ASR-VLA: ADAPTIVE SAFETY REASONING WITH SAFETY-CRITICAL FUTURE PREDICTION FOR VISION-LANGUAGE-ACTION MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Vision-Language-Action (VLA) models have demonstrated strong capabilities in instruction-following robotic manipulation. However, ensuring safety during task execution remains an open challenge. Existing safety-aware VLA methods primarily incorporate safety constraints into policy learning or apply post hoc correction to potentially unsafe actions. Such methods generally neither incorporate potential future hazards into action selection nor determine how the robot should respond to them. Consequently, the resulting policies may struggle to select actions that reduce risk without unnecessarily compromising task progress. To address this issue, we introduce ASR-VLA, an adaptive framework that combines a safety-critical world model with a vision-language reasoning model. The world model predicts a future observation associated with the highest risk. The reasoning model analyzes this observation together with the current scene and task context, producing structured guidance for selecting actions that respond to the potential hazards while preserving task progress. Experiments on SafeLIBERO demonstrate that ASR-VLA substantially improves both task performance and safety, achieving a 93.68% success rate and a 9.46% collision rate. Compared with the strongest baseline, ASR-VLA improves the success rate by 20.72 percentage points and reduces the collision rate by 10.46 percentage points, demonstrating a better balance between task completion and safety.

1 INTRODUCTION

Vision-Language-Action (VLA) models provide a unified framework for instruction-following robotic manipulation. They leverage large-scale vision-language pretraining to map visual observations and natural-language instructions to robot actions (Brohan et al., 2022; Kim et al., 2024; Zitkovich et al., 2023; Black et al., 2024). These models have demonstrated strong performance across a range of manipulation tasks and shown substantial potential for generalization. However, prior studies have mainly focused on task completion and generalization, while execution safety has received less attention. In real-world environments, a VLA-based policy may successfully complete the instructed task while still causing collisions or unintended contact with surrounding objects. Successful task completion does not necessarily imply safe execution, and this gap remains an important challenge for the reliable deployment of VLA models in real-world environments (Zhang et al., 2026a; 2025; Ma et al., 2026).

Recent studies have explored VLA safety through constrained policy learning, failure detection, and deployment-time action intervention (Hu et al., 2025; Zhang et al., 2026a; Gu et al., 2026). These approaches have shown that incorporating safety information can improve the reliability of VLA execution. However, they generally focus on detecting, constraining, or correcting unsafe behavior, without explicitly reasoning about how a potential hazard relates to the current manipulation task or what response would be most appropriate. This distinction is important because similar risk estimates may call for different actions depending on the scene geometry and task stage. Without such task-aware reasoning, safety interventions may become overly conservative or poorly aligned with the ongoing task, leading to unnecessary motion deviations or reduced task performance. Therefore, a key open problem is how VLA policies can interpret task-relevant hazards and translate them into context-specific guidance for action selection.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

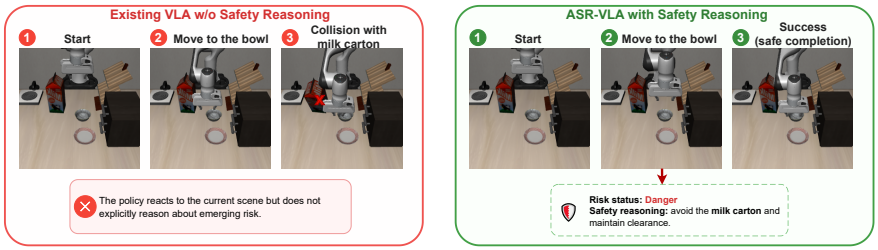


Figure 1: VLA execution without and with safety reasoning

To bridge this gap, we propose **ASR-VLA**, an adaptive safety reasoning framework that enables VLA policies to derive task-aware action guidance from the potential hazards. ASR-VLA combines a vision-language reasoning model with a safety-critical world model. The world model predicts a future observation associated with the highest risk and provides it as auxiliary visual evidence. The reasoning model interprets this evidence together with the current observation to identify the relevant hazards and determine an appropriate response. As illustrated in Fig. 1, a safety-aware VLA without reasoning may detect risk yet still generate a collision-prone trajectory because it lacks task-specific guidance for responding to the hazards. In contrast, ASR-VLA can recommend targeted adjustments, such as increasing clearance before continuing the manipulation, supporting safer execution without unnecessarily interrupting task progress.

We summarize our contributions as follows. **(1)** We introduce ASR-VLA, an adaptive framework that integrates safety reasoning into VLA policies, enabling task-aware responses to predicted future hazards. **(2)** We develop a future-grounded safety reasoning mechanism that augments vision-language reasoning with prospective visual evidence. By jointly analyzing the current scene and a predicted safety-critical future observation, the reasoning model can better identify emerging hazards and generate task-specific guidance that is directly grounded in VLA action generation. **(3)** We construct a safety-reasoning dataset for SafeLIBERO and conduct extensive experiments demonstrating that ASR-VLA achieves a 93.68% success rate and a 9.46% collision rate, improving the success rate by 20.72 percentage points and reducing the collision rate by 10.46 percentage points over the strongest evaluated baseline.

2 Related Work

2.1 Reasoning in Vision-Language-Action Models

Recent studies have incorporated intermediate reasoning stages into VLA policies to move beyond direct action prediction from observations and language instructions (Zhong et al., 2026; Zhou et al., 2026; Mazzaglia et al., 2025; Liu et al., 2026). Embodied Chain-of-Thought reasoning (Zawalski et al., 2024) trains VLA models to generate structured reasoning about task plans, subtasks, visual observations, and robot motions before predicting actions. CoT-VLA (Zhao et al., 2025) introduces visual reasoning by autoregressively predicting a future subgoal image before generating a short sequence of actions to achieve the predicted state. HALO (Shou et al., 2026) further integrates textual task reasoning, visual subgoal prediction, and action generation within a unified multimodal reasoning framework. LaST-R1 (Chen et al., 2026a) explores adaptive latent reasoning and jointly optimizes the latent reasoning process and robot actions through post-training with reinforcement learning. Despite these advances, existing methods rarely reason about the potential hazards or appropriate responses.

2.2 Safety-Aware Vision-Language-Action Models

Safety has recently received increasing attention in VLA research, as successful task completion does not necessarily imply safe execution (Lin et al., 2025; Tölle et al., 2025; Zhang et al., 2026b; Brunke et al., 2025; Chen et al., 2026b; Zhang et al., 2026c). SafeVLA (Zhang et al., 2026a) studies safety alignment through constrained policy learning. VLSA (Hu et al., 2025) introduces a plug-and-play safety constraint layer that modifies unsafe actions during deployment. SafeDojo (Tang et al., 2026) combines an interactive world model with constrained reinforcement learning to improve the safety of VLA policies. Recent work has also introduced benchmarks for evaluating VLA safety. LIBERO-Safety (Cui et al., 2026) covers both physical and semantic safety, and SafeVLA-Bench (Fan et al., 2026) evaluates the gap between task success and safe execution. Despite these

advances, existing safety-oriented VLA methods generally do not explicitly analyze visual hazards in the task context or determine how the robot should respond. In this work, ASR-VLA addresses this limitation by using a predicted future observation to analyze task-relevant hazards and generate structured safety guidance for action selection.

3 METHODOLOGY

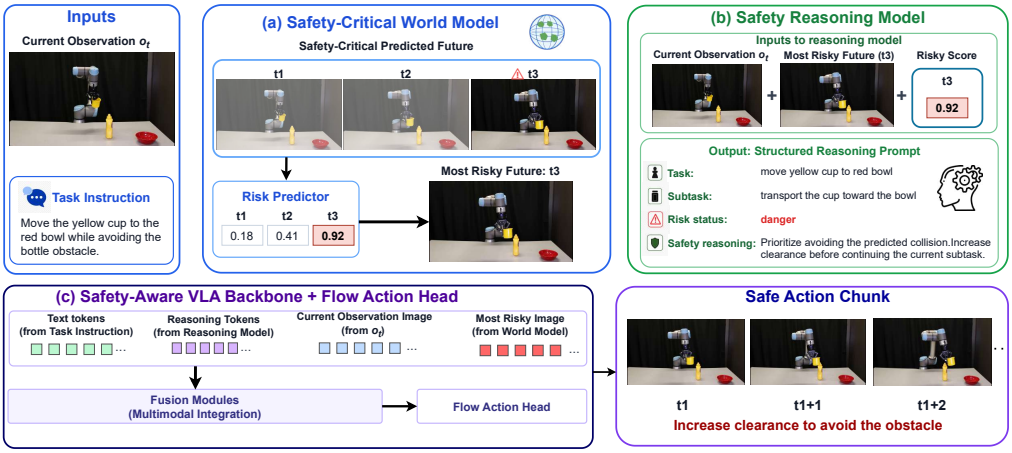


Figure 2: Overview of ASR-VLA. Given the current observation and task instruction, ASR-VLA integrates prospective visual risk evidence and structured safety reasoning to generate a safe action chunk. In panel (a), the safety-critical world model predicts the high-risk future observation, together with its corresponding risk score.

3.1 PRELIMINARIES

In this paper, we consider instruction-following VLA policies for robotic manipulation. At each time step t, the robot receives an observation $o_t = (I_t, x_t)$, where I_t denotes the visual images, possibly comprising multiple camera views, and x_t represents the robot proprioceptive state. The robot is also provided with a task instruction ℓ specifying the manipulation task. Given the current observation and instruction, a standard VLA policy generates an action chunk

$$\boldsymbol{a}_{t:t+H_a-1} = \pi(o_t, \ell), \tag{1}$$

where H_a denotes the action horizon. Although existing VLA policies can achieve strong performance on instruction-following manipulation tasks, standard VLA policies generate actions primarily from the current observation and typically do not model how safety risks may evolve during subsequent execution. Therefore, our objective is to augment the VLA policy with safety reasoning so that it can identify potential hazards and adapt its behavior before an unsafe interaction occurs. As illustrated in Fig. 2, ASR-VLA consists of three major components: a safety-critical world model, a safety reasoning model, and a safety-aware VLA backbone with a flow-based action head.

The safety-critical world model predicts the future visual observation associated with the highest estimated risk, together with its corresponding risk score. The safety reasoning model then analyzes the current observation, task instruction, predicted high-risk future observation, and estimated risk score to generate a structured safety reasoning output R_t. Finally, the VLA backbone integrates the current observation, task instruction, predicted high-risk future observation, and safety reasoning output to generate actions that adjust the robot's motion in response to emerging safety risks while maintaining task progress. Accordingly, the proposed VLA policy generates the action chunk as

$$\boldsymbol{a}_{t:t+H_a-1} = \pi_\theta(o_t, \ell, R_t, \hat{I}_t^{\text{risk}}), \tag{2}$$

where $\hat{I}_t^{\text{risk}}$ denotes the predicted high-risk future visual observation. The details of R_t are shown in Eq. 12.

3.2 SAFETY-CRITICAL WORLD MODEL

The current visual observation captures only the present scene and does not reveal how the spatial relationship among the robot, task-relevant objects, and nearby obstacles may evolve during subsequent execution. To provide the policy with visual evidence of potential future risks, we introduce a safety-critical world model that predicts the future visual observation associated with the highest safety risk. Given I_t, x_t, and ℓ, the world model W_ω predicts a high-risk future observation together with its corresponding risk score

$$(\hat{I}_t^{\text{risk}}, \hat{\rho}_t) = W_\omega\left(I_t, x_t, \ell\right), \tag{3}$$

where $\hat{\rho}_t$ denotes the estimated risk score.

Safety-Critical Target Construction. To construct the training target, we compute the minimum robot-obstacle clearance d_j at each time step in the collected trajectories (Ames et al., 2016). Let

$$d_j = \min_{b \in \mathcal{B}_{\text{obs}}} d(\mathcal{G}_{\text{robot}}(x_j), \mathcal{G}_b(x_j)), \tag{4}$$

where $\mathcal{B}_{\text{obs}}$ denotes the set of obstacles. $\mathcal{G}_{\text{robot}}(x_j)$ and $\mathcal{G}_b(x_j)$ represent the robot and obstacle geometries at time step j, respectively. The function $d(\cdot,\cdot)$ measures the minimum distance between two geometries. We then convert the minimum clearance into a normalized risk score

$$\rho_j = \text{sigmoid}\left(\frac{d_{\text{safe}} - d_j}{\beta}\right), \tag{5}$$

where d_{safe} denotes the desired minimum clearance and β controls the sharpness of the transition. Consequently, the risk score increases as the robot approaches an obstacle and becomes high when the clearance falls below the prescribed safety margin.

For each current time step t, we identify the future time step with the highest risk over $\mathcal{K}_t$

$$k_t^\star = \arg\max_{k \in \mathcal{K}_t} \rho_{t+k}, \quad \mathcal{K}_t = \{1, \ldots, \min(H_w, T-t)\}, \tag{6}$$

where $k_t^\star$ denotes the offset from the current time step to the highest-risk future time step and H_w denotes the number of future observations considered when constructing the target. The resulting supervision is

$$I_t^{\text{risk}} = I_{t+k_t^\star}, \qquad \rho_t^\star = \rho_{t+k_t^\star}. \tag{7}$$

Here, I_t^{risk} denotes the ground-truth high-risk future observation, while $\rho_t^\star$ denotes its corresponding ground-truth risk score.

World-Model Training Data. We use $\pi_{0.5}$ (Intelligence et al., 2025) as the data-collection policy and deploy it in the obstacle-augmented training environments to collect closed-loop trajectories. Since $\pi_{0.5}$ does not incorporate future-risk prediction or structured safety reasoning, the collected trajectories naturally contain low-clearance, near-collision, and collision states, providing safety-critical supervision for training the world model. For each time step t in a collected trajectory, we compute the future risk scores using Eq. 4 and 5 and select the high-risk future observation using Eq. 6 and 7. This produces the world-model training dataset

$$\mathcal{D}_{\text{WM}} = \left\{\left(I_t, x_t, \ell, I_t^{\text{risk}}, \rho_t^\star\right)\right\}. \tag{8}$$

Architecture and Training Objective. Given I_t, x_t, and ℓ, we instantiate the world model W_ω using three modality-specific encoders, followed by a multimodal fusion module and two prediction heads. The world model separately encodes each input modality and fuses them into a shared representation

$$z_t = F_\omega\left(E_I(I_t), E_x(x_t), E_\ell(\ell)\right), \tag{9}$$

where E_I, E_x, and E_ℓ denote the visual, state, and language encoders, respectively, and F_ω denotes the multimodal fusion module. An image decoder D_I then predicts the high-risk future observation, while a risk-score prediction head f_ρ estimates its associated risk score

$$\hat{I}_t^{\text{risk}} = D_I(z_t), \qquad \hat{\rho}_t = f_\rho(z_t). \tag{10}$$

The world model is optimized using a joint objective for future-image prediction and risk-score regression

$$\mathcal{L}_{\text{WM}} = \mathbb{E}_{\mathcal{D}_{\text{WM}}}\left[\lambda_{\text{img}} \left\|\hat{I}_t^{\text{risk}} - I_t^{\text{risk}}\right\|_1 + \lambda_{\text{risk}}\left(\hat{\rho}_t - \rho_t^\star\right)^2\right], \tag{11}$$

where λ_{img} and λ_{risk} control the relative contributions of future-image prediction and risk-score regression.

3.3 SAFETY REASONING MODEL

To enable the policy to respond to potential safety risks during task execution, we introduce a safety reasoning model. The model takes the current observation, the predicted high-risk future observation, the task instruction, and the estimated risk level as inputs, and produces structured, task-aware safety guidance for the VLA policy.

Given I_t, $\hat{I}_t^{\text{risk}}$, ℓ, and $\hat{\rho}_t$, the reasoning model is formulated as

$$R_t = G_\phi \left(I_t, \hat{I}_t^{\text{risk}}, \ell, \hat{\rho}_t \right), \tag{12}$$

where G_ϕ denotes a vision-language reasoning model and R_t is the structured reasoning context provided to the VLA backbone.

Reasoning Supervision. During training, the current subtask c_t is obtained from the phase annotation of the collected trajectory, while the target safety status is constructed from $\rho_t^\star$. Using two thresholds τ_{warn} and τ_{danger}, we define

$$q_t^\star = \begin{cases} \texttt{DANGER}, & \rho_t^\star \geq \tau_{\text{danger}}, \\ \texttt{WARNING}, & \tau_{\text{warn}} \leq \rho_t^\star < \tau_{\text{danger}}, \\ \texttt{SAFE}, & \rho_t^\star < \tau_{\text{warn}}. \end{cases} \tag{13}$$

The reasoning model is trained to predict the structured reasoning sequence using an autoregressive language-modeling objective. Let $R_t^\star = (r_{t,1}^\star, \ldots, r_{t,L_t}^\star)$ denote the tokenized target reasoning sequence, where $L_t = |R_t^\star|$. The reasoning model G_ϕ defines an autoregressive conditional distribution p_ϕ over this sequence. We optimize it using

$$\mathcal{L}_{\text{reason}} = -\sum_{j=1}^{L_t} \log p_\phi \left(r_{t,j}^\star \mid r_{t,<j}^\star, I_t, \hat{I}_t^{\text{risk}}, \ell, \hat{\rho}_t \right), \tag{14}$$

where $r_{t,<j}^\star = (r_{t,1}^\star, \ldots, r_{t,j-1}^\star)$ is the ground-truth token prefix. In our implementation, G_ϕ in Eq. 12 is initialized from Qwen3-VL-8B-Instruct (Bai et al., 2025) and fine-tuned using low-rank adaptation (LoRA) (Hu et al., 2022). The pretrained backbone is kept frozen, while the inserted LoRA parameters are optimized using $\mathcal{L}_{\text{reason}}$.

3.4 SAFETY-GROUNDED VLA POLICY

The safety-critical world model and reasoning model provide two complementary forms of safety information. We incorporate the outputs of both modules into the VLA policy rather than using them as an external action filter. We use $\pi_{0.5}$ as the VLA backbone and augment its original multimodal input with the predicted high-risk future observation and the structured safety reasoning sequence. Given o_t, ℓ, R_t and $\hat{I}_t^{\text{risk}}$, the safety-grounded policy generates the action chunk according to Eq. 2.

Rather than producing unconstrained free-form reasoning, the model generates a structured output consisting of the task instruction, current subtask, safety status, and corresponding guidance. Specifically, the task field retains ℓ, the subtask field identifies the current manipulation phase, the safety-status field q_t takes one of $\{\texttt{SAFE}, \texttt{WARNING}, \texttt{DANGER}\}$, and the guidance field g_t specifies the appropriate behavioral adjustment. This structured format enables the model to relate predicted hazards to the current task stage while avoiding unnecessarily long free-form reasoning, which is particularly important for maintaining low-latency VLA execution. A `SAFE` status indicates that the current subtask can proceed normally, `WARNING` calls for an early adjustment while preserving task progress, and `DANGER` prioritizes hazard avoidance before resuming the subtask. Complete templates and representative examples are provided in Appendix C.

Multimodal Safety Grounding. We incorporate the predicted high-risk future observation and structured reasoning sequence into the VLA policy by augmenting its original multimodal input. The conditioning representations are constructed as

$$\begin{aligned} \mathcal{Z}_t &= \left[E_V(I_t),\, E_V(\hat{I}_t^{\text{risk}}),\, E_L(\ell),\, E_L(R_t),\, E_X(x_t) \right], \\ h_t &= F_\theta(\mathcal{Z}_t), \end{aligned} \tag{15}$$

where E_V, E_L, E_X are the visual, language, and state encoders and F_θ is the multimodal fusion module of the VLA backbone. $\hat{I}_t^{\text{risk}}$ provides additional visual information, while R_t provides safety guidance that complements the task instruction. The fused representation h_t is then passed to the flow-based action head.

Flow-Based Action Generation. We train the safety-grounded VLA policy using our constructed SafeLIBERO safety-reasoning dataset. Each sample provides the current observation, task instruction, robot state, safety-critical future evidence, structured safety reasoning, and an expert action chunk. These inputs are encoded into the fused representation h_t, while the expert action chunk $\boldsymbol{a}_t^\star$ serves as the supervision target for the flow-based action head.

The action head learns a conditional vector field that transports Gaussian noise to the expert action distribution (Lipman et al., 2022; Tong et al., 2023). For a flow time step $s \sim \mathcal{U}(0,1)$, we define the linear interpolation

$$\boldsymbol{a}_t^s = (1-s)\boldsymbol{\epsilon} + s\boldsymbol{a}_t^\star. \tag{16}$$

The corresponding target vector field is

$$\boldsymbol{v}_t^\star = \frac{d\boldsymbol{a}_t^s}{ds} = \boldsymbol{a}_t^\star - \boldsymbol{\epsilon}. \tag{17}$$

The action head v_θ predicts the conditional vector field from the interpolated action, flow time, and fused representation

$$\hat{\boldsymbol{v}}_t^s = v_\theta\left(\boldsymbol{a}_t^s, s, h_t\right). \tag{18}$$

It is optimized using the conditional flow-matching objective

$$\begin{aligned}\mathcal{L}_{\text{flow}} = \mathbb{E}_{(\boldsymbol{a}_t^\star, h_t)\sim\mathcal{D}_{\text{VLA}},\ s\sim\mathcal{U}(0,1),\boldsymbol{\epsilon}\sim\mathcal{N}(\mathbf{0},\boldsymbol{I})} \\ \left[\left\| v_\theta\left(\boldsymbol{a}_t^s, s, h_t\right) - \left(\boldsymbol{a}_t^\star - \boldsymbol{\epsilon}\right)\right\|_2^2\right].\end{aligned} \tag{19}$$

During inference, we initialize the action chunk with $\boldsymbol{a}_t^0 \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$ and solve the conditional flow ordinary differential equation

$$\frac{d\boldsymbol{a}_t^s}{ds} = v_\theta\left(\boldsymbol{a}_t^s, s, h_t\right), \qquad s \in [0,1]. \tag{20}$$

The terminal state $\boldsymbol{a}_t^1$ is used as the predicted action chunk. At the next replanning step, ASR-VLA recomputes the future-risk score, and the reasoning model is invoked only when the predicted safety status changes; otherwise, its most recent output is reused.

4 Experimental Results

We conduct experiments on SafeLIBERO benchmark to answer the following research questions: **Q1: Comparison with Safety-Aware VLA Methods.** How does ASR-VLA compare with representative state-of-the-art safety-aware VLA methods? **Q2: Ablation Study.** What are the contributions of each component to the overall performance of ASR-VLA? **Q3: Safety Grounding.** How effectively does structured safety reasoning guide the policy's action generation? **Q4: Hardware Demonstration.** Can the framework transfer its safety-aware behavior to a physical platform? **Q5: Robustness to Execution Perturbations.** How robust is ASR-VLA to execution perturbations while preserving task progress and avoiding collisions? (Provided in Appendix B) **Q6: Inference Efficiency.** What additional inference latency is introduced by the world model and reasoning model? (Provided in Appendix B)



(a) Goal. (b) Spatial. (c) Object.

Figure 3: Representative simulation scenarios from three SafeLIBERO tasks. Semi-transparent robot poses indicate earlier execution stages, while the opaque pose shows the final configuration.

4.1 EXPERIMENTAL SETUP

Benchmark and Demonstration Dataset. We evaluate ASR-VLA on SafeLIBERO (Hu et al., 2025), which contains 25 manipulation tasks across four task suites: *Spatial*, *Goal*, *Object*, and *Long*. Representative simulation scenes from the *Goal*, *Spatial*, and *Object* suites are shown in Figure 3. For all tasks, we collect simulated demonstrations containing external-view and wrist-view observations, robot proprioceptive states, and task instructions. Based on these demonstrations, we construct a safety-reasoning dataset by augmenting each sample with a safety-critical future observation, a continuous risk score, manipulation-phase annotation, and structured safety reasoning supervision. This dataset is used to train the world model, reasoning model, and safety-grounded VLA policy. Detailed task descriptions, dataset construction, and task-specific configurations are provided in Appendix A. We additionally conduct hardware validation on a robotic platform equipped with a Robotiq 2F-85 gripper. The hardware evaluation includes three manipulation tasks: moving a cup to a bowl from left to right, performing the mirrored cup-to-bowl task from right to left, and moving a cup to a basket from left to right.

Baselines. We compare ASR-VLA with the following methods: $\boldsymbol{\pi}_{0.5}$, which is fine-tuned on the LIBERO dataset (Liu et al., 2023) without safety consideration. **OpenVLA-OFT.** An OpenVLA policy fine-tuned on the same demonstration LIBERO data using supervised learning, without safety-specific training. **VLSA** (Hu et al., 2025). A safety-aware VLA method that uses a plug-and-play safety constraint layer to correct potentially unsafe actions during deployment. **SafeVLA** (Zhang et al., 2026a). A safety-aligned VLA method trained through constrained policy learning.

For a fair comparison, all trainable baselines use the same task demonstrations, observation interface, and evaluation initializations whenever their formulations permit. We report two primary evaluation metrics. The *success rate* (**SR**) measures the percentage of episodes that satisfy the task objective. The *collision rate* (**CR**) measures the percentage of episodes containing at least one collision.

Method	Goal	Spatial	Object	Long	Overall
OpenVLA-OFT	25.43 / 83.14	36.33 / 87.67	29.50 / 82.75	7.00 / 92.50	26.40 / 85.60
$\pi_{0.5}$	58.14 / 75.36	57.82 / 87.21	53.61 / 77.18	34.73 / 85.89	51.07 / 81.41
VLSA	75.71 / 20.29	77.33 / 19.33	80.25 / 25.25	47.00 / **9.50**	72.96 / 19.92
SafeVLA	65.20 / 22.17	86.80 / 24.27	59.35 / 32.15	35.20 / 29.80	63.71 / 27.09
ASR-VLA w/o Safety Reasoning	94.46 / 18.63	95.27 / 17.73	91.25 / 7.10	82.70 / 18.70	91.74 / 14.74
ASR-VLA w/o Risky-Future Image	94.80 / 16.91	94.20 / 16.27	92.85 / 7.25	79.60 / 18.90	91.60 / 13.98
ASR-VLA	**97.03 / 11.71**	**96.33 / 12.27**	**93.20 / 5.25**	**84.80** / 9.70	**93.68 / 9.46**

Table 1: Overall performance on the SafeLIBERO scenarios. Each entry reports success rate and collision rate in the order **SR** (↑) **/ CR** (↓).

4.2 Q1. OVERALL PERFORMANCE

In Table 1, we evaluate all methods on the valid benchmark tasks using 50 initial configurations per task and five random seeds. We report the task-averaged SR and CR (%) across all evaluation episodes. The results compare ASR-VLA with standard and safety-aware VLA baselines. OpenVLA-OFT and $\pi_{0.5}$ are optimized primarily for task execution and do not receive safety supervision. Consequently, both methods exhibit high collision rates across all four suites. Compared with OpenVLA-OFT, $\pi_{0.5}$ achieves substantially higher task success, suggesting stronger task generalization under our evaluation settings. Nevertheless, without future-risk prediction, it frequently contacts nearby obstacles, which can interrupt manipulation progress and contribute to task failures. Among the safety-aware baselines, VLSA substantially reduces CR through deployment-time action correction. However, its external safety filter may introduce conservative or abrupt motion deviations. Such action corrections may shift the robot into states that are rarely encountered during training, where the policy is less reliable and subsequent errors are more likely to accumulate. SafeVLA applies PPO-Lagrangian safety alignment to a policy pretrained on the original LIBERO dataset. Its performance suggests that balancing task rewards and safety constraints during safety alignment may involve a trade-off between task success and collision avoidance. One possible explanation is that the distribution shift from the original LIBERO tasks to safety-critical configurations

limits effective adaptation, while overly conservative policy updates can degrade the manipulation capabilities inherited from pretraining. In contrast, ASR-VLA incorporates safety information before action generation. By grounding safety signals in the VLA policy, ASR-VLA can adjust its motion in response to potential hazards while preserving progress toward the task objective. Compared with other baselines, ASR-VLA improves SR by 20.72 percentage points and reduces CR by 10.46 percentage points.

4.3 Q2. Ablation Study

To examine the contribution of each component, we compare the full ASR-VLA with two ablated variants. *ASR-VLA w/o Safety Reasoning* removes R_t while retaining $\hat{I}_t^{\text{risk}}$. *ASR-VLA w/o Risky-Future Image* retains R_t but removes $\hat{I}_t^{\text{risk}}$ from the VLA policy input. As shown in Table 1, removing R_t reduces the overall SR from 93.68% to 91.74% and increases CR from 9.46% to 14.74%. This indicates that $\hat{I}_t^{\text{risk}}$ alone is insufficient for the policy to determine an appropriate response to potential hazards. Removing $\hat{I}_t^{\text{risk}}$ also degrades performance, yielding 91.60% SR and 13.98% CR. Without $\hat{I}_t^{\text{risk}}$, the policy lacks visual evidence of where and how the hazards will emerge, making R_t less precise. These results demonstrate that both R_t and $\hat{I}_t^{\text{risk}}$ are necessary for effective safety-aware action generation.

4.4 Q3. Safety Grounding

We conduct a controlled prompt-intervention experiment to verify whether the VLA policy actually grounds the safety reasoning in action generation. Using the same trained policy, we replace only the reasoning input while keeping the current observations, and all other policy inputs unchanged. We consider five conditions. *No Reasoning*, which removes the reasoning input; *Always Safe*, *Always Warning* and *Always Danger*, which replace the predicted reasoning with fixed reasoning inputs; *Predicted Reasoning*, which uses the online reasoning generated by ASR-VLA.

As shown in Fig. 4, removing the reasoning input produces the highest CR. Replacing it with a fixed `SAFE` or `DANGER` input reduces the CR to 15.3% and 12.7%, respectively, showing that changes in the reasoning input affect the policy's action generation. However, these fixed inputs cannot adapt to changes during task execution. Predicted reasoning achieves the lowest CR corresponding to a 60% relative reduction compared with the no-reasoning condition. These results show that the reasoning input influences action generation and that adaptive reasoning is more effective than fixed safety inputs. We further examine whether the predicted risk signal becomes active before physical contact. A collision is considered anticipated if $\hat{\rho}_t \geq \tau_{\text{warn}}$ at any step $t < t_{\text{col}}$, where t_{col} denotes the first collision step. As shown in Fig. 8 (in Appendix B), the model performs particularly well when the hazard develops progressively during the approach or transport phase. In these cases, the robot-obstacle clearance decreases over multiple replanning steps, providing the world model with sufficient visual evidence to identify the emerging risk before contact occurs. In contrast, missed warnings are more common in tightly constrained spatial configurations, where a small lateral deviation or brief occlusion can cause the clearance to collapse within a single action chunk. Such abrupt, geometry-sensitive contacts leave less time for the estimated risk score to cross the warning threshold. These results suggest that ASR-VLA is effective at anticipating gradual and visually observable hazards, while sudden contacts in narrow configurations remain more challenging.

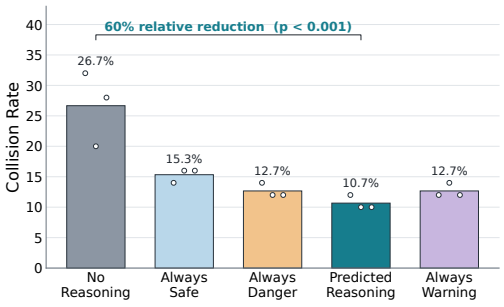
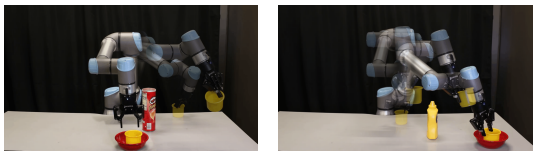


Figure 4: Collision rates under controlled interventions on the safety reasoning input. Each condition is evaluated over three random seeds with 50 episodes per seed. The p-value is from a two-sided two-proportion z-test.

4.5 Q4. Hardware Demonstration



(a) (b)

Figure 5: Real-world robot results. (a) Cup-to-bowl from left to right. (b) Mirrored cup-to-bowl from right to left. Semi-transparent poses indicate earlier execution stages, and the opaque pose indicates the final configuration.

To evaluate the real-world applicability of ASR-VLA, we apply the same training pipeline in a matched UR5e simulation environment and deploy the resulting simulation-trained policy on a physical UR5e manipulator equipped with a Robotiq 2F-85 gripper. Representative execution trajectories are shown in Fig. 5. Additional details of the matched simulation environment and training setup are provided in Appendix D. The sim-to-real transfer mainly requires bridging differences in perception and low-level execution between simulation and the physical platform. Following SafeVLA (Zhang et al., 2026a), we use FoundationPose (Wen et al., 2024) to estimate the 6D poses of task-relevant objects from real-camera observations. These pose estimates are combined with robot proprioception to construct the same structured observation representation used by the policy during simulation. We further maintain a consistent control interface across simulation and deployment. Both platforms use Cartesian end-effector actions, which are converted into robot commands through the same inverse kinematics pipeline. The simulation environment uses the corresponding UR5e-Robotiq embodiment, reducing the kinematic discrepancy and allowing the predicted actions to be executed without an additional retargeting module. The hardware demonstrations show that the simulation-trained policy can reproduce its safety-aware manipulation behavior under real-world perception and control conditions.

4.6 Q5: Robustness to Execution Perturbations



Figure 6: Robustness under large execution perturbations. The evaluation uses scenarios reliably completed under nominal execution, isolating the effect of safety reasoning after unexpected action deviations.

We design this experiment as a controlled stress test of safety-aware recovery. Specifically, we select scenarios in which the policy can reliably complete the task without collisions under nominal execution. We then apply lateral perturbations of the same magnitude in both directions, each lasting four consecutive control steps. This setup isolates the effect of safety reasoning after an unexpected execution deviation, rather than differences in the policy's nominal task-solving capability. As shown in Fig. 6, the perturbation causes the variant without reasoning to enter unsafe states, leading to a substantial increase in CR and, in some cases, failed task execution. In contrast, ASR-VLA consistently suppresses the increase in CR across all four suites while preserving a similar level of SR. This distinction indicates the safety improvement is not achieved by simply stopping the robot or sacrificing task progress. Instead, the adaptive reasoning helps the policy adjust its subsequent motion, recover clearance from nearby obstacles, and continue the task.

5 Conclusion

In this work, we presented ASR-VLA, an adaptive safety reasoning framework that integrates a reasoning module into VLA action generation. By incorporating the predicted high-risk future observation and task-aware safety guidance into the VLA policy, ASR-VLA improves SR while reducing CR. Experiments on 25 SafeLIBERO tasks show that ASR-VLA achieves 93.68% SR with 9.46% CR, outperforming evaluated standard and safety-aware VLA baselines. Ablation and intervention studies further demonstrate the contributions of both future-risk prediction and safety reasoning. Future work will focus on reducing reasoning latency and extending the framework to broader safety conditions.

AI USE STATEMENT

Generative AI tools were used to assist with polishing the writing and improving the clarity of the manuscript. AI tools were also used to assist with code generation for parts of the experimental implementation. All AI-generated code was manually reviewed, verified, and, where necessary, modified by the authors before being used in the experiments. The authors take full responsibility for the final manuscript and implementation.

ETHICS STATEMENT

This work studies safety-aware vision-language-action policies for robotic manipulation in simulation and on a physical robotic platform. The study does not involve human participants, personally identifiable information, or sensitive data. The real-world experiments are conducted in controlled laboratory settings using standard robotic hardware. While the proposed method is designed to reduce collision risk, the reported results should not be interpreted as providing formal safety guarantees for deployment in uncontrolled real-world environments.

REFERENCES

Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International conference on machine learning*, pp. 22–31. Pmlr, 2017.

Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8): 3861–3876, 2016.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*, 2025.

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0: A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.

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

Lukas Brunke, Yanni Zhang, Ralf Römer, Jack Naimer, Nikola Staykov, Siqi Zhou, and Angela P Schoellig. Semantically safe robot manipulation: From semantic scene understanding to motion safeguards. *IEEE Robotics and Automation Letters*, 2025.

Hao Chen, Jiaming Liu, Zhonghao Yan, Nuowei Han, Renrui Zhang, Chenyang Gu, Jialin Gao, Ziyu Guo, Siyuan Qian, Yinxi Wang, et al. Last-r1: Reinforcing robotic manipulation via adaptive physical latent reasoning. *arXiv preprint arXiv:2604.28192*, 2026a.

Zixing Chen, Yifeng Gao, Li Wang, Yunhan Zhao, Yi Liu, Jiayu Li, Xiang Zheng, Zuxuan Wu, Cong Wang, Xingjun Ma, et al. Hazardarena: Evaluating semantic safety in vision-language-action models. *arXiv preprint arXiv:2604.12447*, 2026b.

Rongxu Cui, Zongzheng Zhang, Jingrui Pang, Haohan Chi, Jinbang Guo, Saining Zhang, Shaoxuan Xie, Xin Jin, Yao Mu, Jiaolong Yang, et al. Libero-safety: A comprehensive benchmark for physical and semantic safety in vision-language-action models. *arXiv preprint arXiv:2606.23686*, 2026.

Jialiang Fan, Weizhe Xu, Oleg Sokolsky, Insup Lee, and Fanxin Kong. Safevla-bench: A benchmark for the success-safety gap in vision-language-action models. *arXiv preprint arXiv:2606.00773*, 2026.

Mingjian Gao, Wenqiao Zhang, Yuqian Yuan, Yang Dai, Binhe Yu, Zheqi Lv, Haoyu Zheng, Jiaqi Zhu, Zhiqi Ge, Zixuan Wan, et al. Visualthink-vla: Visual intermediate reasoning for effective and low-latency vision-language-action policies. *arXiv preprint arXiv:2605.30011*, 2026.

Qiao Gu, Yuanliang Ju, Shengxiang Sun, Igor Gilitschenski, Haruki Nishimura, Masha Itkina, and Florian Shkurti. Safe: Multitask failure detection for vision-language-action models. *Advances in Neural Information Processing Systems*, 38:40041–40076, 2026.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Liang Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *Iclr*, 1(2):3, 2022.

Songqiao Hu, Zeyi Liu, Shuang Liu, Jun Cen, Zihan Meng, and Xiao He. Vlsa: Vision-language-action models with plug-and-play safety constraint layer. *arXiv preprint arXiv:2512.11891*, 2025.

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

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

Zijun Lin, Jiafei Duan, Haoquan Fang, Dieter Fox, Ranjay Krishna, Cheston Tan, and Bihan Wen. Failsafe: Reasoning and recovery from failures in vision-language-action models. *arXiv preprint arXiv:2510.01642*, 2025.

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

Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems*, 36:44776–44791, 2023.

Zhuoyang Liu, Jiaming Liu, Hao Chen, Jiale Yu, Ziyu Guo, Chengkai Hou, Chenyang Gu, Xiangju Mi, Renrui Zhang, Kun Wu, et al. Last _$\{0\}$: Latent spatio-temporal chain-of-thought for robotic vision-language-action model. *arXiv preprint arXiv:2601.05248*, 2026.

Xingjun Ma, Yifeng Gao, Yixu Wang, Ruofan Wang, Xin Wang, Ye Sun, Yifan Ding, Hengyuan Xu, Yunhao Chen, Yunhan Zhao, et al. Safety at scale: A comprehensive survey of large model and agent safety. *Foundations and Trends in Privacy and Security*, 8(3-4):1–240, 2026.

Pietro Mazzaglia, Cansu Sancaktar, Markus Peschl, and Daniel Dijkman. Hybrid training for vision-language-action models. *arXiv preprint arXiv:2510.00600*, 2025.

Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the 24th international conference on Machine learning*, pp. 745–750, 2007.

Allen Ren, Justin Lidard, Lars Ankile, Anthony Simeonov, Pulkit Agrawal, Anirudha Majumdar, Benjamin Burchfiel, Hongkai Dai, and Max Simchowitz. Diffusion policy policy optimization. In *International Conference on Learning Representations*, volume 2025, pp. 77288–77329, 2025.

Quanxin Shou, Fangqi Zhu, Shawn Chen, Puxin Yan, Zhengyang Yan, Yikun Miao, Xiaoyi Pang, Zicong Hong, Ruikai Shi, Hao Huang, et al. Halo: A unified vision-language-action model for embodied multimodal chain-of-thought reasoning. *arXiv preprint arXiv:2602.21157*, 2026.

Kai Tang, Peidong Jia, Zhong Chu, Jixian Wu, Rui Ma, Jiajun Cao, Fangyuan Zhao, Sixiang Chen, Yichen Guo, Xiaowei Chi, et al. Safedojo: Safe reinforcement learning for vla via interactive world model. *arXiv preprint arXiv:2606.20698*, 2026.

Maximilian Tölle, Theo Gruner, Daniel Palenicek, Tim Schneider, Jonas Günster, Joe Watson, Davide Tateo, Puze Liu, and Jan Peters. Towards safe robot foundation models using inductive biases. *arXiv preprint arXiv:2505.10219*, 2025.

Alexander Tong, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Kilian Fatras, Guy Wolf, and Yoshua Bengio. Conditional flow matching: Simulation-free dynamic optimal transport. *arXiv preprint arXiv:2302.00482*, 2(3), 2023.

Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. Foundationpose: Unified 6d pose estimation and tracking of novel objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 17868–17879, 2024.

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

Borong Zhang, Yuhao Zhang, Jiaming Ji, Yingshan Lei, Juntao Dai, Yuanpei Chen, and Yaodong Yang. Safevla: Towards safety alignment of vision-language-action model via constrained learning. *Advances in Neural Information Processing Systems*, 38:153335–153373, 2026a.

Dapeng Zhang, Jing Sun, Chenghui Hu, Xiaoyan Wu, Zhenlong Yuan, Rui Zhou, Fei Shen, and Qingguo Zhou. Pure vision language action (vla) models: A comprehensive survey. *arXiv preprint arXiv:2509.19012*, 2025.

Heng Zhang, Wei-Hsing Huang, Qiyi Tong, Gokhan Solak, Puze Liu, Kaidi Zhang, Sheng Liu, Jan Peters, Yu She, and Arash Ajoudani. Compliantvla-adaptor: Vlm-guided variable impedance action for safe contact-rich manipulation. *arXiv preprint arXiv:2601.15541*, 2026b.

Yuhao Zhang, Borong Zhang, Jiaming Fan, Jiachen Shen, Yishuai Cai, Yaodong Yang, and Jiaming Ji. Redvla: Physical red teaming for vision-language-action models. *arXiv preprint arXiv:2604.22591*, 2026c.

Qingqing Zhao, Yao Lu, Moo Jin Kim, Zipeng Fu, Zhuoyang Zhang, Yecheng Wu, Zhaoshuo Li, Qianli Ma, Song Han, Chelsea Finn, et al. Cot-vla: Visual chain-of-thought reasoning for vision-language-action models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 1702–1713, 2025.

Linqing Zhong, Yi Liu, Yifei Wei, Ziyu Xiong, Si Liu, and Guanghui Ren. Acot-vla: Action chain-of-thought for vision-language-action models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8152–8162, 2026.

Zhongyi Zhou, Yichen Zhu, Xiaoyu Liu, Zhibin Tang, Junjie Wen, Yaxin Peng, Chaomin Shen, and Yi Xu. Chatvla-2: Vision-language-action model with open-world reasoning. *Advances in Neural Information Processing Systems*, 38:45537–45559, 2026.

Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, pp. 2165–2183. PMLR, 2023.

A Benchmark and Environmental Settings

A.1 SafeLIBERO Benchmark

We evaluate all policies on SafeLIBERO (Hu et al., 2025), a safety-critical extension of the LIBERO manipulation benchmark. LIBERO consists of four complementary task suites: *Spatial*, which tests spatial-relation understanding; *Goal*, which varies the desired goal while retaining related scene semantics; *Object*, which varies the manipulated object; and *Long*, which requires multi-stage, long-horizon execution. SafeLIBERO augments these environments with movable objects that can obstruct the robot's motion, making it possible to evaluate task completion and collision jointly within each evaluation episode.

SafeLIBERO defines two obstacle-intervention levels. In **Level I**, an obstacle is placed close to the target object and primarily interferes with approach and grasping. In **Level II**, the obstacle is placed farther from the target but obstructs the robot's transport or placement path. The benchmark randomizes the poses of task objects and obstacles within a bounded range. The obstacle set includes visually and geometrically diverse objects, such as a moka pot, storage box, milk carton, wine bottle, coffee mug, and book. These variations produce 50 predefined initial configurations for each task-level scenario.

Our evaluation covers 25 SafeLIBERO task-level scenarios, as listed in Table 2. All aggregate results in the main paper and Appendix B are obtained from evaluations on these 25 scenarios unless explicitly stated otherwise.

Suite	Scenarios	Rollouts/seed	Total rollouts
Goal	7	350	1,750
Spatial	6	300	1,500
Object	8	400	2,000
Long	4	200	1,000
Total	25	1,250	6,250

Table 2: Scope of the standardized local evaluation. Each scenario is evaluated using 50 initial configurations and five outer seeds, i.e., 250 rollouts per method.

A.2 Task Definitions

Table 3 lists every task-level scenario used in our standardized evaluation. The same task instruction appearing at Levels I and II represents two different safety scenarios because their obstacle placements and resulting collision risks differ. Task identifiers are assigned separately within each suite and level, rather than globally across the benchmark.

A.3 Observation and Action Spaces

At time step t, the policy receives two RGB observations: a third-person agent-view image I_t^{agent} and an eye-in-hand wrist-view image I_t^{wrist}. The simulator renders both views at 256×256 pixels; they are resized with padding to 224×224 before being passed to the policy. The proprioceptive state is $x_t = [p_t^{\text{eef}}, r_t^{\text{eef}}, q_t^{\text{gripper}}]$, where $p_t^{\text{eef}} \in \mathbb{R}^3$ is the end-effector position, $r_t^{\text{eef}} \in \mathbb{R}^3$ is its axis-angle orientation, and $q_t^{\text{gripper}} \in \mathbb{R}^2$ contains the two gripper joint positions. A task instruction ℓ specifies the task. Depending on the evaluated variant, the policy may additionally receive the current subtask and structured safety guidance.

The policy predicts a chunk of 7-DoF continuous actions. Each action is represented as

$$a_t = [\Delta x_t, \Delta y_t, \Delta z_t, \Delta r_t^x, \Delta r_t^y, \Delta r_t^z, a_t^{\text{grip}}], \tag{21}$$

where the first six dimensions specify the relative end-effector translation and rotation, and a_t^{grip} controls the gripper. We execute the first five actions of the predicted chunk and then query the policy again, providing closed-loop replanning every five simulator steps.

Table 3: The 25 SafeLIBERO task-level scenarios in our matched evaluation.

Suite	Level	ID	Task instruction
Goal	I	0	Put the bowl on the plate.
Goal	I	1	Put the bowl on top of the cabinet.
Goal	I	2	Put the bowl on the stove.
Goal	II	0	Put the bowl on the plate.
Goal	II	1	Put the bowl on top of the cabinet.
Goal	II	2	Put the bowl on the stove.
Goal	II	3	Put the cream cheese in the bowl.
Spatial	I	0	Pick up the black bowl between the plate and the ramekin and place it on the plate.
Spatial	I	2	Pick up the black bowl on the stove and place it on the plate.
Spatial	II	0	Pick up the black bowl between the plate and the ramekin and place it on the plate.
Spatial	II	1	Pick up the black bowl on the ramekin and place it on the plate.
Spatial	II	2	Pick up the black bowl on the stove and place it on the plate.
Spatial	II	3	Pick up the black bowl on the wooden cabinet and place it on the plate.
Object	I	0	Pick up the orange juice and place it in the basket.
Object	I	1	Pick up the chocolate pudding and place it in the basket.
Object	I	2	Pick up the milk and place it in the basket.
Object	I	3	Pick up the BBQ sauce and place it in the basket.
Object	II	0	Pick up the orange juice and place it in the basket.
Object	II	1	Pick up the chocolate pudding and place it in the basket.
Object	II	2	Pick up the milk and place it in the basket.
Object	II	3	Pick up the BBQ sauce and place it in the basket.
Long	I	0	Put both the alphabet soup and the cream cheese box in the basket.
Long	I	1	Put both the alphabet soup and the tomato sauce in the basket.
Long	II	0	Put both the alphabet soup and the cream cheese box in the basket.
Long	II	1	Put both the alphabet soup and the tomato sauce in the basket.

A.4 Evaluation Protocol and Metrics

For each task-level scenario, we evaluate each of the 50 predefined initial configurations under five random seeds. This yields 250 rollouts per scenario and 6,250 episodes for our method across 25 task-level scenarios.

We use the benchmark's displacement-based collision proxy. Let p_0^{obs} and q_0^{obs} denote the initial position and orientation of the active obstacle. At each simulator step t, we compare the obstacle's current pose with its initial pose. A raw collision flag is raised if the Euclidean distance between its current and initial positions exceeds $0.001\,\mathrm{m}$, or if the angular difference between its current and initial orientations exceeds $0.20\,\mathrm{rad}$. To suppress transient simulator noise, the condition must hold for three consecutive simulator steps. Once triggered, the collision flag remains true for the remainder of the episode.

For a set of N rollouts, let $S_i \in \{0,1\}$ denote task success and $C_i \in \{0,1\}$ denote whether a collision occurred. We report task success rate (SR), collision rate (CR): $\mathrm{SR} = \frac{1}{N}\sum_{i=1}^{N} S_i, \mathrm{CR} = \frac{1}{N}\sum_{i=1}^{N} C_i$.

For suite-level results, we first compute each metric over all tasks and initializations within each suite separately for the five random seeds, and then report the mean and standard deviation across seeds.

B Additional Results

The Additional Results section supplements the main evaluation with analyses of inference efficiency, anticipatory risk detection, per-scenario performance, world-model prediction quality, and representative failure modes, together with comparisons with reward-based post-training baselines.

B.1 Q6. Inference Efficiency

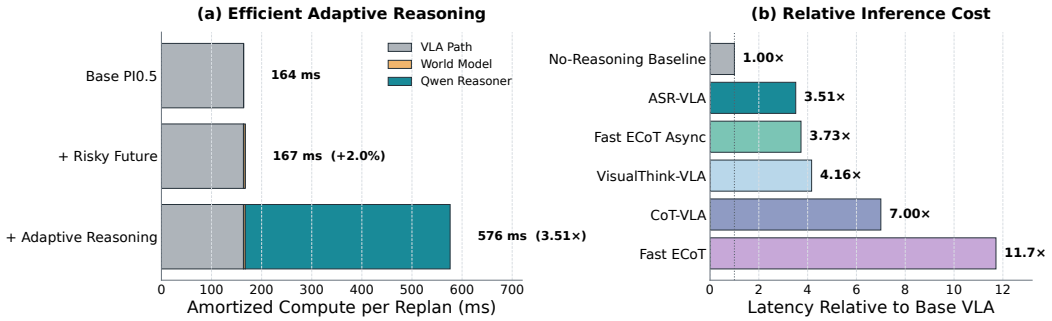


Figure 7: Inference-efficiency analysis of adaptive safety reasoning. (a) Amortized latency per replanning step for the base VLA. (b) End-to-end latency relative to the base VLA.

Using a vision-language reasoning model introduces additional inference latency, which remains an important challenge for reasoning-based robotic policies. ASR-VLA reduces this overhead through event-triggered adaptive reasoning. At every replanning step, the world model predicts the safety-critical future observation and its associated risk score. The estimated risk score is converted into one of three discrete safety states: SAFE, WARNING, or DANGER. The reasoning model is invoked at the beginning of each episode and whenever the predicted safety state differs from that of the preceding replanning step. When the safety state remains unchanged, ASR-VLA reuses the most recently generated structured reasoning output while the world model and VLA policy continue to operate at every replanning step. This design avoids redundant reasoning calls during periods in which the estimated safety condition remains stable, while immediately updating the guidance when the predicted risk level changes.

As shown in Fig. 7(a), the base $\pi_{0.5}$ policy requires 164 ms per replanning step. Adding the safety-critical world model increases the latency to 167 ms, corresponding to only a 1.8% overhead. When adaptive safety reasoning is included, the complete ASR-VLA pipeline requires an average amortized latency of 576 ms per replanning step, or $3.51\times$ that of the base policy. This trajectory-level average includes both replanning steps that trigger a new reasoning update and steps that reuse the cached reasoning output. The results indicate that the vision-language reasoning model remains the dominant source of additional computation, while event-triggered invocation substantially reduces unnecessary repeated reasoning.

Figure 7(b) compares the normalized end-to-end inference cost with representative reasoning-based VLA methods. ASR-VLA has a relative latency of $3.51\times$, compared with $3.73\times$ for Fast ECoT Async (Zawalski et al., 2024), $4.16\times$ for VisualThink-VLA (Gao et al., 2026), $7.00\times$ for CoT-VLA (Zhao et al., 2025), and $11.7\times$ for Fast ECoT (Zawalski et al., 2024). These results demonstrate that structured, event-triggered safety reasoning provides a favorable trade-off between adaptive hazard response and inference efficiency.

B.2 Additional Safety Grounding Results

We retrospectively examine episodes that eventually result in a collision. A collision is considered anticipated if the estimated risk score exceeds the warning threshold ($\tau_{\text{warn}} = 0.65$) at any control step before the first detected collision. Specifically, for each collision episode, we identify the first collision step t_{col} and determine whether $t < t_{\text{col}}$ such that $\hat{\rho}_t \geq \tau_{\text{warn}}$. The green portion denotes collisions preceded by a warning, while the gray portion denotes missed collisions. As shown in Fig. 8, ASR-VLA anticipates a substantial fraction of collisions across six representative Level-II task scenarios from all four suites. The detection rate ranges from 72.5% to 100.0%, and exceeds 84% on five of the six scenarios. In particular, all collision episodes in Object Level-II Task 0 are preceded by a warning. These results indicate that the predicted risk signal frequently becomes active before physical contact, providing advance warning for safety-aware action generation rather than merely identifying collisions after they occur.

Additionally, to further examine whether the policy relies on the temporal alignment between future-risk evidence and structured reasoning, we compare four conditions. *Predicted Alignment* uses

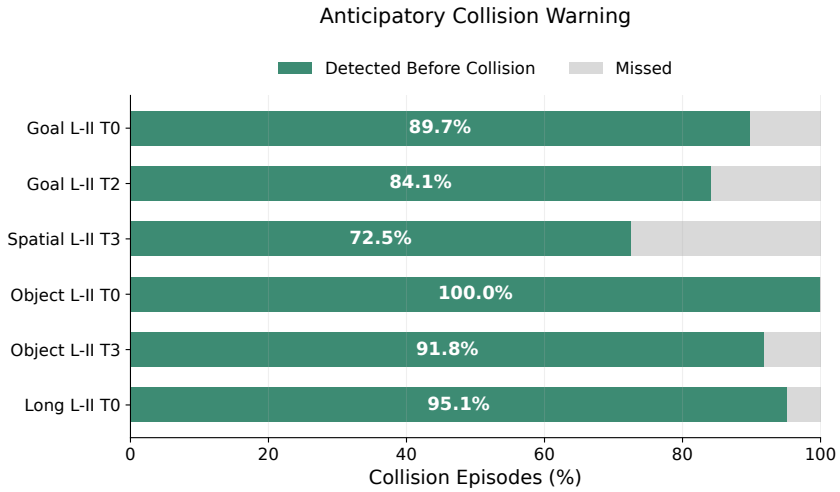


Figure 8: Anticipatory collision-warning performance on representative SafeLIBERO tasks.

the future image and reasoning generated for the current state. *Count-Matched Shuffled* preserves the same number of reasoning updates but randomly reassigns them to different states. *Delayed-5* applies each reasoning output five replanning steps after it is generated, while *Always Danger* replaces the adaptive reasoning with a fixed danger warning throughout the episode.

As shown in Fig. 9, the correctly aligned reasoning consistently achieves the highest average collision-free success rate and best on three of four tasks. Randomly shuffling or delaying the reasoning degrades performance, even though the amount of reasoning remains unchanged. These results indicate that the benefit of safety reasoning does not arise merely from providing additional language context. Instead, the reasoning must remain temporally aligned with the predicted safety-critical future in order to provide effective guidance for action generation.

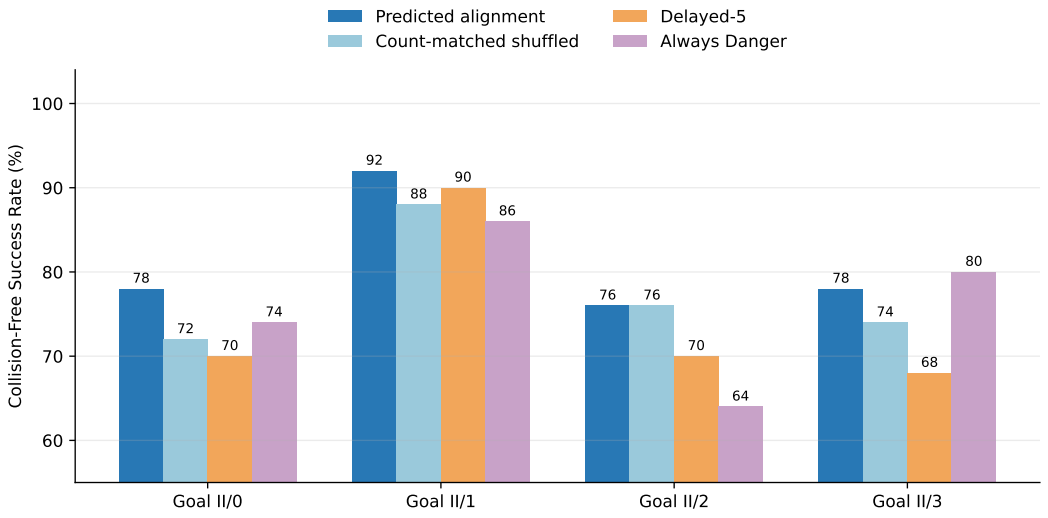


Figure 9: Temporal-alignment study of structured safety reasoning on the four Goal Level-II tasks. *Predicted Alignment* uses reasoning generated for the current state. *Count-Matched Shuffled* randomly reassigns reasoning outputs while preserving the same number of reasoning updates. *Delayed-5* applies each reasoning output five replanning steps after it is generated, and *Always Danger* uses a fixed danger warning throughout the episode. Correct temporal alignment achieves the highest average collision-free success rate and best on three of four tasks.

Suite	Scenario	Short task	N	SR ↑	CR ↓	CSR ↑
Goal	I/0	Bowl → plate	250	97.60±0.80	2.40±0.80	96.40±1.50
Goal	I/1	Bowl → cabinet	250	95.20±2.40	12.40±3.44	82.80±5.31
Goal	I/2	Bowl → stove	250	98.40±1.96	6.80±1.60	92.40±1.96
Goal	II/0	Bowl → plate	250	95.60±1.96	15.60±1.50	83.20±0.98
Goal	II/1	Bowl → cabinet	250	95.60±1.96	2.00±1.26	94.80±1.60
Goal	II/2	Bowl → stove	250	98.00±2.19	17.60±1.50	82.00±1.26
Goal	II/3	Cream cheese → bowl	250	98.80±0.98	25.20±4.83	74.40±4.63
Spatial	I/0	Between objects → plate	250	98.40±1.50	13.60±2.33	86.40±2.33
Spatial	I/2	Stove → plate	250	93.60±4.08	6.40±2.33	88.00±5.37
Spatial	II/0	Between objects → plate	250	99.60±0.80	6.80±2.40	93.20±2.40
Spatial	II/1	Ramekin → plate	250	94.40±4.80	14.40±3.44	82.00±6.20
Spatial	II/2	Stove → plate	250	96.80±1.60	0.40±0.80	96.80±1.60
Spatial	II/3	Cabinet → plate	250	95.20±2.99	32.00±1.79	66.40±1.96
Object	I/0	Orange juice → basket	250	96.00±2.83	1.60±2.33	95.20±3.71
Object	I/1	Chocolate pudding → basket	250	97.20±0.98	0.00±0.00	97.20±0.98
Object	I/2	Milk → basket	250	83.20±4.12	3.60±3.44	80.00±5.22
Object	I/3	BBQ sauce → basket	250	99.20±0.98	0.00±0.00	99.20±0.98
Object	II/0	Orange juice → basket	250	89.20±2.40	12.00±2.83	80.00±4.56
Object	II/1	Chocolate pudding → basket	250	98.80±0.98	0.00±0.00	98.80±0.98
Object	II/2	Milk → basket	250	95.60±1.50	0.40±0.80	95.20±2.04
Object	II/3	BBQ sauce → basket	250	86.40±3.88	24.40±0.80	73.60±1.50
Long	I/0	Soup+cheese → basket	250	97.20±2.99	3.60±2.33	96.40±2.33
Long	I/1	Soup+tomato → basket	250	84.00±4.56	6.40±2.65	82.80±5.00
Long	II/0	Soup+cheese → basket	250	90.40±2.94	16.40±2.33	78.00±3.35
Long	II/1	Soup+tomato → basket	250	67.60±2.65	12.40±6.62	67.60±2.65

Table 4: Complete per-task results for ASR-VLA. SR denotes task success rate, CR denotes collision rate, and CSR denotes collision-free success rate. All values are percentages reported as mean ± population standard deviation over five outer seeds.

B.3 Results per Task

Table 4 reports the performance of ASR-VLA on each of the 25 SafeLIBERO task-level scenarios. Each task-level scenario is evaluated using 50 initial configurations under five random seeds, yielding $N = 250$ rollouts per row. We report task success rate (SR), collision rate (CR), and collision-free success rate (CSR). Values are the mean and standard deviation across the five random seeds. The complete task instructions corresponding to the scenario identifiers are provided in Table 3.

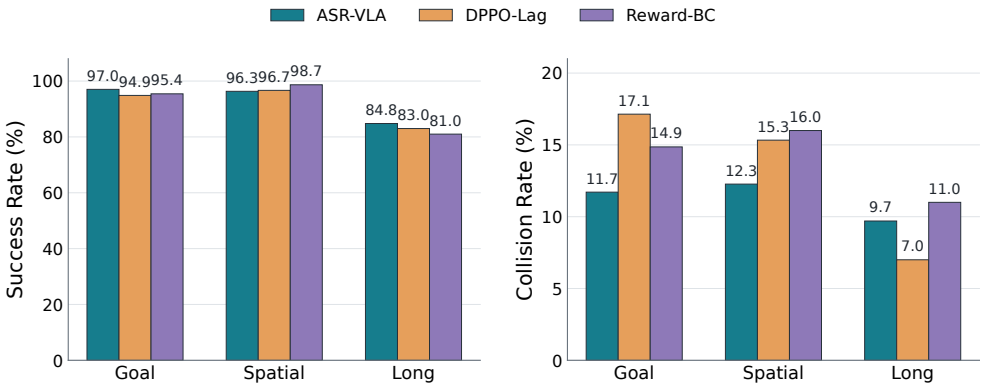


Figure 10: Suite-level comparison between ASR-VLA and two short-budget reward-driven variants. Left: SR (higher is better). Right: CR (lower is better). Object is excluded from this comparison. The RL variants are evaluated using five initializations per task and five outer seeds.

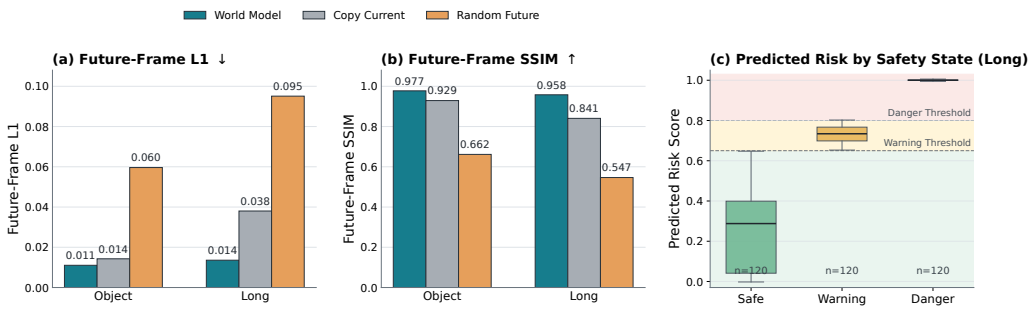


Figure 11: World-model prediction analysis on in-domain validation frames. (a) Future-frame L1 error, where lower is better. (b) Future-frame structural similarity (SSIM) (Wang et al., 2004), where higher is better. The learned world model is compared with copying the current observation and selecting a random future frame. (c) Distribution of estimated risk scores conditioned on the ground-truth Safe, Warning, and Danger states in the balanced Long validation split. Dashed lines denote the 0.65 Warning and 0.80 Danger thresholds. Each risk state contains 120 samples.

B.4 Reinforcement Learning Results

We additionally investigate two RL-based variants to examine the effect of RL post-training.

DPPO-Lagrangian (DPPO-Lag) applies a clipped PPO-style update to the diffusion policy (Achiam et al., 2017; Ren et al., 2025). Its constrained advantage is defined as $A = R - \lambda C$, where $R = 1$ for task success and $R = -0.2$ otherwise, while $C = 1$ denotes a collision. The non-negative Lagrange multiplier λ is updated toward a collision-cost limit of 0.1. Since the diffusion policy does not provide a tractable action likelihood, we use the negative per-sample diffusion loss as a log-probability proxy in the PPO ratio.

Reward-weighted BC (Reward-BC) is a reward-driven baseline rather than a policy-gradient method (Peters & Schaal, 2007). It minimizes the standard diffusion behavior-cloning loss on replayed action chunks, while weighting each sample according to its episode outcome. The weights are 2.0 for collision-free success, 0.8 for success with a collision, 0.4 for failure without a collision, and 0.1 for collision with failure. Consequently, safe successful trajectories contribute most strongly to the update.

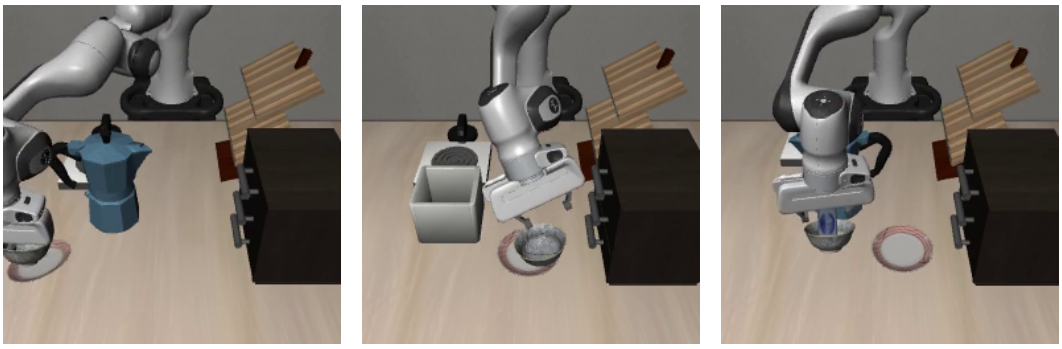
As shown in Fig. 10, neither RL-based variant consistently dominates ASR-VLA. On the *Goal* suite, ASR-VLA achieves both higher SR (97.03%) and lower CR (11.71%) than DPPO-Lag (94.86%, 17.14%) and Reward-BC (95.43%, 14.86%). Overall, RL-based variants optimization improves individual metrics but does not provide a uniformly better success-safety trade-off. ASR-VLA remains competitive across all three suites without requiring online reward-based fine-tuning.

B.5 World Model Prediction Analysis

As shown in Fig. 11(a), the learned world model consistently outperforms both non-learned baselines. The improvement over Copy Current is particularly pronounced on the *Long* suite, indicating that the model captures task-progress-dependent scene evolution rather than relying on the strong temporal similarity between adjacent frames.

The SSIM results in Fig. 11(b) reinforce this observation. The predicted frames preserve the spatial structure of the risk-selected targets more faithfully than either copying the current observation or selecting an arbitrary future frame. This distinction is especially clear on long-horizon tasks, where the safety-critical future may differ substantially from the current scene as the robot progresses through multiple manipulation stages.

Figure 11(c) evaluates whether the predicted scalar risk remains aligned with the decision boundaries used by the downstream reasoning module. The predicted scores exhibit a clear ordering across the three safety states: `SAFE` samples remain largely below the warning threshold, `WARNING` samples concentrate within the intermediate-risk region, and `DANGER` samples are concentrated above the danger threshold. Only one of the 360 samples falls into an incorrect safety-state interval, showing



(a) (b) (c)

Figure 12: Representative failure cases of ASR-VLA.

that the risk-score prediction head reliably maintains the threshold-based safety-state partition on this in-domain split.

B.6 Failure Mode Analysis

In Fig. 12, we present representative failure cases and analyze their causes. Our qualitative analysis reveals two primary failure modes. First, some initial configurations require tightly constrained manipulation trajectories with little tolerance for control error. For example, in Fig. 12(a) and (c), the robot may need to fully extend its arm, transport an object through a narrow collision-free region, and release the object while it is still suspended above the target. In such cases, even a small positional deviation can cause the gripper to contact a nearby obstacle. These configurations remain challenging even when the model correctly identifies the potential risk, because the generated actions may not provide sufficient clearance.

Second, some tasks, such as Fig. 12(b), employ strict geometric success predicates. In several failed episodes, the manipulated object appeared to be correctly placed from visual inspection, but the simulator did not register task completion because the final position or orientation fell slightly outside the predefined success region. Consequently, the robot continued making corrective movements after an apparently successful placement. These unnecessary adjustments could displace the object or introduce a subsequent collision. Thus, a subset of the measured failures reflects not only deficiencies in safety reasoning, but also limited placement precision and sensitivity of the benchmark to small placement errors.

C Reasoning Output and Prompt

C.1 Input and Output

At each reasoning update, the vision-language reasoning model receives the current visual observation I_t, the predicted high-risk future observation $\hat{I}_t^{\text{risk}}$, the task instruction ℓ, and the estimated risk score $\hat{\rho}_t$. Based on these inputs, the model identifies the current manipulation phase, assesses the predicted hazard, and produces structured guidance for the downstream VLA policy.

The reasoning context is represented as

$$R_t = (\ell,\, c_t,\, q_t,\, g_t)\,,$$

where c_t denotes the inferred current subtask, q_t denotes the predicted safety status, and g_t denotes the corresponding safety guidance. The task field retains the original task instruction ℓ so that the complete reasoning sequence remains self-contained when supplied to the grounded policy.

We constrain the output to the following four-field format:

```
Task: {task instruction}
Subtask: {current manipulation phase}
Risk status: {safe, warning, danger}
Safety reasoning: {task- and scene-specific guidance}
```

Field	Symbol	Description
Task	ℓ	The original task instruction specifying the manipulation objective.
Subtask	c_t	The current manipulation phase inferred from the visual scene and task progress.
Risk status	q_t	A discrete assessment in $\{\texttt{safe}, \texttt{warning}, \texttt{danger}\}$ indicating the severity of the predicted future hazard.
Safety reasoning	g_t	Task- and scene-specific guidance describing how the robot should respond while preserving progress whenever possible.

Table 5: Fields in the structured reasoning sequence supplied to the grounded VLA policy.

This structured format separates high-level safety interpretation from low-level action generation. In particular, the reasoning model does not produce continuous actions or action tokens. Instead, it describes the current task phase, the severity of the predicted hazard, and an appropriate response that can be grounded by the downstream $\pi_{0.5}$ policy.

During supervised training, the target structured context is

$$R_t^\star = (\ell,\, c_t^\star,\, q_t^\star,\, g_t^\star)\,,$$

where $c_t^\star$ is obtained from the manipulation-phase annotation of the expert trajectory. The target safety status is constructed from the ground-truth risk score $\rho_t^\star$

$$q_t^\star = \begin{cases} \texttt{safe}, & \rho_t^\star < \tau_{\text{warn}}, \\ \texttt{warning}, & \tau_{\text{warn}} \le \rho_t^\star < \tau_{\text{danger}}, \\ \texttt{danger}, & \rho_t^\star \ge \tau_{\text{danger}}, \end{cases}$$

with $\tau_{\text{warn}} = 0.65$ and $\tau_{\text{danger}} = 0.80$. The guidance target $g_t^\star$ specifies an appropriate response conditioned on the task, current subtask, and predicted hazard. We construct $g_t^\star$ using a set of task- and phase-specific guidance templates defined according to the manipulation stage and safety state. The templates describe high-level corrective behaviors, such as increasing obstacle clearance, adjusting the approach direction, or delaying a motion before continuing the current subtask. The template-generated guidance is then manually reviewed and, when necessary, edited to ensure consistency with the scene geometry, task progress, and intended safety response. The resulting guidance remains at the level of semantic behavioral instruction and does not directly specify continuous robot actions.

Table 5 summarizes the four output fields and their roles.

C.2 Reasoning-Model Prompt

We initialize the reasoning model from Qwen3-VL-8B-Instruct and fine-tune it to generate the structured sequence in Eq. 12. Each training example contains two visual inputs: the current observation and the predicted high-risk future observation. The accompanying textual prompt provides the task instruction and the predicted scalar risk score.

The normalized textual prompt is:

```
You are the safety reasoning module for a robot manipulation
policy.
Analyze the current observation and the predicted high-risk
future observation in the context of the task.
Identify the current manipulation phase, assess the
predicted hazard, and provide safety guidance for the
low-level policy.

<Task>
{task instruction}
</Task>
```

```
<World Model Signal>
Estimated risk score: {estimated risk score}
Warning threshold: 0.65
Danger threshold: 0.80
</World Model Signal>

<Required Output>
Output exactly four fields in the following order:
Task
Subtask
Risk status
Safety reasoning
The safety reasoning must be specific to the current task
phase and the predicted hazard.
Do not output continuous actions, action tokens, or
additional text.
</Required Output>
```

For example, for the task of moving the milk into the basket, a warning-state target may take the form

```
Task: Pick up the milk and place it in the basket.
Subtask: Approach the milk while avoiding the wine bottle.
Risk status: warning
Safety reasoning: The predicted approach passes too close
to the wine bottle. Increase lateral clearance before
descending toward the milk, then continue the grasping
subtask.
```

The reasoning model therefore produces structured textual context only. Continuous robot actions are generated exclusively by the grounded VLA policy.

C.3 Prompt Supplied to the Grounded Policy

The grounded VLA policy receives the current agent-view image, the current wrist-view image, the robot proprioceptive state, and the world model's predicted high-risk future image. Its language input is the complete structured reasoning sequence R_t generated by the reasoning model:

```
Task: {task instruction}
Subtask: {predicted current subtask}
Risk status: {safe / warning / danger}
Safety reasoning: {task- and scene-specific guidance}
```

For example, the reasoning sequence above is supplied to the policy as

```
Task: Pick up the milk and place it in the basket.
Subtask: Approach the milk while avoiding the wine bottle.
Risk status: warning
Safety reasoning: The predicted approach passes too close
to the wine bottle. Increase lateral clearance before
descending toward the milk, then continue the grasping
subtask.
```

D Hyperparameters

This section summarizes the architecture, training, and deployment settings used for ASR-VLA. Hyperparameters shared across tasks are reported once, while task-specific settings are stated explicitly.

D.1 Supervised VLA Training

All locally trained policies use the $\pi_{0.5}$ backbone and predict 10-step chunks of 7-DoF continuous actions. We adapt both the 2B vision-language backbone and the 300M action expert using

Hyperparameter	ASR-VLA
Initialization	Task+Subtask checkpoint
Visual inputs	Agent view + wrist view + predicted risky future
Language inputs	Task + Subtask + Risk status + Safety reasoning
Training steps	1,000
Global batch size	32
FSDP devices	2
Optimizer	AdamW
Warmup steps	100
Learning rate	1×10^{-5}
Schedule horizon	2,000
Gradient clipping	1.0
LoRA rank / scale, 2B backbone	$16/16$
LoRA rank / scale, action expert	$32/32$
Predicted action horizon	10
Training seed	42

Table 6: Training settings for the supervised ASR-VLA policy.

low-rank adaptation (LoRA). Unless otherwise stated, training uses bfloat16 precision, the AdamW optimizer, no exponential moving average, and gradient-norm clipping at 1.0. Table 6 summarizes the corresponding training settings.

The reported batch sizes are global batch sizes. Goal Task+Subtask training uses four GPUs, whereas the remaining supervised policies use two GPUs. We retain the checkpoints after 5,000 optimization steps for Task-only and Task+Subtask and after 1,000 steps for ASR-VLA. LoRA is applied to the attention and feed-forward layers of both the vision-language backbone and the action expert. After warmup, the learning rate remains constant for the remainder of training.

D.2 World Model and Reasoning Model

The safety-critical world model receives the current visual observation I_t, robot proprioceptive state x_t, and task instruction ℓ. It jointly predicts the risk-selected future observation $\hat{I}_t^{\text{risk}}$ and its associated scalar risk estimate $\hat{\rho}_t$. Consistent with the training objective in Eq. 11, the world model is optimized using future-image reconstruction and risk-score regression losses.

The reasoning model receives the current observation, predicted high-risk future observation, task instruction, and predicted risk score. It generates the four-field structured reasoning sequence described in Sec. 3.3, containing the task, current subtask, risk status, and task-aware safety reasoning. The reasoning model generates textual context only and does not contain an action-prediction head. Their architecture and optimization settings are summarized in Table 7.

D.3 Deployment and Evaluation Settings

Table 8 summarizes the control, risk, and evaluation settings used in the local SafeLIBERO experiments.

E Hardware

All training and evaluation jobs were executed across two local compute servers. The primary server contains one AMD EPYC 7352 processor (24 physical cores and 48 hardware threads), 256 GB of system memory, and four NVIDIA RTX A5000 GPUs with 24 GB of memory per GPU. The secondary server contains one Intel Core i9-10980XE processor (18 physical cores and 36 hardware threads), 128 GB of system memory, and three NVIDIA GeForce RTX 3090 GPUs with 24 GB of memory per GPU. SafeLIBERO uses a simulated Franka Emika Panda manipulator, whereas the additional cross-embodiment MuJoCo validation uses a simulated Universal Robots UR5e manipulator equipped with a Robotiq 2F-85 gripper.

Component	Hyperparameter	Value
World model	Inputs	Current observation, robot proprioception, and task instruction
	Prediction targets	Risk-selected future observation and scalar risk score
	Internal image resolution	64×64
	Encoder channels	32, 64, 128, 256
	Fusion hidden dimension	512
	Optimizer	AdamW
	Batch size	64
	Learning rate	1×10^{-4}
	Training steps	2,000
	Split / training seed	5,307
	Gradient clipping	5.0
	Image / risk loss weights	$\lambda_{\text{img}} = 1/\lambda_{\text{risk}} = 1$
Qwen reasoner	Backbone	Qwen3-VL-8B-Instruct
	Visual inputs	Current observation and predicted high-risk future observation
	Textual inputs	Task instruction and predicted scalar risk score
	Output	Task, Subtask, Risk status, and Safety reasoning
	Objective	Autoregressive language modeling over the structured reasoning sequence
	Adaptation	LoRA on all linear layers
	LoRA rank / scale / dropout	$16/32/0.05$
	Model loading / compute precision	4-bit / bfloat16
	Optimizer	AdamW
	Learning rate	2×10^{-5}
	Weight decay	0
	Per-device batch size	1
	Gradient accumulation / effective batch	$64/64$
	Gradient clipping	1.0
	Training epochs	1
	Maximum sequence length	2,048 tokens

Table 7: Architecture and training settings for the safety-critical world model and the Qwen vision-language reasoning model.

Hyperparameter	Value	Hyperparameter	Value
Simulator rendering resolution	256×256	Risk temperature	0.01
Policy input resolution	224×224	Collision mode	Obstacle displacement
World-model internal resolution	64×64	Position threshold	0.001 m
Flow-matching sampling steps	10	Orientation threshold	0.20 rad
Predicted action horizon	10	Persistence requirement	3 consecutive simulator steps
Executed actions per replanning cycle	5	Maximum episode steps, Goal	280
World-model update interval	Every replanning cycle	Maximum episode steps, Spatial	280
Initial simulator settling	30 no-op steps	Maximum episode steps, Object	350
Warning / Danger thresholds	$\tau_{\text{warn}}/\tau_{\text{danger}} = 0.65/0.80$	Maximum episode steps, Long	620
Future-risk horizon	20	Initializations per scenario	50
Safe-clearance threshold	0.02	Independent outer seeds	5

Table 8: Control, safety, and evaluation settings used in the matched local SafeLIBERO experiments.

We evaluate ASR-VLA on a real-world platform consisting of a Universal Robots UR5e manipulator equipped with a Robotiq 2F-85 parallel gripper. The corresponding MuJoCo XML model is used during training, allowing the policy to be developed and validated under the same robot embodiment

before deployment. We consider three manipulation tasks: (1) moving a cup into a bowl from left to right, (2) the mirrored cup-to-bowl task from right to left, and (3) moving a cup into a basket.

For each task, we collect 300 expert demonstrations in the MuJoCo environment. To improve robustness, the demonstrations randomize both the initial gripper pose and the obstacle placement while preserving the manipulation objective. Based on these demonstrations, we fine-tune a nominal task policy initialized from $\pi_{0.5}$. Following the same procedure used in SafeLIBERO, this nominal policy is then executed in the randomized simulation environments to collect closed-loop trajectories for world-model training. Since the nominal policy is optimized for task completion without explicit safety reasoning, the resulting trajectories naturally contain safe, low-clearance, near-collision, and collision states. These trajectories are subsequently used to construct safety-critical future targets, continuous risk scores, and structured reasoning supervision for training the world model, reasoning model, and safety-grounded VLA policy.

During deployment, RGB observations from the real camera are processed using FoundationPose (Wen et al., 2024) to estimate the 6-DoF poses of task-relevant objects. These pose estimates are combined with robot proprioception to construct the same structured observations used during simulation training, thereby substantially reducing the visual domain gap between simulation and the real robot. The policy predicts Cartesian end-effector actions, which are converted into UR5e joint commands through inverse kinematics and executed by the robot controller. No additional policy fine-tuning or domain adaptation is performed during deployment.

F Algorithms

This section summarizes the two principal procedures of ASR-VLA: constructing and training the safety-critical world model, fine-tuning the structured safety reasoning model, and executing the complete closed-loop policy.

F.1 Safety-Critical World-Model Training

The world-model training data are constructed from closed-loop trajectories collected by a nominal VLA policy. For each trajectory state, we compute the minimum clearance between the robot geometry and the surrounding obstacles and convert it into a continuous risk score. Within the future horizon H_w, the observation with the highest risk is selected as the safety-critical prediction target. The world model is then trained to jointly predict this future observation and its associated scalar risk score.

F.2 Reasoning-Model Fine-Tuning

The reasoning model receives the current observation, the predicted high-risk future observation, the task instruction, and the predicted risk score. Its supervised target is a four-field structured reasoning sequence containing the task, current subtask, safety status, and corresponding safety guidance.

During training, the target subtask $c_t^\star$ is obtained from the phase annotation of the expert trajectory. The target safety status $q_t^\star$ is constructed from the ground-truth risk score $\rho_t^\star$ using the warning and danger thresholds. The reasoning model is optimized using an autoregressive language-modeling objective over the complete structured output.

Algorithm 1: Construction of safety-critical future targets and training of the world model.

Require: Closed-loop trajectory dataset $\mathcal{D}$ collected by $\pi_{0.5}$; future horizon H_w; safe-clearance threshold d_{safe}; temperature β; loss weights λ_{img} and λ_{risk}; world model W_ω
Ensure: Trained world-model parameters ω

Initialize the world-model training set $\mathcal{D}_{\text{WM}} \leftarrow \varnothing$
for all trajectories $\tau \in \mathcal{D}$ **do**
  Read the task instruction ℓ and trajectory length T
  for $j = 1, \ldots, T$ **do**
    Compute the risk score by Eq. 5
  end for
  for $t = 1, \ldots, T-1$ **do**
    Select the highest-risk future time step by Eq. 6
    Construct the visual and scalar supervision targets by Eq. 7
    Add $\left(I_t,\ x_t,\ \ell,\ I_t^{\text{risk}},\ \rho_t^\star\right)$ to $\mathcal{D}_{\text{WM}}$
  end for
end for
repeat
  Sample a minibatch $\mathcal{B} \subset \mathcal{D}_{\text{WM}}$
  Predict the safety-critical future and its risk score

$$\left(\hat{I}^{\text{risk}}, \hat{\rho}\right) \leftarrow W_\omega(I, \ell, x)$$

  Compute the minibatch-averaged training objective by Eq. 11.
  Update ω using $\nabla_\omega \mathcal{L}_{\text{WM}}$
until the world-model training budget is exhausted

Algorithm 2: Supervised fine-tuning of the structured safety reasoning model.

Require: Reasoning supervision dataset $\mathcal{D}_{\text{reason}}$; warning threshold τ_{warn}; danger threshold τ_{danger}; reasoning model G_ϕ
Ensure: LoRA-adapted reasoning parameters ϕ

Initialize the reasoning SFT set $\mathcal{S}_{\text{reason}} \leftarrow \varnothing$
for all samples $\left(I_t, \hat{I}_t^{\text{risk}}, \ell, \hat{\rho}_t, \rho_t^\star, c_t^\star, g_t^\star\right) \in \mathcal{D}_{\text{reason}}$ **do**
  if $\rho_t^\star \geq \tau_{\text{danger}}$ **then**
    $q_t^\star \leftarrow$ `danger`
  else if $\rho_t^\star \geq \tau_{\text{warn}}$ **then**
    $q_t^\star \leftarrow$ `warning`
  else
    $q_t^\star \leftarrow$ `safe`
  end if
  Construct the multimodal reasoning input

$$X_t^{\text{reason}} \leftarrow \text{FormatInput}\left(I_t, \hat{I}_t^{\text{risk}}, \ell, \hat{\rho}_t\right)$$

  Construct the four-field target sequence

$$R_t^\star \leftarrow \text{FormatOutput}(\ell, c_t^\star, q_t^\star, g_t^\star)$$

  Add $(X_t^{\text{reason}}, R_t^\star)$ to $\mathcal{S}_{\text{reason}}$
end for
repeat
  Sample $(X_t^{\text{reason}}, R_t^\star) \sim \mathcal{S}_{\text{reason}}$
  Tokenize the target sequence as

$$R_t^\star = \left(r_{t,1}^\star, \ldots, r_{t,L_t}^\star\right)$$

  Compute the autoregressive language-modeling loss

$$\mathcal{L}_{\text{reason}} \leftarrow -\sum_{j=1}^{L_t} \log p_\phi\left(r_{t,j}^\star \mid r_{t,1:j-1}^\star, X_t^{\text{reason}}\right)$$

  Update the Qwen LoRA parameters ϕ using AdamW
until the reasoning-model training budget is exhausted