

Igniting VLMs toward the Embodied Space

X SQUARE ROBOT TEAM

While foundation models show remarkable progress in language and vision, existing vision-language models (VLMs) still have limited spatial and embodiment understanding. Transferring VLMs to embodied domains reveals fundamental mismatches between modalities, pretraining distributions, and training objectives, leaving action comprehension and generation as a central bottleneck on the path to AGI. We introduce **WALL-OSS**, an end-to-end embodied foundation model that leverages large-scale multimodal pretraining to achieve (1) embodiment-aware vision-language understanding, (2) strong language-action association, and (3) robust manipulation capability. Our approach employs a tightly coupled architecture and multi-strategies training curriculum that enables Unified Cross-Level CoT—seamlessly unifying instruction reasoning, subgoal decomposition, and fine-grained action synthesis within a single differentiable framework. Our results show that WALL-OSS attains high success on complex long-horizon manipulations, demonstrates strong instruction-following capabilities, complex understanding and reasoning, and outperforms strong baselines, thereby providing a reliable and scalable path from VLMs to embodied foundation models.

Date: September 8, 2025

Code: <https://github.com/X-Square-Robot/wall-x>

Project Page: <https://x2robot.com/en/research/68bc2cde8497d7f238dde690>



1 Introduction

Foundation models in language and vision have advanced rapidly. The trajectory of LLM progress is now mirrored in VLMs: powerful omni models such as Gemini 2.5 (28) and GPT-5 (56) can jointly process text and visual streams while preserving deliberate reasoning and low error across multimodal domains. However, these systems remain largely disembodied: they neither refine themselves via feedback from interaction with the physical world nor generate executable actions. Consequently, within the embodied space, action understanding and generation remains a core bottleneck on the path toward AGI.

Modalities that have most successfully approached general-purpose generalization typically rely on pretraining over large-scale, diverse, and readily accessible data. Benefiting from the development of human media, the Internet hosts massive, richly distributed text, image, and video data; thus, text and 2D vision have relatively early acquired both single-modality general competence and cross-modal association. In embodied settings, however, aligned Vision-Language-Action (VLA) data are scarce and heterogeneous, making from-scratch training struggle to achieve generalization in the action modality as well as cross-modal association. This status quo has motivated transferring strong VLM backbones to the action space, e.g., OpenVLA (39) and π_0 (14), which model continuous actions while learning cross-modal correspondences to migrate vision and language priors (VL priors) into the action space.

A naive fine-tuning strategy—attaching an action head to a pretrained VLM and supervising it with robot trajectories—often leads to severe weight drift, thereby eroding the original VL prior. “Knowledge insulation” strategies preserve pretrained capabilities by minimizing perturbations to VLM parameters; however, loosely coupled architectural designs restrict the binding between semantics and control, and because VLMs themselves are deficient in the embodied space, they perform poorly when executing actions outside the VL prior. Overall, learning actions with current VLMs is challenging due to three fundamental gaps:

Gap in modalities and data scale. Thanks to the long-term development of computer vision, visual encoders extract features that filter high-frequency noise while retaining high-level semantics; the compressed visual tokens, together with the already highly compressed text modality, enable alignment via CLIP and related multimodal models and VLMs after training on large-scale web text-image/video pairs (59). In contrast,

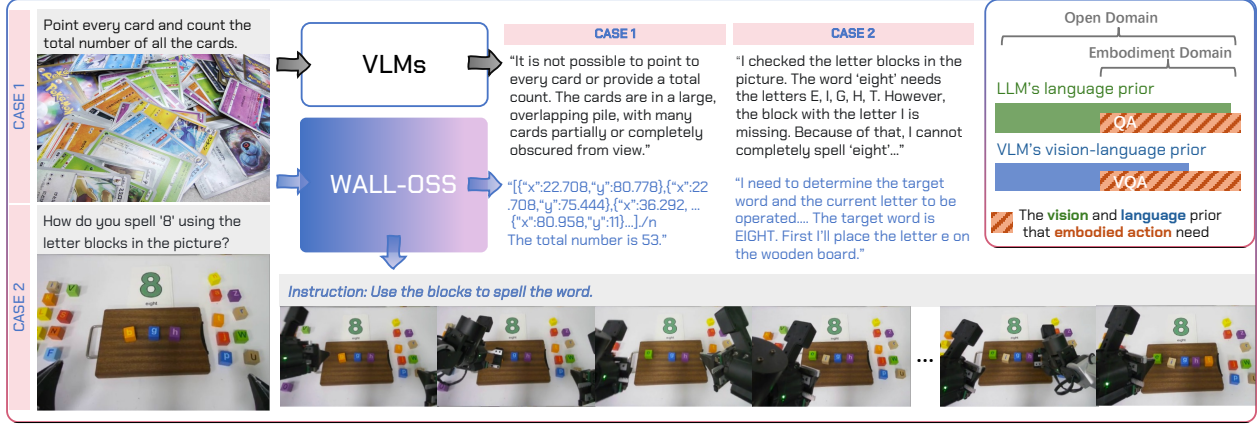


Figure 1 Current VLMs lack a sufficient understanding of space and action within embodied AI. This deficiency stems from a mismatch between the capabilities of existing pre-trained VLMs and the specific knowledge required for embodied tasks. WALL-OSS unleashes the embodied potential of VLMs, leading to enhanced embodied understanding and the ability to generate complex actions.

actions in the embodied space are continuous over 3D space and time; compared to the visual modality, their feature extraction lacks long-term representational study, and their high-level semantic alignment with text lacks the drive of massive data. Moreover, in embodied scenarios, multiple subtasks and concrete scene-level action descriptions are often abstracted into a single high-level instruction, further increasing cross-modal association difficulty. Pipeline-based agents (first decomposing instructions with an LLM, then invoking a skill library (6, 44)) are limited by non-differentiable interfaces and error accumulation, degrading success rates on long-horizon tasks. In addition, some methods such as 3D-VLA (79) and PointVLA (41) attempt to use the 3D visual modality as a bridge between 2D vision and action, but 3D data are likewise difficult to obtain, and current 3D-vision foundation models such as VGGT (69) and π^3 (72) still fall short of the accuracy required for fine-grained action prediction.

Gap in pretraining distributions. As shown in Figure 1, after large-scale pretraining, LLMs largely cover the needs of embodied semantics within the text modality (capable of QA). For the visual modality, however, embodied data typically involve first-person viewpoints, fisheye imaging, and self-occlusions, which differ substantially from Internet imagery. Even after large-scale pretraining, VLMs struggle to cover the needs of embodied scenarios (e.g., embodied VQA), and prior analyses have reported shortcomings in spatial reasoning, embodied-scene understanding, and progress tracking (36, 70).

Gap in training objectives. LLMs/VLMs use next-token likelihood on discrete sequences as their training objective, whereas action trajectories are continuous, high-frequency signals that are more naturally modeled with conditional generative objectives such as diffusion (32, 22) or flow matching (46), which learn score/velocity fields that transport a base distribution to the action distribution. However, grafting these objectives directly onto VLMs magnifies the tokenization gulf and independence assumptions, causing catastrophic degradation that weakens language–action alignment and generalization. This mirrors the challenges faced when unifying generation and understanding in current image/video settings. Some approaches (e.g., π_0) attempt to discretize actions at a high level under an autoregressive paradigm to align action tokens with text tokens, and then generate continuous action signals by letting action noise from diffusion/flow-matching interact with intermediate VLM representations via self-attention (14). Yet such loosely coupled designs still fail to learn tight text–action bindings, leading to inadequate instruction following.

To address the above issues, we propose **WALL-OSS**, a tightly coupled VLA foundation model and training recipe. We first introduce a tightly coupled Mixture-of-Experts (MoE) architecture and training strategy that activates different experts and weights across different stages with different action modeling approaches (discrete or continuous), thereby bridging the training objective gap. We combine multimodal tasks with VQA tasks and corresponding datasets tailored to embodiment, spatial localization, and progress modeling to remedy VLM shortcomings in embodied scenarios, and implant discretized action priors into the VLM output space for coarse action supervision, thereby bridging the VLM pretraining-distribution gap. Building

upon these designs, we model unified CoT forward mapping that enables the model to decompose tasks from high-level semantics to fine-grained actions, compelling the model to understand the relationships between each level in this chain, thereby bridging the modality gap.

Our key contributions establish a principled framework for scaling VLMs to embodied intelligence. Methodologically, we demonstrate how architectural coupling and curriculum design can systematically address the fundamental misalignments between disembodied pretraining and embodied domains. Empirically, WALL-OSS achieves state-of-the-art performance across diverse manipulation benchmarks, with particularly strong results on instruction following, compositional reasoning, and long-horizon control. More broadly, our work validates a scalable paradigm for bridging the semantic-sensorimotor divide, offering both theoretical insights and practical advances toward general-purpose embodied AI systems. Given the current scarcity of embodied general foundation models, we also open-source comprehensive training code and model checkpoints to facilitate reproducible research and community development in this emerging domain.

2 Related Work

Foundation models in language, vision and manipulation. Web-scale pretraining has pushed LLMs and VLMs toward stronger multimodal reasoning and grounding (e.g., Gemini 2.5 (37, 4), Claude 3.7 (3), DeepSeek-R1, Llama 3 and 4 (2, 5), Qwen 2.5 (73); GPT-4o (55), Qwen2.5-VL (10), and the InternVL series (21, 20, 19, 80, 71)). Building on these priors, VLA baselines align VLMs with robot control through encoding actions as tokens or learning continuous heads, spanning continuous models such as diffusion and flow matching (e.g. DP, Octo, π_0 , $\pi_{0.5}$ (23, 65, 14, 33)), or discrete tokenization with compression (e.g. ACT, OpenVLA, FAST, RT-2) (78, 39, 57, 84). Our recipe follows *discrete priors* \rightarrow *continuous control*: we first implant discrete action priors into the VLM output space (Inspiration), then optimize a flow-matching head for high-frequency actions (Integration), while a tightly coupled MoE with static routing enforces strong language–action binding.

Multimodal co-training with robot data. To offset the cost and coverage limits of robot trajectories, multimodal and cross-source co-training is effective. Initializing from a VLM and jointly training on web image–text, dialogs, and long videos together with multi-embodiment robot data preserves open-world language–vision competence while injecting embodied spatial semantics. Cross-embodiment and cross-environment aggregation further boosts transfer (25, 65, 14, 33, 35, 11, 15). Our approach emphasizes embodied cognition supervision in a unified format that links instruction, reasoning, sub-tasks, and continuous actions; systematic unification of open-source action data with distribution mixing and quotas across two stages to limit drift; and a tightly coupled MoE with static routing that converts heterogeneous supervision into controllable gains on the action channel, improving robustness in unfamiliar scenes and across embodiments.

Language reasoning and subtask decomposition. For long-horizon tasks, combining high-level reasoning with low-level control improves success. Unlike pipeline designs that pair a VLM planner with a separate controller (e.g., SayCan (7) and Code-as-Policies (43); hierarchical variants such as Hi Robot (9) and GR00T N1 (8)), we adopt a single-model Uni-CoT formulation that learns end-to-end mappings from instruction to CoT, sub-tasks, and continuous actions. This complements hierarchical and action-grammar approaches (e.g., RT-H (13)), recent CoT-for-control methods (77, 76) and prompting frameworks (53), as well as instruction-tuning for action (54) and hierarchical goal-conditioned policies (42). The model can include or skip intermediates and interleave reasoning with execution. Coupled with our two-stage curriculum and tightly coupled routing, this unified design reduces interface-induced error accumulation and strengthens progress awareness and instruction following on complex, long-horizon manipulation tasks.

3 WALL-OSS

In this section, we present the overall architecture and training methodology of WALL-OSS. Our objective is to develop a transformer-based foundation model for embodied domains that not only enhances spatial understanding of embodied scenarios within the VLM but also achieves high-quality action generation.

Model Architecture. Figure 2(a) and (b) illustrate the structural paradigms of existing VLA models. In (a), the *mixed design* directly extends the original VLM to model actions through either Discrete Action Modeling

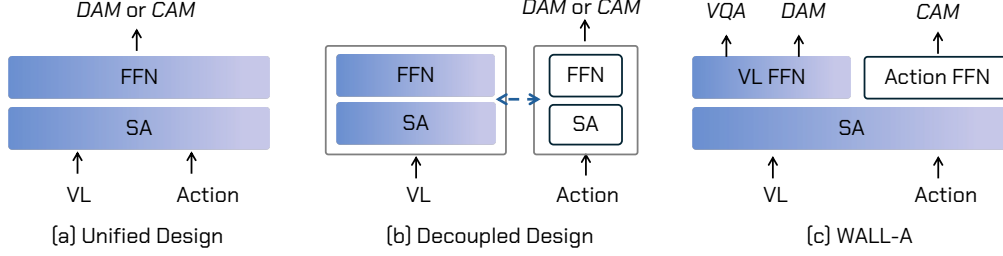


Figure 2 Different paradigms for transferring VLMs to action modeling. The blue parts denote the initialized weights inherited from the pretrained VLM backbone. DAM and CAM refer to Discrete Action Modeling and Continuous Action Modeling, respectively, while VL denotes Vision–Language, and SA denotes Self–Attention.

(DAM) or Continuous Action Modeling (CAM), following the next-token prediction paradigm, as exemplified by RT-2 (84) and OpenVLA (39). However, action supervision significantly perturbs the weight distribution of the original VLM, leading to substantial degradation in action instruction-following and generalization capabilities, resulting in overfitting to actions. In (b), the *decoupled design* employs a separate branch for action prediction that interacts with and extracts information from the VLM, as demonstrated in $\pi 0$ (14). The limitation is that vision and language merely serve as auxiliary signals for action generation; this loosely coupled architecture weakens the instruction-following capability of action prediction. Our architecture addresses these limitations by leveraging a Mixture-of-Experts (MoE) (64) design and assigning different FFNs to corresponding training tasks, thereby forming a tightly coupled cross-modal structure that effectively enhances the model’s cross-modal association capabilities.

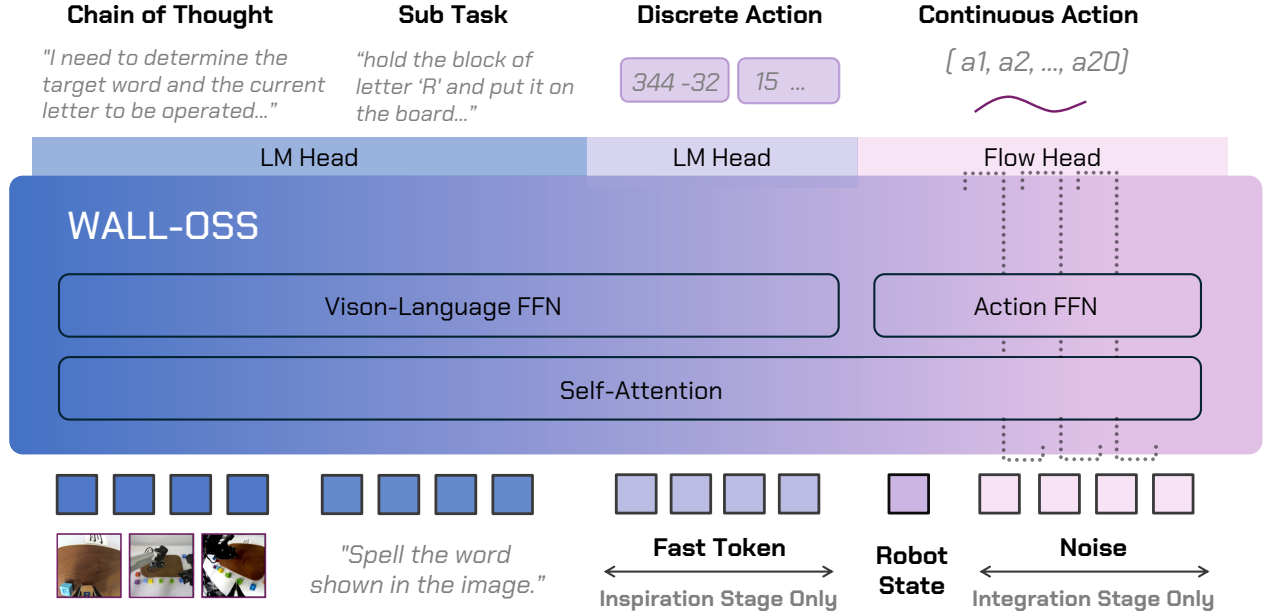


Figure 3 Architecture of **WALL-OSS**.

Figure 3 presents the overall architecture of **WALL-OSS**. We adopt QwenVL2.5-3B as the main backbone. The model takes as input vision (egocentric and arm-mounted camera views) together with text instructions and produces different outputs depending on the training phase, while remaining conditioned on the same multimodal inputs throughout. Let $\mathbf{c} = (\text{vision}, \text{instruction})$ and $\mathbf{h} = F_{\theta}(\mathbf{c})$ denote the input pair and its VLM encoding, respectively.

Following the pretraining of VLMs, our pretraining process comprises two main components: **Inspiration** of the VLM and **Integration** of the three modalities (Vision–Language–Action), as shown in Figure 4.

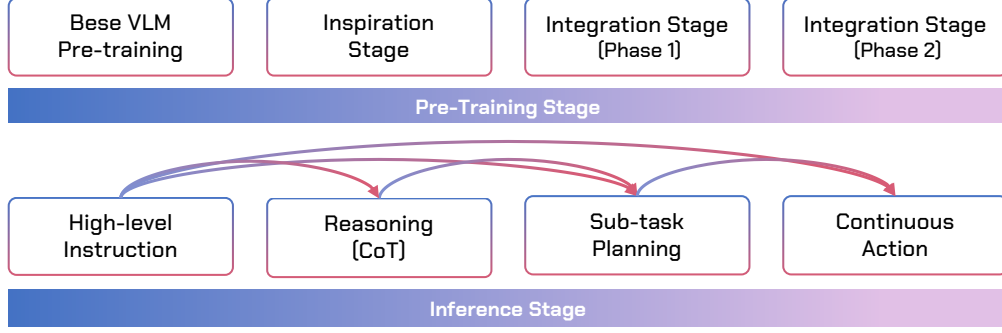


Figure 4 Overview of training and inference pipeline.

Inspiration Stage. We begin by reusing the original feed-forward networks (FFNs) of the pretrained VLM and augmenting them with embodied VQA to strengthen spatial reasoning in robotic environments. The training objectives encompass masked language modeling, image/video-text contrastive learning, instruction following, and temporal order/causality modeling to build strong grounded VL priors. In parallel, we introduce a discrete action objective in which textual tokens are aligned with discrete action tokens obtained via FAST tokenization (cf. $\pi 0$ -FAST (14)):

$$\underbrace{\mathbf{z}_{1:K} = \text{FAST}(\mathbf{a})}_{\text{DCT} \rightarrow \text{Quant} \rightarrow \text{BPE}}, \quad \mathcal{L}_{\text{Inspiration}} = \lambda_{\text{VQA}} \sum_t -\log p_{\theta}(\tau_t | \tau_{<t}, \mathbf{c}) + \lambda_{\text{D}} \sum_k -\log p_{\theta}(z_k | z_{<k}, \mathbf{c}).$$

This phase equips the VLM with coarse, semantically grounded action awareness; its outputs include chain-of-thought reasoning, sub-task predictions, and discrete FAST action tokens. The embodied VQA enhances the VLM’s spatial understanding capabilities in embodied environments, while the FAST token prediction provides coarse action comprehension. This inspires the initial VLM with fundamental embodied reasoning and action awareness.

Integration Stage. Building on these priors, we replace discrete action prediction with continuous action modeling via flow matching, and partition this stage into two phases: (1) freeze the VLM and train only the flow head under the Action FFN, and (2) unfreeze the VLM for joint optimization. For the integration stage, vision, language, and action representations interact through attention, while a static router directs action-centric features to an Action FFN and vision-language features to a Vision-Language FFN, instead of a learned softmax/top-k router. In the first phase, we form noisy samples and regress the velocity field

$$x_t = (1 - \rho(t)) x_0 + \rho(t) \epsilon, \quad \mathcal{L}_{\text{Integration}} = \lambda_{\text{C}} \mathbb{E} \left[w(t) \| v_{\phi}(x_t, \mathbf{h}, t) - (\epsilon - x_0) \|_2^2 \right],$$

with gradient routing $\partial \mathcal{L}_{\text{Integration}} / \partial \theta \approx 0$ (**stop-grad** to the VLM from the previous stage to mitigate catastrophic forgetting) and $\partial \mathcal{L}_{\text{Integration}} / \partial \phi_{\text{action}} \neq 0$ (optimize the flow head and Action FFN only). The supervision from the Inspiration stage stabilizes cross-modal attention here, preserving VL priors and providing reliable initialization for continuous action.

For the second phase, we jointly optimize both modules using the same flow-matching objective with gradient routing $\partial \mathcal{L}_{\text{Integration}} / \partial \theta \neq 0$ and $\partial \mathcal{L}_{\text{Integration}} / \partial \phi \neq 0$ (unfreeze the VLM and train jointly with the action branch). This integration enables fine-grained action output within the tightly coupled architecture, compelling the model to integrate multimodal information and efficiently accomplish multimodal tasks, thereby achieving alignment across different modalities.

Unified Cross-Level CoT. Our approach generalizes the notion of chain-of-thought reasoning beyond traditional narrow-sense CoT (step-by-step textual reasoning in LLMs) to encompass a broad-sense CoT that spans the entire semantic-to-sensorimotor spectrum: instruction \rightarrow reasoning (CoT) \rightarrow subtask plan \rightarrow continuous actions. This unified formulation enables forward arbitrary mapping across hierarchical abstraction levels, allowing the model to seamlessly transition between high-level deliberation and low-level execution within a single differentiable framework.

Existing approaches often decompose these levels into pipeline-style or multi-module systems, where an instruction is first passed to a planner to generate a high-level plan, which is then executed by a controller, such as Hi Robot (9) and GR00T N1 (8). While these paradigms can substantially reduce the difficulty of training and execution in the short term, they introduce non-differentiable interfaces, limit system performance by the capacity of each module, and tend to accumulate errors across stages.

In contrast, **WALL-OSS** employs a single, end-to-end model that jointly learns mappings across abstraction levels, supporting flexible forward arbitrary mapping where the model can adaptively leverage full intermediate reasoning steps or directly map instructions to actions based on task complexity and contextual demands.

Formally, let v denote visual input, x the language instruction, c an optional chain-of-thought, and $a_{1:T} \in \mathbb{R}^{T \times d}$ the target action trajectory. Uni-CoT trains a unified predictor F_θ with a path-drop objective that allows the model to condition on or bypass intermediate reasoning flexibly:

$$\min_{\theta} \mathbb{E}_{(v,x,c,a)} \left[\ell_{\text{act}}(F_\theta(v, x, c), a_{1:T}) + \lambda \ell_{\text{VQA}}(H_\theta(v, x), y) \right],$$

where H_θ is an embodied-aware VQA head with supervision y . This formulation enables both full-chain and direct $x \rightarrow a$ training within a single model, ensuring end-to-end differentiability while retaining flexibility in reasoning.

By covering mappings across semantic levels, Uni-CoT (i) enhances spatial understanding in VLMs through embodied VQA, leading to stronger grounding and action prediction, and (ii) significantly improves long-horizon task success rates and instruction following. Moreover, as shown in the lower part of Figure 4, during inference, Uni-CoT can adaptively decide whether to invoke CoT/subtask decomposition, and even interleave reasoning with execution—emitting actions for completed subtasks while continuing to reason—thus enabling flexible, asynchronous control for real-time human-robot interactions.

4 Data Composition

We construct an embodiment-centric, multisource dataset to address the lack of large-scale, aligned VLA supervision and the spatial understanding gaps of current VLMs. The corpus exceeds 10,000 hours and comprises three complementary components: (1) self-collected robot action data for high quality and task complexity, (2) open-source action data for cross-morphology and cross-environment generalization, and (3) multimodal VQA data to preserve and strengthen language-vision ability while providing additional supervision for spatial-temporal and reasoning, as shown in Figure 5.

To align with our two-stage training recipe (Inspiration and Integration), we functionally orchestrate the data sources: (1) Inspiration focuses on embodied VQA, instruction following, and discrete action priors via FAST to inject coarse action awareness into the VLM while improving spatial reasoning; (2) Integration focuses on high-frequency continuous control with flow matching over real and unified open-source trajectories, first training the action branch and then jointly optimizing with the VLM to tighten language-vision-action alignment and mitigate forgetting.

4.1 Self-collected robot action data

Platforms and scenes. We target real-world daily manipulation across desktop arms, mobile stands, wheeled bi-arm systems, and wheeled humanoids, with egocentric/exocentric and arm-mounted cameras. Scenes span kitchen cleaning, dressing/organizing, mobile pick-and-place, and assembly tasks.

Task spectrum. We include two core families: (1) short-horizon manipulation under explicit instructions emphasizing precision and generalization; and (2) long-horizon reasoning tasks with goal clarity but implicit procedures, requiring task decomposition, progress tracking, and real-time decisions to complete multi-step routines.

Annotation and quality control. We build a multi-model pipeline for fine-grained step annotations with human spot checks, enabling CoT-style stage supervision directly on trajectories. We enforce multi-sensor timestamp synchronization, outlier filtering, low-quality/idle-frame removal, rule-based validation plus manual audits, and



Figure 5 Overview of the multisource dataset. Left: composition across three sources (self-collected actions, open-source actions, and multimodal VQA). Middle (top to bottom): example images from self-collected actions, open-source actions, and multimodal VQA. Top-right: our representative robot hardware.

automatic augmentations (illumination/background). To balance with open-source actions and multimodal corpora, we select a representative self-collected subset for WALL-OSS training.

4.2 Open-source action data

Coverage and sources. We aggregate diverse datasets to complement sensor forms, operation styles, and environmental variation, with representative sources including DROID, BC-Z, BRIDGE, and others (see Table 1 for a complete list).

Standardization. To enable stable joint pretraining across heterogeneous sources, we enforce a unified specification: (1) coordinate frames and units (meters for position; radians for angles); (2) morphology normalization via a maximally expressive DoF template with masking/placeholders for missing joints across single-/dual-arm, wheeled, and humanoid platforms; (3) perception alignment by unifying intrinsics/extrinsics and timestamps, resampling frame rate/resolution, and channel-aligning multi-view video; and (4) action time-base normalization by standardizing control frequencies and resampling/interpolating trajectories to the flow-matching grid. This protocol reduces cross-source friction and improves optimization stability.

4.3 Multimodal VQA: general and embodied

We adopt two VQA streams: a general, perception-centric stream to maintain VLM capability, and an embodied stream that targets spatial-temporal and task reasoning for manipulation. **General VQA for VLM maintenance.** We combine open-source VQA datasets (e.g., CapsFusion, RoboPoint, Robo2VLM, COCO; see Table 1) and general VQA annotations over self-collected trajectories to maintain and enhance language-vision competence and instruction following. These signals are weakly related to manipulation and primarily regularize the vision-language backbone rather than the action head.

Embodied VQA for spatial-temporal and reasoning. We build an automatic generation pipeline over self-collected trajectories to supply supervision that targets spatial understanding and task reasoning: Action Planning VQA (subgoal reasoning), Spatial & Temporal QA (localization, ordering, progress), Perception VQA (attributes), and Cognition & Affordance (interactability). Labels use a unified format for cross-task sharing, e.g., `<box>[x1,y1,x2,y2]</box>`, `<point>[x,y]</point>`, optional `<mask>` encodings, and natural-language

Category	Datasets
Open-source action data	Agibotworld (16); Droid (38); Bc_z (34); RH20T (26); Furniture_bench (31); Fractal (27); Bridge_data_v2 (68); DobbE (62); UMI-biarm (48); Utaustin_mutex (63); Stanford_hydra (12); Austin_sailor (52); Fmb (50); Austin_sirius (47); Taco_play (61, 51); Stanford_kuka_multimodal (40); Berkeley_autolab_ur5 (18); Jaco_play (24); Viola (82); Berkeley_fanuc_manipulation (81); Berkeley_cable_routing (49); Austin_buds (83); Dlr_edan_shared_control (67, 58); Nyu_rot (30)
Open-source image-text data	CapsFusion (74); Cambrian (66); PixMo (1); RoboPoint (75); Robo2VLM (17); COCO (45); VQAv2 (29); VQASynth-SpaceLLaVA (60); SpaceThinker (60); OpenSpaces_MC_R1 (60); OpenSpaces (60); SpaceOm (60)

Table 1 Summary of open-source datasets used in our corpus. Names are shown explicitly; citations appended where available.

QA/description. These signals explicitly link the chain “instruction \rightarrow CoT \rightarrow sub-tasks \rightarrow continuous actions” and directly serve Uni-CoT and our tightly coupled MoE architecture.

4.4 Splits, sampling, and scale

We stratify by scene/object/task/morphology to form cross-environment and cross-morphology validation/test sets. Long-horizon and rare skills receive temperature-controlled resampling and hard-example upsampling to strengthen long-range dependencies and progress modeling. During both training stages, we mix sources with quota control to preserve VLM capabilities and improve continuous control jointly. We apply light domain randomization and occlusion perturbations to the visual stream for robustness. The overall corpus exceeds tens of thousands of hours; source-wise proportions across duration, scenes, and morphologies, as well as task breakdowns. We provide qualitative visualizations and show in ablations that embodied VQA plus discrete action priors significantly improve subsequent flow-matching control quality and instruction adherence, especially in long-horizon tasks.

5 Experiments

5.1 Setup

To comprehensively evaluate the capabilities of our VLA model, we designed an evaluation suite comprising an Embodied VQA benchmark and six robotic manipulation tasks. This suite is centered around three core dimensions: 1) language instruction understanding, reasoning, and generalization; 2) planning and execution for long-horizon, multi-stage tasks; and 3) action accuracy and robustness.

5.1.1 Evaluation Tasks

Embodied Visual Question Answering To evaluate the model’s understanding of embodied scenes, we developed an Embodied VQA benchmark. Based on video frames randomly sampled from our action dataset, this benchmark requires the model to perform three tasks: (1) Scene Captioning, which generates a textual description for a given image; (2) Object Grounding, which locates the 2D coordinates of an object specified by a textual description; and (3) Action Planning, which plans the next action in natural language given an image and a high-level instruction. We compared the performance of our pre-trained WALL-OSS model against the original base VLM (qwen2.5-vl-3b) on this benchmark via manual evaluation.

Robotic Manipulation Tasks We designed six manipulation tasks to assess the model’s performance across our core dimensions. Among these, set-table, tidy-bedroom, and place-by-color are novel tasks unseen during pre-training, designed to evaluate the model’s adaptation capability on novel tasks. The tasks are grouped by the primary capability they target:

For instruction following and reasoning:

Instruction-pick-place: The model is required to follow diverse natural language instructions—such as descriptive commands—to pick up a specified object and place it at a designated target location. This task is excluded from fine-tuning and used to evaluate zero-shot generalization.

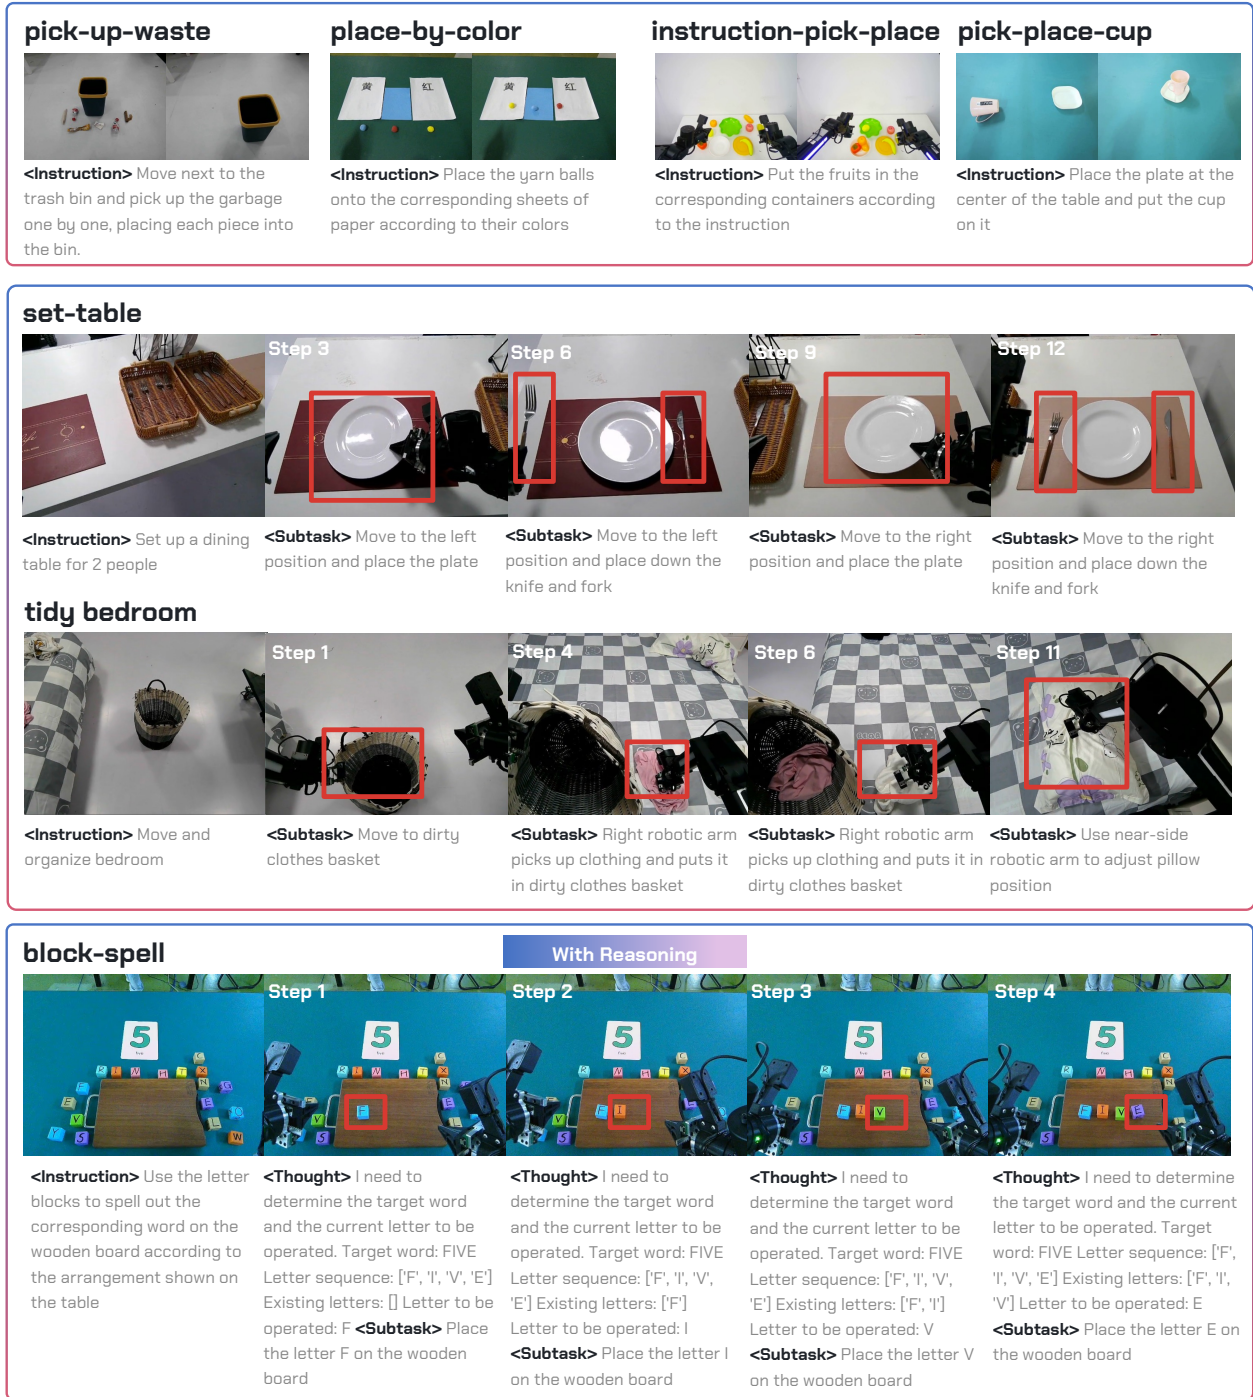


Figure 6 Overview of evaluation tasks used during fine-tuning. Top: single-instruction tasks (Pick-Up-Waste, Place-by-Color, Instruction-Pick-Place, Pick-Place-Cup). Bottom: long-horizon and reasoning tasks (Set-Table, Block-Spell, Tidy-Bedroom).

Place-by-color: The model must place yarn balls onto paper of the corresponding color. This task includes two conditions: (1) direct visual color matching (e.g., placing a red yarn ball on red paper), and (2) text-based reasoning (e.g., placing the ball on white paper printed with the word "red"). 500 episodes are used for fine-tuning.

Block-spell: This task requires the model to spell out answers using toy blocks. It evaluates the model’s reasoning ability, as the answer must be fully inferred before any spelling action can be taken. The prompts fall into two categories: (1) identifying and spelling the name of an object depicted on an image card, and (2) solving a simple arithmetic expression and spelling the resulting answer. 1600 episodes are used for fine-tuning.

For long-horizon planning:

Set-table: The model must set two table placements by taking a plate, a knife, and a fork from the center of the table and arranging them correctly for each setting (plate in the middle, cutlery on either side). 1500 episodes are used for fine-tuning.

Tidy-bedroom: The model needs to tidy a bedroom by collecting dirty clothes from the bed, placing them into a laundry basket, and arranging the pillows correctly. 1000 episodes are used for fine-tuning.

For action accuracy and robustness:

Collect-waste: The model must navigate to a trash can, collect scattered waste from the surrounding area, and perform accurate disposal. This task involves dynamic object interaction and spatial reasoning under mobility constraints. 900 episodes are used for fine-tuning.

Pick-place-cup: The model must place a cup onto a plate centered on the table. This task challenges the model with varied initial states—such as an inverted cup, a misoriented handle, or a misplaced plate—that require a dynamic sequence of manipulations like reorienting the cup and repositioning the plate before the final placement. 500 episodes are used for fine-tuning.

5.1.2 Baselines and Fine-tuning Configuration

We compare WALL-OSS against two representative baselines that capture distinct paradigms in VLA learning: π_0 (14), which leverages pre-trained VLMs with flow matching for action generation, and Diffusion-Policy (23), which employs conditional denoising diffusion for visuomotor control without VLM initialization.

To ensure comprehensive evaluation, we implement two instruction paradigms for baseline models: *Flat* training with high-level task instructions, and *GPT4-Subtask* training using human-annotated subtask decompositions with GPT-4 inference-time generation. This design enables us to isolate the contribution of task decomposition from architectural innovations.

For fair backbone comparison, WALL-OSS and π_0 leverage pre-trained VLM weights, while Diffusion-Policy trains from scratch due to its smaller capacity. All models receive identical task supervision, with WALL-OSS uniquely supporting joint action-VQA training through interleaved sampling (1:15 ratio for non-subtask VQA, 1:100 for subtask samples).

5.1.3 Evaluation Protocol

To ensure objectivity and fairness in evaluation, we implemented a rigorous blind third-party testing protocol. The model trainers provided a detailed testing document that specified the environment setup, initial state for each task, and a standardized scoring rubric. Subsequently, third-party evaluators, who were not involved in the model’s development or training, conducted the assessments. They were blind to the version of the model being tested and strictly adhered to the document provided to perform evaluations and record results. This protocol minimizes subjective bias and ensures the reliability of our evaluation outcomes.

5.2 Results and Analysis

We present our experimental results, starting with an evaluation of the model’s foundational scene understanding and progressing to its performance on increasingly complex manipulation tasks.

model	Object Grounding	Scene Captioning	Action Planning
qwen2.5-vl-3b	46.1%	57.7%	59.8%
WALL-OSS	91.6%	87.6%	69.0%

Table 2 Embodied VQA Benchmark Results.

5.2.1 Pre-training Enhances Embodied Scene Understanding

We first verify that our pre-training strategy not only preserves the VLM’s visual-language alignment but also enhances its deep understanding of embodied scenes before assessing action generation capabilities. As shown in Table 2, we compare WALL-OSS with its original base VLM (qwen2.5-vl-3b) on our self-collected Embodied VQA benchmark. In the Object Grounding task, WALL-OSS achieves significantly superior performance; the base VLM, unfamiliar with embodied manipulation scenes that are rare in its pre-training data, is often misled by elements such as the robot arm, leading to localization failures. A similar trend emerges in the Scene Captioning task, where the base VLM often generates hallucinations irrelevant to the scene, whereas WALL-OSS accurately describes the manipulation actions of the robot arm. Finally, in Action Planning, both models are prone to misjudging the current task stage, yet WALL-OSS provides more specific and on-topic responses than the base model. These findings collectively demonstrate that our pre-training strategy successfully injects robot-centric scene knowledge into the VLM, laying a critical foundation for its performance on subsequent manipulation tasks.

5.2.2 Zero-Shot Instruction Following Emerges After Pre-Training

Having established the model’s enhanced scene understanding, we next assess whether **WALL-OSS** can follow novel instructions without task-specific fine-tuning. This analysis focuses on verifying the model’s ability to generalize to unseen task formulations by leveraging its pre-trained capabilities.

As illustrated in Figure 7, we evaluate the model on a pick-and-place task under two instruction types: (1) those involving objects and containers encountered during pre-training, and (2) those featuring entirely novel items. The model achieves an average task progress of 85% on seen-object instructions and maintains a performance of 61% on novel-object instructions. Most failure cases are attributed to minor pose inaccuracies when selecting grasping or placement positions for unfamiliar objects, rather than semantic misinterpretation of the instructed targets.

These results demonstrate that WALL-OSS can accurately interpret and execute manipulation instructions in a zero-shot setting. This capability is enabled by its multimodal pre-training, which preserves the VLM’s inherent reasoning abilities and supports robust generalization to novel tasks.

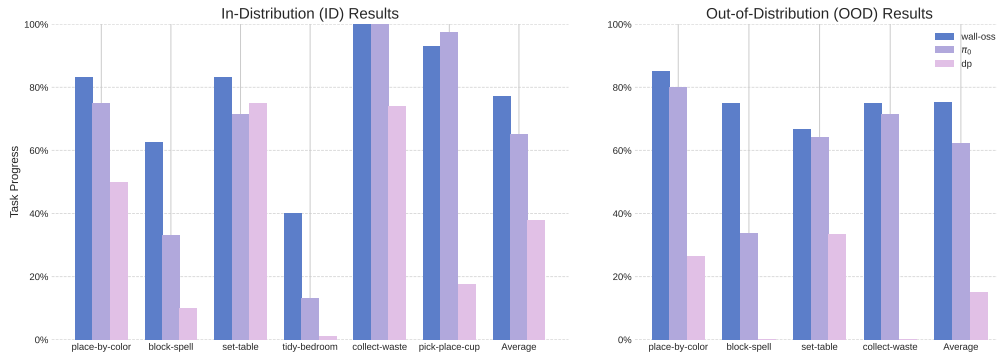


Figure 7 Performance comparison with state-of-the-art policies across all evaluation tasks. Left: In-distribution (ID) performance comparison. Right: Out-of-distribution (OOD) performance

5.2.3 Pre-training Boosts Action Accuracy and Generalization

In addition to instruction comprehension and reasoning, effective embodied intelligence requires precise and robust action execution. We therefore evaluate the impact of pre-training on raw action accuracy and generalization through two tasks: Collect-Waste and Pick-Place-Cup.

The results show that with sufficient demonstration data (Collect-Waste, 1000 demonstrations), the pre-trained models—WALL-OSS and π_0 —achieve a 100% success rate on the in-distribution (ID) task, while the non-pre-trained DP model reaches only 80%. However, when the task complexity increases and data availability is reduced (Pick-Place-Cup, 500 demonstrations), the performance gap widens substantially: the pre-trained models maintain a success rate above 90%, whereas DP drops below 20%. In the out-of-distribution (OOD) generalization test—performing Collect-Waste in a novel environment—DP’s success rate falls from 80% to 0%, completely failing to complete the task. In contrast, both WALL-OSS and π_0 maintain success rates above 80%, demonstrating strong robustness.

These results clearly indicate that pre-training is essential for VLA models, significantly improving action accuracy, data efficiency (i.e., few-shot learning), and generalization to novel environments.

5.2.4 Subtask Generation Improves Performance on Long-Horizon Tasks

Long-horizon tasks present significant challenges for embodied intelligence models due to their inherent complexity, delayed supervision, and the compounding effect of execution errors across stages. To address this, we investigate whether dynamically decomposing high-level goals into intermediate subtasks—generated by a language model—can improve task success rates and execution stability. We design two long-horizon tasks **Set-Table** and **Tidy-Bedroom** that are not included in the pre-training data to evaluate the generalization performance of WALL-OSS. Each task comprises more than five sequential stages, with average execution times exceeding three minutes for Set-Table and five minutes for Tidy-Bedroom.

During fine-tuning, WALL-OSS is jointly trained on action generation and subtask prediction. Despite having only 1% of the training data labeled with subtasks, the model consistently learns to generate high-quality subtask instructions. At inference time, WALL-OSS first generates a subtask instruction and then uses it to guide action generation. As shown in Figure 7, WALL-OSS significantly outperforms both π_0 and Diffusion-Policy across the Set-Table and Tidy-Bedroom tasks. Qualitative analysis indicates that while all models exhibit comparable performance on low-level actions (e.g., picking up a plate), the baselines frequently suffer from stage confusion due to the absence of explicit subtask guidance. This manifests as repetitive or misplaced actions—such as repeatedly placing cutlery at the same location or failing to set the remaining parts of the table—because the model lacks a clear notion of task progression. Moreover, when multiple plausible actions exist under similar visual observation, baseline models tend to struggle with disambiguation and often make inconsistent or suboptimal choices, further compounding execution errors across stages. In the more spatially complex Tidy-Bedroom task, these issues become more pronounced: baseline models often get stuck in repetitive, ineffective behaviors—especially when target objects (e.g., a missed piece of clothing) fall outside the current field of view.

In contrast, WALL-OSS leverages its multimodal reasoning capabilities to generate contextual subtask cues, enabling it to disambiguate between plausible options, maintain coherent task progression, and reliably complete all stages.

5.2.5 Chain-of-Thought (CoT) Empowers Embodied Planning and Reasoning

Complex embodied tasks often require reasoning over spatial relationships, object attributes, and the consequences of sequential actions. While existing VLA models directly map instructions to actions, they often struggle with tasks that demand intermediate decision-making and commonsense reasoning. To address this, we investigate the role of CoT prompting in enabling WALL-OSS to plan and execute multi-step behaviors through explicit intermediate reasoning.

Inspired by the success of large language models (LLMs), we explore transferring CoT reasoning capabilities to VLA models. To evaluate the integration of reasoning and action execution, we design two tasks—Place-by-Color and Block-Spell—that require intermediate logical inference. During fine-tuning, WALL-OSS is

block type	WALL-OSS (Co-training)	WALL-OSS (Action-only)	π_0 (Action-only)
letter	87%	26%	9%
number	95%	80%	35%

Table 3 Instruction-following accuracy in block-spell task.

trained to jointly generate a CoT reasoning trace and a sub-instruction, using supervision on only 1% of the training frames. At inference time, the model autonomously generates both a CoT reasoning trace and a sub-instruction, which are then used as conditional inputs for action generation. In the Place-by-Color task, we observe that CoT provides minimal benefit when the goal involves direct visual matching (e.g., matching a red yarn ball to red paper).

However, in the more challenging condition requiring text-based reasoning—such as placing the yarn ball on paper printed with the word “red”—WALL-OSS achieves a significantly higher task completion rate than all baselines. This demonstrates that direct-action models are limited to visually intuitive tasks, while incorporating CoT is essential for solving instruction-conditioned tasks that require intermediate reasoning.

In the Block-Spell task, baseline models under the flat setting achieve near-zero task progress, reaffirming that the absence of task decomposition is detrimental for reasoning-intensive tasks. Consequently, in Figure 7, we report baseline results only under the GPT-subtask setting. As shown, WALL-OSS significantly outperforms the GPT-subtask baseline in both in-distribution (ID) and out-of-distribution (OOD) scenarios. Our analysis suggests that while GPT-4 can often infer correct sub-instructions, its response lacks the timeliness and contextual adaptability exhibited by WALL-OSS—particularly in complex settings such as obstructed egocentric views. Moreover, baseline models exhibit a disconnect between the high-level policy generated by GPT and the low-level action execution module, leading to poor alignment and degraded instruction following performance.

5.2.6 Multi-modal Co-training Enhances Fine-grained Instruction Following

Beyond general instruction following, many real-world tasks require precise grounding of fine-grained visual and linguistic cues. We therefore examine the effectiveness of multi-modal co-training in improving the model’s ability to interpret and execute detailed, object-specific instructions.

As shown in Table 3, WALL-OSS achieves significantly higher accuracy than π_0 on the Block-Spell task, particularly in identifying and placing the correct sequence of letter blocks. In the same task, we observed that π_0 performed substantially worse than WALL-OSS in identifying and manipulating specific letter blocks. To investigate this gap, we conducted an ablation study comparing three fine-tuning configurations:

- WALL-OSS (Multi-modal Co-training): Fine-tuned by jointly training on action generation, CoT + subtask generation, and 2D referring expression grounding.
- WALL-OSS (Action-only): Fine-tuned using only subtask instructions for action generation.
- π_0 (Action-only): A VLA with a similar architecture to WALL-OSS, fine-tuned only for action generation.

As shown in Table 3, we evaluate each model’s ability to select the correct block when provided with a precise instruction (e.g., “pick up the letter A”). The results indicate that the WALL-OSS (Multi-modal Co-training) configuration achieves the highest instruction-following accuracy. While the performance of WALL-OSS (Action-only) declines relative to the multi-modal setup, it remains significantly better than π_0 (Action-only), whose accuracy is close to random chance.

These findings clearly demonstrate that multi-modal co-training substantially enhances a model’s ability to follow fine-grained instructions. The pre-training phase of WALL-OSS establishes a strong foundation for multimodal action alignment, and maintaining this co-training strategy during fine-tuning further amplifies the model’s instruction-grounded execution capabilities.

6 Discussion

6.1 A Unified Theory of Embodied Foundation Models

The emergence of foundation models has fundamentally reshaped our understanding of artificial intelligence, yet their extension to embodied domains reveals deeper questions about the nature of intelligence itself: How large are the intrinsic gaps across modalities? Can a unified architecture process heterogeneous modalities without degrading any of them? If so, which approximation trajectories reach that state most efficiently?

WALL-OSS is designed to bridge the language–vision–action gap. Our tightly coupled MoE architecture, multimodal curriculum that augments VLMs with embodied understanding, and multi-stage training schedule collectively yield a unified, flexible, differentiable, end-to-end mapping from high-level instructions—through narrow-sense CoT and subtask decomposition—to discrete and then continuous actions. The rich forward mappings learned during training correspondingly confer flexibility at inference.

Early explorations that aggressively discretized nonlinguistic modalities and trained them jointly with text under a next-token objective proved unsatisfactory. The field subsequently shifted toward conservative side-branches for new modalities, during which original VL priors were heavily disrupted. By contrast, WALL-OSS employs a smaller fraction of additional parameters yet demonstrates that simultaneously improving embodied VL understanding and action generation is both natural and feasible, rather than presupposing a zero-sum trade-off between VL and action capabilities.

Overall, WALL-OSS achieves: (i) enhanced multimodal grounding and reasoning ability in embodied environments, thanks to embodied VQA and Uni-CoT design; (ii) competitive action generation with stronger OOD generalization than DP, though Pi-0 remains superior for precise manipulation; (iii) measurable gains from the Inspiration stage, validating our staged training strategy. These findings support our claim that tightly coupled multimodal reasoning and action modeling are key to scalable embodied intelligence.

6.2 Future Directions

While unified text-image architectures reach related conclusions, action remains qualitatively distinct: signals are sparse, lie on high-dimensional manifolds, and a compact trajectory corresponds to vast visual and linguistic representations. These representations span high-level state analysis, intent formation, and task decomposition, whereas execution depends on low-level spatial perception and kinematics. Effective modeling must therefore align high-level textual/visual reasoning with low-level perception and control.

Motivated by the former, we introduce large language models; for the latter, we ask whether intermediate formulations (e.g., future-frame prediction) or intermediate modalities (e.g., 3D perception) are needed to reduce the difficulty of learning the VL→A mapping. We view strictly end-to-end VL→A and the intermediate-route approaches as two convergent paths, focusing on their relative efficiency toward AGI. Video and 3D modalities typically exhibit higher correlation with action than text–action pairs, effectively de-sparsifying supervision; yet 3D data remain scarce. Large-scale video prediction can leverage internet data but introduces substantial redundant supervision—humans act without predicting all details and perceive selectively by intent—prompting bidirectional world-modeling for behavior prediction. With human media largely unchanged, the growth of data and algorithms in each modality is driven by adjacent industries (e.g., AIGC, graphics), which also nurture the "intermediate" priors needed for action modeling. Progress in end-to-end action prediction depends on the embodied industry and likewise amounts to implicit learning of such intermediates. As demand stimulates both academia and industry, the relative efficiency of different routes toward AGI is inherently dynamic.

References

- [1] Pixmo dataset. https://docs.allenai.org/training_data/pixmo, 2024. Dataset website.
- [2] Introducing llama 3.1: Our most capable models to date. <https://ai.meta.com/blog/meta-llama-3-1/>, 2024. Meta AI Blog.
- [3] Claude 3.7 sonnet and claude code. <https://www.anthropic.com/news/claude-3-7-sonnet>, 2025. Anthropic News.
- [4] Gemini models. <https://deepmind.google/models/gemini/>, 2025. Google DeepMind Models page.
- [5] The llama 4 herd: The beginning of a new era of natively multimodal intelligence. <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>, 2025. Meta AI Blog.
- [6] Michael Ahn, Anthony Brohan, Noah Brown, et al. Do as i can, not as i say: Grounding language in robotic affordances. In *Robotics: Science and Systems (RSS)*, 2022.
- [7] Michael Ahn et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022. <https://arxiv.org/abs/2204.01691>.
- [8] First Author and Second Author. Gr00t n1: Modular reasoning and control for embodied agents. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [9] First Author and Second Author. Hi robot: Towards hierarchical instruction-to-action modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [10] Shuai Bai, Keqin Chen, Xuejing Liu, et al. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. <https://arxiv.org/abs/2502.13923>.
- [11] Bowen Baker, Ilge Akkaya, Peter Zhokov, Joost Huizinga, Jie Tang, Adrien Ecoffet, Brandon Houghton, Raul Sampedro, and Jeff Clune. Video pretraining (vpt): Learning to act by watching unlabeled online videos. In *Advances in Neural Information Processing Systems*, 2022.
- [12] Suneel Belkhale, Yuchen Cui, and Dorsa Sadigh. Hydra: Hybrid robot actions for imitation learning. In *Conference on Robot Learning*, pages 2113–2133. PMLR, 2023.
- [13] Suneel Belkhale, Tianli Ding, Ted Xiao, Pierre Sermanet, Quon Vuong, Jonathan Tompson, Yevgen Chebotar, Debidatta Dwibedi, and Dorsa Sadigh. Rt-h: Action hierarchies using language. *arXiv preprint arXiv:2403.01823*, 2024.
- [14] Kevin Black et al. *pio*: A vision-language-action flow model for general robot policies. *arXiv preprint arXiv:2410.24164*, 2024.
- [15] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.
- [16] Qingwen Bu, Jisong Cai, Li Chen, Xiuqi Cui, Yan Ding, Siyuan Feng, Shenyuan Gao, Xindong He, Xuan Hu, Xu Huang, et al. Agibot world colosseum: A large-scale manipulation platform for scalable and intelligent embodied systems. *arXiv preprint arXiv:2503.06669*, 2025.
- [17] Kaiyuan Chen, Shuangyu Xie, Zehan Ma, and Ken Goldberg. Robo2vlm: Visual question answering from large-scale in-the-wild robot manipulation datasets, 2025. <https://arxiv.org/abs/2505.15517>.
- [18] Lawrence Yunliang Chen, Simeon Adebola, and Ken Goldberg. Berkeley UR5 demonstration dataset. <https://sites.google.com/view/berkeley-ur5/home>.
- [19] Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling (internvl 2.5). *arXiv preprint arXiv:2412.05271*, 2024. <https://arxiv.org/abs/2412.05271>.
- [20] Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites (internvl 1.5). *arXiv preprint arXiv:2404.16821*, 2024. <https://arxiv.org/abs/2404.16821>.
- [21] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. https://openaccess.thecvf.com/content/CVPR2024/papers/Chen_InternVL_Scaling_up_Vision_Foundation_Models_and_Aligning_for_Generic_CVPR_2024_paper.pdf.
- [22] Cheng Chi, Zhenjia Xu, Brian Ichter, and Avinash Shah. Diffusion policy: Visuomotor policy learning via action diffusion. In *Robotics: Science and Systems (RSS)*, 2023.
- [23] Cheng Chi et al. Visuomotor policy learning via action diffusion. *arXiv preprint arXiv:2303.04137*, 2023. <https://arxiv.org/abs/2303.04137>.
- [24] Shivin Dass, Jullian Yapeter, Jesse Zhang, Jiahui Zhang, Karl Pertsch, Stefanos Nikolaidis, and Joseph J. Lim. Clvr jaco play dataset, 2023. https://github.com/clvr-ai/clvr_jaco_play_dataset.
- [25] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023.

- [26] Hao-Shu Fang, Hongjie Fang, Zhenyu Tang, Jirong Liu, Chenxi Wang, Junbo Wang, Haoyi Zhu, and Cewu Lu. Rh20t: A comprehensive robotic dataset for learning diverse skills in one-shot. *arXiv preprint arXiv:2307.00595*, 2023.
- [27] Charles Gaydon, Michel Daab, and Floryne Roche. Fractal: An ultra-large-scale aerial lidar dataset for 3d semantic segmentation of diverse landscapes. *arXiv preprint arXiv:2405.04634*, 2024.
- [28] Gemini Team. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next-generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. <https://arxiv.org/abs/2507.06261>.
- [29] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [30] Siddhant Haldar, Vaibhav Mathur, Denis Yarats, and Lerrel Pinto. Watch and match: Supercharging imitation with regularized optimal transport. In *Conference on Robot Learning*, pages 32–43. PMLR, 2023.
- [31] Minh Heo, Youngwoon Lee, Doohyun Lee, and Joseph J Lim. Furniturebench: Reproducible real-world benchmark for long-horizon complex manipulation. *The International Journal of Robotics Research*, page 02783649241304789, 2023.
- [32] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.
- [33] Physical Intelligence, Kevin Black, Noah Brown, et al. $\pi_{0.5}$: a vision-language-action model with open-world generalization. *arXiv preprint arXiv:2504.16054*, 2025. <https://arxiv.org/abs/2504.16054>.
- [34] Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. Bc-z: Zero-shot task generalization with robotic imitation learning. In *Conference on Robot Learning*, pages 991–1002. PMLR, 2022.
- [35] Yunfan Jiang, Agrim Gupta, Zichen Zhang, Guanzhi Wang, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu, and Linxi Fan. Vima: General robot manipulation with multimodal prompts. *arXiv preprint arXiv:2210.03094*, 2022.
- [36] Amita Kamath, Jack Hessel, and Kai-Wei Chang. What’s “up” with vision-language models? measuring spatial reasoning. *arXiv preprint arXiv:2310.19785*, 2023.
- [37] Koray Kavukcuoglu. Gemini 2.5: Our most intelligent ai model. <https://blog.google/technology/google-deepmind/gemini-model-thinking-updates-march-2025/>, 2025. Google DeepMind Blog.
- [38] Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint arXiv:2403.12945*, 2024.
- [39] Moo Jin Kim, Ankur Handa, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- [40] Michelle A Lee, Yuke Zhu, Krishnan Srinivasan, Parth Shah, Silvio Savarese, Li Fei-Fei, Animesh Garg, and Jeannette Bohg. Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks. In *2019 International conference on robotics and automation (ICRA)*, pages 8943–8950. IEEE, 2019.
- [41] Chengmeng Li, Junjie Wen, Yan Peng, Yaxin Peng, Feifei Feng, and Yichen Zhu. Pointvla: Injecting the 3d world into vision-language-action models. *arXiv preprint arXiv:2503.07511*, 2025. doi: 10.48550/arXiv.2503.07511. <https://arxiv.org/abs/2503.07511>.
- [42] Yi Li, Yuquan Deng, Jesse Zhang, Joel Jang, Marius Memmel, Raymond Yu, Caelan Reed Garrett, Fabio Ramos, Dieter Fox, Anqi Li, et al. Hamster: Hierarchical action models for open-world robot manipulation. *arXiv preprint arXiv:2502.05485*, 2025.
- [43] Chen Liang, Michael I. Jordan, et al. Code as policies: Language model programs for embodied control. *arXiv preprint arXiv:2209.07753*, 2022. <https://arxiv.org/abs/2209.07753>.
- [44] Jacky Liang, Scott Emmons, Stefanos Nikolaidis, et al. Code as policies: Language model programs for embodied control. *arXiv preprint arXiv:2209.07753*, 2023.
- [45] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014.
- [46] Yaron Lipman, Ricky Chen, Heli Ben-Hamu, Max Nickel, and Manzil Zaheer Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- [47] Huihan Liu, Soroush Nasiriany, Lance Zhang, Zhiyao Bao, and Yuke Zhu. Robot learning on the job: Human-in-the-loop autonomy and learning during deployment. *The International Journal of Robotics Research*, page 02783649241273901, 2022.
- [48] Kehui Liu, Chuyue Guan, Zhongjie Jia, Ziniu Wu, Xin Liu, Tianyu Wang, Shuai Liang, Peng Chen, Pingrui Zhang, Haoming Song, et al. Fastumi: A scalable and hardware-independent universal manipulation interface with dataset. *arXiv preprint arXiv:2409.19499*, 2024.

- [49] Jianlan Luo, Charles Xu, Xinyang Geng, Gilbert Feng, Kuan Fang, Liam Tan, Stefan Schaal, and Sergey Levine. Multistage cable routing through hierarchical imitation learning. *IEEE Transactions on Robotics*, 40:1476–1491, 2024.
- [50] Jianlan Luo, Charles Xu, Fangchen Liu, Liam Tan, Zipeng Lin, Jeffrey Wu, Pieter Abbeel, and Sergey Levine. Fmb: a functional manipulation benchmark for generalizable robotic learning. *The International Journal of Robotics Research*, 44(4):592–606, 2025.
- [51] Oier Mees, Jessica Borja-Diaz, and Wolfram Burgard. Grounding language with visual affordances over unstructured data. *arXiv preprint arXiv:2210.01911*, 2022.
- [52] Soroush Nasiriany, Tian Gao, Ajay Mandlekar, and Yuke Zhu. Learning and retrieval from prior data for skill-based imitation learning. *arXiv preprint arXiv:2210.11435*, 2022.
- [53] Soroush Nasiriany, Fei Xia, Wenhao Yu, Ted Xiao, Jacky Liang, Ishita Dasgupta, Annie Xie, Danny Driess, Ayzaan Wahid, Zhuo Xu, et al. Pivot: Iterative visual prompting elicits actionable knowledge for vlms. *arXiv preprint arXiv:2402.07872*, 2024.
- [54] Dantong Niu, Yuvan Sharma, Giscard Biamby, Jerome Quenum, Yutong Bai, Baifeng Shi, Trevor Darrell, and Roei Herzig. Llarva: Vision-action instruction tuning enhances robot learning. *arXiv preprint arXiv:2406.11815*, 2024.
- [55] OpenAI. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o/>, 2024.
- [56] OpenAI. Introducing gpt-5 for developers. <https://openai.com/index/introducing-gpt-5-for-developers/>, August 2025.
- [57] Karl Pertsch, Kyle Stachowicz, Brian Ichter, et al. Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747*, 2025. <https://arxiv.org/abs/2501.09747>.
- [58] Gabriel Quere, Annette Hagengruber, Maged Iskandar, Samuel Bustamante, Daniel Leidner, Freek Stulp, and Jörn Vogel. Shared control templates for assistive robotics. In *2020 IEEE international conference on robotics and automation (ICRA)*, pages 1956–1962. IEEE, 2020.
- [59] Alec Radford, Jong Wook Kim, Chris Hallacy, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021.
- [60] remyxai. Vqasynth. GitHub repository, 2024. <https://github.com/remyxai/VQASynth/tree/main>.
- [61] Erick Rosete-Beas, Oier Mees, Gabriel Kalweit, Joschka Boedecker, and Wolfram Burgard. Latent plans for task-agnostic offline reinforcement learning. In *Conference on Robot Learning*, pages 1838–1849. PMLR, 2023.
- [62] Nur Muhammad Mahi Shafiullah, Anant Rai, Haritheja Etukuru, Yiqian Liu, Ishan Misra, Soumith Chintala, and Lerrel Pinto. On bringing robots home. *arXiv preprint arXiv:2311.16098*, 2023.
- [63] Rutav Shah, Roberto Martín-Martín, and Yuke Zhu. Mutex: Learning unified policies from multimodal task specifications. *arXiv preprint arXiv:2309.14320*, 2023.
- [64] Noam Shazeer and et al. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *ICLR*, 2017.
- [65] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024.
- [66] Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, Austin Wang, Rob Fergus, Yann LeCun, and Saining Xie. Cambrian-1: A fully open, vision-centric exploration of multimodal llms, 2024.
- [67] Jörn Vogel, Annette Hagengruber, Maged Iskandar, Gabriel Quere, Ulrike Leipscher, Samuel Bustamante, Alexander Dietrich, Hannes Höppner, Daniel Leidner, and Alin Albu-Schäffer. Edan: An emg-controlled daily assistant to help people with physical disabilities. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4183–4190. IEEE, 2020.
- [68] Homer Rich Walke, Kevin Black, Tony Z Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, et al. Bridgedata v2: A dataset for robot learning at scale. In *Conference on Robot Learning*, pages 1723–1736. PMLR, 2023.
- [69] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. https://openaccess.thecvf.com/content/CVPR2025/papers/Wang_VGGT_Visual_Geometry_Grounded_Transformer_CVPR_2025_paper.pdf.
- [70] Junying Wang et al. A4bench: An affordance benchmark for multimodal large language models. *arXiv preprint*, 2025. Placeholder entry; replace with final arXiv ID when available.
- [71] Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, Zhaokai Wang, Zhe Chen, Hongjie Zhang, Ganlin Yang, Haomin Wang, Qi Wei, Jinhui Yin, Wenhao Li, Erfei Cui, Guanzhou Chen, Zichen Ding, Changyao Tian, Zhenyu Wu, Jingjing Xie, Zehao Li, Bowen Yang, Yuchen Duan, Xuehui Wang, Songze Li, Xiangyu Zhao, Haodong Duan, Nianchen Deng, Bin Fu, Yanan He, Yi Wang, Conghui He, Botian Shi, Junjun He, Yingting Xiong, Han Lv, Lijun Wu, Wenqi Shao, Kaipeng Zhang, Huipeng Deng, Jiaye Ge, Biqing Qi, Jiaye Ge, Qipeng Guo, Wenwei Zhang, Wanli Ouyang, Limin

- Wang, Min Dou, Xizhou Zhu, Tong Lu, Dahua Lin, Jifeng Dai, Weijie Su, Bowen Zhou, Kai Chen, Yu Qiao, Wenhai Wang, and Gen Luo. Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*, 2025. <https://arxiv.org/abs/2508.18265>.
- [72] Yifan Wang, Jianjun Zhou, Haoyi Zhu, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen, Jiangmiao Pang, Chunhua Shen, and Tong He. pt^3 : Scalable permutation-equivariant visual geometry learning. *arXiv preprint arXiv:2507.13347*, 2025. doi: 10.48550/arXiv.2507.13347. <https://arxiv.org/abs/2507.13347>.
- [73] An Yang, Baosong Yang, Beichen Zhang, et al. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024. <https://arxiv.org/abs/2412.15115>.
- [74] Qiying Yu, Quan Sun, Xiaosong Zhang, Yufeng Cui, Fan Zhang, Yue Cao, Xinlong Wang, and Jingjing Liu. Capsfusion: Rethinking image-text data at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14022–14032, 2024.
- [75] Wentao Yuan, Jiafei Duan, Valts Blukis, Wilbert Pumacay, Ranjay Krishna, Adithyavairavan Murali, Arsalan Mousavian, and Dieter Fox. Robopoint: A vision-language model for spatial affordance prediction for robotics. *arXiv preprint arXiv:2406.10721*, 2024.
- [76] Michał Zawalski, William Chen, Karl Pertsch, Oier Mees, Chelsea Finn, and Sergey Levine. Robotic control via embodied chain-of-thought reasoning. In *Conference on Robot Learning (CoRL)*, 2024.
- [77] Qingqing Zhao, Yao Lu, Moo Jin Kim, Zipeng Fu, Zhuoyang Zhang, Yecheng Wu, Zhaoshuo Li, Qianli Ma, Song Han, Chelsea Finn, et al. Cot-vla: Visual chain-of-thought reasoning for vision-language-action models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [78] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware, 2023. <https://arxiv.org/abs/2304.13705>.
- [79] Haoyu Zhen, Xiaowen Qiu, Peihao Chen, Jincheng Yang, Xin Yan, Yilun Du, Yining Hong, and Chuang Gan. 3d-vla: A 3d vision-language-action generative world model. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, volume 235 of *Proceedings of Machine Learning Research*, pages 61229–61245. PMLR, July 2024. <https://proceedings.mlr.press/v235/zhen24a.html>.
- [80] Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Yuchen Duan, Hao Tian, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Yue Cao, Yangzhou Liu, Weiye Xu, Hao Li, Jiahao Wang, Han Lv, Dengnian Chen, Songze Li, Yinan He, Tan Jiang, Jiapeng Luo, Yi Wang, Conghui He, Botian Shi, Xingcheng Zhang, Wenqi Shao, Junjun He, Yingtong Xiong, Wenwen Qu, Peng Sun, Penglong Jiao, Lijun Wu, Kaipeng Zhang, Huipeng Deng, Jiaye Ge, Kai Chen, Limin Wang, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahua Lin, Yu Qiao, Jifeng Dai, and Wenhai Wang. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025. <https://arxiv.org/abs/2504.10479>.
- [81] Xinghao Zhu, Ran Tian, Chenfeng Xu, Mingxiao Huo, Wei Zhan, Masayoshi Tomizuka, and Mingyu Ding. Fanuc manipulation: A dataset for learning-based manipulation with fanuc mate 200id robot, 2023.
- [82] Yifeng Zhu, Abhishek Joshi, Peter Stone, and Yuke Zhu. Viola: Object-centric imitation learning for vision-based robot manipulation. In *6th Annual Conference on Robot Learning*, 2022.
- [83] Yifeng Zhu, Peter Stone, and Yuke Zhu. Bottom-up skill discovery from unsegmented demonstrations for long-horizon robot manipulation. *IEEE Robotics and Automation Letters*, 7(2):4126–4133, 2022.
- [84] Brianna Zitkovich, Anthony Brohan, Noah Brown, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Proceedings of Machine Learning Research*, volume 229, 2023. <https://proceedings.mlr.press/v229/zitkovich23a.html>.