

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 ONETWVLA: A UNIFIED VISION-LANGUAGE-ACTION MODEL WITH ADAPTIVE REASONING

Anonymous authors

Paper under double-blind review

## ABSTRACT

General-purpose robots capable of performing diverse tasks require synergistic reasoning and acting capabilities. However, recent dual-system approaches, which separate high-level reasoning from low-level acting, often suffer from challenges such as limited mutual understanding of capabilities between systems and latency issues. This paper introduces OneTwoVLA, a single unified vision-language-action model that can perform both acting (System One) and reasoning (System Two). Crucially, OneTwoVLA adaptively switches between two modes: explicitly reasoning at critical moments during task execution, and generating actions based on the most recent reasoning at other times. To further unlock OneTwoVLA’s reasoning and generalization capabilities, we design a scalable pipeline for synthesizing embodied reasoning-centric vision-language data, used for co-training with robot data. We validate OneTwoVLA’s effectiveness through extensive experiments, highlighting its superior performance across four key capabilities: long-horizon task planning, error detection and recovery, natural human-robot interaction, and generalizable visual grounding, enabling the model to perform long-horizon, highly dexterous manipulation tasks such as making hotpot or mixing cocktails. Project page: <https://onetwovla-anonymous.github.io>.




Figure 1: **Overview.** OneTwoVLA is a single unified vision-language-action model capable of both reasoning and acting. Crucially, OneTwoVLA can adaptively reason at critical moments during execution (e.g., upon completing subtasks, detecting errors, or requiring human inputs), while generating actions at other times.

## 1 INTRODUCTION

A distinctive characteristic of human physical intelligence is the ability to both *reason* and *act* (Varela Francisco et al., 1991; Anderson, 2003). Crucially, these processes are not separate but flexibly interleaved, creating a powerful synergy—reasoning guides our actions, while actions provide feedback that informs subsequent reasoning. Consider someone preparing a dish: *reasoning* enables them to develop a comprehensive understanding of the scene and goal (e.g., interpreting the recipe, planning the sequence of steps), while *acting* corresponds to the physical execution (e.g., chopping, mixing) that grounds abstract reasoning in the real world. This paper aims to imbue robots with a similar synergistic relationship between reasoning and acting.

054 Current approaches (Ahn et al., 2022; Hu et al., 2023; Shi et al., 2025; Team et al., 2025; Figure, 055 2025) often draw inspiration from Kahneman’s dual-system framework (Kahneman, 2011). Typ- 056 ically, a System Two, such as internet-pretrained vision-language models (VLMs) (Beyer et al., 057 2024; Karamcheti et al., 2024), is dedicated to slow high-level reasoning, generating intermediate 058 reasoning contents. Meanwhile, a System One, such as vision-language-action models (VLAs) (Kim 059 et al., 2024; Black et al., 2024; Bjorck et al., 2025), translates these intermediate contents into precise 060 low-level robot actions. However, this explicit decoupling results in both systems lacking mutual 061 awareness of each other’s capabilities; System Two may produce intermediate contents that System 062 One cannot execute (Shi et al., 2025). Furthermore, in real-world deployment, issues such as latency 063 may cause System Two to respond belatedly, providing outdated or irrelevant guidance.

064 We argue that achieving stronger reasoning-acting synergy demands a unified model. Indeed, the 065 recent trend towards unifying capabilities within single models is proving crucial for advancing 066 AI (Yao et al., 2023; OpenAI, 2025), and we believe this approach holds particular promise for 067 robot learning. In light of this, we introduce OneTwoVLA, a single unified vision-language-action 068 model capable of both acting (System One) and reasoning (System Two). Importantly, it adaptively 069 determines when to engage each mode. As shown in Fig. 1, OneTwoVLA triggers natural language 070 reasoning at key steps — like completing a subtask, detecting an error, or requiring human input 071 — producing outputs such as scene descriptions, task plans, historical summaries, and next-step 072 instructions. Otherwise, it generates actions informed by its most recent reasoning outputs. A key 073 advantage of this unified model is its natural support for co-training with vision-language data, sig- 074 nificantly enhancing reasoning and generalization. To facilitate this, we develop a scalable pipeline 075 for synthesizing high-quality, embodied reasoning-centric vision-language data.

076 Our extensive experiments validate OneTwoVLA’s effectiveness, demonstrating its ability to inte- 077 grate diverse capabilities within a single model: **1) Long-horizon task planning:** OneTwoVLA rea- 078 sons to formulate, track, and dynamically adjust task plans based on execution feedback, signifi- 079 cantly outperforming flat VLA (by 30%) and dual-system VLA (by 24%) baselines. Vision-language 080 co-training further enables generalization to novel task instructions (e.g., planning coffee prepara- 081 tion for “Help me stay awake”). **2) Error detection and recovery:** OneTwoVLA detects execution 082 errors in real time, reasons about corrective strategies, and performs agile recovery actions. **3) Natu- 083 ral human-robot interaction:** OneTwoVLA adjusts actions immediately upon human intervention 084 and proactively seeks clarification when faced with ambiguity. **4) Generalizable visual ground- 085 ing:** OneTwoVLA exhibits superior understanding of spatial relationships, object attributes, and 086 semantic features, even generalizing to objects absent from its robot training data.

## 2 RELATED WORK

089 **Vision-Language-Action Models.** Initialized from pre-trained vision-language models 090 (VLMs) (Chen et al., 2023; Beyer et al., 2024; Liu et al., 2024a; Wang et al., 2024; Lu et al., 2024), 091 vision-language-action models (VLAs) (Driess et al., 2023; Brohan et al., 2023; Kim et al., 2024; 092 Black et al., 2024; Pertsch et al., 2025; Team et al., 2025; Bjorck et al., 2025; Wen et al., 2025; 093 Huang et al., 2025) have emerged as a promising approach for building general-purpose robots. 094 These VLAs, trained on large robot datasets (Mandlekar et al., 2018; Gupta et al., 2018; Dasari et al., 095 2019; Cabi et al., 2019; Fang et al., 2020; Brohan et al., 2022; Jang et al., 2022; Walke et al., 2023; 096 O’Neill et al., 2024; Khazatsky et al., 2024; Lin et al., 2024), can handle a wide range of real-world 097 manipulation tasks. However, these VLAs exhibit limited reasoning capabilities (Hu et al., 2023; Shi 098 et al., 2025; Bjorck et al., 2025), showing vulnerability when confronted with long-horizon tasks or 099 complex dynamic environments. Moreover, their generalization performance degrades substantially 100 when facing novel objects or instructions outside the training distribution (Kim et al., 2024; Black 101 et al., 2024). In contrast, our work enhances reasoning and generalization capabilities through a 102 unified model architecture and a co-training framework.

103 **Reasoning for Robot Control.** Previous works (Stone et al., 2023; Huang et al., 2023; Li et al., 104 2023a; Belkhale et al., 2024; Liu et al., 2024b; Shi et al., 2024; Zhi et al., 2024; Zhao et al., 2025; 105 Li et al., 2025) demonstrate that high-level reasoning can enhance low-level policy performance in 106 robot control. In particular, many studies (Ahn et al., 2022; Huang et al., 2024; Hu et al., 2023; Shi 107 et al., 2025; Team et al., 2025; Bjorck et al., 2025; Figure, 2025) explore dual-system frameworks, 108 where a foundation model (e.g., a VLM) serves as System Two to perform high-level reasoning, 109 while a low-level policy operates as System One to generate actions based on reasoning outputs.

---

108 **Algorithm 1** Inference Pipeline of OneTwoVLA

109 **Require:** VLA model  $\pi_\theta$ , language instruction  $\ell$

110 1:  $t \leftarrow 0$ ,  $I_{\text{ref}}^{1:n} \leftarrow$  initial image,  $R \leftarrow$  none

111 2: **while**  $R \neq$  “Task Finished” **do**

112 3:      $DT \sim \pi_\theta.\text{decide}(\cdot | I_t^{1:n}, I_{\text{ref}}^{1:n}, \ell, R)$

113 4:     **if**  $DT = [\text{BOR}]$  **then**

114 5:          $\hat{R} \sim \pi_\theta.\text{reason}(\cdot | I_t^{1:n}, I_{\text{ref}}^{1:n}, \ell, R)$

115 6:          $R \leftarrow \hat{R}$ ,  $I^{\text{ref}} \leftarrow I_t$

116 7:     **else if**  $DT = [\text{BOA}]$  **then**

117 8:          $A_t \sim \pi_\theta.\text{act}(\cdot | I_t^{1:n}, I_{\text{ref}}^{1:n}, \ell, R, s_t)$

118 9:         Execute  $A_t$

119 10:     **end if**

120 11:      $t \leftarrow t + 1$

121 12: **end while**

---

122 While this dual-system framework proves effective for accomplishing long-horizon manipulation  
 123 tasks, it inherently suffers from limitations such as the two systems lacking mutual awareness of  
 124 each other’s capabilities (Shi et al., 2025) as well as latency issues with System Two. Recently,  
 125  $\pi_{0.5}$  (Intelligence et al., 2025) employs a single model to predict a subtask before each action, but  
 126 this reasoning is simple and information-limited. If this inflexible paradigm generates extensive  
 127 reasoning at every step, it significantly impacts inference efficiency (Zawalski et al., 2024). To  
 128 mitigate this, ECoT-Lite (Chen et al., 2025) avoids producing reasoning during test time, but this  
 129 leads to degraded performance and prevents effective human-robot interaction. To address these  
 130 limitations, we propose a unified model capable of adaptively deciding when to reason versus when  
 131 to act, allowing for both informative reasoning and efficient execution. For related work on co-  
 132 training for robot learning, please refer to the Appendix B.

### 133 3 METHOD

134 In this section, we first introduce the framework of OneTwoVLA in Sec. 3.1, including its formu-  
 135 lation, adaptive inference, and model instantiation. We then describe how we curate robot data to  
 136 enable synergistic reasoning and acting in Sec. 3.2. Finally, we present our scalable pipeline for  
 137 synthesizing vision-language data enriched with embodied reasoning in Sec. 3.3.

#### 138 3.1 FRAMEWORK OF ONETWOLVA

139 **Problem Formulation.** The central problem investigated in this work is how to develop a robotic  
 140 control policy  $\pi_\theta$  capable of both reasoning and acting, with the critical ability to autonomously  
 141 decide at each timestep  $t$  whether to reason or act. Formally, the policy operates in two modes.  
 142 When in reasoning mode, the policy takes as input the current image observations from multiple  
 143 cameras  $I_t^1, \dots, I_t^n$  (denoted as  $I_t^{1:n}$ , where  $n$  is the number of cameras), the reference images from  
 144 the latest reasoning timestep  $I_{\text{ref}}^1, \dots, I_{\text{ref}}^n$  (denoted as  $I_{\text{ref}}^{1:n}$ , which introduces observation histories  
 145 to prevent ambiguous states), the language instruction  $\ell$ , and the latest reasoning content  $R$ . The  
 146 policy performs reasoning in the form of textual output, generating updated reasoning content  $\hat{R} \sim$   
 147  $\pi_\theta(\cdot | I_t^{1:n}, I_{\text{ref}}^{1:n}, \ell, R)$ . Sec. 3.2 provides further details on the specific content of this reasoning  
 148 process. In acting mode, the policy  $\pi$  additionally incorporates the robot’s proprioceptive state  $s_t$  and  
 149 generates an action chunk  $A_t$  based on the latest reasoning content:  $A_t \sim \pi_\theta(\cdot | I_t^{1:n}, I_{\text{ref}}^{1:n}, \ell, R, s_t)$ .

150 **Adaptive Inference of OneTwoVLA.** In Algorithm 1, we present the detailed process of how One-  
 151 TwoVLA autonomously decides whether to reason or act. We introduce two special decision tokens  
 152 ( $DT$ ): *beginning of reasoning* ([BOR]) and *beginning of action* ([BOA]). Given the prefix (com-  
 153 prising image observations  $I_t^{1:n}$ , reference images  $I_{\text{ref}}^{1:n}$ , instruction  $\ell$ , and the latest reasoning  
 154 content  $R$ ), the model first predicts either [BOR] or [BOA]. When [BOR] is predicted, the model  
 155 enters reasoning mode and generates textual reasoning content  $R$  until producing an *end of sentence*  
 156 ([EOS]) token. Conversely, when [BOA] is predicted, the model enters acting mode and directly  
 157 generates the action chunk  $A_t$ . This adaptive framework yields high inference efficiency: during  
 158 task execution, the model operates primarily in acting mode, invoking reasoning only at a few criti-  
 159 cal steps, which adds only minimal overhead. As shown in Fig. 2, for completing a long-horizon  
 160 task, OneTwoVLA achieves total times that match those of a flat VLA without language reason-  
 161 ing ( $\pi_0$  (Black et al., 2024)). In contrast, a Dual-System approach that “always reasons” (e.g., Hi




Figure 2: **Task completion times on Tomato-Egg.**  
 For experimental settings, see Sec. 4.1.

Robot (Shi et al., 2025), ViLa (Hu et al., 2023)) incurs significant latency due to extensive reasoning. Moreover, our framework inherently supports error recovery and human-robot interaction: when the policy detects an error (e.g., failing to grasp an object), it autonomously enters reasoning mode to determine a corrective strategy and execute agile recovery actions. When human interaction occurs, any interaction text will be consistently added to the language instruction  $\ell$  in subsequent steps.

**Model Instantiation.** OneTwoVLA is designed to be general, allowing most existing VLAs to be integrated with minimal modifications. For a specific instance, we employ  $\pi_0$  (Black et al., 2024) as the base VLA, which demonstrates strong performance across various tasks. The vision-language model of  $\pi_0$  auto-regressively generates textual reasoning during inference and is supervised via a cross-entropy loss during training. To model complex continuous action distributions, we inherit the action expert architecture from  $\pi_0$  and train it using a flow matching loss (Lipman et al., 2022; Liu, 2022). OneTwoVLA’s inference flow is detailed in Fig. 3. See Appendix F.2 for more training details.




Figure 3: **Inference flow** of OneTwoVLA in two modes.

### 3.2 CURATING ROBOT DATA WITH EMBODIED REASONING

Most existing robotic manipulation datasets consist primarily of observation-action pairs and lack associated reasoning information. To address this gap, we introduce a novel robot data format. For a given task, we first collect demonstration trajectories provided by human experts. Subsequently, each trajectory is segmented into a sequence of intervals. There are two types of intervals: *reasoning intervals*, which capture key steps requiring model reasoning (e.g., upon completing subtasks, detecting errors, or when human interaction is required), which we further annotate with textual reasoning content; and *acting intervals*, in which the model primarily learns to predict actions based on observations and the latest reasoning content. During training, we supervise the decision tokens according to the interval type and the freshness of reasoning content  $R$ . In reasoning intervals, the ground-truth decision token is `[BOR]` if the current reasoning  $R$  is stale (i.e., needs updating); once  $R$  has been updated, the ground truth becomes `[BOA]`. In acting intervals, the model always learns to predict `[BOA]`. See Appendix F.1 for more details.

Next, we elaborate on the embodied reasoning content. As shown in Fig. 4 left, it consists of four components: 1) a detailed *scene description*, primarily focusing on the locations of task-relevant objects; 2) a *high-level plan* that outlines the sequential steps to accomplish the task; 3) a concise *historical summary* to keep the model informed about the task’s progress; and 4) the immediate *next step* that the robot needs to execute. This comprehensive reasoning content encourages the model to understand the visual world, learn high-level planning, and track task progress. Furthermore, to equip the policy with error detection and recovery capabilities, we specifically collect and label robot data focused on recovery from failure states. To enable natural human-robot interaction, we annotate certain intervals of the demonstrations with interaction context (e.g., the robot’s question and the human’s answer shown in Fig. 4 left).

We design a two-stage *fully automated pipeline* for labeling reasoning intervals and generating their reasoning content. In the first stage, *interval annotation*, we predefine a high-level plan with  $K$  ordered subtasks  $P = (p_1, \dots, p_K)$ . From each demonstration, we uniformly subsample  $N = 32$  frames  $S = \{I_{t_n}\}_{n=1}^N$  and prompt Gemini 2.5 to identify reasoning intervals immediately after each subtask completion, yielding  $K + 1$  intervals  $\mathcal{R} = \{(r_j^s, r_j^e)\}_{j=0}^K$ . In the second stage, *reasoning content generation*, for each interval  $(r_j^s, r_j^e) \in \mathcal{R}$  we take its midpoint frame and denote it as  $\hat{I}_j$ . We then construct four reasoning fields: 1) a scene description  $D_j$  generated by Gemini from  $\hat{I}_j$ ; 2) the high-level plan  $P$ ; 3) a historical summary  $H_j = (p_1, \dots, p_j)$ ; and 4) the next step  $X_j = p_{j+1}$ . The tuple  $(D_j, P, H_j, X_j)$  serves as the reasoning content for interval  $j$ . This automated pipeline produces high-quality annotations. For example, in the Tomato-Egg task, 81.5% of intervals




Figure 4: **Left.** Example of *robot data* with reasoning content. The reasoning content comprises a scene description, a high-level plan, a historical summary, and the next-step instruction. Interaction texts (e.g., the robot question and the human answer) are appended after the instruction. **Right.** Examples of synthetic embodied reasoning-centric *vision-language data*. The top two examples illustrate visual grounding tasks, while the bottom two demonstrate long-horizon tasks. More examples are provided in Appendix E.

are judged correct by human annotators, and 83.3% of scene descriptions are deemed reasonable. Additional prompts and extensions to error recovery are provided in Appendix D.

### 3.3 SCALABLE SYNTHESIS OF VISION-LANGUAGE DATA WITH EMBODIED REASONING

The carefully curated robot data described in Sec. 3.2 allows the model to directly learn the desired task, but its size scales linearly with the costly human effort, making large dataset creation impractical. To endow our model with stronger generalization and the ability to cope with highly varied scenarios, we leverage off-the-shelf foundation models and design a fully scalable pipeline that synthesizes vision-language data enriched with embodied reasoning. This pipeline consists of three steps: 1) We prompt Gemini 2.5 Pro to generate diverse textual descriptions of tabletop layouts featuring common household items; 2) Based on these textual descriptions, we employ the text-to-image generation model FLUX.1-dev (Labs, 2024) to synthesize high-quality images depicting the tabletop layouts. We further augment the synthetic images by randomly applying fisheye distortion or compositing a robot gripper with adaptive brightness, making the visuals more closely resemble real robot observations; 3) Finally, we utilize Gemini again to generate task instructions and corresponding reasoning contents for each synthesized image. Through this pipeline, we automatically generated 16,000 data samples, with examples shown in Fig. 4 right.

The generated task instructions fall into two categories: 1) Visual grounding tasks (Shridhar & Hsu, 2018; Bhat et al., 2024; Kim et al., 2023), where the instruction implicitly refers to an object in the image through spatial relationships, attributes, or semantic features. The accompanying reasoning must reveal the object’s explicit name and, optionally, its location; 2) Long-horizon tasks, where the instruction describes an extended, multi-step objective. The reasoning must supply a high-level, step-by-step plan for completing the task. For part of the dataset, we additionally instruct Gemini to incorporate elements of human–robot interaction into the plan. A detailed quality analysis of the synthetic data, along with additional examples, is provided in Appendix E.

## 4 EXPERIMENTS

In this section, we evaluate OneTwoVLA through extensive real-world experiments, demonstrating its superior performance in versatile capabilities: long-horizon task planning (Sec. 4.1), error detection and recovery (Sec. 4.2), natural human-robot interaction (Sec. 4.3), and visual grounding (Sec. 4.4). Additionally, we show that co-training with our synthetic vision-language data yields generalizable behaviors and open-world visual grounding capabilities on unseen scenarios and tasks.




Figure 5: **Task illustrations and reasoning examples.** In the three leftmost columns, we present three challenging, long-horizon manipulation tasks. Completing these tasks requires not only planning abilities, but also error detection and recovery capabilities, as well as the the ability to interact naturally with humans. In the rightmost column, we demonstrate two tasks drawn from our experiments on generalizable planning. For every task, we include a sample of the model’s reasoning content. See Appendix D for additional reasoning examples.

#### 4.1 LONG-HORIZON TASK PLANNING

**Hardware.** We utilize two robot platforms. The primary platform consists of a single 7-DoF Franka arm equipped with a parallel jaw gripper. A wrist-mounted GoPro camera with fisheye lens provides wide field-of-view observations. Most of our experiments are conducted using this setup. Additionally, we employ a dual-arm platform featuring two 6-DoF ARX arms with three cameras (two wrist and one base), primarily for generalizable planning experiments. See Appendix H for further details.

**Long-horizon Tasks.** We design three challenging long-horizon tasks (shown in Fig. 5), each requiring the robot to understand the scene, plan accordingly, accurately track task progress, and generate precise actions throughout execution. We briefly describe these tasks here, with more details provided in Appendix C.1: 1) Tomato-Egg: The robot pours oil followed by tomato and egg liquid into a cooking machine. Once cooking completes, it uses a spoon to scoop the scramble onto a plate—a contact-rich action demanding fine precision. 2) Hotpot: Four plates containing different food items are placed on the table, and their relative ordering is randomized. The robot must sequentially dip beef and one vegetable type, precisely place them into a strainer, and finally lift the strainer. 3) Cocktail: The robot mixes one of three cocktails (Mojito, Mountain Fuji, or Vodka Sunrise), each requiring 3-4 steps of ingredient pouring. The robot must distinguish between nearly ten visually similar ingredients and pour accurately. For all tasks, the initial placement of all manipulated objects is randomized within a  $10 \times 10 \text{ cm}^2$  area.

**Baselines.** We compare OneTwoVLA against two baselines: 1) a state-of-the-art VLA model  $\pi_0$  (Black et al., 2024), which does not perform reasoning. To ensure fair comparison, we fine-tune  $\pi_0$  on the same dataset used for training OneTwoVLA; and 2) a dual-system approach inspired by Hi Robot (Shi et al., 2025), in which Gemini 2.5 Pro serves as the high-level System Two that decomposes complex instructions into sequences of atomic commands (i.e., the *next step* field in OneTwoVLA’s reasoning content). In practice, we invoke System Two at fixed intervals (i.e., effectively “always reasoning”). For System One, we annotate our dataset with atomic commands and fine-tune  $\pi_0$  to execute the commands produced by System Two.

**Experimental Results.** As shown in Fig. 6 left, OneTwoVLA achieves an average success rate of 87% across the three challenging tasks, outperforming  $\pi_0$  by 30% and the dual-system approach by 24%. OneTwoVLA consistently generates correct plans, accurately tracks task progress, and outputs precise actions. In contrast, lacking explicit reasoning and historical context,  $\pi_0$  sometimes loses track of its current step — such as staying stuck at the initial position when preparing Mojito or repeatedly picking up beef in the Hotpot task. We also observe that explicit reasoning facilitates more fine-grained action learning;  $\pi_0$  sometimes struggles to grasp ingredients precisely in Hotpot task or scoops too lightly in the Tomato-Egg task, whereas OneTwoVLA performs these delicate actions accurately. Regarding the dual-system approach, we found limitations arising




Figure 6: **Left: evaluation results on long-horizon tasks.** OneTwoVLA excels in long-horizon task planning compared to baselines. **Right: evaluation results on generalizable planning tasks.** By co-training with synthetic vision–language data, OneTwoVLA-VL further enhances its generalization to novel tasks. In both figures, each method is evaluated over 20 trials per task.

from the lack of mutual awareness between the two systems’ capabilities. System Two occasionally outputs atomic commands that are infeasible for System One to execute (e.g., instructing to add green onion in the Tomato–Egg task when none is present). Additionally, the significant inference latency of Gemini 2.5 Pro may prevent System Two from promptly updating its reasoning content, causing System One to encounter out-of-distribution states during execution.

**Generalizable Planning.** We investigate how co-training with vision–language (VL) data can improve OneTwoVLA’s ability to generalize in task planning. Specifically, we collect additional demonstration data for various atomic skills (e.g., pick, place, open, etc.) across two robot platforms. We then co-train OneTwoVLA on these robot data together with the with 16,000 VL samples synthesized by the pipeline described in Sec. 3.3. We denote this variant as **OneTwoVLA-VL**, while the model trained exclusively on robot data is denoted as OneTwoVLA. During evaluation, the policy receives instructions that *never appear* in the robot data (e.g., Fig. 5, last column). We test on four challenging tasks that demand deep commonsense reasoning (see Appendix C.2 for details).

As shown in Fig. 6 right, OneTwoVLA-VL exhibits strong generalization, transferring knowledge from VL data to robot control. The model proactively searches for non-visible cola by opening the refrigerator in *Get Icy Cola*, handles complex spatial relationships and occlusion by first removing fruits before retrieving the plate in *Empty Plate*, recognizes the need for a tool and uses a nearby stick to sweep distant objects within reach in *Tool Use*. Furthermore, it exhibits sophisticated scene-aware human intent understanding in *Prepare Drinks*, preparing coffee for “Help me stay awake”, kale juice for “I want something healthy”. In contrast,  $\pi_0$  executes random atomic skills when faced with such novel tasks, and OneTwoVLA without VL co-training produces entirely incorrect plans—both exhibiting only minimal generalizable planning abilities.

## 4.2 ERROR DETECTION AND RECOVERY

Recovering from mistakes is a critical capability for general-purpose robots. OneTwoVLA can detect errors in real-time, rapidly reason about recovery strategies, and subsequently generate corrective actions learned from collected robot recovery data. Table 1 presents a quantitative comparison of error recovery performance on two tasks: Hotpot and Tomato–Egg. In the Hotpot task, the robot occasionally fails to grasp the strainer due to misalignment. We therefore collect 200 demonstrations containing recovery actions (600 demonstrations in total). After retraining, OneTwoVLA reasons to retract, reposition to align with the strainer and try grasping again, subsequently succeeding in lifting it up. In contrast,  $\pi_0$  frequently ignores errors and continues to lift the gripper despite not having grasped the strainer. In the Tomato–Egg task, sometimes the oil bottle slips from the gripper while pouring; we collect 100 recovery demonstrations (200 in total). OneTwoVLA recognizes the error, reasons to adjust its grasp for increased firmness and retry the action. However, the dual-system approach fails to respond promptly due to latency issues. System Two only alerts that the oil bottle is not grasped after the robot has already reached the pouring pose, by which time recovery is hard because the robot has entered an out-of-distribution state.

|             | Hotpot | Tomato–Egg | Total                 |
|-------------|--------|------------|-----------------------|
| OneTwoVLA   | 5 / 6  | 3 / 4      | <b>8 / 10 (80.0%)</b> |
| $\pi_0$     | 3 / 7  | 5 / 7      | 8 / 14 (57.1%)        |
| Dual-System | 4 / 5  | 3 / 7      | 7 / 12 (58.3%)        |

Table 1: **Error Detection and Recovery Results.** Values are reported as # successful recoveries / # error occurrences.

378  
379

## 4.3 NATURAL HUMAN-ROBOT INTERACTION

380 Deploying robots in human-centric scenarios requires natural interaction with people. We conduct 10 human-robot interactions for OneTwoVLA and the Dual-System  
381 ( $\pi_0$  is omitted as it cannot generate language outputs) on the Hotpot and Cocktail tasks, with subtask-level interaction results shown in Table 2.  
382 Due to its adaptive nature and explicit reasoning process, OneTwoVLA is able to engage with humans in a natural way — seamlessly handling human interventions and proactively seek clarification when faced with ambiguities. For example, in the Hotpot task, when a human interrupts by requesting, “Could you also dip another vegetable for me?” OneTwoVLA immediately responds by clarifying, “Sure! Would you like green bok choy, enoki mushrooms, or cabbage?” In the Cocktail task, when the robot is preparing a Vodka Sunrise and the human interrupts with, “I don’t want orange vodka, I want lemon-flavored one,” OneTwoVLA immediately reasons that it needs to put down the orange vodka, retrieve the lemon vodka, and generate action sequences that align with the human’s intent. In contrast, the dual-system approach frequently loses context during interaction and struggles to maintain a coherent reasoning process, merely picking up the lemon vodka without continuing to prepare the cocktail in the example above.

383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431

**Generalizable Human-Robot Interaction.** We further evaluate generalization on the Hotpot and Cocktail tasks by testing 20 novel interaction scenarios that *never appear* in the robot training data. With vision–language data co-training, OneTwoVLA-VL achieves a 72.5% success rate. In the Hotpot task, OneTwoVLA-VL demonstrates proactive reasoning by asking “Which plate of meat would you like me to cook?” when given the instruction “cook meat” with two meat plates present. It also effectively handles unseen dynamic interruptions — when picking up enoki mushrooms, if interrupted with, “I don’t want enoki mushrooms, please cook some green bok choy instead,” it immediately switches to executing the new instruction. In contrast, OneTwoVLA without VL data co-training fail to interpret such unseen interaction commands, exhibiting poor generalization.

## 4.4 ENHANCED VISUAL GROUNDING

410 Grounding objects in language instructions to the visual world is a prerequisite for robots to accomplish more complex tasks. We categorize visual grounding into three key aspects (Shridhar & Hsu, 2018; Bhat et al., 2024; Kim et al., 2023; Shridhar et al., 2020): spatial relationships, object attributes, and semantic features. To validate OneTwoVLA’s effectiveness in these aspects, we design experiments where instruction following requires non-trivial object grounding capabilities. Furthermore, to demonstrate the impact of our synthetic vision-language data, we conduct experiments in open-world settings where diverse items and environments pose additional challenges. The specific experimental settings are described below (shown in Fig. 7):

411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431

1) Single-Env: Four objects are randomly arranged on a tabletop in a single environment. We collect 50 picking-up demonstrations for each object using the UMI (Chi et al., 2024) device, totaling 200 demonstrations. For testing, we perform 40 trials per method in the same environment using the same four objects. 2) Open-World: We collect demonstrations in 16 diverse in-the-wild environments, totaling 933 valid demonstrations using the UMI device. Each demonstration involves moving the gripper to a randomly selected object within the scene, collectively including 180 distinct household items. For testing, we evaluate each method across 8 unseen environments, testing

|             | Hotpot  | Cocktail | Total                 |
|-------------|---------|----------|-----------------------|
| OneTwoVLA   | 10 / 10 | 10 / 10  | <b>20 / 20 (100%)</b> |
| Dual-System | 8 / 10  | 5 / 10   | 13 / 20 (65%)         |

Table 2: **Human-Robot Interaction Results.** Each entry is reported as # successes / # interactions.




Figure 7: **Illustrations of visual grounding tasks.** In the Single-Env setting, we provide task instructions that require understanding of spatial relationships, object attributes, or semantic features. In the Open-World setting, we further evaluate the model’s *generalizable* visual grounding capabilities.

5 times per environment, each time randomly selecting one from 20 objects: 5 objects seen in robot data, 10 objects unseen in robot data but present in synthetic vision-language data, and 5 objects unseen in either dataset. In both settings, training and test instructions refer to target objects using their names or through spatial relationships, attributes, or semantic features. Our annotated reasoning explicitly identifies the target object’s name and includes additional information about it. We compare three methods:  $\pi_0$ , OneTwoVLA, and OneTwoVLA-VL. Both  $\pi_0$  and OneTwoVLA are trained exclusively on robot data, whereas OneTwoVLA-VL is additionally co-trained with 16,000 synthetic vision-language samples. Further details are provided in Appendix C.3.

**Explicit reasoning facilitates visual grounding.** In the Single-Env setting, as shown in Table 3, OneTwoVLA achieves a success rate of 78%, significantly outperforming  $\pi_0$ , which has a success rate of only 5%. In most cases, OneTwoVLA accurately interprets spatial relationships, object attributes, and semantic features described in the instructions, reasons about the correct object, and then successfully picks up the target object. In stark contrast,  $\pi_0$  consistently fails to comprehend the instructions, even when the target object is explicitly named.  $\pi_0$  typically extends the gripper forward aimlessly or randomly picks up the closest object. This clear performance gap demonstrates that explicitly learning to reason helps the model truly understand the visual world rather than attempting to find shortcuts to overfit actions. Moreover, we find that the reasoning content also aids the model in fitting actions, as evidenced by  $\pi_0$ ’s action mean squared error (MSE) on the validation set being 62% higher than OneTwoVLA’s.

**Reasoning-centric vision-language data enables generalizable visual grounding.** In the Open-World setting, OneTwoVLA-VL achieves a 73% success rate, significantly outperforming both OneTwoVLA and  $\pi_0$ . In most cases, OneTwoVLA-VL can correctly handle objects unseen in the robot data but present in vision-language (VL) data, effectively transferring commonsense knowledge from VL data to the robot policy. Remarkably, OneTwoVLA-VL generalizes even to novel objects that appear in neither the robot nor VL training data (e.g., Sprite, GoPro). We attribute this exceptional generalization capability to VL data co-training, which better activates web knowledge already encoded in the pretrained vision-language model. In contrast, OneTwoVLA and  $\pi_0$  frequently exhibit aimless reaching behaviors — even for objects present in the training data — indicating that they merely overfit to action training data without developing genuine understanding of the visual environment in this complex and diverse setting.

## 5 CONCLUSION, LIMITATIONS AND FUTURE WORK

In this paper, we present OneTwoVLA, a single unified model capable of both reasoning and acting, and adaptively switching between these two modes. This synergy is enabled by our meticulously designed framework and reasoning-enriched robot data curation. Moreover, we propose a scalable pipeline for synthesizing embodied reasoning-centric vision-language data to further enhance the model’s reasoning and generalization capabilities. Extensive experiments demonstrate OneTwoVLA’s superior performance across four key abilities: long-horizon task planning, error detection and recovery, natural human-robot interaction, and generalizable visual grounding.

There are several limitations that future work can address. First, OneTwoVLA relies on hand-selected heuristics to determine which steps require reasoning. Analogous to supervised fine-tuning (SFT) in LLMs, this strategy aligns the policy with a set of candidate reasoning steps and content; however, achieving more optimal reasoning likely requires reinforcement learning (RL). Future work could investigate RL-based training to further strengthen the reasoning capabilities of VLA models. Second, although our adaptive framework allows the model to reason only at a few critical steps during task execution, the robot still needs to pause for two to three seconds while reasoning occurs. Future research could explore the design of asynchronous architectures, enabling simultaneous reasoning and action generation. Third, we have not optimized action inference efficiency. As our single unified model scales to larger parameter sizes, the time cost of action inference may become a bottleneck. We anticipate that incorporating advanced inference techniques from the LLM literature could help accelerate action inference in future work. Finally, due to resource constraints, we only investigate the effect of high-quality synthetic vision-language data on VLA reasoning capabilities. Future work could explore the impact of vision-language data from various sources.

|              | Single-Env | Open-World |
|--------------|------------|------------|
| $\pi_0$      | 5%         | 3%         |
| OneTwoVLA    | 78%        | 8%         |
| OneTwoVLA-VL | 88%        | 73%        |

Table 3: **Evaluation results on visual grounding tasks.** OneTwoVLA exhibits strong visual grounding capabilities, and co-training with VL data further enhances its generalization.

486 REFERENCES  
487

488 Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea  
489 Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say:  
490 Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.

491 Michael L Anderson. Embodied cognition: A field guide. *Artificial intelligence*, 149(1):91–130,  
492 2003.

493 Suneel Belkhale, Tianli Ding, Ted Xiao, Pierre Sermanet, Quon Vuong, Jonathan Tompson, Yevgen  
494 Chebotar, Debidatta Dwibedi, and Dorsa Sadigh. Rt-h: Action hierarchies using language. *arXiv  
495 preprint arXiv:2403.01823*, 2024.

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

499 Vineet Bhat, Prashanth Krishnamurthy, Ramesh Karri, and Farshad Khorrami. Hifi-cs: Towards  
500 open vocabulary visual grounding for robotic grasping using vision-language models. *arXiv  
501 preprint arXiv:2409.10419*, 2024.

502 Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi Fan,  
503 Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, et al. Gr0ot n1: An open foundation model  
504 for generalist humanoid robots. *arXiv preprint arXiv:2503.14734*, 2025.

505 Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo  
506 Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al.  $\pi_0$ : A vision-language-action flow  
507 model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.

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

511 Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski,  
512 Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action  
513 models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.

514 Serkan Cabi, Sergio Gómez Colmenarejo, Alexander Novikov, Ksenia Konyushkova, Scott Reed,  
515 Rae Jeong, Konrad Zolna, Yusuf Aytar, David Budden, Mel Vecerik, et al. Scaling data-  
516 driven robotics with reward sketching and batch reinforcement learning. *arXiv preprint  
517 arXiv:1909.12200*, 2019.

518 William Chen, Suneel Belkhale, Suvir Mirchandani, Oier Mees, Danny Driess, Karl Pertsch,  
519 and Sergey Levine. Training strategies for efficient embodied reasoning. *arXiv preprint  
520 arXiv:2505.08243*, 2025.

521 Xi Chen, Xiao Wang, Lucas Beyer, Alexander Kolesnikov, Jialin Wu, Paul Voigtlaender, Basil  
522 Mustafa, Sebastian Goodman, Ibrahim Alabdulmohsin, Piotr Padlewski, et al. Pali-3 vision lan-  
523 guage models: Smaller, faster, stronger. *arXiv preprint arXiv:2310.09199*, 2023.

524 Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake,  
525 and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The Inter-  
526 national Journal of Robotics Research*, pp. 02783649241273668, 2023.

527 Cheng Chi, Zhenjia Xu, Chuer Pan, Eric Cousineau, Benjamin Burchfiel, Siyuan Feng, Russ  
528 Tedrake, and Shuran Song. Universal manipulation interface: In-the-wild robot teaching with-  
529 out in-the-wild robots. *arXiv preprint arXiv:2402.10329*, 2024.

530 Sudeep Dasari, Frederik Ebert, Stephen Tian, Suraj Nair, Bernadette Bucher, Karl Schmeckpeper,  
531 Siddharth Singh, Sergey Levine, and Chelsea Finn. Robonet: Large-scale multi-robot learning.  
532 *arXiv preprint arXiv:1910.11215*, 2019.

533 Ria Doshi, Homer Walke, Oier Mees, Sudeep Dasari, and Sergey Levine. Scaling cross-embodied  
534 learning: One policy for manipulation, navigation, locomotion and aviation. *arXiv preprint  
535 arXiv:2408.11812*, 2024.

540 Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Ayzaan Wahid,  
 541 Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, et al. Palm-e: An embodied mul-  
 542 timodal language model. *arXiv preprint arXiv:2303.03378*, 2023.

543

544 Hao-Shu Fang, Chenxi Wang, Minghao Gou, and Cewu Lu. Grasnet-1billion: A large-scale bench-  
 545 mark for general object grasping. In *Proceedings of the IEEE/CVF conference on computer vision*  
 546 and pattern recognition

547 Figure. Helix: A vision-language-action model for generalist humanoid control, 2025. URL  
 548 <https://www.figure.ai/news/helix>.

549

550 Abhinav Gupta, Adithyavairavan Murali, Dhiraj Prakashchand Gandhi, and Lerrel Pinto. Robot  
 551 learning in homes: Improving generalization and reducing dataset bias. *Advances in neural infor-*  
 552 *mation processing systems*, 31, 2018.

553 Joey Hejna, Chethan Bhateja, Yichen Jiang, Karl Pertsch, and Dorsa Sadigh. Re-mix: Optimizing  
 554 data mixtures for large scale imitation learning. *arXiv preprint arXiv:2408.14037*, 2024.

555

556 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*  
 557 *neural information processing systems*, 33:6840–6851, 2020.

558 Yingdong Hu, Fanqi Lin, Tong Zhang, Li Yi, and Yang Gao. Look before you leap: Unveiling the  
 559 power of gpt-4v in robotic vision-language planning. *arXiv preprint arXiv:2311.17842*, 2023.

560

561 Haoxu Huang, Fanqi Lin, Yingdong Hu, Shengjie Wang, and Yang Gao. Copa: General robotic  
 562 manipulation through spatial constraints of parts with foundation models. In *2024 IEEE/RSJ*  
 563 *International Conference on Intelligent Robots and Systems (IROS)*, pp. 9488–9495. IEEE, 2024.

564 Huang Huang, Fangchen Liu, Letian Fu, Tingfan Wu, Mustafa Mukadam, Jitendra Malik, Ken Gold-  
 565 berg, and Pieter Abbeel. Otter: A vision-language-action model with text-aware visual feature  
 566 extraction. *arXiv preprint arXiv:2503.03734*, 2025.

567

568 Wenlong Huang, Fei Xia, Dhruv Shah, Danny Driess, Andy Zeng, Yao Lu, Pete Florence, Igor  
 569 Mordatch, Sergey Levine, Karol Hausman, et al. Grounded decoding: Guiding text generation  
 570 with grounded models for embodied agents. *Advances in Neural Information Processing Systems*,  
 571 36:59636–59661, 2023.

572

573 Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess,  
 574 Adnan Esmail, Michael Equi, Chelsea Finn, Niccolò Fusai, et al.  $\pi_{0.5}$ : a vision-language-action  
 575 model with open-world generalization. *arXiv preprint arXiv:2504.16054*, 2025.

576

577 Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine,  
 578 and Chelsea Finn. Bc-z: Zero-shot task generalization with robotic imitation learning. In *Confer-*  
 579 *ence on Robot Learning*, pp. 991–1002. PMLR, 2022.

580

581 Daniel Kahneman. *Thinking, fast and slow*. macmillan, 2011.

582

583 Siddharth Karamcheti, Suraj Nair, Ashwin Balakrishna, Percy Liang, Thomas Kollar, and Dorsa  
 584 Sadigh. Prismatic vlms: Investigating the design space of visually-conditioned language models.  
 585 In *Forty-first International Conference on Machine Learning*, 2024.

586

587 Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth  
 588 Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty El-  
 589 lis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint*  
 590 *arXiv:2403.12945*, 2024.

591

592 Junghyun Kim, Gi-Cheon Kang, Jaein Kim, Suyeon Shin, and Byoung-Tak Zhang. Gvcci: Life-  
 593 long learning of visual grounding for language-guided robotic manipulation. In *2023 IEEE/RSJ*  
 594 *International Conference on Intelligent Robots and Systems (IROS)*, pp. 952–959. IEEE, 2023.

595

596 Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair,  
 597 Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source  
 598 vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.

594 Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024.  
 595

596 Boyi Li, Philipp Wu, Pieter Abbeel, and Jitendra Malik. Interactive task planning with language  
 597 models. *arXiv preprint arXiv:2310.10645*, 2023a.

598 Xinghang Li, Minghuan Liu, Hanbo Zhang, Cunjun Yu, Jie Xu, Hongtao Wu, Chilam Cheang,  
 599 Ya Jing, Weinan Zhang, Huaping Liu, et al. Vision-language foundation models as effective robot  
 600 imitators. *arXiv preprint arXiv:2311.01378*, 2023b.

601

602 Yi Li, Yuquan Deng, Jesse Zhang, Joel Jang, Marius Memmel, Raymond Yu, Caelan Reed Garrett,  
 603 Fabio Ramos, Dieter Fox, Anqi Li, et al. Hamster: Hierarchical action models for open-world  
 604 robot manipulation. *arXiv preprint arXiv:2502.05485*, 2025.

605 Fanqi Lin, Yingdong Hu, Pingyue Sheng, Chuan Wen, Jiacheng You, and Yang Gao. Data scaling  
 606 laws in imitation learning for robotic manipulation. *arXiv preprint arXiv:2410.18647*, 2024.

607 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching  
 608 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

609

610 Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction  
 611 tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recogni-  
 612 tion*, pp. 26296–26306, 2024a.

613 Peiqi Liu, Yaswanth Orru, Jay Vakil, Chris Paxton, Nur Muhammad Mahi Shafiullah, and Lerrel  
 614 Pinto. Ok-robot: What really matters in integrating open-knowledge models for robotics. *arXiv  
 615 preprint arXiv:2401.12202*, 2024b.

616

617 Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint  
 618 arXiv:2209.14577*, 2022.

619 Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren,  
 620 Zhuoshu Li, Hao Yang, et al. Deepseek-vl: towards real-world vision-language understanding.  
 621 *arXiv preprint arXiv:2403.05525*, 2024.

622 Abhiram Maddukuri, Zhenyu Jiang, Lawrence Yunliang Chen, Soroush Nasiriany, Yuqi Xie,  
 623 Yu Fang, Wenqi Huang, Zu Wang, Zhenjia Xu, Nikita Chernyadev, et al. Sim-and-real co-training:  
 624 A simple recipe for vision-based robotic manipulation. *arXiv preprint arXiv:2503.24361*, 2025.

625

626 Ajay Mandlekar, Yuke Zhu, Animesh Garg, Jonathan Booher, Max Spero, Albert Tung, Julian Gao,  
 627 John Emmons, Anchit Gupta, Emre Orbay, et al. Roboturk: A crowdsourcing platform for robotic  
 628 skill learning through imitation. In *Conference on Robot Learning*, pp. 879–893. PMLR, 2018.

629

630 Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhui Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng  
 631 Dai, Yu Qiao, and Ping Luo. Embodiedgpt: Vision-language pre-training via embodied chain of  
 632 thought. *Advances in Neural Information Processing Systems*, 36:25081–25094, 2023.

633

634 Soroush Nasiriany, Sean Kirmani, Tianli Ding, Laura Smith, Yuke Zhu, Danny Driess, Dorsa  
 635 Sadigh, and Ted Xiao. Rt-affordance: Reasoning about robotic manipulation with affordances. In  
 636 *CoRL 2024 Workshop on Mastering Robot Manipulation in a World of Abundant Data*, 2024.

637

638 OpenAI. Introducing 4o image generation, 2025. URL <https://openai.com/index/introducing-4o-image-generation/>.

639

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

644

645 Karl Pertsch, Kyle Stachowicz, Brian Ichter, Danny Driess, Suraj Nair, Quan Vuong, Oier Mees,  
 646 Chelsea Finn, and Sergey Levine. Fast: Efficient action tokenization for vision-language-action  
 647 models. *arXiv preprint arXiv:2501.09747*, 2025.

648

649 Lucy Xiaoyang Shi, Zheyuan Hu, Tony Z Zhao, Archit Sharma, Karl Pertsch, Jianlan Luo, Sergey  
 650 Levine, and Chelsea Finn. Yell at your robot: Improving on-the-fly from language corrections.  
 651 *arXiv preprint arXiv:2403.12910*, 2024.

648 Lucy Xiaoyang Shi, Brian Ichter, Michael Equi, Liyiming Ke, Karl Pertsch, Quan Vuong, James  
 649 Tanner, Anna Walling, Haohuan Wang, Niccolò Fusai, et al. Hi robot: Open-ended instruction  
 650 following with hierarchical vision-language-action models. *arXiv preprint arXiv:2502.19417*,  
 651 2025.

652 Mohit Shridhar and David Hsu. Interactive visual grounding of referring expressions for human-  
 653 robot interaction. *arXiv preprint arXiv:1806.03831*, 2018.

654 Mohit Shridhar, Dixant Mittal, and David Hsu. Ingress: Interactive visual grounding of referring  
 655 expressions. *The International Journal of Robotics Research*, 39(2-3):217–232, 2020.

656 Alexander Soare. Does diffusion policy produce multi-modal actions?, 2024. URL  
 657 [https://github.com/alexander-soare/little\\_experiments/blob/main/  
 658 action\\_multimodality.md](https://github.com/alexander-soare/little_experiments/blob/main/action_multimodality.md).

659 Austin Stone, Ted Xiao, Yao Lu, Keerthana Gopalakrishnan, Kuang-Huei Lee, Quan Vuong, Paul  
 660 Wohlhart, Sean Kirmani, Brianna Zitkovich, Fei Xia, et al. Open-world object manipulation using  
 661 pre-trained vision-language models. *arXiv preprint arXiv:2303.00905*, 2023.

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

666 J Varela Francisco, Thompson Evan, and Rosch Eleanor. The embodied mind: Cognitive science  
 667 and human experience, 1991.

668 Quan Vuong, Sergey Levine, Homer Rich Walke, Karl Pertsch, Anikait Singh, Ria Doshi, Charles  
 669 Xu, Jianlan Luo, Liam Tan, Dhruv Shah, et al. Open x-embodiment: Robotic learning datasets and  
 670 rt-x models. In *Towards Generalist Robots: Learning Paradigms for Scalable Skill Acquisition@  
 671 CoRL2023*, 2023.

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

675 Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu,  
 676 Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the  
 677 world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.

678 Junjie Wen, Yichen Zhu, Jinming Li, Zhibin Tang, Chaomin Shen, and Feifei Feng. Dexvla:  
 679 Vision-language model with plug-in diffusion expert for general robot control. *arXiv preprint  
 680 arXiv:2502.05855*, 2025.

681 Jonathan Yang, Catherine Glossop, Arjun Bhorkar, Dhruv Shah, Quan Vuong, Chelsea Finn, Dorsa  
 682 Sadigh, and Sergey Levine. Pushing the limits of cross-embodiment learning for manipulation  
 683 and navigation. *arXiv preprint arXiv:2402.19432*, 2024.

684 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.  
 685 React: Synergizing reasoning and acting in language models. In *International Conference on  
 686 Learning Representations (ICLR)*, 2023.

687 Wentao Yuan, Jiafei Duan, Valts Blukis, Wilbert Pumacay, Ranjay Krishna, Adithyavairavan Murali,  
 688 Arsalan Mousavian, and Dieter Fox. Robopoint: A vision-language model for spatial affordance  
 689 prediction for robotics. *arXiv preprint arXiv:2406.10721*, 2024.

690 Michał Zawalski, William Chen, Karl Pertsch, Oier Mees, Chelsea Finn, and Sergey Levine. Robotic  
 691 control via embodied chain-of-thought reasoning. *arXiv preprint arXiv:2407.08693*, 2024.

692 Qingqing Zhao, Yao Lu, Moo Jin Kim, Zipeng Fu, Zhuoyang Zhang, Yecheng Wu, Zhaoshuo Li,  
 693 Qianli Ma, Song Han, Chelsea Finn, et al. Cot-vla: Visual chain-of-thought reasoning for vision-  
 694 language-action models. *arXiv preprint arXiv:2503.22020*, 2025.

702 Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual  
703 manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.

704  
705 Peiyuan Zhi, Zhiyuan Zhang, Yu Zhao, Muzhi Han, Zeyu Zhang, Zhitian Li, Ziyuan Jiao, Baoxiong  
706 Jia, and Siyuan Huang. Closed-loop open-vocabulary mobile manipulation with gpt-4v. *arXiv*  
707 *preprint arXiv:2404.10220*, 2024.

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

711  
712 Minjie Zhu, Yichen Zhu, Jinming Li, Zhongyi Zhou, Junjie Wen, Xiaoyu Liu, Chaomin Shen, Yixin  
713 Peng, and Feifei Feng. Objectvla: End-to-end open-world object manipulation without demon-  
714 stration. *arXiv preprint arXiv:2502.19250*, 2025.

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751


752


753

754

755

756 APPENDIX  
757758 Please visit our anonymous website to view robot rollout videos: [759  
760 anonymous.github.io](https://onetwovla-).761 A LARGE LANGUAGE MODELS USAGE STATEMENT  
762763 We used large language models only for light grammar checking and language polishing. All content  
764 was written by the authors.  
765766 B MORE RELATED WORK  
767768 **Co-training for Robot Learning.** Co-training with data from diverse sources has been shown to  
769 benefit robot learning (Vuong et al., 2023; Driess et al., 2023; Li et al., 2023b; Nasiriany et al.,  
770 2024; Hejna et al., 2024; Yang et al., 2024; Doshi et al., 2024; Yuan et al., 2024; Maddukuri et al.,  
771 2025). In particular, several prior works (Brohan et al., 2023; Mu et al., 2023; Zhu et al., 2025;  
772 Zhou et al., 2025) explore co-training robot policies with action-free vision-language data alongside  
773 robot data, demonstrating improvements in policy generalization. However, these methods (Brohan  
774 et al., 2023; Mu et al., 2023; Zhou et al., 2025) typically either rely on existing vision-language  
775 datasets, which suffer from limited quality due to their significant domain gap from robot application  
776 scenarios; or manually collect vision-language datasets (Zhu et al., 2025), which are inherently  
777 limited in size and difficult to scale up. To address these limitations, we propose a scalable pipeline  
778 for synthesizing vision-language data rich in embodied reasoning. Our pipeline ensures both high  
779 quality and scalability, significantly enhancing policy’s reasoning and generalization capabilities.  
780781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809

810  
811  
812  
813 C TASKS AND EVALUATIONS  
814  
815  
816  
817  
818  
819820 In this section, we provide a detailed description of the tasks and evaluations.  
821  
822  
823  
824  
825  
826  
827  
828  
829830 Figure 8: **Execution processes of three long-horizon tasks:** Tomato-Egg, Hotpot, and Cocktail (ex-  
831 exemplified by Mountain Fuji preparation).  
832  
833  
834835 Fig. 8 shows the complete execution progress of the three long-horizon tasks. Detailed descriptions  
836 of these tasks are as follows:837 1) Tomato-Egg: The robot first pours oil, then tomato and egg liquid into a cooking machine.  
838 Once cooking is finished, the robot picks up a spoon hanging on a rack, scoops out the tomato-  
839 egg scramble, transfers it onto a plate, and finally places the spoon into the cooking machine. We  
840 observe that sometimes the robot fails to grip the oil bottle firmly enough, causing it to slip from the  
841 gripper. We collect dedicated recovery data for re-grasping the oil bottle more securely after it has  
842 slipped. This enables the robot to automatically perform this recovery if it encounters a bottle slip  
843 during testing. We collect 200 robot demonstrations for this task.844 2) Hotpot: Four plates containing beef, green bok choy, enoki mushrooms, and cabbage are placed  
845 on a table with randomized relative positions. A hotpot with a strainer is positioned to the right of  
846 the plates. For each test, the human instructs the robot to dip beef and one type of vegetable. The  
847 robot must accurately pick up the ingredients sequentially, place them in the strainer, wait for them  
848 to cook, and then lift the strainer. Notably, for OneTwoVLA and the dual-system approach, in 10 of  
849 the experiments, the initial instruction is only to dip the beef. While waiting for the beef to cook,  
850 the human interacts with the robot saying, “Could you also dip another vegetable for me?”, requiring  
851 the robot to ask, “Sure! Would you like green bok choy, enoki mushrooms, or cabbage?” Following  
852 the human’s specification, the robot then proceeds to dip the requested vegetable. This interaction  
853 step is omitted for  $\pi_0$  due to its lack of text output capabilities. Furthermore, we observe instances  
854 where the robot fails to grasp the strainer. To address this, we specifically collect recovery data for  
855 correcting misaligned grasps. This enables the robot to automatically perform this recovery if it fails  
856 to pick up the strainer during testing. We collect 600 robot demonstrations for this task.857 3) Cocktail: The robot is instructed to prepare one of three cocktails: Mojito, Mountain  
858 Fuji, or Vodka Sunrise. Each cocktail requires pouring 3-4 different ingredients. For OneT-  
859 woVLA and the dual-system approach, in 10 trials, the initial human instruction is general: “Make  
860 me a cocktail.” The robot must clarify by asking: “Which cocktail would you like?”, and then pro-  
861 ceed based on the human’s specific cocktail choice. This interaction step is again omitted for  $\pi_0$ .  
862 Additionally, during 3 separate Vodka Sunrise trials, the human interrupts with, “I don’t want  
863 orange vodka, I want lemon-flavored one,” requiring the robot to put down the orange vodka and  
864 pick up lemon vodka instead. We collect 100 robot demonstrations for each type of cocktail, totaling  
865 300 demonstrations.

864 C.2 GENERALIZABLE PLANNING TASKS  
865881 **Figure 9: Execution processes of four generalizable planning tasks:** Get Icy Cola, Empty Plate,  
882 Prepare Drinks (exemplified by kale juice preparation) and Tool Use.  
883

884 We collect 2,000 robot demonstrations using the single-arm Franka system and dual-arm ARX system.  
885 Each demonstration belongs to one category of atomic skill, including pick, place, move, open,  
886 close, and pour. The task instructions and corresponding reasoning contents for these demon-  
887 strations focus on short-horizon atomic skills. Training solely on this data limits the model’s generaliz-  
888 able long-horizon planning capabilities. OneTwoVLA overcomes this limitation through co-training  
889 with our synthesized embodied reasoning-centric vision-language data, which equips it to general-  
890 ize to previously unseen tasks. Fig. 9 shows the complete execution progress of these unseen tasks.  
891 Detailed descriptions of these tasks are as follows:

892 1) **Get Icy Cola:** The instruction is “Get me a can of icy cola.” The challenge is that a cola can  
893 is not directly visible in the scene. The robot must infer that “icy cola” implies the cola is stored in  
894 the fridge and therefore plan the necessary steps to open the fridge, locate the cola, and retrieve it.  
895 2) **Empty Plate:** The instruction is “Pass me an empty plate”. However, the plate in the scene  
896 is not empty, as it contains apples and grapes. The robot needs to remove each fruit from the plate  
897 before finally picking up the empty plate.  
898 3) **Tool Use:** The instruction is “Pick up the cocoa powder can, which is out of reach”. The  
899 primary difficulty here is that the target object is not within the robot’s direct reach. The robot must  
900 recognize the need for a tool (a nearby stick), plan to first grasp the stick, use it to sweep the distant  
901 cocoa powder can within reach, and only then proceed to pick up the can.  
902 4) **Prepare Drinks:** The robot needs to plan and prepare appropriate drinks based on user  
903 intent: such as coconut latte for “Help me stay awake,” kale juice for “I want something  
904 healthy,” and a blue mood cocktail for “I’m feeling down.” This task requires scene-aware user  
905 intent understanding capability.

907 C.3 VISUAL GROUNDING TASKS  
908

910 Task descriptions can be found in Sec. 4.4. In the Single-Env setting, each robot demon-  
911 stration is paired with 11 instruction-reasoning pairs. These instructions refer to target objects using  
912 their names (2 instances), spatial relationships (3 instances), attributes (3 instances), or semantic  
913 features (3 instances). In the Open-World setting, each demonstration includes a total of 17  
914 instruction-reasoning pairs, broken down as 2 using direct names, 5 using spatial relationships, 5  
915 using attributes, and 5 using semantic features. All instruction-reasoning pairs are first generated  
916 with Gemini 2.5 Pro and then verified by human annotators.


917 During testing, we evaluate each method 40 times in both settings. This consists of 10 tests for each  
918 reference type. Table 4 presents the experimental results broken down by these four types.

918 Here we list the objects used in visual grounding tasks. The Single-Env task uses four objects:  
 919 blue cube, eggplant toy, coconut water bottle, and black mouse. For the Open-World task evalua-  
 920 tion, we use the following objects (shown in Fig. 10):  
 921 1) 5 objects seen in robot data: flower, mouse, cardholder, tissue, and glasses case.  
 922 2) 10 objects unseen in robot data but present in synthetic vision-language data: globe, teddy bear,  
 923 straw hat, binoculars, trowel, croissant, map, magnifying glass, VR headset, lantern.  
 924 3) 5 objects unseen in either dataset: GoPro, Sprite, Starbucks Coffee, HDMI cable, Captain Amer-  
 925 ica model.

926 Fig. 11 displays the 16 training environments for the Open-World task, while Fig. 12 shows the 8  
 927 evaluation environments.  
 928

|              | Single-Env |         |           |          |       | Open-World |         |           |          |       |
|--------------|------------|---------|-----------|----------|-------|------------|---------|-----------|----------|-------|
|              | Name       | Spatial | Attribute | Semantic | Total | Name       | Spatial | Attribute | Semantic | Total |
| OneTwoVLA-VL | 10/10      | 8/10    | 8/10      | 9/10     | 35/40 | 8/10       | 6/10    | 7/10      | 8/10     | 29/40 |
| OneTwoVLA    | 10/10      | 5/10    | 8/10      | 8/10     | 31/40 | 2/40       | 0/10    | 1/10      | 0/10     | 3/40  |
| $\pi_0$      | 2/10       | 0/10    | 0/10      | 0/10     | 2/40  | 1/10       | 0/10    | 0/10      | 0/10     | 1/40  |

936 Table 4: **Experimental results for the visual grounding tasks.** Results are broken down by the four instruction  
 937 reference types: direct names, spatial relationships, object attributes, and semantic features.  
 938  
 939



959 Figure 10: **Objects for Open-World task evaluation.**  
 960  
 961  
 962  
 963  
 964  
 965  
 966  
 967  
 968  
 969  
 970  
 971



Figure 11: Training environments for **Open-World** visual grounding task.




Figure 12: Evaluation environments for Open-World visual grounding task.

1080 **D MORE REASONING EXAMPLES**  
10811082 Detailed OneTwoVLA reasoning examples during task execution are presented in this section. These  
1083 include examples for long-horizon task planning (Table 5), generalizable planning (Table 6), error  
1084 detection and recovery (Table 7), natural human-robot interaction (Table 8), Single-Env visual  
1085 grounding (Table 9), and Open-World visual grounding (Table 10).  
10861087 We provide prompt templates for annotating reasoning intervals (Fig. 13) and for generating reasoning  
1088 content (Fig. 14).  
10891089 Our template can readily extend to error detection/recovery and human-robot interaction (HRI) via  
1090 a plan-augmentation step. Given the original high-level plan  $P = (p_1, \dots, p_K)$ , we construct an  
1091 augmented plan  $\tilde{P}$  by inserting explicit substeps that represent anticipated recovery events and in-  
1092 teraction turns (e.g., “detect failure,” “retract,” “retry,” “ask question,” “receive answer”). We use  $\tilde{P}$   
1093 to prompt Gemini in both stages—(i) interval annotation and (ii) reasoning-content generation—so  
1094 that reasoning intervals and their fields align with these events.  
10951095 **Error detection and recovery.** Using  $\tilde{P}$  encourages the annotator to place reasoning intervals  
1096 around failure and recovery micro-steps (e.g., a failed grasp followed by a retraction). During as-  
1097 sembly, however, we map back to the original plan  $P$ . The reasoning tuple remains  $(D_j, P, H_j, X_j)$ ,  
1098 where  $H_j = \text{render}(p_1, \dots, p_j)$  over  $P$ . If a failure occurs on  $p_k$ , we set the next step to a retry ac-  
1099 tion conditioned on  $p_k$  (e.g.,  $X_j = \text{retry}(p_k)$ ), optionally preceded by a retraction, without altering  
1100  $P$  itself.  
11011101 **Human-robot interaction.** In  $\tilde{P}$  we insert explicit interaction substeps, such as the robot asking  
1102 a question and the human providing an answer (e.g., “Which cocktail should I make?” → “Mo-  
1103 jito”). During assembly, we extend the reasoning tuple to  $(D_j, P_j, H_j, X_j, Q_j, A_j)$ , where  $Q_j$  is  
1104 the robot’s question and  $A_j$  is the human’s answer (empty if no interaction). The plan is updated  
1105 around the dialogue: before the interaction we use a provisional plan  $P_j^{\text{pre}}$ , and after observing  $A_j$   
1106 we synthesize a conditioned plan  $P_j^{\text{post}}$  that replaces the remaining substeps; thus  $P_j = P_j^{\text{pre}}$  before  
1107 the dialogue and  $P_j = P_j^{\text{post}}$  after. The historical summary  $H_j$  records the newly acquired informa-  
1108 tion (i.e.,  $A_j$ ), and the next step  $X_j$  is conditioned on  $A_j$  (e.g., proceed to gather ingredients for a  
1109 Mojito). Interaction boundaries are marked as reasoning intervals during annotation.  
1110  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119  
1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187



**Instruction:** Dip the beef and the green bok choy.

**Reasoning:**

**Scene description:** The beef is located on the bottom-right plate. The green bok choy is on the top-right plate. The pot is in the center-right area of the table and contains a hotpot strainer.

**Plan:** 1. Put the beef into the hotpot strainer. 2. Put the green bok choy into the hotpot strainer. 3. Wait for the ingredients to cook and then pick up the hotpot strainer.

**What I have done:** Nothing.

**Now I need to:** Put the beef into the hotpot strainer.



**Instruction:** Make me a plate of tomato-egg scramble.

**Reasoning:**

**Scene description:** The oil bottle is located at the front of the table. The cooking machine is on the left side of the table.

**Plan:** 1. Pour the oil into the cooking machine. 2. Pour the tomatoes and egg liquid into the cooking machine. 3. Wait for the cooking process to finish, then pick up the spoon. 4. Scoop out the tomato-egg scramble, transfer it onto the plate, and place the spoon into the cooking machine.

**What I have done:** Nothing.

**Now I need to:** Pour the oil into the cooking machine.



**Instruction:** Make me a cup of Mountain Fuji cocktail.

**Reasoning:**

**Scene description:** The orange-flavored vodka is the second from the left on the bottom shelf. The Blue Curacao is the second from the right on the bottom shelf. The lemon juice is the second from the left on the top shelf. The yogurt is located in the middle of the top shelf.

**Plan:** 1. Pour the orange-flavored vodka into the glass. 2. Pour the Blue Curacao into the glass. 3. Pour the lemon juice into the glass. 4. Pour the yogurt into the glass.

**What I have done:** Nothing.

**Now I need to:** Pour the orange-flavored vodka into the glass.

Table 5: **Reasoning examples for long-horizon task planning.**

1188

1189

1190

1191

1192



**Instruction:** Give me an icy cola.

**Reasoning:**

Scene description: The cola is out of sight, maybe in the fridge.

Plan: 1. Open the fridge door. 2. Pick up the icy cola bottle and place it on the table. 3. Close the fridge door.

1193

1194

1195

1196

1197



**Instruction:** Pass me an empty plate.

**Reasoning:**

Scene description: The grapes and apple are on the plate.

Plan: 1. Pick up the grapes and place them on the table. 2. Pick up the apple and place it on the table. 3. Pick up the empty plate.

1198

1199

1200

1201

1202



**Instruction:** I'm exhausted — can you make me a drink to keep me awake?

**Reasoning:**

I need to make a coconut latte.

Plan: 1. Add thick coconut milk to the coffee cup. 2. Pour espresso into the cup. 3. Drop in ice cubes.

1203

1204

1205

1206

1207

**Instruction:** I've been eating clean and need a nutritious drink.

**Reasoning:**

I need to make a glass of kale juice with some kale leaves and ice cubes.

Plan: 1. Pour kale juice into the cup. 2. Put a kale leaf in the cup. 3. Add ice cubes to the cup.

1208

1209

1210

1211

1212



**Instruction:** I'm feeling kind of down today. Can you fix me a drink?

**Reasoning:**

I need to make a Blue Mood cocktail.

Plan: 1. Pour in white rum. 2. Add Blue Curacao. 3. Mix in lemon juice. 4. Drop in ice cubes.

1213

1214

1215

1216

1217

**Instruction:** Pick up the cocoa powder bottle that is out of reach.

**Reasoning:**

Scene description: The cocoa power bottle is out of reach. There is a wooden stick nearby.

Plan: 1. Grasp the stick. 2. Sweep the cocoa powder bottle closer with the stick and put down the stick. 3. Pick up the cocoa powder bottle.

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

Table 6: **Reasoning examples for generalizable planning.**

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295



**Instruction:** Dip the beef and the enoki mushrooms.

**Reasoning:**

Scene description: Some of the beef and the green bok choy are in the hotpot strainer. *But I fail to grasp the hotpot strainer due to misalignment.*

Plan: 1. Put the beef into the hotpot strainer. 2. Put the enoki mushrooms into the hotpot strainer. 3. Wait for the ingredients to cook and then pick up the hotpot strainer.

What I have done: 1. Put the beef into the hotpot strainer. 2. Put the enoki mushrooms into the hotpot strainer.

Now I need to: *Retract, reposition to align with the strainer, and try grasping again.*



**Instruction:** Make me a plate of tomato-egg scramble.

**Reasoning:**

Scene description: The oil bottle is located at the front of the table. The cooking machine is on the left side of the table. *But I discover that the oil bottle slips from the gripper.*

Plan: 1. Pour the oil into the cooking machine. 2. Pour the tomatoes and egg liquid into the cooking machine. 3. Wait for the cooking process to finish, then pick up the spoon. 4. Scoop out the tomato-egg scramble, transfer it onto the plate, and place the spoon into the cooking machine.

What I have done: Nothing.

Now I need to: *Retry to grasp the oil bottle with increased firmness.*

Table 7: **Reasoning examples for error detection and recovery.**

|      |  |
|------|--|
| 1296 | <b>Instruction:</b> Dip the beef for me.   |
| 1297 | <b>Human:</b> Could you also dip another vegetable for me?   |
| 1298 | <b>Robot:</b> Sure! Would you like green bok choy, enoki mushrooms, or cabbage?  |
| 1299 | <b>Human:</b> I want some green bok choy.  |
| 1300 | <b>Reasoning:</b>  |
| 1301 | Scene description: Some of the beef is in the hotpot strainer. <i>The green bok choy is on the top-left plate.</i>   |
| 1302 | Plan: 1. Put the beef into the hotpot strainer. 2. <i>Put the green bok choy into the hotpot strainer.</i> 3. Wait for the ingredients to cook and then pick up the hotpot strainer. |
| 1303 | What I have done: 1. Put the beef into the hotpot strainer.  |
| 1304 | Now I need to: <i>Put the green bok choy into the hotpot strainer.</i>   |
| 1305 | <b>Instruction:</b> Make me a cup of Vodka Sunrise.  |
| 1306 | <b>Human:</b> I don't want orange vodka, I want lemon-flavored one.  |
| 1307 | <b>Reasoning:</b>  |
| 1308 | Scene description: The grenadine syrup and orange juice have been poured into the glass. <i>The lemon-flavored vodka is the second from the right on the top shelf.</i>              |
| 1309 | Plan: 1. Pour the grenadine syrup into the glass. 2. Pour the orange juice into the glass. 3. <i>Put down the orange vodka, and retrieve the lemon vodka.</i>                        |
| 1310 | What I have done: 1. Pour the grenadine syrup into the glass. 2. Pour the orange juice into the glass.   |
| 1311 | Now I need to: <i>Put down the orange vodka, and retrieve the lemon vodka.</i>   |

Table 8: Reasoning examples for natural human robot interaction.

|      |   |
|------|---|
| 1316 | <b>Instruction:</b> Pick up the black mouse.  |
| 1317 | <b>Reasoning:</b>   |
| 1318 | I need to pick up the black mouse, which is to the left back of the eggplant, has a smooth, arching shape that fits the palm, and is a key component of a typical office workstation. |
| 1319 | <b>Instruction:</b> Pick up the object which is to the left back of the coconut water bottle.   |
| 1320 | <b>Reasoning:</b>   |
| 1321 | I need to pick up the blue cube, which is identified by its blue shade, and is a solid form useful for creating stable structures.  |
| 1322 | <b>Instruction:</b> Pick up the object which exhibits a purple coloration.  |
| 1323 | <b>Reasoning:</b>   |
| 1324 | I need to pick up the eggplant, which is to the right of the coconut water bottle, and is a common ingredient in ratatouille.   |
| 1325 | <b>Instruction:</b> Pick up the object which is a packaged drink.   |
| 1326 | <b>Reasoning:</b>   |
| 1327 | I need to pick up the coconut water bottle, which is to the left back of the blue cube, and has a cylindrical form.   |
| 1328 |   |
| 1329 |   |
| 1330 |   |
| 1331 |   |
| 1332 |   |
| 1333 |   |
| 1334 |   |
| 1335 |   |
| 1336 |   |
| 1337 |   |
| 1338 |   |
| 1339 |   |
| 1340 |   |

Table 9: Reasoning examples for Single-Env visual grounding.

|      |  |  |
|------|--|--|
| 1350 |  | <b>Instruction:</b> Bring me the Sprite.<br><b>Reasoning:</b> I need to pick up green Sprite, which is a model of a drink, located on the right side of the table.   |
| 1351 |  |  |
| 1352 |  | <b>Instruction:</b> Give me the object behind the seashell.<br><b>Reasoning:</b> I need to pick up the lantern to the left of the croissant, which is black with a metal structure, and holds light and can be carried.          |
| 1353 |  |  |
| 1354 |  |  |
| 1355 |  |  |
| 1356 |  |  |
| 1357 |  |  |
| 1358 |  |  |
| 1359 |  |  |
| 1360 |  | <b>Instruction:</b> Hand me the object which is black and round.<br><b>Reasoning:</b> I need to pick up the alarm clock sitting on the left side of the table.   |
| 1361 |  |  |
| 1362 |  |  |
| 1363 |  |  |
| 1364 |  |  |
| 1365 |  |  |
| 1366 |  | <b>Instruction:</b> I need something to get a closer look at this fine print.<br><b>Reasoning:</b> I need to pick up the magnifying glass, which is to the right of the red pepper, and has a wooden handle and a circular lens. |
| 1367 |  |  |
| 1368 |  |  |
| 1369 |  |  |
| 1370 |  |  |

Table 10: **Reasoning examples for Open-World visual grounding.**

1371  
1372  
1373  
1374  
1375  
1376  
1377  
1378  
1379  
1380  
1381  
1382  
1383  
1384  
1385  
1386  
1387  
1388  
1389  
1390  
1391  
1392  
1393  
1394  
1395  
1396  
1397  
1398  
1399  
1400  
1401  
1402  
1403

1404  
 1405 You are given:  
 1406 A step-by-step plan for a tabletop robot task, recorded by a wrist-mounted camera.  
 1407 A set of sampled frames from the video, named by integer frame IDs.  
 1408 Each image is introduced by a text line “Frame <id>” followed by the image itself.  
 1409  
 1410 Goal:  
 1411 Annotate “boundary intervals” between consecutive steps in the plan. A boundary interval is a  
 1412 contiguous range of frames [start\_id, end\_id] such that:  
 1413 1) From the visual evidence, the previous step has clearly been completed.  
 1414 2) The next step has not yet started.  
 1415 3) The tabletop is substantially visible across the interval (i.e., a relatively complete view of the  
 1416 workspace surface; avoid frames where the table is mostly occluded by the robot or out of view).  
 1417 4) Gripper-state criterion for all selected intervals: the gripper is open, the gripper holds no object  
 1418 (empty), and it is positioned relatively far from the tabletop (i.e., retracted or clearly above the surface  
 1419 rather than near contact). Prefer intervals where this is directly visible.  
 1420  
 1421 Additionally:  
 1422 Typical step structure involves: pick a tool/object, perform an operation, then place/put the tool  
 1423 elsewhere. Favor boundary intervals after “place/release” of the prior step and before the approach to  
 1424 the next target, where the gripper is open, empty, and away from the tabletop.  
 1425 If the gripper is temporarily occluded or out of view, accept frames where other cues strongly indicate  
 1426 open/empty/away (e.g., no contact, no object in jaws across adjacent frames). If not confirmable, keep  
 1427 the interval minimal and lower the confidence, but still follow the schema.  
 1428 Include a pre-task boundary interval that begins at frame 0 (implicit) and ends just before Step 1 visibly  
 1429 begins, while satisfying tabletop visibility and the gripper-state criterion.  
 1430 Include a post-task boundary interval that starts right after the final step is complete and ends at the  
 1431 last frame (implicit), while satisfying tabletop visibility and the gripper-state criterion.  
 1432 Intervals must be non-overlapping, ordered, and as long as possible while respecting the rules above.  
 1433 Prefer intervals with at least 2-3 frames where possible. If only a single frame satisfies the conditions,  
 1434 use a single-frame interval [k, k].  
 1435 If some step transition has no frames that both (a) show previous-step-complete and next-step-not-  
 1436 started and (b) maintain tabletop visibility and (c) meet the gripper-state criterion, pick the closest  
 1437 visually justifiable frame(s) and keep the interval minimal, but never violate chronological order.  
 1438  
 1439 Assumptions and cues:  
 1440 “Previous step completed” cues can include: the target object placed in its goal pose, gripper  
 1441 released/open, robot no longer manipulating that object, or the scene stabilized.  
 1442 “Next step not started” cues can include: the gripper has not yet approached/touched the next target,  
 1443 no new object interaction has begun, or no evident motion initiating the next step.  
 1444 The camera is on the wrist; favor frames where the tabletop surface is broadly visible and not heavily  
 1445 occluded by the manipulator.  
 1446  
 1447 Input you receive:  
 1448 Step plan (numbered 1..N).  
 1449 A list of frames; each is introduced by the line “Frame <id>” followed by its image. Frame IDs are  
 1450 chronological and unique. The first frame ID is 0; the last is given.  
 1451  
 1452 Output format (return ONLY this JSON; no extra text):  
 1453 

```
{"version": "v1", "total_frames": "", "steps": ["<step 1 text>", "<step 2 text>", "..."], "intervals": [{"label": "pre_step_1", "start": "", "end": "", "tabletop_visible": true, "confidence": "<float 0..1>"}, {"label": "between_step_1_2", "start": "", "end": "", "tabletop_visible": true, "confidence": "<float 0..1>"}, {"label": "...", "start": "", "end": "", "tabletop_visible": true, "confidence": "<float 0..1>"}, {"label": "between_step_(N-1)_N", "start": "", "end": "", "tabletop_visible": true, "confidence": "<float 0..1>"}, {"label": "post_step_N", "start": "", "end": "", "tabletop_visible": true, "confidence": "<float 0..1>"}]}
```

 1454

Figure 13: **Prompt template for annotating reasoning interval.**

```

1458
1459 You are given:
1460 A central frame image (the current interval's keyframe).
1461 The task's high-level plan as an ordered list of substeps  $P = (p_1, \dots, p_K)$ .
1462 The list of completed substeps up to the current interval.
1463 The next substep (or "Task Finished" if the task is finished).
1464
1465 Your goal:
1466 1) Produce the "scene description" for the current interval focusing strictly on task-relevant objects for
1467 the given plan and substeps.
1468 2) Output must be valid JSON only, matching the schema below.
1469
1470 Position frame of reference:
1471 "front" = nearest edge of the table to the camera/viewer.
1472 "back" = farthest edge of the table from the camera/viewer.
1473 "left/right" = viewer's left/right.
1474 You may also use "center/middle," "front-left," "front-right," "back-left," "back-right," and proximity
1475 phrases like "near the left edge," "near center," "near the back edge."
1476
1477 Rules:
1478 Write a concise scene summary (2–4 sentences).
1479 List only task-relevant objects on the table (tools, containers, parts, targets, intermediates).
1480 One sentence per object describing its absolute position on the table using the frame above.
1481 Do not hallucinate; if uncertain or partially occluded, lower confidence and note the uncertainty.
1482 If multiple similar objects exist, disambiguate with indices (e.g., "blue block #1", "blue block #2")
1483 based on spatial layout.
1484 If an object is not on the table or not visible, omit it from the objects list.
1485 Use brief, literal attributes (e.g., color, material, state like "open/closed," "full/empty").
1486 Output must be valid JSON only. Do not include any extra text or formatting outside JSON.
1487
1488 Inputs:
1489 TASK_NAME: {TASK_NAME}
1490 HIGH_LEVEL_PLAN (ordered): {HIGH_LEVEL_PLAN_AS_LIST}
1491 COMPLETED_SUBSTEPS (up to current interval): {COMPLETED_SUBSTEPS_AS_LIST}
1492 NEXT_SUBSTEP: {NEXT_SUBSTEP_OR_DONE}
1493 IMAGE: [central frame image provided as input]
1494
1495 Output JSON schema:
1496 { "scene_summary": "string, 2–4 sentences describing the visible scene and its relevance to the task.", "objects": [ { "name": "string, concise object name (e.g., 'red mug', 'blue block #1')", "attributes": ["string", "..."], "position_sentence": "string, exactly one sentence stating the object's absolute position on the table (front/back/left/right/center + optional proximity/edge terms).", "position_tags": ["front|back|left|right|center", "optional tags like 'front-left', 'near-left-edge', 'near-back-edge'"], "confidence": 0.0 }, "uncertain_observations": ["string notes about any ambiguities, occlusions, or visibility issues (optional; omit if none)"] }
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511

```

Figure 14: **Prompt template for generating reasoning content (scene descriptions).**

1512 E SYNTHETIC VISION-LANGUAGE DATA EXAMPLES  
1513

1514 Our 16,000 synthetic images are entirely annotated by Gemini 2.5 Pro, without any human intervention.  
1515 For 6,000 of these images, we generate visual grounding tasks. Each of these images is  
1516 annotated with 17 instruction-reasoning pairs, with the instructions referring to objects using their  
1517 direct names (2 instances), spatial relationships (5 instances), attributes (5 instances), and semantic  
1518 features (5 instances). For the remaining 10,000 images, we annotate a long-horizon planning task  
1519 along with a corresponding high-level, step-by-step plan for task completion. We also attempt to  
1520 use GPT-4o for annotating our synthetic images but find its spatial understanding to be weak. We  
1521 therefore use Gemini 2.5 Pro, which demonstrates strong spatial reasoning capabilities.  
1522

1522 We present illustrative examples synthesized by our embodied reasoning-centric visual-language  
1523 data synthesis pipeline. Table 12 and Table 13 show samples of synthesized data for visual grounding  
1524 and long-horizon tasks, respectively, each including textual descriptions of tabletop layouts, synthe-  
1525 sized images corresponding to these descriptions, and the accompanying instruction-reasoning pairs  
1526 (for visual grounding example, we only show one pair for each of the four reference types). Fig. 15  
1527 illustrates the effects of applying fisheye distortion or compositing a robot gripper with adaptive  
1528 brightness to the synthetic images.

1529 Moreover, Fig. 16, 17, and 18 detail the specific prompts used with Gemini 2.5 Pro throughout our  
1530 pipeline: Fig. 16 shows the prompt for generating diverse tabletop descriptions, while Fig. 17 and 18  
1531 illustrate the prompts for generating visual grounding and long-horizon task instructions and their  
1532 associated reasoning, respectively.

1533 **Data Quality.** We randomly sample 50 instances from the synthetic dataset and identify the follow-  
1534 ing types of errors: 1) *Wrong Image*: the presence of hallucinated objects or implausible physics; 2)  
1535 *Wrong Text*: reasoning that contradicts the corresponding image or common sense. If the generated  
1536 image is incorrect, we do not evaluate its text. The results are summarized in Table 11.  
1537

|  | Wrong Image | Wrong Text | All Correct |
|--|-------------|------------|-------------|
|  | 2           | 8          | 40          |

1541 Table 11: **Frequency of errors and fully correct cases in 50 randomly sampled synthetic data instances.**  
1542

1566  
1567  
1568  
1569  
1570

**Tabletop layout description:** The entire scene is sharp and focused. From a 45-degree overhead view, observe a medium-sized, rectangular driftwood table with a natural, uneven texture. On it are displayed seashells, a coral sculpture, a nautical map under a magnifying glass, a vintage lantern, and a model sailboat, backed by a beach house and ocean view.

**Direct instruction:** Can you pass me the map?  
**Reasoning:** I need to pick up the map, located towards the center of the table.

**Spatial instruction:** Get me the object sitting on the back right area of the paper item.  
**Reasoning:** I need to pick up the miniature sailboat model, which is placed on the map near the top right corner of the map.

**Attribute instruction:** I need the item with the complex branching structure.  
**Reasoning:** I need to pick up the large piece of coral positioned at the left side of the table.

**Semantic instruction:** Could you pass me the device used for seeing small details much larger?  
**Reasoning:** I need to pick up the magnifying glass situated on the map near the bottom right corner of the map.

**Tabletop layout description:** The entire scene is sharp and focused. From a 45-degree overhead view, observe a large, oval, stone table with a rough texture, set outdoors in a mountain retreat. On it, a hearty meal of grilled trout, wild rice, steamed vegetables, a jug of water, and a lantern. The rugged outdoor setting is enhanced by the surrounding pine trees and the fading twilight.



**Direct instruction:** Can you hand me the lantern?  
**Reasoning:** I need to pick up the black lantern situated at the back left of the rock table.

**Spatial instruction:** Pass me the item directly to the left of the grilled trouts.  
**Reasoning:** I need to pick up the fork sitting to the immediate left of the grilled trouts.

**Attribute instruction:** Please give me the plate with the long, silvery food item.  
**Reasoning:** I need to pick up the plate containing the two cooked fish positioned in front of the vegetable bowl.

**Semantic instruction:** Hand me that grain-based side dish, please.  
**Reasoning:** I need to pick up the plate of wild rice located at the back left, to the left of the vegetable bowl.

**Tabletop layout description:** The entire scene is sharp and focused. From a 45-degree overhead view, observe a rectangular maple table with a honey finish, placed in a sunny conservatory. A watering can, a terracotta pot with blooming flowers, a gardening book, gloves, and a small birdhouse are charmingly scattered on it. Glass panels reveal a lush garden outside.



**Direct instruction:** Can you pass me the watering can?  
**Reasoning:** I need to pick up the green watering can positioned on the back left side of the table.

**Spatial instruction:** I want the item located on the front right side of the table.  
**Reasoning:** I need to pick up the pair of gardening gloves, situated on the front right area of the table surface.

**Attribute instruction:** I need the object that's open and has pages with text.  
**Reasoning:** I need to pick up the book, which is open and located between the watering can and the brown gloves.

**Semantic instruction:** Please pass me the item that could provide shelter for small birds.  
**Reasoning:** I need to pick up the wooden birdhouse, sitting on the back right side of the table.

Table 12: Examples of synthetic vision-language data for visual grounding tasks.

1620

1621

1622

1623

1624

1625

1626

1627

1628

1629

1630

1631

1632

1633

1634

1635

1636

1637

1638

1639

1640

1641

1642

1643

1644

1645

1646

1647

1648

1649

1650

1651

1652

1653

1654

1655

1656

1657

1658

1659

1660

1661

1662

1663

1664

1665

1666

1667

1668

1669

1670

1671

1672

1673



**Tabletop layout description:** The entire scene is sharp and focused. From a 45-degree overhead view, directly facing the cluttered table, observe a rustic wooden surface covered in vibrant fresh ingredients like tomatoes, peas, herbs, and spices. A cutting board, knife, spoon, and book are laid out, bathed in bright light, against a backdrop of softly visible shelves with plants.

**Instruction:** Prepare a Fresh Herb and Tomato Relish.

**Reasoning:**

Plan: 1. Chop the tomatoes using the knife and cutting board. 2. Chop the fresh basil and parsley. 3. Combine the chopped tomatoes, basil, and parsley in a bowl. 4. Add the green peas and the dried herbs and spices from the piles on the board to the bowl. 5. Mix all ingredients thoroughly using the wooden spoon.



**Tabletop layout description:** The entire scene is sharp and focused. From a 45-degree overhead view, directly facing the cluttered table, observe a low Japanese tea table with a woven tatami mat beneath it, showcasing a tea set, a small bonsai tree, and a scroll. The scene plays out indoors, in a serene, minimalist room, emphasizing peaceful tranquility.

**Instruction:** Prepare and serve tea.

**Reasoning:**

Plan: 1. Pick up the teacup and place it next to the teapot. 2. Fill the teapot with hot water. 3. Pour tea from the teapot into the teacup. 4. Place the teapot on the table. 5. Serve the teacup to the guest.



**Tabletop layout description:** The entire scene is sharp and focused. From a 45-degree overhead view, directly facing the cluttered table, observe a small potted plant beside an empty terracotta pot. A soil bag, a silver trowel, and various tools are scattered across the worn wooden surface. The scene takes place outdoors, in a garden full of flowers, bathed in warm afternoon light.

**Instruction:** Repot the small plant into the larger terracotta pot.

**Reasoning:**

Plan: 1. Remove the plant from its current small pot. 2. Add soil from the bag into the bottom of the larger terracotta pot using the trowel. 3. Place the plant in the center of the larger pot. 4. Fill the remaining space in the larger pot with soil from the bag using the trowel. 5. Use the trowel to gently firm the soil around the base of the plant.



**Tabletop layout description:** The entire scene is sharp and focused. From a 45-degree overhead view, directly facing the cluttered table, observe a white folding table at an outdoor market. Assorted fruits, vegetables, and price tags cover the table. The bright, natural light enhances the colors of the produce.

**Instruction:** Prioritize fresh produce to boost your daily vitamin intake.

**Reasoning:**

Plan: 1. Pick up some tomatoes and place them in the basket. 2. Pick up some lemons and place them in the basket. 3. Pick up some oranges and place them in the basket. 4. Carry the basket to checkout.

Table 13: Examples of synthetic vision-language data for long-horizon tasks.




Figure 15: **Augmentations for our synthetic images.** From left to right: original synthetic images, synthetic images with fisheye distortion, synthetic images with a robot gripper composed with adaptive brightness, and synthetic images with both fisheye distortion and compositing a robot gripper with adaptive brightness.

1728

1729

1730

1731

1732

1733

1734

1735

1736

1737

1738

1739

1740

1741

1742

1743

1744

1745

1746

1747

1748

1749

1750

1751

1752

1753

1754

1755

1756

1757

1758

1759

1760

1761

1762

1763

1764

1765

1766

1767

1768

1769

1770

1771

1772

1773

1774

1775

1776

1777

1778

1779

1780

1781

Create 30 detailed 50-word prompts that describe scenes from a 45-degree top-down view of a table. The table should have a clear description of its shape, size, texture, and color. On the table, place around five objects, describing each object in detail and their positions relative to each other (e.g., object A is placed above object B). The background and environment should be clearly defined, either indoor or outdoor, and the scene should be rich in detail. Ensure there is no blurriness or out-of-focus areas, and the lighting and atmosphere should enhance the realism.

Please ensure each of the following prompts is unique and creatively different, varying the table, objects, environment (like indoor or outdoor), lighting, and overall atmosphere.

Each prompt should start with "The entire scene is sharp and focused. From a 45-degree overhead view, observe ...", followed by a description of the table's \*\*COLOR\*\* (this could be diverse across different prompts), shape, texture, size, etc.

Use the following format to separate each prompt:

\*\*START Prompt <Prompt ID>\*\*

[Detailed description of the scene]

\*\*END Prompt <Prompt ID>\*\*

Repeat this process for 30 distinct prompts. Request to generate all at once.

Figure 16: **Prompt used to generate tabletop descriptions.**

1782  
 1783 In the provided image, you will notice several items placed on a table. Your task is to come up with 17  
 1784 different instructions based on these items. These tasks will be categorized into three types based on object  
 1785 properties: spatial, semantic, and attribute.  
 1786 Spatial pertains to the object's position in space (e.g., on top of the plate, to the right of the book, or at the  
 1787 bottom right corner of the table).  
 1788 Semantic refers to the object's general, high-level meaning (e.g., sushi is a type of Japanese food, a kettle is  
 1789 used for boiling water, a book is meant for reading, etc.).  
 1790 Attribute is concerned with the object's specific features or characteristics (e.g., a ball is round, a handle is  
 1791 made of wood, etc.).  
 1792 For the objects on the table in the image, your task is to create 17 instructions, which can either directly ask  
 1793 for an object or describe it using its spatial, semantic, or attribute properties (e.g., "pass me the item on the  
 1794 plate," "give me something that helps with drying hair," or "hand me the yellow object").  
 1795 Each task is essentially a "pick" task, but the instruction should sound natural and realistic.  
 1796 After giving the instruction, provide a more specific description that starts with "I need to pick up," and then  
 1797 clearly name the object, possibly with some additional spatial details to help locate it.  
 1798 When describing a location, try to be as accurate as possible. Avoid using vague descriptions such as "in  
 1799 the middle/center of the table," "near," "beside," or "next to," as these could apply to many objects.  
 1800 Instead, use precise relative positioning, such as "to the left front of an object," "on top of an object,"  
 1801 "between object A and object B," "to the right back of an object," or "behind an object."  
 1802 When giving instructions, avoid mentioning the specific name of the object and instead use pronouns like  
 1803 "item," "object," or "device."  
 1804 When providing attribute instructions, only list 1 or 2 properties of the object.  
 1805 Your Tasks:  
 1806 First, generating 2 tasks with direct references to the object name.  
 1807 Then, generate 5 tasks \*\*only\*\* related to spatial properties (focusing on the location of the objects).  
 1808 Next, generate 5 tasks related to semantic properties (focusing on the general meaning or purpose of the  
 1809 objects).  
 1810 Finally, generate 5 tasks related to attribute properties (focusing on specific features of the objects).  
 1811 For each task, follow this format:  
 1812 \*\*Start Task <task id>\*\*  
 1813 Instruction: ...  
 1814 I need to pick up ...  
 1815 \*\*End Task <task id>\*\*  
 1816 Separate these 4 types of tasks by  
 1817   ### Tasks Related to Spatial Properties  
 1818   ### Tasks Related to Semantic Properties  
 1819   ### Tasks Related to Attribute Properties

1820  
 1821 Figure 17: **Prompt used to generate visual grounding task instructions and reasoning.**  
 1822  
 1823  
 1824  
 1825  
 1826  
 1827  
 1828  
 1829  
 1830  
 1831  
 1832  
 1833  
 1834  
 1835

1836  
 1837     In the given image, there is a table with several items placed on it in a messy manner.  
 1838     Your task is to first imagine a long-horizon task based on the items in the image (such as organizing the  
 1839     table, making a sandwich, etc.). This task needs to be relatively long-term, meaning it should require about  
 1840     several steps to complete.  
 1841     The second step is to provide a plan, where each step is a brief action description (e.g., Pick up sth and  
 1842     place it somewhere, Close sth, Open sth, Move sth to somewhere, etc.).  
 1843     Output in the following format:  
 1844        \*\*Start Task\*\*  
 1845        Instruction: ...  
 1846        1.  
 1847        2.  
 1848        ...  
 1849        N.  
 1850        \*\*End Task\*\*  
 1851     If you cannot think of an interesting task, simply output "Fail to think of a plan."  
 1852     Note that the instruction and plan should be brief and precise.

Figure 18: **Prompt used to generate long-horizon task instructions and reasoning.**

1852  
 1853  
 1854  
 1855  
 1856  
 1857  
 1858  
 1859  
 1860  
 1861  
 1862  
 1863  
 1864  
 1865  
 1866  
 1867  
 1868  
 1869  
 1870  
 1871  
 1872  
 1873  
 1874  
 1875  
 1876  
 1877  
 1878  
 1879  
 1880  
 1881  
 1882  
 1883  
 1884  
 1885  
 1886  
 1887  
 1888  
 1889

1890 **F IMPLEMENTATION DETAILS**  
18911892 **F.1 ROBOT DATA INTERVALS**  
18931894 As mentioned in Sec. 3.2, we segment robot demonstrations into two types of intervals: *reasoning*  
1895 *intervals* and *acting intervals*. Below, we detail what OneTwoVLA learns in each interval type.  
18961897 1) *Reasoning intervals*, OneTwoVLA learns to:  
18981899 

- 1900 • Predict [BOR] and the updated reasoning content  $\hat{R}$  based on the latest reasoning content  $R$ .
- 1901 • Predict [BOA] and actions based on the updated reasoning content  $\hat{R}$ .
- 1902 • Predict actions based on the latest reasoning content  $R$  without supervising [BOA]. This is to  
1903 prevent incorrect action prediction if the model fails to update the reasoning promptly during  
1904 deployment.

1905 2) *Acting intervals*, OneTwoVLA learns to:  
19061907 

- 1908 • Predict [BOA] and actions based on the latest reasoning content  $R$ .
- 1909 • (Optional) Predict [BOR] based on outdated reasoning without supervising the reasoning content.  
1910 This is included because we observe that during deployment, the model sometimes fails  
1911 to enter the reasoning mode. Since predicting decision tokens is essentially a binary classifica-  
1912 tion problem, and *acting intervals* are typically significantly longer than *reasoning intervals*, the  
1913 model predominantly learns to predict [BOA], leading to an imbalanced classification problem.  
1914 This optional training helps to increase the proportion of [BOR] predictions.

1915 Additionally, it is important to note that *reasoning interval* during training is designed to encourage  
1916 the model to learn the reasoning process more effectively. In real-world deployment, the robot  
1917 only reasons at a small number of steps (rather than continuous intervals), ensuring that the overall  
1918 operational efficiency is almost unaffected.1919 **F.2 POLICY TRAINING**  
19201921 As shown in Sec. 3.1, we use  $\pi_0$  as our base model. For each task, we train the model for 30,000  
1922 steps on 8xH100 GPUs, requiring approximately 10 hours. We adopt training hyperparameters from  
1923  $\pi_0$ . We make two modifications to the original  $\pi_0$ ’s input. Firstly, we use the current image  $I_t$   
1924 and the reference image  $I_{\text{ref}}$  as image observations. We incorporate  $I_{\text{ref}}$  because the textual scene  
1925 descriptions in reasoning may become outdated as the task progresses (e.g., an object’s position  
1926 described relative to the gripper becomes invalid upon gripper movement). Including  $I_{\text{ref}}$ , which  
1927 corresponds to the image observation for the current reasoning content, helps prevent model confu-  
1928 sion that might arise from potentially outdated textual descriptions. Second, we input not only the  
1929 current robot proprioceptive states but also the proprioceptive states from 0.05 and 0.25 seconds ear-  
1930 lier. This temporal context allows the model to generate more consistent and smooth actions during  
1931 execution.1932 **F.3 DEPLOYMENT**  
19331934 In real-world deployment, we use the temporal ensemble (Zhao et al., 2023) technique to ensure  
1935 smooth action execution. Specifically, in acting mode, the policy generates temporally overlapping  
1936 action sequences every 0.2 seconds. At any given timestep, multiple predicted actions are averaged  
1937 using exponential weighting to determine the actual executed actions.1938 Table 14 lists the computation time for  $\pi_0$ , along with the computation time for OneTwoVLA in  
1939 acting mode for varying input token counts and in reasoning mode for varying output token counts,  
1940 all of which are tested while processing two image inputs on an NVIDIA 4090 GPU. In acting  
1941 mode, although OneTwoVLA has additional reasoning content as input and outputs an extra [BOA]  
1942 compared to  $\pi_0$ , this has minimal impact on computation time and remains well below 0.2 seconds,  
1943 thus execution efficiency is not affected in this mode. In reasoning mode, when the reasoning token  
1944 count is low (less than 20 tokens), execution efficiency is unaffected; however, when reasoning  
1945 content is lengthy (exceeding 100 tokens), the robot needs to pause for a few seconds. Nevertheless,

1944 reasoning only occurs at a few critical moments, resulting in minimal impact on overall execution  
 1945 efficiency. For example, in one trial of the Tomato-Egg task, the entire long-horizon task takes 183  
 1946 seconds, with reasoning occurring 5 times, totaling 16 seconds of reasoning time, which accounts  
 1947 for 8.7% of the total duration. Similarly, in one trial of the preparing Mountain Fuji task, the  
 1948 entire long-horizon task takes 135 seconds, with reasoning occurring 5 times, totaling 14 seconds of  
 1949 reasoning time, which accounts for 10.4% of the total duration.

1950

|                      | # input tokens | # output tokens | computation time |
|----------------------|----------------|-----------------|------------------|
| $\pi_0$              | 20             |                 | 0.082 s          |
| OneTwoVLA-Act-20     | 20             | 1               | 0.102 s          |
| OneTwoVLA-Act-200    | 200            | 1               | 0.104 s          |
| OneTwoVLA-Reason-20  | 200            | 20              | 0.853 s          |
| OneTwoVLA-Reason-100 | 200            | 100             | 2.346 s          |
| OneTwoVLA-Reason-200 | 200            | 200             | 4.361 s          |

1951

1952 Table 14: **Computation times of  $\pi_0$  and OneTwoVLA.**  $\pi_0$ ’s input tokens consist solely of instruction  $\ell$ .  
 1953 OneTwoVLA’s input tokens are typically longer, including instruction and latest reasoning content ( $\ell$  and  $R$ ).  
 1954 In acting mode (OneTwoVLA-Act rows), OneTwoVLA’s output token is a single [BOA]. While in reason-  
 1955 ing mode (OneTwoVLA-Reason rows), OneTwoVLA outputs [BOR] and updated reasoning content,  $\hat{R}$ . We  
 1956 showcase computation times when its output token length is 20, 100, and 200.  
 1957

1958

1959

1960

1961

1962

1963

1964

1965

1966

1967

1968

1969

1970

1971

1972

1973

1974

1975

1976

1977

1978

1979

1980

1981

1982

1983

1984

1985

1986

1987

1988

1989

1990

1991

1992

1993

1994

1995

1996

1997

## 1998 G OTHER FINDINGS




Figure 19: **Multi-modality task illustration.** Two cubes and two bottles are symmetrically placed on the table. When the instruction doesn’t specify grasping the left or right object, OneTwoVLA can reason to grasp either the left or the right object, producing multi-modal actions.

### G.1 ONETWOLVA PRODUCES MULTI-MODAL ACTIONS

In this section, we design experiments to show OneTwoVLA’s capability to produce multi-modal actions.

**Tasks and Evaluations.** Two identical cubes are symmetrically placed on a table, each with an identical bottle positioned symmetrically behind it. Using the UMI device, we collect 50 demonstrations for each of these four objects (totaling 200 demonstrations). Each demonstration instruction is either “Grasp the cube” or “Grasp the bottle,” without specifying left or right. During testing, the object positions and the robotic gripper’s initial pose remain fixed. Each method is tested 20 times per instruction.

**Comparative Methods.** 1) OneTwoVLA: For each demonstration, we explicitly include disambiguating reasoning content (e.g., specifying picking up the left or right object) to resolve the ambiguity. 2)  $\pi_0$ : The model receives the original instruction directly, without explicit disambiguation.


**Experimental Results.** As shown in Fig. 19, OneTwoVLA demonstrates multi-modal action capability by alternating between reasoning to grasp objects from either side. Specifically, in the “grasp cube” experiment, OneTwoVLA grasps the left cube 9 times and the right cube 11 times. In the “grasp bottle” experiment, it grasps the left bottle 8 times and the right bottle 12 times. OneTwoVLA achieves this balanced left-right performance because its reasoning process is probabilistic, which means the model can sample different decisions (such as whether to grasp from the left or right) based on predicted token probabilities, much like language models generate varied responses from the same input. In contrast, although flow matching (Lipman et al., 2022; Liu, 2022) or diffusion (Ho et al., 2020; Chi et al., 2023) algorithms theoretically enable multi-modality,  $\pi_0$  consistently selects only the right-side objects, exhibiting only unimodal behavior, similar to observations in some other studies (Soare, 2024). Additionally, the disambiguating reasoning content helps the model fit actions more accurately. This is evidenced by  $\pi_0$  occasionally failing to grasp the block, while OneTwoVLA consistently achieves precise grasps. Moreover,  $\pi_0$ ’s action mean squared error (MSE) on the validation dataset is 56% higher than OneTwoVLA’s. This interesting finding suggests that when training on large-scale, variable-quality robot datasets, detailed annotation of reasoning content may enhance action learning.

### G.2 ONETWOLVA PRODUCES REASONING-COMPLIANT ACTIONS

Our experiments show that the actions generated by OneTwoVLA consistently align with its reasoning, even when the reasoning itself is incorrect. This finding is similar to observations in previous

2052 work (Zawalski et al., 2024). For example, in the `Hotpot` task, if OneTwoVLA occasionally reacts  
 2053 incorrectly about food locations, it proceeds to reach toward those incorrect positions. Similarly,  
 2054 in the `Open-World` experiment, OneTwoVLA moves to the object specified in its reasoning,  
 2055 even if that object does not align with the instruction. This indicates that OneTwoVLA’s cognition  
 2056 and behavior are highly unified, showcasing synergistic reasoning and acting. Additionally, this  
 2057 interesting phenomenon may indicate that improving the model’s reasoning ability (e.g., through  
 2058 additional vision-language data, using more powerful VLM as the base model, or more precise reasoning  
 2059 annotations) may contribute to generating more appropriate actions.

## H HARDWARE SETUP



2083 **Figure 20: Robot platform overview.** We employ two robot platforms: a single-arm Franka system (left)  
 2084 and a dual-arm ARX system (right).

2085  
 2086 We utilize two robot platforms. The primary platform (Fig. 20, left) is a single 7-DoF Franka arm  
 2087 equipped with a Weiss WSG-50 parallel-jaw gripper. A wrist-mounted GoPro camera with fisheye  
 2088 lens provides wide-angle observations. The arm is mounted on a custom height-adjustable table that  
 2089 can be pushed by a person—while not autonomous, this mobility allows us to evaluate the policy  
 2090 beyond traditional laboratory environments. The action space is 7-dimensional (6-DoF end-effector  
 2091 pose plus gripper width). Expert demonstrations for this platform are collected using UMI (Chi  
 2092 et al., 2024).

2093 The second platform (Fig. 20, right) features two 6-DoF ARX arms with parallel-jaw grippers and  
 2094 a three-camera system (two wrist-mounted and one base-mounted). It also includes a holonomic  
 2095 wheeled base and a 1-DoF torso lift mechanism, though these components have not yet been utilized  
 2096 in our experiments. The resulting action space is 14-dimensional ( $2 \times 7$ ). Expert demonstrations  
 2097 are collected via teleoperation using a Meta Quest headset.

2106  
2107  
2108  
2109  
2110  
2111  
2112  
2113  
2114  
2115  
2116  
2117  
I FAILURE CASESInstruction: "Prepare a plate of tomato-egg scramble."  
Instruction: "Dip the beef and the green bok choy."  
Instruction: "Make me a cup of Vodka Sunrise."  
Instruction: "Pick up the small basketball toy."  
Figure 21: Failure cases of OneTwoVLA.

Despite the promising performance of OneTwoVLA, it still makes mistakes. Fig. 21 illustrates the main failure cases of OneTwoVLA. In the Tomato-Egg task, OneTwoVLA occasionally fails to grip the yellow plate containing tomato and egg liquid firmly enough, resulting in the plate being dropped (see the first column in Fig. 21). In the Hotpot task, OneTwoVLA sometimes misidentifies the location of the target ingredient. For instance, as shown in the Fig. 21 second column, the robot is instructed to pick up green bok choy but instead it attempts to pick up enoki mushrooms. The third column of Fig. 21 shows a case in Cocktail task, where OneTwoVLA fails to pour the orange juice accurately while preparing the Vodka Sunrise, causing the juice to spill. In the Open-world experiments, OneTwoVLA shows vulnerability when encountering objects that are not present in either the robot data or our synthesized vision-language data. For instance, as illustrated in the Fig. 21 fourth column, the robot consistently moves toward the chessboard despite being instructed to pick up the small basketball toy. We believe that training on larger robot datasets, as well as co-training with richer vision-language data, can further facilitate OneTwoVLA in learning fine-grained actions and improve generalization capabilities.