

Revealing Multi-View Hallucination in Large Vision-Language Models

Wooje Park^{1*}, Insu Lee^{1*}, Soohyun Kim¹, Jaeyun Jang¹,
Minyoung Noh¹, Kyuhong Shim², Byonghyo Shim¹

¹Seoul National University, ²Sungkyunkwan University

{wjpark, islee, soohyunkim, jyjjang, mynoh, bshim}@islab.snu.ac.kr, khshim@skku.edu

Abstract

Large vision-language models (LVLMs) are increasingly being applied to multi-view image inputs captured from diverse viewpoints. Despite this growing use, current LVLMs often generate incorrect responses due to visual interference from non-target instances or viewpoints, a phenomenon we term multi-view hallucination (MVH). To systematically analyze this problem, we construct MVH-Bench, a benchmark comprising 4.8k question-answer pairs targeting two types of hallucination: *cross-instance* and *cross-view*. Empirical results show that MVH is prevalent across recent LVLMs. To address this issue, we propose Reference Shift Contrastive Decoding (RSCD), a training-free decoding technique that suppresses visual interference by generating negative logits through attention masking. Experiments on MVH-Bench with LLaVA-OneVision and Qwen2.5-VL demonstrate that RSCD improves performance by up to 25.7 and 94.8 points over existing hallucination mitigation methods, highlighting the effectiveness of RSCD.

1 Introduction

Large vision-language models (LVLMs) have shown remarkable advances in interpreting both visual and linguistic modalities and generating responses based on the given information (Hurst et al., 2024; Liu et al., 2024a; Bordes et al., 2024; Li et al., 2025d; Lee et al., 2026). Due to their capability to perceive complex visual scenes and describe them in natural language, LVLMs are widely employed as visual assistants in embodied AI, digital twins, and surveillance systems (Majumdar et al., 2024; Yuan et al., 2024; Yang et al., 2024; Gholizadeh Hamlabadi et al., 2025; Kim et al., 2026). A salient feature of these applications is that images captured from diverse viewpoints (i.e., multi-view images) are used as visual inputs to provide broader and more comprehensive information

about a scene that a single-view image cannot offer. Nevertheless, variations across different views of the same scene complicate the interpretation of rich visual information, making it difficult for LVLMs to produce precise responses (Hong et al., 2023; Park et al., 2025b; Lee et al., 2025).

In recent years, several approaches to enhance inter-view image correspondence have been proposed. One line of work aligns view-specific embeddings within a shared feature space, whereas another focuses on constructing a unified scene representation for comprehensive scene understanding (Islam et al., 2024; Hong et al., 2023; Park et al., 2025b; Lee et al., 2025). While these approaches enhance the information aggregation capabilities of LVLMs across multiple views, they fail to adequately judge which visual evidence corresponds to specific instances or viewpoints, resulting in significant perceptual errors.

Consider a scenario where LVLMs are queried about an instance in multi-view images (e.g., “What is the person wearing a white top doing?”) (see Figure 1). Although the correct answer is “pointing at flowers”, the model may confuse the queried instance with another instance, “the person wearing a black top”, and incorrectly respond with “holding a watering can”. Similarly, when LVLMs are queried about a specific viewpoint, such as “In view 1, which direction is the watering can pointing?”, the model may rely on another viewpoint (view 2), and answer “pointing upward” rather than the correct answer, “pointing to the left”.

We refer to this phenomenon as multi-view hallucination (MVH), in which LVLMs generate incorrect responses due to visual interference across multiple viewpoints. Building upon the examples described above, we categorize MVH into two types, depending on the cause of interference: *cross-instance hallucination*, where the model’s responses include mixed details from other instances, and *cross-view hallucination*, where the model

* Co-first authors.




Figure 1: Illustration of two types of multi-view hallucination in LVLMs, categorized by the source of the interference. *Cross-Instance Hallucination*: The model relies on information from another instance. *Cross-View Hallucination*: The model relies on information from another view.

references unintended viewpoints to produce responses.

To analyze MVH in LVLMs systematically, we design a new evaluation dataset dubbed MVH-Bench, comprising 4.8k question-answer pairs drawn from diverse multi-view scenarios. For each multi-view image set, MVH-Bench includes two pairs of binary questions and one pair of multiple-choice questions generated by interchanging descriptions across instances and/or viewpoints. Leveraging this paired-question structure, we design a set of dedicated evaluation metrics that measure the ability to provide answers grounded in sophisticated reasoning.

When we evaluate state-of-the-art (SOTA) closed and open-source LVLMs using the proposed MVH-Bench, we find that they are all confused by complex visual cues and thus struggle to generate correct answers. For example, we observe that models tend to generate incorrect responses based on visual cues associated with certain words or partial phrases in the query. To better investigate these phenomena, we compare the model’s responses before and after disrupting query context formation by blocking interactions between text tokens at each layer. Interestingly, we observe under this condition that the model’s responses shift toward non-target visual cues. This implies that information exchange between text tokens in the intermediate layers plays a key role in accurate visual grounding.

To alleviate visual grounding errors arising from insufficient query understanding, we propose a novel contrastive decoding technique referred to

as Reference Shift Contrastive Decoding (RSCD). The core idea of RSCD is to amplify the contrast between the original logits and those obtained with incomplete textual understanding, to suppress predictions grounded in non-target visual cues. To do so, we first obtain negative logits by intentionally disrupting text-to-text attention pathways in the intermediate layers, and then adjust the original logits in the direction opposite to the difference. Our experiments on MVH-Bench using LLaVA-OneVision (Li et al., 2025a) and Qwen2.5-VL (Bai et al., 2025b) show that RSCD consistently outperforms the recent decoding technique AvisC (Woo et al., 2025) by 25.7 and 94.8 points in overall benchmark score, respectively.

In summary, our main contributions are as follows:

- In this work, we characterize the MVH problem as an error caused by visual interference across multiple viewpoints and categorize it into two error types based on the source of the confusion.
- In order to scrutinize the MVH problem in detail, we design a new benchmark named MVH-Bench and propose dedicated accuracy metrics, viz., p-Acc and q-Acc, using which we evaluate 23 SOTA LVLMs.
- Through extensive experiments, we show that RSCD outperforms recent hallucination mitigation methods, improving MVH-Bench scores by up to 25.7 and 94.8 for LLaVA-OneVision and Qwen2.5-VL, respectively.





Figure 2: Overview of the MVH-Bench construction pipeline: (a) instance-descriptor pair extraction, and (b) automated question-answer generation followed by human verification.

2 Multi-View Hallucination Benchmark

2.1 Benchmark Design

In order to construct MVH-Bench, we use Ego-Exo4D (Grauman et al., 2024) and LEMMA (Jia et al., 2020), both of which provide synchronized multi-view videos of diverse scenes involving multi-person interactions. Following the procedure of Lee et al. (2025), we extract image frames from Ego-Exo4D and use the frames provided by the authors for LEMMA. We then pair images from diverse combinations of first-person and third-person views to cover a broad range of multi-view scenarios. The dataset construction process consists of three steps: 1) Instance-descriptor pair extraction, 2) Automated question-answer generation, and 3) Human verification (see Figure 2 for an overview).

2.1.1 Instance-Descriptor Pair Extraction

MVH-Bench includes four subcategories that closely reflect real-world scenarios: *action*, *object*, *numerical*, and *spatial* (Lee et al., 2025). For each subcategory, we extract subcategory-specific *instance-descriptor* pairs (I, D) from each view of every image pair using GPT-4o. Here, an *instance* uniquely identifies a specific person or object in an image, while a *descriptor* describes the state

or attribute of that instance. For example, in the action subcategory, an instance may be “a person wearing a black shirt” and the corresponding descriptor may be “holding a smartphone”. Given a multi-view image pair, let $\mathcal{S}_1 = \{(I_i, D_i)\}_{i=1}^{n_1}$ and $\mathcal{S}_2 = \{(I_j, D_j)\}_{j=1}^{n_2}$ denote the sets of extracted instance-descriptor pairs from the first and second views, respectively. The prompts used to extract instance-descriptor pairs are provided in Appendix F.5.

2.1.2 Automated Question-Answer Generation

We design MVH-Bench using two types of questions, binary and multiple-choice, each serving a distinct purpose. Using a binary question-answer (QA) format, we evaluate the model’s ability to correctly understand the relationships between instances and their corresponding descriptors or views. With a multiple-choice QA format, we assess its ability to distinguish the correct descriptor from plausible distractors. Using the instance-descriptor pairs from $\mathcal{S}_1$ and $\mathcal{S}_2$, we automatically generate questions for the two MVH categories: *cross-instance* and *cross-view*. For clarity, we describe the process using a single image pair from the action subcategory.

Cross-Instance Hallucination. We observe that models often confuse the queried instance with other instances and incorrectly respond based on their descriptors. To systematically diagnose this behavior, we evaluate the model’s ability to correctly ground the queried instance among other instances. Concretely, we sample (I_i, D_i) from $\mathcal{S}_1$ and (I_j, D_j) from $\mathcal{S}_2$ such that the instances and descriptors are mutually distinct (i.e., $I_i \neq I_j$ and $D_i \neq D_j$). We then construct binary and multiple-choice QA pairs using these. For the binary questions, we generate four QAs following the template:

Q: “Am/Is I_x D_y ?”

A: “Yes” if $x = y$, “No” otherwise.

For the multiple-choice questions, we generate two QA pairs using the template:

Q: “What is I_x doing?”

A) “ D_x ” B) “ D_y ”

C) “Neither D_x nor D_y ”

A: A) “ D_x ”

Cross-View Hallucination. Another failure mode is that the model tends to focus on instances from the non-queried view. To rigorously diagnose this phenomenon, we evaluate the model’s ability to correctly ground instance-descriptor pairs in the appropriate view. Specifically, in our benchmark, we deliberately design binary and multiple-choice QAs using pairs (I_i, D_i) and (I_j, D_j) with identical instances yet having distinct descriptors (i.e., $I_i = I_j$ and $D_i \neq D_j$). This design ensures that the same instance can be queried across different views. For the binary case, we generate four QA pairs following the template:

Q: “In view x , is I_x D_y ?”

A: “Yes” if $x = y$, “No” otherwise.

For the multiple-choice case, we generate two QA pairs using the template:

Q: “In view x , what is I_x doing?”

A) “ D_x ” B) “ D_y ”

C) “Neither D_x nor D_y ”

A: A) “ D_x ”

During evaluation, answer options in multiple-choice questions are randomly shuffled to make sure correct answers are distributed uniformly.

2.1.3 Human Verification

After automated QA generation, human annotators review every QA pair and remove those describing non-existent and/or incorrect content, as well as pairs that are trivially easy to solve. The resulting MVH-Bench consists of 3,200 binary and 1,600 multiple-choice QAs, evenly distributed across hallucination types and subcategories. The dataset is split into test and validation sets in a 9:1 ratio while preserving the original distribution. The validation set is used for hyperparameter search and analysis, while the test set is reserved for the final evaluation. Additional details of the human verification stage are provided in Appendix F.3.

2.2 MVH-Bench Evaluation

Utilizing the paired-question format in MVH-Bench, we propose a dedicated metric to assess whether a model correctly grounds its responses. For a formal expression of the evaluation metrics, please refer to Appendix F.2.

2.2.1 Binary Question

We evaluate the performance on binary questions using three accuracy metrics. The first is the mean accuracy (Acc) across all four QA pairs. In our setup, even questions whose correct answer is “No” contain elements that are present in the multi-view images. Consequently, it is highly likely that the model generates the answer “Yes” exclusively. To assess the model’s ability to truly ground the queried instance, we measure pair accuracy (p-Acc), which evaluates if the model correctly answers both the Yes and No questions associated with each view. Additionally, a model may appear to perform well by relying on a single view while failing to understand others. To verify consistent understanding across all views, we further measure quadruplet accuracy (q-Acc). In this metric, we consider a prediction correct only when all four binary questions are answered correctly. To quantify the model’s tendency toward biased responses, we also compute the yes-error ratio (YER), which is the fraction of Yes predictions among the incorrect cases.

2.2.2 Multiple-Choice Question

For the multiple-choice setting, we measure the mean accuracy (Acc) and pair accuracy (p-Acc). In pair accuracy, to assess model consistency across views, we consider a prediction correct when both questions derived from each multi-view image are




Figure 3: (a) Comparison between conventional and multi-view hallucination by their underlying causes. (b) Overview of the proposed RSCD, illustrating its core idea and underlying intuition.

answered correctly. In addition, to analyze the model’s error tendencies, we report the fraction of incorrect answers that choose the adversarial option B (D_y) over the neutral option C. We refer to this as the adversarial error ratio (AER). A high AER value indicates that the model tends to select the adversarial distractor.

Finally, we define the overall model performance, denoted as MVH Score, as the sum of all evaluation metrics except YER and AER.

3 Reference Shift Contrastive Decoding

3.1 MVH Problem Statement

Multi-view hallucination (MVH) is fundamentally distinct from conventional hallucination in that it arises when models ground their responses in visual information from non-target instances or viewpoints. As illustrated in Figure 3(a), conventional hallucination occurs when the model relies on prior knowledge (e.g., *the color of a banana is yellow*). In contrast, MVH occurs when the model is distracted by spurious visual cues (e.g., *the color of the banana on the plate is green*). Therefore, an approach that simply makes the model attend more to *all* visual cues would not work in most cases. A better way to address MVH is to guide the model

to attend to the *correct* visual evidence. We argue that understanding the full context of the query is essential for the model to identify the correct visual evidence. When the model understands the query only partially, it may attend to visual content that is consistent with only part of the query. To empirically examine this hypothesis, we analyze how query understanding emerges in LVLMS and also contributes to visual grounding.

3.2 Identifying Layers for Query Understanding

In LVLMS, system tokens, image tokens, and text tokens are concatenated into a single sequence of length T and processed by a Transformer decoder. At each decoder layer $l \in \{0, \dots, L-1\}$, these tokens are updated through self-attention, with the attention matrix $A \in \mathbb{R}^{T \times T}$ computed as follows:

$$A = \text{softmax}(QK^\top + M^c), \quad (1)$$

where Q and K are the query and key matrices, respectively. The causal mask M^c is defined as

$$M_{i,j}^c = \begin{cases} -\infty, & j > i, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

The attention matrix A determines how token representations are updated, with $A_{i,j}$ representing the

LVLMs	Cross-Instance						Cross-View						MVH Score
	Binary			M.C.		Score	Binary			M.C.		Score	
	Acc	p-Acc	q-Acc	Acc	p-Acc		Acc	p-Acc	q-Acc	Acc	p-Acc		
Closed-Source													
GPT-4o	73.66	52.41	32.04	57.27	36.11	251.49 ± 0.32	68.87	46.48	23.43	55.74	32.31	226.83 ± 0.16	478.32 ± 0.47
GPT-4o mini	62.20	31.53	12.87	66.99	44.81	218.40 ± 0.77	61.92	35.83	14.72	61.06	33.06	206.59 ± 0.72	424.99 ± 0.67
Gemini 2.5 Flash	64.42	39.68	14.91	60.83	38.15	217.99 ± 2.34	71.00	51.16	28.06	69.54	48.89	268.65 ± 2.47	486.64 ± 2.90
Gemini 2.0 Flash	60.72	29.58	9.63	68.15	44.72	212.80 ± 3.18	69.91	44.54	23.24	73.06	53.80	264.55 ± 2.49	477.35 ± 5.66
Claude Sonnet 4.5	61.44	36.94	13.61	60.00	36.02	208.01 ± 0.87	69.26	49.35	23.61	72.03	49.42	263.67 ± 3.20	471.68 ± 4.05
Open-Source													
InternVL3-14B	64.68	40.69	18.33	68.47	47.59	239.76 ± 3.70	55.37	29.03	4.48	56.90	26.11	171.89 ± 6.53	411.65 ± 9.67
Gemma3-12B	55.76	18.52	3.89	59.21	33.24	170.62 ± 5.12	58.35	24.91	5.09	58.89	30.37	177.61 ± 0.94	348.23 ± 5.98
Llama-3.2-11B-V-I	54.60	20.74	4.81	59.72	33.52	173.39 ± 10.38	49.38	18.01	1.94	47.41	16.11	132.85 ± 4.80	306.24 ± 5.70
InternVL3-8B	64.21	38.43	16.76	69.58	48.43	237.41 ± 1.58	50.86	22.69	1.76	50.93	15.74	141.98 ± 0.80	379.39 ± 2.38
Qwen2.5-VL-7B	64.10	41.30	16.29	66.57	43.52	231.78 ± 6.48	67.45	47.45	21.95	70.70	50.56	258.11 ± 4.05	489.89 ± 9.85
LLaVA-OneVision-7B	58.36	32.36	10.93	62.78	36.29	200.72 ± 5.99	52.36	27.17	4.26	59.26	29.72	172.77 ± 6.18	373.49 ± 6.16
Mantis-8B-Idetics2	55.21	28.29	7.96	56.94	28.98	177.38 ± 4.11	58.94	35.05	11.76	58.93	33.70	198.38 ± 5.24	375.76 ± 7.04
Deepseek-VL-7B	58.17	29.30	7.22	60.69	27.96	183.35 ± 3.36	50.28	20.14	0.93	49.31	8.24	128.89 ± 3.58	312.24 ± 2.29
Qwen2-VL-7B	61.07	35.18	12.87	61.99	37.87	208.98 ± 8.73	61.97	37.92	14.26	58.71	31.30	204.16 ± 2.05	413.14 ± 10.21
LLaVA-NeXT-I-7B	56.85	30.05	8.34	60.14	33.05	188.43 ± 7.60	61.02	37.31	15.09	62.69	38.33	214.44 ± 9.44	402.87 ± 15.00

Table 1: Performance comparison of recent closed and open-source LVLMs on MVH-Bench. The best and second-best results are highlighted in **bold** and underline, respectively. All results are averaged over three independent runs, with standard deviations reported after $\pm$.

amount of information flow from the j -th token to the i -th token (Zhang et al., 2025). Since self-attention is the only mechanism for information exchange across token positions, query understanding likely emerges through interactions among text tokens within self-attention. Under this view, blocking attention between text tokens may hinder the model’s ability to understand the query. To examine this effect, we employ an additional mask M^{t2t} that blocks attention from text tokens to text tokens:

$$M_{i,j}^{t2t} = \begin{cases} -\infty, & i, j \in \mathcal{T}, \\ 0, & \text{otherwise,} \end{cases} \quad (3)$$

where $\mathcal{T}$ denotes the set of text-token indices. We apply this mask over a sliding window of w consecutive decoder layers to analyze how the effect of blocking text-to-text information flow varies across layer ranges. Concretely, we perform a captioning task on image pairs, asking the model to describe either the first or the second image in each pair (see Appendix B.2 for details). To quantify the resulting response changes, we define *reference accuracy*, which considers a response correct if the generated caption is grounded in the corresponding visual evidence. As shown in Figure 4, reference accuracy drops when the mask is applied to middle layers (e.g., layers 13–20 in LLaVA-OneVision). This suggests that blocking text-to-text attention in these layers may lead the model to reference the incorrect

LVLMs	Cross-Instance		Cross-View	
	YER	AER	YER	AER
Closed-Source				
GPT-4o	70.09	22.10	73.51	25.42
Gemini 2.5 Flash	52.14	66.67	41.47	73.71
Claude Sonnet 4.5	39.05	64.58	35.04	76.50
Open-Source				
InternVL3-14B	74.64	73.26	68.72	82.60
Gemma3-12B	86.54	75.25	76.11	72.51
Qwen2.5-VL-7B	45.50	86.98	54.80	94.13
LLaVA-OneVision-7B	74.02	97.88	70.18	99.24
LLaVA-NeXT-I-7B	71.54	97.50	68.62	98.55

Table 2: Comparison of bias evaluation metrics for recent closed and open-source LVLMs on MVH-Bench.

regions of the image. Based on this observation, we identify this layer range, denoted by Λ^* , as important for query understanding and visual grounding. We provide additional analysis in Appendix D that further supports our hypothesis.

3.3 Contrastive Decoding via Reference Shift

To mitigate MVH, we deliberately amplify the effects of incomplete query understanding and subtract the resulting bias from the original logits. Specifically, for each layer $l \in \Lambda^*$, we mask the top $\rho\%$ of attention entries in each row corresponding to a text token $i \in \mathcal{T}$. Here, ρ controls how

LVLMs	Cross-Instance						Cross-View						MVH Score
	Binary			M.C.		Score	Binary			M.C.		Score	
	Acc	p-Acc	q-Acc	Acc	p-Acc		Acc	p-Acc	q-Acc	Acc	p-Acc		
LLaVA-OneVision-7B													
Base model	58.36	32.36	10.93	62.78	36.29	200.72 $\pm$ 5.99	52.36	27.17	4.26	59.26	29.72	172.77 $\pm$ 6.18	373.49 $\pm$ 6.16
VCD	<u>60.60</u>	32.59	<u>11.94</u>	68.29	43.98	217.40 $\pm$ 3.04	<u>55.07</u>	<u>29.31</u>	2.59	<u>61.62</u>	28.43	<u>177.02</u> $\pm$ 1.86	<u>394.42</u> $\pm$ 1.34
ICD	60.58	<u>33.19</u>	11.30	68.06	44.81	<u>217.94</u> $\pm$ 2.49	54.89	28.47	3.61	60.42	27.87	175.26 $\pm$ 4.30	393.20 $\pm$ 5.95
AvisC	60.46	29.31	10.28	<u>68.15</u>	<u>44.35</u>	212.54 $\pm$ 3.49	53.47	26.43	2.04	61.81	<u>28.61</u>	172.36 $\pm$ 1.86	384.90 $\pm$ 2.39
RSCD (Ours)	62.34	38.75	14.35	67.31	43.61	226.36 $\pm$ 4.43	56.74	33.06	5.65	60.79	27.96	184.19 $\pm$ 1.99	410.55 $\pm$ 6.41
Qwen2.5-VL-7B													
Base model	<u>64.10</u>	<u>41.30</u>	<u>16.29</u>	66.57	43.52	<u>231.78</u> $\pm$ 6.48	<u>67.45</u>	<u>47.45</u>	<u>21.95</u>	<u>70.70</u>	<u>50.56</u>	<u>258.11</u> $\pm$ 4.05	<u>489.89</u> $\pm$ 9.85
VCD	62.75	38.52	13.52	67.87	44.72	227.38 $\pm$ 23.72	66.04	45.51	19.63	70.60	46.48	248.27 $\pm$ 24.03	475.64 $\pm$ 47.62
ICD	61.69	32.13	10.46	65.55	42.04	211.88 $\pm$ 19.12	65.77	39.31	16.39	69.67	47.41	238.54 $\pm$ 26.04	450.42 $\pm$ 44.95
AvisC	61.81	36.43	11.39	<u>68.84</u>	<u>45.93</u>	224.39 $\pm$ 15.60	64.44	38.66	16.02	69.35	45.65	234.12 $\pm$ 17.49	458.51 $\pm$ 33.08
RSCD (Ours)	67.50	45.19	17.69	73.66	53.70	257.73 $\pm$ 2.90	72.87	56.02	31.02	77.36	58.33	295.60 $\pm$ 0.34	553.33 $\pm$ 2.73

Table 3: Performance of various methods on MVH-Bench using LLaVA-OneVision-7B and Qwen2.5-VL-7B. The best and second-best results are highlighted in **bold** and underline, respectively. All results are averaged over three independent runs, with standard deviations reported after $\pm$.

strongly textual context formation is disrupted. We define the resulting mask as the reference shift mask (RSM), denoted by M^{rs} :

$$M_{i,j}^{rs} = \begin{cases} -\infty, & i \in \mathcal{T} \text{ and } j \in \mathcal{J}_i, \\ 0, & \text{otherwise,} \end{cases} \quad (4)$$

where $\mathcal{J}_i = \{j \mid A_{i,j} \in \text{TopP}(A_{i,\mathcal{T}}, \rho)\}$ is the set of text-token indices that receive the top $\rho\%$ attention from token i within $A_{i,\mathcal{T}}$.

A forward pass with M^{rs} yields the negative logit $\text{logit}(y_{n+1} \mid y_{\leq n}, M^{rs})$, where $y_{\leq n}$ denotes the input sequence up to position n . Let $\text{logit}(y_{n+1} \mid y_{\leq n})$ be the base logit from the unmodified model. We then contrast the base and negative logit at each decoding step as follows:

$$\text{logit}^{rscd} = (1 + \alpha) \text{logit}(y_{n+1} \mid y_{\leq n}) - \alpha \text{logit}(y_{n+1} \mid y_{\leq n}, M^{rs}), \quad (5)$$

where α scales the amplification of the differences between the two logits. Following previous work, we apply adaptive plausibility constraints during decoding to consider only high-probability tokens under the original distribution (Li et al., 2023b). The overall decoding process of RSCD is illustrated in Figure 3(b).

4 Experimental Results

4.1 LVLMs' Performance on MVH-Bench

To assess multi-view hallucination in recent LVLMs, we evaluate five closed-source and ten




Figure 4: Layer-wise analysis of text-to-text attention blocking for LLaVA-OneVision-7B and Qwen2.5-VL on a captioning task. The yellow shaded region indicates the layer range selected by RSCD.

open-source LVLMs on MVH-Bench (see Table 1). All models achieve consistently low p-Acc and q-Acc across all categories, indicating that they struggle to resolve the MVH problem. Furthermore, the models that performed best differed between the Cross-Instance and Cross-View categories, suggesting that these two settings pose distinct challenges.

To quantify the extent to which models are affected by distracting visual cues in MVH-Bench, we report YER and AER in Table 2 (see Appendix C.3 for results on the remaining models). Most models exhibit high YER and AER scores,

indicating a strong tendency to answer “Yes” to binary questions and to select adversarial options in multiple-choice questions. Notably, open-source models show even higher AER than their closed-source counterparts. These results suggest that current LVLMs are highly susceptible to multi-view interference.

4.2 Overview of Baseline Methods

To evaluate the effectiveness of RSCD, we compare it with recent representative contrastive decoding-based hallucination mitigation approaches, namely VCD (Leng et al., 2024), ICD (Wang et al., 2024c), and AvisC (Woo et al., 2025). While all three follow the same idea of constructing a hallucination-prone *negative logits* distribution, they differ in how they obtain the negative logits. VCD deliberately corrupts the visual input with noise, weakening the model’s access to accurate visual signals and thereby inducing a visually misaligned distribution. ICD appends role prefixes to the instruction to form disturbance instructions, thereby amplifying hallucinations. AvisC constructs its negative distribution by conditioning only on image tokens that receive high attention yet are semantically uninformative or query-irrelevant.

4.3 Performance Evaluation of RSCD

We assess the effectiveness of RSCD by comparing it with recent hallucination mitigation methods, namely VCD, ICD, and AvisC, on LLaVA-OneVision-7B and Qwen2.5-VL-7B (see Table 3). RSCD consistently outperforms all competing methods on both models. In particular, it outperforms AvisC by 25.7 points on LLaVA-OneVision-7B and by 94.8 points on Qwen2.5-VL-7B. Moreover, improvements are observed across most metrics and categories. These results suggest that RSCD effectively mitigates MVH by encouraging LVLMs to focus on the queried instance or view.

5 Analysis

5.1 Probing the Causes of MVH

To further investigate whether MVH arises from visual interference, we modify the MVH-Bench validation set by replacing the distractors. Specifically, we replace the descriptor phrases in questions whose correct answer is “No” with alternatives that do not appear in the images. As shown in Figure 5, the MVH Score increases substantially, while both YER and AER decrease. This demonstrates that




Figure 5: Effect of replacing visual distractors with random/non-visual distractors on MVH-Bench performance.

the poor performance of current models on MVH-Bench is largely driven by visual distraction.

5.2 Analysis of Layer Range Selection

To examine the effectiveness of the selected layer range in RSCD, we analyze the results obtained by applying the mask to different layer ranges (see Figure 6(a)). Specifically, using Λ^* (the layer range 13–20 in LLaVA-OneVision) as a reference, we evaluate the performance of ranges that precede and follow Λ^* , as well as the full set of layers. RSCD achieves the best performance when masking is applied at Λ^* , and still shows slight improvement over the base model even when extended to later layers. In contrast, applying RSCD to early layers significantly degrades performance.

In addition, we analyze the sensitivity of RSCD to the selected layer range by varying the interval around Λ^* (see Figure 6(b)). Starting from Λ^* , we reduce or extend the range by one layer and measure the resulting MVH Score. The performance remains stable across these variations, indicating that RSCD is not highly sensitive to the exact layer boundaries.

5.3 Ablation Study on Hyperparameters

We analyze the effect of the two RSCD hyperparameters, α and ρ . As shown in Figure 6(c), RSCD consistently outperforms the baseline across a wide range of hyperparameter values. This indicates that RSCD is insensitive to the choice of hyperparameters. We also observe a clear trend: as the masking ratio ρ increases, performance gradually improves. This indicates that stronger perturbations to the text tokens produce negative logits that are further shifted away from the correct prediction. However, in the extreme case where all tokens are masked ($\rho = 1.0$), performance drops substantially. This shows that partially masking attention over the text




Figure 6: Analysis of RSCD across different layer ranges and hyperparameter settings.

Method	Time (sec./query)
Base model	0.81
VCD	1.38
ICD	1.39
AvisC	1.37
RSCD (Ours)	0.89

Table 4: Inference cost comparison between RSCD and other methods used in our experiments.

tokens is more effective than complete masking for inducing partially grounded responses.

5.4 Inference Cost of RSCD

Conventional contrastive decoding methods for mitigating visual hallucination generate negative logits by modifying the image itself, altering the preceding prompt, or masking a subset of visual tokens. Such approaches incur additional computational overhead due to input manipulation and the re-computation of intermediate visual representations across transformer layers. In contrast, RSCD produces negative logits without altering the image input or its token embeddings, thereby enabling direct reuse of cached image key-value states. Due to this feature, RSCD is particularly efficient in multi-view settings where the input consists of a substantial number of image tokens (see Table 4).

5.5 Qualitative Analysis of RSCD

To investigate the effect of the reference shift mask, we visualize the changes in the text-to-image attention maps before and after applying the mask (see Figure 7). Specifically, we focus on the attention from the token corresponding to the phrase “to




Figure 7: Qualitative comparison of attention maps with and without reference shift mask. Top: attention map for the original prediction. Bottom: attention map after applying the reference shift mask.

the right of the person”. To answer the question correctly, the model must verify whether “the sink” is actually to the right of the person. In the top row (base model), the token attends strongly to the sink region (solid yellow line). In contrast, in the bottom row (with reference shift mask), attention to the sink is suppressed and instead shifts toward the image region corresponding to the phrase “to the right of the person” (solid white line). This shift shows that our masking strategy perturbs the full query context, causing the model to rely on a partially interpreted phrase and resulting in an incorrect answer. For additional qualitative examples, please refer to Appendix E.6.

6 Conclusion

In this work, we identified multi-view hallucination as a key challenge for LVLMs and introduced MVH-Bench to systematically evaluate their robustness against cross-instance and cross-view interference. Experiments on recent LVLMs show that MVH remains a persistent challenge. To address this issue, we proposed RSCD, a simple yet effective training-free decoding method that suppresses predictions driven by incomplete query understanding. Extensive experiments demonstrate that RSCD consistently mitigates MVH and outperforms existing hallucination mitigation methods.

Limitations

Although MVH-Bench provides a systematic benchmark for evaluating multi-view hallucination, it currently focuses on paired-view settings and does not cover more complex configurations involving a larger number of viewpoints. While RSCD performs consistently across the five LVLMs evaluated in this work, its effectiveness on other LVLMs remains to be further investigated. Extending the benchmark to more diverse multi-view configurations and evaluating RSCD across a broader range of LVLMs would therefore be a promising direction for future work.

Acknowledgments

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) under the Graduate School of Artificial Intelligence Semiconductor (IITP-2025-RS-2023-00256081), and by the National Research Foundation of Korea (NRF) (2022M3C1A3099336), both funded by the Korea government (MSIT). The authors employed ChatGPT for language editing, including proofreading and stylistic refinement. All AI-assisted revisions were subsequently reviewed and validated by the authors.

References

- Wenbin An, Feng Tian, Sicong Leng, Jiahao Nie, Haonan Lin, Qianying Wang, Ping Chen, Xiaoqin Zhang, and Shijian Lu. 2025. Mitigating object hallucinations in large vision-language models with assembly of global and local attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 29915–29926.
- Anthropic. 2025. [System Card: Claude Sonnet 4.5](#).
- Anthropic. 2026. [System Card: Claude Sonnet 4.6](#).
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, and 1 others. 2025a. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025b. Qwen2.5-VL technical report. *arXiv preprint arXiv:2502.13923*.
- Florian Bordes, Richard Yuanzhe Pang, Anurag Ajay, Alexander C Li, Adrien Bardes, Suzanne Petryk, Oscar Mañas, Zhiqiu Lin, Anas Mahmoud, Bargav Jayaraman, and 1 others. 2024. An introduction to vision-language modeling. *arXiv preprint arXiv:2405.17247*.
- Yohann Cabon, Lucas Stofl, Leonid Antsfeld, Gabriela Csurka, Boris Chidlovskii, Jerome Revaud, and Vincent Leroy. 2025. Must3r: Multi-view network for stereo 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1050–1060.
- Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. 2017. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*.
- Wey Yeh Choong, Yangyang Guo, and Mohan Kankanhalli. 2024. VidHal: Benchmarking Temporal Hallucinations in Vision LLMs. *arXiv preprint arXiv:2411.16771*.
- Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Naveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. 2018. Embodied question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ailin Deng, Zhirui Chen, and Bryan Hooi. 2024. Seeing is believing: Mitigating hallucination in large vision-language models via clip-guided decoding. *arXiv preprint arXiv:2402.15300*.
- Hongcheng Gao, Jiashu Qu, Jingyi Tang, Baolong Bi, Yue Liu, Hongyu Chen, Li Liang, Li Su, and Qingming Huang. 2025. Exploring hallucination of large multimodal models in video understanding: Benchmark, analysis and mitigation. *arXiv preprint arXiv:2503.19622*.
- Kamran Gholizadeh HamAbadi, Monica Vahdati, Haiwei Dong, and Abdulmotaleb El Saddik. 2025. Ai-enhanced creation of digital twins from iphone lidar for immersive xr experiences in nvidia omniverse. In *Proceedings of the International Workshop on Intelligent Immersification in the Metaverse: AI-Driven Immersive Multimedia*, pages 25–33.
- Xuan Gong, Tianshi Ming, Xinpeng Wang, and Zhihua Wei. 2024. [DAMRO: Dive into the attention mechanism of LVLM to reduce object hallucination](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7696–7712, Miami, Florida, USA. Association for Computational Linguistics.

- Google DeepMind. 2025a. [Gemini 2.0 Flash Model Card](#).
- Google DeepMind. 2025b. [Gemini 3 Flash Model Card](#).
- Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, Eugene Byrne, Zach Chavis, Joya Chen, Feng Cheng, Fu-Jen Chu, Sean Crane, Avijit Dasgupta, Jing Dong, Maria Escobar, and 5 others. 2024. [Ego-exo4d: Understanding skilled human activity from first- and third-person perspectives](#). In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19383–19400.
- Yuliang Guo, Sparsh Garg, S Mahdi H Miangoleh, Xinyu Huang, and Liu Ren. 2025. Depth any camera: Zero-shot metric depth estimation from any camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26996–27006.
- Jinghan He, Kuan Zhu, Haiyun Guo, Junfeng Fang, Zhenglin Hua, Yuheng Jia, Ming Tang, Tat-Seng Chua, and Jinqiao Wang. 2025. [Cracking the code of hallucination in LVLMS with vision-aware head divergence](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3488–3501, Vienna, Austria. Association for Computational Linguistics.
- Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 2023. 3d-llm: Injecting the 3d world into large language models. *Advances in Neural Information Processing Systems*, 36:20482–20494.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Md Mofijul Islam, Alexi Gladstone, Riashat Islam, and Tariq Iqbal. 2024. [EQA-MX: Embodied question answering using multimodal expression](#). In *The Twelfth International Conference on Learning Representations*.
- Baoxiong Jia, Yixin Chen, Siyuan Huang, Yixin Zhu, and Song-Chun Zhu. 2020. Lemma: A multiview dataset for learning multi-agent multi-view activities. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Dongfu Jiang, Xuan He, Huaye Zeng, Cong Wei, Max Ku, Qian Liu, and Wenhui Chen. 2024. [Mantis: Interleaved multi-image instruction tuning](#). *Transactions on Machine Learning Research*.
- Zhangqi Jiang, Junkai Chen, Beier Zhu, Tingjin Luo, Yankun Shen, and Xu Yang. 2025. Devils in middle layers of large vision-language models: Interpreting, detecting and mitigating object hallucinations via attention lens. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25004–25014.
- Sihyeon Kim, Boryeong Cho, Sangmin Bae, Sumyeong Ahn, and Se-Young Yun. 2024. Vacode: Visual augmented contrastive decoding. *arXiv preprint arXiv:2408.05337*.
- Soohyun Kim, Gusang Lee, Kyuhong Shim, and Byonghyo Shim. 2026. [Mitigating object and relationship hallucination in large vision language model with multi-agent guidance](#). In *ICASSP 2026 - 2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4836–4840.
- Ming Kong, Xianzhou Zeng, Luyuan Chen, Yadong Li, Bo Yan, and Qiang Zhu. 2025. Mhbench: Demystifying motion hallucination in videollms. In *AAAI*, volume 39, pages 4401–4409.
- Insu Lee, Wooje Park, Jaeyun Jang, Minyoung Noh, Kyuhong Shim, and Byonghyo Shim. 2025. [Towards Comprehensive Scene Understanding: Integrating First and Third-Person Views for LVLMS](#). In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pages 18576–18628. Curran Associates, Inc.
- Insu Lee, Wooje Park, Wonseok Shin, Jinwoo Son, and Byonghyo Shim. 2026. Visual information-guided parallel decoding for diffusion multimodal large language models. *arXiv preprint arXiv:2608.26580*.
- Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. 2024. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13872–13882.
- Vincent Leroy, Yohann Cabon, and Jérôme Revaud. 2024. Grounding image matching in 3d with mast3r. In *European Conference on Computer Vision*, pages 71–91. Springer.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2025a. [LLaVA-OneVision: Easy Visual Task Transfer](#). *Transactions on Machine Learning Research*.
- Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. 2023a. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*.
- Chaoyu Li, Eun Woo Im, and Pooyan Fazli. 2025b. Vid-halluc: Evaluating temporal hallucinations in multimodal large language models for video understanding. In *CVPR*, pages 13723–13733.
- Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li.

2024. LLaVA-NeXT-Interleave: Tackling Multi-image, Video, and 3D in Large Multimodal Models. *arXiv preprint arXiv:2407.07895*.
- Jiale Li, Mingrui Wu, Zixiang Jin, Hao Chen, Jiayi Ji, Xiaoshuai Sun, Liujuan Cao, and Rongrong Ji. 2025c. Mihbench: Benchmarking and mitigating multi-image hallucinations in multimodal large language models. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 3143–3152.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. 2023b. [Contrastive decoding: Open-ended text generation as optimization](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12286–12312, Toronto, Canada. Association for Computational Linguistics.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Xin Zhao, and Ji-Rong Wen. 2023c. [Evaluating object hallucination in large vision-language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 292–305, Singapore. Association for Computational Linguistics.
- Zongxia Li, Xiyang Wu, Hongyang Du, Fuxiao Liu, Huy Nghiem, and Guangyao Shi. 2025d. A survey of state of the art large vision language models: Alignment, benchmark, evaluations and challenges. *arXiv preprint arXiv:2501.02189*.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 740–755. Springer.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024a. [LLaVA-NeXT: Improved reasoning, OCR, and world knowledge](#).
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. 2024b. [Mmbench: Is your multi-modal model an all-around player?](#) In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part VI*, page 216–233, Berlin, Heidelberg. Springer-Verlag.
- Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Hao Yang, Yaofeng Sun, Chengqi Deng, Hanwei Xu, Zhenda Xie, and Chong Ruan. 2024. DeepSeek-VL: Towards Real-World Vision-Language Understanding. *arXiv preprint arXiv:2403.05525*.
- Mi Luo, Zihui Xue, Alex Dimakis, and Kristen Grauman. 2024. Put myself in your shoes: Lifting the egocentric perspective from exocentric videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 407–425. Springer.
- Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mccvay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent Berges, Shiqi Zhang, Pulkit Agrawal, Yonatan Bisk, Dhruv Batra, and 5 others. 2024. [OpenEQA: Embodied Question Answering in the Era of Foundation Models](#). In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16488–16498.
- Meta AI. 2024. [Llama 3.2-11B-Vision-Instruct](#).
- Lorenzo Mur-Labadia, Maria Santos-Villafranca, Jesus Bermudez-Cameo, Alejandro Perez-Yus, Ruben Martinez-Cantin, and Jose J Guerrero. 2025. Omama: Learning object mask matching between egocentric and exocentric views. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6892–6903.
- OpenAI. 2026a. [GPT-5.4 Thinking System card](#).
- OpenAI. 2026b. [GPT-5.5 System Card](#).
- Hansol Park, Hoseong Ahn, Junwon Moon, Yejin Lee, and Kyuhong Shim. 2025a. Evaluating hallucinations in multimodal llms with spoken queries under diverse acoustic conditions. *arXiv preprint arXiv:2510.08581*.
- Jungin Park, Jiyoung Lee, and Kwanghoon Sohn. 2025b. Bootstrap your own views: Masked ego-exo modeling for fine-grained view-invariant video representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13661–13670.
- Wooje Park, Insu Lee, Minyoung Noh, Jaeyun Jang, Sungmin Lee, Kyuhong Shim, and Byonghyo Shim. 2026. Dependency-aware revocable decoding for efficient diffusion large language model inference. *arXiv preprint arXiv:2608.26574*.
- Luigi Piccinelli, Christos Sakaridis, Yung-Hsu Yang, Mattia Segu, Siyuan Li, Wim Abbeloos, and Luc Van Gool. 2025. Unidepthv2: Universal monocular metric depth estimation made simpler. *arXiv preprint arXiv:2502.20110*.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. [Object hallucination in image captioning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4035–4045, Brussels, Belgium. Association for Computational Linguistics.
- Gunnar A Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, and Karteek Alahari. 2018. Charades-ego: A large-scale dataset of paired third and first person videos. *arXiv preprint arXiv:1804.09626*.

- Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. 2012. Indoor segmentation and support inference from rgbd images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 746–760. Springer.
- 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. 2024. [Aligning large multimodal models with factually augmented RLHF](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 13088–13110, Bangkok, Thailand. Association for Computational Linguistics.
- Lexiang Tang, Xianwei Zhuang, Bang Yang, Zhiyuan Hu, Hongxiang Li, Lu Ma, Jinghan Ru, and Yuexian Zou. 2025. Not all tokens and heads are equally important: Dual-level attention intervention for hallucination mitigation. *arXiv preprint arXiv:2506.12609*.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivi re, and 1 others. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Junyang Wang, Yuhang Wang, Guohai Xu, Jing Zhang, Yukai Gu, Haitao Jia, Jiaqi Wang, Haiyang Xu, Ming Yan, Ji Zhang, and Jitao Sang. 2024a. [Amber: An llm-free multi-dimensional benchmark for mllms hallucination evaluation](#). *Preprint*, arXiv:2311.07397.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, and 1 others. 2024b. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Qianqian Wang, Yifei Zhang, Aleksander Holynski, Alexei A Efros, and Angjoo Kanazawa. 2025a. Continuous 3d perception model with persistent state. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10510–10522.
- Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, and 1 others. 2025b. Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*.
- Xintong Wang, Jingheng Pan, Liang Ding, and Chris Biemann. 2024c. [Mitigating hallucinations in large vision-language models with instruction contrastive decoding](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 15840–15853, Bangkok, Thailand. Association for Computational Linguistics.
- Sangmin Woo, Donguk Kim, Jaehyuk Jang, Yubin Choi, and Changick Kim. 2025. Don’t miss the forest for the trees: Attentional vision calibration for large vision language models. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 1927–1951.
- Mingrui Wu, Jiayi Ji, Oucheng Huang, Jiale Li, Yuhang Wu, Xiaoshuai Sun, and Rongrong Ji. 2024. [Evaluating and analyzing relationship hallucinations in large vision-language models](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 53553–53570. PMLR.
- Tsung-Han Patrick Wu, Heekyung Lee, Jiaxin Ge, Joseph Gonzalez, Trevor Darrell, and David Chan. 2026. Generate, but verify: Reducing hallucination in vision-language models with retrospective resampling. *Advances in Neural Information Processing Systems*, 38:65749–65777.
- Jianing Yang, Alexander Sax, Kevin J Liang, Mikael Henaff, Hao Tang, Ang Cao, Joyce Chai, Franziska Meier, and Matt Feiszli. 2025. Fast3r: Towards 3d reconstruction of 1000+ images in one forward pass. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21924–21935.
- Yijun Yang, Tianyi Zhou, Kanxue Li, Dapeng Tao, Lu-song Li, Li Shen, Xiaodong He, Jing Jiang, and Yuhui Shi. 2024. Embodied multi-modal agent trained by an llm from a parallel textworld. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26275–26285.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Chun-Hsiao Yeh, Chenyu Wang, Shengbang Tong, Ta-Ying Cheng, Ruoyu Wang, Tianzhe Chu, Yuexiang Zhai, Yubei Chen, Shenghua Gao, and Yi Ma. 2025. Seeing from another perspective: Evaluating multi-view understanding in mllms. *arXiv preprint arXiv:2504.15280*.
- Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nie ner, and Angela Dai. 2023. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–22.
- Hao Yin, Guangzong Si, and Zilei Wang. 2025. Clear-sight: Visual signal enhancement for object hallucination mitigation in multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14625–14634.
- Tongtong Yuan, Xuange Zhang, Bo Liu, Kun Liu, Jian Jin, and Zhenzhen Jiao. 2024. Surveillance video-and-language understanding: from small to large multimodal models. *IEEE Transactions on Circuits and Systems for Video Technology*.

- Zhi Zhang, Srishti Yadav, Fengze Han, and Ekaterina Shutova. 2025. Cross-modal information flow in multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19781–19791.
- Kening Zheng, Junkai Chen, Yibo Yan, Xin Zou, Huiyu Zhou, and Xuming Hu. 2025. Reefknot: A comprehensive benchmark for relation hallucination evaluation, analysis and mitigation in multimodal large language models. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 6193–6212.
- Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, and 1 others. 2025. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*.

A Related Work

A.1 Multi-View Datasets and Tasks

Multi-view images capture complementary visual cues that provide richer information than a single view, leading to a more complete understanding of the scene. This advantage has motivated the development of diverse multi-view datasets. Among these, several works have introduced datasets that capture the spatial layout of indoor environments from diverse viewpoints (Silberman et al., 2012; Chang et al., 2017; Yeshwanth et al., 2023).

In parallel, other works have introduced datasets that capture real-world human activities and interactions from first and third-person perspectives (Sigurdsson et al., 2018; Jia et al., 2020; Grauman et al., 2024). These datasets have been widely used to support various vision tasks, including 3D reconstruction (Yang et al., 2025; Cabon et al., 2025; Wang et al., 2025a), depth estimation (Piccinelli et al., 2025; Guo et al., 2025), correspondence matching (Leroy et al., 2024; Mur-Labadia et al., 2025), and viewpoint transformation (Luo et al., 2024).

In addition to these applications, these datasets have also been explored in LVLM research, where multi-view information is leveraged to generate responses grounded in extended environmental context and to support human-AI interaction. One line of research examines whether LVLMs can fully interpret the 3D structure and spatial relationships of environments from multi-view data (Hong et al., 2023; Yeh et al., 2025). Another line of work investigates whether LVLMs can function as visual assistants capable of holistic scene understanding (Das et al., 2018; Islam et al., 2024; Lee et al., 2025). Despite rapid progress in multi-view understanding, the problem of multi-view hallucination largely remains unaddressed.

A.2 Hallucination in LVLMs

Hallucination in LVLMs refers to the generation of responses that are inconsistent with the underlying visual evidence. To quantitatively assess the susceptibility and robustness of LVLMs to hallucination, several benchmarks have been proposed. Some benchmarks focus on object-level hallucinations by evaluating whether mentioned objects are actually present in the image (Rohrbach et al., 2018; Li et al., 2023c), while others extend this perspective to hallucinations about object attributes and inter-object relationships (Wang et al., 2024a; Sun et al., 2024; Zheng et al., 2025; Wu et al., 2024).

Beyond single-image settings, MIHBench (Li et al., 2025c) evaluates a model’s ability to aggregate visual information across multiple images to form a unified judgment. More recent works have also considered video settings, introducing datasets that capture dynamic and temporal hallucinations in LVLMs (Choong et al., 2024; Li et al., 2025b; Kong et al., 2025; Gao et al., 2025). For example, MH-Bench (Kong et al., 2025) and VidHalluc (Li et al., 2025b) assess the robustness of reasoning over semantically similar actions or the temporal order of frames.

Building on these benchmarks, recent mitigation approaches encourage LVLMs to rely more on the visual input when generating responses. To this end, recent methods aim to directly strengthen LVLMs’ reliance on visual input. These approaches steer more attention toward image tokens during inference (Yin et al., 2025; He et al., 2025; Tang et al., 2025; Jiang et al., 2025), leverage auxiliary signals to highlight informative regions of the visual input (Li et al., 2025b), or refine logits with image-conditioned signals (An et al., 2025; Deng et al., 2024). Other methods mitigate hallucinations by explicitly inducing and then subtracting hallucination-driving factors during decoding. For example, some methods deliberately amplify language priors or hallucination-inducing visual cues and then remove their contribution at inference time (Leng et al., 2024; Wang et al., 2024c; Woo et al., 2025; Kong et al., 2025; Gong et al., 2024; Kim et al., 2024; Park et al., 2025a). Some approaches instead retrospectively verify generated responses and regenerate them when potential hallucinations are detected (Wu et al., 2026; Park et al., 2026). These benchmarks and methods have advanced the analysis and mitigation of hallucinations in LVLMs, but to the best of our knowledge, hallucinations in multi-view settings remain largely unexplored.

B Experimental Details

B.1 LVLMs and Experimental Setup

We evaluate both closed-source and open-source LVLMs in our experiments. For closed-source models, we include GPT-5.5 (OpenAI, 2026b), GPT-5.4 (OpenAI, 2026a), GPT-5.4-mini (OpenAI, 2026a), GPT-4o (Hurst et al., 2024), GPT-4o-mini (Hurst et al., 2024), Gemini-3-Flash (Google DeepMind, 2025b), Gemini 2.5 Flash (Comanici et al., 2025), Gemini 2.0 Flash (Google DeepMind,

LVL	Vision Encoder	LLM Backbone
InternVL3-14B (Zhu et al., 2025)	InternViT-300M	Qwen2.5-14B
Gemma3-12B (Team et al., 2025)	SigLIP-400M	Gemma3-12B
Llama-3.2-11B-V-I (Meta AI, 2024)	-	Llama-3.1-8B
InternVL3-8B (Zhu et al., 2025)	InternViT-300M	Qwen2.5-7B
Qwen2.5-VL-7B (Bai et al., 2025b)	ViT (customized)	Qwen2.5-7B
LLaVA-OneVision-7B (Li et al., 2025a)	SigLIP-SO	Qwen2-7B
Mantis-8B-Idetics2 (Jiang et al., 2024)	SigLIP	Mistral-7B-v0.1
Deepseek-VL-7B (Lu et al., 2024)	SigLIP-L, SAM-B	DeepSeek-LLM-7B
Qwen2-VL-7B (Wang et al., 2024b)	ViT-L	Qwen2-7B
LLaVA-NeXT-I-7B (Li et al., 2024)	SigLIP-SO	Qwen1.5-7B
InternVL3.5-14B (Wang et al., 2025b)	InternViT-300M	Qwen3-14B
InternVL3.5-8B (Wang et al., 2025b)	InternViT-300M	Qwen3-8B
Qwen3-VL-8B (Bai et al., 2025a)	SigLIP2-SO-400M	Qwen3-8B

Table 5: Comparison of open-source LVLs across vision encoders and LLM architectures.

2025a), Claude Sonnet 4.6 (Anthropic, 2026), and Claude Sonnet 4.5 (Anthropic, 2025). For open-source LVLs, we evaluate the models summarized in Table 5, which provides an overview of their architectures, including vision encoders and LLM backbones. By evaluating a diverse set of vision-language architectures, we can thoroughly assess the robustness of recent LVLs and their ability to handle multi-view hallucinations. Experiments conducted with three independent runs are reported with the mean and standard deviation, while those conducted with a single run are reported as a single result. All evaluations use each model’s default generation settings. For the RSCD experiments on LLaVA-OneVision-7B and Qwen2.5-VL-7B, we set $(\alpha, \rho, \Lambda^*)$ to $(1.0, 0.7, \{13, \dots, 20\})$ and $(1.5, 0.8, \{12, \dots, 20\})$ respectively. For Qwen3-VL-8B, InternVL3.5-14B, and Llama-3.2-11B-Vision-Instruct, we set $(\alpha, \rho, \Lambda^*)$ to $(1.0, 0.8, \{11, \dots, 25\})$, $(1.0, 0.8, \{13, \dots, 29\})$, and $(1.5, 0.9, \{5, \dots, 23\})$, respectively. Unless otherwise specified, all analyses are conducted with LLaVA-OneVision-7B and run on NVIDIA RTX A6000 GPUs.

B.2 Details of the Captioning Analysis

In this section, we provide additional details on the captioning task described in Section 3.2. Here, we adopt a captioning-based analysis to explicitly examine which image content the model relies on when generating its response. Specifically, we construct image pairs using samples from the COCO dataset (Lin et al., 2014) and generate captions using the prompt “Provide a one-sentence caption for Image 1/2”. We primarily use COCO images for this analysis because the high degree of visual

similarity across views in MVH-Bench makes it difficult to clearly attribute the generated response to a specific visual input. When blocking text-to-text attention in intermediate layers, we observe that the model often generates captions describing the non-target image or produces sentences that mix information from both images. These behaviors indicate that disrupting the formation of textual context can cause the model to rely on irrelevant visual information.

In the main paper, we introduce a metric called *reference accuracy* that assesses whether the generated caption is well grounded in the target image. To compute this metric, we first obtain reference captions by prompting the model with only a single image. We then compare the generated caption with the reference captions using CLIP embedding similarity and consider the prediction correct if it is more similar to the reference caption of the target image.

C Additional Experiments on MVH-Bench

C.1 Evaluation of Recently Released LVLs on MVH-Bench

We additionally evaluate recently released LVLs on MVH-Bench, with the results presented in Table 6. While newer models generally achieve higher MVH Scores than the models originally evaluated in the main paper, they still struggle with MVH. These results indicate that MVH remains a challenge for recent LVLs and also underscore the importance of MVH-Bench as a diagnostic benchmark.

LVLM	Cross-Instance						Cross-View						MVH Score
	Binary			M.C.		Score	Binary			M.C.		Score	
	Acc	p-Acc	q-Acc	Acc	p-Acc		Acc	p-Acc	q-Acc	Acc	p-Acc		
Closed-Source													
GPT-5.5	69.10	45.83	23.06	73.89	55.56	<u>267.43</u>	76.60	59.17	37.50	82.78	68.33	324.38	591.81
GPT-5.4	<u>70.00</u>	44.86	24.72	75.28	57.78	272.64	73.33	51.81	31.11	<u>80.69</u>	<u>64.44</u>	<u>301.39</u>	<u>574.03</u>
GPT-5.4 mini	62.36	37.92	14.44	64.72	42.50	221.94	64.03	40.69	13.61	65.28	38.33	221.94	443.89
Gemini 3.0 Flash	75.14	<u>57.92</u>	<u>35.56</u>	47.50	26.67	242.78	81.53	<u>69.58</u>	<u>49.17</u>	58.33	39.72	298.33	541.11
Claude Sonnet 4.6	68.96	59.44	38.61	50.56	39.17	256.74	<u>77.08</u>	70.00	53.06	48.89	38.61	287.64	544.38
Open-Source													
InternVL3.5-14B	62.36	34.17	12.78	75.69	<u>58.89</u>	243.89	<u>70.28</u>	47.08	25.56	80.83	66.11	<u>289.86</u>	533.75
InternVL3.5-8B	<u>63.75</u>	<u>38.89</u>	<u>19.17</u>	73.33	55.28	<u>250.42</u>	71.94	51.67	<u>27.50</u>	<u>78.75</u>	<u>61.39</u>	291.25	<u>541.67</u>
Qwen3-VL-8B	71.46	53.19	28.33	<u>75.69</u>	59.44	288.12	70.76	<u>52.92</u>	29.17	76.81	58.33	287.99	576.11

Table 6: Performance comparison of recently released LVLMs on MVH-Bench. The best and second-best results are highlighted in **bold** and underline, respectively.

LVLM	POPE Acc. ↑	CHAIRs ↓	CHAIRi ↓	MVH Score ↑
Qwen2.5-VL-7B	86.77	47.40	13.12	489.89
LLaVA-OV-7B	87.38	43.20	11.07	373.49
InternVL3-8B	89.63	31.20	8.80	379.39

Table 7: Comparison of model performance on conventional hallucination benchmarks and MVH-Bench.

LVLMs	Cross-Instance		Cross-View	
	YER	AER	YER	AER
Closed-Source				
GPT-4o mini	82.52	68.72	67.85	71.34
Gemini 2.0 Flash	72.01	61.19	73.57	64.43
Open-Source				
Llama-3.2-11B-V-I	85.94	85.95	79.75	94.51
InternVL3-8B	76.66	81.58	71.90	88.58
Mantis-8B-Idefics2	62.50	71.55	61.89	77.73
Deepseek-VL-7B	66.46	88.47	74.75	95.86
Qwen2-VL-7B	66.84	91.48	61.03	95.37

Table 8: Comparison of bias evaluation metrics for recent closed and open-source LVLMs on MVH-Bench.

C.2 Comparison with Conventional Hallucination Benchmarks

To examine the relationship between conventional hallucination and MVH, we compare model performance across standard hallucination benchmarks (POPE (Li et al., 2023c), CHAIR (Rohrbach et al., 2018)) and MVH-Bench. As shown in Table 7, we observe that the model rankings are not consistent across these benchmarks. For example, InternVL3-8B achieves the best performance on conventional hallucination metrics (e.g., highest POPE accuracy and lowest CHAIR scores), but

does not achieve the best performance on MVH-Bench. Conversely, Qwen2.5-VL-7B achieves the highest MVH Score, despite showing relatively lower performance among the compared models on both POPE and CHAIR. These discrepancies indicate that lower hallucination rates on conventional benchmarks do not directly translate to better MVH performance, suggesting that MVH captures a distinct failure mode requiring dedicated evaluation.

C.3 Additional Results on Bias Evaluation

We additionally report YER and AER for the remaining models not covered in the main text (see Table 8).

C.4 Effect of Query Paraphrasing

To examine the effect of query paraphrasing on MVH-Bench, we paraphrase the validation-set questions using GPT-4o and evaluate model performance. As shown in Table 9, paraphrasing affects model performance, but the direction and magnitude of the change vary across models. Nevertheless, MVH is consistently observed under both the original and paraphrased queries, suggesting that the phenomenon is not specific to a particular wording of the benchmark questions.

Model	Query Type	Cross-Instance Score	Cross-View Score	MVH Score
LLaVA-OV-7B	Original	216.25	156.88	373.12
	Paraphrase	211.25	173.12	384.38
Qwen2.5-VL-7B	Original	243.12	286.25	529.38
	Paraphrase	225.00	248.12	473.12

Table 9: Performance of LVLMs on original and paraphrased queries from the MVH-Bench validation set.




Figure 8: Layer-wise analysis of text-to-text attention blocking for Qwen3-VL-8B, InternVL3.5-14B, and Llama-3.2-11B-V-I on the captioning task. The yellow shaded region indicates the layer range selected by RSCD.

D Further Analysis of Text-to-Text Attention Masking

D.1 Text-to-Text Attention Masking in a Text-Only Setting

Our text-to-text attention masking analysis on captioning tasks suggests that text-to-text attention plays an important role in visual grounding. However, this observation alone may not be sufficient to conclude that the effect is caused by incomplete textual understanding, since textual interpretation and visual grounding are tightly coupled in multimodal question answering.

To further substantiate our findings, we conduct an analysis in a text-only setting to isolate textual reasoning from multimodal interactions. Specifically, we perform a layer-wise text-to-text attention blocking experiment with LLaVA-OV-7B on HotpotQA (Yang et al., 2018), a benchmark that requires multi-hop reasoning over evidence distributed across multiple paragraphs. We evaluate the quality of the predicted answer and supporting facts generated by the model, using exact match (EM) for answer accuracy and the F1 score (Sup-F1) for supporting facts. Notably, as shown in Table 10, Sup-F1 drops more sharply than EM and remains low in the middle layer range (13–20). This pattern is consistent with the captioning analysis presented in the main paper. The discrepancy between EM and Sup-F1 suggests that, even when producing the same answer, the model may rely on incorrect yet superficially relevant evidence. This reflects a reduced ability to distinguish gold sup-

Center Layer	EM	Sup-F1
1	n/a	n/a
3	29.50	n/a
5	34.50	16.24
7	35.00	12.39
9	33.50	11.77
11	26.50	8.62
13	27.00	1.69
15	23.00	1.79
17	17.50	0.75
19	27.50	2.78
21	32.00	7.27
23	25.00	15.46
25	38.00	28.86
27	34.50	27.08

Table 10: Effect of layer-wise text-to-text attention masking on HotpotQA performance.

porting facts from partially related cues. Taken together, these results provide empirical support that masking text-to-text attention disrupts complete textual understanding and causes the model to rely more on partial matches.

D.2 Text-to-Text Attention Masking across Architectures and Model Scales

To demonstrate the generality of our findings, we extend the layer-wise text-to-text masking analysis to three additional models¹: Qwen3-VL-8B-Instruct, InternVL3.5-14B, and Llama-3.2-11B-Vision-Instruct (see Figure 8). Despite differences in model scale, training pipeline, and multimodal

¹Llama-3.2-11B-V-I integrates visual and textual modalities through dedicated cross-attention layers, whereas Qwen3-VL and InternVL3.5 concatenate visual tokens with text tokens and perform joint processing via self-attention.

Center Layer	Ref. Acc.
1	n/a
3	n/a
5	0.69
7	0.70
9	0.72
11	0.74
13	0.65
15	0.51
17	0.54
19	0.59
21	0.69
23	0.69
25	0.69
27	0.69

Table 11: Reference accuracy of LLaVA-OV-7B under layer-wise text-to-text attention blocking on the MVH-Bench validation subset.

integration strategy, all three models exhibit a consistent U-shaped pattern in reference accuracy across the middle layers. This consistency, together with the results in the main paper, suggests that the observed behavior is not a heuristic assumption, but an empirically grounded phenomenon.

D.3 Text-to-Text Attention Masking on MVH-Bench

We additionally perform the same layer-wise text-to-text attention blocking analysis on the MVH-Bench validation subset. In MVH-Bench, the two views often contain highly overlapping visual information, such that reference captions generated from each view can include similar objects or descriptions. Moreover, even without blocking text-to-text attention, the model sometimes generates hallucinated descriptions of the non-queried view, making it difficult to isolate the effect of text-to-text masking. Despite these limitations, we observe the same trend: while the overall reference accuracy is lower, it drops clearly in the middle layers, particularly around center layers 13–20 (see Table 11).

E Additional Evaluation and Analysis of RSCD

E.1 Generalization Across Model Families and Scales

To further validate the generalizability of RSCD, we evaluate it across recent LVLMs with diverse model families, architectures, and scales (Qwen3-VL-8B, InternVL3.5-14B, and Llama-3.2-11B-Vision-Instruct). As shown in Table 12, we observe that RSCD consistently outperforms competing

methods across all evaluated models, demonstrating its effectiveness and generalizability.

E.2 Generalization Across Other Multimodal Benchmarks

To assess the generalization capability of RSCD, we additionally evaluate it on widely used multimodal benchmarks, including MMBench (Liu et al., 2024b), SEED-Bench² (Li et al., 2023a), POPE, and CHAIR, using LLaVA-OV-7B. Note that we use the same hyperparameters for RSCD as reported in the manuscript across all benchmarks. We do not perform any task-specific tuning.

As shown in Tables 13 to 15, RSCD improves performance on most benchmarks, demonstrating its effectiveness across diverse multimodal evaluation settings. The only exception is CHAIRs, where RSCD shows a slight performance degradation. RSCD generates captions with more object mentions per caption (from 8.78 to 10.58 object mentions per caption). Since CHAIRs is a sentence-level metric that marks a caption as incorrect if it contains any hallucinated object, mentioning more objects increases the chance that a caption is counted as hallucinated. Notably, the improved CHAIRi and Recall indicate that RSCD reduces the proportion of hallucinated objects while capturing more visual content.

E.3 Evaluation Beyond Two-View Settings

To examine MVH in settings with more than two views, we conduct an additional evaluation on the MVH-Bench validation set by using all available synchronized multi-view images from the source dataset. We denote this setting as the max-view setting, where the average number of input images is 4.3.

As shown in Table 16, compared to the original two-view setting, we observe performance degradation under the max-view setting in both cross-instance and cross-view categories, indicating that MVH becomes more challenging as the number of views increases. Using the same hyperparameters as those used in the main paper, RSCD achieves the best overall performance among decoding-based methods, demonstrating its robustness as the number of input views increases.

²Due to its large scale, we randomly sample 50 examples from each of the 12 categories, including 3 video categories.

LVLM	Cross-Instance						Cross-View						MVH Score
	Binary			M.C.		Score	Binary			M.C.		Score	
	Acc	p-Acc	q-Acc	Acc	p-Acc		Acc	p-Acc	q-Acc	Acc	p-Acc		
Qwen3-VL-8B													
Base model	71.46	53.19	28.33	75.69	59.44	288.12	70.76	52.92	29.17	<u>76.81</u>	<u>58.33</u>	287.99	576.11
VCD	72.01	51.53	28.06	77.36	61.67	<u>290.62</u>	<u>72.50</u>	55.28	<u>31.11</u>	78.33	59.44	296.67	<u>587.29</u>
ICD	<u>72.36</u>	<u>53.75</u>	<u>31.11</u>	74.03	57.78	289.03	70.97	52.08	29.17	74.72	56.11	283.06	572.08
RSCD (Ours)	73.47	56.39	33.33	<u>76.11</u>	<u>60.00</u>	299.31	73.68	<u>54.17</u>	32.78	75.69	57.78	<u>294.10</u>	593.40
InternVL3.5-14B													
Base model	62.36	34.17	12.78	75.69	58.89	243.89	70.28	47.08	25.56	80.83	<u>66.11</u>	289.86	533.75
VCD	<u>65.07</u>	<u>39.31</u>	<u>16.94</u>	<u>74.44</u>	<u>56.67</u>	<u>252.43</u>	<u>72.01</u>	<u>51.39</u>	<u>29.17</u>	80.83	65.00	<u>298.40</u>	<u>550.83</u>
ICD	62.99	34.72	14.72	74.17	55.83	242.43	68.06	44.72	20.56	80.00	64.17	277.50	519.93
RSCD (Ours)	67.85	46.67	23.33	73.89	55.00	266.74	73.75	54.86	32.50	<u>80.56</u>	67.50	309.17	575.90
Llama-3.2-11B-V-I													
Base model	54.72	20.83	5.00	56.39	29.44	166.39	50.83	<u>18.47</u>	3.33	45.28	13.89	<u>131.81</u>	298.19
VCD	57.92	<u>24.17</u>	<u>5.28</u>	<u>64.44</u>	37.50	<u>189.31</u>	<u>50.76</u>	17.08	0.83	<u>48.47</u>	13.89	131.04	<u>320.35</u>
ICD	53.68	14.03	2.50	63.47	<u>38.33</u>	172.01	50.14	15.97	<u>1.39</u>	47.64	14.72	129.86	301.88
RSCD (Ours)	<u>57.85</u>	31.53	7.22	65.56	41.11	203.26	49.65	20.69	<u>1.39</u>	49.44	<u>14.17</u>	135.35	338.61

Table 12: Performance comparison of RSCD across different LVLM families and scales on MVH-Bench. The best and second-best results are highlighted in **bold** and underline, respectively.

Benchmark	Method	Acc.
MMBench	Baseline	85.22
	RSCD	85.63
SEED-Bench	Baseline	68.50
	RSCD	70.00

Table 13: Evaluation of RSCD on MMBench and SEED-Bench.

Benchmark	Method	Acc.	F1	Yes
POPE	Baseline	87.38	86.52	43.62
	RSCD	87.98	86.79	40.98

Table 14: Evaluation of RSCD on POPE.

E.4 Evaluation of Free-Form Generation

To examine MVH beyond constrained question-answering formats, we conduct an additional captioning evaluation with LLaVA-OV-7B on the MVH-Bench validation set. Given multi-view images, we ask the model to generate a caption for a queried view. As shown in Table 17, the base model often grounds its captions in non-queried views, resulting in relatively low reference accuracy. When RSCD is applied, reference accuracy improves from 69% to 77% without reducing the response length. In addition, when we reformulate

Benchmark	Method	CHAIRs ↓	CHAIRi ↓	Recall ↑
CHAIR	Baseline	43.20	11.07	65.15
	RSCD	47.80	10.14	67.74

Table 15: Evaluation of RSCD on CHAIR.

MVH-Bench into an open-ended format, RSCD outperforms the base model (Table 18). This suggests that the MVH problem goes beyond rigid QA formats and that RSCD remains effective in free-form generation.

E.5 Analysis of Layer Range Selection Cost

RSCD requires a lightweight layer selection procedure to identify the layer range Λ^* for each model. In our experiments, we perform this selection by probing only 30 paired-image samples, generating a short caption for each image. Notably, the middle-layer degradation pattern is clearly observable with as few as 3 paired-image samples, suggesting that the layer selection process can be performed with much lower cost if needed. Furthermore, experiments across diverse datasets and layer-wise analyses show that RSCD remains effective within the selected layer range for each model and is robust to minor variations in the selection. Therefore, once the layer range is identified, RSCD does not require

Method	Cross-Instance						Cross-View						MVH Score
	Binary			M.C.		Score	Binary			M.C.		Score	
	Acc	p-Acc	q-Acc	Acc	p-Acc		Acc	p-Acc	q-Acc	Acc	p-Acc		
Base model (two-view)	54.22	27.26	0.00	64.17	35.83	181.48	50.00	23.75	<u>2.50</u>	60.00	27.50	163.75	345.23
Base model (max-view)	52.24	23.33	0.00	59.17	23.61	158.35	51.25	23.75	0.00	56.25	27.50	158.75	317.10
VCD	<u>58.44</u>	28.12	5.56	<u>65.56</u>	<u>41.11</u>	<u>198.79</u>	51.25	20.00	<u>2.50</u>	<u>58.75</u>	27.50	160.00	<u>358.79</u>
ICD	57.95	<u>30.52</u>	<u>5.62</u>	61.67	35.83	191.59	48.12	21.25	0.00	56.25	22.50	148.12	339.71
AvisC	51.86	20.87	0.00	58.09	23.96	154.78	56.88	32.50	5.00	53.75	17.50	<u>165.63</u>	320.41
RSCD (Ours)	61.18	32.81	8.68	66.94	41.39	210.99	<u>55.00</u>	<u>31.25</u>	0.00	<u>58.75</u>	<u>25.00</u>	170.00	380.99

Table 16: Performance comparison on MVH-Bench under the original two-view and max-view settings. The best and second-best results are highlighted in **bold** and underline, respectively.

Method	Ref. Acc.	Avg. # Words
Base model	69	10.01
RSCD (Ours)	77	10.50

Table 17: Free-form captioning performance of RSCD on the MVH-Bench validation set.

Method	Cross-Inst. Score	Cross-View Score	MVH Score
Base model	139.38	96.25	235.62
RSCD (Ours)	153.12	150.00	303.12

Table 18: Open-ended evaluation performance of RSCD on the MVH-Bench validation set.

dataset-specific calibration.

E.6 Additional Qualitative Examples

In Figure 10, we present additional qualitative examples which are consistent with the observations in Section 5.5. As shown in the left part of Figure 10, the baseline model attends primarily to the “object located in the upper left of the image” in “view 1” (solid yellow line). However, after applying RSM, the attention spreads to the corresponding object in both views. Similarly, in the right part of Figure 10, the baseline model correctly attends to the “person in a gray shirt” when answering whether they are “standing with their hands behind their back” (solid yellow line). When RSM is employed, the model shifts its attention to another person who is “standing with their hands behind their back”, rather than the referenced “person in a gray shirt”. These examples clearly illustrate that RSM distorts the intended semantics of the query, so that the model might attend to semantically mismatched regions.

F Additional Details of MVH-Bench

F.1 Benchmark Statistics

Figure 9 summarizes the statistics of MVH-Bench. Figure 9(a) and (b) show the distribution of video categories for image pairs extracted from the Ego-Exo4D and LEMMA datasets, respectively. Figure 9(c) illustrates the distribution of image pair view configurations formed by combining first-person view (FPV) and third-person view (TPV): FPV&TPV, TPV&TPV, and FPV&FPV. Figure 9(d) presents the source dataset composition, where 66% of the samples originate from Ego-Exo4D and the remaining 34% from LEMMA. These statistics demonstrate that MVH-Bench covers diverse activity scenarios and maintains a balanced distribution of multi-view configurations.

F.2 Formal Definition of Evaluation Metrics

For completeness, we provide the formal definitions of the evaluation metrics introduced in Section 2.2. Let $m(q)$ be the model’s prediction for the question q , and $\mathbb{I}[\cdot]$ be the indicator function. Here, $x, y \in \{i, j\}$ denote indices of the sampled instance-descriptor pairs, where x specifies the instance (I_x) and y specifies the descriptor (D_y). MVH-Bench contains N multi-view image pairs, with the superscript n indexing the n -th image pair ($n \in \{1, \dots, N\}$).

For binary questions, the mean accuracy Acc, pair accuracy p-Acc, and quadruplet accuracy q-Acc are defined as

$$\text{Acc} = \frac{1}{4N} \sum_{n,x,y} \mathbb{I}[m(q_{xy}^n) = a_{xy}^n], \quad (6)$$

$$\text{p-Acc} = \frac{1}{2N} \sum_{n,x} \prod_y \mathbb{I}[m(q_{xy}^n) = a_{xy}^n], \quad (7)$$

$$\text{q-Acc} = \frac{1}{N} \sum_n \prod_{x,y} \mathbb{I}[m(q_{xy}^n) = a_{xy}^n]. \quad (8)$$

We additionally define the yes-error ratio (YER) as

$$\text{YER} = \frac{\sum_{n,x,y} \mathbb{I}[m(q_{xy}^n) = \text{Yes}] \mathbb{I}[m(q_{xy}^n) \neq a_{xy}^n]}{\sum_{n,x,y} \mathbb{I}[m(q_{xy}^n) \neq a_{xy}^n]}. \quad (9)$$

For multiple-choice questions, the mean accuracy Acc and pair accuracy p-Acc are defined as

$$\text{Acc} = \frac{1}{2N} \sum_{n,(x,y)} \mathbb{I}[m(q_{xy}^n) = a_{xy}^n], \quad (10)$$

$$\text{p-Acc} = \frac{1}{N} \sum_n \prod_{(x,y)} \mathbb{I}[m(q_{xy}^n) = a_{xy}^n]. \quad (11)$$

Finally, we define the adversarial error ratio (AER) as

$$\text{AER} = \frac{\sum_{n,(x,y)} \mathbb{I}[m(q_{xy}^n) = \text{B}] \mathbb{I}[m(q_{xy}^n) \neq a_{xy}^n]}{\sum_{n,(x,y)} \mathbb{I}[m(q_{xy}^n) \neq a_{xy}^n]}. \quad (12)$$

F.3 Data Construction and Human Verification Details

We provide additional details on the data construction and human verification process introduced in the main paper.

First, in the instance-descriptor pair extraction stage, we obtained 68,298 instance-descriptor pairs from 24,564 images. Based on these pairs, the automated QA generation process produced 111,671 QA pairs. During human verification, we filtered out questions with ambiguous references, such as those referring to non-specific instances like “background”, “something”, or “item”. To improve dataset diversity, we kept at most two QA sets per category for each image pair.

The human verification stage was conducted by five annotators with expertise in vision-language reasoning. As shown in Figure 11, each annotator was provided with a pair of multi-view images, the instance-descriptor pairs used to generate questions for each view, and the automatically generated QA pairs. We adopted a hierarchical verification process to ensure annotation quality. Specifically, four annotators independently verified the QA pairs for one subcategory, while the remaining annotator

reviewed the verified QA pairs together with the annotator responsible for the corresponding subcategory. Through this joint review process, disagreements were resolved, and samples that did not reach consensus were discarded.

During verification, we observed several failure patterns in the QA pairs generated by GPT-4o. For spatial descriptors, ambiguities frequently arise when the same relation applies to multiple objects; for example, descriptors such as “on the table” or “beside the wall” may refer to multiple instances. For action descriptors, overly generic verbs such as “standing” or “looking” are frequently generated. In addition, some numerical questions can be answered without visual grounding, such as questions asking for the “number of hands”.

After verification, the final dataset consists of 4,800 QA pairs and 1,274 images, corresponding to approximately 96% of the generated samples being filtered out during the human verification stage. Among the samples retained in the final benchmark, we estimate that around 70% involved partial human editing, including refining descriptors, specifying instances more clearly, correcting inaccurate answers, or fixing grammatical issues.

F.4 Rationale Behind the Template-Based Design of MVH-Bench

The goal of MVH-Bench is to provide a controlled diagnostic evaluation of multi-view hallucination. By holding the question format constant and varying only the instances, views, and/or descriptors, we ensured that any model errors could be attributed to incorrect instance/view grounding rather than to the question format itself. To balance this control with generalizability, we designed diverse templates across categories, subcategories, question types, and perspectives (TPV and FPV).

The purpose of instance-descriptor swapping is to construct visually plausible hard negatives. This ensures that success requires precise grounding rather than reliance on superficial cues. Each hard negative is paired with its corresponding positive example to enable pairwise evaluation for each question. This pairwise evaluation protocol, widely used in visual grounding benchmarks such as POPE and MME, provides a stricter test of consistent grounding and true reasoning. In MVH-Bench, we further strengthen this evaluation by adding q-Acc, which measures whether the model answers matched quadruples consistently. In short, the essence of our approach is to make it harder




Figure 9: MVH-Bench statistics. Video category distribution of image pairs extracted from (a) Ego-Exo4D and (b) LEMMA. (c) Distribution of image pair combinations across FPV and TPV. (d) Source dataset composition.

for models to obtain high scores through shortcuts, such as relying on only one view or part of the multi-view input.

F.5 Prompts for Instance-Descriptor Pair Extraction

We design tailored prompts for instance-descriptor pair extraction for each subcategory (action, object, numerical, and spatial) and for each viewpoint (FPV and TPV). Figures 12 to 15 show the prompts used for the first-person viewpoint for the four subcategories, while Figures 16 to 19 illustrate the corresponding prompts for the third-person viewpoint.

G Ethics Statement

To ensure responsible data use, we obtained the necessary licenses from the contributing institutions to use the Ego-Exo4D and LEMMA datasets in this study.




Figure 10: Qualitative example illustrating the effect of reference shift masking. Top: attention map of the base model. Bottom: attention map with reference shift mask.




Figure 11: User interface used during the human verification process for constructing MVH-Bench.




Figure 12: Instance-descriptor pair extraction prompt: action (first-person view image).

Instance-descriptor pair extraction prompt: object (first-person view image)

{First-Person View Image}

You are given the visual input from the camera worn by the user (referred to as 'I').

Based on this visual input, you must generate all instance-attribute description pairs that describe instance along with their color, pattern, or object class, strictly grounded in the visual content.

Instructions:

- Identify all visible people and objects in the scene.
- For each, describe one relevant visual attribute: color, pattern, or object class.
- Each pair must be a tuple: ('Instance Description', 'Attribute Description').
- Do not assign more than one attribute to the same instance; represent each object or person with a single, most visually supported description.
- Avoid vague descriptions like 'person', 'something', 'object', or 'background' unless they are detailed with context.
- Only generate up to five of the most clearly supported and valid descriptions based on the visual content.

Description Guidelines:

- 'Instance Description' should clearly specify the instance (a visible person or object).
 - 'the shirt I am wearing'
 - 'the sweater worn by the person holding a cup'
 - 'the shoes I am wearing'
 - 'the object placed on the table'
 - 'the object I am holding in my right hand'
- 'Attribute Description' must describe one of the following, based on what is visible:
 - A color such as:
 - 'navy blue'
 - 'green'
 - A pattern such as:
 - 'striped pattern'
 - 'checkered'
 - An object class or category such as:
 - 'sneakers'
 - 'book'
 - 'mug cup'

Output Format:

- Your final output must be a single Python list.
- Each item in the list must be a tuple: ('Instance Description', 'Attribute Description')

Validation Tips:

Use this question to check each pair:

- Is/Are [Instance Description] [Attribute Description]?

If the answer is yes and it is visually supported, the tuple is likely valid.

Strictly follow the format shown in the examples below.

Examples:

```
[
  ('the shirt I am wearing', 'navy blue'),
  ('the shirt worn by the man wearing a cap', 'white'),
  ...
]
```

Figure 13: Instance-descriptor pair extraction prompt: object (first-person view image).

Instance-descriptor pair extraction prompt: numerical (first-person view image)

{First-Person View Image}

You are given visual input from a camera worn by the user (referred to as 'I').

Based on this visual input, you must generate all count-based description pairs that describe the number of instances or comparative counts of visible instances, strictly grounded in the visual content.

Instructions:

- Identify all visible people and objects that can be counted.
- For each case, generate outputs in the following formats:
A count pair: ('Instance Description', 'Number')
- Avoid vague descriptions such as 'things', 'items', 'hand', 'person', 'object', or 'background' unless they are detailed with context.
- All numbers must be integers from 1 to 5.
- Only generate up to five of the most clearly supported and valid descriptions based on the visual content.

Description Guidelines:

- 'Instance Descriptions' should clearly specify the group of people or objects.
'people near the window'
'books on the shelf'
'cups on the table'

Output Format:

- Your final output must be a single Python list.
- Each item in the list must be a tuple: ('Instance Description', 'Number')

Validation Tips:

Use these questions to check each item:

- Is/Are there [Number] [Instance Description]?

If the answer is yes and it is visually supported, the entry is likely valid.

Strictly follow the format shown in the examples below.

Examples:

```
[
  ('people standing near the counter', 3),
  ('cups on the table', 2),
  ...
]
```

Figure 14: Instance-descriptor pair extraction prompt: numerical (first-person view image).

Instance-descriptor pair extraction prompt: spatial (first-person view image)

{First-Person View Image}

You are given visual input from a camera worn by the user (referred to as 'I').

Based on this visual input, you must generate all valid location-based description pairs that describe the location of an instance, strictly grounded in the visual content.

Instructions:

- Identify all visible people or objects whose positions are clearly identifiable, or whose relative positions can be meaningfully compared.
- For each case, generate outputs in the following formats:
A location pair: ('Instance Description', 'Location Description')
- Avoid vague descriptions such as 'things', 'items', 'hand', 'person', 'object', or 'background' unless they are detailed with context.
- Only generate up to five of the most clearly supported and valid descriptions based on the visual content.

Description Guidelines:

- 'Instance Description' should be a clearly visible person or object in the scene, such as:
'the blue mug cup'
'the book on the shelf'
'the person wearing [color/item]'
- 'Location Description' refers to the absolute position of an instance, such as:
'positioned above/below [object]'
'to the left/right of [person]'
'in front of [person]'
'behind [person]'
'in the top/bottom/left/right of the image'

Output Format:

- Your final output must be a single Python list.
- Each item in the list must be a tuple: ('Instance Description', 'Location Description')

Validation Tips:

Use these questions to check each item:

- Is [Instance Description] [Location Description]?

If the answer is yes and it is visually supported, the entry is likely valid.

Strictly follow the format shown in the examples below.

Examples:

```
[
  ('the scissors', 'in the right side of the image'),
  ('the tissue', 'to my right'),
  ...
]
```

Figure 15: Instance-descriptor pair extraction prompt: spatial (first-person view image).

Instance-descriptor pair extraction prompt: action (third-person view image)

{Third-Person View Image}

You are given the visual input from a fixed-position camera capturing a scene.

Based on this visual input, you must generate all person-action/posture description pairs that describe people and their actions or postures, strictly grounded in the visual content.

Instructions:

- Describe all people visible in the image using specific action or posture descriptions, grounded in visual evidence.
- Each pair must be a tuple: ('Person Description', 'Action/Posture Description').
- Do not assign more than one action or posture to the same person; represent each individual with a single, most visually supported description.
- Avoid vague descriptions like 'standing', 'moving', 'looking', or 'background' unless they are detailed with context.
- Only generate up to five of the most clearly supported and valid descriptions based on the visual content.

Description Guidelines:

- 'Person Description' should clearly refer to someone visible in the image, such as:
 - 'The person wearing [color/item]'
 - 'The man standing near the sink'
- 'Action/Posture Description' must be:
 - In present participle (verb+ing) form for actions such as:
 - 'holding a cup'
 - 'typing on a keyboard'
 - Or a posture-related phrase such as:
 - 'sitting cross-legged'
 - 'leaning back in a chair'

Output Format:

- Your final output must be a single Python list.
- Each item in the list must be a tuple:('Person Description', 'Action/Posture Description')

Validation Tips:

Use the following template questions to help verify the correctness of each tuple:

- Is [Person Description] [Action Description]?
- Is [Person Description] [Posture Description]?

If the answer is yes and it is visually supported, the tuple is likely valid.

Strictly follow the format shown in the examples below.

Examples:

```
[
  ('the person wearing a watch', 'sitting cross-legged'),
  ('the person wearing a green shirt', 'raising one arm'),
  ...
]
```

Figure 16: Instance-descriptor pair extraction prompt: action (third-person view image).

Instance-descriptor pair extraction prompt: object (third-person view image)

[\(Third-Person View Image\)](#)

You are given the visual input from a fixed-position camera capturing a scene.

Based on this visual input, you must generate all instance-attribute description pairs that describe instances along with their color, pattern, or object class, strictly grounded in the visual content.

Instructions:

- Identify all visible people and objects in the scene.
- For each, describe one relevant visual attribute: color, pattern, or object class.
- Each pair must be a tuple: ('Instance Description', 'Attribute Description').
- Do not assign more than one attribute to the same instance; represent each object or person with a single, most visually supported description.
- Avoid vague descriptions like 'person', 'something', 'object', or 'background' unless they are detailed with context.
- Only generate up to five of the most clearly supported and valid descriptions based on the visual content.

Description Guidelines:

- 'Instance Description' should clearly specify the instance (a visible person or object).
 - 'the shirt worn by the person taking the photo'
 - 'the sweater worn by the person holding a cup'
 - 'the shoes worn by the woman walking on the path'
 - 'the object placed on the table'
 - 'the object held in the right hand of the person wearing a black watch'
- 'Attribute Description' must describe one of the following, based on what is visible:
 - A color such as:
 - 'navy blue'
 - 'green'
 - A pattern such as:
 - 'striped pattern'
 - 'checkered'
 - An object class or category such as:
 - 'sneakers'
 - 'book'
 - 'mug cup'

Output Format:

- Your final output must be a single Python list.
- Each item in the list must be a tuple: ('Instance Description', 'Attribute Description')

Validation Tips:

Use this question to check each pair:

- Is/Are [Instance Description] [Attribute Description]?

If the answer is yes and it is visually supported, the tuple is likely valid.

Strictly follow the format shown in the examples below.

Examples:

```
[
  ('the shirt worn by the person taking the photo', 'navy blue'),
  ('the shirt worn by the man wearing a cap', 'white'),
  ...
]
```

Figure 17: Instance-descriptor pair extraction prompt: object (third-person view image).




Figure 18: Instance-descriptor pair extraction prompt: numerical (third-person view image).




Figure 19: Instance-descriptor pair extraction prompt: spatial (third-person view image).