

ReViP: Mitigating False Completion in Vision-Language-Action Models with Vision-Proprioception Rebalance

Zhuohao Li^{1,2}, Yinghao Li^{1,2}, Jian-Jian Jiang¹, Lang Zhou^{1,2}, Tianyu Zhang^{3,2}, Jiadong Yin¹, Mu Lin¹, Yi-Lin Wei¹, and Wei-Shi Zheng^{1,4,5,6*}

¹ Sun Yat-sen University, China

² Shenzhen Loop Area Institute, China

³ Beijing Institute of Technology, China

⁴ Peng Cheng Laboratory, Shenzhen, China

⁵ Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, China

⁶ Guangdong Province Key Laboratory of Information Security Technology, China
lizhh268@mail2.sysu.edu.cn; wszheng@ieee.org

Project Page: <https://isee-laboratory.github.io/ReViP/>

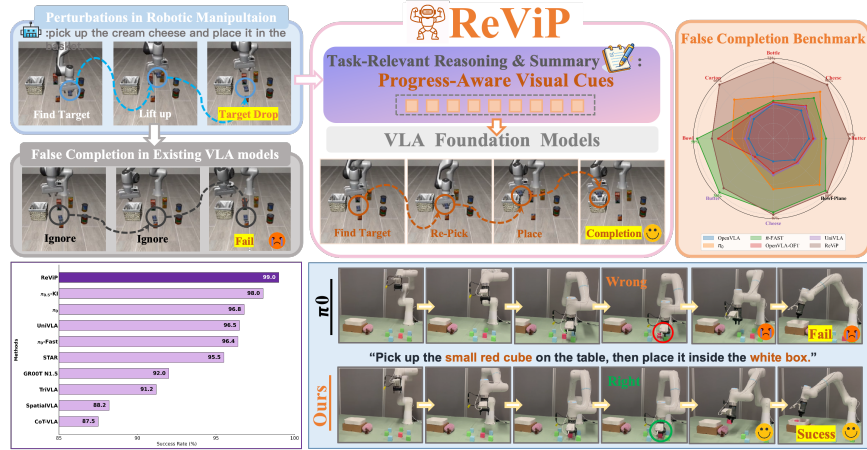


Fig. 1: From False Completion to Vision-Proprioception Rebalance. **Top:** False completion in VLA models occurs when policies prematurely fail despite unmet goals, often due to over-prioritizing internal state progression over visual feedback. *ReViP* addresses this by injecting progress-aware visual cues, rebalancing visual and proprioceptive dynamics for successful task completion; To evaluate and reproduce false completion, we introduce the **first False-Completion Benchmark Suite**, featuring eight tasks with controlled perturbations such as *Object Drop*, *Distractor Swap*, and *Relayout*. **Bottom:** Extensive experiments in both *simulation* (False-Completion, LIBERO, RoboTwin 2.0) and *real-world* settings show that *ReViP* significantly reduces false completion and improves task success rates, consistently outperforming existing methods.

*: corresponding author.

Abstract. Vision-Language-Action (VLA) models have advanced robotic manipulation by combining vision, language, and proprioception to predict actions. However, previous methods fuse proprioceptive signals directly with vision-language features, resulting in state-dominant bias and **false completions** despite visible execution failures. We systematically analyze this failure mode, attributing it to modality imbalance, where policies overly rely on internal state progression and underuse visual evidence. To address this, we introduce the first **False-Completion Benchmark Suite**, featuring eight tasks with three controlled perturbations (*Object Drop*, *Distractor Swap*, *Relayout*) to comprehensively evaluate false completion. Moreover, we propose **ReViP**, a novel VLA framework with **Vision-Proprioception Rebalance** to enhance visual grounding and robustness under perturbations. The key insight is to introduce auxiliary *progress-aware visual cues* to adaptively modulate the coupling between semantic perception and proprioceptive dynamics. Specifically, progress-aware visual cues are extracted by an external Task-Stage Observer, which performs task-relevant reasoning on real-time observations to drive task-stage feature-wise linear modulation, enhancing environmental awareness and mitigating state-driven errors. Extensive experiments show that ReViP effectively mitigates false completion and improves success rates over strong VLA baselines, achieving a **26%** gain over π_0 model on our suite, with gains extending to LIBERO, RoboTwin 2.0, and real-world evaluations.

Keywords: VLA Models · False Completion · Robotic Manipulation

1 Introduction

Recently, Vision-Language-Action (VLA) models [14, 15, 17, 22, 36, 37, 48, 51] have demonstrated impressive progress in robotic manipulation. Despite encouraging progress, previous methods typically encode visual-language inputs via Visual-Language Models (VLMs) [3, 10, 27, 31, 43, 47] and directly fuse proprioceptive signals [6, 21], leading to state-dominant bias and false completions even in the presence of clear visual failure cues. We attribute this behavior to modality imbalance, where policies over-rely on internal state while underusing visual evidence. For instance, when a target object drops during the execution, the policy continues the planned placement action towards the goal region, despite the target remaining visible and requiring retrieval. As shown in Fig. 1, the policy may prematurely halt or incorrectly declare success when the goal is unmet, causing the robot to terminate the operation without further action. We term this phenomenon “**False Completion**”. This behavior suggests that visual feedback is underutilized, with decision-making processes relying on proprioceptive cues, which conflicts with human common-sense task completion reasoning.

To investigate this, we first conduct systematic modality-controlled studies to comprehensively understand false completion and its underlying causes. We find that naively removing state inputs reduces performance, highlighting the necessity of rebalancing visual and proprioceptive signals. To address the lack of a

benchmark for false completion, we then introduce the first **False-Completion Benchmark Suite** on the LIBERO [30], enabling evaluation across VLA models. The suite features eight distinct tasks and three controlled perturbation sources (*Object Drop*, *Distractor Swap*, and *Relayout*), enabling a comprehensive evaluation of model performance in response to false completion across a diverse set of perturbations. Each perturbation addresses a unique aspect of model robustness: *Object Drop* evaluates the model’s response to dynamic disruptions, while *Distractor Swap* challenges its capacity to distinguish between targets and distractors. And *Relayout* tests the model’s adaptability to changes in the spatial configuration of the environment.

Additionally, to address false completion arising from modality imbalance, we present ReViP, a VLA framework that enhances visual grounding and robustness under perturbations by rebalancing vision perception with proprioceptive dynamics. To be specific, ReViP constructs a **Task-Stage Observer** based on an external VLM to process the current observation and instruction and perform task-relevant reasoning to extract **progress-aware visual cues** that reflect task progress and environmental state. Building on this, the **Task-Stage Enhancer** is proposed to inject these priors into the policy via a *Vision-Proprioception Feature-wise Linear Modulation*, adaptively rebalancing semantic perception and proprioceptive dynamics at the feature level. This design not only promotes feedback awareness but also mitigates state-driven errors under perturbations, such as object drops, distractor-induced mismatches, and environment relayout.

Finally, we conduct extensive evaluations across multiple simulation benchmarks, including false-completion assessment, the standard LIBERO, and the dual-arm RoboTwin 2.0 benchmark [9], spanning a diverse set of robotic manipulation tasks. In addition, real-world experiments are performed to further validate our hypothesis. The results reveal that ReViP significantly reduces false completions under previously unseen disturbances and achieves state-of-the-art (SOTA) task success rates, consistently outperforming existing methods on these challenging tasks. We will open-source our code and benchmark for the development of the community. Our contributions can be summarized as follows:

- We identify *false completion* as a critical failure mode in VLA models, arising from modality imbalance which leads to state-dominant bias. We systematically study this phenomenon through modality-controlled experiments to reveal its underlying causes.
- We introduce the *first False-Completion Benchmark Suite* tailored for comprehensively evaluating model performance under diverse perturbations (*Object Drop*, *Distractor Swap*, *Relayout*), emphasizing robustness across various false-completion settings.
- We propose *ReViP*, inspired by vision-proprioception rebalance, integrating a Task-Stage Observer for progress-aware feedback and a Task Stage Enhancer to adaptively strengthen the visual stream at the feature level.
- We demonstrate that ReViP significantly mitigates false completions and achieves SOTA task success rates, consistently outperforming strong VLA baselines across LIBERO, RoboTwin 2.0, and real-world tasks.

2 Related Work

Vision-Language-Action Models. VLA models [6, 14, 24, 36, 41, 44, 49, 51] connect perception to control by mapping pretrained vision-language representations to executable actions. Early approaches in the RT line (RT-1 [7], RT-2 [51], RT-H [4]) introduced action tokenization and trained transformer policies from scratch on large-scale demonstration corpora such as Open X-Embodiment [35] and DROID [2]. Recently, the π series (FAST [36], π_0 [6], and $\pi_{0.5}$ [17]) proposed heterogeneous co-training across robots and used flow matching-based decoders [29] for continuous action generation. Beyond these, a range of VLA variants [14, 20, 23, 26, 33, 37, 46] explore performing better robot manipulation, including structured reasoning, spatial grounding, and efficient architectures. For example, CoT-VLA [46] integrates chain-of-thought style reasoning with a lightweight policy to improve efficiency and the GR00T series [5] targets generalist humanoid robots by developing open foundation VLA models for complex bimanual and whole-body tasks. Few studies explicitly explore the modality mechanism in VLA models [45], and the phenomenon of false completion remains underexplored. In contrast, in this work, we focus on this modality imbalance and target false completion by explicitly rebalancing visual and proprioceptive signals.

VLA Models with External VLMs. Directly fine-tuning the VLM inside a VLA on robot data often degrades its pretrained multimodal understanding because of distribution shift and forgetting [6, 17, 50], which motivates keeping a strong external VLM for task understanding [13, 25]. Recent methods fall into two patterns. The first uses the VLM to decompose long-horizon instructions into subgoals, leaving the VLM only as a planner without exposing its internal semantic features to the VLA backbone [1, 16]. The second uses the VLM as a goal or step judge to decide success and trigger handcrafted retries or resets, which requires task-specific flows and incurs expensive restarts [11, 12]. While generic external success verification remains underexplored, frameworks like VLATest [40] use auxiliary evaluators to probe and stress-test VLA behavior across variations. Differing from them, ReViP injects internal VLM representations into the VLA via task-stage modeling feedback to rebalance vision and proprioception.

3 False Completion

In this section, we first define false completion, a failure mode where VLA policies prematurely proceed or terminate despite unmet task goals. We then explore the impact of modality imbalance through two key questions, and introduce our False-Completion Benchmark.

3.1 Problem Setup

False completion definition. Here, we formalize the setting of false completion. The observation I_t at time step t consists of synchronized images from a

set of cameras \mathcal{C} that include egocentric and exocentric views:

$$I_t = \{ I_t^{(c)} \in \mathbb{R}^{H_c \times W_c \times 3} \mid c \in \mathcal{C} \}. \quad (1)$$

Given (I_t, S_t, l) , where S_t is proprioception and l is the language instruction, a policy π_θ then predicts a block of n future actions (action chunk) as follows:

$$A_t = \pi_\theta [a_t, a_{t+1}, \dots, a_{t+n-1} \mid (I_t, S_t, l)] \quad (2)$$

Then, a termination flag $d_t \in \{0, 1\}$ is defined, where $d_t = 1$ is a completion declaration by the policy. Note that $d_t = 1$ is a conceptual indicator, reflecting a lack of significant action change over time, rather than a direct decision from the model itself. Let $G(I_t) \in \{0, 1\}$ be a visual goal predicate decided from the multi-view images. Following that, a task instance is considered truly completed only when both the policy declares completion and the visual predicate confirms goal satisfaction. However, when the policy issues a completion declaration while a visual predicate indicates the goal is unmet, we obtain a **false completion**:

$$\begin{aligned} \text{True Completion} &= [d_t = 1 \wedge G(I_t) = 1], \\ \text{False Completion} &= [d_t = 1 \wedge G(I_t) = 0]. \end{aligned} \quad (3)$$

3.2 False Completion in VLAs

In this paper, we attribute this failure mode to a modality imbalance within VLA models. Specifically, we identify a **state-dominant bias**: a tendency for the policy to prioritize internal proprioceptive sequences over external perceptual feedback. This bias causes the robot to blindly persist in task execution even when visual evidence signaling failure (e.g., an object being dropped) is present, ultimately leading to false completions.

Building on our hypothesis, to understand the effect of modality imbalance, we aim to answer two questions: **Q1**: *Is false completion caused by modality imbalance (state-dominant bias)* and **Q2**: *Can false completion be mitigated by simply removing state inputs*. To answer these, we conduct targeted modality-controlled experiments on both real robots.

(1) Modality imbalance leads

to false completion. We first investigate whether false completion arises from state-dominant modality imbalance by reproducing a state-vision mismatch in a real-world *cup pick-and-place* task (detailed in Section 5.4). By prematurely closing the gripper near the

object, we induce a state-vision mismatch: the state signal indicates the object “grasped” while the object is not visually grasped. In over 50 experimental trials, we record the robot’s behavioral response, categorizing actions as either *Proceed-to-Goal* or *Back-to-Object*. We mask proprioceptive signals by zeroing out the state input (*masking state*), forcing the model to rely solely on visual inputs.

Table 1: Real-world modality-controlled behavior. #Proceed-to-Goal and #Back-to-Object reflect the policy’s response to a state-vision mismatch with Success Rate for normal deployment.

Modality	#Proceed-to-Goal ↓	#Back-to-Object ↑	Success Rate (Normal) †
State-Enabled	46 / 50	4 / 50	70%
State-Masked	17 / 50	33 / 50	40%

As shown in Table 1, with state enabled (i.e., π_0), the robot predominantly follows proprioception, proceeding to the goal (*46 out of 50*) despite the visual mismatch, exhibiting false completion. This demonstrates that false completion is primarily driven by over-reliance on proprioceptive signals, i.e., state-dominant bias.

(2) Naively removing state helps false completion at a cost. Building on the above analysis that state-dominant bias causes false completion, we further explore whether removing state inputs can mitigate the issues. Table 1 shows that when state signals are masked, the robot exhibits markedly different behavior, returning to the object in 33 out of 50 trials. This confirms that removing state inputs increases reliance on vision and partially alleviates false completion. However, naive removal comes at a cost: overall success during standard task execution drops sharply from 70% to 40%, demonstrating that proprioceptive signals provide critical task-relevant information that vision alone cannot recover. **Naively removing state inputs fundamentally compromises the model’s capabilities, highlighting the need to rebalance visual and proprioceptive signals**, which motivates the design of ReViP (Section 4).

We also conduct modality-ablation experiments on simulation benchmarks to further validate our findings. Results are consistent with the real-world study, and detailed settings are provided in the supplementary materials.

3.3 False-Completion Benchmark Suite

To comprehensively evaluate false completion and in the absence of a suitable benchmark, we introduce the **False-Completion Benchmark Suite** on LIBERO.

Its design principle is simple: keep the instruction fixed while perturbing the scene at initialization or during execution, so that replaying demonstration-like proprioceptive trajectories is insufficient and policies must rely on current visual evidence.

The suite contains eight distinct tasks from three perturbation families (Fig. 2), each probing a different failure surface of state-dominant policies. All perturbations are deterministic (including a reproducible drop trigger) to ensure fair model comparisons. Full details of the benchmark settings are provided in the supplementary materials.




Fig. 2: Illustrative examples of the False-Completion Benchmark. Object Drop tests response to execution-time failures. Distractor Swap tests instance-level grounding under similar objects. Relayout tests spatial adaptation when targets and goals appear in new configurations.

Object Drop. This family directly stresses recovery from unexpected displacement (objects of varying sizes and cabinet-contact cases), where policies must detect failure and regrasp instead of continuing the original placement trajectory.

Distractor Swap. This family swaps target and distractor poses while keeping language unchanged, forcing instance-level visual grounding and exposing failures caused by state-driven trajectory replay.

Object Relayout. This family jointly relocates the target and goal region, breaking demonstration-specific spatial priors and requiring replanning from current visual configuration.

In summary, the benchmark treats false completion as a primary robustness challenge, with controlled perturbations that diagnose visual grounding, progress monitoring, and recovery behavior in a unified setting.

4 ReViP

To address false completion arising from state-dominant bias, we design ReViP, whose core idea is to rebalance visual and proprioceptive signals. Simply removing state inputs does not resolve this bias and can degrade task performance, motivating a more careful feature-level modulation. In this section, we introduce the Task-Stage Observer (TSO), which extracts progress-aware visual cues (Section 4.1), followed by the Task-Stage Enhancer (TSE), which adaptively rebalances visual-proprioceptive signals (Section 4.2). We then describe the action prediction process (Section 4.3) and summarize the overall structure of ReViP.

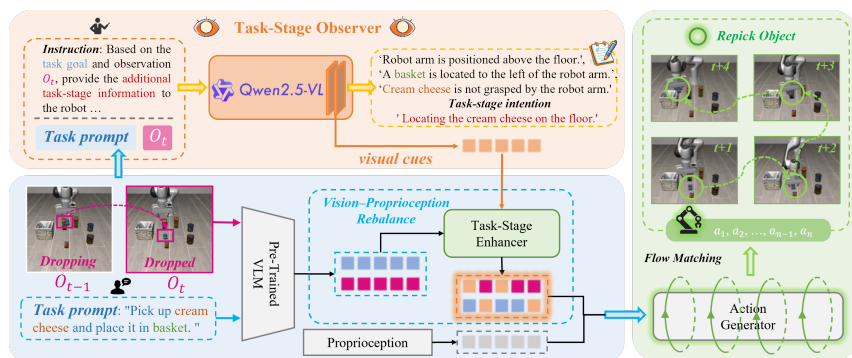


Fig. 3: Overview of ReViP. It consists of two stages: (above) Task-Stage Observer for cues extraction and (below) Task-Stage Enhancer for rebalancing. Through a frozen vision-language model, progress-aware visual cues are extracted from the observation O_t and task instructions. These cues are then injected into the VLA backbone by the Task-Stage Enhancer to adaptively rebalance the visual and proprioceptive streams.

4.1 Task-Stage Observer

To deal with visual information being underused, we propose TSO to convert observation I_t and instruction l into task-centric visual cues while explicitly fil-

tering task-irrelevant content. We employ a Qwen 2.5-VL [3] as the representative vision-language backbone for the TSO.

At time step t , TSO takes (I_t, l) as inputs, and performs task-relevant reasoning that first identifies the robot’s currently visible physical state and the spatial location and status of the task-relevant objects, then summarizes the immediate stage intention aligned with the given goal. The resulting cues provide explicit visual evidence for decision-making and generate a concise hypothesis for the subsequent task stage. As illustrated in the representative *object drop* scenario (Fig. 3), when the target *cream cheese* accidentally falls while the robot is moving toward the basket from I_{t-1} to I_t , TSO analyzes (I_t, l) to summarize critical task-stage evidence such as “*Cream cheese is not grasped by the robot arm*”. These cues indicate the target is no longer held by the gripper and localize it within the scene. By localizing the dropped object and generating a new stage intention, e.g., “*Locating the cream cheese on the floor.*”, the TSO guides the policy to initiate a re-pick action instead of continuing toward the basket.

To inject these discrete language cues into the VLA backbone, we adopt an LLM-based embedding extraction strategy [19, 28, 34, 38, 39] that converts the last layer hidden states into a compact continuous feature:

$$z_t = \text{Proj}\left(\text{Pool}(\text{TSO}_L(I_t, l))\right) \in \mathbb{R}^d, \quad (4)$$

where L denotes the last layer of the external VLM, $\text{Pool}(\cdot)$ aggregates the selected token representations into a single vector, and $\text{Proj}(\cdot)$ is a linear projection that maps this vector to the VLA semantic space of dimension d .

4.2 Task-Stage Enhancer

Building upon the TSO, we introduce TSE to strengthen task-relevant visual features and to rebalance the relative influence of vision and proprioception during action generation.

The TSE first translates the extracted task-relevant cues z_t into adaptively controlled signals via TS-FiLM. This mechanism ensures that decision-making is driven by current task-centric visual cues rather than an inertial reliance on internal state, which directly targets the false-to-true completion shift. Specifically, given z_t , the TSE produces feature-wise modulation parameters through a compact bottleneck mapping $h(\cdot)$, which can be expressed as:

$$[\gamma_t, \beta_t] = h(z_t), \quad (5)$$

where $\gamma_t, \beta_t \in \mathbb{R}^D$, and D is the token hidden size of the VLM backbone in VLA models. In the ReViP framework, let $P_t \in \mathbb{R}^{B \times S \times D}$ represent the concatenated vision and language prefix tokens, accompanied by a binary validity mask $M_t \in \{0, 1\}^{B \times S}$. We then apply token-wise TS-FiLM to the vision-language prefix before action generation. Therefore, TSE generates modulated prefix tokens \tilde{P}_t before action prediction:

$$\tilde{P}_t = \left(P_t + \alpha(\gamma_t \odot P_t + \beta_t)\right) \odot M_t, \quad (6)$$

where $\alpha \in \mathbb{R}_+$ is a learnable modulation factor and \odot denotes Hadamard product over tokens. This operation performs feature-level injection of progress-aware visual cues, effectively amplifying channels aligned with visual evidence while attenuating distractors that contribute to state-dominant bias.

4.3 Action Prediction.

The action head is trained to predict a conditional velocity field $v_\theta(\cdot, \cdot | \tilde{F}_t)$ via flow matching on the modulated prefix \tilde{F}_t . For a ground-truth action chunk A_t and random noise ε , we sample $\tau \sim \mathcal{U}(0, 1)$ and construct a noisy interpolation along the straight-line path:

$$v_\tau = (1 - \tau)A_t + \tau\varepsilon, \quad (7)$$

The target velocity (i.e., the time derivative of this path) is:

$$u_\tau = \frac{dv_\tau}{d\tau} = \varepsilon - A_t, \quad (8)$$

Above all, the training objective regresses the predicted velocity onto this target:

$$\begin{aligned} \tilde{F}_t &= \text{Fusion}(\tilde{P}_t, S_t), \\ \mathcal{L}_{\text{FM}} &= \|v_\theta(v_\tau, \tau | \tilde{F}_t) - u_\tau\|_2^2. \end{aligned} \quad (9)$$

When a perturbation such as an object drop occurs, z_t steers $[\gamma_t, \beta_t]$ to emphasize evidence and suppress irrelevant content, which guides the action generator to approach and regrasp the object instead of proceeding to the goal, reducing false completion.

Overall Framework. Integrating these components, ReViP processes the observation and instruction with the TSO, distilling them into progress-aware visual cues that capture relevant objects and stage intent. TSE injects these cues into the VLA backbone via TS-FiLM, enhancing VLM-derived features while adaptively rebalancing visual and proprioceptive signals. Finally, the modulated representation conditions a flow matching-based action generator to predict action chunks in parallel.

5 Experiments

5.1 Simulation Benchmark Experiment.

We evaluate ReViP in two key aspects: First, we test robustness using our False-Completion Benchmark Suite, which applies controlled perturbations to induce state-dominant bias. Second, we assess generalization on the LIBERO [30] and RoboTwin 2.0 benchmark [9], covering diverse scenes and tasks.

Implementation Details. Training is performed on $8 \times$ H100 GPUs (80GB) with a total batch size of 32 for 60k training steps. ReViP uses π_0 , a representative model, as the backbone policy, with TSO instantiated using Qwen2.5-VL-3B (**ReViP**). We also evaluate a stronger observer with Qwen2.5-VL-72B

(**ReViP***) to study how VLM capacity affects cue quality, providing a comprehensive validation mechanism. Note that TSO runs asynchronously, triggered once per action chunk, without disrupting real-time closed-loop control. Episodes succeed only if the instructed goal is reached within the step budget. Perturbation triggers are deterministic (e.g., reproducible drop rule) for fair comparison. Full evaluation protocols are provided in the supplementary materials.

Baselines. We compare ReViP with representative VLAs, including OpenVLA [22], OpenVLA-OFT [21], SpatialVLA [37], π_0 [6], π_0 -Fast [36], CoT-VLA [46], TriVLA [33], GR00T N1.5 [5], and UniVLA [8] in single-arm settings. For bimanual-arm setting, we further evaluate DP3 [42], RDT [32] and π_0 on the RoboTwin 2.0 benchmark [9].

5.2 Evaluation on False-Completion Benchmark









Quantitative Results. Table 2 presents a comparative evaluation of existing VLA methods on False-Completion benchmark. Both ReViP and ReViP* achieve the highest average success rate and the best average rank across all eight tasks, outperforming the strongest baseline, π_0 -Fast, by 18% and the baseline model, π_0 , by 26%. This aggregate improvement demonstrates that rebalancing visual and proprioceptive inputs with visual cues consistently enhances performance across diverse perturbations.

In Object Drop tasks, where policies must detect failures rather than follow an inertial plan, ReViP achieves 62.4% across five drop tasks, with ReViP* reaching 65.2%. This demonstrates that TSO enhances visual failure detection and re-planning, preventing the blind continuation of proprioception-driven progress. In Distractor Swap tasks, where object recognition and manipulation are confused, ReViP substantially improves over π_0 (15% \rightarrow 37%), with ReViP* reaching 39%, highlighting progress-aware visual cues’ role in disambiguating similar instances. In Relay task, ReViP achieves 84% with ReViP* improving to 88%, surpassing π_0 (70%) and UniVLA (20%). These results underscore ReViP*’s superior performance and validate the benefit of higher VLM capacities in enhancing cue quality, aligning with our objective of rebalancing vision and proprioception.

5.3 Evaluation on General Simulation Benchmark

Qualitative Results. Fig. 4 presents a representative comparison between ReViP and π_0 under the Object Drop perturbation. When the bottle slips from the gripper, π_0 continues its proprioception-driven trajectory, ignoring the visible failure and ultimately declaring a false completion. In contrast, ReViP immediately responds to the drop event, triggers a re-pick attempt, and guides the policy to replan, recover the object, and successfully complete the task. This qualitative behavior highlights the mechanism behind ReViP’s performance gains on False-Completion Benchmark: decisions are based on the current visual observation rather than on internal proprioceptive progress.

Table 2: Experimental Results of False-Completion Benchmark. Success rates (SR) and average ranks (RK) across eight tasks grouped into three perturbation sources: Object Drop, Distractor Swap, and Relayout. * indicates non-tabletop cabinet scenes (i.e., cabinet scenes). The best and second-best results are shown in **bold** and underline, respectively. Our method achieves the highest performance across the benchmark.

Methods	Object-Drop					Distractor-Swap		Relayout	Average	
	Butter	Cheese	Bottle	Carton	Bowl*	Butter	Cheese	Bowl-Plane	SR \uparrow	RK \downarrow
										
OpenVLA [22]	12%	30%	16%	6%	2%	0%	0%	12%	10%	7
π_0 [6]	24%	78%	26%	40%	20%	6%	24%	70%	36%	4
π_0 -Fast [36]	28%	62%	20%	16%	58%	36%	44%	84%	44%	3
OpenVLA-OFT [21]	16%	42%	18%	12%	35%	2%	10%	18%	19%	5
UniVLA [8]	18%	40%	16%	10%	6%	4%	12%	20%	18%	6
ReViP	46%	92%	66%	62%	46%	26%	48%	84%	59%	2
ReViP*	50%	96%	72%	66%	42%	<u>32%</u>	<u>46%</u>	88%	62%	1

Along with the numerical improvements in Table 2, these qualitative findings show that ReViP consistently transforms failure cases into successful completions through vision-proprioception rebalancing.

Performance on LIBERO. As shown in Table 4, we evaluate ReViP on the LIBERO benchmark, including LIBERO-Spatial, LIBERO-Object, LIBERO-Goal, and LIBERO-10 (also called LIBERO-Long), each with 10 tasks and 50 demonstrations. Both ReViP and ReViP* consistently outperform all other models, achieving the highest average task success rate of 96.7%, surpassing UniVLA and the backbone model π_0 . Specifically, on LIBERO-Spatial, ReViP reaches 99.0% success, effectively saturating the suite. On the more challenging LIBERO-10 suite which involves long-horizon manipulation across mixed task families, ReViP improves π_0 's success rate from 85.2% to 89.8%, and ReViP* further reaches 92.2%, demonstrating better stability over extended executions.

These results indicate that the proposed vision and proprioception rebalance with visual cues not only mitigates false completion but also improves performance on standard manipulation tasks.

Performance on RoboTwin 2.0 (Dual Arm). To further evaluate ReViP in dual-arm scenarios, we test

on the RoboTwin 2.0 benchmark, featuring contact-rich manipulation (Table 3). We consider four representative tasks, each trained with 50 demonstrations collected in clean environments. Evaluation under *hard mode conditions* includes 100 trials per task with domain randomization, lighting variations, object clutter, and table height shifts. ReViP achieves the highest success rates, outperforming RDT and π_0 , indicating the task-stage feedback mechanism scales effectively to dual-arm tasks and remains robust under severe environmental perturbations.

Table 3: Challenging Dual-Arm RoboTwin 2.0 (Easy→Hard) Results. Success rates (%), best in **bold**.

Task	DP3	RDT	π_0	ReViP
Place Object Stand	0	5	11	22
Place Bread Basket	1	2	4	15
Pick Dual Bottles	1	13	12	21
Put Bottles Dustbin	21	4	13	24
Average	6	6	10	21

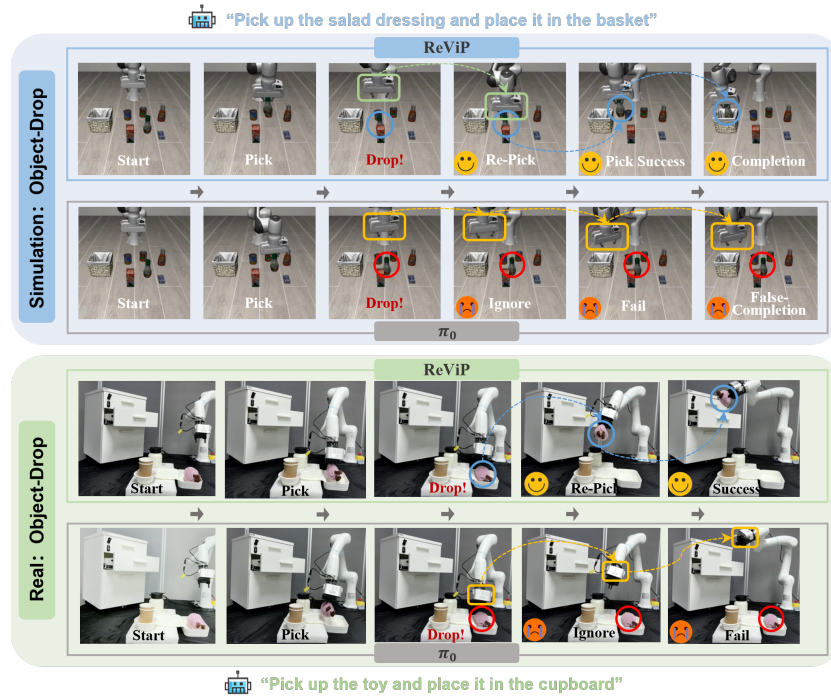


Fig. 4: Qualitative comparisons on our False Completion benchmark with Object Drop settings (above: simulation, below: real-world). For simulation, ReViP detects the object drop during execution, and successfully re-picks the target (salad dressing), achieving a true completion ☺, while π_0 fails to react to the clear visual failure and continues executing a state-dominant bias, resulting in false completion ☹. In the real world, ReViP can also retrieve the toy when it falls, but π_0 still fails to respond correctly.

5.4 Real-World Experiment

Robot platform. As shown in Fig. 5, our real-world experiments use a 6-dof ROKAE robotic arm equipped with a JODELL gripper. Two Orbbec Femto Bolt RGB-D cameras capture both first-person and third-person views.

During data collection, the robot is guided through kinesthetic teaching, with trajectories manually generated and recorded via interpolated playback.

Tasks and Datasets. Our datasets include two manipulation tasks with three different perturbations: Drop, Swap, and Relay. (1) *Pick-and-Place*: This task evaluates VLA robustness against visual distractions, with multiple distractor objects re-



Fig. 5: Real-world experiments use a robotic arm with grippers and cameras.

Table 4: Experimental Results of Simulation Experiments on the LIBERO benchmark. Success rates (SR) and average ranks (RK) across four task suites: Spatial, Object, Goal, and 10. The best and second-best results are shown in **bold** and underline, respectively. Our method achieves the highest performance across all suites.

Methods	Spatial		Object		Goal		10		Average	
	SR ↑	RK ↓	SR ↑	RK ↓	SR ↑	RK ↓	SR ↑	RK ↓	SR ↑	RK ↓
SpatialVLA [37]	88.2	9	89.9	9	78.6	10	55.5	11	78.1	10
π_0 [6]	96.8	3	98.8	1	95.8	3	85.2	6	94.2	5
π_0 -Fast [36]	96.4	5	96.8	5	88.6	7	60.2	10	85.5	8
CoT-VLA [46]	87.5	10	91.6	8	87.6	8	69.0	9	83.9	9
STAR [23]	95.5	6	98.3	3	95.0	5	88.5	4	94.3	4
$\pi_{0.5}$ -KI [17]	<u>98.0</u>	2	97.8	4	95.6	4	85.8	5	94.3	4
TriVLA [33]	91.2	8	93.8	6	89.8	6	73.2	8	87.0	6
GR00T N1.5 [5]	92.0	7	92.0	7	86.0	9	76.0	7	86.5	7
UniVLA [8]	96.5	4	96.8	5	95.6	4	<u>92.0</u>	2	95.2	3
ReViP	99.0	1	<u>98.6</u>	2	<u>96.2</u>	2	89.8	3	<u>95.9</u>	2
ReViP*	99.0	1	98.8	1	96.6	1	92.2	1	96.7	1

Table 5: Experimental Results of Real-World Experiments. Success rates (SR) are reported. The best result is shown in **bold**.

Methods	Object Placement		Drawer Manipulation			Average
	Cup → Box	+Drop	Toy → Drawer	+Close	+Drop	SR ↑
π_0 [6]	8/10	4/10	7/10	9/10	3/10	62%
π_0 -Fast [36]	7/10	3/10	7/10	8/10	3/10	56%
ReViP	10/10	8/10	9/10	10/10	7/10	88%

quiring accurate target localization and manipulation. (2) *Drawer-Manipulation*: This task tests the VLA’s ability to handle long-horizon sequences, involving dynamic scene interactions and a series of actions in the correct order. Further details are in the supplementary materials.

Baseline. We compare ReViP with two SOTA methods, π_0 and π_0 -Fast, to assess its performance in handling multi-distractor scenarios and its error-correction ability in long-horizon tasks.

Results. As shown in Table 5 and Fig. 4, ReViP achieves the highest success rate in both sub-tasks and overall performance. Under disturbances, ReViP detects the dropped object and autonomously re-grasps it, showing successful completion.

5.5 Extended Real-World Evaluation

We further evaluate ReViP under new, challenging real-world conditions to test robustness against false completion. The extended tasks include high-contrast backgrounds, alternative cameras (e.g., RealSense), small objects, and long-horizon multi-step sequences. These extended results (Table 6 and Fig. 1) confirm ReViP’s robustness under challenging settings. Detailed task designs, scene and hardware variations, and qualitative results are in the supplementary materials.

Table 6: Extended Real-World Evaluation: New, Challenging Settings Success rates (SR) are reported. The best result is shown in **bold**.

Methods	Cup Placement			Cube Placement		Drawer Manipulation	Average SR \uparrow
	Cup \rightarrow Box	+Drop	+Distractor	Cube \rightarrow Box	+Distractor		
π_0 [6]	5/10	4/10	3/10	4/10	2/10	4/10	34%
π_0 -Fast [36]	3/10	1/10	2/10	2/10	0/10	6/10	23%
ReViP	9/10	8/10	6/10	7/10	6/10	8/10	73%

5.6 Ablation Study

Effect of VLMs for TSO. To assess TSO’s robustness, we evaluate various VLM backbones (Table 7), including LLaVA-1.5-7B [31], InternVL2.5-8B [10], and Qwen2.5-VL [3]. Performance is consistent across VLMs, indicating TSO is not tied to a specific backbone. Stronger VLMs generally lead to higher success rates, with Qwen2.5-VL-72B achieving the best average SR (62%, +26%).

Table 7: Ablation Study: Effect of VLMs for TSO. Success rates (SR) on the False-Completion benchmark are reported. \clubsuit indicates non-tabletop cabinet scenes.

Methods	Object-Drop				Distractor-Swap		Relayout Bowl-Plane	Average SR \uparrow	
	Butter	Cheese	Bottle	Carton	Bowl \clubsuit	Cheese			
Baseline [6]	24%	78%	26%	40%	20%	6%	24%	70%	36%
LLaVA-1.5-7B [31]	32%	94%	70%	36%	64%	14%	40%	74%	55% (\uparrow 19%)
InternVL2.5-8B [10]	48%	80%	72%	64%	60%	16%	32%	80%	57% (\uparrow 21%)
Qwen2.5-VL-3B [3]	46%	92%	66%	62%	46%	26%	148%	84%	59% (\uparrow 23%)
Qwen2.5-VL-72B [3]	50%	96%	72%	66%	42%	32%	46%	88%	62% (\uparrow 26%)

Effect of FiLM factor α . Table 8 shows that fixed α performs reasonably across a range of values, while learnable α in ReViP adaptively selects modulation strength and achieves higher performance, supporting the effectiveness of our designs.

Table 8: Ablation Study: FiLM factor α . Success rates (SR) on the False-Completion-Drop benchmark are reported.

α	0	0.25	0.5	1.0	Learnable
SR \uparrow	38.6%	56.4%	62.4%	46.8%	65.2%

5.7 Comparison of Efficiency and Performance

Table 9 reports real-time efficiency and performance on our False-Completion benchmark using a single NVIDIA H100 GPU (30 warm-up runs and mean of the next 50). OpenVLA-OFT is slow and inaccurate, while π_0 is fast but lower in performance. ReViP achieves a favorable efficiency–performance balance with modest overhead. TSO is triggered once per action chunk and pipelined asynchronously, ensuring a control frequency of 16 Hz, sufficient for real-time operation with only slight latency compared to π_0 .

Table 9: Comparison of efficiency and performance. Real-time latency, control frequency, and success rate (SR) on the False-Completion benchmark are reported.

Methods	Latency (ms) \downarrow	Frequency (Hz) \uparrow	SR (%) \uparrow
OpenVLA-OFT [21]	175	5.71	19
π_0 [6]	44.6	22.42	36
ReViP	62.4	16.03	59

6 Conclusion

In this paper, we identify *false completion* as a critical failure mode in VLA models that conflicts with human commonsense task evaluation, and systematically study modality imbalance to reveal its causes. We further introduce the first False-Completion benchmark suite for evaluating false completion, and propose *ReViP*, a VLA framework that adaptively rebalances vision and proprioception using visual cues. Extensive simulation and real-world experiments demonstrate that ReViP effectively mitigates false completion, highlighting the benefits of vision-proprioception rebalancing.

References

1. Ahn, M., Brohan, A., Brown, N., Chebotar, Y., Cortes, O., David, B., Finn, C., Fu, C., Gopalakrishnan, K., Hausman, K., Herzog, A., Ho, D., Hsu, J., Ibarz, J., Ichter, B., Irpan, A., Jang, E., Ruano, R.J., Jeffrey, K., Jesmonth, S., Joshi, N.J., Julian, R., Kalashnikov, D., Kuang, Y., Lee, K.H., Levine, S., Lu, Y., Luu, L., Parada, C., Pastor, P., Quiambao, J., Rao, K., Rettinghouse, J., Reyes, D., Sermanet, P., Sievers, N., Tan, C., Toshev, A., Vanhoucke, V., Xia, F., Xiao, T., Xu, P., Xu, S., Yan, M., Zeng, A.: Do as i can, not as i say: Grounding language in robotic affordances. In: International Conference on Robot Learning (2023)
2. Alexander, Pertsch, K., Nair, S., Balakrishna, A., Dasari, S., Karamcheti, S., Nasiriany, S., Srirama, M.K., Chen, L.Y., Ellis, K., et al.: DROID: A large-scale in-the-wild robot manipulation dataset. In: robotics: science and systems (2024)
3. Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., Zhong, H., Zhu, Y., Yang, M., Li, Z., Wan, J., Wang, P., Ding, W., Fu, Z., Xu, Y., Ye, J., Zhang, X., Xie, T., Cheng, Z., Zhang, H., Yang, Z., Xu, H., Lin, J.: Qwen2.5-vl technical report (2025)
4. Belkhale, S., Ding, T., Xiao, T., Sermanet, P., Vuong, Q., Tompson, J., Chebotar, Y., Dwibedi, D., Sadigh, D.: Rt-h: Action hierarchies using language. In: Robotics: Science and Systems (2024)
5. Bjorck, J., Castañeda, F., Cherniadev, N., Da, X., Ding, R., Fan, L., Fang, Y., Fox, D., Hu, F., Huang, S., et al.: Gr00t n1: An open foundation model for generalist humanoid robots. arXiv preprint arXiv:2503.14734 (2025)
6. Black, K., Brown, N., Driess, D., Esmail, A., Equi, M., Finn, C., Fusai, N., Groom, L., Hausman, K., Ichter, B., Jakubczak, S., Jones, T., Ke, L., Levine, S., Li-Bell, A., Mothukuri, M., Nair, S., Pertsch, K., Shi, L.X., Tanner, J., Vuong, Q., Walling, A., Wang, H., Zhilinsky, U.: π_0 : A vision-language-action flow model for general robot control. Robotics: Science and Systems (2025)
7. Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., Gopalakrishnan, K., Hausman, K., Herzog, A., Hsu, J., Ibarz, J., Ichter, B., Irpan, A., Jackson, T., Jesmonth, S., Joshi, N.J., Julian, R., Kalashnikov, D., Kuang, Y., Leal, I., Lee, K., Levine, S., Lu, Y., Malla, U., Manjunath, D., Mordatch, I., Nachum, O., Parada, C., Peralta, J., Perez, E., Pertsch, K., Quiambao, J., Rao, K., Ryoo, M.S., Salazar, G., Sanketi, P.R., Sayed, K., Singh, J., Sontakke, S., Stone, A., Tan, C., Tran, H.T., Vanhoucke, V., Vega, S., Vuong, Q., Xia, F., Xiao, T., Xu, P., Xu, S., Yu, T., Zitkovich, B.: Rt-1: Robotics transformer for real-world control at scale. Robotics: Science and Systems (2023)

8. Bu, Q., Yang, Y., Cai, J., Gao, S., Ren, G., Yao, M., Luo, P., Li, H.: Univla: Learning to act anywhere with task-centric latent actions. In: *Robotics: Science and Systems* (2025)
9. Chen, T., Chen, Z., Chen, B., Cai, Z., Liu, Y., Li, Z., Liang, Q., Lin, X., Ge, Y., Gu, Z., et al.: Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation (2025)
10. Chen, Z., Wu, J., Wang, W., Su, W., Chen, G., Xing, S., Zhong, M., Zhang, Q., Zhu, X., Lu, L., et al.: Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In: *Conference on Computer Vision and Pattern Recognition*. pp. 24185–24198 (2024)
11. Du, Y., Konyushkova, K., Denil, M., Raju, A., Landon, J., Hill, F., de Freitas, N., Cabi, S.: Vision-language models as success detectors. In: *Conference on Lifelong Learning Agents* (2023)
12. Duan, J., Pumacay, W., Kumar, N., Wang, Y.R., Tian, S., Yuan, W., Krishna, R., Fox, D., Mandlekar, A., Guo, Y.: Aha: A vision-language-model for detecting and reasoning over failures in robotic manipulation. In: *International Conference on Learning Representations* (2025)
13. Hancock, A.J., Wu, X., Zha, L., Russakovsky, O., Majumdar, A.: Actions as language: Fine-tuning vlms into vlas without catastrophic forgetting (2025)
14. Hou, Z., Zhang, T., Xiong, Y., Duan, H., Pu, H., Tong, R., Zhao, C., Zhu, X., Qiao, Y., Dai, J., Chen, Y.: Dita: Scaling diffusion transformer for generalist vision-language-action policy. In: *International Conference on Computer Vision* (2025)
15. Huang, C.P., Wu, Y.H., Chen, M.H., Wang, Y.C.F., Yang, F.E.: Thinkact: Vision-language-action reasoning via reinforced visual latent planning. In: *arXiv preprint arXiv:2507.16815* (2025)
16. Huang, W., Xia, F., Xiao, T., Chan, H., Liang, J., Florence, P., Zeng, A., Tompson, J., Mordatch, I., Chebotar, Y., Sermanet, P., Brown, N., Jackson, T., Luu, L., Levine, S., Hausman, K., Ichter, B.: Inner monologue: Embodied reasoning through planning with language models. In: *International Conference on Robot Learning* (2023)
17. Intelligence, P., Black, K., Brown, N., Darpinian, J., Dhabalia, K., Driess, D., Esmail, A., Equi, M., Finn, C., Fusai, N., et al.: $\pi_{0.5}$: a vision-language-action model with open-world generalization. In: *Conference on Robot Learning* (2025)
18. James, S., Ma, Z., Rovick Arrojo, D., Davison, A.J.: Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters* (2020)
19. Jiang, Z., Meng, R., Yang, X., Yavuz, S., Zhou, Y., Chen, W.: Vlm2vec: Training vision-language models for massive multimodal embedding tasks. In: *International Conference on Learning Representations* (2024)
20. Khan, M.H., Asfaw, S., Iarchuk, D., Cabrera, M.A., Moreno, L., Tokmurziyev, I., Tsetserukou, D.: Shake-vla: Vision-language-action model-based system for bimanual robotic manipulations and liquid mixing. In: *arXiv preprint arXiv:2501.06919* (2025)
21. Kim, M.J., Finn, C., Liang, P.: Fine-tuning vision-language-action models: Optimizing speed and success (2025)
22. Kim, M.J., Pertsch, K., Karamcheti, S., Xiao, T., Balakrishna, A., Nair, S., Rafailov, R., Foster, E.P., Sanketi, P.R., Vuong, Q., Kollar, T., Burchfiel, B., Tedrake, R., Sadigh, D., Levine, S., Liang, P., Finn, C.: Openvla: An open-source vision-language-action model. In: *Conference on Robot Learning* (2024)

23. Li, H., Lv, Q., Shao, R., Deng, X., Li, Y., HAO, J., Nie, L.: Star: Learning diverse robot skill abstractions through rotation-augmented vector quantization. In: Forty-second International Conference on Machine Learning (2025)
24. Li, H., Yang, S., Chen, Y., Tian, Y., Yang, X., Chen, X., Wang, H., Wang, T., Zhao, F., Lin, D., et al.: Cronusvla: Transferring latent motion across time for multi-frame prediction in manipulation. In: The Association for the Advancement of Artificial Intelligence (2026)
25. Li, Q., Liang, Y., Wang, Z., Luo, L., Chen, X., Liao, M., Wei, F., Deng, Y., Xu, S., Zhang, Y., et al.: Cogact: A foundational vision-language-action model for synergizing cognition and action in robotic manipulation. arXiv preprint arXiv:2411.19650 (2024)
26. Li, W., Zhang, R., Shao, R., He, J., Nie, L.: Cogvla: Cognition-aligned vision-language-action model via instruction-driven routing & sparsification. In: Advances in neural information processing systems (2025)
27. Li, X., Li, P., Liu, M., Wang, D., Liu, J., Kang, B., Ma, X., Kong, T., Zhang, H., Liu, H.: Towards generalist robot policies: What matters in building vision-language-action models. In: arXiv preprint arXiv:2412.14058 (2024)
28. Li, Z., Chen, W., Yu, J., Lu, Z.: Functional consistency of llm code embeddings: A self-evolving data synthesis framework for benchmarking. *Expert Systems with Applications* p. 129523 (2025)
29. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: International Conference on Learning Representations (2023)
30. Liu, B., Zhu, Y., Gao, C., Feng, Y., Liu, Q., Zhu, Y., Stone, P.: Libero: Benchmarking knowledge transfer for lifelong robot learning. In: Advances in Neural Information Processing Systems (2023)
31. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. In: Advances in Neural Information Processing Systems (2023)
32. Liu, S., Wu, L., Li, B., Tan, H., Chen, H., Wang, Z., Xu, K., Su, H., Zhu, J.: RDT-1b: a diffusion foundation model for bimanual manipulation. In: The Thirteenth International Conference on Learning Representations (2025)
33. Liu, Z., Gu, Y., Zheng, S., Fu, Y., Xue, X., Jiang, Y.G.: Trivla: A triple-system-based unified vision-language-action model with episodic world modeling for general robot control (2025)
34. Meng, R., Jiang, Z., Liu, Y., Su, M., Yang, X., Fu, Y., Qin, C., Chen, Z., Xu, R., Xiong, C., Zhou, Y., Chen, W., Yavuz, S.: Vlm2vec-v2: Advancing multimodal embedding for videos, images, and visual documents (2025)
35. O'Neill, A., Rehman, A., Maddukuri, A., Gupta, A., Padalkar, A., Lee, A., Pooley, A., Gupta, A., Mandlekar, A., Jain, A., et al.: Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration0. In: International Conference on Robotics and Automation (2024)
36. Pertsch, K., Stachowicz, K., Ichter, B., Driess, D., Nair, S., Vuong, Q., Mees, O., Finn, C., Levine, S.: Fast: Efficient action tokenization for vision-language-action models. In: Robotics: Science and Systems (2025)
37. Qu, D., Song, H., Chen, Q., Yao, Y., Ye, X., Ding, Y., Wang, Z., Gu, J., Zhao, B., Wang, D., et al.: Spatialvla: Exploring spatial representations for visual-language-action model. In: Robotics: Science and Systems (2025)
38. Tang, E., Yang, B., Song, X.: Understanding LLM embeddings for regression. *Transactions on Machine Learning Research* (2025)
39. Tao, C., Shen, T., Gao, S., Zhang, J., Li, Z., Hua, K., Hu, W., Tao, Z., Ma, S.: Llms are also effective embedding models: An in-depth overview (2025)

40. Wang, Z., Zhou, Z., Song, J., Huang, Y., Shu, Z., Ma, L.: Vlatest: Testing and evaluating vision-language-action models for robotic manipulation. *Proceedings of ACM Software Engineering* (2025)
41. Yang, S., Li, H., Chen, Y., Wang, B., Tian, Y., Wang, T., Wang, H., Zhao, F., Liao, Y., Pang, J.: Instructvla: Vision-language-action instruction tuning from understanding to manipulation. In: *International Conference on Learning Representations* (2026)
42. Ze, Y., Zhang, G., Zhang, K., Hu, C., Wang, M., Xu, H.: 3d diffusion policy. *arXiv preprint arXiv:2403.03954* (2024)
43. Zhai, S., Zhang, Q., Zhang, T., Huang, F., Zhang, H., Zhou, M., Zhang, S., Liu, L., Lin, S., Pang, J.: A vision-language-action-critic model for robotic real-world reinforcement learning. In: *arXiv preprint arXiv:2509.15937* (2025)
44. Zhang, W., Liu, H., Qi, Z., Wang, Y., Yu, X., Zhang, J., Dong, R., He, J., Wang, H., Zhang, Z., Li, Y., Zeng, W., Jin, X.: DreamVLA: A vision-language-action model dreamed with comprehensive world knowledge. In: *Conference on Neural Information Processing Systems* (2025)
45. Zhao, J., Lu, W., Zhang, D., Liu, Y., Liang, Y., Zhang, T., Cao, Y., Xie, J., Hu, Y., Wang, S., Guo, J., Wang, D., Gao, Y.: Do you need proprioceptive states in visuomotor policies? *arXiv preprint arXiv:2509.18644* (2025)
46. Zhao, Q., Lu, Y., Kim, M.J., Fu, Z., Zhang, Z., Wu, Y., Li, Z., Ma, Q., Han, S., Finn, C., Handa, A., Lin, T., Wetzstein, G., Liu, M., Xiang, D.: Cot-vla: Visual chain-of-thought reasoning for vision-language-action models. In: *Conference on Computer Vision and Pattern Recognition* (2025)
47. Zhong, Y., Huang, X., Li, R., Zhang, C., Chen, Z., Guan, T., Zeng, F., Lui, K.N., Ye, Y., Liang, Y., Yang, Y., Chen, Y.: Dexgraspvla: A vision-language-action framework towards general dexterous grasping. In: *The Association for the Advancement of Artificial Intelligence* (2026)
48. Zhong, Z., Yan, H., Li, J., Liu, X., Gong, X., Zhang, T., Song, W., Chen, J., Zheng, X., Wang, H., Li, H.: Flowvla: Visual chain of thought-based motion reasoning for vision-language-action models. In: *arXiv preprint arXiv:2508.18269* (2025)
49. Zhou, X., Han, X., Yang, F., Ma, Y., Knoll, A.C.: Opendrivelva: Towards end-to-end autonomous driving with large vision language action model. In: *The Association for the Advancement of Artificial Intelligence* (2026)
50. Zhou, Z., Zhu, Y., Wen, J., Shen, C., Xu, Y.: Chatvla-2: Vision-language-action model with open-world embodied reasoning from pretrained knowledge. In: *Advances in Neural Information Processing Systems* (2025)
51. Zitkovich, B., Yu, T., Xu, S., Xu, P., Xiao, T., Xia, F., Wu, J., Wohlhart, P., Welker, S., Wahid, A., Vuong, Q., Vanhoucke, V., Tran, H.T., Soricut, R., Singh, A., Singh, J., Sermanet, P., Sanketi, P.R., Salazar, G., Ryoo, M.S., Reymann, K., Rao, K., Pertsch, K., Mordatch, I., Michalewski, H., Lu, Y., Levine, S., Lee, L., Lee, T.E., Leal, I., Kuang, Y., Kalashnikov, D., Julian, R., Joshi, N.J., Irpan, A., Ichter, B., Hsu, J., Herzog, A., Hausman, K., Gopalakrishnan, K., Fu, C., Florence, P., Finn, C., Dubey, K.A., Driess, D., Ding, T., Choromanski, K.M., Chen, X., Chebotar, Y., Carbajal, J., Brown, N., Brohan, A., Arenas, M.G., Han, K.: RT-2: vision-language-action models transfer web knowledge to robotic control. In: *Conference on Robot Learning* (2023)

Supplementary Material

This appendix provides comprehensive supplementary material to support the methodology, analysis, and findings presented in the main paper. It includes implementation details covering model architectures and training configurations, along with additional experimental results and analyses.

A Additional Implementation Details

A.1 Model Details

Task-Stage Observer. The Task-Stage Observer (TSO) is built upon the Qwen 2.5 VL model [3] and is designed to perform task-relevant reasoning that first identifies the robot’s currently visible physical state and the spatial location and status of the task-relevant objects, then summarizes the immediate stage intention aligned with the given goal. The TSO takes three inputs: the task instruction, the task goal prompt, and a third-person observation image that covers the full workspace. This global viewpoint enables reliable retrieval of object-level evidence even under perturbations.

To ensure that the TSO produces concise and structured semantics rather than free-form scene descriptions, we design a dedicated instruction protocol for the model, as shown in Table 10. This protocol guides the model to enumerate visual evidence, specify stage-wise completion conditions, and infer an explicit task stage description that is aligned with the ongoing trajectory. Following the hidden state aggregation process detailed in Section 3.2 of the main paper, the TSO output is transformed into a single embedding vector of dimension 2048 for ReViP (8192 for ReViP*). This vector serves as the task stage representation z_t and is subsequently fed into the Task-Stage Enhancer to enable task stage feature-wise modulation in ReViP.

Task-Stage Enhancer. In ReViP, the Task-Stage Observer outputs a task stage representation $z_t \in \mathbb{R}^d$, which is processed by a lightweight bottleneck network to produce the modulation parameters (γ_t, β_t) . The mapping $h(\cdot)$ is implemented as a two-layer MLP with SiLU activations. To maintain stable training, the final projection uses a Xavier uniform initializer scaled by 0.01 with zero-initialized biases, so that TS-FiLM behaves close to identity at the early stage while the learnable scalar α gradually adjusts the modulation strength. The resulting $\gamma_t, \beta_t \in \mathbb{R}^D$ are broadcast across tokens and applied to the concatenated vision and language prefix P_t following Eq. (5) in the main paper. Importantly, the binary prefix mask $M_t \in \{0, 1\}^{B \times S}$ is applied after modulation to ensure that TS-FiLM only affects valid prefix tokens and fully preserves the padding structure used by the attention backend.

Attention Mechanism. After applying TS-FiLM to obtain the modulated prefix \tilde{P}_t , the tokens are fed into the vision language backbone together with a suffix sequence that contains the state token and the diffusion-conditioned noisy action tokens. Let $S_t \in \mathbb{R}^{B \times 1 \times D}$ denote the state token and $A_t \in \mathbb{R}^{B \times H \times D}$ denote the

Table 10: Instruction details for the Task-Stage Observer.

Component	Content
System Prompt	You are an expert robot task assistant, integrated with a robot arm capable of executing physical actions in the real world. Based on the task goal and observation, provide the additional task-stage information to the robot in a structured and precise way.
Instruction Rules	<ul style="list-style-type: none"> – LIST the current state of all relevant objects (robot arm, scene objects) as visual evidence. – DEFINE clear completion conditions for each task stage. – Determine the current <code>task_stage_cues</code> based on object states and describe it as an ongoing action.
Input Fields	Task Goal: { <code>task_goal</code> } Observation: {Third person view image}
Output Format	<pre>{ "visual_evidence": ["<evidence_1>", "..."], "task_stage_cues": "<task_stage_cues>"}</pre>

sequence of action and timestep tokens. Then, the action generator’s input X_t can be defined as follows:

$$X_t = [\tilde{P}_t; S_t; A_t] \in \mathbb{R}^{B \times L \times D}, \quad (10)$$

where $L = S + 1 + H$ is the total token length.

Self attention is then applied over X_t with a block structured attention mask \mathcal{M}_t^A that encodes the prefix style interaction pattern inherited from the VLA backbone:

$$H_t = \text{Attn}(Q(X_t), K(X_t), V(X_t); \mathcal{M}_t^A). \quad (11)$$

Prefix tokens in \tilde{P}_t attend to each other across vision and language, but do not attend to the suffix tokens. In contrast, the suffix tokens in S_t and A_t are allowed to attend to all prefix tokens as well as to the appropriate subset of previous suffix tokens according to the autoregressive part of \mathcal{M}_t^A . The binary prefix mask M_t is applied after TS-FiLM so that only valid prefix positions contribute to Q , K , and V , and padding tokens remain completely excluded. This design lets the diffusion head integrate task stage modulated visual context with proprioceptive state and time-dependent action information, while keeping the original routing and padding structure of the attention mechanism unchanged during both training and inference.

A.2 Training Details.

Training Setup. We adopt π_0 [6] as the backbone model with full parameter fine-tuning and set the action chunk size to $K = 50$. The model is trained for 60k

steps with a batch size of 32. The learning rate follows a cosine decay schedule with 3,000 warmup steps, increasing to a peak of 2.0×10^{-5} and then decaying to a final learning rate of 2.0×10^{-6} . We use the AdamW optimizer with gradient clipping set to 1.0 and apply an exponential moving average (EMA) with a decay value of 0.999.

To avoid excessive training latency caused by online task-stage extraction, we preprocess the entire dataset using the Task-Stage Observer before training. For every action step, we extract the corresponding task-stage intention and store it as part of the training annotations. During training, the progress-aware visual cues are injected into ReViP to provide explicit task-stage guidance, enabling the model to more effectively rebalance visual and proprioceptive features. At inference time, task-stage intention is computed online by the Task-Stage Observer and fed into ReViP in real time.

B Additional Experiments Details

B.1 False Completion Benchmark Suite Details

Table 11: Overview of the False Completion Benchmark Suite. The suite contains eight LIBERO manipulation tasks constructed from three perturbation families: Object Drop, Distractor Swap, and Object Relayout. For each perturbation type, we provide representative visual examples together with the corresponding language instructions. The target object is highlighted in purple and the goal region is highlighted in teal.

Type	Instruction	Visual Example
Object Drop	<ol style="list-style-type: none"> ❶ Pick up the butter and place it in the basket ❷ Pick up the cream cheese and place it in the basket ❸ Pick up the salad dressing and place it in the basket ❹ Pick up the orange juice and place it in the basket ❺ Pick up the black bowl in the top drawer of the wooden cabinet and place it on the plate 	
Distractor Swap	<ol style="list-style-type: none"> ❻ Pick up the butter and place it in the basket. ❼ Pick up the cream cheese and place it in the basket. 	
Object Relayout	<ol style="list-style-type: none"> ❽ Pick up the black bowl next to the plate and place it on the plate. 	

False completion arises when external disturbances or unintended robot actions invalidate the ongoing plan while the robot still declares success. To study state-dominant bias and visually apparent false completions in a controlled manner, we construct a False Completion Benchmark Suite on top of LIBERO that

injects targeted perturbations and probes complementary aspects of the phenomenon.

We consider three perturbation sources that commonly induce false completion: *Object Drop*, *Distractor Swap*, and *Object Relayout*. Together, these perturbations form a suite of **8 manipulation tasks** that constitute the proposed False Completion Benchmark. The instruction remains fixed, and only the environment state is perturbed either at the start of the episode or during execution. This disturbance design prevents success via replaying demonstrations like proprioceptive trajectories and instead compels policies to rely on current visual evidence for progress checking and replanning. An overview of all tasks, including representative visualizations and the corresponding language instructions, is provided in Table 11. Below, we detail the three perturbation families.

Object Drop. This setting tests recovery from unexpected displacement and exposes false completion most directly. We instantiate three variants that differ in object scale and contact context, which in total yield five tasks:

- *Small Object Drop*: a piece of cheese and a block of butter are grasped and then forcibly dropped onto the table, requiring the policy to visually re-detect the fallen small object and regrasp it.
- *Large Object Drop*: a bottle of salad dressing and a carton of orange juice are lifted and then dropped, which introduces longer trajectories and larger motion amplitude compared with small object cases.
- *Bowl Drop in Cabinet*: a bowl is dropped inside a cabinet with occlusion and contact constraints, and the policy needs to visually localize the fallen bowl and complete the task under clutter and limited visibility.

Distractor Swap. This setting evaluates whether the policy detects that two visually similar objects have exchanged positions rather than replaying a proprioception-driven trajectory. The initial poses of the target and a visually similar distractor are swapped while the instruction remains fixed. Concretely, cheese and butter exchange locations after an initial movement phase, and the policy must complete the original instruction with respect to the correct target object. This produces two tasks that differ only in which object is specified as the linguistic target.

Object Relayout. This setting breaks demonstration-specific spatial priors while preserving the linguistic target. The target object and its goal region are jointly relocated to new feasible positions within the workspace. The policy is therefore required to update its plan based on the current visual configuration rather than relying on stale internal state or memorized spatial layouts. In our benchmark, this gives one relayout task that relocates both the target and its goal surface.

Overall, the three perturbation families jointly emphasize visual grounding, continuous progress monitoring, and effective replanning under distribution shifts that frequently result in false completions. The suite provides a focused and diagnostic evaluation framework for VLA models, isolating failure modes that arise when policies overly trust internal state progression in the presence of visually apparent disturbances.

B.2 Task Definitions and Evaluation Settings

For all tasks in the False Completion Benchmark Suite, an episode is considered successful only when the agent completes the instructed goal within the predefined step budget. Any premature termination or incorrect placement is counted as failure, even if the policy internally declares success.

Step Budget. The three perturbation families differ in whether the disturbance introduces additional action requirements. For Distractor Swap and Object Relayout, the perturbation only alters the initial scene configuration without increasing manipulation complexity. The number of required actions remains comparable to the original LIBERO tasks, and therefore, the maximum episode length is kept unchanged. In contrast, Object Drop requires the agent to detect the unexpected displacement, relocate the fallen object, and perform an additional re-pick action. This introduces extra motion beyond the demonstration regime. To ensure that policies have sufficient budget to perform recovery, we extend the maximum step limit by an additional 30 to 50 steps, depending on the specific drop variant.

Drop Trigger Rule. For Object Drop tasks, we design a deterministic and reproducible drop triggering mechanism to guarantee fair comparisons across models. A drop event is initiated only when the agent maintains a continuous gripper closing state for a fixed duration, indicating that the object has been securely grasped and lifted. Once this condition is satisfied, the environment forces the gripper to open, causing the object to fall from a consistent spatial location. This rule ensures that all models experience the same drop timing and position, removes stochasticity, and isolates the evaluation of false completion behavior to perception and decision making rather than random disturbances.

Fairness Across Perturbations. For Distractor Swap and Object Relayout, perturbations are applied deterministically at environment initialization. Since the instruction remains unchanged and the agent starts from the same initial state distribution, all models operate under identical scene configurations. For Object Drop, the deterministic drop trigger further ensures that each model receives an identical disturbance during execution. Together, these settings guarantee that the benchmark provides a controlled and strictly fair evaluation of policies under visually apparent distribution shifts that may induce false completion.

B.3 Supplementary Simulation Study: State Masking on LIBERO

To complement the real-robot modality-mismatch study in the main paper, we conduct a modality-ablation experiment in simulation on LIBERO [30] to further validate our findings. LIBERO covers 40 diverse manipulation tasks across four sub-suites, providing a representative large-scale testbed for evaluating the effect of modality ablation under a unified protocol (Section B.4). This supplementary experiment extends the empirical support for our conclusion beyond the real-world setting and tests whether the same trend holds across diverse manipulation scenarios.

As shown in Table 12, the simulation results exhibit a trend consistent with the real-world observations. When the state input is directly masked, performance drops across all four sub-suites. Specifically, compared with the original π_0 , the average suc-

cess rate decreases from 94.2% to 90.5%, with declines of 2.0% on Spatial, 3.0% on Object, 5.2% on Goal, and 4.6% on LIBERO-10. These results suggest that although removing state input can reduce the policy’s reliance on proprioceptive cues, it also removes task-relevant information that is important for normal task execution. Notably, the degradation is more pronounced on Goal and LIBERO-10, indicating that proprioceptive information remains particularly important for tasks requiring stable progress tracking and accurate action execution.

Overall, the simulation study is consistent with the real-robot analysis: naively removing state inputs may partially alleviate state-dominant behavior, but it introduces a clear cost to task performance. These findings further strengthen **our motivation for ReViP**, suggesting that *addressing false completion should shift from directly removing state inputs to rebalancing visual and proprioceptive signals, so as to preserve the complementary benefits of multimodal information for robust task execution.*

B.4 LIBERO & RoboTwin 2.0 Benchmark Details

LIBERO Simulation Benchmark. We evaluate ReViP in the LIBERO simulation benchmark [30], a standardized collection of language-conditioned robotic manipulation tasks designed to probe multimodal grounding, visual reasoning, and action generation. Compared with earlier platforms such as RL Bench [18], LIBERO provides richer language instructions, more diverse scene layouts, and a broader range of object-level interactions, making it a more suitable testbed for evaluating vision language action models.

LIBERO consists of four curated task suites: Spatial, Object, Goal, and Long. Each suite contains ten manipulation tasks accompanied by 50 human teleoperated demonstrations. These suites are designed to decouple different aspects of visuomotor reasoning:

- **LIBERO Spatial** tests spatial reasoning by placing identical objects in varying spatial configurations. Success requires interpreting relations such as left, right, in front of, or behind as described in the instruction.
- **LIBERO Object** evaluates generalization across object categories. While spatial layouts are fixed, the target and distractor objects vary in type, shape, and color, requiring accurate object-level grounding from both vision and language.
- **LIBERO Goal** probes goal-oriented understanding by varying the goal specification while holding object identities and spatial layouts constant. Fine-grained semantic distinctions in instructions must be mapped to distinct manipulation outcomes.

Table 12: Experimental Results of existing VLA models under modality control on LIBERO. Success Rates (%) are reported.

Methods	Spatial \uparrow	Object \uparrow	Goal \uparrow	10 \uparrow	Average \uparrow
π_0 [6]	96.8	98.8	95.8	85.2	94.2
w/ State-Masked	94.8 (\downarrow 2.0)	95.8 (\downarrow 3.0)	90.6 (\downarrow 5.2)	80.6 (\downarrow 4.6)	90.5 (\downarrow 3.7)

- **LIBERO Long** introduces long-horizon tasks involving multiple intermediate steps, diverse objects, and broader spatial compositions. Successful execution requires perception grounding, semantic tracking, and sequential planning.

ReViP is trained and evaluated under the same standard LIBERO protocol as OpenVLA [22] and π_0 to ensure comparability. By covering spatial reasoning, object identification, goal disambiguation, and long-horizon sequencing, the LIBERO benchmark offers a comprehensive and controlled environment for assessing the robustness and visual grounding capabilities of vision-language action models. This makes LIBERO an appropriate platform for evaluating the contributions of task stage modulation in ReViP.

RoboTwin 2.0 Benchmark. RoboTwin 2.0 [9] is a scalable simulation framework and benchmark for robust bimanual robotic manipulation. It instantiates 50 dual-arm tasks across five robot embodiments and provides a large-scale dataset of clean expert trajectories. At the core of RoboTwin 2.0 is a diverse object library and an automatic expert data synthesis pipeline that uses multimodal language models and simulation in the loop refinement to generate realistic task-level execution code. To enhance sim-to-real transfer and robustness, RoboTwin 2.0 applies structured domain randomization along five axes, including clutter, lighting, background, tabletop height, and language instructions, and defines unified evaluation protocols for dual arm policies [9].

In our study, RoboTwin 2.0 serves as a complementary benchmark to LIBERO and is used to assess whether ReViP extends effectively from single-arm settings to more challenging dual-arm coordination. We select four representative tasks from the 50-task suite that cover typical bimanual interaction patterns and object configurations. Following the official RoboTwin 2.0 protocol, we fine-tune each model with 50 clean expert demonstrations per task and evaluate it with 100 rollouts under the Hard setting with domain randomization. This evaluation regime imposes a strong robustness requirement and provides a stringent test of the compared vision language action policies.

B.5 Real-World Setup

Our real-world experiments are performed on a 6-DoF ROKAE robotic arm equipped with a JODELL parallel-jaw gripper. We collect real-world successful demonstration trajectories for the Robot Arm via manually moving the robot along the desired path. For all tasks, we gather 50 expert demonstrations. At each control step, the policy receives a main image and wrist image and a 7-dimensional proprioceptive state in joint space (6 joint positions and 1 gripper scalar). The model outputs action joints in an action chunk of 50.

The main paper reports results on two representative manipulation Tasks 1-2, while this appendix extends the evaluation with additional extended Tasks 3-5 designed to assess hardware robustness, scene generalization, and multi-step reasoning ability (please see Section C.1).

The instructions and descriptions for Tasks 1-2 are provided below:

- “*Pick up the cup on the white plate, then place it inside the white plastic square box*”

This is a structured pick-and-place scenario where the plate serves as a visually distinctive anchor for the target object. The task evaluates basic visual grounding and manipulation stability. In the main paper, this task further includes Drop perturbations to assess the model’s ability to recover from false-completion cases.

- “*Pick up the toy and place it in the open drawer, then close the drawer*”

A multi-stage drawer manipulation task requiring correct sequencing: grasping, placing, and completing the closing motion. This task tests long-horizon reasoning, stage transitions, and re-localization when object visibility changes near drawer boundaries.

Drop Trigger Mechanism We evaluate the model’s response to false-completion scenarios using a controlled Drop perturbation. To ensure **repeatability, fairness, and naturalistic failure behavior**, we implement Drop using a hardware-in-the-loop condition:

First, during execution, we continuously monitor the gripper width in real time. Once the gripper closes beyond a threshold, indicating that it has successfully grasped the object, the robot proceeds with the placement trajectory. Then, at a predefined early stage of this motion, we command the gripper to open, causing the object to fall naturally along its original motion path.

C Supplementary Quantitative Analysis

C.1 Extended Real-World Task Results

Beyond the standard configuration described above, the appendix includes additional hardware and scene generalization experiments to further evaluate the robustness of ReViP under real-world variations not covered in the main paper. To this end, we introduce three additional tasks that involve new backgrounds, different camera hardware, new object categories, and extended multi-step action sequences. In particular, we replace the wrist-mounted cameras with the ORBBEC DaBai DCW module and substitute the tabletop background with a high-contrast green cloth. These changes alter imaging characteristics, depth quality, and background complexity, creating a significantly more challenging perception environment.

The instructions and descriptions for the extended Tasks 3-5 are provided below:

- “*Pick up the cup on the table, then place it inside the white plastic square box*”. This task removes the plate anchor used in the main paper and employs a plain tabletop decorated with a high-contrast green cloth. The modified background introduces strong color bias and lower object-background separability, making target localization more difficult.

- “Pick up the small red cube on the table, then place it inside the white box”. A fine-grained manipulation task involving a small, low-profile cube whose visual features are harder to extract under depth noise. Surrounding distractors further increase the challenge, requiring precise visual grounding and accurate grasping.
- “Pick up the purple toy on the table, then place it into the middle drawer of the cabinet, and finally close the drawer”. This extended long-horizon task involves grasping, precise placement, and drawer closing. The scene contains up to ten objects of diverse shapes and colors, creating substantial visual clutter. This setting rigorously tests the model’s ability to maintain stable grounding and execute multi-step actions under complex external interference.

Table 6 in the main paper summarizes the success rates across all extended tasks. The results reveal several consistent trends:

Cup Placement. ReViP substantially outperforms both π_0 and π_0 -Fast across all variants, especially under Drop and Distractor settings. The improvement highlights ReViP’s stronger ability to recover from false-completion risks (e.g., object drop) and to reject visually similar distractors in cluttered scenes.

Cup Placement. The cube task poses significant difficulty due to the small size and weak geometric features of the target object. While both baselines suffer heavily under distractor interference, ReViP maintains considerably higher success rates, demonstrating improved fine-grained localization and stable visual grounding.

Drawer Manipulation. ReViP also achieves the highest performance on the long-horizon drawer task. Baseline policies frequently exhibit state-dominant behavior, prematurely closing the drawer after empty grasps or proceeding despite misalignment. In contrast, ReViP more reliably detects errors, re-localizes the target, and completes the full multi-step sequence.

Across all real-world settings, ReViP maintains substantially higher robustness under perturbations and complex visual clutter. These results confirm that the proposed vision-proprioception rebalance improves both false-completion resilience and long-horizon execution in challenging, unconstrained environments.

C.2 Study of the Model’s Plug-and-Play Capability

To further evaluate the plug-and-play capability of the *Vision-Proprioception Rebalance* beyond the π_0 [6] backbone used in ReViP, we extend our framework to another state-of-the-art VLA architecture, $\pi_{0.5}$ [17]. We instantiate ReViP on $\pi_{0.5}$, denoted ReViP_{0.5} by integrating the same Task-Stage Observer, Task-Stage Enhancer, and TS-FiLM pipeline, without adding any extra modules or backbone-specific tuning. This setup examines whether the proposed progress-aware visual cues injection functions in a plug-and-play manner across different policy formulations and whether the benefits observed in ReViP naturally transfer to new architectures.

Table 13: Experimental Results on the Plug-and-Play Capability of Vision-Proprioception Rebalance. Success rates (SR) across two task suites: Object Drop in False Completion benchmarks & LIBERO 10. ReViP and ReViP_{0.5} consistently improve success rates across both suites. * indicates results from the official GitHub implementation.

Methods	Object-Drop	LIBERO-10
π_0 [6]	37.6%	85.2%
$\pi_{0.5}$ * [17]	54.4%	92.4%
ReViP	65.2% (\uparrow 27.6%)	92.2% (\uparrow 7.0%)
ReViP _{0.5}	68.2% (\uparrow 13.8%)	95.8% (\uparrow 3.4%)

For evaluation, we adopt two challenging task suites that probe complementary aspects of robustness: (i) *Object Drop*, containing five representative false-completion tasks requiring failure recovery, and (ii) *LIBERO-10*, a long-horizon manipulation suite demanding consistent grounding and precise action sequencing. Both suites remain identical across backbones to ensure fair comparison.

As shown in Table 13, ReViP and ReViP_{0.5} consistently improve success rates across both task suites on their respective backbones, π_0 and $\pi_{0.5}$. On the π_0 backbone, ReViP achieves improvements of 27.6% on Object Drop and 7.0% on LIBERO-10. When transferred to $\pi_{0.5}$, the same design further yields gains of 13.8% on Object Drop and 7.0% on LIBERO-10. These results demonstrate that task-stage conditioning through structured visual evidence functions effectively in a plug-and-play fashion across different VLA architectures, indicating that the proposed vision-proprioception rebalance mechanism is broadly applicable and **not tied to any specific backbone**.

D Supplementary Qualitative Analysis

We present additional qualitative results from both simulation and real-world experiments, focusing on challenging false-completion scenarios and demonstrating ReViP’s improved robustness, generalization, and execution quality.

D.1 Additional Visualizations of Simulation Results

As shown in Figure 6, we present additional simulation results on both the false-completion benchmark and LIBERO. The model consistently completes multi-step tasks under diverse false-completion disturbances, environments, and object configurations. These visualizations highlight the complexity of the tasks and further demonstrate ReViP’s generalization and execution capabilities.

D.2 Additional Visualizations of Real-World Results

Figure 7 provides qualitative comparisons between ReViP and the baseline π_0 across two representative real-world settings: false-completion scenarios and long-



Fig. 6: Additional qualitative results in simulation on the False-Completion Benchmark, LIBERO-10, and LIBERO-Goal. In the False-Completion Benchmark, ReViP detects the object drop during execution and re-grasps the target, achieving true completion instead of prematurely terminating. On LIBERO-10 and LIBERO-Goal, ReViP consistently completes long-horizon tasks under diverse disturbances, environments, and object configurations. These visualizations highlight the complexity of the tasks and further demonstrate ReViP’s robustness, generalization, and execution capabilities.

horizon manipulation tasks. The visualizations highlight how ReViP maintains correct visual grounding and robust task progression.

In cluttered environments containing many distractors with highly similar appearance, ReViP accurately identifies the true target object, such as distinguishing the intended cup or red cube from visually similar items. This allows ReViP to execute the correct grasping action. In contrast, π_0 is frequently misled and grasps an incorrect object, which triggers an incorrect stage transition and results in a typical false-completion failure despite clear visual evidence of mis-grasp.

For complex long-horizon manipulation tasks, π_0 often follows internal state progression rather than visual feedback. This leads to errors such as pushing or closing the cabinet even when the gripper is empty, ultimately causing task failure. ReViP detects these inconsistencies, re-localizes the missing target, and then completes the full sequence of grasping, placing, and cabinet operations.

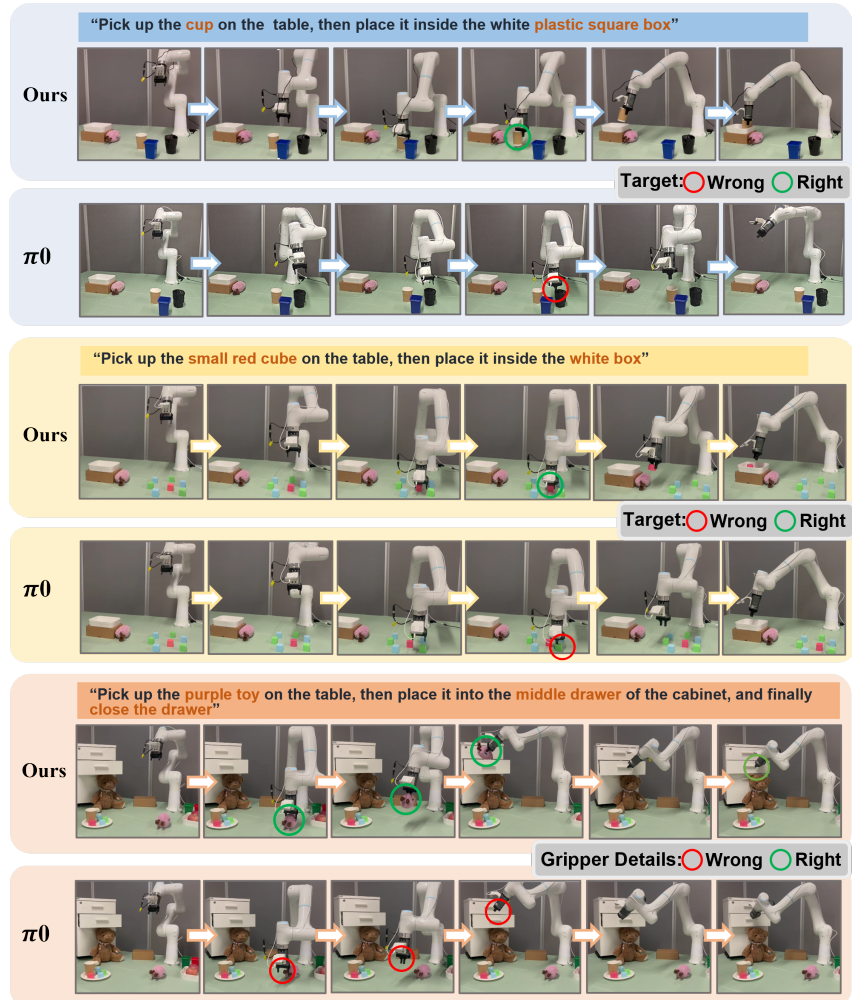


Fig. 7: Additional qualitative real-world comparisons between ReViP and π_0 on challenging false-completion scenarios and long-horizon manipulation tasks. ReViP accurately identifies the correct target in cluttered scenes and avoids mis-grasp-induced false completion, while π_0 is easily confused by distractors. For long-horizon drawer tasks, ReViP re-localizes the target and completes the full sequence, whereas π_0 often proceeds with incorrect state-driven actions. The corresponding videos are provided in the supplementary materials.

These qualitative examples show that ReViP mitigates false-completion behaviors and significantly improves long-horizon task reliability through better integration of visual observations. The corresponding real-world videos are provided in the supplementary materials.