ASYM-OPSD: MITIGATING MLLM HALLUCINATIONS VIA ASYMMETRIC ON-POLICY SELF-DISTILLATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Hallucination in Multimodal Large Language Models (MLLMs) remains a critical barrier to trustworthy deployment. Existing mitigation strategies, whether inference-time interventions such as contrastive decoding and attention reallocation, or training-based preference alignment, share a fundamental limitation: they rely on coarse-grained, sequence-level signals that cannot precisely localize the specific tokens responsible for hallucinations. We introduce ASYM-OPSD, an asymmetric on-policy self-distillation framework that derives dense, token-level reward signals from hierarchical textual descriptions of the input image that serve as privileged visual anchors. When a teacher model conditioned on anchored input evaluates the student's own rollouts, the per-token log-probability gap between teacher and student reveals which tokens are visually grounded and which are hallucinated. Our pipeline comprises three stages: visual anchor construction, anchor-guided warm-up that repairs hallucinated responses to initialize a teacher model, and reinforcement learning with group-relative token-level advantages and curriculum scheduling on anchor granularity. Experiments on Qwen3-VL-4B and InternVL3.5-8B show that ASYM-OPSD surpasses training-based and inference-time baselines across four hallucination benchmarks, mitigating long-form hallucinations by up to $3.2\times$ without compromising general VQA performance.

1 INTRODUCTION

Multimodal Large Language Models (MLLMs) have demonstrated unprecedented capabilities in multimodal tasks, ranging from grounded image captioning to complex visual reasoning Liu et al. (2023); Bai et al. (2025). However, these models remain susceptible to hallucination: they confidently fabricate non-existent objects based on strong language priors, and misattribute colors and spatial relationships despite visual evidence to the contrary Guan et al. (2024b); Sun et al. (2024). This unreliability significantly undermines the trustworthiness of these systems and poses a critical barrier to their deployment in real-world, safety-critical applications.

Prior mitigation strategies fall into three families. (i) **Contrastive decoding** constructs perturbed inputs and suppresses tokens whose probability is inflated by visual-unaware priors Fang et al. (2025); Chen et al. (2025); recent re-examinations attribute its reported gains to decoding artifacts rather than improved visual grounding— gains stem from blunt, unidirectional shifts of the output distribution that merely rebalance the model's yes/no bias, and from an adaptive plausibility constraint that degrades sampling into greedy search Yin et al. (2025). (ii) **Attention reallocation** methods Zhong et al. (2026); Wang et al. (2026) amplify attention weights toward visual tokens, yet when the image is misinterpreted during encoding, amplification only magnifies the erroneous signal rather than correcting it Yang et al. (2026). (iii) **Preference-based alignment** methods Zheng et al. (2026); Pan et al. (2026a) optimize the model with sequence-level or uniform reward signals. Because hallucinations are localized to a few tokens within an otherwise correct response, uniform supervision dilutes the penalty over grounded tokens, and the model learns only very slowly which tokens were at fault. All three families share one bottleneck: **none of them provides a fine-grained, token-level signal that separates visually grounded tokens from hallucinated ones**.

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.

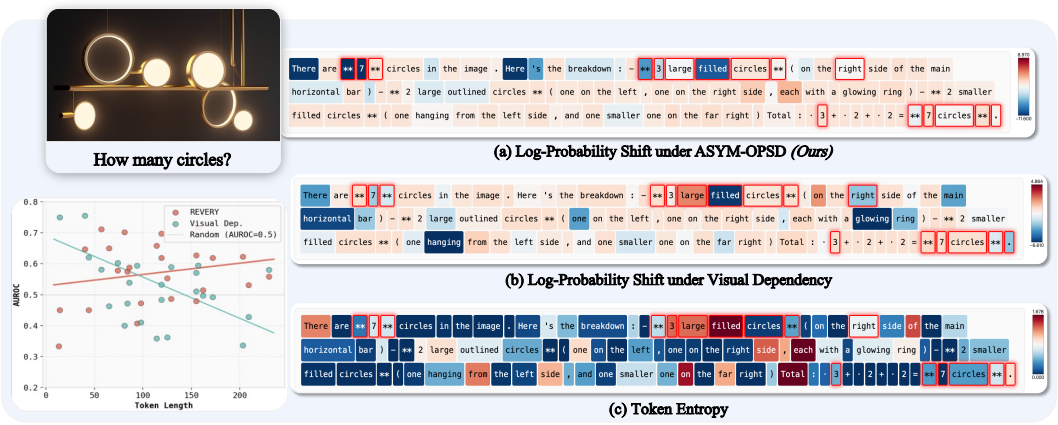

Figure 1: (Left) Per-sample AUROC of token-level hallucination detection v.s. response length. Each point represents one of 25 manually annotated samples from HaloQuest Challenge, evaluated on Qwen3-VL-4B. (Right) Comparison of token-level indicators on a hallucinated response. Red boxes denote human-annotated hallucinations. ASYM-OPSD (a) yields more localized and precise signals on hallucinated tokens than visual dependency (b) and entropy (c).

We recover this signal from a simple change of the input. When a vision-language model scores the same response twice, once with and once without a textual description of the image, its per-token probabilities shift: tokens consistent with the description gain probability, and contradicted tokens lose it. We take the per-token log-probability difference between these two conditions as a *groundedness score* over all tokens of the response. Figure 1 validates the signal on Qwen3-VL-4B over 25 manually annotated HaloQuest responses: per-sample detection AUROC (area under the Receiver Operating Characteristic curve) stays high across response lengths, and the gap concentrates on human-annotated hallucinated segments more precisely than visual-dependency and entropy cues. Because the score is computed from model log-probabilities alone, it requires no human annotation and is available at training time: computed between a teacher that reads the description and a student that does not, it supervises every rollout at the token level.

We propose ASYM-OPSD, an asymmetric on-policy self-distillation framework built on this signal, with the pipeline illustrated in Figure 2. **(1) Visual anchor construction**: a frontier model describes every training image at three granularities—brief, medium, and full—and a second frontier model cross-checks these *visual anchors* against the image. **(2) Anchor-guided warm-up**: training directly on the gap fails—the anchored answers differ too widely from the model's own, turning most token rewards negative while the anchor's descriptive phrasing leaks through the teacher's probabilities; repairing the hallucinated answers with the anchors and fine-tuning on them instead yields a frozen teacher that reads anchors natively. **(3) RL with token-level advantages**: the student samples rollouts from the standard input, the teacher re-scores each one under anchored inputs Lu & Lab (2025); Agarwal et al. (2024), and each per-token gap is normalized into a group-relative advantage Shao et al. (2024), crediting each token by its own groundedness. Finally, a curriculum fades the anchor from the paragraph level to the brief level, bridging the teacher's training condition and the student's anchor-free deployment.

Extensive experiments on two backbones, Qwen3-VL-4B and InternVL3.5-8B, demonstrate the effectiveness of ASYM-OPSD: across HallusionBench, HaloQuest, CHAIR, and MMHal-Bench, it achieves the best score on 13 of the 14 hallucination metrics in both model groups, mitigating hallucination in long-form responses by up to $3.2\times$; general capability is largely preserved, with slight gains on mathematical reasoning and scene-text QA (Section 4.1). A length-stratified analysis shows that ASYM-OPSD is the only method whose hallucination rate does not grow with response length (Section 4.3); a token-level localization study confirms that the anchored gap separates grounded from hallucinated tokens far more precisely than the model's internal signals (Section 4.4); and ablations quantify the contribution of teacher repair and the anchor-granularity curriculum (Section 4.5). Because anchors never appear at inference time, inference is a single standard forward pass, with none of the extra computation that decoding-time baselines impose on every generated token Leng et al. (2023); Zhong et al. (2026).

Our contributions are:

- **A multi-granularity visual anchor dataset**: 22,664 image–question pairs spanning seven public visual-linguistic datasets, each annotated with verified anchors at three granularities—brief, medium, and full descriptions of the image—generated by a frontier vision-language model and cross-checked by a second one, providing reliable textual surrogates of image content for training.
- **ASYM-OPSD**: an asymmetric on-policy self-distillation framework that derives dense, token-level groundedness rewards from the log-probability gap between anchored and standard conditions, made trainable through anchor-guided warm-up, group-relative token-level advantages, and a curriculum on anchor granularity.
- **Comprehensive experiments**: ASYM-OPSD outperforms training-based and inference-time baselines across four hallucination benchmarks on Qwen3-VL-4B and InternVL3.5-8B, while preserving general VQA performance at zero additional inference cost.

2 Related Work

2.1 On Policy distillation.

Reinforcement learning with verifiable rewards (RLVR) Guo et al. (2025a); Zhou et al. (2025); Huang et al. (2026) uses algorithms such as PPO and GRPO to improve reasoning in math and coding, but its reward is sparse: feedback applies only to the final response, not individual tokens Zhang et al. (2025); Guo et al. (2025b). On-policy distillation Lu & Lab (2025); Agarwal et al. (2024); Jin et al. (2026); Ko et al. (2026) addresses this by training on the student's own rollouts and replacing outcome rewards with dense, token-level supervision from a teacher. On-policy self-distillation (OPSD) Ye et al. (2026a); Zhao et al. (2026); Hübotter et al. (2026) takes this a step further: a single model serves as both teacher and student under different contexts, removing the need for an external teacher. Despite this simplification, OPSD remains limited in two respects. When the teacher's context is richer than the student's, it can suppress epistemic verbalization, pushing reasoning toward compressed, overconfident patterns that hurt out-of-distribution generalization Kim et al. (2026). More directly relevant to our work, OPSD has been studied predominantly in text-only reasoning tasks. The recent Vision-OPD **?** extends self-distillation to multimodal LLMs, transferring a model's privileged perception of cropped regions to its full-image policy, and improving fine-grained visual understanding without external teachers. ASYM-OPSD instead privileges a multi-granularity *textual* description of the image and interprets the per-token gap as a groundedness reward, targeting hallucination mitigation rather than fine-grained perception.

2.2 Hallucination Mitigation in MLLMs

Hallucinations in MLLMs are generally defined as generated content lacking visual grounding, where the model produces objects, attributes, or claims unsupported by the input image Wang et al. (2024); Ji et al. (2025); Sun et al. (2024); Guan et al. (2024a). Existing mitigation strategies broadly fall into two categories: training-free inference interventions and training-based preference alignment. Training-free approaches modify the decoding process directly. For instance, contrastive decoding Leng et al. (2023); Jung et al. (2025); Manevich & Tsarfaty (2024); Chen et al. (2025) compares token distributions between original and perturbed inputs, while attention reallocation Yang et al. (2026); Zhong et al. (2026); Wang et al. (2026); Yu et al. (2026); Xu et al. (2026) artificially amplifies attention weights toward salient visual regions. Although these methods avoid additional training, their heuristic nature can distort the model's native dynamics, suppress valid textual context, and degrade generation coherence. Conversely, training-based methods like RLHF or DPO optimize the model using preference pairs to favor factual outputs Pan et al. (2026a); Zheng et al. (2026); Ye et al. (2026b); Pan et al. (2026b). However, these approaches incur substantial annotation costs and typically rely on sequence-level supervision. This coarse-grained reward structure makes it hard to localize the tokens responsible for hallucinations, limiting precise correction.

3 Method

Given an image I and a question q, we aim to generate a visually grounded response $y = (y_1, \ldots, y_T)$. Conditioning a vision-language model on privileged visual information raises the

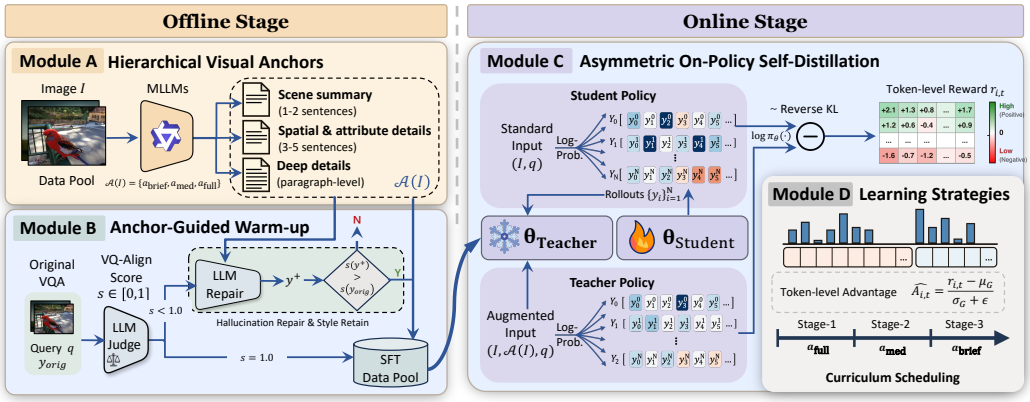

Figure 2: Overview of ASYM-OPSD. Left: offline data construction—hierarchical anchors and anchor-guided repair yield the SFT teacher. Right: online RL—the frozen teacher re-scores the student's rollouts under anchored inputs, and the per-token log-probability gap serves as the reward.

probability of tokens consistent with the image and lowers that of hallucinated ones, so the per-token difference between augmented and standard inputs identifies hallucinations densely at the token level. We exploit this signal via asymmetric on-policy self-distillation: a teacher (θ_t, SFT-initialized) scores each rollout conditioned on visual anchors, the student (θ_s, the base model) generates from the standard input, and their log-probability gap serves as the reward. Three stages follow—visual anchor construction (Section 3.1), anchor-guided warm-up (Section 3.2), and RL with curriculum scheduling (Section 3.3); Figure 2 gives an overview.

3.1 Visual Anchors

Training data is drawn from existing hallucination evaluation benchmarks, where each sample consists of an image and a question. For every image I, we generate hierarchical textual descriptions called visual anchors using the frontier Qwen3.5-Max model at three granularities: a_{brief}, a one-to-two-sentence scene summary; a_{med}, a three-to-five-sentence description of spatial relations and object attributes; and a_{full}, a paragraph-level detailed description. Anchors are cross-verified by two independent frontier models (Gemini-3.5-Flash and GPT-5.6-SOL); the full construction pipeline is in Appendix A.1. We denote the anchor set as $\mathcal{A}(I) = \{a_{\text{brief}}, a_{\text{med}}, a_{\text{full}}\}$, where the three levels also define the curriculum of Section 3.3. During training, the teacher receives the augmented input $(I, \mathcal{A}(I), q)$ while the policy operates on the standard input (I, q); anchors appear only during training, so inference is a single standard forward pass.

3.2 Anchor-Guided Warm-Up

Direct application of on-policy self-distillation from a base model suffers from two failure modes, both rooted in the large gap between the anchor-conditioned teacher's distribution and the raw policy's own answers. First, because the two distributions differ widely in both style and length, the teacher assigns negative rewards to a large fraction of tokens, and the policy drifts along this gap—collapsing into extremely short responses to minimize the cumulative penalty. Second, the teacher's input contains detailed textual anchors whose descriptive style (e.g., "The image shows…", "There is…") propagates into the teacher's log-probability distribution. As the policy mimics this distribution during optimization, it gradually absorbs the descriptive style and drifts away from the concise, conversational tone expected of a visual question-answering system.

To address both issues, we construct SFT data via anchor-guided repair before initiating RL. The data construction proceeds as follows. First, the base policy model generates an initial response for each training sample. Qwen3.5-Max then scores these responses for hallucination with a judge score $s \in [0, 1]$. For responses containing hallucinations, the same model repairs them using the visual anchors $\mathcal{A}(I)$: it rewrites the response to correct visual inconsistencies while preserving the original conversational style. Each repaired response y^+ is re-evaluated by the judge, and only

samples showing strict improvement ($s(y^+) > s(y_{\text{orig}})$) are retained. We also retain 20% of responses that received perfect scores from the initial scoring round. The final SFT dataset pairs inputs $(I, \mathcal{A}(I), q)$ with outputs consisting of 80% repaired responses and 20% original perfect responses. The policy model is fine-tuned on this dataset with standard next-token prediction, producing checkpoint θ_{SFT}. This checkpoint is then frozen and used as the teacher θ_t for the RL stage: having learned to interpret visual anchors, it can accurately evaluate token-level grounding when conditioned on augmented input.

3.3 ASYM-OPSD

In the RL stage, the student model θ_s is trained while the teacher model $\theta_t = \theta_{\text{SFT}}$ (from Section 3.2) is frozen. The two models share the same architecture but differ in both initialization and input: the teacher receives the augmented input $(I, \mathcal{A}(I), q)$ and has been fine-tuned to interpret visual anchors, while the student generates from the standard input (I, q). For each prompt, the student generates n rollouts $\{y_i\}_{i=1}^{n}$. The teacher then evaluates each rollout via forced decoding, computing token-level log-probabilities over the same tokens. The per-token reward is defined as the log-probability difference between teacher and student:

$$r_{i,t} = \log \pi_{\theta_t}(y_{i,t} \mid I, \mathcal{A}(I), q, y_{i,<t}) - \log \pi_{\theta_s}(y_{i,t} \mid I, q, y_{i,<t}), \tag{1}$$

where positive $r_{i,t}$ indicates the teacher is more confident than the student in the token, suggesting it is visually grounded, and negative $r_{i,t}$ suggests potential hallucination. The reward is always computed on the student's own rollouts rather than on sequences generated independently, so the signal lands on the responses the policy actually produces Lu & Lab (2025). Raw rewards are clamped to $[-8, 8]$ for numerical stability.

We convert the per-token reward into a policy gradient signal using group-relative normalization inspired by GRPO Shao et al. (2024). For each prompt, we compute the mean reward $\bar{r}_i$ of each rollout and derive group statistics μ_G, σ_G across the n rollouts. The per-token advantage is:

$$\hat{A}_{i,t} = \frac{r_{i,t} - \mu_G}{\sigma_G + \epsilon}. \tag{2}$$

Unlike standard GRPO, which assigns a single scalar advantage to all tokens in a rollout, we preserve token-level granularity so that each token receives a distinct advantage based on its own groundedness. The student is then optimized with the standard clipped GRPO objective: a low-variance KL estimate toward a frozen reference copy of the student is folded into the per-token reward with coefficient β_{KL}, and the surrogate

$$\mathcal{J}(\theta_s) = \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} \min\big(\rho_{i,t}\,\hat{A}_{i,t},\ \text{clip}(\rho_{i,t}, 1-\epsilon, 1+\epsilon)\,\hat{A}_{i,t}\big)\right], \tag{3}$$

is maximized, where $\rho_{i,t} = \pi_{\theta_s}(y_{i,t} \mid I, q, y_{i,<t}) \,/\, \pi_{\theta_{\text{old}}}(y_{i,t} \mid I, q, y_{i,<t})$ is the importance ratio under PPO clipping with $\epsilon = 0.2$.

We schedule the anchor granularity across training to progressively narrow the information gap between teacher and student. In the first third of training, the teacher uses a_{full} for the strongest signal. In the middle third, it switches to a_{med}. In the final third, it uses a_{brief}, closely approximating the inference-time condition where no anchors are provided. This curriculum bridges the train–test distribution mismatch while maintaining strong supervision early in training.

4 EXPERIMENTS

We evaluate ASYM-OPSD on two backbones from different families, Qwen3-VL-4B and InternVL3.5-8B, with four hallucination benchmarks and six general VQA benchmarks (Section 4.1). The experiments ask four questions: whether anchored self-distillation reduces hallucination below training-based and inference-time alternatives (Table 2); whether that reduction costs general capability (Figure 3); whether the per-token reward separates grounded from hallucinated tokens better than the model's internal signals (Figure 5); and where the gains come from and where the remaining hallucinations live (Figure 4).

4.1 TRAINING SETTINGS

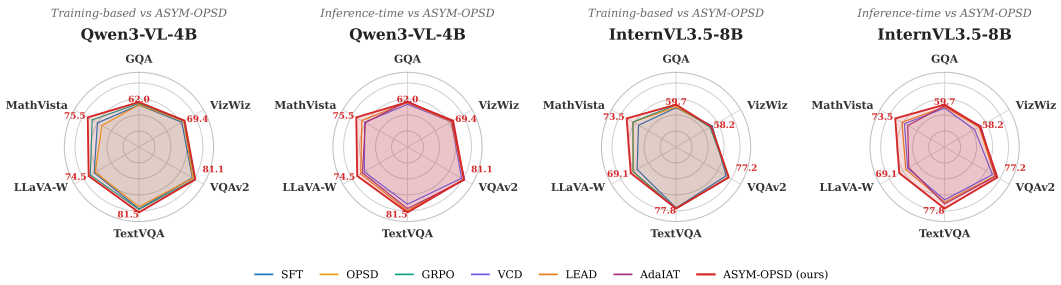

Figure 3: General VQA comparison across six benchmarks. Each radar overlays training-based baselines (SFT, OPSD, GRPO) and inference-time baselines (VCD, LEAD, AdaIAT) against ASYM-OPSD, on Qwen3-VL-4B (left half) and InternVL3.5-8B (right half).

Data. We construct visual anchors on a stratified root drawn from seven vision-language datasets spanning visually challenging QA, chart and document understanding, spatial reasoning, geometry, instruction following, and map reading. After RL finalization—anchor verification drops rejected samples and the purely descriptive DocCI source is capped at 200 so that caption-style outputs cannot dominate rollout batches—the RL corpus is fixed at 19,398 prompts, with per-source counts in Table 1. Anchor-guided warm-up data is generated by repairing hallucinated base-model responses, yielding 8,749 fine-tuning samples for each of the Qwen3-VL and InternVL3.5 teachers; Appendix A.1 details the full pipeline.

Table 1: Composition of the final 19,398 RL training prompts across seven vision-language sources, after anchor verification and the RL-stage DocCI cap; Appendix A.1 details how the stratified root is finalized.

Source	#	%
HaloQuest	5,265	27.1%
ChartQA	4,823	24.9%
VSR	4,270	22.0%
Geo3K	2,384	12.3%
LRV	1,942	10.0%
MapQA	514	2.6%
DocCI	200	1.0%
Total	**19,398**	**100%**

Hyperparameters. All methods are implemented on the verl RL framework and trained on a single node with $8\times$ NVIDIA A100-SXM4-80GB GPUs; rollout generation and teacher scoring run on vLLM, and a full 3-epoch ASYM-OPSD run takes 1–4 GPU-days (98.3 hours for Qwen3-VL-4B, 31.4 hours for InternVL3.5-8B). Key hyper-parameters of ASYM-OPSD: constant learning rate 2×10^{-6} with AdamW, $n = 32$ (Qwen3-VL-4B) or $n = 20$ (InternVL3.5-8B) rollouts per prompt at temperature 1.0, token-level rewards clamped to $[-8, 8]$, and a KL penalty toward the frozen reference (coefficient 0.02; 0.005 for InternVL3.5-8B) folded into the reward. Full hyper-parameters of every run are listed in Appendix A.2.

Baselines. We compare against six baselines. Training-based: (1) **SFT**, the anchor-guided warm-up checkpoint without RL; (2) **GRPO**, which optimizes with an outcome-level GraphHallu reward and standard scalar advantage; and (3) **OPSD**, on-policy self-distillation from the warm-up SFT teacher. Inference-time decoding: (4) **VCD**, which contrasts output distributions from the original and perturbed images Leng et al. (2023); (5) **LEAD**, which anneals between discrete-token and continuous latent decoding guided by latent entropy Xu et al. (2026); and (6) **AdaIAT**, which adaptively amplifies attention from generated text toward visual tokens Zhong et al. (2026). All six are applied to the same base models; each decoding baseline runs its own intervention following its official implementation with the recommended sampler and generation budget, and every configuration is listed in Appendix A.2.

Evaluation. We assess hallucination on four benchmarks—**CHAIR**, **HallusionBench**, **HaloQuest**, and **MMHal-Bench**—covering caption-level object hallucination, insufficient-context queries, and visual grounding and reasoning; rule-based matching scores CHAIR, and the VLM-based API evaluator Qwen3.5-Max judges the open-ended responses. General capabilities are measured on six general-VQA benchmarks—**GQA**, **VizWiz**, **VQAv2**, **TextVQA**, **LLaVA-Wilder**, and **MathVista**—from scene-graph QA to open-ended dialogue and mathematical reasoning. Full metric definitions and judging details are in Appendix A.4.

Table 2: Hallucination results on four benchmarks. aAcc/fAcc/qAcc: image-, figure-, and question-level accuracy on HallusionBench; acc: HaloQuest accuracy; CHAIR_i/CHAIR_s: object hallucination rate per object mention / per caption; the MMHal column reports the hallucination rate (fraction of responses rated hallucinated). **Bold**/underline: best/second-best per model group. Subscripts show raw deltas vs. the raw model (blue: ASYM-OPSD); negative deltas on ↓ metrics are improvements.

		HallusionBench			HaloQuest	CHAIR (COCO)		MMHal
Model	**Method**	aAcc↑	fAcc↑	qAcc↑	acc↑	CHAIR_i ↓	CHAIR_s ↓	hall. rate↓
Qwen3-VL-4B	Raw	64.66 $_{0}$	45.38 $_{0}$	35.60 $_{0}$	82.00 $_{0}$	9.05 $_{0}$	42.40 $_{0}$	15.33 $_{0}$
	SFT	60.67 $_{-4.0}$	38.73 $_{-6.7}$	26.54 $_{-9.1}$	83.83 $_{+1.8}$	7.46 $_{-1.6}$	44.00 $_{+1.6}$	16.20 $_{+0.9}$
	OPSD	65.46 $_{+0.8}$	47.40 $_{+2.0}$	35.38 $_{-0.2}$	85.67 $_{+3.7}$	8.64 $_{-0.4}$	43.40 $_{+1.0}$	14.58 $_{-0.8}$
	GRPO	66.43 $_{+1.8}$	46.82 $_{+1.4}$	37.14 $_{+1.5}$	85.83 $_{+3.8}$	8.36 $_{-0.7}$	40.40 $_{-2.0}$	8.42 $_{-6.9}$
	AdaIAT	64.39 $_{-0.3}$	44.51 $_{-0.9}$	34.73 $_{-0.9}$	81.33 $_{-0.7}$	8.30 $_{-0.8}$	41.60 $_{-0.8}$	13.54 $_{-1.8}$
	VCD	64.39 $_{-0.3}$	43.64 $_{-1.7}$	34.51 $_{-1.1}$	83.17 $_{+1.2}$	9.46 $_{+0.4}$	44.80 $_{+2.4}$	11.11 $_{-4.2}$
	LEAD	66.25 $_{+1.6}$	48.55 $_{+3.2}$	33.85 $_{-1.8}$	80.63 $_{-1.4}$	8.92 $_{-0.1}$	44.60 $_{+2.2}$	11.46 $_{-3.9}$
	ASYM-OPSD	**69.80** $_{+5.1}$	**52.31** $_{+6.9}$	**39.56** $_{+4.0}$	**86.50** $_{+4.5}$	**6.90** $_{-2.2}$	**38.80** $_{-3.6}$	**5.21** $_{-10.1}$
InternVL3.5-8B	Raw	62.53 $_{0}$	39.02 $_{0}$	27.91 $_{0}$	64.76 $_{0}$	8.69 $_{0}$	30.00 $_{0}$	23.96 $_{0}$
	SFT	62.98 $_{+0.5}$	40.75 $_{+1.7}$	31.21 $_{+3.3}$	66.94 $_{+2.2}$	9.12 $_{+0.4}$	31.50 $_{+1.5}$	21.05 $_{-2.9}$
	OPSD	62.09 $_{-0.4}$	39.85 $_{+0.8}$	27.47 $_{-0.4}$	68.33 $_{+3.6}$	8.38 $_{-0.3}$	28.20 $_{-1.8}$	19.79 $_{-4.2}$
	GRPO	63.33 $_{+0.8}$	40.46 $_{+1.4}$	29.67 $_{+1.8}$	67.33 $_{+2.6}$	8.32 $_{-0.4}$	**27.40** $_{-2.6}$	19.79 $_{-4.2}$
	AdaIAT	60.42 $_{-2.1}$	36.92 $_{-2.1}$	27.91 $_{0}$	64.33 $_{-0.4}$	8.66 $_{0}$	27.80 $_{-2.2}$	20.83 $_{-3.1}$
	VCD	60.14 $_{-2.4}$	32.95 $_{-6.1}$	25.49 $_{-2.4}$	57.00 $_{-7.8}$	10.57 $_{+1.9}$	34.20 $_{+4.2}$	22.92 $_{-1.0}$
	LEAD	61.03 $_{-1.5}$	33.82 $_{-5.2}$	26.15 $_{-1.8}$	57.50 $_{-7.3}$	9.25 $_{+0.6}$	28.40 $_{-1.6}$	17.71 $_{-6.3}$
	ASYM-OPSD	**65.81** $_{+3.3}$	**43.93** $_{+4.9}$	**32.97** $_{+5.1}$	**69.83** $_{+5.1}$	**8.30** $_{-0.4}$	27.60 $_{-2.4}$	**15.62** $_{-8.3}$

4.2 Main Results

Table 2 reports the four hallucination benchmarks on both backbones, where ASYM-OPSD is the best method on 13 of 14 metrics. Its largest gains land on open-ended generation: the MMHal hallucination rate drops from 15.33 to 5.21 on Qwen3-VL-4B and from 23.96 to 15.62 on InternVL3.5-8B, caption hallucination shrinks accordingly (CHAIR_s $42.40 \rightarrow 38.80$ on Qwen3-VL-4B), and QA-style metrics improve by smaller but consistent margins (HallusionBench aAcc $+5.1$/$+3.3$). Two observations follow. First, sequence-level signals capture only part of the benefit: the strongest training-based baseline, GRPO, also reduces hallucination but saturates well above ASYM-OPSD (MMHal 8.42 vs. 5.21 on Qwen3-VL-4B), although ASYM-OPSD is tuned more conservatively than GRPO because its token-level dense rewards produce higher-variance gradients. Second, inference-time interventions are unstable across backbones: VCD and LEAD improve several Qwen3-VL-4B metrics yet fall below the raw model on most InternVL3.5-8B ones (HaloQuest -7.8/-7.3), as their interventions act on every decoded token, hallucinated or not.

Figure 3 reports the six general VQA benchmarks. General capability is largely preserved: ASYM-OPSD is the strongest method on 8 of the 12 model–benchmark combinations, stays within 0.9 points of the raw model elsewhere, and even improves slightly on mathematical reasoning and scene-text QA—MathVista ($74.8 \rightarrow 75.5$ on Qwen3-VL-4B, $70.3 \rightarrow 73.5$ on InternVL3.5-8B) and TextVQA ($79.7 \rightarrow 81.5$ on Qwen3-VL-4B)—axes on which training-based baselines lose 4.1–14.7 points and decoding interventions up to 10.3.

4.3 How Hallucination Scales with Response Length

Figure 4 stratifies the 600 HaloQuest evaluation items by response length, plotting each method's hallucination rate against response length, together with a linear slope approximating the marginal change in hallucination probability per additional 1,000 tokens. The raw model (gray) dips at intermediate lengths, but this dip is an artifact of item composition rather than a length effect: that band pools more wrong-premise questions, whose correct answers are short premise-rejection statements (5 of the 17 shortest responses, versus 6 of 46 at other lengths), so many short, correct answers pull the band's hallucination rate down.

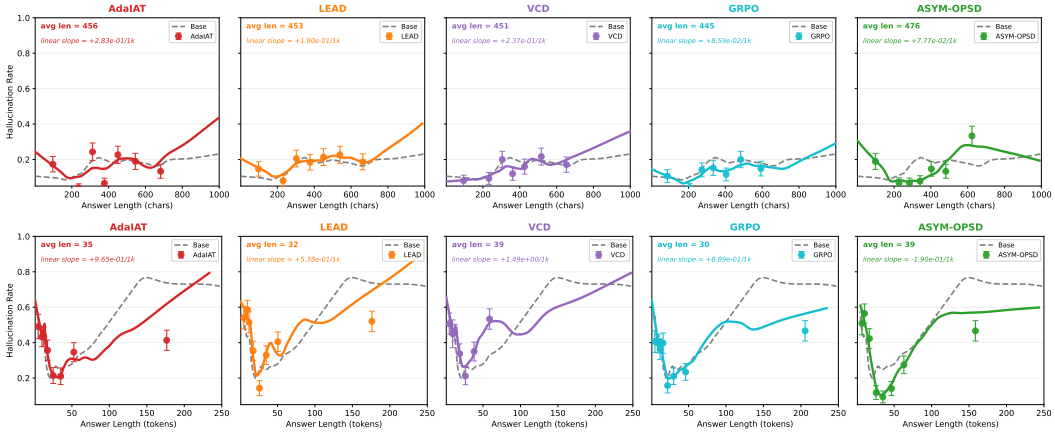

Figure 4: Response length versus hallucination rate, stratified over the 600 HaloQuest evaluation items (top: Qwen3-VL-4B; bottom: InternVL3.5-8B). Each curve plots one method's hallucination rate as a function of response length; the gray curve is the raw model. The linear slope ($\times 1000$) approximates the change in hallucination probability per additional 1,000 tokens.

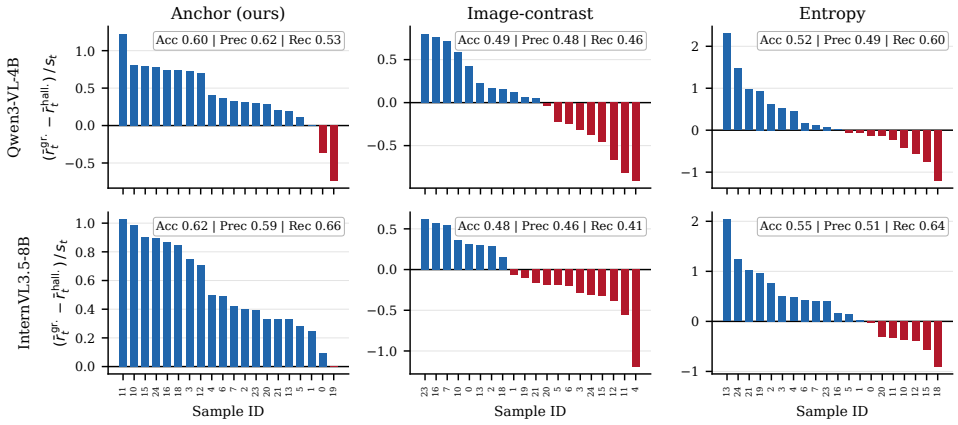

Figure 5: Token-level hallucination localization by three signals (anchor gap of Equation 1, image contrast, negative entropy) over the same 20 human-annotated responses, on Qwen3-VL-4B (top) and InternVL3.5-8B (bottom). Bars: per-response normalized difference between the mean reward of grounded tokens and that of hallucinated tokens; corner: token-level accuracy, precision, and recall of distinguishing grounded from hallucinated tokens.

Under this length-stratified view, ASYM-OPSD is the only method whose hallucination rate does not grow with response length; every other method fits a positive slope. The effect concentrates in long answers, precisely where models pile up visual details that the image may not support: in the longest band on InternVL3.5-8B, ASYM-OPSD compresses the base model's hallucination rate from 0.303 to 0.095, a 3.2$\times$ reduction. Other methods lift average scores while leaving long-answer hallucinations in place; the gains of ASYM-OPSD land on exactly those long answers.

4.4 Token-Level Localization Accuracy

We test how precisely three token-level reward signals localize hallucinated tokens: (i) our anchored log-probability gap of Equation 1, computed between the anchor-conditioned and the standard input; (ii) the *image-contrast* gap $r_t = \log P(y_t \mid I, q) - \log P(y_t \mid q)$, which underlies contrastive decoding Leng et al. (2023); and (iii) per-token *negative entropy*, on which latent-entropy decoding builds Xu et al. (2026). All three signals are computed on the same raw backbone via forced scoring,

Table 3: Ablation on curriculum learning and teacher repair. **OPSD** uses the SFT teacher without repair; **ASYM-OPSD$_{\text{w/o cur.}}$** removes the curriculum schedule while keeping repair; **ASYM-OPSD** adds curriculum-guided teacher repair. **Bold**/underline: best/second-best among the three variants. Subscripts show raw deltas vs. the raw model (blue: ASYM-OPSD); negative deltas on ↓ metrics are improvements.

Model	Variant	HB aAcc↑	HB fAcc↑	HB qAcc↑	HQ acc↑	CHAIR$_i$↓	CHAIR$_s$↓	MMHal↓
Qwen3-VL-4B	Raw	64.66 $_{0}$	45.38 $_{0}$	35.60 $_{0}$	82.00 $_{0}$	9.05 $_{0}$	42.40 $_{0}$	15.33 $_{0}$
	OPSD (w/o repair)	65.46 $_{+0.8}$	47.40 $_{+2.0}$	35.38 $_{-0.2}$	85.67 $_{+3.7}$	8.64 $_{-0.4}$	43.40 $_{+1.0}$	14.58 $_{-0.8}$
	ASYM-OPSD$_{\text{w/o cur.}}$	68.20 $_{+3.5}$	48.55 $_{+3.2}$	39.34 $_{+3.7}$	84.50 $_{+2.5}$	9.31 $_{+0.3}$	45.20 $_{+2.8}$	13.54 $_{-1.8}$
	ASYM-OPSD	**69.80** $_{+5.1}$	**52.31** $_{+6.9}$	**39.56** $_{+4.0}$	**86.50** $_{+4.5}$	**6.90** $_{-2.2}$	**38.80** $_{-3.6}$	**5.21** $_{-10.1}$
InternVL3.5-8B	Raw	62.53 $_{0}$	39.02 $_{0}$	27.91 $_{0}$	64.76 $_{0}$	8.69 $_{0}$	30.00 $_{0}$	23.96 $_{0}$
	OPSD (w/o repair)	62.09 $_{-0.4}$	39.85 $_{+0.8}$	27.47 $_{-0.4}$	68.33 $_{+3.6}$	8.38 $_{-0.3}$	28.20 $_{-1.8}$	19.79 $_{-4.2}$
	ASYM-OPSD$_{\text{w/o cur.}}$	64.22 $_{+1.7}$	41.04 $_{+2.0}$	**32.97** $_{+5.1}$	68.33 $_{+3.6}$	8.86 $_{+0.2}$	30.20 $_{+0.2}$	18.75 $_{-5.2}$
	ASYM-OPSD	**65.81** $_{+3.3}$	**43.93** $_{+4.9}$	**32.97** $_{+5.1}$	**69.83** $_{+5.1}$	**8.30** $_{-0.4}$	**27.60** $_{-2.4}$	**15.62** $_{-8.3}$

the quantity the reward service scores during RL (Section 3.3). The evaluation set consists of 20 raw-model responses hand-annotated at the entity level into *grounded* and *hallucinated* tokens: 10 from HaloQuest and 10 from DocCI.

For each response, we compare the mean reward of its grounded tokens with that of its hallucinated tokens, normalized per response (bars in Figure 5), and report token-level accuracy, precision, and recall for distinguishing grounded from hallucinated tokens (corner of each panel; full statistics in Table A7, Appendix A.3).

The anchor signal is the only one that separates the two classes on both backbones: precision $0.62/0.59$, accuracy $0.60/0.62$, and recall $0.53/0.66$ (Qwen3-VL-4B/InternVL3.5-8B). The other two signals stay near chance: image contrast never exceeds precision 0.48, since hallucinated tokens are language-plausible collocations rated highly with or without the image, so the gain from seeing the image carries little information about whether a token contradicts it; entropy attains recall $0.60/0.64$ only by flagging 57–58% of all tokens, with precision falling back to $0.49/0.51$.

4.5 Ablation Studies

Table 3 isolates the two training-time choices. Distilling without the repaired teacher (OPSD) recovers only part of what ASYM-OPSD achieves—MMHal stays at 14.58 on Qwen3-VL-4B and 19.79 on InternVL3.5-8B, against 5.21 and 15.62 with repair—because the residual hallucinations in the teacher's own predictions bleed into the per-token signal it scores. Removing only the anchor-granularity curriculum (ASYM-OPSD$_{\text{w/o cur.}}$) preserves most QA-side gains but gives back the open-ended ones: on Qwen3-VL-4B, MMHal rises to 13.54, CHAIR$_i$ to 9.31, and CHAIR$_s$ to 45.20, versus 5.21, 6.90, and 38.80 for the full method. Without fading the anchor from paragraph-level to brief, the gap between the teacher's anchored condition and the student's anchor-free deployment stays open, and long-form generation pays for it. Both trends replicate across the two backbones.

5 Conclusion

This paper introduced ASYM-OPSD: a frozen teacher conditioned on hierarchical visual anchors rescores the student's own rollouts, and the per-token probability gap between the two supplies a dense groundedness signal; anchor-guided warm-up, token-level advantages, and an anchor-granularity curriculum make the signal trainable while keeping inference identical to the raw model. On Qwen3-VL-4B and InternVL3.5-8B, ASYM-OPSD mitigates hallucination more consistently than training-based and inference-time baselines while preserving general VQA capability. Future work includes replacing the frontier models that currently build anchors and repairs with smaller verifiers, and extending the signal beyond question answering.

AI USE STATEMENT

In this work, we used generative AI tools to assist with language polishing and manuscript editing. Generative models were also employed in the data and evaluation pipeline, as described in the main text and appendix: hierarchical visual anchors were generated by Qwen3.5-Max and cross-checked or repaired by Gemini-3.5-Flash and GPT-5.6-SOL; hallucinated responses were rewritten to build the anchor-guided warm-up data; and open-ended responses were evaluated by the VLM-based API evaluator (Qwen3.5-Max), while CHAIR is scored purely by rule-based matching.

We did not use generative AI tools to fabricate experimental results or human annotations, or to substitute for the human labeling reported in this paper: the entity-level grounded and hallucinated token labels of the token-localization study were produced manually by the authors.

We reviewed the AI-assisted manuscript edits and the AI-generated data artifacts against our methods and reported findings. We take responsibility for the final content of this work, including text, claims, data, and artifacts produced with the assistance of generative AI.

REFERENCES

Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos, Matthieu Geist, and Olivier Bachem. On-policy distillation of language models: Learning from self-generated mistakes, 2024. URL https://arxiv.org/abs/2306.13649.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, Mei Li, Kaixin Li, Zicheng Lin, Junyang Lin, Xuejing Liu, Jiawei Liu, Chenglong Liu, Yang Liu, Dayiheng Liu, Shixuan Liu, Dunjie Lu, Ruilin Luo, Chenxu Lv, Rui Men, Lingchen Meng, Xuancheng Ren, Xingzhang Ren, Sibo Song, Yuchong Sun, Jun Tang, Jianhong Tu, Jianqiang Wan, Peng Wang, Pengfei Wang, Qiuyue Wang, Yuxuan Wang, Tianbao Xie, Yiheng Xu, Haiyang Xu, Jin Xu, Zhibo Yang, Mingkun Yang, Jianxin Yang, An Yang, Bowen Yu, Fei Zhang, Hang Zhang, Xi Zhang, Bo Zheng, Humen Zhong, Jingren Zhou, Fan Zhou, Jing Zhou, Yuanzhi Zhu, and Ke Zhu. Qwen3-vl technical report, 2025. URL https://arxiv.org/abs/2511.21631.

Wei Chen, Xin Yan, Bin Wen, Fan Yang, Tingting Gao, Di ZHANG, and Long Chen. Decoupling contrastive decoding: Robust hallucination mitigation in multimodal large language models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=PAYYkls0WD.

Hao Fang, Changle Zhou, Jiawei Kong, Kuofeng Gao, Bin Chen, and Shu-Tao Xia. Grounding language with vision: A conditional mutual information calibrated decoding strategy for reducing hallucinations in lvlms, 2025. URL https://arxiv.org/abs/2505.19678.

Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. Hallusionbench: An Advanced Diagnostic Suite for Entangled Language Hallucination and Visual Illusion in Large Vision-Language Models . In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14375–14385, Los Alamitos, CA, USA, June 2024a. IEEE Computer Society. doi: 10.1109/CVPR52733.2024.01363. URL https://doi.ieeecomputersociety.org/10.1109/CVPR52733.2024.01363.

Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. Hallusionbench: An advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14375–14385, June 2024b.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Hanwei Xu, Honghui Ding, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jingchang Chen, Jingyang Yuan, Jinhao Tu, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaichao You, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingxu Zhou, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan,

Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638, Sep 2025a. ISSN 1476-4687. doi: 10.1038/s41586-025-09422-z. URL https://doi.org/10.1038/s41586-025-09422-z.

Yiran Guo, Lijie Xu, Jie Liu, Dan Ye, and Shuang Qiu. Segment policy optimization: Effective segment-level credit assignment in rl for large language models, 2025b. URL https://arxiv.org/abs/2505.23564.

Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Xu Tang, Yao Hu, and Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models, 2026. URL https://arxiv.org/abs/2503.06749.

Jonas Hübotter, Frederike Lübeck, Lejs Behric, Anton Baumann, Marco Bagatella, Daniel Marta, Ido Hakimi, Idan Shenfeld, Thomas Kleine Buening, Carlos Guestrin, and Andreas Krause. Reinforcement learning via self-distillation, 2026. URL https://arxiv.org/abs/2601.20802.

Jiaming Ji, Sitong Fang, Wenjing Cao, Jiahao Li, Xuyao Wang, Juntao Dai, Chi-Min Chan, Sirui Han, Yike Guo, and Yaodong Yang. The mirage of multimodality: Where truth is tested and honesty unravels, 2025. URL https://arxiv.org/abs/2505.20214.

Woogyeol Jin, Taywon Min, Yongjin Yang, Swanand Ravindra Kadhe, Yi Zhou, Dennis Wei, Nathalie Baracaldo, and Kimin Lee. Entropy-aware on-policy distillation of language models, 2026. URL https://arxiv.org/abs/2603.07079.

Chaeyoung Jung, Youngjoon Jang, and Joon Son Chung. AVCD: Mitigating hallucinations in audio-visual large language models through contrastive decoding. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=3zx087XPtz.

Jeonghye Kim, Xufang Luo, Minbeom Kim, Sangmook Lee, Dohyung Kim, Jiwon Jeon, Dongsheng Li, and Yuqing Yang. Why does self-distillation (sometimes) degrade the reasoning capability of llms?, 2026. URL https://arxiv.org/abs/2603.24472.

Jongwoo Ko, Sara Abdali, Young Jin Kim, Tianyi Chen, and Pashmina Cameron. Scaling reasoning efficiently via relaxed on-policy distillation, 2026. URL https://arxiv.org/abs/2603.11137.

Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. Mitigating object hallucinations in large vision-language models through visual contrastive decoding, 2023. URL https://arxiv.org/abs/2311.16922.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 34892–34916. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/6dcf277ea32ce3288914faf369fe6de0-Paper-Conference.pdf.

Kevin Lu and Thinking Machines Lab. On-policy distillation. *Thinking Machines Lab: Connectionism*, 2025. doi: 10.64434/tml.20251026. https://thinkingmachines.ai/blog/on-policy-distillation.

Avshalom Manevich and Reut Tsarfaty. Mitigating hallucinations in large vision-language models (LVLMs) via language-contrastive decoding (LCD). In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 6008–6022, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.359. URL https://aclanthology.org/2024.findings-acl.359/.

Miao Pan, Wangjie Gan, Jintao Chen, Wenqi Zhang, Bing Sun, Jianwei Yin, and Xuhong Zhang. Ground what you see: Hallucination-resistant mllms via caption feedback, diversity-aware sampling, and conflict regularization, 2026a. URL https://arxiv.org/abs/2601.06224.

Zhiyu Pan, Yizheng Wu, Jiashen Hua, Junyi Feng, Shaotian Yan, Bing Deng, Zhiguo Cao, and Jieping Ye. Through the lens of contrast: Self-improving visual reasoning in VLMs. In *The Fourteenth International Conference on Learning Representations*, 2026b. URL https://openreview.net/forum?id=ZymCPON45y.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. URL https://arxiv.org/abs/2402.03300.

Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liangyan Gui, Yu-Xiong Wang, Yiming Yang, Kurt Keutzer, and Trevor Darrell. Aligning large multimodal models with factually augmented RLHF. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 13088–13110, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.775. URL https://aclanthology.org/2024.findings-acl.775/.

Tianbo Wang, Yuqing Ma, Kewei Liao, Zhange Zhang, Simin Li, Jinyang Guo, and Xianglong Liu. AFTER: Mitigating the object hallucination of LVLM via adaptive factual-guided activation editing. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=ggycXmhrrG.

Zhecan Wang, Garrett Bingham, Adams Wei Yu, Quoc V. Le, Thang Luong, and Golnaz Ghiasi. Haloquest: A visual hallucination dataset for advancing multimodal reasoning. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LXXVII*, pp. 288–304, Berlin, Heidelberg, 2024. Springer-Verlag. ISBN 978-3-031-72979-9. doi: 10.1007/978-3-031-72980-5_17. URL https://doi.org/10.1007/978-3-031-72980-5_17.

Zhongxing Xu, Zhonghua Wang, Zhe Qian, Dachuan Shi, Feilong Tang, Ming Hu, Shiyan Su, Xiaocheng Zou, Wei Feng, Dwarikanath Mahapatra, Yifan Peng, Mingquan Lin, and Zongyuan Ge. Thinking in uncertainty: Mitigating hallucinations in mlrms with latent entropy-aware decoding, 2026. URL https://arxiv.org/abs/2603.13366.

Chengxu Yang, Jingling Yuan, Chuang Hu, and Jiawei Jiang. Visual attention drifts,but anchors hold:mitigating hallucination in multimodal large language models via cross-layer visual anchors, 2026. URL https://arxiv.org/abs/2603.25088.

Tianzhu Ye, Li Dong, Xun Wu, Shaohan Huang, and Furu Wei. On-policy context distillation for language models, 2026a. URL https://arxiv.org/abs/2602.12275.

Zekai Ye, Qiming Li, Xiaocheng Feng, Ruihan Chen, Ziming Li, Haoyu Ren, Kun Chen, Dandan Tu, and Bing Qin. Not all tokens see equally: Perception-grounded policy optimization for large vision-language models, 2026b. URL https://arxiv.org/abs/2604.01840.

Hao Yin, Guangzong Si, and Zilei Wang. The mirage of performance gains: Why contrastive decoding fails to mitigate object hallucinations in MLLMs? In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=wGeQr5cARm.

Xinlei Yu, Chengming Xu, Guibin Zhang, Yongbo He, Zhangquan Chen, Zhucun Xue, Jiangning Zhang, Yue Liao, Xiaobin Hu, Yu-Gang Jiang, and Shuicheng YAN. Visual multi-agent system: Mitigating hallucination snowballing via visual flow. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=Ii4HBlERix.

Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu, Jingren Zhou, and Junyang Lin. The lessons of developing process reward models in mathematical reasoning, 2025. URL https://arxiv.org/abs/2501.07301.

Siyan Zhao, Zhihui Xie, Mengchen Liu, Jing Huang, Guan Pang, Feiyu Chen, and Aditya Grover. Self-distilled reasoner: On-policy self-distillation for large language models, 2026. URL https://arxiv.org/abs/2601.18734.

Zhixiao Zheng, Zheren Fu, Zhiyuan Yao, Dongming Zhang, and Zhendong Mao. Cat-PO: Cross-modal adaptive token-rewards for preference optimization in truthful multimodal LLMs. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=iIbe6qDN0A.

Li'an Zhong, Ziqiang He, Jibin Zheng, Jin Li, Z. Jane Wang, and Xiangui Kang. Adaiat: Adaptively increasing attention to generated text to alleviate hallucinations in lvlm, 2026. URL https://arxiv.org/abs/2603.04908.

Hengguang Zhou, Xirui Li, Ruochen Wang, Minhao Cheng, Tianyi Zhou, and Cho-Jui Hsieh. R1-zero's "aha moment" in visual reasoning on a 2b non-sft model, 2025. URL https://arxiv.org/abs/2503.05132.

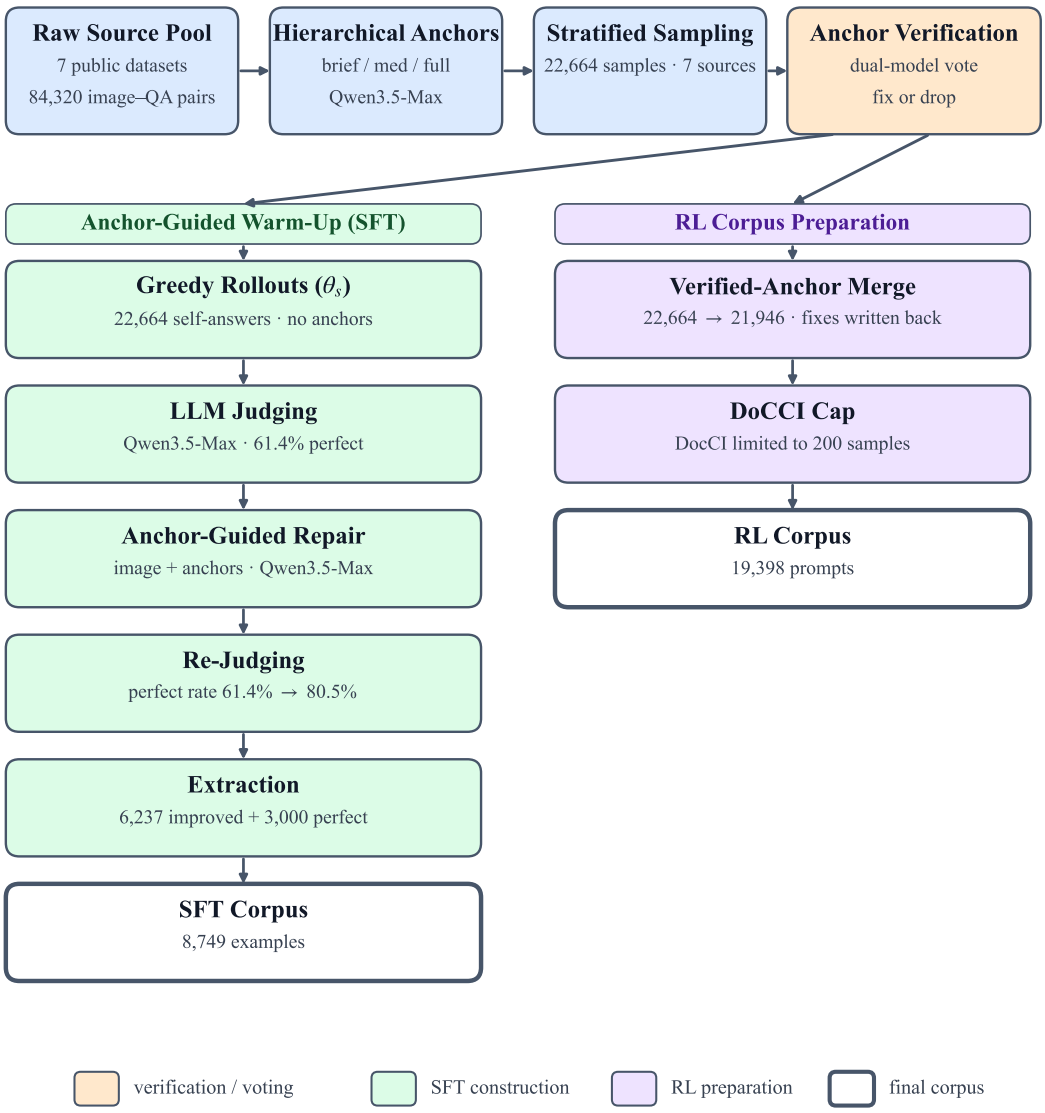

Figure A1: Data-construction pipeline of ASYM-OPSD. A shared trunk pools 84,320 image–question–answer pairs, generates hierarchical anchors, draws a stratified 22,664-sample root, and verifies every anchor by dual-model voting before splitting into the anchor-guided warm-up lane (left) and the RL-preparation lane (right). Both final corpora are outlined in white; all counts are measured from the materialized artifacts.

APPENDIX

A.1 DATASETS CONSTRUCTION DETAIL

This section documents the complete data-construction pipeline of ASYM-OPSD: how the sources are pooled and stratified, how the hierarchical visual anchors are generated and verified, how the anchor-guided warm-up corpus is distilled from the policy's own rollouts, and how the RL corpus is finalized. Figure A1 gives an overview. All counts reported below are measured from the materialized artifacts of each pipeline stage rather than estimated, and every sampling decision uses fixed seeds.

Table A1: Composition of the source pool and the stratified sampling used for the training root. †One imageless HaloQuest entry is discarded before sampling; the source is otherwise retained in full.

Source	Task	Pool	Sampled	Rate (%)
LRV	reasoning verification	29,238	2,000	6.8
ChartQA	chart QA	25,000	5,000	20.0
DocCI	detailed captioning	10,000	2,000	20.0
VSR	spatial reasoning	10,000	4,500	45.0
HaloQuest	hallucination QA	5,453	5,452	99.98†
Geo3K	geometry	3,917	3,000	76.6
MapQA	map QA	712	712	100.0
Total	—	**84,320**	**22,664**	26.9

Table A2: Anchor length statistics over the 22,664-sample root, in characters; every sampled item has a non-empty anchor at all four levels.

Granularity	Mean	Median	P90
Brief (a_{brief})	235	218	354
Medium (a_{med})	666	652	837
Full (a_{full})	1,923	2,033	2,426
Regular (source-native)	1,068	1,116	1,848

Source pool and stratified sampling. The source pool aggregates 84,320 single-image question–answer pairs from seven public datasets spanning complementary visual-linguistic competencies (Table A1): LRV for visual reasoning verification, ChartQA for chart question answering, DocCI for detailed captioning, VSR for spatial-relation reasoning, HaloQuest for hallucination-prone queries involving visual challenges, false premises, and insufficient context, Geo3K for geometric problem solving, and MapQA for map reading. A pre-filtering pass discards one HaloQuest entry whose image is missing. From this pool we draw a 22,664-sample stratified subset that forms the common root of both training lanes: hallucination-prone sources with limited data (HaloQuest, MapQA) are retained in full; VSR and Geo3K contribute 45% and 76.6% of their pools; the remaining sources are sampled at fixed quotas of 6.8%–20% to balance domain representation; and the LRV split is balanced 1,000/1,000 between positive and negative verification pairs.

Hierarchical anchor generation. Visual anchors at all granularities are generated by the Qwen3.5-Max LLM. Every generation prompt couples the image with a persona matched to the source (*e.g.*, chart-analysis expert for ChartQA, geometry expert for Geo3K) and, to suppress fabrication, conditions on the gold question–answer pair as well as on source-specific ground-truth evidence where available: ChartQA prompts embed the reference data extracted from the underlying chart; VSR, Geo3K, and MapQA embed the author-provided reasoning chains; DocCI reuses its official long ground-truth description as the full-level anchor and only derives the medium- and brief-level variants from it. Following the design in Section 3.1, each sample receives brief (a_{brief}), medium (a_{med}), and full (a_{full}) anchors; for hallucination-prone HaloQuest, the brief persona additionally doubles as a verification expert that must state explicitly what is and is not present, matching the dataset's false-premise queries. Each item also retains a source-native *regular* description when the parent dataset provides one, which serves as a fallback in the anchor-selection order (brief, regular, medium, full) used by the repair and warm-up prompts. All 22,664 sampled items carry non-empty anchors at every granularity, and the resulting length distributions separate cleanly across levels (Table A2).

Dual-model anchor verification. Because the anchor generator itself can hallucinate, every anchor entering either training lane is audited before use by two independent commercial frontier VLMs, Gemini-3.5-Flash and GPT-5.6-SOL, whose checkpoints are held fixed throughout—each returning a verdict with the located errors and a suggested revision. To keep the bar aligned with the anchors' role, an anchor is ruled incorrect only when it fabricates an element absent from the image or directly contradicts visible content; minor imprecisions of wording (*e.g.*, "white" for "light gray") are tolerated, since an anchor is a summary rather than an exhaustive inventory. Voting is deliberately lenient toward the data: an anchor passes if either model deems it correct, and only anchors rejected

Table A3: Outcomes of dual-model anchor verification. "Pass": accepted without modification. "Repaired": accepted after one repair round. "Dropped": the sample is removed from the training data. The brief- and medium-level rows audit the same RL corpus at two anchor granularities; their drops overlap, and the union of the two sets (1,466 records) is what the RL corpus finalization removes.

Setting	#	Pass	Repaired	Dropped
RL corpus, brief level	22,664	15,382	6,594	688
RL corpus, medium level	22,664	17,102	4,710	852
SFT + RL, brief level	31,901	26,938	4,900	63

by both models proceed to repair, in which the primary auditor, conditioned on the image, the offending anchor, and both models' error reports, rewrites only the identified claims while preserving the verified content and phrasing. After one repair round the anchor faces a final dual vote; if it is still rejected, or the repair degenerates (*e.g.*, fewer than ten characters), the underlying sample is removed from the training data. Table A3 reports the outcomes: anchors pass untouched in 67.9% of audits at the brief level and 75.5% at the medium level; after one repair attempt, only 3.0–3.8% of samples are eventually dropped. The extracted warm-up records are audited downstream by the answer-level stage of the next paragraph rather than here.

Self-answered rollouts, judging, and repair. The anchor-guided warm-up corpus is distilled from the policy's own rollouts (Section 3.2). The base model (Qwen3-VL-4B-Instruct) answers all 22,664 prompts greedily (temperature 0, up to 2,048 tokens) from the standard input of image and question—without anchors, mirroring the deployment condition. A fixed text-only judge (Qwen3.5-Max) then scores every response against the reference answer on a $[0, 1]$ scale with few-shot calibration: 13,923 responses (61.4%) receive a perfect 1.0, 2,302 (10.2%) score 0, and the rest fall in between. Every response below 1.0 is rewritten by Qwen3.5-Max conditioned on the image and the selected anchor, in one of two modes. *Alignment* mode, for sources whose flawed but well-formed drafts are worth preserving (HaloQuest, ChartQA, DocCI, LRV, MapQA), acts as a factual-alignment editor bound by hard constraints—no truncation, no broken reasoning, no evaluative meta-text—so that the output still reads as a first-time direct answer to the question. *Analyst* mode, for sources best answered by reconstruction (Geo3K, VSR), builds a complete answer exclusively from the anchors. Re-scoring the repaired responses with the same judge raises the perfect rate from 61.4% to 80.5%: 6,237 responses (27.5%) strictly improve, 1,034 (4.6%) regress, and regressions are discarded from the warm-up set. Per-source improvement rates (Table A4) are highest for open-ended captioning and geometric reconstruction (DocCI 61.6%, Geo3K 58.3%), where hallucination leaves the most room for correction, and lowest for ChartQA (9.1%), whose rollouts are already largely accurate.

Extraction and answer-level auditing. Extraction keeps every strictly improved response (6,237) plus a fixed-seed random subsample of 3,000 perfect rollouts, yielding 9,237 warm-up candidates for the Qwen3-VL-4B teacher (9,658 for the InternVL3.5-8B teacher, produced by the identical procedure on that model's own rollouts; Table A4). Each record conditions the student on the anchor—input: image, brief-level anchors, and question; target: the repaired (or originally perfect) response—while special vision tokens and unescaped angle brackets are sanitized to prevent prompt injection. A second dual-model stage then audits the *answers* for three failure modes: *meta-text leakage* (anchor, draft, or reasoning prefixes surfacing in the response), *residual hallucinated details* (a correct main claim retaining an uncorrected color, count, or entity inherited from the flawed draft), and *visual factual errors* that contradict the image. The two auditors—an image-grounded VLM and a text-only LLM—have complementary blind spots, and an answer passes if either deems it correct; answers rejected by both undergo minimal-edit repair subject to hard style-preservation constraints: the rewrite must match the original's length, structure, and tone, and any repair that substantially shortens the answer is rejected and re-attempted, which prevents the corpus from drifting toward terse responses over successive repairs. A final adversarial filter applies a stricter policy—any incorrect verdict on leakage or residual details discards the record, while visual factual errors require both models to agree—and replaces every anchor with its verified version. This pass rewrites 2,735 responses and, by the identical audit on each model's own candidates, removes 488 of the 9,237 Qwen3-VL-4B candidates and 909 of the 9,658 InternVL3.5-8B candidates, leaving 8,749 examples in the final warm-up corpus used for training for each of the two teachers.

Table A4: Per-source results of anchor-guided repair and the resulting warm-up composition. "Improved" counts rollouts whose repaired version strictly raises the judge score; "Rate" is the improvement fraction per source; "SFT (%)" is the number of extracted examples (improved plus sampled perfect) with corpus share in parentheses, shown for the Qwen3-VL-4B teacher before the answer-level audit; the audit removes 488 candidates (909 for InternVL3.5-8B) and leaves 8,749 examples per teacher for training.

Source	Sampled	Improved	Rate (%)	SFT (%)
HaloQuest	5,452	1,304	23.9	2,109 (22.8)
Geo3K	3,000	1,750	58.3	1,799 (19.5)
VSR	4,500	834	18.5	1,625 (17.6)
ChartQA	5,000	453	9.1	1,382 (15.0)
DocCI	2,000	1,231	61.6	1,271 (13.8)
LRV	2,000	458	22.9	750 (8.1)
MapQA	712	207	29.1	301 (3.3)
Total	**22,664**	**6,237**	**27.5**	**9,237 (100)**

RL corpus finalization. The RL corpus inherits the 22,664-sample root. Verification outcomes are written back: anchors judged correct pass through unchanged, repaired anchors replace their flawed originals, and the 1,466 samples rejected by the anchor-verification passes (the union of the brief- and medium-level drops of Table A3) are removed. A final registry cap then limits DocCI—the only purely descriptive source—to 200 examples so that its initial 2,000-sample allocation cannot dominate rollout batches with captioning-style outputs ($22{,}664 - 1{,}466 - 1{,}800 = 19{,}398$), leaving 19,398 prompts. At training time, the policy samples from the anchor-free input (I, q); anchors enter only the teacher stream via token-level forced decoding (Section 3.3), so the RL corpus stores all four anchor levels per prompt without ever exposing them to the policy.

A.2 Training Detail

This section records the implementation of both training stages of ASYM-OPSD—the anchor-guided warm-up (Section 3.2) and the asymmetric self-distillation RL stage (Section 3.3)—built on verl for RL and LLaMA-Factory for SFT. Every value below is measured from the deployed training scripts, framework defaults, and checkpoint metadata; hyper-parameters that a launch script does not override follow the framework defaults verbatim.

System stack. All experiments run on a single node with $8\times$NVIDIA A100-SXM4-80GB GPUs. The RL stack couples fully sharded data parallelism (FSDP full-sharding, no Megatron) for the actor updates with a vLLM engine for rollout generation and teacher log-probability evaluation; the warm-up stack runs LLaMA-Factory with DeepSpeed ZeRO-3. Software versions are PyTorch 2.10.0 (CUDA 12.8) and vLLM 0.17.0. Base weights are the official HuggingFace releases—Qwen3-VL-4B-Instruct and InternVL3.5-8B—held as local `bf16` copies without further revision pinning.

Warm-up (SFT). Table A5 lists the warm-up configuration. Both backbones are fine-tuned in full-parameter mode with the vision tower and projector frozen, so only the language model updates—matching the RL-side freezing below. The chat template disables the thinking mode for Qwen3-VL-4B (`qwen3_vl_nothink`) so that targets match deployment-style direct answers; InternVL3.5-8B uses its native template. Training runs 2 epochs at a peak learning rate of 1.0×10^{-5} with a cosine schedule and 10% warmup in `bf16` under ZeRO-3, with images capped at 262,144 pixels; successive data generations follow Appendix A.1.

RL configuration. Table A6 compares the five configurations behind the reported training-based methods; all five runs optimize on the 19,398-prompt RL corpus of Appendix A.1. The advantage estimator `grpo_dense` implements Equation 2 with group statistics computed jointly over the response tokens of all n rollouts of a prompt, and the per-token reward is the clamped log-probability gap of Equation 1. The GRPO baseline replaces the token-level teacher signal with an outcome reward from an internal Qwen3-8B judge (paired True/False verdicts mapped to $\pm 1/0$) under the standard sequence-level estimator.

Table A5: Warm-up (SFT) configuration, measured from the LLaMA-Factory training recipes. The vision tower and projector are frozen in both runs.

	Qwen3-VL-4B	**InternVL3.5-8B**
Stage / precision	full-parameter SFT, `bf16`, ZeRO-3	
Chat template	`qwen3_vl_nothink`	`intern_vl`
Per-device batch × grad. accum.	1×2	2×2
Learning rate	1.0×10^{-5}, cosine, 10% warmup	
Epochs	2	
Image max pixels	262,144	
Frozen modules	vision tower, projector	
Save frequency	every 500 steps	

Table A6: RL configurations of the reported training-based methods, measured from the launch scripts. "4B"/"8B" abbreviate Qwen3-VL-4B/InternVL3.5-8B. All five runs share the optimization and rollout defaults listed in the text and freeze the vision tower throughout.

	GRPO 4B	**OPSD 4B**	**OPSD 8B**	**ASYM-OPSD 4B**	**ASYM-OPSD 8B**
Advantage estimator	`grpo`	`grpo_dense`	`grpo_dense`	`grpo_dense`	`grpo_dense`
Learning rate	2e−6	5e−7	2e−6	2e−6	2e−6
Train batch size	256	128	256	256	256
Rollout group n	16	16	16	32	20
β_{KL} (low-var. KL)	0.02	0.1	0.02	0.02	0.005
KL applied	in reward	in loss	in reward	in reward	in reward
Entropy coefficient	0	−0.01	0	−0.01	0
Reward source	outcome (LLM judge)	token-level log-prob gap (Eq. 1)			
Reward clamp	—	$[-8, 8]$			
Epochs / save frequency	3 epochs / every 20 steps				

Shared optimization and rollout defaults. All runs inherit the framework defaults: PPO clipping at $\epsilon = 0.2$ (symmetric), token-mean loss aggregation, gradient clipping at 1.0, AdamW with betas $(0.9, 0.999)$ and weight decay 0.01, and a constant learning-rate schedule without warmup. Prompts and responses are capped at 2,048 tokens each, with one train batch split into a PPO mini-batch of 128 and micro-batches of 4 sequences per GPU. Rollouts are drawn by vLLM with sample-based decoding at temperature 1.0 and top-p 1.0; log-probability recomputation uses micro-batches of 4 sequences per GPU; the KV cache is freed after every step, and the multimodal preprocessing cache is disabled to prevent image–text misalignment. The low-variance KL toward the frozen reference model of Equation 3—initialized identically to the actor—is applied with the coefficient of Table A6, folded into the per-token reward. Both backbones train with gradient checkpointing and frozen vision towers; InternVL3.5-8B additionally wraps its decoder layers outside the FSDP sharding scope for compatibility. Online validation is disabled throughout (`test_freq` $= 0$).

Teacher log-probability serving. Teacher log-probabilities are computed by a dedicated service rather than co-located forward calls: a single asynchronous gateway load-balances `/compute_logprobs` requests over 8 independent vLLM workers (tensor-parallel size 1), each holding the frozen warm-up checkpoint of the corresponding backbone in `bf16` with a 10,240-token context window and up to 60% GPU utilization. Requests are chunked into groups of 32 sequences with at most 8 in flight and 720 s timeouts, and a keep-alive thread keeps the training GPUs busy during the round-trip. For each rollout, the teacher prompt is assembled at the token-ID level: the anchor text is spliced into the original prompt sequence around its preserved vision placeholder tokens—no decode-re-encode round-trip is involved—and anchors are chosen with the priority *verified* → *brief* → *regular* → *medium* → *full*, the Appendix A.1 selection order with verified anchors taking precedence. Raw per-token rewards are sanitized (NaN → 0, infinite values clamped) and then clamped to $[-8, 8]$ as stated in Section 3.3.

Evaluation and decoding-baseline configuration. The base models and all trained checkpoints are evaluated with vLLM under greedy decoding (temperature 0, up to 4,096 new tokens, generation batch 16, `bf16`), following the protocol of Section 4.1. The three inference-time baselines run on custom HuggingFace decoding loops (batch size 1 per sample, with a processor shim for

Table A7: Signal-quality summary of Figure 5: token-level detection over the 20 annotated responses, median per-response standardized (grounded − hallucinated) gap, and a two-sided sign test over the 20 per-response gaps. **Bold**: best per backbone group.

Backbone	Signal	Acc↑	Prec↑	Rec↑	Med. gap↑	sign p
Qwen3-VL-4B	Anchor (ours)	**0.60**	**0.62**	0.53	+0.35	<0.001
	Image-contrast	0.49	0.48	0.46	+0.06	0.82
	Entropy	0.52	0.49	**0.60**	+0.05	0.82
InternVL3.5-8B	Anchor (ours)	**0.62**	**0.59**	0.66	+0.46	<0.001
	Image-contrast	0.48	0.46	0.41	−0.13	0.50
	Entropy	0.55	0.51	**0.64**	+0.29	0.26

Table A8: Exact scores underlying the radar comparison of Figure 3 on the six general-VQA benchmarks. The gray row is the raw base model; the blue row is ASYM-OPSD.

	GQA↑	VizWiz↑	VQAv2↑	TextVQA↑	LLaVA-W↑	MathV↑
Qwen3-VL-4B						
Raw	**62.89**	68.90	81.19	79.71	**74.30**	74.80
SFT	59.33	66.36	77.87	78.16	68.50	65.02
OPSD	59.54	68.37	78.50	76.15	67.40	60.10
GRPO	60.81	68.97	80.31	78.63	72.70	70.70
VCD	60.02	67.48	78.63	73.82	66.80	66.30
LEAD	61.82	67.34	81.11	79.80	70.70	69.50
AdaIAT	**62.89**	68.14	**81.89**	77.86	68.80	65.70
ASYM-OPSD	62.01	**69.41**	81.11	**81.47**	**74.50**	**75.50**
InternVL3.5-8B						
Raw	59.26	56.81	76.73	**77.98**	62.10	70.30
SFT	56.79	**59.30**	74.23	77.24	62.20	60.70
OPSD	57.52	57.93	76.82	77.92	66.60	66.70
GRPO	59.42	56.85	76.91	77.91	66.50	66.70
VCD	56.44	52.53	71.88	69.72	60.10	63.50
LEAD	57.50	57.52	74.48	71.90	62.10	66.20
AdaIAT	58.69	**59.30**	75.31	73.45	58.80	60.00
ASYM-OPSD	**59.70**	58.20	**77.20**	77.83	**69.10**	**73.50**

InternVL3.5-8B) under a shared sampler of temperature 0.6, top-p 0.95, and top-k 20 with up to 4,096 new tokens and seed 42, each following its original implementation:

- **VCD** perturbs the image with 500 forward diffusion-noise steps and contrasts the original and noised logits at strength $\alpha = 1.0$ with adaptive plausibility truncation at $\beta = 0.1$.
- **LEAD** alternates between discrete-token and continuous latent modes through entropy hysteresis ($b_1 = 1.0$, $b_2 = 0.2$, dwell window of 256 steps, at most 5 mode switches), annealing the two mixing coefficients from initial values $(\alpha_0, \beta_0) = (0.6, 0.7)$ toward 1 as generation proceeds.
- **AdaIAT** amplifies attention from generated text toward visual tokens with gain $\alpha = 6.0$ over the middle 15%–55% of layers, threshold interpolation $\beta = 0.5$, and uniform per-head direction coefficients, requiring eager attention (attention outputs are recomputed from the KV cache after each reweighting step).

A.3 More Experiment Results

Table A7 completes the localization analysis of Section 4.4: it tabulates, for each signal and backbone of Figure 5, the token-level detection statistics of the parameter-free mean-threshold detector, the median per-response standardized gap, and the two-sided sign test over the 20 per-response gaps.

Table A8 lists the exact scores behind the radar comparison of Figure 3.

Table A9: Overview of the four evaluation protocols.

Benchmark	Judge	Judge input	Special rules
CHAIR	none (rules)	—	mention outside GT = hallucination
HallusionBench	LLM, three-way	text only	imageless VS: "unclear" counts correct
HaloQuest	LLM, three-way	text only	only "correct" counts; subgroups
MMHal-Bench	VLM, rating	image + text	0–6 rating; < 3 is hallucination

A.4 Evaluation Protocol

This section documents the final protocol under which the four hallucination benchmarks of Table 2 are scored. All models answer greedily (temperature 0, at most 4,096 tokens, batch 16) with vLLM. Verdicts on the three judged benchmarks come from a fixed VLM-based API evaluator, Qwen3.5-Max; CHAIR is fully rule-based. Judgments are cached per item, so interrupted runs resume without re-querying. The evaluated sets are the full splits: CHAIR 500, HallusionBench 1,129, HaloQuest 600, and MMHal-Bench 96 samples (Table A9).

CHAIR (MSCOCO). 500 images are captioned under the prompt *"Please describe this image in detail."* and scored against MSCOCO ground-truth object lists, with no judge involved. The word-extraction pipeline tokenizes and part-of-speech-tags each caption (NLTK), lemmatizes WordNet-style, merges two-word compounds (*hot dog*, *traffic light*), removes known compositive false positives (*toilet* suppressing *seat/bowl*), disambiguates color homonyms (*orange* preceding a noun reads as a color, not a fruit), and maps about 400 synonyms onto the 80 MSCOCO object classes. Every object mentioned outside its caption's ground-truth list is a hallucinated instance:

$$\mathrm{CHAIR}_i = 100 \cdot \frac{\#\text{ hallucinated object mentions}}{\#\text{ object mentions}}, \qquad \mathrm{CHAIR}_s = 100 \cdot \frac{\#\text{ captions containing hallucination}}{\#\text{ captions}},$$

with Recall, Precision, and F1 computed against the ground-truth object lists in the usual way.

HallusionBench. The 1,129 samples are judged three-way (*correct/incorrect/unclear*) by a text-only judge that receives the question, an explanatory reference answer, and the prediction. The verbatim prompt is:

HallusionBench judge prompt (verbatim; the HaloQuest variant replaces the reference line — see below)

Imagine you are an intelligent teacher. Thoroughly read the question, reference answer and the prediction answer to ensure a clear understanding of the information provided. Assess the correctness of the predictions.

If the prediction answer does not conflict with the reference answer, please generate 'correct'. If the prediction answer conflict with the reference answer, please generate 'incorrect'. If the prediction answer is unclear about the answer, please generate 'unclear'.

Question: {question}
Reference answer: {gt_answer_details}
Prediction answer: {model_pred}
Output:

The lower-cased verdict maps "incorrect" $\to 0$, any remaining "correct" $\to 1$, else "unclear" $\to 2$. A prediction is credited only when judged *correct*, with one official exception: imageless visual-scene items (category VS with an empty figure) also accept *unclear*. Reported metrics: **aAcc** (simple per-sample accuracy), **fAcc** (group samples by source figure; credit the figure only if every question of it is correct, imageless items excluded), and **qAcc** (group samples by question identity; credit the question only if all its phrasings are correct).

HaloQuest. The 600 samples are judged with the same verbatim prompt as above, replacing the reference line by `Reference answer: {groundtruth_responses}`, which concatenates all equivalent ground-truth phrasings separated by semicolons. Only a *correct* verdict is credited; *incorrect* and *unclear* are both rejected. Reported metrics: the overall accuracy, plus subgroup accuracies by hallucination type (*false premises*, *insufficient context*) and by image origin (*generated* vs. *real*).

MMHal-Bench. The 96 samples are the only image-conditioned ones: the judge sees the original image and rates the answer 0–6 under the official rubric template, augmented as in our deployment with guidelines that separate *core* (wrong main answer, rated 0–2) from *peripheral* hallucinations (correct main answer with minor inaccuracies, rated ≥ 4). The verbatim prompt (braces mark injected fields) is:

MMHal-Bench judge prompt (verbatim)

Please act as an impartial and objective judge and evaluate the quality of the response provided by a Large Multimodal Model (LMM) to the user question. Your evaluation should be mainly based on whether the response is informative, and whether the response contains any hallucination. Hallucination, in this context, refers to a situation where the LMM generates a response that includes information not present or implied in the image or previous conversation. A hallucination could be a false claim about an object, action, emotion, or any other detail that is not grounded in the image.

For clarity, consider these examples:

Example 1:
- Image Contents: Lynx, Cat
- Question: How many cats can you see in this image?
- Standard Human-Generated Answer: There are two cats visible in the image.
- LMM Response to Evaluate: There is one cat in the image.
- Analysis: The LMM's response is false because it claims there is only one cat in the image, when in fact there are two.
- Hallucination: Yes.

Example 2:
- Image Contents: Tree, Plant, Panda, Animal
- Question: What can you see from this photo?
- Standard Human-Generated Answer: In this photo we can see branches. We can see a panda on a branch. The background is blurry.
- LMM Response to Evaluate: In the photo, I can see a large black and white panda sitting on a tree branch. The image captures the panda's unique appearance and its natural habitat, showcasing the beauty and rarity of these animals.
- Analysis: The LMM's response is not false because it does not make any false claims about the image contents. For example, it does not claim that the panda is eating bamboo.
- Hallucination: No.

Example 3:
- Image Contents: Fashion accessory, Mammal, Dog, Umbrella
- Question: Describe this image in details.
- Standard Human-Generated Answer: In this picture I can see a dog laying under a yellow umbrella on a beach, and there is a rope on the right side. In the background I can see few persons talking or lying. At the top of the image, I can see the blue and clear sky.
- LMM Response to Evaluate: The image features a brown dog laying on the sand at the beach, underneath a yellow umbrella. The dog is positioned near the water, and there are people in the background. The beach is sandy, and the sun is shining, creating a pleasant atmosphere. The dog is enjoying the shade provided by the umbrella, which is a common way to protect pets from the sun's harmful rays. The presence of people in the background suggests that the beach is a popular spot for leisure activities, such as sunbathing, swimming, and relaxing.
- Analysis: The LMM's response is false because it claims that the dog is positioned near the water, but there is no water in the image.
- Hallucination: Yes.

Example 4:
- Image Contents: Clothing, Luggage and bags, Man, Building, Woman, Person
- Question: Where is this photo taken?
- Standard Human-Generated Answer: From the slogan "Welcome to Houston, Texas", we can infer that this photo is taken at an airport in Houston, Texas.
- LMM Response to Evaluate: This photo is taken at an airport, specifically at the entrance of a terminal. However, the specific location of this airport cannot be determined from the image.
- Analysis: Although the LMM's response does not explicitly mention the location "Houston, Texas" and does not perfectly answer the question, it makes no false claims about the image contents. Therefore, this response is not considered hallucination. It is OK if the LMM fails to directly answer the question, expresses uncertainty, or cannot include all the objects or details in the human-generated answer.
- Hallucination: No.

Example 5:

- Image Contents: Tree, Tower
- Question: What environmental concerns are involved in this photo?
- Standard Human-Generated Answer: This photo shows a smokestack emitting smoke. It may cause severe air pollution which threatens human health.
- LMM Response to Evaluate: In the photo, there are several environmental concerns related to the smokestack emitting smoke. The smoke from the smokestack is a byproduct of industrial processes, which can contribute to air pollution, climate change, and negative impacts on human health. The smoke contains harmful pollutants, such as particulate matter, sulfur dioxide, nitrogen oxides, and carbon monoxide, which can lead to respiratory problems, heart disease, and other health issues. Additionally, the smoke contributes to the greenhouse effect, which can lead to global warming and climate change. The smokestack's emissions also affect the environment, as they can harm wildlife, vegetation, and ecosystems. Therefore, it is essential to address these environmental concerns by implementing measures to reduce emissions and promote sustainable practices in industrial processes.
- Analysis: Although the LMM's response is significantly longer than the standard human-generated answer, it does not contain any false claims about the image contents. Instead, it provides additional general information about the environmental concerns, which can be inferred from the smoke emission. Such detailed analysis or reasoning should be considered as a positive aspect, as long as it contains no false claims.
- Hallucination: No.

With these examples in mind, please help me evaluate whether the response by the LMM is informative, and whether hallucination exists in it, based on the comparison between the LMM's response and the factual information provided in the image contents, question, and the standard human-generated answer below.

Please note that the standard human-generated answer may only contain factual information but may not give a detailed analysis. Also, the standard human-generated answer may not be completely comprehensive in describing all the objects and their attributes, so please be a bit more cautious during evalutation. LMM's detailed analysis or reasoning should be encouraged.

IMPORTANT GUIDELINES for distinguishing core vs. peripheral hallucinations:
1. Core hallucination: The LMM provides a directly wrong answer to the main question asked by the user. This is a severe hallucination and should be rated 0-2.
2. Peripheral hallucination: The LMM correctly answers the main question but contains minor inaccuracies or over-descriptions in supplementary details. This should NOT be treated as a severe hallucination. If the main answer is correct, the rating should be at least 4, even if some minor supplementary details are inaccurate.
3. Speculative language: Claims qualified with uncertainty markers (e.g., "possibly", "likely", "appears to be", "may be") express uncertainty rather than asserting false information. Such qualified claims should NOT be considered hallucinations.
4. Reasonable inferences: Detailed descriptions that provide additional plausible context (e.g., describing the setting, atmosphere, or likely purpose of objects visible in the image) should be treated as informative analysis, not hallucination, as long as they do not directly contradict the image.
5. Proportionality: The rating should reflect the overall quality of the response. A minor incorrect detail in an otherwise accurate and informative response should not reduce the rating below 4.

To evaluate the LMM responses, first, begin your evaluation by providing a short explanation. Focus on whether the LMM correctly answers the main question. Second, after providing your explanation, you must rate the response by choosing from the following options:
- Rating: 6, very informative with good analysis or reasoning, no hallucination
- Rating: 5, very informative, no hallucination
- Rating: 4, correctly answers the main question, may have minor peripheral inaccuracies
- Rating: 3, not informative, no hallucination
- Rating: 2, very informative, but with core hallucination (wrong main answer)
- Rating: 1, somewhat informative, but with core hallucination (wrong main answer)
- Rating: 0, not informative, with core hallucination

Image Contents
{image_content}

Question
{question}

Standard Human-Generated Answer
{gt_answer}

LMM Response to Evaluate
{model_pred}

The rating is parsed with the regular expression `rating:\s*([0-6])` (parse failures default to 0), and a sample counts as hallucinated when its rating is below 3. The protocol yields the average rating over the 96 samples, the hallucination rate (the fraction of ratings below 3), and per-type average ratings over the nine MMHal question types (object, attribute, adversarial, comparison, count, cultural, commonsense, temporal, text).

A.5 Case Study

We close with three HallusionBench cases in which the raw backbone answers incorrectly and the final ASYM-OPSD checkpoint flips the verdict to correct under the judge of Appendix A.4 (selection: judge-*incorrect* base answer and judge-*correct* ASYM-OPSD answer). The three cases span both backbones and three HallusionBench subcategories—figure understanding, chart reading, and video reasoning—one on Qwen3-VL-4B and two on InternVL3.5-8B. Each panel reproduces the image, the question with its ground-truth answer, and both responses verbatim.

Case Study 1 — Figure understanding: vegetables or fruits (VD/figure, Qwen3-VL-4B)



User Question:
Are all the items in this figure vegetables or fruits?

Ground Truth: No, there is some meat in this figure.

Base Response: Qwen3-VL-4B-Instruct (judged **incorrect**)

No, not all the items in this figure are vegetables — some are fruits.

Here's a breakdown:

Vegetables (botanically, these are the edible parts of plants, often stems, leaves, roots, or flowers):
- Broccoli
- Kale (or curly lettuce)
- Red and yellow bell peppers
- Tomatoes (though botanically fruits, they are commonly used as vegetables in cooking)
- Cucumbers
- Celery
- Onions
- Radishes
- Carrots (partially visible)
- Beets (partially visible)
- Lettuce (especially the red leaf variety)

Fruits (botanically, these are the seed-bearing structures of plants):
- Tomatoes (as mentioned, botanically fruits, but often treated as vegetables in culinary contexts)
- Lemons (a citrus fruit)
- Green and red cherry tomatoes (also fruits)
- A small green lime (a citrus fruit)

Important note: In culinary contexts, tomatoes, cucumbers, and peppers are often classified as vegetables, even though botanically they are fruits. But since the question is asking whether *all* items are vegetables or fruits, and we have **lemons and limes** — which are clearly fruits — the answer is **no, not all items are vegetables**.

So, the correct answer is: **No, not all items are vegetables — some are fruits (like lemons, limes, and tomatoes).**

Final Answer: **No.**

ASYM-OPSD Response (judged **correct**)

No, not all the items in this figure are vegetables or fruits — there is one item that is **not** a vegetable or fruit.

The image displays a variety of colorful produce, including:

- **Vegetables**: broccoli, red and yellow bell peppers, tomatoes, cucumbers, celery, red onions, radishes, and leafy greens like kale and lettuce.
- **Fruits**: a lemon and a lime.

However, there is also a **piece of cooked meat** visible on the right side of the image — a slice of what appears to be **cooked steak** or **roast meat** — which is neither a vegetable nor a fruit.

So, the answer is: **No, not all items are vegetables or fruits — there is meat (cooked steak) in the image.**

Case Study 2 — Chart reading: export recipients in 2013 (VS/chart, InternVL3.5-8B)

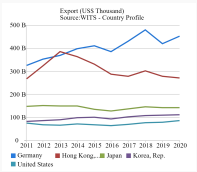

User Question:
According to the image, in 2013, was Hong Kong the largest recipient of exports from China in terms of US dollars, followed by the United States and Japan?

Ground Truth: No. According to the image, in 2013, Hong Kong was the largest recipient of exports from China in terms of US dollars, followed by Germany and Japan.

Base Response: InternVL3.5-8B (judged **incorrect**)

No. According to the image, in 2013, Germany was the largest recipient of exports from China in terms of US dollars, followed by Hong Kong, Japan, Korea, Rep., and the United States.

ASYM-OPSD Response (judged **correct**)

No, according to the image, in 2013, Hong Kong was the largest recipient of exports from China in terms of US dollars, followed by Germany and Japan. The United States was not the second-largest recipient in 2013.

Case Study 3 — Video reasoning: three-pointer or dunk (VD/video, InternVL3.5-8B)



User Question:
According to the positive sequence of the images, does the basketball player make a three-point shooting?

Ground Truth: The basketball player makes a dunk.

Base Response: InternVL3.5-8B (judged **incorrect**)

Yes, the basketball player appears to make a three-point shot based on the sequence of images. The player is shown dribbling, jumping, and releasing the ball towards the basket, with the final image suggesting the ball is heading towards the hoop.

ASYM-OPSD Response (judged **correct**)

No, the basketball player does not make a three-point shooting according to the sequence of the images.

Here's the sequence of events depicted in the images:

1. The player in white is dribbling the ball while being guarded by a player in red.
2. The player in white is driving towards the basket, attempting to get past the defender.
3. The player in white is jumping and appears to be attempting a layup or dunk, with the defender trying to block the shot.
4. The ball is seen going through the hoop, indicating a successful shot.

Since the player is attempting a layup or dunk, not a three-point shot, the answer is no, the player does not make a three-point shooting.