a pre-trained checkpoint, SFT is conducted on fewer GPUs with fewer training steps than pre-training.

4.2.2 Co-Training in Post-Training

Domain-specific SFT has become an important protocol for quantitatively evaluating the quality of a pre-trained VLA model. A strong pre-trained model is expected to adapt efficiently to a target domain after fine-tuning on the benchmark training set. However, this evaluation protocol can also expose a critical failure mode. After extensive benchmark-specific SFT, a VLA model may improve task performance by exploiting repeated visual and task patterns in the benchmark, while becoming less sensitive to the language instruction. In this case, the policy is no longer strongly conditioned on both vision and language. Instead, it behaves more like a vision-action pattern matcher. We refer to this phenomenon as VLA-to-VA degradation.

This degradation stems from three compounding factors. First, domain-specific SFT datasets are limited in diversity with concentrated visual layouts and instruction expressions, making shortcut correlations between scene patterns and actions easy to learn. Second, training and test splits share similar visual patterns, allowing high benchmark scores to be achieved through pattern memorization rather than genuine language grounding. As a result, such benchmarks overestimate instruction-following ability and fail to distinguish true language-conditioned control from benchmark-pattern overfitting. Third, models without strong compositional grounding ability tend to treat language as a weak context signal during SFT, with actions increasingly dominated by visual shortcuts learned from the benchmark data.

To mitigate this risk, we propose a mixed post-training strategy as an optional enhancement to the standard domain SFT described in §§ 4.2.1: rather than fine-tuning solely on the benchmark training set, the model is co-trained with a subset of pre-training data filtered by distributional proximity to the target domain. This provides broader adaptation signals while preserving the base model’s robust execution capabilities, without introducing unrelated data that might dilute domain-specific learning. In our main experiments, we follow the standard domain-specific SFT protocol to ensure fair comparison with baselines; the mixed post-training strategy is validated as an additional enhancement in §§ 6.5.1.

To directly evaluate whether the language-following capability is preserved, we construct a new benchmark, RoboTwin-IF (Instruction Following), based on RoboTwin 2.0 (Chen et al., 2025b). The benchmark examines whether the model performs the instructed action in the same or similar visual scenes with different instructions, rather than relying on visual pattern matching to select a default behavior. The detailed benchmark design is presented in §§ 6.2.1.

5 Deployment

In our deployment setup, inference is performed on a remote server with observations and actions transmitted between the robot and the server over a WiFi connection. To mitigate the latency introduced by cloud-based inference and network transmission, we adopt Real-Time Chunking (RTC) (Black et al., 2026), which asynchronously generates the next action chunk while the robot executes the current one, effectively hiding the round-trip latency and enabling smooth real-time control.

6 Experiments

6.1 Are Standard Benchmarks Enough?

We evaluate QWEN-ROBOTMANIP on robotic manipulation benchmarks that span a diverse range of embodiments and task types, and compare against several recent VLA models (Black et al., 2024; Bjorck et al., 2025; Community, 2026; Black et al., 2025). We begin with LIBERO (Liu et al., 2023) and RoboTwin (Mu et al., 2025), two standard benchmarks that are widely used for VLA evaluation. LIBERO (Liu et al., 2023) comprises four single-arm tabletop manipulation suites across 130 combinations of tasks and scenes. RoboTwin (Mu et al., 2025) presents 50 dual-arm manipulation tasks in easy and hard modes, requiring adaptation to varied backgrounds, objects, and spatial layouts.

Figure 4 (left) reports the results. A notable pattern is that models trained from scratch—StarVLA and Ours-scratch—attain competitive or even superior results compared to well-pretrained models such as $\pi_{0.5}$ and Abot-M0 on LIBERO and RoboTwin, despite lacking large-scale robotic pretraining. This is not a coincidence but a structural property of these benchmarks. Because training and evaluation data are drawn from the same environment and task distributions, high success rates can be achieved through in-distribution pattern matching alone. Models that lack genuine generalization can perform well simply by memorizing recurring visual and behavioral patterns, and the benchmark cannot distinguish this from real capability. Prior work has further shown that fine-tuning a pretrained VLA on a benchmark’s own training split yields performance comparable to training from scratch (Yan et al., 2025), confirming that

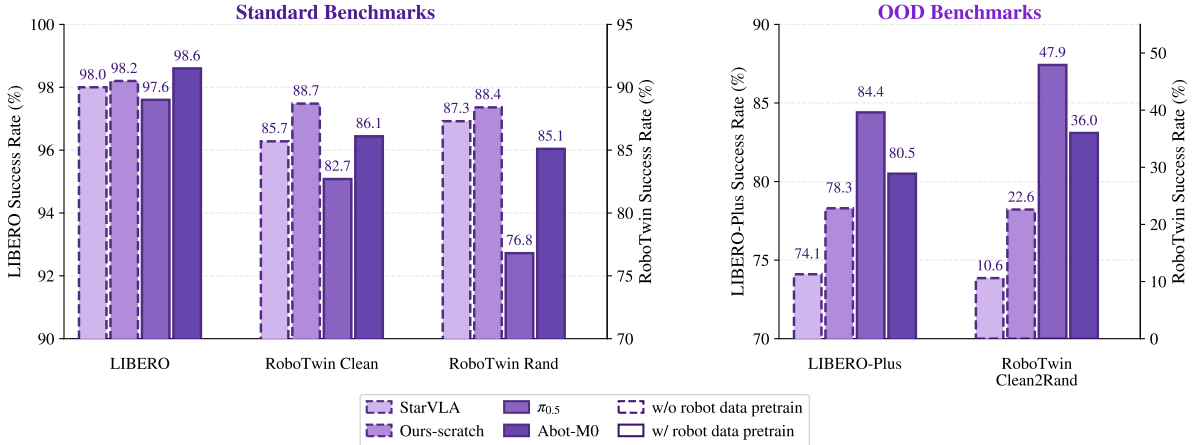


Figure 4: **Standard in-distribution benchmarks cannot reveal whether a model benefits from large-scale robot data pretraining.** Left: on in-distribution benchmarks (LIBERO, RoboTwin), models without large-scale robot pretraining (dashed borders) match or exceed pretrained ones. Right: on OOD benchmarks (LIBERO-Plus, RoboTwin-Clean2Rand), a clear separation emerges—pretraining provides genuine generalization that training-from-scratch cannot replicate.

the pretrained prior contributes negligible transferable value under in-domain evaluation.

Figure 4 (right) tells a different story. On OOD benchmarks—LIBERO-Plus and RoboTwin-Clean2Rand—where evaluation conditions diverge from training, a clear separation emerges: $\pi_{0.5}$ substantially outperforms StarVLA and Ours-scratch, with the gap widening as perturbation severity increases. StarVLA collapses from 85.7% (RoboTwin Easy, IID) to 10.6% (RoboTwin-Clean2Rand, OOD). This confirms that OOD evaluation is the correct north star for measuring foundation model quality: it reveals the transferable structure that pretraining provides and that in-domain metrics systematically fail to capture.

The way practitioners actually use a foundation model reinforces this conclusion. A researcher or engineer deploying a VLA model does not have access to the same distribution of tasks, objects, and environments present in any benchmark. They have a handful of demonstrations collected on their own hardware, in their own workspace, for their own task. What determines whether the foundation model helps them is not its in-domain benchmark rank but how much generalizable structure it has internalized, and how efficiently that structure transfers under minimal fine-tuning on unfamiliar data. The following section therefore adopts OOD evaluation as the primary measure and reports comprehensive benchmarking results for QWEN-ROBOTMANIP across both in-distribution and out-of-distribution settings.

6.2 Generalization Capabilities

The following sections evaluate QWEN-ROBOTMANIP on a suite of out-of-distribution settings that directly measure the generalization capabilities a foundation model is expected to provide. We organize these evaluations around three axes (Figure 5): task and scene generalization under controlled perturbations, instruction following with novel language and tasks, and zero-shot cross-embodiment transfer.

6.2.1 Evaluation Protocol

Task and Scene Generalization. We evaluate robustness of VLA models to visual and physical changes on four benchmarks. **LIBERO-Plus** (Fei et al., 2025) introduces controlled perturbations along seven orthogonal dimensions on top of the original LIBERO benchmark—background textures, camera viewpoints, language instructions, lighting conditions, object layouts, robot initial states, and sensor noise—with models fine-tuned on the standard LIBERO training set and evaluated under each perturbation. **RoboTwin-Clean2Rand** constructs an OOD evaluation protocol on top of RoboTwin (Mu et al., 2025): all models are fine-tuned exclusively on a Clean dataset with fixed white background, default lighting, no distractors, and a fixed table height, then evaluated under controlled randomizations along individual axes (background, lighting, clutter, table height) as well as a Hard setting that applies all randomizations simultaneously. **RoboCasa365** (Nasiriany et al., 2026) provides a large-scale kitchen manipulation benchmark with three evaluation suites of increasing difficulty: Atomic (18 basic manipulation skills with diverse object and layout variations), Composite-Seen (multi-step long-horizon tasks seen during

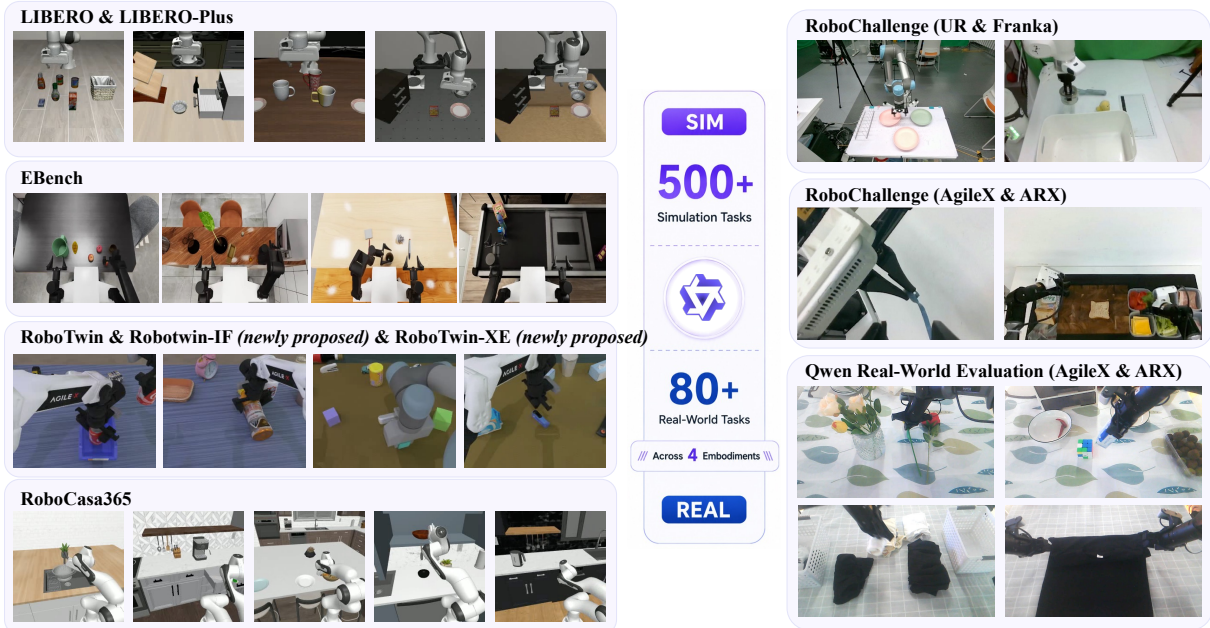


Figure 5: **Evaluation settings for QWEN-ROBOTMANIP**, spanning 500+ simulation tasks across LIBERO, LIBERO-Plus, EBench, RoboTwin, RoboTwin-IF, RoboTwin-XE, and RoboCasa365, and 80+ real-world tasks across 4 embodiments including UR, Franka, AgileX (ALOHA), and ARX.

training), and Composite-Unseen (long-horizon tasks absent from the training set). **EBench** (Laboratory, 2026) is an indoor mobile manipulation benchmark built on NVIDIA Isaac Sim, spanning 26 task types and 794 evaluation instances. The benchmark evaluates generalization along perturbation dimensions: background, instruction, object, and a mixed setting that combines all perturbations, reporting both success rate and a process score for each.

Instruction Following. Existing OOD benchmarks primarily probe robustness to visual and physical perturbations, while leaving generalization to unseen language instructions largely untested. We develop **RoboTwin-IF** (Instruction Following), a benchmark built on RoboTwin (Mu et al., 2025) that systematically evaluates instruction-following capabilities across five task suites, each targeting a distinct dimension of language grounding:

- **Pick-Diverse-Object:** Four objects are randomly sampled from a pool of 12 everyday items. The instruction names one object by color and noun. The robot must identify and lift the correct target among three distractors, testing *target-object grounding*.
- **Place-Relative:** Two named objects and 1–3 distractors are on the table. The instruction specifies picking up object A and placing it in a spatial relation (“beside” or “on top of”) with respect to object B, testing *spatial-relation understanding*.
- **Operate-Mic-Drawer:** A microphone and a cabinet with a functional drawer are present. The instruction specifies a multi-step bimanual sequence: open the drawer with one arm, then pick up the microphone with the other arm and place it inside. Some instructions further specify which arm performs which sub-task, testing *multi-step sequencing and bimanual coordination*.
- **Operate-Stapler:** A stapler, a colored pad, and 1–2 distractors are on the table. The instruction specifies either pressing the stapler or moving it onto the colored pad. The pad is always present regardless of the verb, acting as a distractor in press episodes and as the placement target in move episodes, testing *verb discrimination with shared scene elements*.
- **Operate-Tabletop:** A bell, a stapler, and 1–2 pickable objects are all present simultaneously. The instruction specifies one of three actions: ring the bell, press the stapler, or pick up a named object. Only one action is correct; the other interactive objects are distractors, testing *three-way verb-and-target discrimination* in a multi-affordance scene.

All models are fine-tuned exclusively on RoboTwin Clean. Each suite uses a two-tier language diversity system: “seen” instruction templates (used during training data collection) and “unseen” templates (held out for evaluation), combined with per-object description variants that further diversify noun phrases. At evaluation time, each episode is paired with a held-out unseen instruction template, ensuring zero

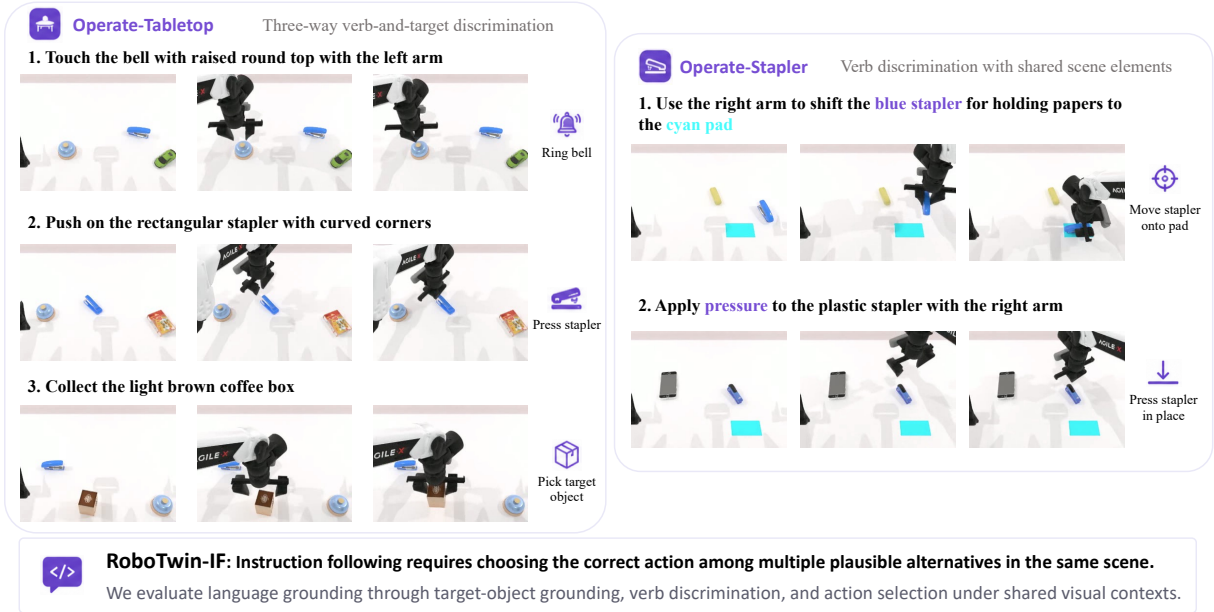


Figure 6: **Representative task suites from RoboTwin-IF.** Operate-Tabletop (left) requires three-way verb-and-target discrimination where a bell, stapler, and pickable objects are all present and only the instruction-specified action is correct. Operate-Stapler (right) requires verb discrimination under shared visual context, where the colored pad and stapler are always present regardless of the instructed verb. Both suites are evaluated on unseen instruction templates held out from training.

overlap with the training distribution.

Cross-Embodiment Generalization. Camera-frame relative EEF actions express deltas in the camera coordinate frame rather than a robot-specific joint space, enabling a single policy to potentially control morphologically distinct robots without re-training. We develop **RoboTwin-XE**, a benchmark based on RoboTwin that evaluates zero-shot transfer to unseen robot embodiments. The model is fine-tuned on the RoboTwin Clean dataset and tested under the RoboTwin Hard setting, replacing the default AgileX platform with ARX-X5, UR5-WSG, and Franka Panda. Initial EEF poses are aligned via IK, head and wrist camera extrinsics are kept identical, and task scenes, object layouts, and perturbation seeds are shared across embodiments. The model must therefore generalize across two types of embodiment-specific differences: the visual appearance of the robot arm in camera observations, and kinematic differences (DOF count, joint arrangement, link lengths, and workspace geometry) that affect reachability and motion dynamics. The model is fine-tuned exclusively on AgileX demonstrations; no target-embodiment data is used.

6.2.2 Summary of Results

On standard in-distribution benchmarks, QWEN-ROBOTMANIP achieves state-of-the-art or competitive performance (Table 3) on LIBERO (99.2%) and RoboTwin Easy/Hard (93.7%/94.0%). However, as argued in § 6.1, these metrics do not reliably distinguish genuine generalization from in-distribution pattern matching. We therefore focus on OOD evaluation as the primary measure of robotic foundation model capabilities.

Figure 7 summarizes results across benchmarks evaluated under OOD settings, where QWEN-ROBOTMANIP consistently outperforms prior state-of-the-art models by a substantial margin. QWEN-ROBOTMANIP outperforms $\pi_{0.5}$ on all benchmarks evaluated, with the gap widening as the evaluation becomes more challenging.

On task and scene generalization, QWEN-ROBOTMANIP surpasses $\pi_{0.5}$ by 7.0 points on LIBERO-Plus (91.4 vs. 84.4), by 21.5 points on the most demanding RoboTwin-C2R Hard setting (69.4 vs. 47.9), and by 18.5 points on EBench (45.6 vs. 27.1). On RoboCasa365, the advantage reaches 19.0 points (35.9 vs. 16.9).

On instruction following, QWEN-ROBOTMANIP scores 72.2% on RoboTwin-IF against 49.6% for $\pi_{0.5}$ —a large improvement of 22.6 points. This result is particularly significant because instruction following is the capability most prone to degradation during VLA training: fine-tuning a VLM on action prediction can erode the language conditioning pathway, causing the policy to collapse into a visually-triggered

Table 3: **In-distribution benchmark results.** QWEN-ROBOTMANIP matches or exceeds prior state-of-the-art across all standard benchmarks.

	LIBERO	RoboTwin-Easy	RoboTwin-Hard
π_0 (Black et al., 2024)	94.4	65.9	58.4
$\pi_{0.5}$ (Black et al., 2025)	97.6	82.7	76.8
StarVLA (Community, 2026)	98.0	85.7	87.3
Abot-M0 (Yang et al., 2026)	98.6	86.1	85.1
Being-H0.7 (Luo et al., 2026c)	99.2	90.2	89.6
QWEN-ROBOTMANIP-scratch	98.2	88.7	88.4
QWEN-ROBOTMANIP	99.1	93.4	92.5
QWEN-ROBOTMANIP-Context	99.2	93.7	94.0




Figure 7: **OOD generalization summary.** QWEN-ROBOTMANIP vs. previous state-of-the-art VLA model across three generalization axes. (a) Task and scene generalization under controlled perturbations. (b) Instruction following with held-out language templates. (c) Zero-shot cross-embodiment transfer. QWEN-ROBOTMANIP outperforms previous models on every OOD benchmark, with the gap widening on harder settings.

default behavior that ignores the instruction entirely (Section 4). The fact that QWEN-ROBOTMANIP maintains strong instruction following across all five RoboTwin-IF suites—including tasks requiring target-object grounding, spatial-relation understanding, and multi-way verb discrimination—indicates that the dual-stream co-training strategy and the diversity of the pretraining corpus together preserve genuine language-conditioned control.

The most revealing comparison is cross-embodiment transfer on RoboTwin-XE. When both models are trained exclusively on demonstrations collected on the AgileX ALOHA platform and evaluated zero-shot on unseen robots, QWEN-ROBOTMANIP achieves 23.9% using camera-frame EEF actions—3.2× the 7.5% achieved by $\pi_{0.5}$. This gap validates our camera-frame alignment strategy for action representation: by expressing actions in the visual domain, physically similar motions become numerically proximate across morphologically distinct robots, enabling effective cross-embodiment transfer.

A complementary signal comes from the comparison with models trained from scratch. On LIBERO-Plus, QWEN-ROBOTMANIP-scratch scores 78.3 versus QWEN-ROBOTMANIP’s 89.0 and QWEN-ROBOTMANIP-Context’s 91.4. The gap is even more pronounced on RoboTwin-C2R, where QWEN-ROBOTMANIP-scratch collapses from 71.6 (Easy) to 22.6 (Hard), retaining about 30% of its Easy performance, while QWEN-ROBOTMANIP degrades far more gracefully from 73.2 to 62.6, retaining roughly 86% (compared to 66% for $\pi_{0.5}$). A consistent pattern emerges across benchmarks. The pretrained VLM backbone already provides robustness to some of the visual and linguistic variations such as background, lighting, and language perturbations. However, the capabilities that distinguish genuine generalization from in-distribution memorization, including spatial reasoning under novel viewpoints, robustness to unseen robot states, and attention to task-relevant objects in cluttered scenes, specifically require the large-scale cross-embodiment pretraining that QWEN-ROBOTMANIP provides.

Table 4: **Out-of-distribution robustness evaluation on LIBERO-Plus across seven perturbation dimensions.**

	Camera	Robot	Language	Light	Background	Noise	Layout	Total
π_0 (Black et al., 2024)	13.8	6.0	58.8	85.0	81.4	79.0	68.9	53.6
$\pi_{0.5}$ (Black et al., 2025)	78.4	73.6	80.8	96.2	94.1	89.0	84.5	84.4
StarVLA (Community, 2026)	52.5	49.8	88.5	95.7	95.7	73.0	76.9	74.1
Abot-M0 (Yang et al., 2026)	60.4	67.9	86.4	96.2	91.6	86.4	82.6	80.5
Cosmos-Policy (Kim et al., 2026b)	75.8	63.3	81.7	96.5	88.9	92.7	82.2	82.2
Being-H0.7 (Luo et al., 2026c)	82.0	59.0	82.8	97.8	90.0	93.5	88.5	84.8
QWEN-ROBOTMANIP-scratch	70.4	44.9	88.1	95.8	95.5	84.4	79.1	78.3
QWEN-ROBOTMANIP	87.2	75.5	85.6	96.6	97.7	97.7	87.3	89.0
QWEN-ROBOTMANIP-Context	89.9	83.9	86.5	98.6	99.9	97.9	87.5	91.4

Table 5: **Out-of-distribution evaluation on RoboTwin-Clean2Rand.** Models are fine-tuned on the Clean dataset only and tested under various environmental randomizations.

	Easy	Background	Light	Clutter	Height	Hard
StarVLA (Community, 2026)	58.1	27.1	50.9	24.2	48.4	10.6
GR00T-N1.7 (Bjorck et al., 2025)	43.6	40.4	41.9	27.1	39.0	20.7
$\pi_{0.5}$ (Black et al., 2025)	73.1	67.0	69.2	57.9	67.6	47.9
Abot-M0 (Yang et al., 2026)	70.7	56.5	68.8	46.0	56.3	36.0
QWEN-ROBOTMANIP-scratch	71.6	60.6	70.7	24.6	63.6	22.6
QWEN-ROBOTMANIP (joint)	73.2	74.6	68.4	61.3	71.0	62.6
QWEN-ROBOTMANIP (eef)	74.0	75.8	70.1	59.8	69.4	60.8
QWEN-ROBOTMANIP-Context (joint)	84.7	82.4	84.2	75.4	79.5	69.4
QWEN-ROBOTMANIP-Context (eef)	85.0	82.4	84.7	66.8	82.9	64.0

6.2.3 Detailed Analysis

LIBERO-Plus. Table 4 reports per-dimension results. QWEN-ROBOTMANIP achieves 89% and QWEN-ROBOTMANIP-Context achieves 91.4% overall, surpassing all baselines. Beyond the aggregate score, the per-dimension breakdown reveals which capabilities benefit from large-scale robot data pretraining and which are already provided by the pretrained VLM backbone. Models without robot data pretraining (StarVLA, QWEN-ROBOTMANIP-scratch) match most pretrained models on Language, Light, and Background perturbations, indicating that the VLM’s visual and linguistic representations are inherently robust to these variations. In contrast, these same models suffer sharp degradation under Robot perturbation (49.8% and 44.9% vs. 75.5% for QWEN-ROBOTMANIP), where the policy must generalize across unseen initial robot states that lie outside the VLM’s purview. QWEN-ROBOTMANIP-Context further raises Robot perturbation robustness to 83.9%, a +8.4 point gain over QWEN-ROBOTMANIP, indicating that in-context history provides an implicit kinematic prior that helps the policy adapt to unfamiliar initial robot configurations within the episode. Camera viewpoint perturbation exhibits a similar pattern (+34.7 and +16.8 over the scratch models), suggesting that robust spatial reasoning under novel camera poses requires the grounding that VLA pretraining on diverse robot setups provides, rather than the appearance-level robustness already captured by the VLM.

RoboTwin-Clean2Rand. Table 5 reports results across individual and compound perturbation axes. QWEN-ROBOTMANIP achieves the highest success rate on the Hard setting under both joint-space control (62.6%) and end-effector control (60.8%), with both modes performing similarly across all perturbation dimensions. Among all methods, QWEN-ROBOTMANIP exhibits the smallest degradation from Easy to Hard, retaining roughly 86% of its Easy performance compared to 66% for $\pi_{0.5}$ and about 30% for models without robot data pretraining. The Clutter perturbation further highlights this gap: QWEN-ROBOTMANIP-scratch collapses from 71.6% to 24.6% when distractor objects are introduced, suggesting that the ability to attend to task-relevant objects amid clutter requires large-scale robotic data pretraining rather than in-domain fine-tuning alone. Notably, QWEN-ROBOTMANIP scores higher under Background randomization (74.6%) than under the Easy setting (73.2%), likely because diverse real-world scenes in the pretraining data make randomized backgrounds more in-distribution than the plain white background of the Easy mode. QWEN-ROBOTMANIP-Context further amplifies these gains, reaching 84.7% (Easy) and 69.4% (Hard) under joint control—an improvement of +11.5 and +6.8 points over QWEN-ROBOTMANIP. This substantial boost suggests that conditioning on intra-episode execution history provides complementary robustness to visual perturbations, as the policy can dynamically calibrate its actions based on observed outcomes rather than relying solely on the current observation.

Table 6: Evaluation on RoboCasa365 across atomic and long-horizon manipulation tasks.

	Atomic	Composite-Seen	Composite-Unseen	Total
π_0 (Black et al., 2024)	36.3	5.2	0.7	15.0
$\pi_{0.5}$ (Black et al., 2025)	39.6	7.1	1.2	16.9
GR00T-N1.5 (Bjorck et al., 2025)	50.7	14.8	2.7	23.9
GR00T-N1.6 (Bjorck et al., 2025)	51.1	9.4	1.7	21.9
RLDX-1 (Kim et al., 2026a)	63.0	27.5	5.4	33.2
QWEN-ROBOTMANIP	68.6	20.1	14.9	35.9
QWEN-ROBOTMANIP-Context	63.9	22.6	11.2	33.8

Table 7: Evaluation on EBench across three splits. Each split reports both success rate (SR) and the EBench composite score (Score).

	Table Top		Simple PnP		Long Horizon		Overall	
	SR	Score	SR	Score	SR	Score	SR	Score
π_0 (Black et al., 2024)	15.7	30	35.0	39	17.0	41	23.6	37
$\pi_{0.5}$ (Black et al., 2025)	12.9	32	45.0	50	18.1	39	27.1	41
X-VLA (Zheng et al., 2025)	8.6	24	50.0	54	6.2	25	23.7	36
InternVLA-A1 (Cai et al., 2026)	4.3	11	43.0	47	17.9	46	23.9	36
QWEN-ROBOTMANIP	50.0	70	56.5	60	29.9	55	45.6	60
QWEN-ROBOTMANIP-Context	49.3	56	55.0	66	26.6	55	43.6	59




Figure 8: EBench performance across operating mode, horizon, and precision.

RoboCasa365. Table 6 reports results across the three evaluation suites. QWEN-ROBOTMANIP achieves 35.9% overall, surpassing the previous state-of-the-art RLDX-1 (33.2%). On the Atomic suite, QWEN-ROBOTMANIP achieves the highest score (68.6%) among all methods, reflecting strong generalization across diverse manipulation primitives. On Composite-Unseen—where the robot must complete long-horizon tasks in OOD scenes—QWEN-ROBOTMANIP achieves 14.9%, nearly tripling the next-best result (5.4% for RLDX-1), demonstrating substantially stronger out-of-distribution compositional generalization.

EBench. Table 7 and Figure 8 report the results on EBench (Laboratory, 2026), an Isaac Sim-based indoor manipulation benchmark that evaluates 26 task types on a dual-arm mobile platform (Lift2 + R5a) across three splits: dexterous tabletop tasks (Table Top), pick-and-place with mobile manipulation (Simple PnP), and extended multi-step sequences (Long Horizon). QWEN-ROBOTMANIP achieves 45.6% overall success rate and a composite score of 60, outperforming $\pi_{0.5}$ (27.1% / 41) and all other baselines by a large margin across every split.

The gains are most striking on the dexterous split. On Table Top, QWEN-ROBOTMANIP achieves 50.0% SR and a score of 70, nearly quadrupling $\pi_{0.5}$'s success rate (12.9%) and more than doubling its score (32). On Simple PnP, QWEN-ROBOTMANIP leads at 56.5% SR while the next-best baseline is X-VLA at 50.0%. On Long Horizon, which requires mobile manipulation over extended sequences, QWEN-ROBOTMANIP reaches 29.9% SR and a score of 55, improving over $\pi_{0.5}$ by +11.8 SR and +16 score. Figure 8 further breaks down performance by operating mode, horizon length, and precision level, showing consistent

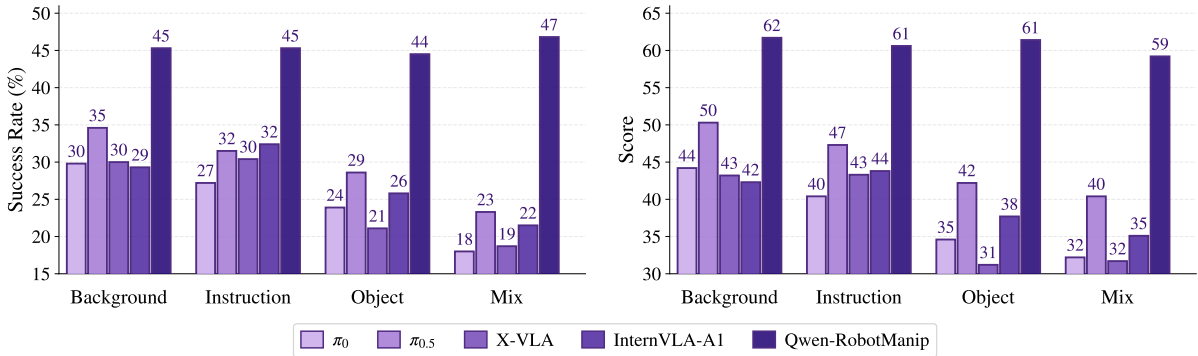


Figure 9: Per-dimension generalization breakdown on EBench.

Table 8: **Instruction-following evaluation on RoboTwin-IF.** Models are fine-tuned on RoboTwin Clean dataset and evaluated with held-out unseen instruction templates.

	Pick-Diverse	Place-Rel.	Ope.-Mic-Dr.	Ope.-Stapler	Ope.-Table	Average
StarVLA (Community, 2026)	11	13	0	49	74	29.4
GR00T-N1.7 (Bjorck et al., 2025)	20	17	0	14	32	16.6
$\pi_{0.5}$ (Black et al., 2025)	44	20	15	92	66	49.6
QWEN-ROBOTMANIP	79	57	42	90	93	72.2
QWEN-ROBOTMANIP-Context	77	71	33	89	90	72.0

advantages across all seven dimensions.

The per-dimension generalization breakdown (Figure 9) further confirms robustness to distribution shift. As perturbation complexity increases, baseline performance degrades substantially— $\pi_{0.5}$ ’s SR drops from 34.6 under Background to 23.3 under Mix (a 33% decline). In contrast, QWEN-ROBOTMANIP remains remarkably stable across all dimensions (SR 44.5–46.8), with virtually no degradation under compounded perturbations. The gap over the next-best method widens from Background (+10.7 SR) and Instruction (+12.9), to Object (+15.9) and Mix (+23.5). Under the most demanding Mix condition—where background, instruction, and object perturbations are applied simultaneously—QWEN-ROBOTMANIP achieves 46.8%, even surpassing its own Background score (45.3%), whereas $\pi_{0.5}$ drops from 34.6 to 23.3 (−33%). This near-uniform performance under diverse distribution shifts demonstrates strong generalization robustness.

RoboTwin-IF. Table 8 reports per-suite results. QWEN-ROBOTMANIP achieves 72.2% average versus $\pi_{0.5}$ ’s 49.6%, a gap of 22.6 points. The improvement happens on four of five suites, with large gains on Pick-Diverse (+35), Place-Relative (+37), Operated-Mic-Drawer (+27), and Operated-Tabletop (+27)—tasks where the instruction must be parsed to determine the correct action among multiple plausible alternatives in the same scene. This comprehensive advantage confirms that QWEN-ROBOTMANIP has learned genuine language-conditioned control rather than relying on visual shortcuts.

Zero-Shot Cross-Embodiment. Table 9 reports zero-shot transfer for both joint-space and camera-frame relative EEF representations on RoboTwin-XE. Joint-space control transfers poorly: joint configurations are robot-specific, so actions meaningful for one morphology produce near-random behavior on another—UR5 and Franka stay below 5% for both methods. Switching to camera-frame EEF dramatically improves transfer: QWEN-ROBOTMANIP reaches 42.9% on ARX, 22.8% on UR5 (5.6× the joint result), and 5.9% on Franka, for a cross-embodiment average of 23.9%. This confirms that the camera-frame representation successfully abstracts away morphological differences, allowing the policy to reason in shared Cartesian space. Our model consistently outperforms $\pi_{0.5}$ in both action spaces, with the largest gap on UR5 (22.8 vs. 10.0). The performance gradient (ARX > UR5 > Franka) correlates with visual and kinematic similarity to the training embodiment: ARX-X5 shares a visually similar 6-DOF form factor and similar reach to AgileX, UR5 differs in appearance and joint arrangement but has a similar workspace, while Franka’s distinct 7-DOF morphology and larger reach present the greatest mismatch.



Figure 10: **Zero-shot cross-embodiment evaluation on RoboTwin-XE.** Left: ARX-X5. Middle: UR5-WSG. Right: Franka Panda.

Table 9: **Zero-shot performance on RoboTwin-XE**, where models are trained on RoboTwin Clean and tested on Hard settings with unseen robot embodiments.

	ARX-X5	UR5-WSG	Franka Panda	Total
$\pi_{0.5}$ (joint) (Black et al., 2025)	24.6	2.2	0.9	9.2
$\pi_{0.5}$ (eef) (Black et al., 2025)	11.5	10.0	1.1	7.5
QWEN-ROBOTMANIP (joint)	37.6	4.1	1.8	14.5
QWEN-ROBOTMANIP (eef)	42.9	22.8	5.9	23.9



Figure 11: **In-domain and out-of-domain tasks of real-world CobotMagic ALOHA platform.**

6.3 Real-World Evaluation

6.3.1 Evaluation on ALOHA Platforms

In-Domain (ID) and Out-of-Domain (OOD) Evaluation on CobotMagic ALOHA. We fine-tune QWEN-ROBOTMANIP on 22.9 hours of teleoperated demonstrations collected on the CobotMagic ALOHA platform, covering multiple bimanual manipulation tasks. We then evaluate the fine-tuned policy on real-world ID and OOD benchmarks (Figure 11) constructed on the same platform.

The ID benchmark includes seven tasks: *table-cleanup* (clearing objects from the table), *three-bowl-stacking* (stacking three bowls in the specified order), *melon-in-bowl* (placing a small spherical melon toy into a bowl), *towel-folding* (folding a towel), *place-block-into-drawer* (opening a drawer, placing a block inside, and closing the drawer), *yellow-disc-insertion* (picking up a disc, performing a bimanual handover, and inserting it into a matching slot), and *three-block-stacking* (stacking three blocks in sequence). These tasks span a wide range of manipulation difficulty.

Table 10: **In-domain real-world evaluation on the CobotMagic ALOHA platform.** We report task success over 5 trials for each task.

Task	$\pi_{0.5}$	StarVLA	QWEN-ROBOTMANIP
table-cleanup	4/5	0/5	5/5
three-bowl-stacking	5/5	4/5	5/5
melon-in-bowl	2/5	0/5	5/5
towel-folding	4/5	3/5	4/5
block-in-drawer-compartment	0/5	0/5	5/5
yellow-disc-insertion	0/5	0/5	2/5
three-block-stacking	0/5	0/5	5/5
Average success rate	42.9%	20.0%	88.6%

Table 11: **Out-of-domain real-world evaluation on the CobotMagic ALOHA platform.** We report task success over 10 trials for each task.

Task	OOD factors	$\pi_{0.5}$	StarVLA	QWEN-ROBOTMANIP
target-object-in-basket	cluttered background, unseen objects, randomized targets	8/10	0/10	10/10
left-right-bowl-stacking	cluttered background, unseen objects, left-right reference	1/10	0/10	10/10
tool-on-towel	cluttered background, unseen small objects, distractors	0/10	0/10	6/10
banana-on-towel	dynamic lighting disturbance (disco light)	6/10	0/10	9/10
Average success rate	–	37.5%	0.0%	87.5%

At the easier end, *table-cleanup*, *three-bowl-stacking*, and *melon-in-bowl* primarily test basic language grounding and sequential pick-and-place execution. *Towel-folding* is more challenging due to the deformable nature of the object. The remaining tasks are substantially harder: *place-block-into-drawer* is particularly challenging because it requires the robot to accurately grasp the small drawer handle and perform precise pulling and pushing motions to open and close the drawer, *yellow-disc-insertion* requires fine-grained pose alignment for contact-rich insertion, and *three-block-stacking* demands reliable relational reasoning and robust multi-step execution, while also testing the model’s ability to recover from intermediate failures, since blocks may slip or collapse during the stacking process.

As shown in Table 10, QWEN-ROBOTMANIP achieves an average success rate of 88.6%, significantly outperforming $\pi_{0.5}$ (42.9%) and StarVLA (20.0%). It succeeds in all 5 trials on five tasks and remains strong on *towel-folding* (4/5). The only task with noticeable room for improvement is *yellow-disc-insertion* (2/5), highlighting the difficulty of precise insertion on real hardware. In comparison, $\pi_{0.5}$ performs reasonably well on relatively simple tasks such as *table-cleanup*, *three-bowl-stacking*, and *towel-folding*, but its performance degrades sharply on more challenging tasks that require long-horizon planning, precise contact-rich manipulation, or recovery from intermediate failures. StarVLA performs poorly across almost all tasks, with limited success even on the easier ones. These results indicate that QWEN-ROBOTMANIP not only handles basic real-world manipulation reliably, but also scales much better to more challenging task settings.

We further evaluate real-world generalization in an out-of-domain (OOD) setting on CobotMagic ALOHA. Compared with the ID benchmark, these tasks introduce distribution shifts in visual scenes, object instances, and task instructions, and are designed to test whether the model can robustly generalize beyond the seen distribution.

The OOD benchmark contains four tasks: *target-object-in-basket* (placing the instructed objects into a basket), *left-right-bowl-stacking* (stacking one bowl onto another based on left-right references), *tool-on-towel* (placing a specified tool onto a towel), and *banana-on-towel* (placing a banana onto a towel). Each task targets a different aspect of generalization. *Target-object-in-basket* introduces cluttered backgrounds, varied objects, and randomized target objects, testing whether the model can correctly identify and manipulate the instructed object under visual variation. *Left-right-bowl-stacking* combines cluttered scenes and left-right relational references, requiring robust spatial language understanding. *Tool-on-towel* further increases difficulty by introducing unseen and physically small objects (i.e., knife and spoon) together with distractors, making both target identification and stable grasping more difficult. Finally, *banana-on-towel* evaluates robustness to severe illumination changes induced by a disco light in the real-world environment.

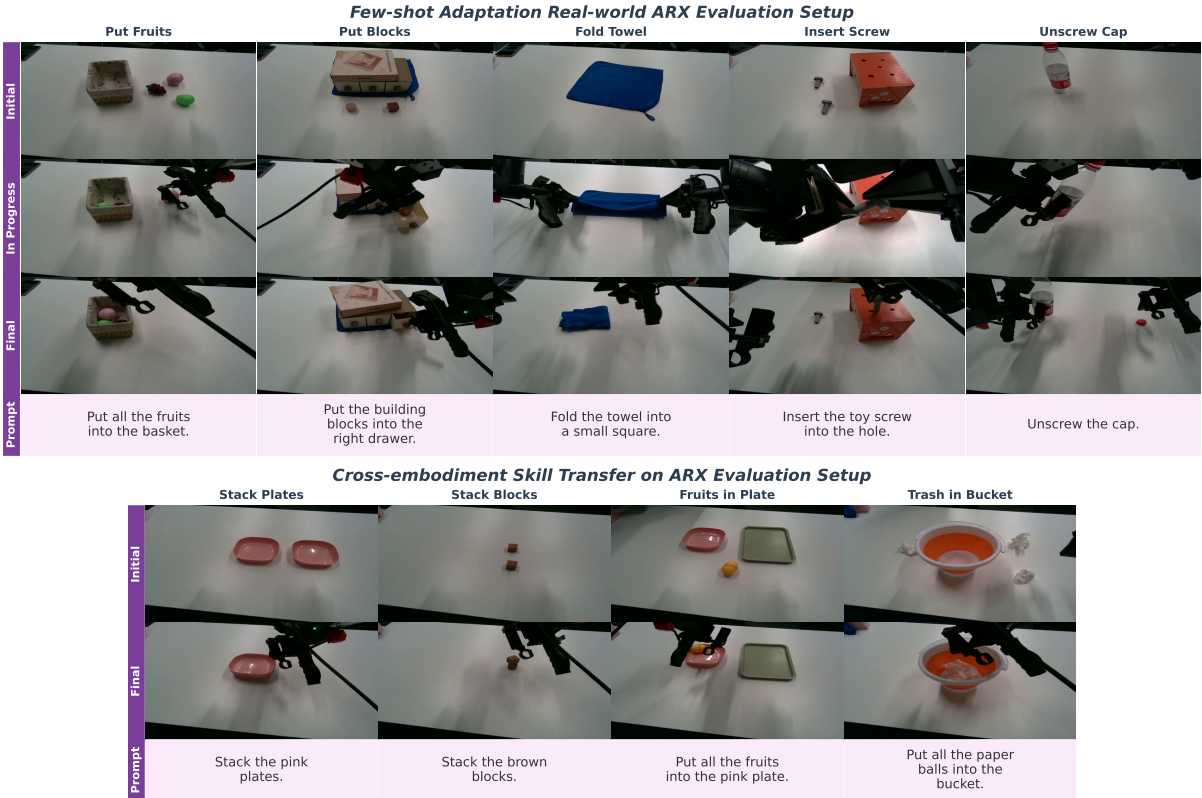


Figure 12: Real-world evaluation setup on the ARX ALOHA platform.

As shown in Table 11, QWEN-ROBOTMANIP generalizes substantially better than the baselines in the OOD setting, achieving an average success rate of 87.5%, compared with 37.5% for $\pi_{0.5}$ and 0.0% for StarVLA. In particular, QWEN-ROBOTMANIP attains perfect success on *target-object-in-basket* and *left-right-bowl-stacking*, demonstrating strong robustness to cluttered scenes, attribute-based grounding, and left-right spatial references. It also remains effective on the more difficult *tool-on-towel* task (6/10), where the robot must identify novel target objects under object distractors, and on *banana-on-towel* (9/10) under lighting variation.

In contrast, $\pi_{0.5}$ retains some robustness on relatively simpler OOD variations, achieving 8/10 on *target-object-in-basket* and 6/10 on *banana-on-towel*. However, its performance collapses on tasks requiring stronger compositional and relational generalization, obtaining only 1/10 on *left-right-bowl-stacking* and 0/10 on *tool-on-towel*. StarVLA fails on all four OOD tasks. These results suggest that QWEN-ROBOTMANIP preserves much stronger visual-linguistic representation and is considerably more robust to scene clutter, novel objects, spatial references, and lighting changes in real-world settings.

Few-shot Adaptation on ARX ALOHA. We compare QWEN-ROBOTMANIP against two baselines on five real-world manipulation tasks (Figure 12, top; Table 12): $\pi_{0.5}$ (Black et al., 2025), a pretrained open-source VLA, and StarVLA (Community, 2026) trained from scratch without pretraining. All methods use joint positions and EEF pose as state and predict actions in EEF space, jointly fine-tuned on the same 130 teleoperated demonstrations (50 for Unscrew Cap, 20 for each other task). The five tasks span multi-object pick-and-place (Put Fruits), long-horizon sequential manipulation (Put Blocks), deformable object handling (Fold Towel), bimanual precision assembly (Insert Screw), and fine-grained rotational control (Unscrew Cap).

QWEN-ROBOTMANIP outperforms both baselines on four of five tasks. The largest gain appears on Put Blocks (15/40 vs. 10/40 for $\pi_{0.5}$), where QWEN-ROBOTMANIP maintains higher sub-step success through all four stages—open drawer, place two blocks, and close drawer—indicating more robust long-horizon execution. On Unscrew Cap, QWEN-ROBOTMANIP doubles the cap-removal rate over $\pi_{0.5}$ (4/10 vs. 2/10) and triples the full-task completion (3/10 vs. 1/10), showing stronger fine-grained rotational control. Fold Towel improves notably on the second fold (3/10 vs. 1/10). Insert Screw remains challenging for all methods, with no model completing a full insertion (0/10), though both $\pi_{0.5}$ and QWEN-ROBOTMANIP succeed at the handover sub-step (2/10). StarVLA, without pretraining, achieves near-zero across tasks, confirming the importance of large-scale pretraining for real-world manipulation.

Table 12: **Few-shot adaptation on ARX ALOHA.** Sub-step success over 10 trials. All methods are jointly fine-tuned on 130 teleoperated demonstrations (50 for Unscrew Cap, 20 for each other task).

Task	Sub-step	StarVLA	$\pi_{0.5}$	QWEN-ROBOTMANIP
Put Fruits	Place fruit 1	3/10	9/10	9/10
	Place fruit 2	1/10	5/10	5/10
	Place fruit 3	0/10	2/10	3/10
	<i>Avg. success</i>	13.3%	53.3%	56.7%
Put Blocks	Open drawer	1/10	4/10	5/10
	Place block 1	1/10	2/10	4/10
	Place block 2	0/10	2/10	3/10
	Close drawer	0/10	2/10	3/10
	<i>Avg. success</i>	5.0%	25.0%	37.5%
Fold Towel	First fold	0/10	3/10	3/10
	Second fold	0/10	1/10	3/10
	<i>Avg. success</i>	0.0%	20.0%	30.0%
Insert Screw	Handover screw	0/10	2/10	2/10
	Insert screw	0/10	0/10	0/10
	<i>Avg. success</i>	0.0%	10.0%	10.0%
Unscrew Cap	Grasp bottle	4/10	9/10	9/10
	Unscrew cap	0/10	2/10	4/10
	Place down	0/10	1/10	3/10
	<i>Avg. success</i>	13.3%	40.0%	53.3%

Table 13: **Cross-embodiment skill transfer on ARX.**

	Stack Plates	Stack Blocks	Fruits in Plate	Trash in Bucket	Avg.
QWEN-ROBOTMANIP w/o UnifiedSpace	0/10	0/10	3/10	0/10	7.5%
QWEN-ROBOTMANIP w/o UnifiedEEF	0/10	0/10	5/10	0/10	12.5%
QWEN-ROBOTMANIP	3/10	5/10	7/10	7/10	55.0%

Cross-embodiment transfer. We also investigate cross-embodiment skill transfer (Figure 12, bottom; Table 13). A single policy is jointly fine-tuned on 6K CobotMagic and 130 ARX demonstrations, using the same state-action parameterization as above, then evaluated on four novel tasks on the ARX platform: stacking two plates, stacking blocks, placing fruits into a designated pink plate with distractor plates, and collecting paper balls into a bucket. ARX has *zero demonstrations* for any of these tasks—the relevant manipulation skills (stacking, precise placement, object collection) must generalize from related but not identical CobotMagic behaviors across a kinematically different embodiment. We ablate two key components: QWEN-ROBOTMANIP w/o UnifiedSpace removes the unified action-space mapping and instead naively concatenates and zero-pads each robot’s action dimensions without semantic alignment; QWEN-ROBOTMANIP w/o UnifiedEEF removes the unified EEF representation while retaining a basic slot layout with rotation unification and gripper normalization. Both variants fail almost entirely (7.5% and 12.5%), showing that neither surface-level dimension alignment nor partial unification can bridge the combined embodiment and task gap. QWEN-ROBOTMANIP achieves 55.0%—over $4\times$ the best variant—succeeding on all four tasks including Stack Blocks (5/10) and Trash in Bucket (7/10). The full unified action space and EEF representation enable QWEN-ROBOTMANIP to learn a stronger manipulation representation through large-scale diverse pretraining, enabling skill-level transfer: a policy can leverage skills learned from one embodiment to execute novel tasks on another embodiment with minimal training data and no task-specific demonstrations.

6.3.2 Table30-v1 Challenge

To evaluate the generalization capability of QWEN-ROBOTMANIP, we submit to the RoboChallenge Table30 v1 benchmark under the generalist track. This benchmark comprises 30 manipulation tasks distributed across 4 robot embodiments, and offers two evaluation tracks: the *specialist* track, where a dedicated policy is trained for each individual task (requiring 30 separate models), and the *generalist* track, where a single unified policy is trained per embodiment to handle all tasks associated with that robot. The generalist setting is substantially more challenging, as it demands that the policy generalize across diverse manipulation skills within each embodiment rather than overfitting to a single task. We focus on this track as it better reflects real-world desideratum of building versatile manipulation policies.

Table 14: **Per-task results on the RoboChallenge Table30 v1 Generalist Track.** Each cell reports success rate (%) / process score. Best results per task are in **bold**.

Robot	Task	QWEN-ROBOTMANIP	DM0_generalist	pi05_generalist	GR00T-MULTI	pi0_generalist
Average (all)		45 / 59.83	37 / 48.43	17.67 / 31.27	15.33 / 32.29	9 / 20.22
ARX5	arrange flowers	30 / 64	20 / 49	0 / 30.5	20 / 57	0 / 13.5
	arrange paper cups	70 / 83	10 / 54	0 / 31	0 / 19	0 / 15
	fold dishcloth	30 / 48	10 / 10.5	0 / 0	0 / 17	0 / 0
	open the drawer	0 / 47	90 / 95	50 / 80	0 / 50	0 / 20
	place shoes on rack	70 / 85	100 / 98.5	0 / 20	40 / 54.5	0 / 16.5
	put cup on coaster	100 / 99	100 / 100	70 / 63	80 / 92	0 / 0
	search green boxes	90 / 92	100 / 95.5	0 / 3	30 / 37.5	0 / 0
	sort electronic products	50 / 60.4	0 / 18.4	0 / 22.5	10 / 34.8	0 / 22.5
	turn on light switch	70 / 69.5	70 / 70.5	10 / 25	0 / 5	20 / 29
	water potted plant	0 / 9	0 / 33.5	0 / 0	0 / 6	0 / 0
wipe the table	0 / 72.5	0 / 47.5	10 / 28	0 / 66	0 / 29	
<i>Avg (ARX5)</i>		46.4 / 66.3	45.5 / 61.1	12.7 / 27.5	16.4 / 39.9	1.8 / 13.2
ALOHA	clean dining table	20 / 57.5	0 / 12	30 / 62	10 / 11.5	0 / 25.5
	make vegetarian sandwich	10 / 46.5	0 / 15	0 / 0	0 / 7	0 / 0
	plug in network cable	30 / 51	10 / 26	0 / 0	0 / 3	0 / 0
	pour fries into plate	30 / 56	0 / 6	0 / 0	0 / 36	0 / 0
	put opener in drawer	0 / 0	10 / 10	20 / 38	0 / 0	0 / 0
	put pen into pencil case	90 / 95	20 / 40	50 / 63.5	20 / 58	0 / 14.5
	scan QR code	10 / 13	0 / 0	0 / 7	30 / 26.5	0 / 3
	stack bowls	100 / 98.5	70 / 71	80 / 83	40 / 55.5	40 / 53.5
	stick tape to box	0 / 22.5	0 / 14	0 / 16	0 / 2	0 / 0
	sweep the rubbish	70 / 84.5	30 / 40	10 / 46	10 / 10	0 / 17
turn on faucet	100 / 100	70 / 84.5	60 / 56	20 / 44	60 / 67.5	
<i>Avg (ALOHA)</i>		41.8 / 56.8	19.1 / 29.0	22.7 / 33.8	11.8 / 23.0	9.1 / 16.5
UR5	arrange fruits in basket	80 / 89.5	70 / 87	0 / 9	30 / 54.5	0 / 11.5
	hang toothbrush cup	60 / 80	90 / 95	50 / 71	70 / 85	20 / 62
	set the plates	100 / 87.5	60 / 62	40 / 49.5	0 / 24	50 / 69.5
	shred scrap paper	0 / 15	30 / 45	20 / 36	0 / 0	20 / 27
	sort books	0 / 9	0 / 8.5	0 / 24	0 / 6.5	10 / 26.5
	stack color blocks	70 / 84.5	100 / 100	10 / 30	0 / 44.5	30 / 39
<i>Avg (UR5)</i>		51.7 / 60.9	58.3 / 66.2	20.0 / 36.6	16.7 / 35.8	21.7 / 39.2
Franka	move objects into box	70 / 75.5	50 / 64.5	20 / 40	50 / 62	20 / 44.5
	press three buttons	0 / 0	0 / 0	0 / 4	0 / 0	0 / 0
<i>Avg (Franka)</i>		35.0 / 37.8	25.0 / 32.2	10.0 / 22.0	25.0 / 31.0	10.0 / 22.2

Specifically, we post-train QWEN-ROBOTMANIP on the demonstration data provided by Table30 v1 using joint control, and submit under the anonymous identity Lira_generalist². QWEN-ROBOTMANIP achieves a success rate of 45% and a process score of 59.83. Notably, QWEN-ROBOTMANIP surpasses DM0_generalist (37% success rate, 48.43 process score) by 8 percentage points in success rate and 11.4 in process score, demonstrating its strong multi-task generalization ability. We highlight three key findings from detailed analysis of the benchmark results.

Strong bimanual coordination. Among the 30 benchmark tasks, 8 require tight bimanual coordination on the ALOHA platform³, where the two arms must jointly stabilize, transport, and manipulate objects. QWEN-ROBOTMANIP achieves an average success rate of 40% on these tasks, far exceeding $\pi_{0.5}$ (21.2%), DM0 (16.2%), GR00T-MULTI (7.5%), and π_0 (7.5%) (Figure 14, left). Notably, QWEN-ROBOTMANIP is the *only* model to succeed on *pour fries into plate* (30% vs. 0% for all baselines), a task demanding sequential bimanual steps—stabilizing the fries box with the left arm, opening it with the right arm, picking it up, and pouring the contents onto the plate. As illustrated in Figure 13, QWEN-ROBOTMANIP completes this full sequence, while DM0_generalist fails at the initial coordination stage: the right arm never reaches a graspable position on the box, resulting in immediate task failure. We attribute this strong bimanual performance to two factors: (1) our pre-training corpus contains a substantial proportion of bimanual demonstration robot data, enabling the model to effectively learn coordinated dual-arm control primitives; and (2) our Human2Robot pipeline, which synthesizes bimanual robot data from egocentric human videos, further expands the effective bimanual pre-training data and exposes the model to diverse manipulation strategies beyond those captured in teleoperated demonstrations alone.

²See Lira_generalist in <https://robochallenge.cn/home>.

³The 8 bimanual tasks: *clean dining table*, *make vegetarian sandwich*, *pour fries into plate*, *put opener in drawer*, *put pen into pencil case*, *stick tape to box*, *sweep the rubbish*, *turn on faucet*.

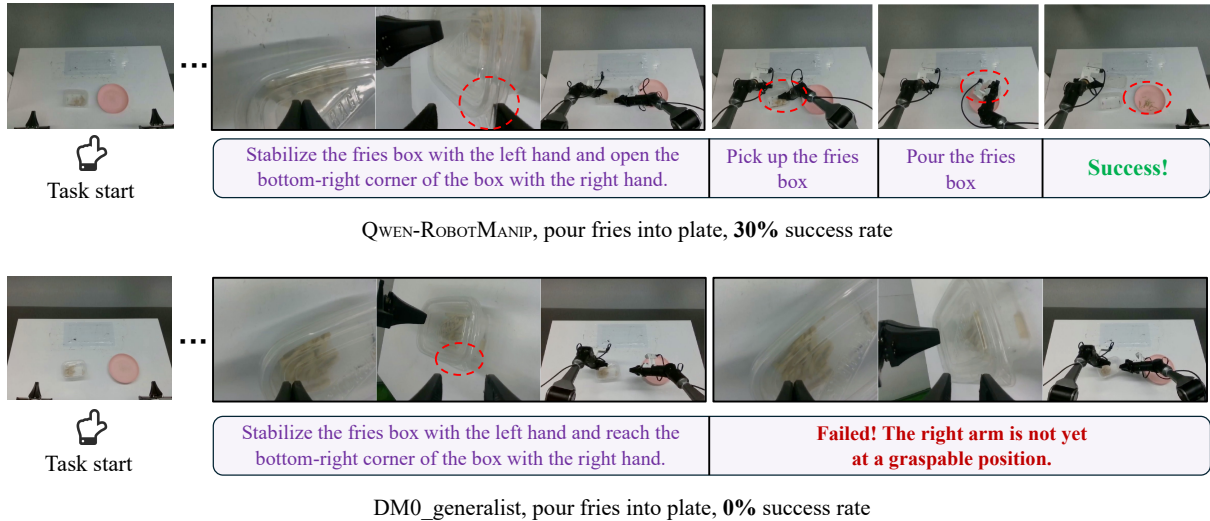


Figure 13: **Case study on *pour fries into plate (ALOHA)***. QWEN-ROBOTMANIP (top) coordinates both arms to stabilize, open, pick up, and pour the fries box, completing the task successfully. DM0_generalist (bottom) fails at the initial stage as the right arm never reaches a graspable position.

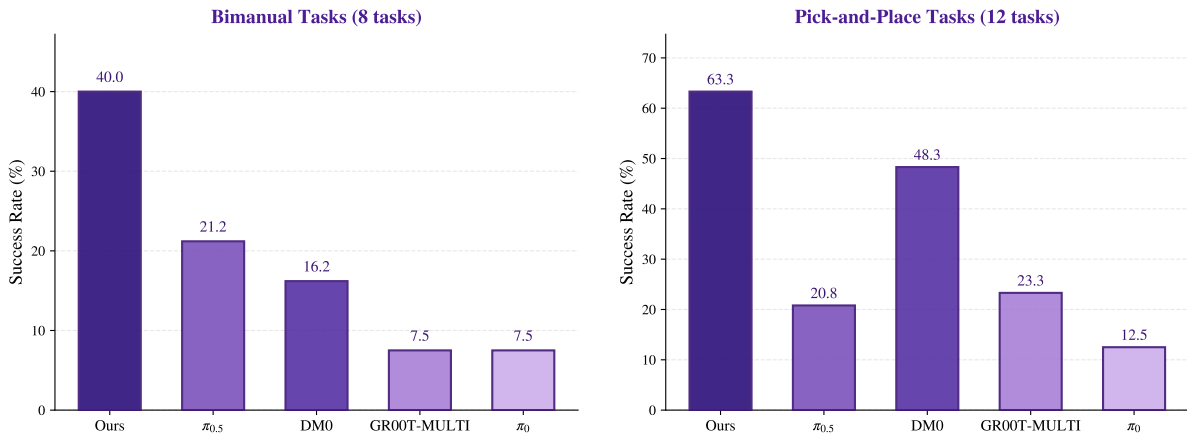


Figure 14: **Average success rate on bimanual coordination tasks (left, 8 tasks) and pick-and-place tasks (right, 12 tasks)**. QWEN-ROBOTMANIP substantially outperforms all baselines in both categories.

Robust pick-and-place across embodiments. We identify 12 tasks across all four platforms that center on pick-and-place primitives⁴, ranging from single-object grasping (*put cup on coaster* and *stack color blocks*) to multi-step sequential manipulation involving 4–5 objects (*arrange paper cups* and *sort electronic products*). As shown in Figure 14 (right), QWEN-ROBOTMANIP achieves 63.3% average success rate on these tasks, surpassing the next-best baseline DM0 (48.3%) by 15.0 percentage points. We attribute this capability to two factors: (1) the large-scale cross-embodiment pre-training data encodes abundant pick-and-place patterns, and (2) the unified action space enables knowledge sharing of fundamental spatial skills across different robot morphologies. Figure 15 shows a representative comparison on *arrange paper cups*, a task requiring sequential pick-and-place of multiple cups followed by stacking. QWEN-ROBOTMANIP (70% SR) successfully picks each cup in sequence, stacks them precisely, while $\pi_{0.5}$ _generalist (0% SR) encounters a stuck cup during stacking and fails to recover, leading to cascading errors in subsequent steps.

Emergent retry behavior. A recurring pattern we observe during real-robot evaluation is that QWEN-ROBOTMANIP exhibits spontaneous retry behavior: when an initial manipulation attempt fails (e.g., a grasp slips or a placement misses) the policy autonomously re-attempts the action rather than proceeding to the next step or stalling. While this behavior is difficult to quantify with a single metric, we observe

⁴The 12 pick-and-place tasks are: *arrange flowers*, *arrange fruits in basket*, *arrange paper cups*, *clean dining table*, *move objects into box*, *set the plates*, *sort books*, *sort electronic products*, *stack bowls*, *place shoes on rack*, *put cup on coaster*, and *stack color blocks*.

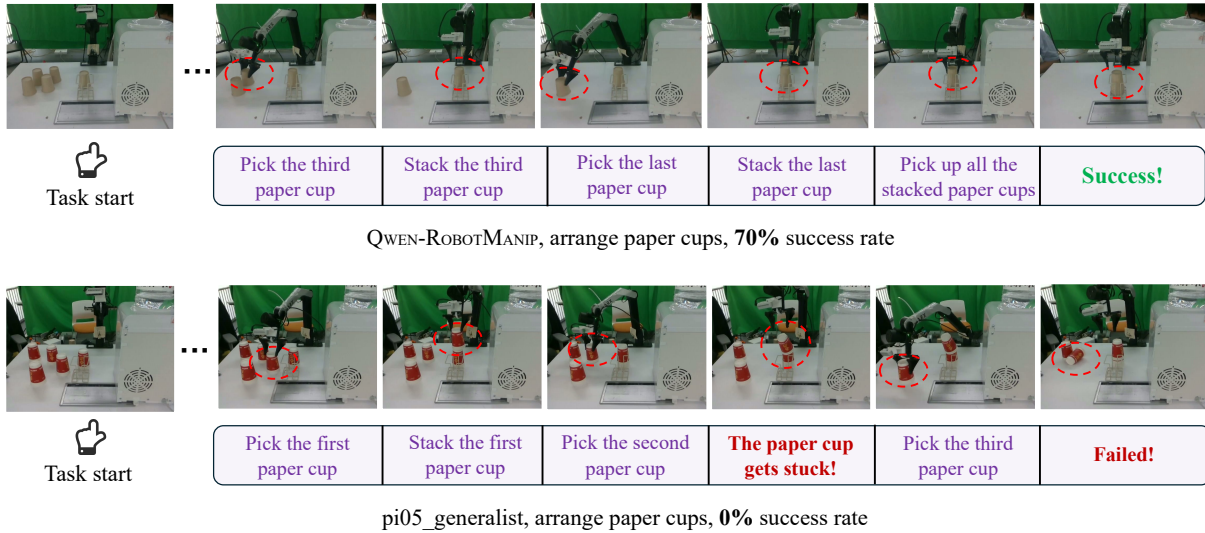


Figure 15: **Case study on *arrange paper cups* (ARX5).** QWEN-ROBOTMANIP (top) sequentially picks, stacks, and collects all paper cups with precise placement. $\pi_{0.5_generalist}$ (bottom) encounters a stuck cup during stacking and fails to recover.

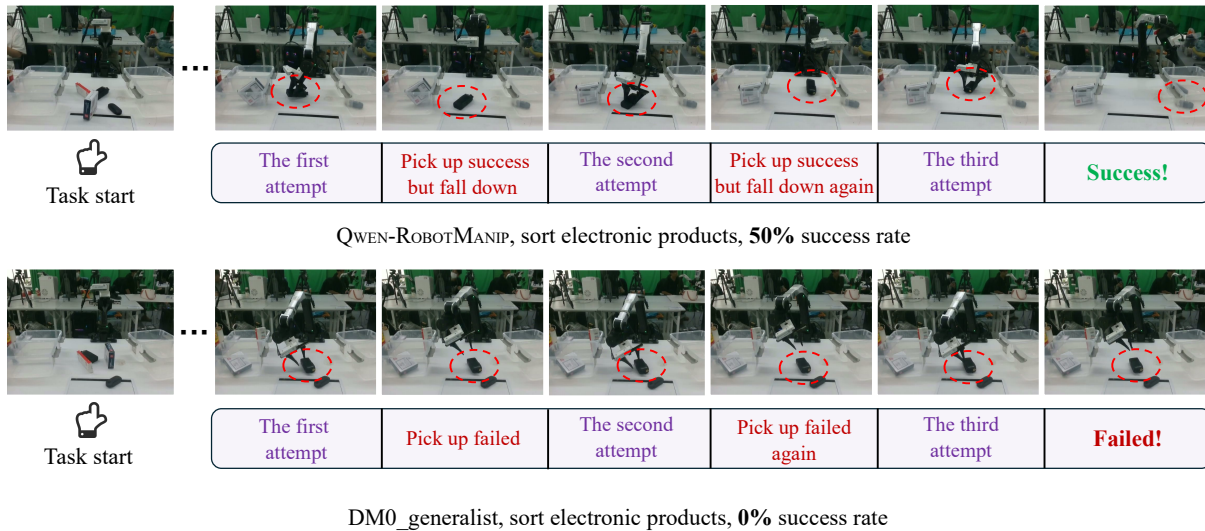


Figure 16: **Case study on *sort electronic products* (ARX5).** QWEN-ROBOTMANIP (top) picks up the object but it falls twice; the policy autonomously retries and succeeds on the third attempt. DM0_generalist (bottom) also attempts three times but never achieves a secure grasp.

it consistently across diverse action primitives including picking, placing, pouring, folding, wiping, and sweeping. This self-corrective capability significantly improves fault tolerance, allowing QWEN-ROBOTMANIP to recover from intermediate failures that would otherwise cause task-level failures. Figure 16 provides a vivid example on *sort electronic products*: QWEN-ROBOTMANIP successfully picks up the object on its first attempt, but the object falls; it then re-attempts, and the object falls again; on the third attempt, QWEN-ROBOTMANIP finally completes the grasp and places the object into the target bin, achieving task success (50% SR). In contrast, DM0_generalist (0% SR) also attempts the pick three times but fails every attempt—the gripper never achieves a secure grasp, leading to task failure. We hypothesize that this retry behavior emerges from the diversity of the pre-training data, where demonstrations naturally contain imperfect attempts followed by corrections, enabling the model to learn recovery strategies as part of its policy.

Breakthrough on challenging tasks. The above three capabilities jointly enable QWEN-ROBOTMANIP to achieve leading performance on several high-difficulty tasks where most generalist baselines fail entirely (see Figure 17). On the ALOHA platform, *plug in network cable* (30%) and *pour fries into plate* (30%) both demand precise bimanual coordination with fine posture adjustment, yet QWEN-ROBOTMANIP

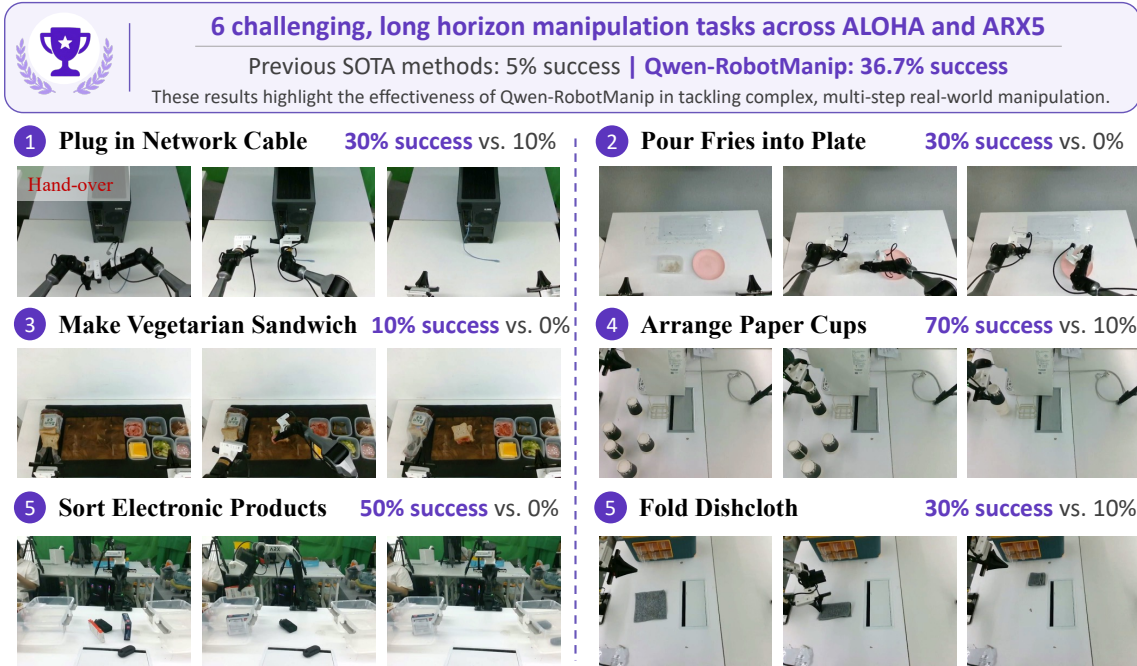


Figure 17: QWEN-ROBOTMANIP on six challenging long-horizon tasks from RoboChallenge Table30-v1, spanning bimanual dexterous manipulation, sequential ingredient stacking, multi-object arrangement, and deformable object handling. Prior SOTA generalist methods achieve an average of 5% success across these tasks. QWEN-ROBOTMANIP achieves 36.7%, with substantial margins on every task.

is the only model achieving $\geq 30\%$ success rate. On *make vegetarian sandwich* (10%)—a long-horizon task requiring sequential ingredient stacking—QWEN-ROBOTMANIP is the only model to achieve a non-zero success rate. On ARX5, spatial precision and self-correction combine to yield strong results on *arrange paper cups* (70%), *sort electronic products* (50%), and *fold dishcloth* (30%), outperforming the next-best method by 60, 40, and 20 percentage points respectively. These results demonstrate that cross-embodiment pre-training enables the interleaved acquisition of bimanual coordination, spatial precision, and self-corrective strategies, which collectively unlock complex manipulation skills that remain out of reach for existing generalist policies.

6.4 Ablation Study

We ablate the core design choices of QWEN-ROBOTMANIP: alignment strategies (state-action representation, in-context adaptation, architecture) and data recipes for scaling (human-to-robot synthesis, vision-language co-training). All variants within each group share identical training conditions at reduced pretraining scale; configurations may differ from the main results.

Action space alignment for data scaling. A central promise of foundation models is that performance should improve predictably as training data grows. For vision-language-action models trained on heterogeneous cross-embodiment data, verifying this property requires careful design of how data from distinct robots is combined. If the action representations across embodiments are misaligned, simply adding more data may not yield increasing performance, because the model must spend capacity reconciling conflicting conventions rather than learning transferable manipulation structure.

To investigate this, we construct a controlled data-scaling experiment. Starting from the full cross-embodiment pre-training mixture, we create nested subsets at 1%, 5%, 10%, 25%, 50%, and 100% of the original dataset by sampling at the per-embodiment task level, ensuring that smaller subsets are strict subsets of larger ones. All models are evaluated on a fixed held-out OOD evaluation set spanning 15 embodiment types and 154 tasks that are entirely absent from any training set. For each data percentage and model variant, we report the best validation MSE achieved across all training checkpoints, measured in a unified action representation to ensure fair comparison across model variants.

Figure 18 compares the scaling behavior of three action space designs. **Ours w/o UnifiedSpace** concatenates each embodiment’s raw action fields and zero-pads to 80 dimensions without any cross-embodiment structural alignment. **Ours w/o UnifiedEEF** maps joints, end-effectors, grippers, and hands into semantically fixed positions within the 80-dimensional canonical vector (§§ 3.2), with end-effector actions

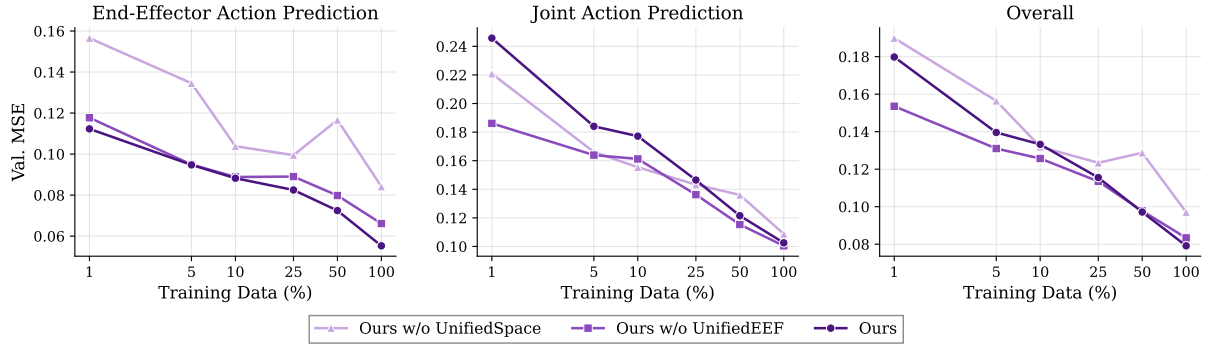


Figure 18: **Data scaling curves for models with varied state and action representations**, evaluated on held-out validation datasets. Each point reports the best validation MSE across all training checkpoints for a given training data percentage.

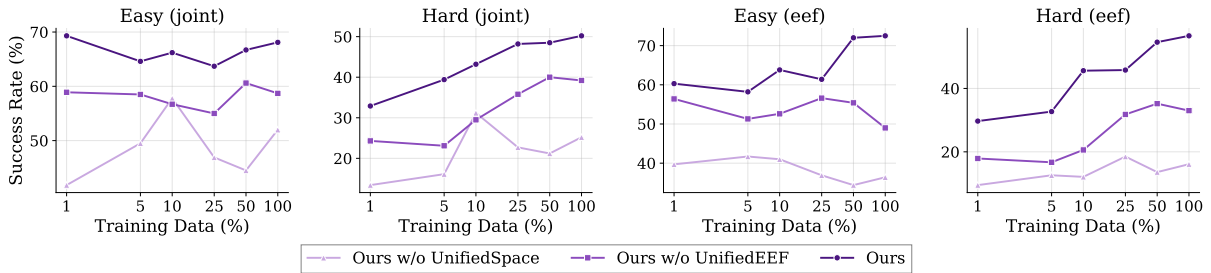


Figure 19: **Downstream performance on RoboTwin-C2R after fine-tuning models pre-trained with varied data percentages and action representations**. Each subplot reports success rate under a different control mode (joint or EEF) and evaluation setting (Easy or Hard).

expressed as axis-angle deltas relative to the initial end-effector pose. **Ours** further unifies end-effector motion by expressing it as camera-frame delta poses (§§ 3.3), grounding actions in the visual observation frame.

Models with unified representations, **Ours w/o UnifiedEEF** and **Ours**, both exhibit a clear data scaling law: their best validation MSE decreases approximately log-linearly with training data volume across the full 1%–100% range, confirming that with proper action-space alignment, scaling up cross-embodiment data consistently reduces OOD prediction error. **Ours w/o UnifiedSpace**, by contrast, shows notably different behavior. For action prediction on end-effector dimensions (Figure 18, left), **Ours w/o UnifiedSpace** produces an unstable scaling curve with substantially higher MSE than the two unified variants. **Ours** achieves the lowest end-effector MSE on average, demonstrating that camera-frame alignment enables the most effective cross-embodiment transfer for end-effector control.

These offline prediction improvements translate directly into downstream task performance. Figure 19 reports success rates on RoboTwin-C2R after fine-tuning each pre-trained variant on the RoboTwin Clean dataset for 80k steps, evaluated under four settings that combine two control modes (joint-space and end-effector) with two perturbation settings (Easy and Hard). Colors from dark to light correspond to progressively less aligned action representations: **Ours**, **Ours w/o UnifiedEEF**, and **Ours w/o UnifiedSpace**. On the OOD Hard settings (second and fourth subplots), **Ours** exhibits a clear data scaling property: downstream success rate increases steadily as pretraining data grows from 1% to 100%, reaching 50.2% (joint) and 56.6% (EEF) at full scale. **Ours** consistently outperforms both ablated variants across all data percentages, while **Ours w/o UnifiedEEF** and **Ours w/o UnifiedSpace** show noisier scaling curves with less consistent improvement. In contrast, on the in-distribution Easy settings (first and third subplots), none of the three variants shows a clear upward trend with increasing pretraining data. This corroborates the finding in §§ 6.1 that standard in-distribution evaluation fails to capture the benefits of large-scale pretraining, and that OOD evaluation is needed to reveal the genuine scaling behavior. A further notable pattern emerges in the control-mode comparison: **Ours** achieves comparable or higher success rates under EEF control than under joint control on both Easy and Hard, while both ablated variants show the opposite trend. This confirms that camera-frame alignment produces a strong EEF-space policy, a capability essential for cross-embodiment transfer.

Ablation study on embodiment prompt design and in-context policy adaptation. We ablate the

Table 15: **Ablation study on embodiment prompt design and in-context policy adaptation.** Out-of-distribution evaluation on RoboTwin-Clean2Rand (joint). The speed is set to 500.

	Emb. Tag	FPS	Context	Denoise Steps	Easy	Hard	Avg
QWEN-ROBOTMANIP w/o UnifiedEEF	×	×	×	4	71.2	54.2	62.7
+ Soft-Prompt	soft	×	×	4	70.2	52.1	61.2
+ Language Prompt	lang.	15	×	4	71.7	55.1	63.4
+ Structure Prompt	✓	15	×	4	73.4	58.3	65.9
QWEN-ROBOTMANIP-Context	✓	15	✓	4	72.1	54.4	63.3
QWEN-ROBOTMANIP-Context	✓	15	✓	10	80.1	61.6	70.9
QWEN-ROBOTMANIP-Context	✓	15	✓	20	79.8	62.1	71.0

embodiment prompt design and in-context policy adaptation mechanism using an early checkpoint of QWEN-ROBOTMANIP trained on a smaller data subset, evaluated on RoboTwin-Clean2Rand under out-of-distribution settings. Results are reported in Table 15.

On the prompt side, encoding embodiment identity as a learnable soft prompt slightly degrades performance relative to the naïve baseline. A natural language prompt recovers this and yields a modest gain, suggesting that semantic embodiment labels carry some useful signal, but the improvement remains small, indicating that coarse identity information alone is insufficient to meaningfully adjust behavior across embodiments. The structured embodiment prompt, which additionally conditions the model on the temporal properties of the episode via the FPS field, produces consistent improvements of 2.2-3.2 points. The structured prompt provides robustness under distribution shift rather than overfitting to a fixed temporal configuration, but is most effective when its fields faithfully reflect the deployment context.

Prompt-level conditioning, however, captures only static episode metadata. The in-context policy adaptation mechanism introduces dynamic intra-episode information in the form of a single historical observation-action chunk, enabling the model to observe the robot’s actual behavioral dynamics rather than relying on coarse category labels. Without context, the policy action distribution is smooth and 4 denoising steps is already sufficient, producing stable behavior with no signs of jitter. Introducing execution history increases the complexity of the action distribution, causing the policy to produce jittery motion at the same 4-step budget, which largely negates the benefit of the context signal and yields performance near the naïve baseline. Increasing the denoising budget to 10 steps resolves the instability and unlocks the full benefit of the context mechanism, reaching 70.9 average and a 5.0-point improvement over the structure prompt baseline, a margin that dwarfs the contribution of any prompt design variant. Further increasing to 20 steps yields no additional gain. These results support the view that in-context history functions as an implicit embodiment identifier reflecting intra-episode kinematic signatures rather than as episodic memory, and that this signal is qualitatively distinct from what static prompt conditioning can provide, provided the action head has sufficient capacity to decode it faithfully.

We do note one practical limitation observed in real-robot deployment: at the start of an episode, the context consists entirely of zero-padded placeholders, and the model, having learned to condition on quiescent history, tends to hesitate before initiating motion. We therefore release both the context-conditioned and context-free variants of QWEN-ROBOTMANIP, allowing users to select appropriate configuration depending on whether rapid response or richer intra-episode adaptation is the priority.

Ablation study on human-to-robot synthetic data. To isolate the contribution of egocentric human data, we compare three pretraining configurations on RoboTwin-Clean2Rand (eef) and LIBERO-Plus at a fixed 7:3 robot-to-auxiliary ratio (Tables 16 and 17): **Robot-only** uses robot data only; **+Ego** mixes in raw egocentric data; **+H2R** replaces the raw data with pipeline-synthesized robot demonstrations from the same ego sources. All three share identical training steps, hyperparameters, and finetuning data, so any performance difference is attributable solely to the auxiliary data source.

On RoboTwin-Clean2Rand, the Hard setting—where all perturbations are applied simultaneously—shows the clearest separation: +H2R reaches 58.7%, a +4.0 gain over Robot-only (54.7%) and +3.7 over +Ego (55.0%). Per-dimension gains are consistent, with the largest on Light (+3.0) and Height (+3.2), where the varied camera perspectives and lighting in ego data naturally augment robustness. The Easy setting also sees a modest improvement (72.9 → 73.4 → 74.2).

On LIBERO-Plus, +H2R raises the average success rate from 87.1% to 89.0%. The Camera dimension shows the largest improvement (+7.2 over Robot-only, 72.8 → 80.0), as egocentric data provides diverse viewpoint coverage that robot-only data lacks. The Robot dimension also benefits (+2.0), suggesting that the diverse manipulation trajectories in H2R data improve robustness to initial pose variation. The

Table 16: Human2Robot ablation on RoboTwin-Clean2Rand (eef).

	Easy	Background	Light	Clutter	Height	Hard
QWEN-ROBOTMANIP (robot-only)	72.9	70.4	70.3	57.2	67.8	54.7
QWEN-ROBOTMANIP (+ego)	73.4	70.6	71.7	59.2	70.2	55.0
QWEN-ROBOTMANIP (+h2r)	74.2	71.4	73.3	58.1	71.0	58.7

Table 17: Human2Robot ablation on LIBERO-Plus.

	Camera	Robot	Language	Light	Background	Noise	Layout	Total
QWEN-ROBOTMANIP (robot-only)	72.8	78.2	88.7	97.5	97.6	95.4	85.4	87.1
QWEN-ROBOTMANIP (+ego)	77.7	79.0	88.5	98.8	98.7	97.3	84.9	88.4
QWEN-ROBOTMANIP (+h2r)	80.0	80.2	89.3	98.5	98.2	96.6	85.2	89.0

Table 18: Ablations on VL data co-training in the pre-training and post-training stage. By default, QWEN-ROBOTMANIP is pre-trained with VL data and post-training without it. (“RT” is the abbreviation of “RoboTwin”.)

	LIBERO	LIBERO-Plus	RT-C2R (easy)	RT-C2R (hard)	RT-IF
QWEN-ROBOTMANIP	99.1	90.1	73.2	62.6	71.6
- without VL data in pre-training	98.2	88.9	66.5	54.4	64.6
- with VL data in post-training	98.6	91.4	74.0	62.5	73.1

monotonic Robot-only \rightarrow +Ego \rightarrow +H2R progression across both benchmarks confirms that raw ego data contributes through visual diversity, while the H2R pipeline unlocks additional gains via action and visual alignment.

Effect of VL data co-training during pre-training. We study the role of VL data in the pre-training stage by comparing our full model with a variant pre-trained without any VL data mixture (Table 18). On LIBERO and LIBERO-Plus, removing VL data causes relatively small drops of 0.9 and 1.2 points, respectively. In contrast, the performance degradation is much larger on the more challenging RoboTwin benchmarks: RoboTwin-Clean2Rand (easy) drops by 6.7 points, RoboTwin-Clean2Rand (hard) by 8.2 points, and RoboTwin-IF by 7.0 points. This trend indicates that VL co-training becomes increasingly important as task diversity, scene complexity, and distribution shift grow.

We attribute these gains to the complementary benefits of different VL data sources. General-domain VL data helps preserve the VLM’s broad perceptual and linguistic capabilities, reducing catastrophic forgetting during VLA training. Spatial grounding and reasoning data improves the model’s ability to localize objects, understand spatial relations, and reason over complex scenes, which is crucial for precise manipulation. Embodied-centric VL data further aligns the VLM’s representations with the semantic structure of embodied tasks. Together, these data sources yield VLM representations that are better suited for downstream alignment with the action expert.

Architecture design. We conduct ablation experiments on three network architecture variants (Figure 20):

1. The first variant replicates the VLM’s vision and language hidden states as input to the DiT, and employs a pure self-attention architecture within the DiT. After each self-attention layer in the DiT, the vision and language features are fused with the corresponding per-layer vision and language features from the VLM via a weighted residual combination.
2. The second variant feeds only the VLM’s *last-layer* hidden states into the DiT. The DiT likewise uses a pure self-attention architecture, but without any per-layer feature fusion with the VLM’s intermediate representations.
3. The third variant introduces a small set of learnable query tokens (Darcet et al., 2024) that act as condensed proxies for the VLM’s vision and language tokens. These query tokens are concatenated with the state and action tokens; all tokens jointly cross-attend to the VLM’s last-layer hidden states and self-attend among themselves. The query tokens’ outputs are discarded after processing.

All three variants are pretrained exclusively on robot data and then fine-tuned on the LIBERO dataset. During both pretraining and fine-tuning, we exclude the structured embodiment prompt (retaining only the original task instruction), ECoT data, and in-context policy adaptation, isolating the camera-frame delta pose representation as the sole alignment design retained in this comparison.

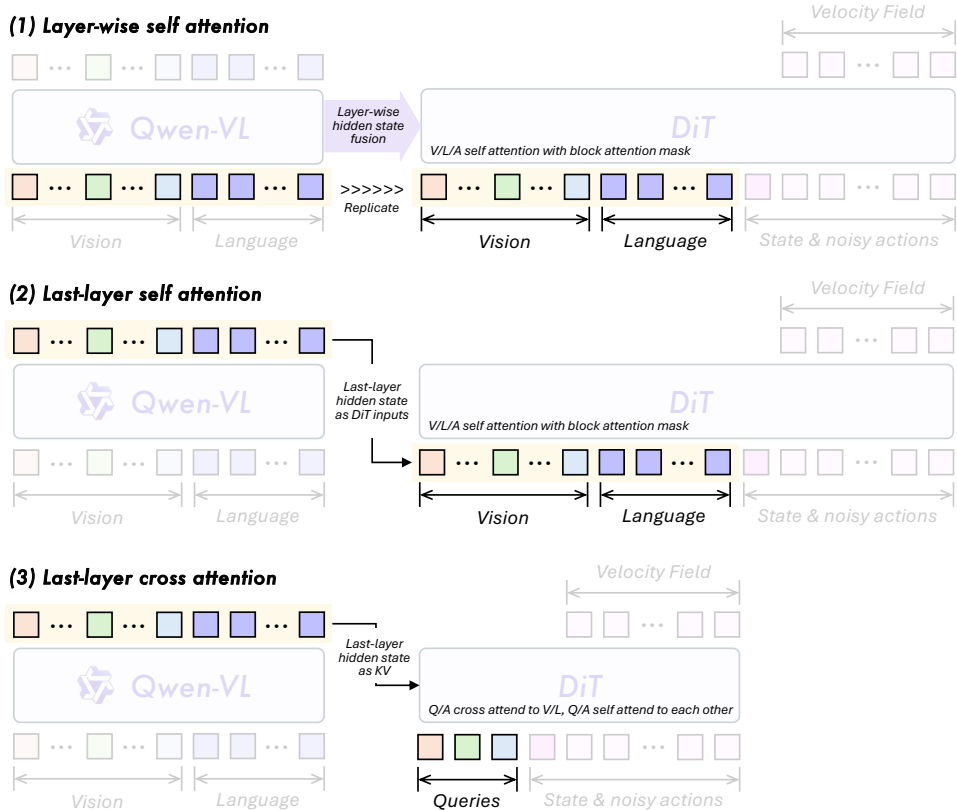


Figure 20: Ablation study of model architecture variants.

Table 19: Evaluation on LIBERO-Plus for different network architecture variants.

	Camera	Robot	Language	Light	Background	Noise	Layout	Total
layer-wise self attention	74.5	74.3	86.8	98.6	96.0	94.9	86.1	86.4
last-layer self attention	75.5	71.4	87.2	98.3	98.9	96.1	88.1	87.0
last-layer cross attention	76.9	73.4	86.9	98.9	97.7	97.3	87.3	87.5

The results are presented in Table 19. On LIBERO-Plus, the third variant (last-layer cross-attention) achieves the highest average success rate (87.5%) while incurring the lowest computational cost, as it avoids both per-layer feature fusion and storing the full set of VLM vision-language tokens in the DiT. The pure self-attention architectures, whether with per-layer or last-layer fusion, do not exhibit a clear advantage for vision-language-action feature interaction. We therefore adopt the third architecture in all subsequent experiments.

6.5 New Features after Alignment

6.5.1 Post-train with VL Data and VLA Data in Pre-training Recipe

Setting 1: Post-training with VL data co-training. In our main experiments, to match the standard fine-tuning setting, VL data is used only during pre-training and excluded from post-training. Here, we further investigate whether incorporating VL data into the post-training stage can improve generalization.

As shown in Table 18 (see “- with VL data in post-training”), adding VL data during post-training leads to improved performance on several benchmarks. In particular, it improves LIBERO-Plus from 90.1 to **91.4**, RoboTwin-Clean2Rand (easy) from 73.2 to **74.0**, and RoboTwin-IF from 71.6 to **73.1**, while the performance on RoboTwin-Clean2Rand (hard) remains nearly unchanged (62.6 vs. 62.5). These results suggest that post-training with VL data is particularly helpful for out-of-distribution generalization.

A closer analysis shows that the gain is most evident on *language-related generalization benchmarks*. For example, on LIBERO-Plus, the success rate under language perturbations increases from 86.9% for QWEN-ROBOTMANIP to 93.9% when VL data is included during post-training. Similarly, on RoboTwin-IF,




Figure 21: **Performance comparison under different proportions of training data in RoboTwin-IF.** We report the score at 10k, 30k, 50k, 70k, and 90k training steps. Compared with the baseline $\pi_{0.5}$, our method benefits from incorporating VLM data and further improves when co-training with both VLM data and VLA pretraining data.

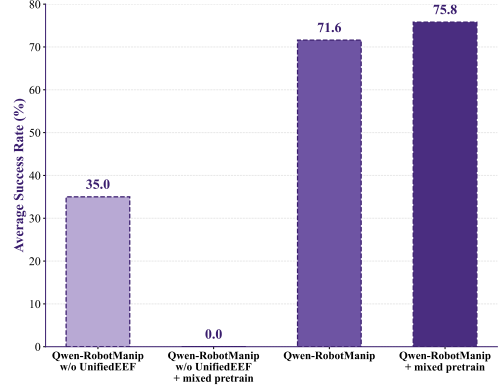


Figure 22: **Instruction-following success rates across configurations.** The chart highlights the critical necessity of the UnifiedEEF module, showing severe degradation without it (35.0%) and total failure when combined with mixed pretraining (0.0%). The optimal configuration utilizes both the base architecture and mixed pretraining (75.8%).

post-training with VL data improves the success rate from 76% to 81% on the “Pick-Diverse-Object” suite. In contrast, RoboTwin-Clean2Rand (hard) keeps the instructions fixed and instead introduces visual distractors, background changes, and lighting variations, so the benefit of VL data is less pronounced in this setting.

We attribute these improvements to better preservation of the base VLM’s foundational capabilities. Fine-tuning solely on action prediction can erode the pretrained VLM’s language understanding and visual grounding abilities due to catastrophic forgetting, thereby weakening its capacity to interpret novel instructions at test time. In contrast, mixing VL data into post-training helps maintain these capabilities, enabling the model to better parse unseen instructions, ground referring expressions to visual entities, and translate linguistic intent into appropriate actions.

Setting 2: Post-training with VL and auxiliary VLA data co-training. When transferring VLA models to a specific downstream domain via post-training such as a novel simulation environment, robot embodiment, or operational scene—it is standard practice to rely exclusively on clean, domain-specific datasets. However, this paradigm introduces a prominent bottleneck in VLA adaptation: severe domain overfitting. In most scenarios, rather than acquiring genuine, generalizable instruction-following capabilities, the model essentially memorizes the specific scene dynamics. Leveraging our proposed unified action space representation, we introduce a mixed post-training strategy to systematically enhance the model’s generalization capacity within the target domain. Specifically, during the curation of the post-training dataset, we augment the target-domain data by integrating a mixture of general Vision-Language (VL) data and pre-training VLA data. Crucially, this mixed-data approach effectively mitigates overfitting without compromising the model’s low-level physical execution proficiency in the target domain.

Specifically, our augmented dataset comprises two primary categories. The first consists of Vision-Language (VL) data, predominantly derived from standard Vision-Language Model (VLM) training corpora, encompassing tasks such as object detection, spatial pointing, and trajectory prediction. The second category entails VLA data, which aggregates trajectories from auxiliary simulators alongside large-scale, real-world demonstrations from morphologically similar embodiments. To facilitate effective cross-embodiment transfer, we align and optimize this auxiliary VLA data within the highly transferable end-effector (EEF) action space.

We evaluate our approach on RoboTwin-IF, the Out-of-Distribution benchmark within our target domain, RoboTwin. As shown in Figure 21, when fine-tuning exclusively on the domain-specific RoboTwin-Clean dataset, our model initially achieves performance that significantly surpasses that of $\pi_{0.5}$. However, as the number of training steps increases, severe overfitting emerges, leading to a progressive performance decay on the RoboTwin-IF evaluation. By integrating VL data (accounting for 10% of the training mixture), both the instability of the peak performance and the subsequent degradation are substantially mitigated. Furthermore, when we expand the mixture to include the auxiliary VLA data (comprising 75% of the total vla data in the training mixture), the overfitting phenomenon is entirely eradicated. Notably, as training progresses, the model’s performance on RoboTwin-IF exhibits continuous improvement rather than decay, ultimately achieving comparable peak performance while maintaining robust OOD generalization.

Table 20: **Systematic advantages of camera-frame delta EEF** across three evaluation dimensions, progressing from in-distribution to zero-shot settings.

Dimension	Evaluation	Best baseline	QWEN-ROBOTMANIP	Gain
EEF control quality	RT-C2R Easy / Hard (eef mode)	49.0 / 33.0 (w/o UnifiedEEF)	72.5 / 56.6	+23.5 / +23.6
Skill composition	CobotMagic→ARX, 4 novel tasks	12.5% (w/o UnifiedEEF)	55.0%	4.4×
Zero-shot transfer	AgileX→ARX, UR5, Franka (avg)	14.5% (joint)	23.9% (eef)	1.65×

Besides, as shown in Figure 22, the most critical insight is that the *UnifiedEEF* module serves as the **exclusive prerequisite** for mixed post-training. While simply removing the module degrades baseline performance (from 71.6% to 35.0%), attempting to introduce mixed post-training data without it triggers a complete representation collapse (0.0% success rate). This catastrophic failure fundamentally proves that the model is entirely incapable of assimilating mixed VLM/VLA data on its own. It is **only** through the structural alignment provided by *UnifiedEEF* that the architecture can unlock the benefits of mixed post-training, successfully pushing the performance ceiling to 75.8%.

6.5.2 EEF control & Cross-Embodiment Transfer

Camera-frame delta EEF expresses actions as pose deltas in the visual observation coordinate system, making the action representation inherently shared across embodiments regardless of their underlying kinematics. We evaluate this design along three complementary dimensions—in-distribution control quality, compositional skill transfer, and zero-shot cross-embodiment generalization—and consolidate the key results in Table 20.

In-distribution EEF control. As the first row of Table 20 shows, QWEN-ROBOTMANIP with camera-frame EEF achieves 72.5% / 56.6% on RoboTwin-C2R Easy / Hard under EEF-mode execution, compared with 49.0% / 33.0% for the best alternative action-space design (QWEN-ROBOTMANIP w/o UnifiedEEF). More notably, QWEN-ROBOTMANIP is the *only* variant where EEF-mode execution surpasses its own joint-mode performance (72.5% vs. 68.1% Easy, 56.6% vs. 50.2% Hard, see Figure 19); all other action-space designs degrade when switching from joint to EEF control. This reversal indicates that camera-frame alignment produces a genuinely strong EEF-space policy rather than merely an alternative control interface. Experiments for data scaling (Figure 18) corroborates this finding: under the unified EEF representation, cross-embodiment data follows a clean log-linear scaling law, whereas the ablated method yields an erratic curve with substantially higher prediction error.

Compositional skill transfer. Camera-frame EEF decouples manipulation skills from robot-specific kinematics, enabling skill-level composition across embodiment–task combinations. In a joint-training experiment with 6K CobotMagic and 130 ARX demonstrations, the policy is evaluated on four novel ARX tasks for which zero target-task demonstrations exist. As the second row of Table 20 shows, QWEN-ROBOTMANIP achieves 55.0%—4.4× the best ablated variant (QWEN-ROBOTMANIP w/o UnifiedEEF, 12.5%)—while QWEN-ROBOTMANIP w/o UnifiedSpace manages only 7.5%. Because camera-frame deltas map the same manipulation primitive to a consistent numerical pattern regardless of the executing robot, the model acquires embodiment-agnostic skill representations that compose freely with new embodiment–task pairings (per-task breakdown in Table 13).

Zero-shot cross-embodiment transfer. The most demanding test deploys a policy trained exclusively on AgileX to three unseen embodiments—ARX, UR5, and Franka—without any target-embodiment data. As the third row of Table 20 shows, EEF control achieves an average success rate of 23.9%, far exceeding the 14.5% of joint control. The improvement is especially pronounced on UR5, where EEF mode reaches 22.8% versus 4.1% for joint mode (5.6×). Joint-space actions are inherently robot-specific and produce near-random behavior on unseen morphologies, whereas camera-frame deltas abstract away kinematic differences and enable meaningful transfer in shared Cartesian space (per-embodiment results in Table 9).

Taken together, the three rows of Table 20 reveal a coherent progression: camera-frame delta EEF first strengthens in-distribution EEF control, then enables cross-embodiment skill composition, and finally supports zero-shot deployment on unseen robots—confirming that it provides an embodiment-agnostic action interface whose benefits compound across increasingly challenging transfer settings.

7 Conclusion

This report set out to investigate whether the scaling recipe behind large language and multimodal models, aligning heterogeneous data under a unified formulation and training at scale, can be applied to robotic manipulation to achieve genuine generalization. We present QWEN-ROBOTMANIP as an affirmative answer, built on the principle that alignment and scale are not independent engineering challenges but tightly coupled prerequisites: without a unified cross-embodiment formulation, scaling data produces conflicts rather than synergy; without sufficient data diversity, even a well-aligned model cannot generalize beyond its training distribution.

Several findings from this work carry implications beyond the specific system we describe. First, the unified alignment framework, spanning a canonical state-action representation, camera-frame delta pose parameterization, and in-context policy adaptation, proves critical not merely for accommodating diverse embodiments but for enabling data scaling itself. Our ablations show that naïve representations fail to exhibit scaling behavior; alignment is what converts additional data volume into improved capability. Second, the fact that QWEN-ROBOTMANIP constructs a $\sim 38,100$ -hour corpus and achieves emergent generalization capabilities using *only* open-source robotic manipulation datasets and egocentric human videos, without any proprietary data collection, suggests that the data barrier for manipulation foundation models may be lower than commonly assumed, provided the right synthesis and curation infrastructure is in place. Third, our systematic comparison between in-domain and out-of-distribution evaluation reveals that standard benchmarks consistently fail to distinguish models whose pretraining contributes genuine generalizable structure from those that succeed through in-distribution pattern matching. We believe the OOD evaluation settings introduced in this work provide a more faithful measure of robotic foundation model capability, and we hope they serve as useful diagnostics for the community.

Limitations remain. The human-to-robot synthesis pipeline, while scalable, introduces distributional gaps from retargeting approximations and inpainting artifacts that bound the effective quality of synthesized data. Our OOD evaluations, though substantially more challenging than standard benchmarks, are still predominantly simulation-based. Broader real-world evaluation across deployment conditions is needed. The fixed action chunk length and inference latency of the current system also constrain applicability to tasks requiring reactive sub-second control.

Looking forward, the alignment-then-scale paradigm demonstrated here naturally extends in several directions: incorporating more robot morphologies and task domains into the pretraining corpus, improving synthesis fidelity through more accurate hand-robot retargeting and physically grounded rendering, and incorporating agentic systems toward longer-horizon reasoning and manipulation. Moreover, we hope this work contributes to a shift in how the community evaluates VLA models, from in-domain benchmark rank toward the out-of-distribution generalization that ultimately determines whether a model can serve as a genuine foundation for real-world deployment.

8 Authors

Core Contributors: Haoqi Yuan*, Zhixuan Liang*, Anzhe Chen*, Ye Wang*, Haoyang Li*, Pei Lin*, Yiyang Huang*, Zixing Lei*, Tong Zhang*, Jiazhao Zhang, Jie Zhang, Jingyang Fan, Gengze Zhou, Qihang Peng, Chenxu Lv, Xiaoyue Chen, An Yang, Fei Huang, Junyang Lin, Dayiheng Liu, Jingren Zhou, Chenfei Wu[†], Xiong-Hui Chen[†]

*Equal Contribution. † Corresponding authors.

Contributors: Jinhui Ye, Sicheng Xie, Hale Yin, Xudong Guo, Shuai Bai, Lulu Hu, Minying Zhang, Shurui Li, Wenhui Xiao, Yue Wang, Kun Yan, Xiao Xu, Jiahao Li, Xuancheng Ren

Acknowledgments: We acknowledge the National Pilot Base for Embodied Intelligence Applications for providing the real-robot experimental environment and equipment. We also thank AgileX Robotics for their hardware support.

References

- AgiBot-World-Contributors. AgiBot World Colosseo: A large-scale manipulation platform for scalable and intelligent embodied systems. *arXiv preprint arXiv:2503.06669*, 2025.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. Qwen3-v1 technical report. *arXiv preprint arXiv:2511.21631*, 2025.
- Johan Bjorck, Fernando Castaneda, Linxi Fan, Dieter Fox, et al. GR00T N1: An open foundation model for generalist humanoid robots. *arXiv preprint arXiv:2503.14734*, 2025.
- 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.
- Kevin Black, Noah Brown, James Darphinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Manuel Y Galliker, et al. $\pi_{0.5}$: a vision-language-action model with open-world generalization. In *9th Annual Conference on Robot Learning*, 2025.
- Kevin Black, Manuel Galliker, and Sergey Levine. Real-time execution of action chunking flow policies. *Advances in Neural Information Processing Systems*, 38:33383–33407, 2026.
- 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. In *NeurIPS*, 2020.
- Junhao Cai, Zetao Cai, Jiafei Cao, Yilun Chen, Zeyu He, Lei Jiang, Hang Li, Hengjie Li, Yang Li, Yufei Liu, et al. Internvla-a1: Unifying understanding, generation and action for robotic manipulation. *arXiv preprint arXiv:2601.02456*, 2026.
- Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, Jie Lei, Tengyu Ma, Baishan Guo, Arpit Kalla, Markus Marks, Joseph Greer, Meng Wang, Peize Sun, Roman Rädle, Triantafyllos Afouras, Effrosyni Mavroudi, Katherine Xu, Tsung-Han Wu, Yu Zhou, Liliane Momeni, Rishi Hazra, Shuangrui Ding, Sagar Vaze, Francois Porcher, Feng Li, Siyuan Li, Aishwarya Kamath, Ho Kei Cheng, Piotr Dollár, Nikhila Ravi, Kate Saenko, Pengchuan Zhang, and Christoph Feichtenhofer. Sam 3: Segment anything with concepts, 2025. URL <https://arxiv.org/abs/2511.16719>.
- Justin Carpentier, Guilhem Saurel, Gabriele Buondonno, Joseph Mirabel, Florent Lamiroux, Olivier Stasse, and Nicolas Mansard. The Pinocchio C++ library: A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives. In *IEEE/SICE International Symposium on System Integration (SII)*, pp. 614–619. IEEE, 2019.
- Anzhe Chen, Yifei Yang, Zhenjie Zhu, Kechun Xu, Zhongxiang Zhou, Rong Xiong, and Yue Wang. Toward embodiment equivariant vision-language-action policy. *arXiv preprint arXiv:2509.14630*, 2025a.
- Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Zixuan Li, Qiwei Liang, Xianliang Lin, Yiheng Ge, Zhenyu Gu, et al. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. *arXiv preprint arXiv:2506.18088*, 2025b.
- William Chen, Suneel Belkhale, Suvir Mirchandani, Oier Mees, Danny Driess, Karl Pertsch, and Sergey Levine. Training strategies for efficient embodied reasoning. *ArXiv*, abs/2505.08243, 2025c.
- Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Robotics: Science and Systems (RSS)*, 2023.
- StarVLA Community. Starvla: A lego-like codebase for vision-language-action model developing. *arXiv preprint arXiv:2604.05014*, 2026.
- Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The EPIC-KITCHENS dataset. In *European Conference on Computer Vision (ECCV)*, 2018.

-
- Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun (eds.), *International Conference on Learning Representations*, volume 2024, pp. 2632–2652, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/0b408293619f725fd30162af057e531a-Paper-Conference.pdf.
- Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, Jiasen Lu, Taira Anderson, Erin Bransom, Kiana Ehsani, Huong Ngo, Yen-Sung Chen, Ajay Patel, Mark Yatskar, Chris Callison-Burch, Andrew Head, Rose Hendrix, Favyen Bastani, Eli VanderBilt, Nathan Lambert, Yvonne Chou, Arnavi Chheda, Jenna Sparks, Sam Skjonsberg, Michael Schmitz, Aaron Sarnat, Byron Bischoff, Pete Walsh, Chris Newell, Piper Wolters, Tanmay Gupta, Kuo-Hao Zeng, Jon Borchardt, Dirk Groeneveld, Crystal Nam, Sophie Lebrecht, Caitlin Wittlif, Carissa Schoenick, Oscar Michel, Ranjay Krishna, Luca Weihs, Noah A. Smith, Hannaneh Hajishirzi, Ross B. Girshick, Ali Farhadi, and Aniruddha Kembhavi. Molmo and pixmo: Open weights and open data for state-of-the-art vision-language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, 2025.
- Danny Driess, Jost Tobias Springenberg, Brian Ichter, Lili Yu, Adrian Li-Bell, Karl Pertsch, Allen Z Ren, Homer Walke, Quan Vuong, Lucy Xiaoyang Shi, et al. Knowledge insulating vision-language-action models: Train fast, run fast, generalize better. *arXiv preprint arXiv:2505.23705*, 2025.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv:2407.21783*, 2024.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *ICML*, 2024.
- 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. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- Haoquan Fang, Jiafei Duan, Donovan Clay, Sam Wang, Shuo Liu, Weikai Huang, Xiang Fan, Wei-Chuan Tsai, Shirui Chen, Yi Ru Wang, et al. Molmoact2: Action reasoning models for real-world deployment. *arXiv preprint arXiv:2605.02881*, 2026.
- Senyu Fei, Siyin Wang, Junhao Shi, Zihao Dai, Jikun Cai, Pengfang Qian, Li Ji, Xinzhe He, Shiduo Zhang, Zhaoye Fei, et al. Libero-plus: In-depth robustness analysis of vision-language-action models. *arXiv preprint arXiv:2510.13626*, 2025.
- Youhe Feng, Hansen Shi, Haoyang Li, Xinlei Guo, Yang Wang, Chengyang Zhang, Jinkai Zhang, Xiaohan Zhang, Jie Tang, and Jing Zhang. Procvlm: Learning procedure-grounded progress rewards for robotic manipulation. *CoRR*, abs/2605.08774, 2026.
- Galaxea AI. Galaxea open-world dataset and G0 dual-system VLA model. *arXiv preprint arXiv:2509.00576*, 2025.
- Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The "something something" video database for learning and evaluating visual common sense. In *Proceedings of the IEEE international conference on computer vision*, pp. 5842–5850, 2017.
- Kristen Grauman, Andrew Westbury, Eugene Byrne, et al. Ego4D: Around the world in 3,000 hours of egocentric video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- 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*, pp. 19383–19400, 2024.
- Byeongho Heo, Song Park, Dongyoon Han, and Sangdoo Yun. Rotary position embedding for vision transformer. In *European Conference on Computer Vision*, pp. 289–305. Springer, 2024.
- 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.

-
- Yiyang Huang, Yuhui Hao, Bo Yu, Feng Yan, Yuxin Yang, Feng Min, Yinhe Han, Lin Ma, Shaoshan Liu, Qiang Liu, et al. Dadu-corki: Algorithm-architecture co-design for embodied ai-powered robotic manipulation. In *Proceedings of the 52nd Annual International Symposium on Computer Architecture*, pp. 327–343, 2025.
- Simar Kareer, Dhruv Patel, Ryan Punamiya, Pranay 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)*, pp. 13226–13233. IEEE, 2025.
- Alexander Khazatsky, Karl Pertsch, Suraj Nair, et al. DROID: A large-scale in-the-wild robot manipulation dataset. In *Robotics: Science and Systems (RSS)*, 2024.
- Dongyoung Kim, Huiwon Jang, Myungkyu Koo, Suhyeok Jang, Taeyoung Kim, Beomjun Kim, Byungjun Yoon, Changsung Jang, Daewon Choi, Dongsu Han, et al. Rldx-1 technical report. *arXiv preprint arXiv:2605.03269*, 2026a.
- 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.
- Moo Jin Kim, Chelsea Finn, and Percy Liang. Fine-tuning vision-language-action models: Optimizing speed and success. *arXiv preprint arXiv:2502.19645*, 2025.
- Moo Jin Kim, Yihuai Gao, Tsung-Yi Lin, Yen-Chen Lin, Yunhao Ge, Grace Lam, Percy Liang, Shuran Song, Ming-Yu Liu, Chelsea Finn, et al. Cosmos policy: Fine-tuning video models for visuomotor control and planning. *arXiv preprint arXiv:2601.16163*, 2026b.
- Xin Kong, Shikun Liu, Xiaoyang Lyu, Marwan Taher, Xiaojuan Qi, and Andrew J. Davison. EscherNet: A generative model for scalable view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9503–9513, 2024.
- Shanghai AI Laboratory. Ebench: Elemental mobile manipulation benchmark, 2026. URL <https://internrobotics.github.io/EBench-doc/>.
- Zixing Lei, Changxing Liu, Yichen Xiong, Minhao Xiong, Yuanzhuo Ding, Zhipeng Zhang, Weixin Li, and Siheng Chen. Towards long-horizon embodied agents with tool-aligned vision-language-action models. *arXiv preprint arXiv:2605.13119*, 2026.
- Marion Lepert, Jiaying Fang, and Jeannette Bohg. Masquerade: Learning from in-the-wild human videos using data-editing. *arXiv preprint arXiv:2508.09976*, 2025a.
- Marion Lepert, Jiaying Fang, and Jeannette Bohg. Phantom: Training robots without robots using only human videos. *arXiv preprint arXiv:2503.00779*, 2025b.
- Hao Li, Ziqin Wang, Zi-Han Ding, Shuai Yang, Yilun Chen, Yang Tian, Xiaolin Hu, Tai Wang, Dahua Lin, Feng Zhao, Si Liu, and Jiangmiao Pang. Robointer: A holistic intermediate representation suite towards robotic manipulation. *CoRR*, abs/2602.09973, 2026.
- 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*, 2025a.
- Ruilong Li, Brent Yi, Junchen Liu, Hang Gao, Yi Ma, and Angjoo Kanazawa. Cameras as relative positional encoding. *Advances in Neural Information Processing Systems*, 2025b.
- Zhixuan Liang, Yao Mu, Mingyu Ding, Fei Ni, Masayoshi Tomizuka, and Ping Luo. Adaptdiffuser: diffusion models as adaptive self-evolving planners. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 20725–20745, 2023.
- Zhixuan Liang, Yizhuo Li, Tianshuo Yang, Chengyue Wu, Sitong Mao, Liua Pei, Tian Nian, Shunbo Zhou, Xiaokang Yang, Jiangmiao Pang, et al. Discrete diffusion vla: Bringing discrete diffusion to action decoding in vision-language-action policies. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026.
- Haotong Lin, Sili Chen, Junhao Liew, Donny Y Chen, Zhenyu Li, Guang Shi, Jiashi Feng, and Bingyi Kang. Depth anything 3: Recovering the visual space from any views. *arXiv preprint arXiv:2511.10647*, 2025.

-
- Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations (ICLR)*, 2023.
- Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. *arXiv preprint arXiv:2306.03310*, 2023.
- Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang Su, and Jun Zhu. RDT-1B: a diffusion foundation model for bimanual manipulation. In *International Conference on Learning Representations (ICLR)*, 2025.
- 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.
- Hao Luo, Ye Wang, Wanpeng Zhang, Haoqi Yuan, Yicheng Feng, Haiweng Xu, Sipeng Zheng, and Zongqing Lu. Joint-aligned latent action: Towards scalable vla pretraining in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 35047–35058, 2026a.
- Hao Luo, Ye Wang, Wanpeng Zhang, Sipeng Zheng, Ziheng Xi, Chaoyi Xu, Haiweng Xu, Haoqi Yuan, Chi Zhang, Yiqing Wang, et al. Being-h0. 5: Scaling human-centric robot learning for cross-embodiment generalization. *arXiv preprint arXiv:2601.12993*, 2026b.
- Hao Luo, Wanpeng Zhang, Yicheng Feng, Sipeng Zheng, Haiweng Xu, Chaoyi Xu, Ziheng Xi, Yuhui Fu, and Zongqing Lu. Being-h0. 7: A latent world-action model from egocentric videos. *arXiv preprint arXiv:2605.00078*, 2026c.
- Jiangran Lyu, Kai Liu, Xuheng Zhang, Haoran Liao, Yusen Feng, Wenxuan Zhu, Tingrui Shen, Jiayi Chen, Jiazhao Zhang, Yifei Dong, et al. Lda-1b: Scaling latent dynamics action model via universal embodied data ingestion. *arXiv preprint arXiv:2602.12215*, 2026.
- Takeru Miyato, Bernhard Jaeger, Max Welling, and Andreas Geiger. GTA: A geometry-aware attention mechanism for multi-view transformers. In *International Conference on Learning Representations (ICLR)*, 2024.
- Yao Mu, Tianxing Chen, Zan Ding, et al. Robotwin: Dual-arm robot benchmark with generative digital twins. *arXiv preprint arXiv:2501.00062*, 2025.
- 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. In *Robotics: Science and Systems (RSS)*, 2024.
- Soroush Nasiriany, Sepehr Nasiriany, Abhiram Maddukuri, and Yuke Zhu. Robocasa365: A large-scale simulation framework for training and benchmarking generalist robots. In *International Conference on Learning Representations (ICLR)*, 2026.
- OpenAI. Gpt-4 technical report. *arXiv:2303.08774*, 2023.
- Abhishek Padalkar, Acorn Pooley, Ajinkya Jain, Alex Bewley, Alex Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anikait Singh, Anthony Brohan, et al. Open X-Embodiment: Robotic learning datasets and RT-X models. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4195–4205, 2023.
- Ryan Punamiya, Simar Kareer, Zeyi Liu, Josh Citron, Ri-Zhao Qiu, Xiongyi Cai, Alexey Gavryushin, Jiaqi Chen, Davide Liconti, Lawrence Y Zhu, et al. Egoverse: An egocentric human dataset for robot learning from around the world. *arXiv preprint arXiv:2604.07607*, 2026.
- Ri-Zhao Qiu, Shiqi Yang, Xuxin Cheng, Chaitanya Chawla, Jialong Li, Tairan He, Ge Yan, David J Yoon, Ryan Hoque, Lars Paulsen, et al. Humanoid policy~ human policy. *arXiv preprint arXiv:2503.13441*, 2025.
- Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *arXiv preprint arXiv:2201.02610*, 2022.
- Abraham Savitzky and Marcel JE Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry*, 36(8):1627–1639, 1964.

-
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: A family of highly capable multimodal models. *arXiv:2312.11805*, 2023.
- Qwen Team. Qwen3.5: Accelerating productivity with native multimodal agents, February 2026. URL <https://qwen.ai/blog?id=qwen3.5>.
- Yang Tian, Yuyin Yang, Yiman Xie, Zetao Cai, Xu Shi, Ning Gao, Hangxu Liu, Xuekun Jiang, Zherui Qiu, Feng Yuan, Yaping Li, Ping Wang, Junhao Cai, Jia Zeng, Hao Dong, and Jiangmiao Pang. InternData-A1: Pioneering high-fidelity synthetic data for pre-training generalist policy. *arXiv preprint arXiv:2511.16651*, 2025.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pp. 5026–5033. IEEE, 2012.
- Ye Wang, Sipeng Zheng, Hao Luo, Wanpeng Zhang, Haoqi Yuan, Chaoyi Xu, Haiweng Xu, Yicheng Feng, Mingyang Yu, Zhiyu Kang, et al. Rethinking visual-language-action model scaling: Alignment, mixture, and regularization. *arXiv preprint arXiv:2602.09722*, 2026.
- Kun Wu, Chengkai Hou, Jiaming Liu, Zhengping Che, et al. RoboMIND: Benchmark on multi-embodiment intelligence normative data for robot manipulation. In *Robotics: Science and Systems (RSS)*, 2025a.
- Kun Wu et al. RoboMIND 2.0: A multimodal, bimanual mobile manipulation dataset for generalizable embodied intelligence. *arXiv preprint arXiv:2512.24653*, 2025b.
- Shihan Wu et al. RoboCOIN: An open-sourced bimanual robotic data collection for integrated manipulation. *arXiv preprint arXiv:2511.17441*, 2025c.
- Feng Yan, Fanfan Liu, Yiyang Huang, Zechao Guan, Liming Zheng, Yufeng Zhong, Chengjian Feng, and Lin Ma. Robotron-mani: All-in-one multimodal large model for robotic manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 13707–13718, October 2025.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- Yandan Yang, Shuang Zeng, Tong Lin, Xinyuan Chang, Dekang Qi, Junjin Xiao, Haoyun Liu, Ronghan Chen, Yuzhi Chen, Dongjie Huo, et al. Abot-m0: Vla foundation model for robotic manipulation with action manifold learning. *arXiv preprint arXiv:2602.11236*, 2026.
- Qiyang Yu, Quan Sun, Xiaosong Zhang, Yufeng Cui, Fan Zhang, Yue Cao, Xinlong Wang, and Jingjing Liu. Capsfusion: Rethinking image-text data at scale. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, 2024.
- 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 in robotics. In *Conference on Robot Learning, 6-9 November 2024, Munich, Germany, 2024*.
- Kevin Zakka. Mink: Python inverse kinematics based on MuJoCo, February 2026. URL <https://github.com/kevinzakka/mink>.
- 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.
- Tao Zhang, Song Xia, Ye Wang, and Qin Jin. Easymimic: A low-cost framework for robot imitation learning from human videos. *arXiv preprint arXiv:2602.11464*, 2026a.
- Tianyi Zhang, Haonan Duan, Haoran Hao, Yu Qiao, Jifeng Dai, and Zhi Hou. Grounding actions in camera space: Observation-centric vision-language-action policy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 18782–18790, 2026b.
- Zhanguang Zhang, Zhiyuan Li, Behnam Rahmati, Rui Heng Yang, Yintao Ma, Amir Rasouli, Sajjad Pakdamansavoji, Yangzheng Wu, Lingfeng Zhang, Tongtong Cao, et al. Do world action models generalize better than vlas? a robustness study. *arXiv preprint arXiv:2603.22078*, 2026c.
- Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. In *Robotics: Science and Systems (RSS)*, 2023.

-
- Jinliang Zheng, Jianxiong Li, Zhihao Wang, Dongxiu Liu, Xirui Kang, Yuchun Feng, Yinan Zheng, Jiayin Zou, Yilun Chen, Jia Zeng, et al. X-vla: Soft-prompted transformer as scalable cross-embodiment vision-language-action model. *arXiv preprint arXiv:2510.10274*, 2025.
- Ruijie Zheng, Dantong Niu, Yuqi Xie, Jing Wang, Mengda Xu, Yunfan Jiang, Fernando Castañeda, Fengyuan Hu, You Liang Tan, Letian Fu, et al. Egoscale: Scaling dexterous manipulation with diverse egocentric human data. *arXiv preprint arXiv:2602.16710*, 2026.
- Enshen Zhou, Jingkun An, Cheng Chi, Yi Han, Shanyu Rong, Chi Zhang, Pengwei Wang, Zhongyuan Wang, Tie-Jun Huang, Lu Sheng, and Shanghang Zhang. Roborefer: Towards spatial referring with reasoning in vision-language models for robotics. *CoRR*, abs/2506.04308, 2025.
- Shangchen Zhou, Chongyi Li, Kelvin CK Chan, and Chen Change Loy. Propainter: Improving propagation and transformer for video inpainting. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 10477–10486, 2023.
- Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.