




EO-1: Interleaved Vision-Text-Action Pretraining for General Robot Control

Delin Qu^{*}, Haoming Song^{*}, Qizhi Chen^{*}, Zhaoqing Chen^{*}, Xianqiang Gao^{*}, Modi Shi, Guanghui Ren
Maoqing Yao, Bin Zhao, Dong Wang

 <https://eo-robotics.ai/eo-1>
 <https://github.com/eo-robotics>
 <https://huggingface.co/IPEC-COMMUNITY>

The human ability to seamlessly perform multimodal reasoning and physical interaction in the open world is a core goal for general-purpose embodied intelligent systems. Recent vision-language-action (VLA) models, which are co-trained on large-scale robot and visual-text data, have demonstrated notable progress in general robot control. However, they still fail to achieve human-level flexibility in interleaved reasoning and interaction. In this work, we introduce **EO-Robotics**, which consists of **EO-1** model and **EO-Data1.5M** dataset. **EO-1** is a unified embodied foundation model that achieves superior performance in multimodal embodied reasoning and robot control through interleaved vision-text-action pre-training. The development of **EO-1** is based on two key pillars: (i) a unified architecture that processes multimodal inputs indiscriminately (image, text, video, and action), and (ii) a massive, high-quality multimodal embodied reasoning dataset, **EO-Data1.5M**, which contains over 1.5 million samples with emphasis on interleaved vision-text-action comprehension. **EO-1** is trained through synergies between auto-regressive decoding and flow matching denoising on **EO-Data1.5M**, enabling seamless robot action generation and multimodal embodied reasoning. Extensive experiments demonstrate the effectiveness of interleaved vision-text-action learning for open-world understanding and generalization, validated through a variety of long-horizon, dexterous manipulation tasks across multiple embodiments. This paper details the architecture of **EO-1**, the data construction strategy of **EO-Data1.5M**, and the training methodology, offering valuable insights for developing advanced embodied foundation models.

1. Introduction

Generalist robot policies with open-world capabilities is essential for deploying autonomous robots in real scenarios, where they need to tackle diverse tasks, follow various instructions, adapt to different situations, and even handle unforeseen events. In the realm of digital domains, the paradigm of learning from large-scale multimodal datasets has demonstrated broad proficiency and generalization across various interactive and assistive applications. Extending this idea to the physical world, researchers are now training large vision-language-action models on extensive robotic datasets to realize generalist robots with similar versatility. Although recent efforts have demonstrated impressive progress in achieving dexterous manipulation and long-horizon physical control within constrained environments, open world generalization remains the central challenge for developing general-purpose autonomous AI agents in the physical world (Black et al., 2025, Gemini-Robotics, 2025). To address this challenge, robots must be capable of acquiring comprehensive world knowledge, engaging in human-level reasoning, and executing dexterous actions.

Early generalist robot policies (Kim et al., 2024, Black et al., 2024, Pertsch et al., 2025) primarily extend vision-language models (VLMs) into vision-language-action (VLA) models with domain-specific robotic data, either through auto-regressive decoding of discrete action tokens or by incorporating additional continuous flow matching modules. However, because these VLA models are trained exclusively on robotic datasets, they are restricted to narrow task domains and specific environments. As a result, they suffer from diminished general semantic knowledge inherited from VLMs and exhibit limited instruction-following capabilities.

Recently, several studies (Black et al., 2025, Driess et al., 2025, Lin et al., 2025) have explored co-training VLA models with both web data and robotic data, showing promising generality when interacting with new objects and unseen backgrounds. Nevertheless, existing approaches remain predominantly generating robot action at the end of the VLA output sequence, overlooking the rich temporal dynamics and causal dependencies among vision, language, and action modalities inherent in open-world embodied interactions. Humans, in contrast, exhibit a flexible and interleaved synergy between multimodal embodied reasoning and physical action, enabling highly generalizable and dexterous manipulation in the open world, *i.e.*, reasoning guides action and action results inform subsequent reasoning. This raises a fundamental research question: **How can we design an effective training paradigm for generalist robot policies that support flexible and mutual-informed reasoning-acting integration?**

Inspired by the impressive results of advanced multimodal understanding and generation systems (Deng et al., 2025, Ma et al., 2025b, Xie et al., 2025, NextStep-1, 2025), which demonstrate the superiority in interleaved multimodal pretraining, we propose a unified embodied model to enable flexible and powerful multimodal embodied reasoning and action generation through interleaved embodied pretraining. Achieving this requires not only an effective and unified architecture that excels in mix-modality generation, but also carefully structured multimodal embodied data that jointly integrate texts, images, videos, and robotic actions.

To realize this vision, we first established a new protocol for scalable data curating, filtering, and construction of high-quality multimodal interleaved embodied data. As for data source, we integrate web vision-language data with real robot episodes, the latter naturally providing action-level, temporal, and physical continuity. Then, we employ VLMs and human to annotate real robot episodes with diverse embodied temporal and spatial QA pairs, including physical common sense understanding, task planning, object localization, affordance pointing, and multi-view correspondence. These annotations enable the learning of fine-grained geometric and spatial-temporal representations of the physical world. Finally, the interleaved vision-text-action data is constructed by concatenating these multimodal QA and robot control actions in temporal order. Importantly, while robot actions are fixed at each timestep, we design three flexible interleaved formats to associated random embodied reasoning QA pairs. Consequently, the interleaved embodied data capture rich world knowledge and nuanced cross-modal interactions, providing models with foundational capabilities of action prediction, scene understanding, and complex multimodal reasoning for generalist robot policies.

Regarding model architecture, we adopt a single unified decoder-only transformer that integrates discrete auto-regressive decoding with continuous flow matching denoising. The unified model is built on a pre-trained VLM, thereby inheriting broad visual-language knowledge, and its shared parameters are further optimized with modality-specific objectives: next-token prediction for text and flow matching for robotic actions. The model applies causal attention cross the entire interleaved vision-text-action sequence to capture the sequential dependencies between reasoning and acting. In addition, two MLPs are added to encode and decode continuous robotic actions, complementing the original text and visual tokenizers. Unlike prior VLA models (Black et al., 2024, 2025, Driess et al., 2025) that introduce extra action-specific modules to learn action generation, our design enables easier alignment between vision/language and action modalities through circumventing train new action-specific parameters from scratch, thereby achieving more effective cross-modal knowledge transfer in generalist robot policies.

Overall, we present the EO-Robotics toolchain for advanced embodied foundation model development, where **EO-1**, a unified embodied foundation model comprising 3B parameters, is trained on the carefully curated interleaved embodied dataset **EO-Data1.5M**. As a generalist VLA, **EO-1** exhibits strong generalization capabilities in multimodal embodied reasoning and real robot control across

a diverse set of challenging tasks, including pick-and-place, articulated manipulation, long-horizon planning and dexterous skills. The impressive results validate the effectiveness of unified multimodal modeling for generalist embodied intelligence. Furthermore, EO-Robotics is released with full openness, including model weights, training code, and all components of the interleaved embodied dataset, to serve as a forward step for general-purpose autonomous robots and facilitate further research in the community. This paper highlights the following contributions:

- **Unified Architecture:** EO-1 is a unified model that integrates multimodal embodied reasoning and real robot control within a shared backbone, enabling seamless cross-modal interaction without introducing extra bottleneck components or action-specific parameters.
- **Interleaved Embodied Dataset:** EO-Data1.5M is a comprehensive dataset derived from the largest robotic datasets, featuring interleaved embodied reasoning and robot control through a scalable data construction pipeline, thereby enabling interleaved vision-text-action pretraining.
- **Real-world Generalization:** Our model surpasses existing open-source models across multiple embodied reasoning and robot control benchmarks, including ERQA, LIBERO, SimplerEnv, and the self-constructed EO-Bench, meanwhile, extensive real-robot evaluations demonstrate its substantially stronger reasoning capabilities and dexterous control in open-world generalization.

2. The EO-1 Model and Training Paradigm

In this section, we introduce the architecture of EO-1 and its training strategies. As illustrated in Figure 1, the main architectural idea of EO-1 is to capture the inherent temporal dynamics and causal relationships between vision-text-action modalities in embodied interactions. This is realized by modeling multimodal understanding and robot control with a single unified decoder-only transformer, enabling effective cross-modal knowledge transfer and alignment for generalist policies.

2.1. Problem Formulation

We formulate the generalist robot policy π_θ as a unified model EO-1 that integrates *Multimodal Embodied Reasoning* and *Robot Control* through a synergized autoregressive and denoising paradigm. The EO-1 architecture can flexibly decode both continuous action chunks and tokenized text outputs, enabling seamless text-based reasoning and physical robot control within one model. The distribution represented by the model can be written as $\pi_\theta(\hat{\mathbf{a}}_{t:t+h}, \hat{\ell}_t | \mathbf{o}_t, \ell_t, \mathbf{a}_{t-h:t})$, where $\mathbf{o}_t = [\mathbf{I}_t^1, \dots, \mathbf{I}_t^n, \mathbf{q}_t]$ consists of multi-view image observations and robot state, language contexts ℓ are embodied reasoning question-answer text data with respect to current observation (e.g., “Q: Based on current image observation, what is the next step to do for the robot cleaning the table. A: pick up the yellow box and place it in the trash box”) or overall task prompts (e.g., “clear the table”), and action sequence $\mathbf{a}_{t:t+h}$ is an action chunk. The unified model is trained on an interleaved multimodal dataset \mathcal{D} to alternately generate text $\hat{\ell}_t$ and robot action $\hat{\mathbf{a}}_{t:t+h}$ for seamless embodied reasoning and action. Note that the unified model does not generate outputs on vision tokens and robot state tokens.

The generalist robot policy π_θ is instantiated by a transformer that takes interleaved multimodal token sequence $x_{1:N}$ as input and generates a sequence of multimodal outputs $\hat{x}_{1:N}$ in an autoregressive manner, i.e., $p(\hat{x}_{1:N}) = \prod_{i=1}^N p(\hat{x}_i | x_{<i})$. For inputs, each token can be a discretized text token x_i^w , a continuous image patch token x_i^I , a continuous robot state token x_i^q , or a partially denoised robot action token x_i^a . For output, discrete text tokens are sampled via a language-modeling head, while continuous action tokens are sampled by a flow-matching head. The language modeling head is implemented with a classic logits head appended to the unified transformer to decode text token output \hat{x}_i^w . For the flow-matching head, rectified flow (Lipman et al., 2022) is utilized to generate

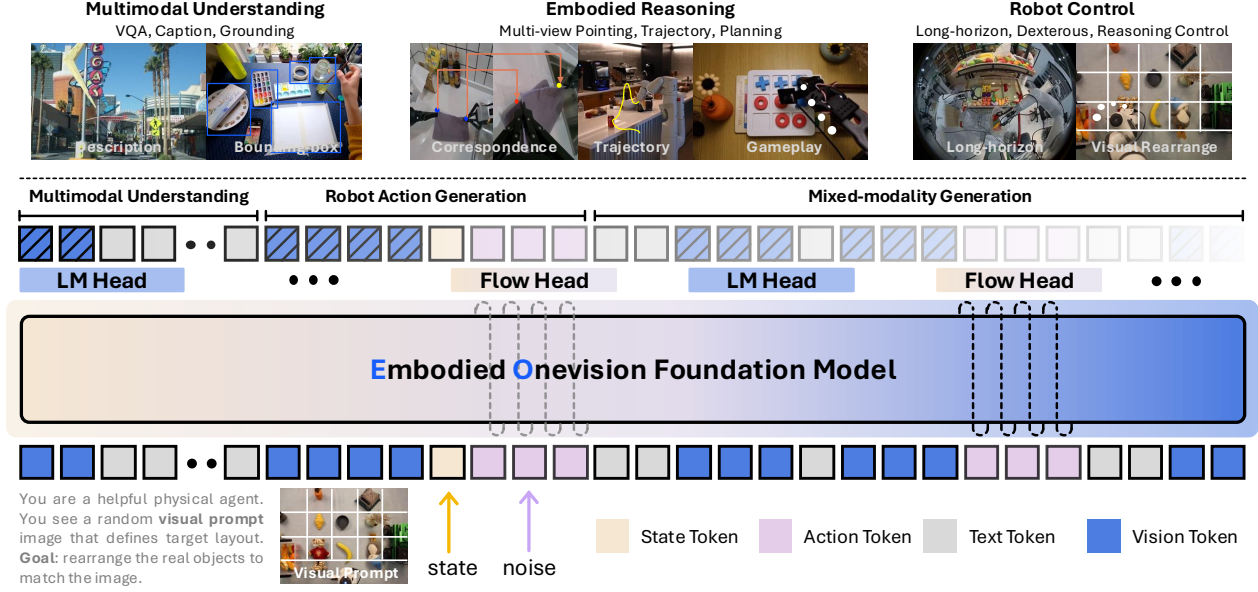


Figure 1: EO-1 Model Architecture. EO-1 model is a Vision-Language-Action (VLA) model that adopts a single unified decoder-only transformer, equipping with discrete language-modeling head for multimodal embodied reasoning and continuous flow-matching head for robot action generation. The language instruction, image observations, robot state, and noisy action are encoded into an interleaved token sequence of tokens to be processed by the shared transformer backbone, whose weights are initialized from Qwen2.5-VL. The model is trained on interleaved vision-text-action data with a combination of flow-matching objective and next-token-prediction objective and capable of seamless embodied reasoning and acting.

continuous action signals $\hat{\mathbf{a}}_t$ according to the forward Euler integration rule:

$$\hat{\mathbf{a}}_t^{\tau+\delta} = \hat{\mathbf{a}}_t^\tau + \delta V_{\pi_\theta}(\hat{\mathbf{a}}_t^\tau, |\mathbf{o}_t, \ell_t, \mathbf{a}_{t-h:t}), \quad (1)$$

where V_{π_θ} represents the parameters of π_θ to predict the vector field for robot action generation. The actions are generated by integrating the predicted vector field from $\tau = 0$ to $\tau = 1$, starting with random noise $z_t^0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and δ is the integration step size. The shared parameters in the unified transformer enable the seamless transfer of semantic knowledge from vision-language understanding to action generation, leading to more general reasoning capabilities and control capabilities in embodied scenarios (Deng et al., 2025, Black et al., 2025, Driess et al., 2025).

2.2. Model and Training

Model Architecture. The proposed unified model architecture processes interleaved multimodal inputs (*i.e.*, text, images, videos, and actions) through a shared transformer backbone that generates both discrete token sequences and continuous action tokens. The model employs a text tokenizer and visual encoder to convert text and image patches into input tokens, while the robot state \mathbf{q}_t is linearly projected into the same transformer embedding space, *i.e.*, $x_i^w, x_i^l, x_i^q \in \mathbb{R}^d$. Note that the text tokenizer and visual encoder are inherited from pre-trained VLM and the state projector is random initialized. For flow matching action denoising, the input “noisy action” \mathbf{a}_t^τ is obtained by $\mathbf{a}_t^\tau = \tau \mathbf{a}_t + (1 - \tau)z_\tau$, where $z_\tau \sim \mathcal{N}(0, I)$ is random noise. Then, another noisy action linear projector is utilized to embed noisy action and flow matching timestep τ into noisy action token $x_i^a \in \mathbb{R}^d$. The unified transformer backbone, initialized from Qwen 2.5 VL (Bai et al., 2025), processes this interleaved multimodal input sequence with shared parameters and generates output sequence through two separate heads: a language head for text token decoding and a flow head for continuous

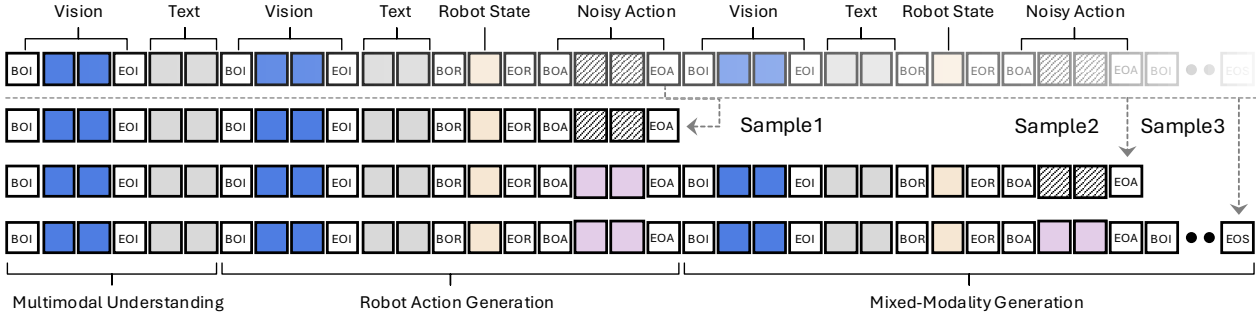


Figure 2: Interleaved rectifying sampling strategy. Our method samples variable-length subsequences from robot action generation segments, enabling efficient training of mixed-modality generation while preserving causal relationships.

action denoising. Note that we do not introduce extra action-specific parameters to separately model action denoising as prior VLA models, enabling seamless integration of discrete multimodal embodied reasoning and continuous robot control.

Unified Prompting and Attention. Training **EO-1** involves three types of data: multimodal understanding data, robot action generation data, and mixed-modality generation data (Section 3). The mixed-modality generation data is formatted as interleaved vision-text-action data, *i.e.*, “[BOI] {image patches} [EOI] [BOS] {text} [EOS] [BOR] {robot state} [EOR] [BOA] {noisy action} [EOA] [BOI] {image patches} [EOI] [BOS] {text} [EOS] [BOR] ...”, where [BOS], [EOS], [BOI], [EOI], [BOR], [EOR], [BOA], and [EOA] denote the beginning and end of sentence, image, robot state, action, respectively. The multimodal understanding format is “[BOI] {image patches} [EOI] [BOS] {text} [EOS]” and the robot control data format is “[BOI] {image patches} [EOI] [BOS] {text} [EOS] [BOR] {robot state} [EOR] [BOA] {noisy action} [EOA]”. During training, we employ an omnidirectional attention mask mechanism to process these three type data with causal attention masks and cache key-value (KV) pairs of the generated multimodal context to accelerate mixed-modality generation at inference. Note that prior VLA models are either trained with solo robot control data or co-trained with multimodal understanding and robot control data, missing interleaved vision-text-action data.

Interleaved Rectifying Sampling. When training on interleaved vision-text-action data for mixed-modality generation, the denoising process of action generation in interleaved data can disrupt the causal relationships in multimodal token sequences, because the subsequent text, image, or action tokens should attend to the clean action tokens and preceding text/image tokens, rather than noisy action tokens. To address this challenge, we propose a rectifying sampling strategy for mixed-modality generation training. As illustrated in Figure 2, we sample $N + 1$ (*e.g.*, $N = 2$) training subsequences from an interleaved sequence containing N action generation segments by splitting the sequence with respect to continuous action generation segments. For the subsequences ending with the first action generation segment, there is no additional operation on it. For the subsequences ending with the interleaved action generation segment, *i.e.*, containing another action generation segment in the middle of the sequence, we replace the noisy action tokens in the middle action generation segment with clean action tokens. With this interleaved rectifying sampling, each action generation segment is trained with both flow-matching denoising on noisy action tokens and indirect gradients backpropagation on clear action tokens for subsequent interleaved generation.

Training Objective. To learn both auto-regressive decoding and flow-matching denoising with one unified model, we employ two learning objectives: i) next token prediction and ii) denoising vector field prediction. For multimodal understanding and embodied reasoning tasks, the output token \hat{x}_i contains only text tokens \hat{x}_i^{tw} and is obtained through original language decoding head. The

language head and transformer backbone are trained on decoded text tokens using the cross-entropy loss between ground-truth text tokens and predicted logits $\mathcal{L}_{ar}(x_{gt}^w, \hat{x}_i^w)$, where $\hat{x}_i^w = \pi_\theta(\hat{x}|x_{<i})$ is conditioning on all preceding multimodal tokens. For robot action generation, the output action $\hat{\mathbf{a}}_t$ is generated by mapping intermediate hidden features of the last transformer layer to denoise vector field $z_0 - \mathbf{a}_t$ through a separate flow head. The flow head and transformer backbone are trained on action tokens by minimizing the following loss:

$$\mathcal{L}_{fm}(\theta) = \mathbb{E}_\tau[\|V_\theta(\mathbf{a}_t^\tau, \tau|x_{<a}) - (z_0 - \mathbf{a}_t)\|^2], \quad (2)$$

where $\mathbf{a}_t^\tau = \tau\mathbf{a}_t + (1 - \tau)z_0$ is the input noisy action. $x_{<a}$ represents the all preceding multimodal tokens. As in (Black et al., 2024), we sample the flow matching timestep τ from a beta distribution that emphasizes lower (noisier) timesteps. With the interleaved multimodal token sequence, we apply next token prediction $\mathcal{L}_{ar}(\theta)$ to the language head and flow matching $\mathcal{L}_{fm}(\theta)$ to the flow head for predicting denoising vector field. The unified model is trained end-to-end by optimizing sum of these two objectives: $\mathcal{L} = \mathcal{L}_{ar}(\theta) + \mathcal{L}_{fm}(\theta)$.

3. Dataset and Benchmark

EO-1 is trained on a diverse range of datasets across multiple modalities, including text, image, video, and robot control data, to perform embodied reasoning and dextrous control, all through a unified multimodal interface. In addition to standard robot control datasets and existing large-scale vision-language datasets, we design a scalable data construction pipeline and build interleaved embodied datasets from large-scale robot control episodes to capture the rich temporal dynamics and causal relationships inherent in embodied interactions. As illustrated in Table 1, The pre-training data corpus is structured into three main categories: web multimodal data, robot control data, and interleaved embodied data. We summarize the scale and composition of our training data across different modalities. In the following sections, we detail our data sources, interleaved data construction pipeline, and show some samples used in training.

Modality	Source	#Data	#Tokens
Web Multimodal Data	LLaVA-1.5 (Liu et al., 2024a), LLaVA-Video-178K (Zhang et al., 2024b), PixMo-Points (Deitke et al., 2024a), RefCOCO (Kazemzadeh et al., 2014a), RoboVQA (Sermanet et al., 2024)	5.7M	7.1B
Robot Control Data	AgiBotWorld (Bu et al., 2025), Open X-Embodiment (OpenX-Embodiment et al., 2024), RoboMIND (Wu et al., 2024), SO100-Community (Shukor et al., 2025), IPEC-Franka (Qu et al., 2025)	1.2M (episode)	127.3B
Interleaved Embodied Data	EO-Data1.5M from AgiBotWorld, Open X-Embodiment and RoboMIND	1.5M	1.0B

Table 1: Overview of the data used in EO-1 training. More details in Section B.

3.1. Data Source

Web Multimodal Data. The web text-image paired data is essential for multimodal understanding. We curated a web multimodal dataset comprising *5.7 Million samples with 7.1 Billion tokens*, integrating six public sources across three task domains: i) **Visual instruction following.** We incorporate the *LLaVA Visual Instruct Series* (Liu et al., 2023b, Sermanet et al., 2024) datasets, which contain GPT-generated multimodal instructions and robotics-oriented visual question answering examples. These sources cover diverse interaction patterns, including image-grounded conversation, multi-step reasoning, and long-horizon task planning. ii) **Video understanding and reasoning.** We leverage the *LLaVA-Video* (Zhang et al., 2024a) and *RoboVQA* (Sermanet et al., 2024), which support temporal multimodal comprehension through detailed video captioning, open-ended QA, and multiple-choice

question answering. iii) **Referring expression and spatial grounding.** To support fine-grained object understanding, we include *RefCOCO (subset)* (Kazemzadeh et al., 2014b) for natural language referring expression comprehension and *PixMo-Points* (Deitke et al., 2024b) for spatial localization via human-annotated coordinate grounding.

Robot Control Data. Robot control data provides critical supervision for policy learning across embodiment generalization and fine-grained manipulation. We aggregate a large-scale real robot control dataset comprising *1.2 Million episodes with 0.13 Trillion tokens* from five public sources, forming one of the most comprehensive collections for robot learning to date. The dataset includes *AgiBot-World* (Bu et al., 2025), featuring dual-arm humanoid demonstrations across diverse bimanual manipulation tasks with rich visuomotor sequences; *Open X-Embodiment* (OpenX-Embodiment et al., 2024), which spans 22 robot types and 527 skills to support cross-platform generalization; and *SO100-Community* (Shukor et al., 2025), a collection of long-horizon community demonstrations using SO100 platform. It also incorporates *RoboMind* (Wu et al., 2024), covering a wide range of tasks across multiple embodiments (e.g., Panda, UR5e, AgileX, humanoids) with diverse object classes.

3.2. Interleaved Embodied Data Construction

Our interleaved embodied data is obtained by annotating existing real robot control data on two aspects, resulting in two components: i) embodied temporal and spatial reasoning data focusing on physical dynamic and spatial relationship understanding of robot execution video (as shown in Figure 3) and ii) interleaved vision-text-action data connecting temporal/spatial reasoning data with robot control data for learning multimodal causal relationships in embodied interactions.


3.2.1. Embodied Reasoning Data Pipeline

To enhance the model’s vision-language capabilities on embodied intelligence, we develop a specialized pipeline in Figure 3, to carefully curate robot-specific question-answer data for both temporal and spatial embodied reasoning. Real robot manipulation videos serve as our primary data source. They perform a variety of tasks in real-world environments and showcase multiple embodiments. To ensure our curated dataset is rich in terms of diversity, quality, embodiments, and scenarios, the curation pipeline are presented as follows:

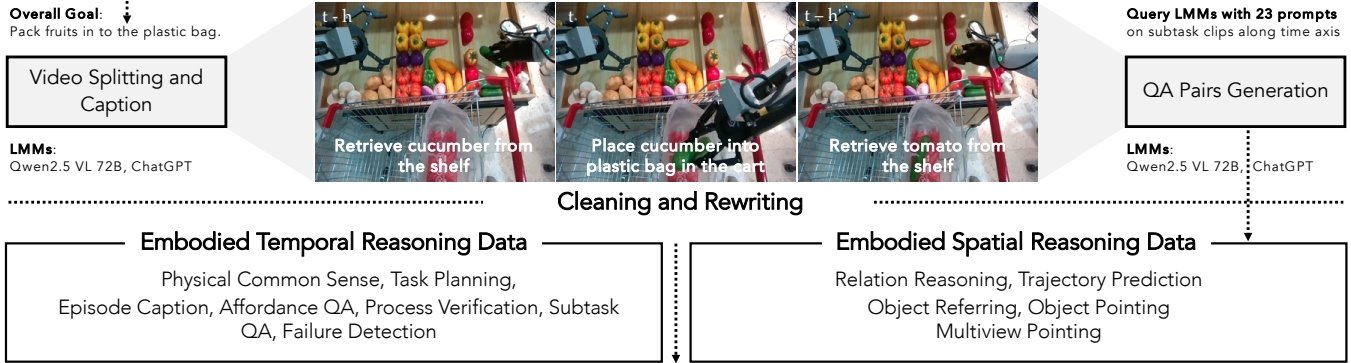
1. **Robot Video Filter and Curation.** As illustrated in Figure 3 and Section D, we first filter and sample a set of robot videos based on visual similarity to improve the data source diversity. This is because existing robot control datasets comprise thousands of videos collected in several fixed environments performing limited tasks, resulting in high visual similarity between videos. We extract visual features with a pretrained vision backbone and cluster them according to feature similarity. The diverse video set is curated by sampling a fixed number of videos from each cluster.
2. **Video Splitting and Captioning.** We employ both human annotators and pretrained VLMs to split videos into short clips containing individual subtasks. Each clip is then processed to extract detailed descriptions of the robot’s actions. These captions are not only used to construct video captioning QA data, but also serve as prompts for subsequent embodied reasoning annotations.
3. **Generating QA Pairs.** We prompt VLM to construct free-form questions according to designed templates, using subtask clips with previously generated captions. We construct two kinds of questions: i) “temporal reasoning” questions that focus on task planning and physical common sense understanding, and ii) “spatial reasoning” questions that require spatial understanding and object referring. As for answers, we first adopt VLM to generate multiple answers and employ human annotators to select or rewrite the correct answers.

EO-1: Interleaved Vision-Text-Action Pretraining for General Robot Control

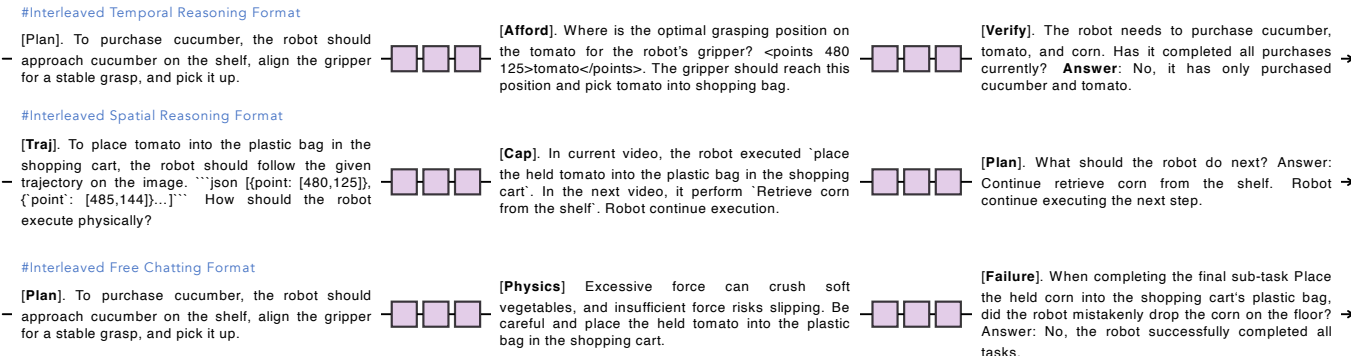

(a) Embodied Onevision Dataset (EO-Data1.5M) and Benchmark Statistics



Robot Video Filter and Curation



(b) Embodied Data Construction Pipeline on Robot Episodes



(c) Interleaved Vision-Text-Action Data Example

Figure 3: (a) Statistics of EO-Robotics Dataset (EO-Data1.5M) and Benchmark (EO-Bench). (b) Dataset curation pipeline on robot episodes. (c) Interleaved Vision-Text-Action Data Example, including three interleaved formats to concatenate embodied reasoning QA pairs and robot control data along temporal order.

4. **Cleaning and Rewriting.** Finally, we apply rule-based cleaning to ensure valid answers in spatial reasoning, and prompt an LLM to rewrite question-answer pairs to improve text diversity and inject fuzzy semantics into data.

Using the aforementioned pipeline, we curated embodied temporal and spatial reasoning datasets specializing on temporal dynamic understanding and spatial relationship reasoning. The annotation details and dataset statistics are presented as below: **Embodied Temporal Reasoning Data.** The temporal reasoning data aims to equip the pre-trained vision-language model with capabilities on physical common sense understanding and task planning, specializing on physical commonsense reasoning and task planning. As shown in Figure 3 (a), the **Physical Commonsense Understanding** data contain 70 thousand QA pairs for describing the effect of various robot actions in the physical world. The task planning data comprises 0.7 million QA pairs across four sub-categories to develop the model’s capacities including i) **Task Planning:** generating subtask sequences to complete long-horizon tasks. ii) **Episode Captioning:** captioning of robot video clips. iii) **Affordance QA:** assess whether executing a specific action is possible. iv) **Process Verification:** recognizing completed actions in video clips. v) **Subtask QA:** determining whether a subtask has been successfully completed. vi) **Failure Detection:** identifying unsuccessful subtask executions. Section D.3.1 shows question-construction prompt templates used to produce these questions and answers.

Embodied Spatial Reasoning Data This kind of data focus on enhancing the model’s abilities in spatial understanding and reasoning, which contains 1.5 million QA pairs and is organized into five sub-categories: relation reasoning, trajectory prediction, object referring, object pointing, and multiview pointing, as illustrated in Figure 3 (a). Specifically, i) **Relation Reasoning** data are annotated to understand relative spatial relationships among multiple objects in the scene. ii) **Trajectory Prediction** focuses on anticipating the future motion trajectories of objects or robot gripper within dynamic environments. iii) **Object Referring** data are constructed by grounding referred objects with bounding boxes among multiple candidates. iv) **Object Pointing** data aim to identify and point specific objects in multi-object scenarios according to the task instruction. v) **Multiview Pointing** is labeled by locating the same objects across different view images of the same scene. The question-construction prompt for embodied spatial reasoning is shown in Section D.3.2.

3.2.2. Interleaved Vision-Text-Action Data Pipeline

To learn the natural sequential coherence among visual, text, and action modalities, we design three flexible formats to concatenate embodied reasoning QA data and robot control data along temporal order in robot videos. They are illustrated in Figure 3 (b), and detailed as follows:

- **Interleaved Temporal Reasoning Format.** For a certain frame from the robot video, the interleaved temporal reasoning data is constructed as. “... [image tokens] [next subtask planning QA] [subtask instruction] [robot action] [image tokens] [task-completion verification QA] ...”. We use a predefined template merge next subtask planning answer into the instruction for action generation, and a task-completion verification QA is appended to new image tokens to determine the task progress.
- **Interleaved Spatial Reasoning Format.** As for spatial reasoning data, we use the trajectory prediction QA and introduce a “[trajectory instruction]” to connect spatial reasoning QA with robot control data, *i.e.*, “... [image tokens] [robot trajectory prediction QA] [trajectory instruction] [robot action] [image tokens] [task-completion verification QA] ...”. The robot trajectory prediction’s answers contain a sequence of [x,y] coordinates of robot gripper trajectory for finish task, and its results are merged into “[trajectory instruction]” for next action generation, *i.e.*, “How should the robot execute the trajectory [x1,y1],..., [x6,y6] physically?”
- **Interleaved Free Chatting Format.** For other temporal and spatial reasoning QA data, we randomly

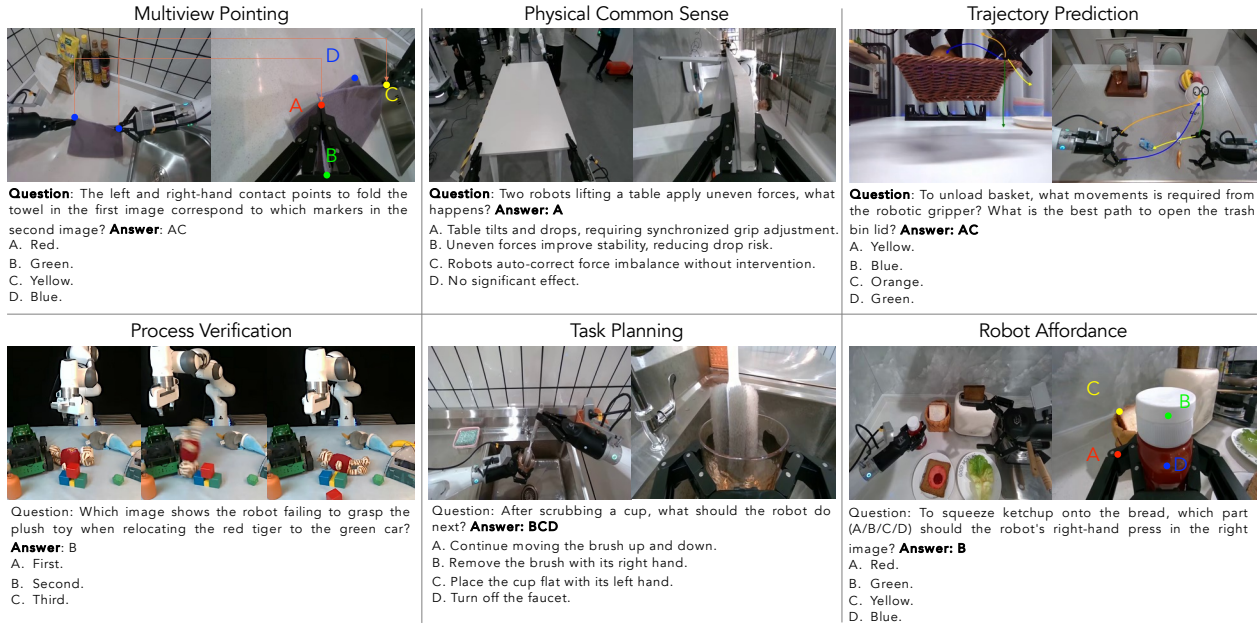


Figure 4: Examples of EO-Bench, including Multiview Pointing, Physical Common Sense, Trajectory Prediction, Process Verification, Task Planning, and Robot Affordance.

select QA pairs and connect them with robot control data through original task instruction: “... [image tokens] [random reasoning QA] [task instruction] [robot action] [image tokens] [task-completion verification QA] ...”. The task instruction is the prompt used in existing VLA models, *i.e.*, “What should the robot do to finish the task {task instruction}”, where {task instruction} is the original task label from robot control data.

With these flexible interleaved formats, we sample the first, last, one random middle image from each segmented subtask video clips and annotate next subtask planning to construct interleaved temporal reasoning data. The interleaved spatial reasoning data is constructed using all robot trajectory prediction QA in embodied spatial reasoning data. The interleaved free chatting format randomly remaining QA data to connect embodied reasoning QA data and robot control data. Totally, we curated 122 thousand interleaved vision-text-action data for mixed-modality generation training.

3.3. Embodied Onevision Reasoning Benchmark

Embodied Onevision Benchmark (EO-Bench) aims to construct a comprehensive and balanced evaluation suite for open-world embodied reasoning, covering both challenging and accessible tasks. Existing benchmarks often suffer from the drawback that a single QA instance may conflate multiple reasoning aspects, such as combining spatial trajectories with extensive commonsense knowledge. This design leads to ambiguous evaluations, making it difficult to attribute performance to specific capabilities. In contrast, our benchmark ensures that each question targets a clearly defined reasoning angle rather than mixing multiple dimensions. This enables a more faithful diagnosis of models’ strengths and weaknesses in an interpretable and disentangled manner.

As shown in Figure 3(a), EO-Bench is organized into four major categories: spatial understanding, physical commonsense, task reasoning, and state estimation. In total, the benchmark comprises 648 QA pairs manually labeled on diverse robot control data, distributed across categories as follows: 370 for spatial understanding, 140 for task reasoning, 84 for physical dynamic reasoning, and 48 for physical commonsense. This distribution reflects our intention to emphasize spatial reasoning, which

is an essential component and the bottleneck of embodied intelligence. Together, these four categories provide a structured yet broad evaluation of a model’s ability to reason about space, physics, tasks, and states in embodied intelligence.

Each major category is further divided into sub-categories, as illustrated in Figures 3 and 4.

- **Spatial Understanding** covers trajectory reasoning (predicting object motion), visual grounding (localizing objects in context), relation reasoning (understanding relative positions among objects), and multiview pointing (identifying objects consistently across multiple viewpoints).
- **Physical Commonsense** includes direct influence and counterfactual reasoning, both of which require models to apply fundamental physical principles in embodied interactions.
- **Task Reasoning** spans task planning (deriving action sequences toward goals), episode captioning (summarizing task-level events), and process verification (evaluating task completion and progress).
- **State Estimation** involves recognizing object states (e.g., open/closed, full/empty), reasoning about robot-object interactions, and assessing action-level consequences.

By structuring the benchmark into these categories and sub-categories, we provide a principled way to evaluate embodied models across complementary aspects of embodied intelligence, while maintaining interpretability of each evaluation dimension.

4. Experimental Results

4.1. Implementation Details

The EO-1 model is trained on a large-scale corpus that integrates 1.2M real-robot demonstrations from AgiBotWorld (Bu et al., 2025), Open X-Embodiment (OpenX-Embodiment et al., 2024), RoboMIND (Wu et al., 2024), and SO100-Community (Shukor et al., 2025), together with 5.7M web multimodal samples and 1.5M interleaved embodied data pairs, yielding a total of 135B tokens. To optimize on this scale, we train for five epochs with Flash-Attention variable-length packing (average sequence length of 16,384) and a batch size of 1. The backbone learning rates are set to 5×10^{-5} for both the language model and multimodal projector, 1×10^{-5} for the vision encoder, and a fixed balance factor of 1.0 is applied between language regression and action flow matching losses. Additional settings include an action denoising chunk size of 16, resolution bounds of 50176–100352, and the DeepSpeed ZeRO-1 optimizer. At inference time, the policy is conditioned on multi-view camera observations and sub-task instructions. It predicts a 16-step action chunk through 10 denoising iterations for robot execution. This design achieves a balance between stability and efficiency, while requiring only 6 GB of GPU memory on a single NVIDIA RTX 4090, thus enabling real-time evaluation in both simulation and real-world environments.

4.2. Open-sourced Benchmark Evaluation Setups

To thoroughly evaluate the embodied reasoning and dexterous manipulation capabilities, we assess our model across two key categories of benchmarks: *Embodied Reasoning* and *Robot Control*. They encompass a broad spectrum of perceptual, cognitive, and control challenges, spanning from high-level visual reasoning to low-level motor skill execution, across both simulated and visually complex real-world environments.

Embodied Reasoning. This category consists of three reasoning-centric benchmarks: **RoboVQA**, **ERQA**, and our self-constructed **EO-Bench**. They are described in the following items. Specifically, these benchmarks are designed to evaluate distinct and complementary cognitive capabilities. To ensure standardized and consistent evaluation, we employ VLMEvalKit, a unified framework that

facilitates rigorous benchmarking and reliable performance comparison across these tasks.

- **RoboVQA** (Sermanet et al., 2024): A free-form visual question answering (VQA) benchmark designed to assess high-level reasoning in egocentric embodied scenes. Following prior works, model performance is evaluated using BLEU-4 scores.
- **ERQA** (Gemini-Robotics, 2025): A curated benchmark that emphasizes spatial reasoning and grounded world knowledge. The tasks are designed around realistic robotic scenarios requiring inference over object relationships, geometry, and agent-centric spatial cues.
- **EO-Bench**: We propose **EO-Bench** to evaluate models in *spatial-temporal reasoning*. The benchmark consists of 700 multiple-choice VQA tasks, distributed evenly across four reasoning categories: **Physical Commonsense**, **Spatial Understanding**, **State Estimation**, and **Task Reasoning**. All tasks are grounded on egocentric visual inputs, meticulously constructed to reflect the real-world challenges of robot perception and decision-making.

Robot Control. This category includes two manipulation benchmarks: **SimplerEnv** and **LIBERO**, which collectively assess a wide spectrum of control capabilities across visually diverse, long-horizon, and goal-conditioned tasks.

- **SimplerEnv** (Li et al., 2024b): A manipulation benchmark featuring WidowX and Google Robots operating in visually diverse environments with variations in lighting, surface textures, and camera viewpoints. It leverages the Bridge Data V2 (Walke et al., 2023) for evaluating Real-to-Sim transfer. SimplerEnv includes both short-horizon atomic tasks and visually challenging variants, specifically designed to assess robustness under appearance shifts between real and simulated settings.
- **LIBERO** (Liu et al., 2023a): A long-horizon manipulation benchmark consisting of temporally extended, multi-stage tasks across various object categories and interaction types. LIBERO emphasizes planning, memory, and semantic understanding in complex simulated environments.

4.3. Benchmark Evaluation Results

Embodied Reasoning. To evaluate the embodied reasoning capabilities, we benchmark it on three key tasks: **RoboVQA** (Sermanet et al., 2024), **ERQA** (Gemini-Robotics, 2025), and **EO-Bench**. These benchmarks assess a diverse range of reasoning skills, including spatial understanding, physical commonsense, and multistep task planning. We compare **EO-1** against three categories of baselines: *public VLMs* (e.g., Qwen2.5 VL (Bai et al., 2025), InternVL2.5 (Chen et al., 2024b)), *private VLMs* (e.g., GPT-4o, Gemini 1.5 Flash (Gemini et al., 2023b), Claude 3.5), and *co-trained visual-language-action models* (e.g., ChatVLA (Zhou et al., 2025), Magma (Yang et al., 2025)), with all models evaluated under consistent embodied conditions. Across all the benchmarks presented in Table 2, **EO-1** consistently demonstrates robust reasoning capabilities in embodied scenarios in the open world. On **RoboVQA**, which evaluates long-horizon visuospatial inference in egocentric settings, **EO-1** achieves the BLEU-4 score of 58.5, significantly outperforming leading closed-source models, such as GPT-4o (47.2). For **ERQA**, a spatially grounded benchmark derived from real-world visual scenes, **EO-1** achieves the accuracy of 45.5, surpassing InternVL2.5 8B (45.2) and outperforming other public and private VLMs. In the suite **EO-Bench** introduced in this paper, mainstream VLMs achieve a modest average score of 32, highlighting the substantial limitations of current methods in handling embodied tasks. On the **EO-Bench @ Spatial** leaderboard, which evaluates performance across multiview reasoning, trajectory prediction, and visual grounding tasks, **EO-1** outperforms other open-source models of comparable size, securing a score of 36.4. This highlights **EO-1**'s improved capability in multi-modal scene understanding and spatial reasoning. Furthermore, in the **EO-Bench @ Temporal** subset, dedicated to multimodal and temporal reasoning in manipulation tasks, **EO-1** achieves 38.9, underscoring its strong aptitude for scene-level abstraction and temporal prediction.

Multi-modal Benchmark	RoboVQA	ERQA	EO-Bench @ Spatial	EO-Bench @ Temporal	Overall
Gemini 1.5 Flash (Gemini et al., 2023a)	46.0	46.3	46.3	37.8	44.1
Claude 3.5	26.7	35.5	24.0	34.8	30.3
GPT-4o 2024-11-20	47.2	40.0	35.6	39.3	40.5
Qwen2.5 VL 3B (Bai et al., 2025)	55.9	35.3	20.0	22.6	33.5
Qwen2.5 VL 7B (Bai et al., 2025)	56.9	39.3	26.7	31.5	38.6
InternVL2.5 4B (Chen et al., 2024a)	45.7	36.8	37.0	31.9	37.9
InternVL2.5 8B (Chen et al., 2024a)	43.9	45.2	32.8	38.1	40.0
ChatVLA 2B (Zhou et al., 2025)	33.7	34.3	26.4	21.9	29.1
Magma 8B (Yang et al., 2025)	30.3	29.3	29.4	36.7	31.4
EO-1 (3B)	58.5	45.5	36.4	38.9	44.8

Table 2: Performance comparison with standard VLMs and co-trained visual-language-action models on **RoboVQA** (Sermanet et al., 2024), **ERQA** (Gemini-Robotics, 2025), and **EO-Bench** benchmarks.

Model	LIBERO-Spatial		LIBERO-Object		LIBERO-Goal		LIBERO-Long		Overall	
	SR	Rank	SR	Rank	SR	Rank	SR	Rank	SR	Rank
OpenVLA (Kim et al., 2024)	84.7 ± 0.9%	5	88.4 ± 0.8%	5	79.2 ± 1.0%	4	53.7 ± 1.3%	5	76.5 ± 0.6%	5
SpatialVLA (Qu et al., 2025)	88.2 ± 0.5%	2	89.9 ± 0.7%	4	78.6 ± 0.6%	5	55.5 ± 1.0%	4	78.1 ± 0.7%	4
OpenVLA-OFT (Kim et al., 2025)	97.6 ± 0.9%	1	98.4 ± 0.8%	1	97.9 ± 1.0%	1	94.5 ± 1.3%	1	97.1 ± 0.6%	1
ThinkAct (Huang et al., 2025)	88.3 ± -%	3	91.4 ± -%	3	87.1 ± -%	3	70.9 ± -%	3	84.4 ± -%	3
MolmoAct-7B-D (Lee et al., 2025)	87.0 ± -%	4	95.4 ± -%	2	87.6 ± -%	2	77.2 ± -%	2	86.6 ± -%	2
Diffusion Policy (Chi et al., 2023)	78.5 ± 1.1%	6	87.5 ± 0.7%	5	73.5 ± 1.2%	6	64.8 ± 1.3%	4	76.1 ± 0.7%	5
Octo (Octo et al., 2024)	78.9 ± 1.0%	5	85.7 ± 0.9%	6	84.6 ± 0.9%	5	51.1 ± 1.3%	6	75.1 ± 0.6%	6
π_0 (Black et al., 2024)	96.8 ± 0.8%	2	98.8 ± 0.9%	2	95.8 ± 1.1%	2	85.2 ± 1.2%	3	94.2 ± 0.9%	2
π_0 -FAST (Pertsch et al., 2025)	96.4 ± 0.7%	3	96.8 ± 0.7%	4	88.6 ± 1.0%	4	60.2 ± 1.4%	5	85.5 ± 1.0%	4
GR00T N1 (Bjorck et al., 2025)	94.4 ± 0.9%	4	97.6 ± 1.0%	3	93.0 ± 1.2%	3	90.6 ± 1.0%	2	93.9 ± 1.1%	3
EO-1 (Ours)	99.7 ± 0.2%	1	99.8 ± 0.1%	1	99.2 ± 0.3%	1	94.8 ± 0.6%	1	98.2 ± 0.3%	1

Table 3: Performance comparison with state-of-the-art policies on **LIBERO Benchmark** (Liu et al., 2023a).

Robot Control. We evaluated manipulation performance on two simulation benchmarks: **LIBERO** (Liu et al., 2023a) and **SimplerEnv** (Li et al., 2024b). LIBERO spans four subsets, Spatial, Object, Goal, and Long, designed to test generalization across spatial layouts, object categories, goal semantics, and long-horizon planning. SimplerEnv includes Google-VM (Visual Matching), Google-VA (Visual Aggregation), and WindowX setups with controlled variations in color, texture, lighting, and camera pose, assessing robustness under visual distribution shift. We compare against leading robot foundation models, including the **RT series** (Brohan et al., 2022, OpenX-Embodiment et al., 2024), **Octo** (Octo et al., 2024), **OpenVLA** (Kim et al., 2024), the π -series (Black et al., 2024, Pertsch et al., 2025), and **ThinkAct** (Huang et al., 2025), among others.

As shown in Table 3 and Table 4, our method achieves superior performance across various robotic manipulation tasks. On the **LIBERO** benchmark, it attains an average success rate of 98.2%, outperforming recent state-of-the-art models, including OpenVLA-OFT (Kim et al., 2025), π_0 (Black et al., 2024), and GR00t-N1 (Bjorck et al., 2025) by 1.1%, 4.0%, and 4.3%, respectively, demonstrating robust generalization across spatial, semantic, and long-horizon tasks. On the **SimplerEnv** benchmarks, it surpasses π_0 (Black et al., 2024) by 3.5%, 5.1%, and 8.3% in the WindowX, Google-VM, and Google-VA variants, respectively, achieving the highest success rates of 72.7%, 76.5%, and 63.0%. These results are obtained with lightweight fine-tuning on a modest mixed-modality dataset, underscoring the method’s data efficiency, dexterous control, and precise visuomotor grounding.

4.4. Real-world Experiment Evaluation Setups

We introduce a comprehensive set of robotic manipulation tasks to evaluate the generalist capabilities of the VLA model across a diverse array of real-world scenarios. These tasks span pick-and-place, articulated object manipulation, long-horizon and dexterous actions, while also challenging the robots’

WidowX Benchmark	Put Spoon on Towel	Put Carrot on Plate	Stack Blocks	Put Eggplant in Basket	Overall
RT-1-X (OpenX-Embodiment et al., 2024)	0%	4.2%	0%	0%	1.1%
OpenVLA (Kim et al., 2024)	0%	0%	0%	4.1%	1.0%
SpatialVLA (Qu et al., 2025)	20.8%	20.8%	25.0%	70.8%	34.4%
Magma (Yang et al., 2025)	37.5%	29.2%	20.8%	91.7%	44.8%
π_0 -FAST (Pertsch et al., 2025)	29.1%	21.9%	10.8%	66.6%	32.1%
Octo-Base (Octo et al., 2024)	12.5%	8.3%	0%	43.1%	16.0%
Octo-Small (Octo et al., 2024)	47.2%	9.7%	4.2%	56.9%	29.5%
RoboVLM (Li et al., 2024a)	45.8%	20.8%	4.2%	79.2%	37.5%
π_0 (Black et al., 2024)	83.8%	52.5%	52.5%	87.9%	69.2%
ThinkAct (Huang et al., 2025)	58.3%	37.5%	8.7%	70.8%	43.8%
EO-1 (Ours)	63.6%	54.5%	81.8%	90.9%	72.7%
Google Robot Benchmark (Matching)	Pick Coke Can	Move Near	Open-Close Drawer	Drawer Apple	Average
RT-1 (Brohan et al., 2022)	85.7%	44.2%	73.0%	6.5%	52.4%
OpenVLA (Kim et al., 2024)	16.3%	46.2%	35.6%	0%	24.5%
TraceVLA (Zheng et al., 2024)	28.0%	53.7%	57.0%	0%	34.7%
SpatialVLA (Qu et al., 2025)	86.0%	77.9%	57.4%	0%	55.3%
Magma (Yang et al., 2025)	75.0%	53.0%	58.9%	8.3%	48.8%
π_0 -FAST (Pertsch et al., 2025)	75.3%	67.5%	42.9%	0	46.4%
RT-1-X (OpenX-Embodiment et al., 2024)	56.7%	31.7%	59.7%	40.7%	47.2%
RT-2-X (OpenX-Embodiment et al., 2024)	78.7%	77.9%	25.0%	7.4%	47.3%
Octo-Base (Octo et al., 2024)	17.0%	4.2%	22.7%	0	11.0%
RoboVLM Li et al. (2024a)	77.3%	61.7%	43.5%	24.1%	51.7%
π_0 (Black et al., 2024)	97.9%	78.7%	62.25%	46.6%	71.4%
ThinkAct (Huang et al., 2025)	92.0%	72.4%	50.0%	-	-
MolmoAct (Lee et al., 2025)	77.7%	77.1%	60.0%	-	-
EO-1 (Ours)	98.0%	83.8%	71.3%	52.8%	76.5%
Google Robot Benchmark (Aggregation)	Pick Coke Can	Move Near	Open-Close Drawer	Drawer Apple	Average
RT-1 (Brohan et al., 2022)	89.8%	50.0%	32.3%	2.6%	43.7%
OpenVLA (Kim et al., 2024)	54.5%	47.7%	17.7%	0.0%	30.0%
TraceVLA (Zheng et al., 2024)	60.0%	56.4%	31.0%	0%	36.9%
SpatialVLA (Qu et al., 2025)	88.0%	72.7%	41.8%	6.3%	52.2%
Magma (Yang et al., 2025)	68.6%	78.5%	59.0%	24.0%	57.5%
π_0 -FAST (Pertsch et al., 2025)	77.6%	68.2%	31.3%	0	44.3%
RT-1-X (OpenX-Embodiment et al., 2024)	49.0%	32.3%	29.4%	10.1%	30.2%
RT-2-X (OpenX-Embodiment et al., 2024)	82.3%	79.2%	35.3%	20.6%	54.4%
Octo-Base (Octo et al., 2024)	0.6%	3.1%	1.1%	0	1.2%
RoboVLM Li et al. (2024a)	75.6%	60.0%	10.6%	0	36.6%
π_0 (Black et al., 2024)	90.1%	80.7%	27.6%	20.5%	54.7%
ThinkAct (Huang et al., 2025)	84.0%	63.8%	47.6%	-	-
MolmoAct	76.1%	61.3%	78.8%	-	-
EO-1 (Ours)	91.6%	81.7%	55.0%	23.8%	63.0%

Table 4: Performance comparison with state-of-the-art policies on **SimplerEnv Benchmark** (Li et al., 2024b).

abilities to perform embodied reasoning, engage in complex planning, and adapt to varied control types, as illustrated in Figure 5. These tasks are described as follows.

- 1. Franka Panda Pick-and-Place (7 Tasks).** Household tasks such as brewing tea, storing kitchen items, and sorting cubes, demonstrating proficiency in object-to-container pick-and-place and articulated object manipulation. These tasks test fine motor control and adaptability to diverse object geometries in basic household settings.
- 2. WidowX 250 S Out-of-Box (13 Tasks).** Out-of-box evaluation in a kitchen environment includes tasks involving both object-to-container manipulation (e.g., vegetable preparation, cup arrangement) and articulated handling (e.g., door closing), highlighting the robot’s flexibility in managing diverse objects and complex actions.
- 3. Agibot G-1 Long-Horizon Dexterity (4 Tasks).** Long-horizon, dexterous tasks requiring sequential planning and fine manipulation, such as folding clothes, making sandwiches, and sorting groceries. These tasks highlight the robot’s proficiency in multi-step operations in multi-step tasks, demonstrating advanced dexterity and adaptability in long-term execution scenarios.
- 4. Embodied Reasoning Control (4 Tasks).** Physical and abstract tasks including visual object rearrangement, tic-tac-toe, and long-horizon planning tasks (e.g., making breakfast sandwich, roasting beef steak), showcasing its embodied reasoning capabilities. These tasks require not only

EO-1: Interleaved Vision-Text-Action Pretraining for General Robot Control



Figure 5: Example of real-world evaluation tasks in diverse robots, including Agibot G-1 Long-horizon Dexterous (row 1-4), Franka Panda Pick-and-Place (row 5), WidowX 250 S Out-of-Box (row 6), and Embodied Reasoning Control in Franka, Agibot G-1, and Lerobot SO100 (row 1,2,7,8).

precise manipulation but also higher-level reasoning, testing the intersection of manipulation and reasoning in real-world contexts.

These 28 task settings serve to assess the model’s performance in four key dimensions: mastering diverse manipulations across multiple robot platforms, specializing in long-horizon dexterous tasks, emerging open-world embodied generalization, and enhanced generalization through unified reasoning capabilities. The results of the four dimensions are presented successively in the following subsections.

4.5. Mastering Diverse Manipulations on Multiple Embodiments




Figure 6: Performance comparison on diverse robot platforms and task categories.

In this section, we demonstrate that **EO-1** can perform a wide range of dexterous manipulation tasks across various robotic platforms, showcasing its robustness and adaptability. We evaluate its performance on both short-horizon and long-horizon tasks using different robots, including Franka Panda, WidowX 250 S, Agibot G-1, and Lerobot SO100. These tasks range from basic pick-and-place to complex articulated object manipulation and long-horizon tasks. We compare **EO-1** against state-of-the-art baselines, including π_0 -Fast, π_0 , and GR00T-N1.5 (Pertsch et al., 2025, Black et al., 2024, Bjorck et al., 2025), highlighting its superior performance across these varied tasks.

The results reported in Figure 6 show that **EO-1** consistently outperforms the baselines across all robot platforms and task categories, achieving a performance of 86.0%, compared with 43.0% for Fast, 71.0% for GR00T-N1.5, and 68.0% for π_0 . Notably, in long-horizon tasks with Agibot G-1, **EO-1** achieves a remarkable completion score of 81.0%, surpassing the performance of π_0 at 67.0%. Similarly, on Franka Panda’s pick-and-place tasks, **EO-1** scores 94.0%, exceeding GR00T-N1.5’s 86.0%. On more complex embodied reasoning control tasks, *e.g.*, Tic-Tac-Toe and Visual Rearrangement, **EO-1** maintains a completion score of 83.0%, which again surpasses that of π_0 at 53.0% and GR00T-N1.5 at 62.0%. These results indicate that **EO-1** excels at managing multi-step and dexterous tasks that demand high precision and adaptability across various environments and robot platforms. The strong performance of **EO-1** across such a diverse set of tasks underscores its potential for real-world deployment, where robotic systems are required to operate reliably in dynamic and varied environments.

4.6. Specializing to Long-horizon Dexterity

In this section, we investigate the **EO-1** model’s ability to specialize in long-horizon dexterous tasks, which require precise coordination and sustained execution over extended time periods. While the

model demonstrated strong performance on short-horizon tasks in previous evaluations, these more complex tasks present unique challenges, as they involve a sequence of coordinated actions across multiple steps. To explore the model’s specialization potential, we fine-tune **EO-1** using high-quality demonstration data for several real-world tasks that push the limits of both dexterous and long-horizon manipulation. As illustrated in Figure 5 (row 1-4), we select four tasks that require intricate multi-step decisions and fine manipulation across multiple sequential actions:

1. **Make Breakfast Sandwich:** Tasks involve the precise placement of bread, ham, lettuce, and other ingredients in a multi-step sequence using both arms to assemble the sandwich.
2. **Roast Beef Steak:** Tasks Involve a series of steps requiring fine motor skills, including brushing oil on the steak, placing it in the oven, and operating the oven with both arms.
3. **Fold Household Clothes:** Tasks involve grasping and folding shorts by coordinating both arms for precise control over the fabric, requiring dexterity and spatial awareness.
4. **Sort Grocery Items:** Tasks involve sort items (e.g., shampoo and green onions) into designated containers based on specific instructions, testing the robot’s instruction-following and object manipulation skills.

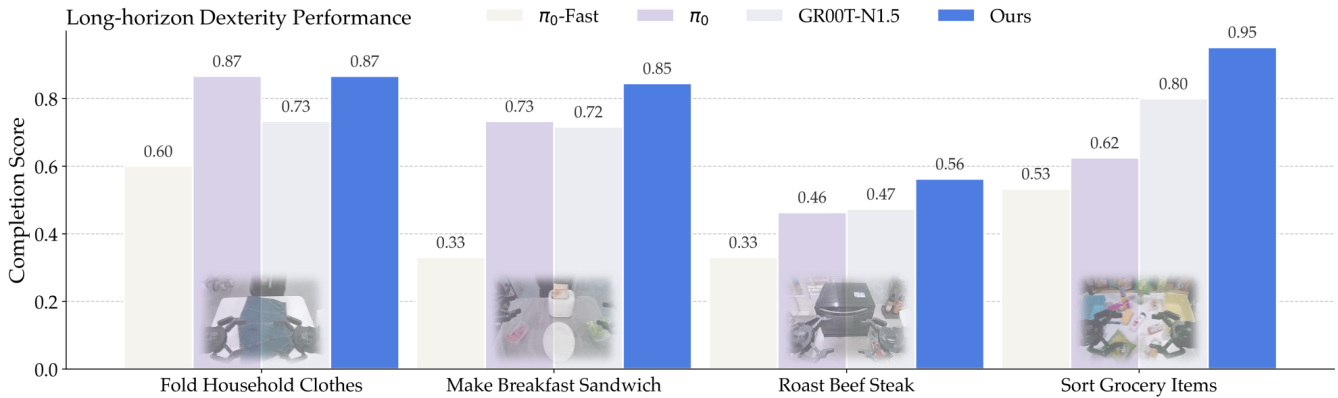


Figure 7: Long-horizon dexterity completion rate comparison on the AgiBot G-1 platform.

Specializing fine-tuning is performed using 150 demonstration trajectories per task, with each suite comprising approximately 0.5 million frames across 8 steps. The training process spans 25 epochs, with a 9:1 sub-task-to-overall task mixture ratio, where 90% of the data is from sub-tasks and 10% from the overall task. During evaluation, 10 test runs per task are conducted, using sub-task instructions to compute the average completion score.

As shown in Figure 7, we assess the specialized **EO-1** model on four long-horizon tasks, comparing its performance against leading methods. Specifically, on the Roast Beef Steak task, which requires fine motor control and interaction with kitchen appliances, **EO-1** reaches a completion score of 56.0%, substantially higher than both GR00T-N1.5 (47.0%) and π_0 (46.0%). On the Make Breakfast Sandwich task, **EO-1** demonstrates an 85.0% completion score, far surpassing π_0 (73.0%) and GR00T-N1.5 (72.0%), showcasing its ability to handle complex, multi-step manipulations. Similarly, in the Sorting Grocery Items task, which tests both object manipulation and instruction-following abilities, **EO-1** excels with a completion score of 95.0%, while π_0 and GR00T-N1.5 achieve notably lower success rates (62.0% and 80.0%, respectively). Overall, as shown in Figure 6(c), **EO-1** achieves an impressive average completion score of 81.0%, significantly outperforming both GR00T-N1.5 (68.0%) and π_0 (67.0%). These results highlight **EO-1**’ ability to excel in long-horizon tasks, showing significant improvements over baseline models, and its capacity to manage complex, multi-step actions in real-world scenarios. By seamlessly integrating regression and flow matching within a unified backbone, the model demonstrates the potential for more stable inferences and precise




Figure 8: Instruction-following in open-world settings. Left: Example scenes from three task types: **Instructed Visual Rearrangement** (Franka Panda), **Sort Grocery Items** (Agibot G-1), and **WidowX Out-of-Box**. Right: Success rates over 18 tasks.

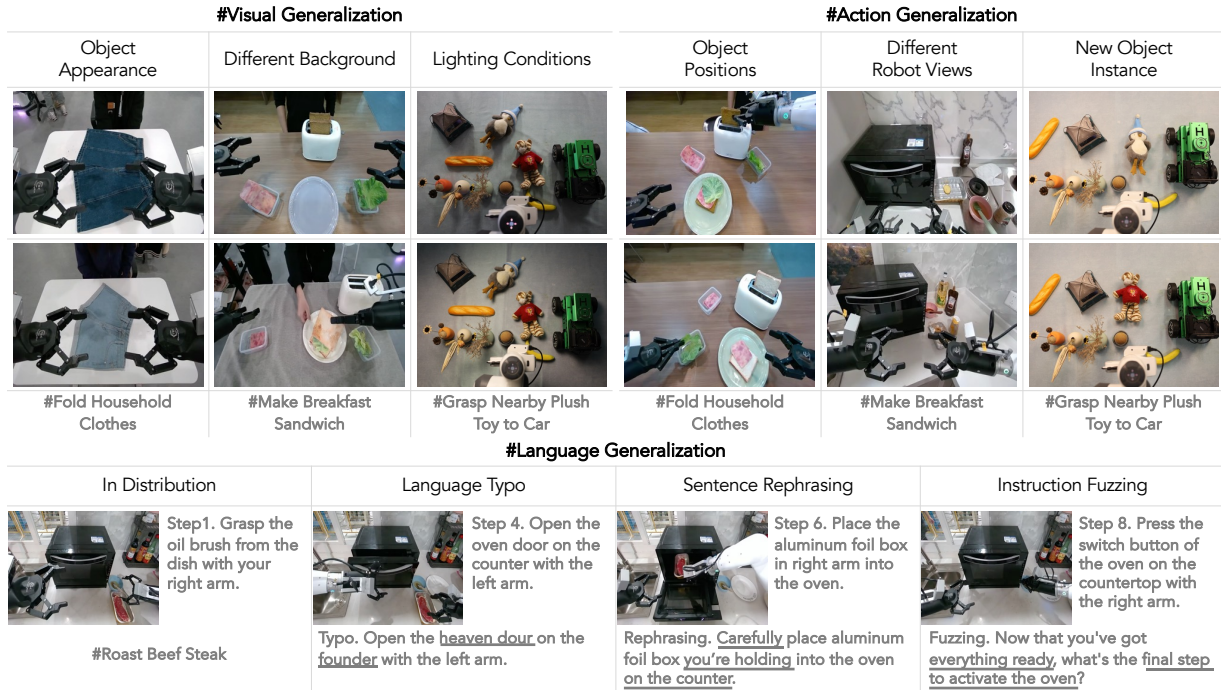
decision-making, especially in long-sequence tasks, without relying on action-specific parameters or introducing heuristic bottlenecks.

4.7. Emerging Open-world Embodied Generalization


The key challenge for embodied foundation models is generalizing to open-world scenarios where natural language instructions must be grounded into precise, executable actions. To evaluate this capability, we perform a generalization assessment using a suite of 15 instruction-following tasks across three different embodiments: **Franka Panda**, **Agibot G-1**, and **WidowX 250 S**, all operating in unstructured, object-rich environments. The tasks include **Visual Rearrangement (Instructed)**, which involves 7 instructions over 5 objects to test spatial instruction following and sequential manipulation; **Sort Grocery Items**, requiring fine-grained recognition and placement of 4 objects under cluttered conditions; and **WidowX Out-of-Box**, which consists of 13 multi-instruction tasks involving both rigid and articulated objects, assessing adaptability in compact workspaces.

Against strong baselines GR00T-N1.5 and π_0 (Pertsch et al., 2025, Black et al., 2024, Bjorck et al., 2025), **EO-1** achieves an average completion score of **87.0%**, outperforming GR00T-N1.5 (75.0%) and π_0 (66.0%). Performance gains are consistent across embodiments and task types, with **EO-1** maintaining high accuracy even under precise spatial references and unseen object-instruction combinations, demonstrating robustness to both language and visual distribution shifts. We observe that while π_0 executes tasks quickly and accurately, it struggles to follow instructions, often prioritizing objects closer to itself instead of adhering to the specified commands. This behavior may be due to the negative impact of pre-training on the entire backbone, which can lead to suboptimal decision-making in dynamic environments. In contrast, GR00T-N1.5, which freezes the language backbone during training, shows improved instruction-following capabilities.

As shown in Figure 9a, we further extend the evaluation to **12 tasks** designed to probe generalization along three orthogonal variation axes—**Visual Generalization**: invariance to changes in object appearance, background, and lighting, tested on WidowX Out-of-Box, Franka Panda Pick-and-Place, and Agibot G-1 household tasks (Fold Clothes, Make Sandwich, Roast Beef). **Action Generalization**: adaptation to altered object positions, new viewpoints, and unseen object instances with different physical properties. **Language Generalization**: robustness to typos, paraphrases, and vague or imprecise instruction formulations. The results are shown in Figure 9b, **EO-1** attains the highest overall completion score (**73.0%**) across all settings, outperforming GR00T-N1.5 (60.0%) and π_0 (51.0%). The advantage holds across all variation types: 67.0% in **Action**, 79% in **Language**, and 72.0% in **Visual**. The largest relative gain appears in language variations, indicating strong linguistic robustness, which is a persistent weakness of many current VLM-based policies. These findings confirm that interleaved vision-text-action training not only strengthens in-distribution manipulation



(a) Examples of the three generalization axes. **Visual:** Changes in object appearance, background, and lighting. **Action:** Altered object positions, novel viewpoints, or new object instances. **Language:** Typos, rephrasings, and fuzzy descriptions.



(b) Generalization performance breakdown. Success rates for EO-1, GR00T-N1.5, and π_0 across visual, action, and language variations.

Figure 9: Generalization illustration and performance evaluation with visual, action, and language variations.

performance but also provides resilience to distribution shifts in perception, control, and instructions, which is critical for real-world deployment of generalist robot policies.

We observe that EO-1 demonstrates strong robustness *In Action Domain* tasks, such as instruction variations like "Carefully place the aluminum foil box you're holding into the oven on the counter" versus "Place the aluminum foil box in right arm into the oven.", or changes in lighting conditions. This is attributed to the model's co-training on interleaved vision-text-action datasets, including grounding boxes, 2D trajectories, instruction fuzzing, and task reasoning corpora. This multi-modal training approach enhances EO-1's generalization ability in open-world scenarios, allowing it to adapt to diverse and dynamic environments. Despite these strengths, we also encounter challenges in generalizing *Out of the Action Domain*, particularly when working with limited robot control data. For example, in the Make Breakfast Sandwich task, variations in the position of the bread lead to a significant drop in success rates. On the other hand, in the WidowX Out-of-Box task with large-scale

Bridge data, where densely annotated interleaved data is available, we see a marked improvement in success rates. This suggests that interleaved vision-text-action co-training can fine-tune the model at the feature level, allowing it to map actions better and thus enhancing its generalization potential.

4.8. Enhanced Generalization with Unified Reasoning

To evaluate whether a *sole* interleaved vision–text–action policy can seamlessly integrate high-level reasoning with low-level control in real environments, we design a reasoning-control benchmark comprising four tasks: **Visual Rearrangement**, **Tic-Tac-Toe**, **Make Breakfast Sandwich**, and **Roast Beef Steak**. These tasks require joint perception, spatial reasoning, multi-step planning, and bimanual manipulation under real-world constraints. Unlike hierarchical baselines, GROOT-N1.5 and π_0 , which use an external LLM planner with a separate controller (GPT-4o), **EO-1** integrates planning and control within a single decoder, aligning reasoning and motor execution in an interleaved sequence. The four tasks are described as follows.

1. **Visual Rearrangement.** The robot must arrange objects in a target layout according to a reference image (*visual prompt*). It retrieves and places items sequentially, considering spatial relationships and maintaining collision-free motion. Success depends on interpreting visual cues and planning precise placements, such as positioning large objects (*e.g.*, a penguin) to anchor the scene.
2. **Tic-Tac-Toe.** The robot perceives the game state, reasons about optimal moves, and places markers to either win or block the opponent. This task evaluates the integration of visual perception, strategic decision-making, and precise placement, requiring real-time dynamic control.
3. **Make Breakfast Sandwich (Reasoning Control).** The robot assembles a sandwich by coordinating both arms to manipulate bread, ham, and lettuce. It plans the order of ingredients, retrieves items, and completes the sandwich. The task emphasizes long-horizon sequencing, bimanual coordination, and maintaining object stability throughout the process.
4. **Roast Beef Steak (Reasoning Control).** The robot performs a cooking sequence involving brushing oil, placing the steak in the oven, and operating the oven controls. It coordinates both arms and times actions to ensure proper cooking, evaluating temporal planning and multi-step manipulation.

We carefully constructed reasoning control data for model training. In Visual Rearrangement, the task involves arranging 5 objects across 16 grids, with 800 demonstrations collected for all 80 unique action steps. In Tic-Tac-Toe, the robot plays as red, performing 50 actions per step, resulting in 450 data samples that challenge real-time decision-making and strategic planning. The Make Breakfast Sandwich and Roast Beef Steak tasks rely on control data from section 4.6, emphasizing bimanual coordination and long-horizon task execution. All tasks utilize the *interleaved data construction pipeline* from Figure 3(b)(c), annotating spatial data (bounding boxes, points, and trajectories) to track object positions and scene state, verify task completion, and guide task planning. During training, **EO-1** is fine-tuned with pre-trained weights, combining regression and flow matching.

Reasoning Control with EO-1. Figure 10 illustrates qualitative rollouts of **EO-1** performing four tasks. In Tic-Tac-Toe, the model perceives the board state, moves to create threats or block the opponent, such as placing a piece at $(0, 2)$ to set up circle and later blocking $(1, 0)$ to prevent loss. It executes precise piece placements without hesitation. In Visual Rearrangement, **EO-1** parses a target layout image, identifies key spatial anchors like placing the gray penguin at $(0, 2)$ to define workspace boundaries, and incrementally matches the arrangement while avoiding collisions. In Roast Beef Steak, the model coordinates both arms: brushing oil on the steak, placing it in the oven at the optimal stage (Step 6), closing the door, and pressing the power button. This showcases its temporal reasoning and efficient role allocation between arms. Finally, in Make Breakfast Sandwich, **EO-1** plans the assembly by retrieving ingredients—bread, ham, and lettuce—in the correct sequence, finishing with the final bread slice to secure the sandwich fillings.

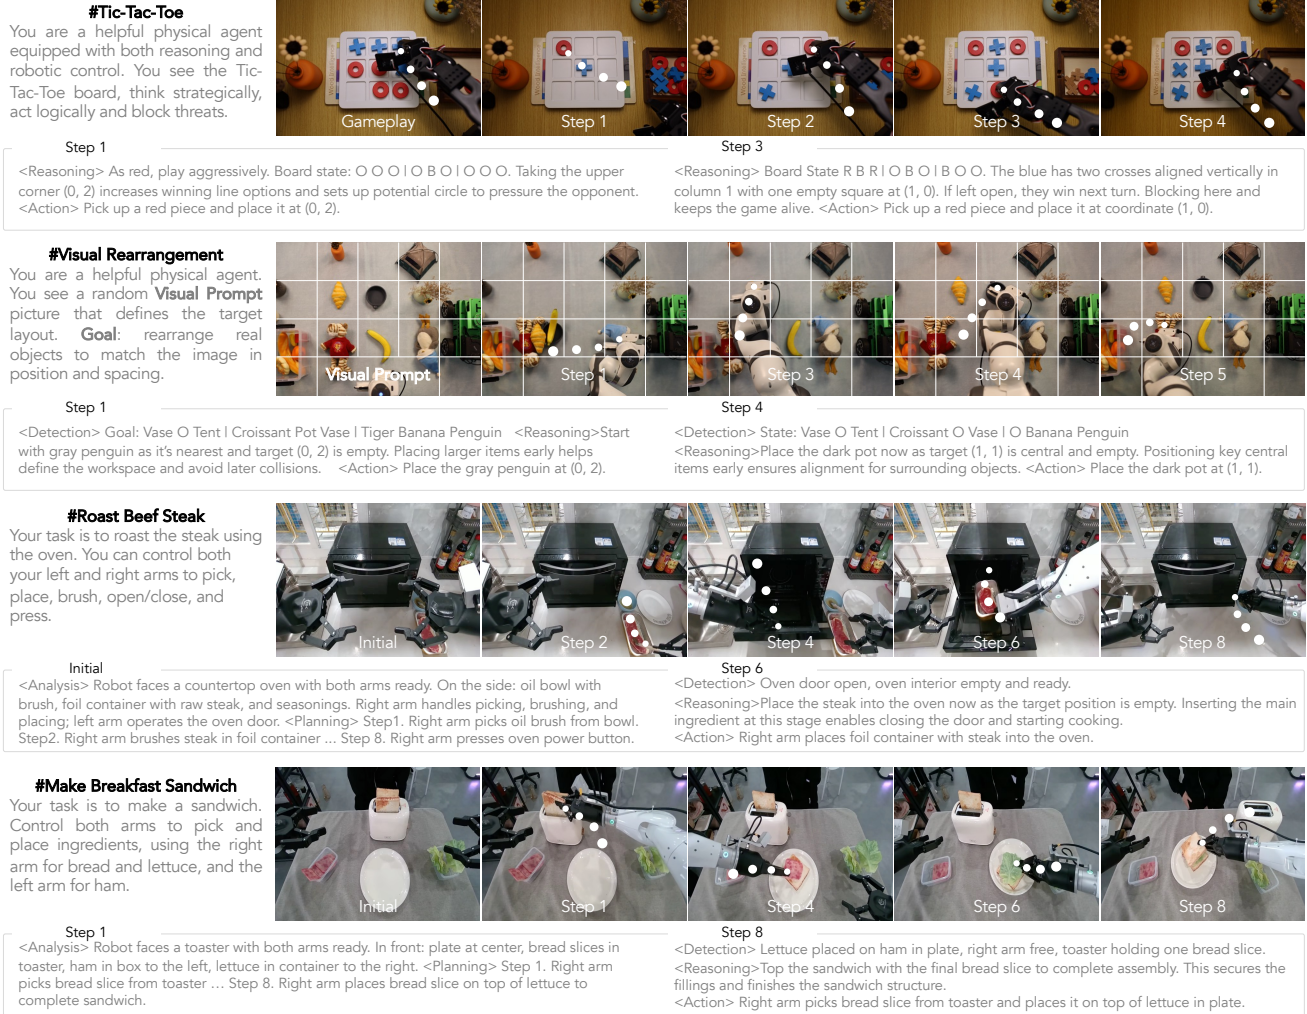


Figure 10: Qualitative rollouts for unified reasoning and control. Each sequence shows how EO-1 perceives the scene, reasons about next steps, and executes precise actions under a single interleaved policy: (a) **Tic-Tac-Toe**, (b) **Visual Rearrangement**, (c) **Roast Beef Steak**, (d) **Make Breakfast Sandwich**. Planning and execution remain aligned throughout, avoiding plan-act mismatches common in hierarchical pipelines.

Quantitative results. As shown in Figure 11, EO-1 achieves the highest average completion score across the four tasks, outperforming GR00T-N1.5 and π_0 . Per-task gains are consistent: Make Breakfast Sandwich (84.0% vs. 70.0%/ π_0 and 66.0%/GR00T-N1.5), Roast Beef Steak (55.0% vs. 41.0% and 46.0%), Visual Rearrangement (79.0% vs. 66.0% and 47.0%), and Tic-Tac-Toe (76.0% vs. 24.0% and 36.0%). The largest margin occurs in Tic-Tac-Toe (+40.0 points over the best baseline), reflecting superior game-state reasoning and real-time execution. The next largest is in Make Breakfast Sandwich (+14.0 points), showing strong temporal sequencing and precise control.

We attribute these gains to the alignment of **natural characteristics** achieved by the interleaved vision-text-action training, which unifies symbolic planning and continuous control within a cohesive backbone. This eliminates the interface gap between planner and controller, reducing error propagation and enabling smooth, context-aware action execution. Qualitative sequences confirm that EO-1 maintains coherent task strategies while performing fine-grained and stable manipulations, which is an essential capability for robust agents embodied in the open world.

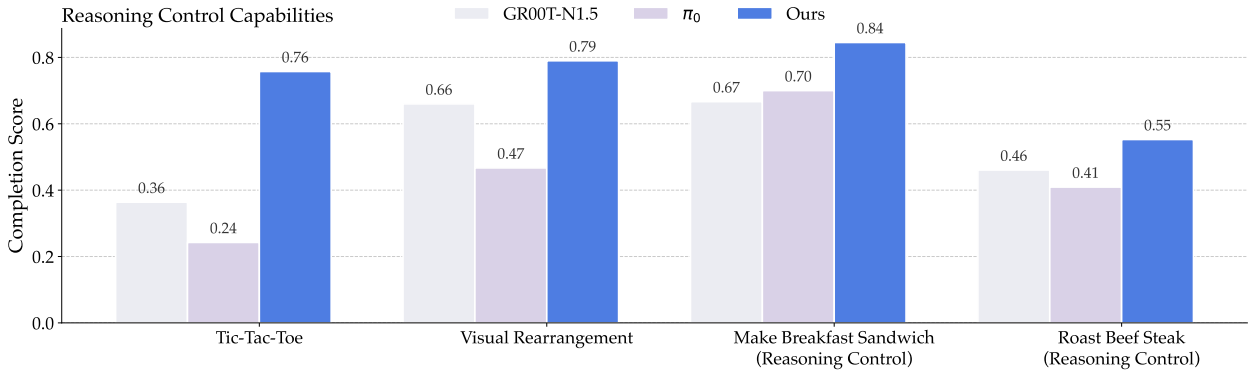


Figure 11: Success rates on the Reasoning-control benchmark. EO-1 outperforms both hierarchical baselines on all tasks and on average, with the largest gains in *Tic-Tac-Toe* and *Roast Beef Steak*.

4.9. Discussion: Architecture and Data Contributions to Generalization

In this section, we examine how the unified architecture and interleaved multimodal pre-training affect model generalization along the three axes defined in Section 4.7: *visual*, *action*, and *language*. Our objective is to isolate the contributions of (i) a hybrid decoding mechanism that combines discrete autoregression with continuous flow matching under shared parameters, and (ii) large-scale, interleaved vision-text-action supervision, to overall robustness and scaling behavior in the open world. The **benchmark settings** include:

1. **LIBERO** (Liu et al., 2023a): Evaluates the model’s capacity to fit training distributions and execute long-horizon, multi-stage manipulation, stressing policy learning and temporal credit assignment.
2. **Agibot Sandwich**, a single real-world bimanual task requiring precise sequencing of bread, ham, and lettuce placement—used to assess the impact of interleaved embodied data on single-task.
3. **WidowX Generalization** with BridgeData V2 (Walke et al., 2023): Five task groups selected from the WidowX 250 S **Prepare Vegetables** and **Arrange Cups** suites. We adopt the generalization protocols in Section 4.7: *Visual* (changes in object appearance, background, lighting), *Action* (altered object positions, novel viewpoints, new instances), and *Language* (typos, rephrasings, fuzzy descriptions). For training, we randomly sample relevant demonstrations from BridgeData at four scales: 50, 500, 5k, and 50k.

For the **model configurations**, we explore the following settings:

1. **EO-1 (fast)**: A fully autoregressive baseline using the Fast-tokenizer (Pertsch et al., 2025), trained on Bridge Data with interleaved vision-text-action pairs.
2. **EO-1 (base)**: Our proposed unified hybrid decoding model that is trained solely on robot control data. discrete autoregressive decoding for vision language tokens and continuous flow-matching denoising for robotic actions, sharing parameters across modalities.
3. **EO-1 (llava)**: **EO-1** co-trained on robot control data with general vision-language instruction datasets (LLaVA-Instruct-150K, RefCOCO) to improve grounding and natural language understanding.
4. **EO-1 (interleaved)**: **EO-1** trained with interleaved vision-text-action sequences from corresponding data subsets, emphasizing *Physical Common Sense*, *Spatial Understanding*, *State Estimation*, and *Task Reasoning*.

Unified Decoding Across Modalities Outperforms Autoregression. As shown in Figure 12, our hybrid model (**EO-1 (base)**) consistently outperforms fully autoregressive **EO-1 (fast)** across all generalization regimes and data scales. With only 50 demonstrations, **EO-1** achieves up to +0.25




Figure 12: WidowX generalization performance across data scales for EO-1 (base), EO-1 (interleaved), and EO-1 (fast) with interleaved vision-text-action data.

LIBERO Benchmark (Arch)	LIBERO-Spatial	LIBERO-Object	LIBERO-Goal	LIBERO-Long	Overall
EO-1 (fast)	93.3%	95.5%	88.0%	74.0%	88.0%
EO-1 (base)	99.7%	99.8%	99.2%	94.8%	98.2 %
Agibot Sandwich (Data)	In-distribution	Visual Generalization	Language Generalization	Action Generalization	Overall
EO-1 (base)	0.85	0.83	0.66	0.75	0.77
EO-1 (llava)	0.84	0.44	0.62	0.73	0.66
EO-1 (interleaved)	0.89	0.88	0.83	0.75	0.84
WidowX Generalization (Data)	Out-of-Box	Visual Generalization	Language Generalization	Action Generalization	Overall
EO-1 (base)	0.87	0.73	0.59	0.67	0.71
EO-1 (llava)	0.37	0.37	0.35	0.28	0.34
EO-1 (interleaved)	0.91	0.73	0.82	0.73	0.80

Table 5: LIBERO, WidowX and Agibot evaluation for EO-1 (base), EO-1 (interleaved), and EO-1 (fast) with interleaved vision-text-action data.

absolute gains in *Action Generalization*, and the margin persists at scale, reaching 0.91 Out-of-Box at 50k demos while pure autoregression plateaus. In particular, pure AR does not show emerging *Language Generalization*, stalling at 0.43 accuracy even as data grow, and its *Action Generalization* saturates at 0.23 from 5 to 50k demonstrations. Similar patterns appear in Table 5: in LIBERO, EO-1 (fast) achieves 88.0% overall success rate compared to 98.2% for EO-1 (base), underscoring the precision gains from incorporating flow-matching into action decoding. Across tasks, discrete-only decoding struggles to capture low-frequency action modes and fails to generalize robustly under distribution shift, even with large-scale data. In contrast, our hybrid architecture unifies discrete autoregression with continuous flow matching, enabling more accurate control and scaling-driven improvements that emerge strongly in open-world scenarios.

Scaling Interleaved Vision-Text-Action Data Boosts Generalization. Scaling interleaved demonstrations from 50 to 50k (Figure 12) shows that EO-1 (interleaved) consistently outperforms EO-1 (base), with faster and more robust generalization as the data grows. This is especially evident in *Language Generalization*, where EO-1 (interleaved) improves significantly from 0.27 to **0.82**, compared to EO-1 (base), which increases from 0.01 to **0.59**. This gap highlights the advantage of incorporating multimodal data aligned with both language and control tasks, enabling faster learning and more effective generalization. However, not all multimodal data are beneficial. As shown in Table 5, incorporating generic instruction-following data, such as EO-1 (llava), leads to a performance drop, from **0.77** to **0.66** on Agibot Sandwich and from **0.71** to **0.34** on WidowX Generalization. This suggests that *multimodal data misaligned with embodied control tasks biases the model toward linguistic priors, weakening its physical grounding*. We observe that **Out of the Action Domain** generalization struggles on limited datasets, despite dense multimodal annotations. In Agibot Sandwich, performance stays at 0.75 for both the base model and EO-1 (interleaved). In contrast, WidowX, trained on large-scale, diverse Bridge Data, improves significantly (0.73 vs. 0.67),

highlighting how multimodal data in multi-task settings enhances action generalization. These results underscore the need for finely annotated, task-aligned multimodal datasets incorporating physical knowledge for effective embodied control.

Finally, The superior performance of our unified architecture, which combines discrete autoregression for vision–language understanding with continuous flow matching for action generation, demonstrates how joint modeling enhances both control precision and generalization. This approach not only strengthens temporal alignment between multimodal understanding and motor control but also ensures stable, high-performance outcomes even in open-world, distribution-shifted scenarios, proving the power of integrating multiple modalities under a unified framework.

5. Related Work

Vision-Language-Action Models. Vision-language-action models (VLAs) have emerged as a promising approach for generalizable robot control by extending pre-trained vision-language models (VLMs) to action prediction. Autoregressive approaches (Brohan et al., 2023, Kim et al., 2024, Qu et al., 2025, Pertsch et al., 2025, Liu et al., 2024b) discretize continuous actions into tokens with linear bins, spatial grids or quantized DCT byte-pair encoding, naturally aligning with VLM training paradigms and exhibiting inherent advantages in knowledge transfer from web-scale pre-training. Despite these advantages, this discrete paradigm suffers from inference latency and limited resolution for high-frequency control tasks. Alternative works (Octo et al., 2024, Liu et al., 2024b, Black et al., 2024, BJORCK et al., 2025) introduce continuous action experts via diffusion or flow-matching (Ho et al., 2020, Lipman et al., 2022, Song et al., 2025), achieving fast inference and dexterous robotic control. Several efforts (Liu et al., 2025, Kim et al., 2025) attempt to combine discrete autoregression and continuous denoising, employing collaborative action ensemble or two-stage tuning procedure. However, these models suffer from the challenges of knowledge preservation and training efficiency, as the gradients from the action expert degrade the pre-trained VLM backbone (Driess et al., 2025, Black et al., 2025).

Co-training strategies have emerged as a promising solution that integrates diverse data sources and enables generalization to new environments. $\pi_{0.5}$ (Black et al., 2025) employs a two-stage co-training approach: first pre-training in multimodal data sources with FAST tokenized actions, then tuning the action expert for low-level control. Building on this foundation, (Driess et al., 2025) formalizes a single-stage recipe insulating the VLM backbone during VLA training. Complementary works (Lin et al., 2025, Zhou et al., 2025) further explore reasoning and generalization through π_0 -style co-training or purely autoregressive paradigms. Nevertheless, instruction following capabilities remain suboptimal, fundamentally due to architectural bottlenecks and the fact that current approaches are predominantly trained on separate image-text pairs or robotic episodes data, overlooking the rich temporal dynamics and causal relationships inherent in embodied interactions. Addressing these challenges necessitates better co-training strategies, improved cross-modal transfer mechanisms, and larger-scale interleaved datasets that comprehensively capture the full spectrum of robotic episodes.

Unified Multimodal Models. Unifying understanding and generation (Deng et al., 2025, Chen et al., 2025b, Liao et al., 2025) has witnessed remarkable progress and emergent properties beyond individual specialized models. Pioneering unified multimodal models (Zhou et al., 2024, Team, 2024, Xie et al., 2024) integrate these capabilities through autoregressive or diffusion modeling in a single architecture, fusing both text autoregression and visual generation in the shared backbone. Recent advancements (Deng et al., 2025, Ma et al., 2025a, Chen et al., 2025b, Xie et al., 2025) introduce modality-specific parameters or dual visual encoding to optimize training pipelines, facilitate joint

token semantic and superior performance. Alternative studies (Chen et al., 2025a, Dong et al., 2023, Ge et al., 2024, Pan et al., 2025) explored unified models with frozen LLMs, assemble LMMs backbones and visual generative models via adapters or learnable queries. While assembling frozen LLMs struggles to demonstrate promising capabilities in robotic learning (Kim et al., 2024, Bjorck et al., 2025, Qu et al., 2025), unifying understanding and generation inspires architecture design in representative works (Black et al., 2024, 2025, Bjorck et al., 2025) and exhibits significant emergent open-world capabilities.

6. Conclusion and Future Work

In this article, we present EO-Robotics, a fully open training recipe to foster the research community to develop an advanced embodied foundation model. EO-Robotics is composed of **EO-1** model, **EO-Data1.5M** dataset, and **EO-Bench** benchmark, where **EO-Data1.5M** is collected by a scalable pipeline for curating interleaved vision-text-action data to improve **EO-1** model’s open-world generalization. **EO-1** effectively synergizes auto-regressive decoding and flow matching denoising in single unified model that is co-trained with web data and curated interleaved embodied data **EO-Data1.5M**, enabling seamless perception, planning, reasoning, and acting in complex physical environments. Our evaluations across embodied reasoning and robot control experiments demonstrate that **EO-1** consistently outperforms strong vision–language–action baselines, exhibiting superior multimodal reasoning and dexterous robot control capabilities in open-world generalization. Moreover, **EO-Bench** presents a new benchmark to evaluate embodied reasoning capabilities of embodied foundation models. In summary, EO-Robotics is fully open-sourced to the community and achieves a forward step for equipping general-purpose autonomous robots with human-like reasoning and acting abilities.

EO-Robotics has made a solid step toward developing general embodied AI, featuring with seamless embodied reasoning and action in the open world. While the results with **EO-1** demonstrate promising robot control capabilities, there remains key aspects that future work can address. First, we aim to enhance **EO-1**’s reasoning and action ability to handle complex scenarios involving navigation, obstacle avoidance, failure detection and analysis, human intent recognition, and human-robot interaction/cooperation. This requires seamlessly integrate more general multimodal understanding, embodied reasoning, and action abilities into one system, leading to more easy-to-use robots in real life. Second, we plan to explore more efficient design of unified model architecture or asynchronous inference pipeline to enable human-like simultaneous reasoning and action. Finally, we will incorporate more data and robot embodiments (e.g., human data and humanoid robot) into training recipe to enable EO-Robotics to be a boarder foundation for general-purpose autonomous robots.

References

- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Siboz Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, et al. Gr00t n1: An open foundation model for generalist humanoid robots. *arXiv preprint arXiv:2503.14734*, 2025.
- 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

- Equi, Chelsea Finn, Niccolo Fusai, et al. $\pi_{0.5}$: a vision-language-action model with open-world generalization. *arXiv preprint arXiv:2504.16054*, 2025.
- 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.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.
- Qingwen Bu, Jisong Cai, Li Chen, Xiuqi Cui, Yan Ding, Siyuan Feng, Shenyuan Gao, Xindong He, Xu Huang, Shu Jiang, et al. Agibot world colosseo: A large-scale manipulation platform for scalable and intelligent embodied systems. *arXiv preprint arXiv:2503.06669*, 2025.
- Jiuhai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi Zhou, Saining Xie, Silvio Savarese, et al. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*, 2025a.
- Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025b.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024a.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 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*, pp. 24185–24198, 2024b.
- Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- 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, YenSung 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, Jen Dumas, Crystal Nam, Sophie Lebrecht, Caitlin Wittlif, Carissa Schoenick, Oscar Michel, Ranjay Krishna, Luca Weihs, Noah A. Smith, Hannaneh Hajishirzi, Ross Girshick, Ali Farhadi, and Aniruddha Kembhavi. Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models. *arXiv preprint arXiv:2409.17146*, 2024a.
- Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, et al. Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models. *arXiv preprint arXiv:2409.17146*, 2024b.
- Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, et al. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025.

- Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, et al. Dreamllm: Synergistic multimodal comprehension and creation. *arXiv preprint arXiv:2309.11499*, 2023.
- 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.
- Yuying Ge, Sijie Zhao, Jinguo Zhu, Yixiao Ge, Kun Yi, Lin Song, Chen Li, Xiaohan Ding, and Ying Shan. Seed-x: Multimodal models with unified multi-granularity comprehension and generation. *arXiv preprint arXiv:2404.14396*, 2024.
- Gemini, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023a.
- Gemini Gemini, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023b.
- Gemini-Robotics. Gemini robotics: Bringing ai into the physical world. *arXiv preprint arXiv:2503.20020*, 2025.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Chi-Pin Huang, Yueh-Hua Wu, Min-Hung Chen, Yu-Chiang Frank Wang, and Fu-En Yang. Thinkact: Vision-language-action reasoning via reinforced visual latent planning. *arXiv preprint arXiv:2507.16815*, 2025.
- Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. ReferItGame: Referring to objects in photographs of natural scenes. In Alessandro Moschitti, Bo Pang, and Walter Daelemans (eds.), *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 787–798, Doha, Qatar, October 2014a. Association for Computational Linguistics. doi: 10.3115/v1/D14-1086. URL <https://aclanthology.org/D14-1086>.
- Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. ReferItGame: Referring to objects in photographs of natural scenes. In Alessandro Moschitti, Bo Pang, and Walter Daelemans (eds.), *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 787–798, Doha, Qatar, October 2014b. Association for Computational Linguistics. doi: 10.3115/v1/D14-1086. URL <https://aclanthology.org/D14-1086>.
- 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.
- Jason Lee, Jiafei Duan, Haoquan Fang, Yuquan Deng, Shuo Liu, Boyang Li, Bohan Fang, Jieyu Zhang, Yi Ru Wang, Sangho Lee, et al. Molmoact: Action reasoning models that can reason in space. *arXiv preprint arXiv:2508.07917*, 2025.

- Xinghang Li, Peiyan Li, Minghuan Liu, Dong Wang, Jirong Liu, Bingyi Kang, Xiao Ma, Tao Kong, Hanbo Zhang, and Huaping Liu. Towards generalist robot policies: What matters in building vision-language-action models. *arXiv preprint arXiv:2412.14058*, 2024a.
- Xuanlin Li, Kyle Hsu, Jiayuan Gu, Karl Pertsch, Oier Mees, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, Sergey Levine, Jiajun Wu, Chelsea Finn, Hao Su, Quan Vuong, and Ted Xiao. Evaluating real-world robot manipulation policies in simulation. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2024b.
- Chao Liao, Liyang Liu, Xun Wang, Zhengxiong Luo, Xinyu Zhang, Wenliang Zhao, Jie Wu, Liang Li, Zhi Tian, and Weilin Huang. Mogao: An omni foundation model for interleaved multi-modal generation. *arXiv preprint arXiv:2505.05472*, 2025.
- Fanqi Lin, Ruiqian Nai, Yingdong Hu, Jiacheng You, Junming Zhao, and Yang Gao. Onetwovla: A unified vision-language-action model with adaptive reasoning. *arXiv preprint arXiv:2505.11917*, 2025.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- 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*, 2023a.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023b.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 26296–26306, 2024a.
- Jiaming Liu, Hao Chen, Pengju An, Zhuoyang Liu, Renrui Zhang, Chenyang Gu, Xiaoqi Li, Ziyu Guo, Sixiang Chen, Mengzhen Liu, et al. Hybridvla: Collaborative diffusion and autoregression in a unified vision-language-action model. *arXiv preprint arXiv:2503.10631*, 2025.
- Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang Su, and Jun Zhu. Rdt-1b: a diffusion foundation model for bimanual manipulation. *arXiv preprint arXiv:2410.07864*, 2024b.
- Chuofan Ma, Yi Jiang, Junfeng Wu, Jihan Yang, Xin Yu, Zehuan Yuan, Bingyue Peng, and Xiaojuan Qi. Unitok: A unified tokenizer for visual generation and understanding. *arXiv preprint arXiv:2502.20321*, 2025a.
- Yiyang Ma, Xingchao Liu, Xiaokang Chen, Wen Liu, Chengyue Wu, Zhiyu Wu, Zizheng Pan, Zhenda Xie, Haowei Zhang, Xingkai Yu, et al. Janusflow: Harmonizing autoregression and rectified flow for unified multimodal understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 7739–7751, 2025b.
- NextStep-1. Nextstep-1: Toward autoregressive image generation with continuous tokens at scale. *arXiv preprint arXiv:2508.10711*, 2025.
- Octo, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Charles Xu, Jianlan Luo, Tobias Kreiman, You Liang Tan, Lawrence Yunliang Chen, Pannag Sanketi, Quan Vuong, Ted Xiao, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024.

- OpenX-Embodiment, Abby O’Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, et al. Open x-embodiment: Robotic learning datasets and rt-x models. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- Xichen Pan, Satya Narayan Shukla, Aashu Singh, Zhuokai Zhao, Shlok Kumar Mishra, Jialiang Wang, Zhiyang Xu, Jiu hai Chen, Kunpeng Li, Felix Juefei-Xu, et al. Transfer between modalities with metaqueries. *arXiv preprint arXiv:2504.06256*, 2025.
- Karl Pertsch, Kyle Stachowicz, Brian Ichter, Danny Driess, Suraj Nair, Quan Vuong, Oier Mees, Chelsea Finn, and Sergey Levine. Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747*, 2025.
- Delin Qu, Haoming Song, Qizhi Chen, Yuanqi Yao, Xinyi Ye, Yan Ding, Zhigang Wang, JiaYuan Gu, Bin Zhao, Dong Wang, et al. Spatialvla: Exploring spatial representations for visual-language-action model. *arXiv preprint arXiv:2501.15830*, 2025.
- Pierre Sermanet, Tianli Ding, Jeffrey Zhao, Fei Xia, Debidatta Dwibedi, Keerthana Gopalakrishnan, Christine Chan, Gabriel Dulac-Arnold, Sharath Maddineni, Nikhil J Joshi, et al. Robovqa: Multi-modal long-horizon reasoning for robotics. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 645–652. IEEE, 2024.
- Mustafa Shukor, Dana Aubakirova, Francesco Capuano, Pepijn Kooijmans, Steven Palma, Adil Zouitine, Michel Aractingi, Caroline Pascal, Martino Russi, Andres Marafioti, et al. Smolvla: A vision-language-action model for affordable and efficient robotics. *arXiv preprint arXiv:2506.01844*, 2025.
- Haoming Song, Delin Qu, Yuanqi Yao, Qizhi Chen, Qi Lv, Yiwen Tang, Modi Shi, Guanghui Ren, Maoqing Yao, Bin Zhao, et al. Hume: Introducing system-2 thinking in visual-language-action model. *arXiv preprint arXiv:2505.21432*, 2025.
- Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024.
- Homer Walke, Kevin Black, Abraham Lee, Moo Jin Kim, Max Du, Chongyi Zheng, Tony Zhao, Philippe Hansen-Estruch, Quan Vuong, Andre He, Vivek Myers, Kuan Fang, Chelsea Finn, and Sergey Levine. Bridgedata v2: A dataset for robot learning at scale. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2023.
- Kun Wu, Chengkai Hou, Jiaming Liu, Zhengping Che, Xiaozhu Ju, Zhuqin Yang, Meng Li, Yinuo Zhao, Zhiyuan Xu, Guang Yang, et al. Robomind: Benchmark on multi-embodiment intelligence normative data for robot manipulation. *arXiv preprint arXiv:2412.13877*, 2024.
- Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin, Yuchao Gu, Zhijie Chen, Zhenheng Yang, and Mike Zheng Shou. Show-o: One single transformer to unify multimodal understanding and generation. *arXiv preprint arXiv:2408.12528*, 2024.
- Jinheng Xie, Zhenheng Yang, and Mike Zheng Shou. Show-o2: Improved native unified multimodal models. *arXiv preprint arXiv:2506.15564*, 2025.
- Jianwei Yang, Reuben Tan, Qianhui Wu, Ruijie Zheng, Baolin Peng, Yongyuan Liang, Yu Gu, Mu Cai, Seonghyeon Ye, Joel Jang, et al. Magma: A foundation model for multimodal ai agents. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 14203–14214, 2025.

Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu, and Chunyuan Li. Llava-next: A strong zero-shot video understanding model, April 2024a. URL <https://llava-vl.github.io/blog/2024-04-30-llava-next-video/>.

Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*, 2024b.

Ruijie Zheng, Yongyuan Liang, Shuaiyi Huang, Jianfeng Gao, Hal Daumé III, Andrey Kolobov, Furong Huang, and Jianwei Yang. Tracevla: Visual trace prompting enhances spatial-temporal awareness for generalist robotic policies. *arXiv preprint arXiv:2412.10345*, 2024.

Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Jacob Kahn, Xuezhe Ma, Luke Zettlemoyer, and Omer Levy. Transfusion: Predict the next token and diffuse images with one multi-modal model. *arXiv preprint arXiv:2408.11039*, 2024.

Zhongyi Zhou, Yichen Zhu, Minjie Zhu, Junjie Wen, Ning Liu, Zhiyuan Xu, Weibin Meng, Ran Cheng, Yaxin Peng, Chaomin Shen, et al. Chatvla: Unified multimodal understanding and robot control with vision-language-action model. *arXiv preprint arXiv:2502.14420*, 2025.

A. Contributors and Acknowledgments

A.1. Core Contributors

Delin Qu, Haoming Song, Qizhi Chen, Zhaoqing Chen, Xianqiang Gao, Dong Wang, Bin Zhao

Model Architecture and Training Delin Qu, Haoming Song, Qizhi Chen

Real-Robot Experiments Haoming Song, Delin Qu, Haoran Yang, Jiacheng Bao, Qizhi Chen

Simulation and Benchmarking Delin Qu, Haoming Song, Xianqiang Gao, Xinyi Ye

Data and Benchmark Curation Zhaoqing Chen, Xianqiang Gao, Qizhi Chen, Delin Qu

Research Leads Dong Wang, Bin Zhao

A.2. Contributors

Qi Lv, Xinyi Ye, Modi Shi, Guanghui Ren, Cheng Ruan, Maoqing Yao, Haoran Yang, Jiacheng Bao.

A.3. Acknowledgments


This work is supported by Shanghai Artificial Intelligence Laboratory. We acknowledge the AgiBot team, including Modi Shi, Cheng Ruan, Guanghui Ren, and Maoqing Yao, for providing hardware support and maintenance of the AgiBot G-1 robots. Special thanks to Xinyi Ye for video editing contributions and Qizhi Chen for webpage design.

B. Training Dataset Statistics

Dataset	Samples	Tokens	Duration (hr)	FPS	Camera View	Category
LLaVA-Video-178K	2.6M	5.95B	–	–	Third-person	Web multimodal data
LLaVA-1.5	1.1M	390.45M	–	–	Third-person	Web multimodal data
Pixmo-Points	1.8M	538.92M	–	–	Third-person	Web multimodal data
RoboVQA	0.2M	218.29M	–	–	Egocentric	Web multimodal data
RefCOCO	50.1K	9.46M	–	–	Third-person	Web multimodal data
Agibot-Beta	213.8M	116.17B	2560.7	30	Egocentric, left, right	Real robot
DROID (OXE)	23.1M	8.41B	428.3	15	Left, Right, wrist	Real robot
RT-1 (OXE)	3.7M	579.32M	338.4	3	Egocentric	Real robot
Bridge-v2 (OXE)	2.0M	1.89M	111.1	5	Shoulder, left, right, wrist	Real robot
Robo-Mind	2.1M	2.09M	19.3	30	Egocentric, left, right	Real robot
SO100-Community	7.9M	1.19B	72.1	30	Third-person, wrist	Real robot
IPEC-Franka	0.8M	108.52M	19.7	10	Third-person, wrist	Real robot
EO-Data1.5M	1.5M	1.0B	13.9	30	Egocentric, left, right	Interleaved Embodied Data
Total multimodal data	6.2M	7.1B	–	–	–	–
Total robot data	253.4M	127.3B	–	–	–	–
Total Interleaved data	1.5M	1.0B	13.9	30	Egocentric, left, right	Interleaved Embodied Data
Total	260.6M	135.4B	–	–	–	–

Table 6: Pre-training Dataset Statistics.

C. EO-1 Reasoning Examples



Where should the gripper grasp the pot to move it?




Indicate the interior of the cubic drawer.




Predict the grasp location on the right side of the cloth.



Which nearest point should the gripper grasp to open the drawer?



Indicate the object grasped by the robotic arm.




Where should the blue plate be placed on the dish rack?

Figure 13: Object pointing examples.




Mark the trajectory from the left gripper to the cleanser.




How can the gripper place the air column bag into the carton?




What's the path for placing the carrot into the sink?



What rotational trajectory does the arm need to grasp the pepper?




How to move the robotic arm to let the apple drop naturally?



What trajectory does placing the pot lid on the pot follow?

Figure 14: Trajectory prediction examples.



Draw a box around the vegetables in the image.



Locate the pot on the burner.




Identify the green snack bag and the coke.



Box the objects that can contain items.



Mark as many vegetables as possible in the image.



Box the objects that can be used to drink water.

Figure 15: Object referring examples.





Figure 16: Multiview reasoning examples.

Task: Put the blue cube on the right side of the table on top of the rectangular block.



Task: Put the red object into the pot.





Figure 17: Task planning examples.



Question: What is the necessary next step in realizing **retrieve a brown chip bag from the top drawer and place it on the counter?**


Answer: Place the brown chip bag on the counter.



Question: What should be your next priority in order to fulfill **pick an apple from a white bowl?**

Answer: Lift the apple out of the bowl.

Figure 18: Subtask QA examples.



Video frames

Question: Provide a comprehensive overview of the video's contents.

Answer: This video showcases a robotic arm delivering a menu to a seated guest at a dining table. The video begins with the robotic arm holding a red menu and moving towards the table. It carefully extends the menu towards the guest, who reaches out to take it. After handing over the menu, the robotic arm steps back, maintaining a respectful distance as the guest examines the menu. The setting is a cozy dining area with a wooden floor and simple furniture. The table is covered with a green tablecloth, and various items like a glass bottle and a striped napkin are visible. The atmosphere suggests a casual yet modern dining experience, enhanced by the presence of the robotic arm.



Video frames

Question: Please provide a comprehensive caption for this video.

Answer: This video showcases a robotic arm efficiently washing a pot and spatula in a kitchen setting. The sequence begins with the robotic arm picking up the light-colored wooden spatula from the sink. It then moves to position the spatula under the running water, using its right arm to turn on the faucet. The robotic arm rinses the spatula, ensuring all sides are clean. After rinsing, the robotic arm uses both arms to pass the spatula to the other hand and places it in a nearby storage box. Next, the arm reaches for the pot, lifting it out of the sink. It proceeds to rinse the pot, moving it around to ensure all sides are clean. The arm then drains any excess water and places the pot beside the storage area. The video takes place in a clean, organized kitchen environment. A blue storage box sits on the counter, holding various utensils. A bottle of dish soap is visible near the sink, along with a sponge holder. The background features a white wall, emphasizing the cleanliness and simplicity of the space.

Figure 19: Episode caption examples.



Previous frames

Current frame

Question: What action is available to undertake at this moment?

Answer: Lift the yellow cube from the stack.



Previous frames

Current frame

Question: What can be undertaken at this very moment?

Answer: Push the oven door closed.

Figure 20: Affordance QA examples.

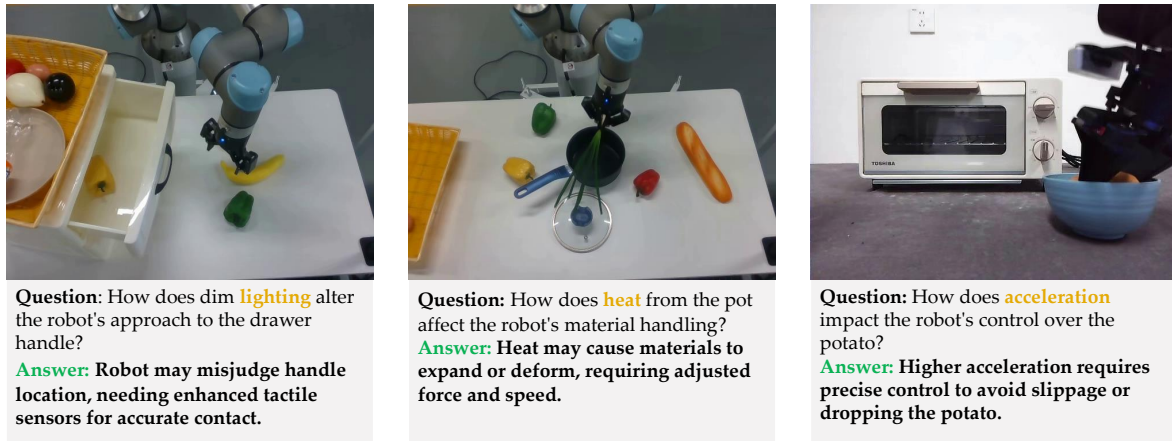


Figure 21: Physical common sense examples.

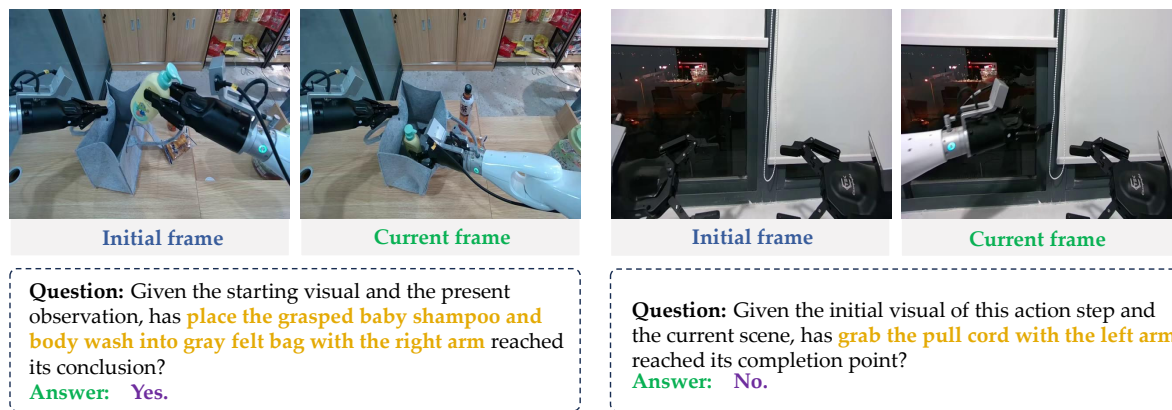


Figure 22: Failure detection examples.

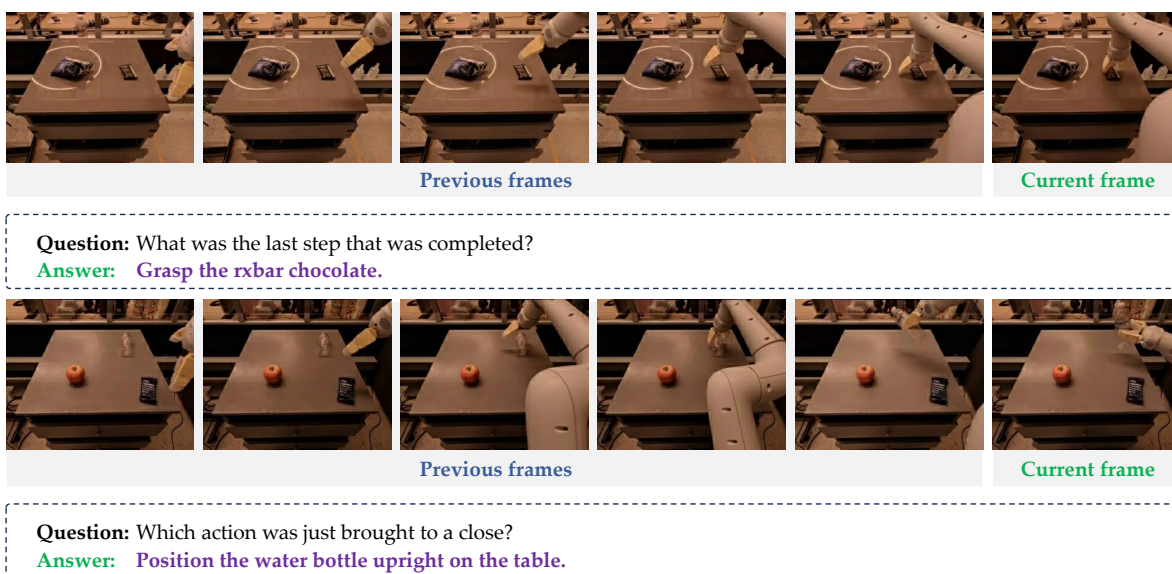


Figure 23: Process verification examples.

D. Data Annotation Pipeline Details

In this section, we provide additional details of the data preprocessing pipeline, which complements the description in the main paper. The process is divided into four major components: i) video filtering and curation, ii) video splitting and captioning, iii) QA generation for temporal and spatial reasoning, and iv) cleaning and rewriting.

D.1. Video Filtering and Curation

The first step focuses on ensuring the diversity and quality of robot videos. Since existing robot datasets often consist of highly redundant scenes collected in constrained environments, we employ a feature-based clustering strategy to enhance visual variety. Specifically, pretrained visual backbones are used to extract video features, followed by K-means clustering. From each cluster, we select a balanced set of videos, and further apply human-assisted inspection to remove samples from the same scene. Figure 24 illustrates the interface for this process, where clusters are visualized and manually pruned to improve dataset coverage.

To further enrich the data for tasks requiring multiview or cross-scene reasoning, we organize all available observations into a unified interface (Figure 25). This allows selecting videos with sufficient viewpoint diversity, ensuring that multiview reasoning tasks are well-supported.

D.2. Video Splitting and Captioning

Once a diverse video set is curated, we process them into finer-grained clips. Annotators or pretrained VLMs are instructed to segment videos into meaningful short clips, each containing a single subtask. For each clip, descriptive captions are generated to capture the specific robot actions. These captions serve dual purposes: they support video captioning QA data and also provide contextual prompts for subsequent reasoning annotations. Figure 26 and 27 show examples of the prompts used to guide caption generation.

D.3. QA Generation for Temporal and Spatial Reasoning

The next stage is to construct question-answer pairs that capture both temporal and spatial reasoning abilities. We adopt a prompting strategy, where VLMs are guided with carefully designed templates to generate diverse questions, and human annotators refine the final answers.

D.3.1. Temporal Reasoning

For temporal reasoning, we focus on two key aspects: *task planning* and *physical common sense*. Task planning questions require understanding the sequential structure of subtasks, while physical common sense questions evaluate the model’s ability to reason about object dynamics and feasibility in the physical world. Figure 28 and 29 provide examples of the prompts used to elicit these two categories of questions.

D.3.2. Spatial Reasoning

For spatial reasoning, we construct four sub-tasks: trajectory prediction, object pointing, object referring, and multiview reasoning. The multiview data are generated using the same prompting strategy as object pointing, without a separate prompt design. Figure 30, 31, and 32 show representative prompts used to construct these spatial reasoning tasks. Together, they ensure that the dataset

fractal20220817_data_0.1.0_lerobot

Create Merge Group

Merge name (e.g., 'graspin'): Create Merge

Exclude Selected Clear Selection

Existing Merge Groups

drawer
Clusters: 0, 2, 4, 6, 8, 9, 10, 11, 14, 13, 15, 16, 18, 21, 22, 24, 25, 26, 28, 29, 31, 33, 34, 35, 37, 38, 40, 41, 42, 43, 44, 47, 46, 45, 48, 49, 51, 52, 53, 54, 55, 57, 60, 58, 61, 59, 62, 63, 64, 66, 67, 65, 68, 69, 72, 71, 75, 77 Add Delete

refrigerator
Clusters: 79, 76, 78, 56, 27, 23, 3 Add Delete

table
Clusters: 1, 5, 7, 12, 17, 19, 32, 36, 39, 70, 73, 74 Add Delete

Excluded Clusters

Export

Export Current Dataset Export All Merges

Clusters

Cluster 0 drawer 950 videos

Cluster 1 table 866 videos

Cluster 2 drawer 1918 videos

Cluster 3 refrigerator 925 videos

Cluster 4 drawer 981 videos

Cluster 5 table 755 videos

Cluster 6 drawer 2032 videos

Cluster 7 table 1179 videos

Cluster 8 drawer 2283 videos

Figure 24: Web interface for video filtering and clustering. We extract features using a pretrained backbone and apply K-means clustering to group visually similar videos. Human verification is then used to remove redundant samples and ensure diversity.

captures fine-grained spatial relations necessary for embodied intelligence.

Cleaning and Rewriting Finally, to enhance linguistic quality and ensure answer validity, we perform a post-processing step. Rule-based filtering is applied to remove invalid or ambiguous spatial answers, and an LLM is prompted to rewrite QA pairs. This step increases textual diversity and introduces natural fuzziness, making the dataset more robust. An example of the rewriting prompt is shown in Figure 33.



Figure 25: Manual selection of multiview videos. All observations from different datasets are displayed, and the interface allows selecting samples with sufficient viewpoint diversity for multiview reasoning tasks.

Video Caption Prompt

Task Description

You are an expert in visual analysis. Your task is to generate a concise video caption, under 150 words, for a video depicting a robotic arm performing a specific task. You will be provided with a brief description of the task. Your caption should comprehensively describe the overall task environment, the actions performed in the video, and the environment where the task takes place.

Output Format

Organize the output in three sections:

1. Briefly describe the general content of the video.
2. Detail the actions in the video in temporal order.
3. Elaborate on the environmental factors present in the video.

Output Example:

This video showcases noodles being prepared.

The video begins with someone standing by a large boiling pot and putting noodles into one of six small circular compartments dipped into boiling water in a large circular stainless steel pot. The person, using their right hand, stirs the noodles in a circular motion using light brown wooden chopsticks, and their left holds the compartment into place using a dark metallic stick-like object.

Figure 26: Video caption prompt.

Clip Caption Prompt

Task Description

You are an expert in visual analysis. Your task is to generate a concise video caption, under 40 words, for a video depicting a robotic arm performing a specific task.

You will be provided with a brief description of the task. Your caption should describe the spatial relation and state of the objects related to the given task, including the robotic arm.

Additional Rules

1. Avoid objects unrelated to the task.
2. Avoid motion description. You should only give the static description.

Example:

Input task: Dip the beef and green bok choy.

Output caption: The beef is on the bottom-right plate. Green bok choy is on the top-right plate. The pot, containing a hotpot strainer, is in the center-right.

Figure 27: Clip caption prompt.

Task Planning Prompt

Task Description

You are an AI visual assistant specialized in analyzing videos performing robotic tasks.

You will receive a video of a robotic arm performing a task, and it's represented by a series of frame images. Robotic arm is always the main subject of the video, so your answer should revolve around it. Additionally, you will receive the main task of the video. Your task is to identify the primary task performed by the robotic arm during the video, extract the necessary steps to complete it, and specify the frame range for each step. You should strictly follow the output format, in reference of the following output format and output examples.

Input

Task name: The performed task by the robotic arm in the video. This task name is annotated by human, so sometimes it may be incorrect or null. You should write your task summary in reference of it.

Video frames: All frame images form the robotic video. Each image represents a video frame. The input frames follow the temporal \ sequence in the robotic video.

Target

1. Task identification: Identify the main task the robotic arm is performing. This task could be a clear goal or a series of related activities (e.g., assembling furniture, repairing equipment, preparing food, etc.). Briefly describe the primary task in one sentence.
2. Step Extraction: Extract the key steps required to complete the primary task, ensuring that each step is clearly described and logically ordered. Each step may include:
 - Specific actions with corresponding scene (e.g., retrieve pear from the shelf, place the held pear into the plastic bag).
 - Frame window: Specify the start and end frame for each step. The start and end frame are both included in the step. Make sure that your last end_frame is consistent with the input frame number. Also, due to the unskilled operation, in some tasks (but not all tasks), robotic arm may swings in a position during the step, resulting in a abnormally long frame range. You should include it into the frame range of your step.

Output Format:

Provide the task description and steps in two parts, formatted as JSON:

1. Task Summary: A string summarizing the primary task in the video without mentioning the subjects - the robotic arm.
2. Steps: An array where each element represents a step, containing:
 - step_description: A concise description of the step which the action being performed in the format of verb phrases without mentioning the subjects - the robotic arm (e.g., "Add syrup in the glass").
 - start_frame: The start frame of the step.
 - end_frame: The end frame of the step.

Additional Rules:

1. You should only output the JSON content in the example output, and avoid the extra thinking content in your output, such as "Okay, here's my break down of the video...".
2. You should avoid too detailed step description for each step. Keep it as a description of a general step in the human view, not a basic action for the robotic arm. E.g., avoid step descriptions like "move towards ...", "grasp ...", "lift the apple" ..., use "retrieve the apple from the shelf" instead.
3. Avoid using only one frame for a step. A step should include at least two frames.

Example output:

```
{
  "task_summary": "Pick up an apple in the supermarket.",
  "steps": [
    {
      "step_description": "Retrieve the apple from the shelf.",
      "start_frame": 0,
      "end_frame": 17
    },
    {
      "step_description": "Place the held apple into the plastic bag in the shopping cart.",
      "start_frame": 18,
      "end_frame": 30
    }
  ]
}

{"task_summary": "Open the fridge to get food",
"steps": [{"step_description": "Open the refrigerator door with the left arm.",
"start_frame": 0,
"end_frame": 7}, {"step_description": "Pick up the dried strawberries from the fridge with the left arm.",
"start_frame": 8,
"end_frame": 14}, {"step_description": "Place the dried strawberries held in the left arm on the table",
"start_frame": 15,
"end_frame": 20}, {"step_description": "Close the refrigerator with the right arm",
"start_frame": 21,
"end_frame": 25}]}
```

Figure 28: Task planning prompt.

Physical Question Prompt

Task Description

You are an AI visual assistant specializing in robotic task video analysis. You will receive a video of a robotic arm (provided as a series of frames), which is the main subject; your response must focus on it. Input also includes the robot's main task and its current operational step.

Generate 6 extremely challenging questions designed to evaluate a state-of-the-art model's embodied reasoning for real-world robotic tasks. These questions must probe deep embodiment abilities. They should focus on what the robot must do or what it needs to pay attention to for successful task completion. You must use your utmost creativity to make these questions exceptionally challenging and diverse in their perspectives.

Question Categories & Guidelines

Common Sense Learning:

Focus on how current video elements or conditions, governed by common sense, impact the task. Questions should address factors critical for task success and the robot's actions/trajjectory. Consider physical principles (e.g., gravity, balance, stability, support, elasticity, deformation, lighting, heat, motion, acceleration), object existence and states, other ongoing actions, and robot conditions.

Example: "How does the target block's softness affect grasping?"

Example: "Will the closed refrigerator door impede grasping the tomato?"

Example: "What must the arm consider, with its camera above its head?"

Counterfactual Reasoning:

If a factor (e.g., environment, object existence/condition, main task, robot condition) changes, how must the robot adapt its movement to still complete the task?

Example: "If the table is wet, can the current trajectory still complete the task?"

Example: "To grasp a durian instead of a cube, how must the robot's action change?"

Spatial Understanding:

Address spatial relationships between objects in the video or how these relationships affect specific robotic actions.

Example: "What is the spatial relation between the fork and knife?"

Example: "Can the arm currently reach the chocolate in the pot?"

Additional Rules

1. Combined Knowledge: Questions must require both video information and external common sense/physical knowledge. You should AVOID questions answerable solely by external knowledge.
2. Minimalism in Questions: You should give as little information as possible in your questions. The model being evaluated must extract necessary details from the video.
3. Diversity in Perspectives: Generate questions from perspectives as diverse as possible. Avoid repeatedly asking questions focusing on the same aspect (e.g., weight) across different videos.

Output Requirements:

Generate exactly 6 questions and the corresponding answers. Distribute them as: 2 Common Sense Learning, 2 Counterfactual Reasoning, and 2 Spatial Understanding questions. Each generated question must be concise, with a maximum of 15 words.

Output Example:

```
{
  "common_sense_learning": [
    {"question": "How does spoon's reflective surface affect visual servoing for grasping?", "answer": "Reflections can mislead visual sensors, requiring robust algorithms or alternative sensing for accurate grasping."},
    {"question": "What utensil property dictates robot's lifting trajectory after grasping?", "answer": "Its elongated shape and center of mass require a primarily vertical lift to maintain stability and avoid collisions."}
  ],
  "counterfactual_reasoning": [
    {"question": "If basket was full of water, how must spoon retrieval adapt?", "answer": "Robot must use slower, controlled movements, accounting for buoyancy, water resistance, and avoiding splashes."},
    {"question": "If the spoon was extremely hot, how must robot's action change?", "answer": "The robot would need a heat-resistant gripper or abort the task if not equipped to handle it safely."}
  ],
  "spatial_understanding": [
    {"question": "Which compartment offers easiest spoon access for current gripper pose?", "answer": "The compartment with the most exposed spoon handle, minimal surrounding clutter, and direct, unobstructed gripper access."},
    {"question": "What spatial clearance is critical when extracting the spoon from dense basket?", "answer": "Sufficient vertical and lateral clearance to avoid collision with other utensils and the basket rim during lift-out."}
  ]
}
```

Figure 29: Physical question prompt.

Trajectory Prediction Prompt

Task Description

Your task is to locate the position of {object_A} and the position of {object_B} in the given image as points. Then, connect these two points with 6 evenly distributed trajectory points.

Input

An image containing two objects: {object_A} and {object_B}.

Target

1. Identify the position of {object_A} (start point) in the image.
2. Identify the position of {object_B} (end point) in the image.
3. Connect the start point and end point using 6 evenly distributed trajectory points. These trajectory points should form a curve or line.

Output Format

The coordinates of the 6 trajectory points should be output in the following XML format:

```
<points x1=xxx y1=xxx x2=xxx y2=xxx x3=xxx y3=xxx x4=xxx y4=xxx x5=xxx y5=xxx x6=xxx y6=xxx></points>
```

Additional Rules

1. Both location annotations of the start point and end point should be precise, following the object descriptions.
2. The first and last of the 6 trajectory points are the two located points themselves.
3. The trajectory points must be evenly distributed.
4. The 6 trajectory points should form a curve or line.
5. If {object_A} is in the gripper, locate the center of the object in the gripper. If not, locate the object itself on the desk or on the ground.

Example Output

```
<points x1=100 y1=200 x2=150 y2=220 x3=200 y3=240 x4=250 y4=260 x5=300 y5=280 x6=350 y6=300></points>
```

Figure 30: Trajectory prediction prompt.

Object Pointing Prompt

Task Description

Your task is to locate all objects (along with a gripper of the robotic arm if present) mentioned in the given sentence as points and output all coordinates in XML format.

Input

A sentence containing a list of objects to be located in the image. The sentence is provided for analysis.

Target

1. Read and interpret the sentence carefully.
2. Locate all objects mentioned in the sentence as points in the image. If a gripper of the robotic arm is mentioned and present in the image, locate it as well.

Output Format

The coordinates of each located object should be output in the following XML format:

```
<points x=xxx y=xxx>object_name</points>
```

Additional Rules

1. Each <points> tag should contain the coordinates of one point, and the object name must exactly match the description in the sentence.
2. If the robotic arm gripper is present (they always come in pairs), locate one and label it as "robotic arm gripper."
3. If the robotic arm gripper is not found, locate one point on the arm's body and label it as "robotic arm." If no arm is found, ignore it.

Example Output

```
<points x=100 y=150>black robotic arm gripper</points>
<points x=200 y=250>oven</points>
```

Figure 31: Object pointing prompt.

Object Referring Prompt

Task Description

Your task is to outline the position of every small object in the provided image. The goal is to accurately mark the location of specific objects for further analysis.

Input

An image containing various objects, including robotic arm grippers, kitchenware, food items, geometric shapes, containers, boxes, clothing items, and other objects on the desktop.

Target

1. Identify and outline the position of the grippers of the robotic arms (if present). Only outline the gripper portions, excluding the main bodies of the arms. There may be one or two robotic arm grippers in the image.
2. Identify and outline the position of all other objects, including kitchenware, fruit, vegetables, food items, snacks, geometric shapes, containers, boxes, clothing items (e.g., T-shirts, pants), ovens, and any other objects on the desktop (if any).
3. Provide fine-grained and comprehensive annotations for each identified object.

Output Format

The coordinates of all the outlined objects should be output in JSON format:

```
[{"bbox_2d": [x1, y1, x2, y2], "label": "object_label"}, {"bbox_2d": [x1, y1, x2, y2], "label": "object_label"}, ...]
```

Additional Rules

1. The annotations should be as fine-grained and comprehensive as possible.
2. Only the gripper portions of the robotic arms should be outlined, excluding the main bodies of the arms.
3. Include all identified objects in the image, from kitchenware to clothing items and more, with precise annotations for each.

Example Output

```
[{"bbox_2d": [420, 130, 528, 236], "label": "grippers of the robotic arms"}, {"bbox_2d": [97, 163, 405, 405], "label": "oven"}, {"bbox_2d": [101, 390, 479, 479], "label": "vegetable"}, {"bbox_2d": [0, 374, 482, 482], "label": "vegetable"}]
```

Figure 32: Object referring prompt.

Semantic Deblurring Prompt

Task Description

You are an AI visual assistant specializing in robotic task analysis. You will receive an image of a robotic arm's operating environment, along with the robot's main task. Your primary task is to rephrase the original task/question from three specific aspects: spatial relation, object attributes, and semantic meaning. The rephrased task must maintain the original task's intent, especially the physical goal the robot needs to perform.

Rephrasing Aspects

Rephrase the task by focusing on one of the following aspects. Ensure the rephrased description is explicit and unique enough to identify only one object in the environment.

Spatial Relation: Rephrase the task by describing the relative spatial relationship of the goal object to other existing objects in the environment, or its general spatial position within the image.

Example: "Grasp the object to the left of the bottle."

Example: "Pick up the object in the middle of the image."

Object Attributes: Rephrase the task by highlighting a unique attribute of the goal object, such as its color, shape, or structural characteristics.

Example: "Retrieve the purple object."

Example: "Move the rectangular object."

Example: "Pick up the object with a brim."

Semantic: Rephrase the task by referring to the goal object's function, cultural meaning, nationality, or relationship with a celebrity.

Example: "Pick up the tool used to feed a baby."

Example: "Get the Chinese food."

Example: "Retrieve the photo of Michael Jackson."

Additional Rules

Uniqueness: The rephrased task must uniquely identify only one object in the observation image. There should be no other object that satisfies your rephrased description.

Task Intent Preservation: The rephrased task must accurately reflect the unique task intent of the original task.

Output Example

```
["spatial_relationship": ["Grasp the object to the left of the pot.", "Retrieve the object in the upper right."],
```

```
"object_attribute": ["Retrieve the purple object.", "Move the round object to the right."],
```

```
"semantic": ["The environment is dark now. Do something to make it light.", "Grasp the second letter of the alphabet."]
```

Figure 33: Semantic deblurring prompt.