

Cracking the Code of Hallucination in LVLMs with Vision-aware Head Divergence

Jinghan He^{1,2}, Kuan Zhu^{1,2}, Haiyun Guo^{1,2*}, Junfeng Fang³, Zhenglin Hua⁴,
Yuheng Jia⁴, Ming Tang¹, Tat-Seng Chua³, Jinqiao Wang^{1,2,5*}

¹Foundation Model Research Center, Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³National University of Singapore ⁴Southeast University ⁵Wuhan AI Research

hejinghan2022@ia.ac.cn, {kuan.zhu, haiyun.guo, jqwang}@nlpr.ia.ac.cn

Abstract

Large vision-language models (LVLMs) have made substantial progress in integrating large language models (LLMs) with visual inputs, enabling advanced multimodal reasoning. Despite their success, a persistent challenge is hallucination—where generated text fails to accurately reflect visual content—undermining both accuracy and reliability. Existing methods focus on alignment training or decoding refinements but primarily address symptoms at the generation stage without probing the underlying causes. In this work, we investigate the internal mechanisms driving hallucination in LVLMs, with an emphasis on the multi-head attention module. Specifically, we introduce Vision-aware Head Divergence (VHD), a metric that quantifies the sensitivity of attention head outputs to visual context. Based on this, our findings reveal the presence of vision-aware attention heads that are more attuned to visual information; however, the model’s overreliance on its prior language patterns is closely related to hallucinations. Building on these insights, we propose Vision-aware Head Reinforcement (VHR), a training-free approach to mitigate hallucination by enhancing the role of vision-aware attention heads. Extensive experiments demonstrate that our method achieves superior performance compared to state-of-the-art approaches in mitigating hallucinations, while maintaining high efficiency with negligible additional time overhead. The code is available at <https://github.com/jinghan1he/VHR>.

1 Introduction

Large vision-language models (LVLMs) (Dai et al., 2023; Liu et al., 2024b) represent a notable advancement in artificial intelligence by enabling large language models (LLMs) to understand visual inputs. However, LVLMs still face the challenge of

* Corresponding author.

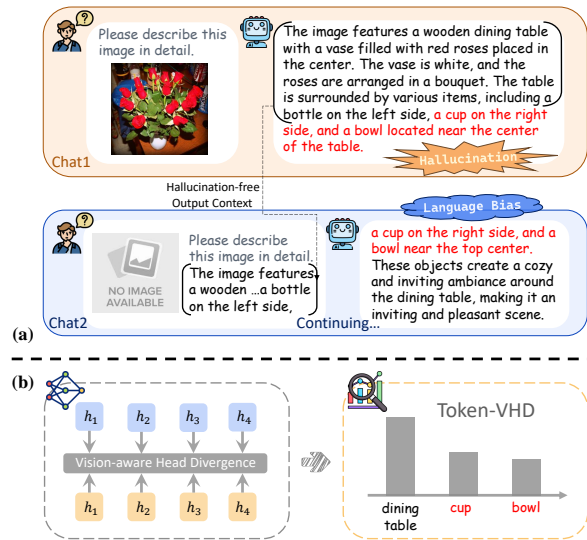


Figure 1: (a) An example indicating the connection between hallucination in LVLMs and language bias. When hallucination occurs (chat 1), we remove the image input and prompt the model to complete the description (chat 2). The output closely resembles the hallucinated content¹. (b) The proposed VHD metric measures the sensitivity of the attention head outputs to image inputs, reflecting the degree of language bias. Hallucinated words generally correspond to lower T-VHD scores.

hallucination (Rohrbach et al., 2018), where generated text does not accurately correspond to visual content. This misalignment can compromise the accuracy and reliability of LVLMs across a wide range of vision and language tasks, limiting their practical applications (You et al., 2024).

To mitigate this issue, several approaches incorporate additional information or models for alignment training (Zhao et al., 2023; Yu et al., 2024) or post-processing (Zhou et al., 2023; Yin et al., 2023), which incur higher training or inference costs. Recently, another line of research focuses on refining decoding strategies, employing methods like contrastive decoding (Leng et al., 2024; Zhu

¹Details on this example are exhibited in Appendix A.

et al., 2024; Kim et al., 2024; Gong et al., 2024) or beam search (Huang et al., 2024) to adjust the logits distribution during inference. However, these approaches merely intervene at the output level to rectify hallucinations after they occur, without directly targeting and adjusting the internal mechanisms that drive hallucinations. This work aims to fill this research gap.

One of the factors contributing to hallucination in LVLMs is their tendency to prioritize language patterns (Ghosh et al., 2024; Parcalabescu and Frank, 2024), which can lead to the generation of fluent but inaccurate content. We further investigate this phenomenon and present an example in Figure 1. Specifically, when prompted to continue generating an image description, the model generates highly consistent outputs, irrespective of whether an image is provided. This problem may arise from biased language patterns in the training data, which are incorporated into the model’s parameters (Liu et al., 2024a), causing output to rely more on internal knowledge than image context. Yu et al., 2023 analyzed similar biases in language models and revealed that the multi-head attention module contains both in-context and memory heads. Manipulating these heads can influence whether the output is driven by contextual information or internal knowledge.

Building on previous findings, we are inspired to investigate the relationship between hallucination in LVLMs and the multi-head attention mechanism. To this end, we introduce a novel metric, Vision-aware Head Divergence (VHD), to quantify how the output of each attention head changes when the image context is removed in a generation step of LVLMs. Our analysis reveals that only a few heads show significant sensitivity to the image context, while the majority exhibit minimal variation. Based on this, we aggregate the VHD values from the most prominent attention heads in a generation step, resulting in the Token-VHD (T-VHD) metric. This metric allows us to evaluate the model’s reliance on visual content versus language priors when predicting each token, as illustrated in the bottom-right part of Figure 1. By examining the T-VHD scores, we observe that words and sentences associated with hallucinations generally correspond to lower values, further supporting the role of language bias in hallucination in LVLMs.

Leveraging the insights above, we propose Vision-aware Head Reinforcement (VHR), a training-free approach aimed at enhancing the

model’s reliance on visual context rather than language priors. This method proactively mitigates hallucination in LVLMs by first identifying key attention heads based on their VHD scores and then amplifying their contributions during generation. Theoretical analysis demonstrates that this scaling-up operation effectively re-orientates the output of the multi-head attention module towards the reinforced head component, improving the alignment of the model’s output with visual context. Experiments on established LVLM hallucination benchmarks show that VHR outperforms existing decoding strategies, validating its effectiveness and efficiency in alleviating hallucinations.

Our main contributions can be summarized as follows:

- We propose the VHD metric to probe the attention heads in LVLMs for the language-bias tendency, and the T-VHD metric to analyze the relationship between language-biased generation and hallucination in LVLMs.
- We propose VHR, a training-free method that proactively mitigates hallucinations by adaptively identifying and reinforcing key attention heads during generation.
- Extensive experiments demonstrate that VHR outperforms existing decoding methods on widely-adopted hallucination benchmarks with negligible additional time cost.

2 Preliminary

LVLM generation. The LVLMs take both image and text as input. The image is encoded into vision tokens using an image encoder and projected to the text embedding space through a connector. These vision tokens x_V are then combined with tokenized text input x_T and passed into the LLM component for autoregressive generation:

$$y_t = \arg \max p_\theta(y_t | y_{<t}, x_V, x_T), \quad (1)$$

where $y_{<t}$ and y_t denote the earlier and the currently generated text tokens, respectively.

Multi-head attention. The multi-head attention mechanism is a core component of transformer models with each attention head performing the self-attention operation among tokens:

$$A_{l,i}(X_{l,i}) = \text{Attention}(X_{l,i}W_{l,i}^Q, X_{l,i}W_{l,i}^K, X_{l,i}W_{l,i}^V),$$

where $\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$. (2)




Figure 2: VHD scores of all the attention heads among all layers at one generation step.

$X_{l,i}$ and $A_{l,i}$ represent the input and output of the i -th attention head in the l -th layer, respectively. W^Q , W^K , and W^V denote the learned weight matrices for the query, key, and value transformations, respectively. d_k is the dimension of the query (Q) and key (K) vectors. The outputs of all the attention heads in the l -th layer are then concatenated and linearly transformed into the output space of this module:

$$\text{MHA}_l(X_l) = [A_{l,1}(X_{l,1}), \dots, A_{l,n_h}(X_{l,n_h})]W_l^O, \quad (3)$$

where n_h denotes the number of attention heads in each layer, X_l is the input to the MHA module in the l -th layer, and W^O is the learned weight matrices for the output linear transformation.

Attention head output during generation. To more clearly delineate the correspondence between the model’s intermediate outputs and its inputs, we combine Equation 1 and 2 to introduce the notation $A_{l,i}(y_t|y_{<t}, x_V, x_T)$. This notation represents the output of the i -th head in the l -th layer for generation step t , given the inputs x_V and x_T , and the generation history $y_{<t}$.

3 Method

3.1 Vision-aware Head Identification

Vision-aware head divergence (VHD). Inspired by the presence of in-context and memory heads in the model (Yu et al., 2023), we investigate whether different attention heads exhibit significantly different degrees of sensitivity to visual content. Specifically, we propose the vision-aware head divergence metric, which measures the change in the output of attention head for generation step t when the image

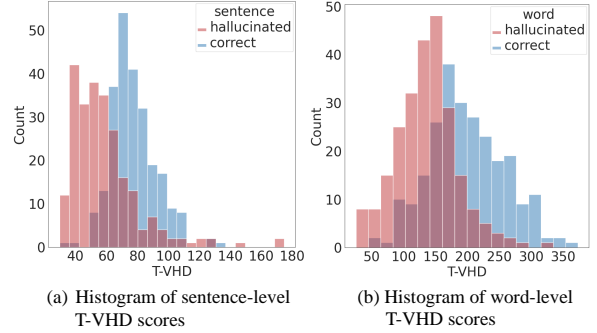


Figure 3: Relationship between T-VHD scores and hallucinations in LVLMs. Sentences and words associated with hallucinations generally correspond to lower T-VHD scores. Best viewed in color.

context is removed:

$$\text{VHD}_{l,i} = d(A_{l,i}(y_t|y_{<t}, x_V, x_T), A_{l,i}(y_t|y_{<t}, x_T)), \quad (4)$$

where d represents the Euclidean distance measure (Tabak, 2014).

Figure 2 visualizes the VHD scores for each attention head in the model. Specifically, we prompt LLaVA-1.5 with an image and the instruction "Please describe the image in detail" to generate descriptions, calculating the VHD scores when predicting the first token. The results show that a few attention heads exhibit notably higher VHD scores, while the others show minimal sensitivity. This suggests the presence of vision-aware attention heads that are more attuned to visual information. More examples of VHD scores during the generation process are presented in Appendix E.

Token-VHD. Beyond the varying degrees of vision awareness within the model, we further explore whether the VHD scores vary across different token generation steps. To this end, the VHD scores of the most prominent attention heads in each layer of the model are aggregated into the Token-VHD metric:

$$\text{T-VHD} = \sum_l \sum_i \text{topk}_i(\text{VHD}_{l,i}, k). \quad (5)$$

Note that we only consider the top k VHD scores at each layer to ensure that the large number of insensitive attention heads does not dilute the aggregation metric. Eventually, T-VHD metric serves as an indicator of the model’s reliance on visual information over language priors when predicting a specific token.

Leveraging this metric, we can quantitatively analyze the relationship between hallucination in

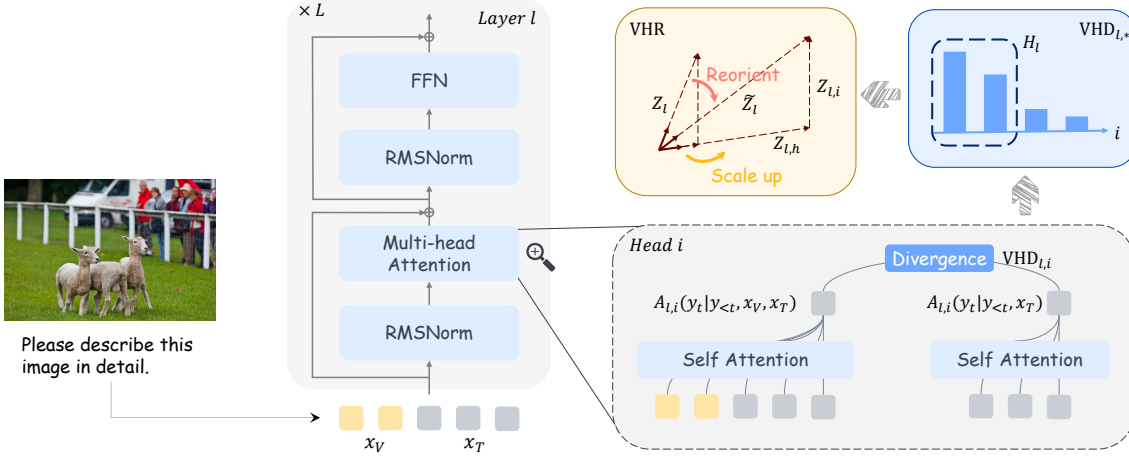


Figure 4: The illustration of the proposed VHD metric and the VHR approach to mitigate hallucinations in LVLm. We select the attention heads that are sensitive to visual information for a given layer based on the VHD metric, *i.e.* H_l , and then amplify their outputs to reinforce their contributions.

LVLms and language bias at different levels of granularity, specifically at the sentence and word levels. To this end, we conduct an experiment on a random sample of 500 images from the CHAIR benchmark (Rohrbach et al., 2018), tracking the T-VHD scores at each generation step. Object-related words in the generated descriptions are classified as either hallucinated or correct, depending on whether they appear in the annotated object set for the given image. Sentences are then labeled according to whether they contain any hallucinated words. Figure 3 presents the experimental results, highlighting the distributional differences in T-VHD scores between hallucinated and correct instances. These findings provide statistical evidence that language bias is closely related to hallucinations in LVLms.

3.2 Vision-aware Head Reinforcement

Since only a small subset of attention heads within the model are sensitive to visual information, we can amplify their contributions during generation to strengthen the model’s reliance on visual cues and counteract language bias. As discussed in Section 3.1, the VHD metric effectively captures the sensitivity of attention heads to visual information, making it a suitable indicator for selecting key attention heads for reinforcement. However, we observed that some high VHD values stem from a surge in the activation of attention heads upon the removal of visual context, indicating negative vision sensitivity. Amplifying the contributions of such heads would diverge from our objective. Therefore, we propose to zero out these undesired outliers, *i.e.*

$VHD_{l,i} = 0$, if the following condition satisfies:

$$\begin{cases} VHD_{l,i} > \mu(VHD_{l,*}) + \sigma(VHD_{l,*}), \\ TA_{l,i} > \mu(TA_{l,*}) + \sigma(TA_{l,*}), \end{cases} \quad (6)$$

where $TA_{l,i} = \|A_{l,i}(y_t|y_{<t}, x_T)\|^2$.

μ and σ represent the mean and standard deviation. Next, for the multi-head attention module in a given layer of the model, we select the first half of the attention heads based on their VHD scores and directly scale up their outputs by a factor of α :

$$\tilde{A}_{l,i} = \begin{cases} \alpha \cdot A_{l,i}, & \text{if } i \in H_l, \\ A_{l,i}, & \text{otherwise,} \end{cases} \quad (7)$$

where $H_l = \{i | VHD_{l,i} > \text{median}(VHD_{l,*})\}$.

Apply VHR layer by layer. This specific implementation allows for the selection and reinforcement of attention heads within a single forward pass, as opposed to first selecting the heads in all layers and then reinforcing them in two separate passes. Additionally, when VHR is applied across multiple layers simultaneously, the reinforcement in earlier layers can influence the VHD scores of subsequent layers. The layer-by-layer VHR strategy helps to avoid such inconsistencies, as the previous layers are already reinforced when calculating the VHD scores for a given layer.

Determine the heads at the first generation step. Although we can compute the VHD scores and select the key heads at each generation step, reinforcing different heads at different steps may bring negative effects. Specifically, LVLms rely on KV caching to speed up inference, which means

that the keys and values of the previous tokens will not be recalculated in subsequent generation steps. Therefore, the important heads should be determined at the beginning of the generation process to ensure consistency in the Q , K , and V of all tokens in the attention module. Our experimental results show that this approach is sufficient to mitigate hallucinations.

Comparison with other head identification methods. Different from existing attention head identification methods in the field of model interpretability (Yu et al., 2023; Li et al., 2024; Zhou et al., 2024; Fang et al., 2024a; He et al., 2024), VHR does not require any annotation and can adaptively detect the key heads for each sample. Furthermore, rather than identifying and reinforcing all key attention heads in the model in two forward passes, VHR ensures computational efficiency and metric consistency by iteratively applying the *select-then-reinforce* approach across the layers. The complete procedure of VHR is provided in Algorithm 1.

3.3 Attention Output Reorientation

Scaling up the outputs of certain attention heads within a layer to reinforce its contribution is a straightforward and intuitive operation, and we present a theoretical analysis to substantiate its rationale. Consider the input to the FFN module following the MHA module in layer l , which can be expressed as follows:

$$\begin{aligned} Z_l &= \text{RMSNorm}(\hat{X}_l + \text{MHA}_l(X_l)) \\ &= \hat{g}_l \cdot \frac{\hat{X}_l + \text{MHA}_l(X_l)}{\|\hat{X}_l + \text{MHA}_l(X_l)\|}, \end{aligned} \quad (8)$$

where \hat{g}_l is a fixed constant after training, and \hat{X}_l is the input to the l -th layer before RMSNorm. Due to the normalization operation, only the direction of the overall output from earlier modules is crucial.

Proposition 1 Consider a layer l within an LVLm, and let h be the index of the attention head to be reinforced. Let \tilde{Z}_l be the input to the FFN module obtained with $\tilde{A}_{l,h} = \alpha \cdot A_{l,h}$ ($\alpha > 1$), Z_l be the original input obtained with $A_{l,h}$, $Z_{l,h}$ be the pseudo-input obtained with only the $A_{l,h}$ component. Then it holds that $\cos(\tilde{Z}_l, Z_{l,h}) > \cos(Z_l, Z_{l,h})$.

The proof is detailed in Appendix B. Proposition 1 implies that amplifying the output of a specific head in the MHA module effectively reorients the direction of Z_l towards the output direction of the

Algorithm 1 VHR

Input image x_V , instruction x_T , generation step t , scale factor α , layers to reinforce L_r

```

1: for layer  $l \in L_r$  do
2:   if  $t = 0$  then
3:     Compute  $\text{VHD}_{l,*}$   $\triangleright$  Equation 4
4:     Zero out  $\text{VHD}_{l,i}$  if Equation 6 holds
5:     Select the heads as  $H_l$   $\triangleright$  Equation 7
6:   end if
7:   Reinforce the heads in  $H_l$   $\triangleright$  Equation 7
8: end for

```

reinforced head component. This provides theoretical support for the mechanism underlying the reinforcement of the key attention head. An overview of the proposed VHD metric and the VHR method is presented in Figure 4.

4 Experiments

4.1 LVLms

We conduct experiments on three of the most representative LVLms, *i.e.* InstructBLIP-7b (Dai et al., 2023), LLaVA-1.5-7b (Liu et al., 2024b), and LLaVA-NeXT-7b (Liu et al., 2024c). LVLms are typically composed of an image encoder, a connector, and an LLM. Specifically, LLaVA-1.5-7b and LLaVA-NeXT-7b leverage MLP to align the visual and textual embedding space and feed all the image tokens from the image encoder to the LLM. In contrast, InstructBLIP uses Q-Former to reduce the number of image tokens before passing them to the LLM. LLaVA-NeXT differs from LLaVA-1.5 by offering a higher image resolution, allowing it to capture more visual details.

4.2 Benchmarks

CHAIR. The Caption Hallucination Assessment with Image Relevance (CHAIR) metric (Rohrbach et al., 2018) evaluates object hallucination in image captioning by comparing generated captions with ground truth data. It identifies objects mentioned in captions but absent in images and calculates their proportion to quantify hallucination. Specifically, CHAIR includes two metrics at both caption level (CHAIR_S) and object level (CHAIR_I):

$$\begin{aligned} \text{CHAIR}_S &= \frac{|\{\text{caption w/ hallucinated objects}\}|}{|\{\text{all captions}\}|}, \\ \text{CHAIR}_I &= \frac{|\{\text{hallucinated objects}\}|}{|\{\text{all mentioned objects}\}|}. \end{aligned} \quad (9)$$

	InstructBLIP			LLaVA-1.5			LLaVA-NeXT		
	CHAIR _S ↓	CHAIR _I ↓	Len	CHAIR _S ↓	CHAIR _I ↓	Len	CHAIR _S ↓	CHAIR _I ↓	Len
Greedy	45.32 \pm 2.24	12.98 \pm 0.76	91.06	49.68 \pm 1.47	14.32 \pm 0.78	83.06	29.08 \pm 2.09	8.08 \pm 0.74	157.06
Beam	48.56 \pm 1.66	13.50 \pm 0.44	94.87	53.84 \pm 2.41	15.60 \pm 0.46	87.47	<u>25.72</u> \pm 2.17	6.92 \pm 0.88	160.64
DoLa	46.00 \pm 1.87	13.00 \pm 0.91	90.75	50.88 \pm 2.34	14.64 \pm 0.90	82.41	28.76 \pm 2.58	8.12 \pm 0.78	155.75
VCD	50.72 \pm 2.44	14.42 \pm 0.99	90.39	51.92 \pm 1.87	15.42 \pm 0.84	83.12	30.80 \pm 2.48	8.72 \pm 0.94	157.72
OPERA	45.76 \pm 2.32	13.06 \pm 0.88	92.46	44.28 \pm 0.95	13.36 \pm 0.47	75.88	-	-	-
CODE	50.76 \pm 2.06	14.12 \pm 0.93	88.57	47.96 \pm 0.80	14.26 \pm 0.57	78.52	27.84 \pm 2.73	7.98 \pm 0.92	151.51
EAH	46.40 \pm 1.15	13.13 \pm 0.60	92.33	<u>38.76</u> \pm 2.47	<u>11.05</u> \pm 0.81	86.28	28.13 \pm 1.13	6.62 \pm 0.49	142.75
VHR	37.76 \pm 2.76	9.75 \pm 0.98	106.49	33.32 \pm 1.31	9.71 \pm 0.45	81.33	24.96 \pm 2.09	<u>6.80</u> \pm 0.59	156.92

Table 1: CHAIR evaluation results on MSCOCO dataset averaged over 5 random splits, with best in **bold** and second-best underlined. *Len* represents the average number of words in the generated descriptions.

We randomly sample 500 images from the COCO 2014 validation set and repeat the experiments for five times with different random seeds. The LVLMs are prompted with "Please describe this image in detail." to get the descriptions. We report the average results for each metric along with the standard deviation.

POPE. POPE (Li et al., 2023) is a dataset for evaluating object hallucinations by having models answer true or false questions about the presence of objects in images. The dataset includes 500 images from MSCOCO (Lin et al., 2014), with each image paired with questions like "Is there a <object> in the image?". The evaluation consists of three splits—random, popular, and adversarial—where objects are sampled in different ways. The evaluation metrics include Accuracy, Precision, Recall, and F1 scores, with the results averaged across all three splits.

LLaVA-Bench. LLaVA-Bench (In-the-Wild) (Liu et al., 2024b) is a comprehensive benchmark designed to evaluate the performance of vision-language models on a wide range of challenging tasks. It includes 24 images across diverse domains, such as indoor and outdoor scenes, memes, accompanied by 60 carefully crafted questions covering simple question answering, detailed descriptions, and complex reasoning. Due to the open-ended nature and complexity of the responses, we prompt the GPT-4V model to evaluate the LVLMs’ outputs in terms of accuracy, detailedness, and naturalness.

4.3 Baselines

We compare VHR with the popular training-free methods that do not introduce external information or models: DoLa (Chuang et al., 2023) derives the next-token distribution by contrasting the logits from later and earlier layers; VCD (Leng et al.,

2024) contrasts the output distribution generated from the original and distorted image; OPERA (Huang et al., 2024) mitigates over-trust in previous summary tokens in beam-search decoding; CODE (Kim et al., 2024) uses self-generated descriptions as contrast references to improve alignment with the actual visual content; EAH (Zhang et al., 2024) enhances the attention sinks on image tokens in shallow layers. In addition, we also compare the performance of base LVLMs using greedy and beam search decoding.

4.4 Implementation Details

We set α to 2 to strike a balance between effectively correcting hallucinations and minimizing the invasiveness of hidden states manipulation. VHR is applied to the second and last 14 layers for LLaVA series and the last 18 layers for InstructBLIP. We faithfully reproduced all baseline methods based on their open-source repositories and set the hyperparameters according to the values reported in the papers. The results of all methods are reported under consistent conditions of base models, prompts, and generation parameters to ensure a fair comparison. Specifically, the `max_new_token` is set to 512, and the number of beams is set to 5 for all methods involving beam search.

4.5 Results

CHAIR. Table 1 presents the performance of VHR in comparison to all baseline approaches on the CHAIR benchmark. The results for OPERA on LLaVA-NeXT are absent due to its excessive memory requirements. VHR demonstrates robust performance across all three LVLMs, achieving reductions of up to 16.36 in CHAIR_S and 4.61 in CHAIR_I on LLaVA-1.5. Notably, with increased image resolution and enhanced model capabili-

	InstructBLIP	LLaVA-1.5	LLaVA-NeXT
Greedy	<u>85.36</u>	84.98	<u>88.51</u>
Beam	84.40	85.30	87.97
DoLa	85.21	85.07	88.46
VCD	84.67	84.41	88.11
OPERA	84.41	<u>85.45</u>	-
CODE	84.80	84.63	88.44
EAH	85.18	85.03	84.28
VHR	85.52	85.47	88.87

Table 2: F1 scores on POPE averaged over popular, adversarial, and random splits, with best in **bold** and second-best underlined.

	Accuracy	Detailedness	Naturalness
InstructBLIP	4.917	5.017	6.717
w/ VHR	5.250	5.117	6.733
LLaVA-1.5	6.017	6.100	7.400
w/ VHR	6.333	6.217	7.333
LLaVA-NeXT	5.383	6.750	7.900
w/ VHR	5.783	6.700	7.833

Table 3: LLaVA-Bench (In-the-Wild) evaluation results, scored by GPT-4o via pairwise response comparison.

ties, LLaVA-NeXT already exhibits a significant reduction in hallucinations compared to other base LLMs, but VHR continues to exhibit notable effectiveness in mitigating its hallucinations. Moreover, VHR consistently outperforms baseline methods with greater stability, requiring only minor trade-offs in the length or richness of the generated description.

POPE. As shown in Table 2, VHR outperforms all other decoding methods across all LLMs. While the binary (yes/no) benchmark limits the full demonstration of VHR’s strength in handling language bias, our method still consistently improves performance across models of varying capabilities. This indicates that VHR is a robust and effective training-free strategy for enhancing models at different performance levels.

LLaVA-Bench. The GPT-4V evaluation results on LLaVA-Bench (In-the-Wild) are presented in Table 3. These results demonstrate that VHR improves model accuracy on highly diverse and challenging tasks while preserving a consistent level of detailedness and naturalness. Additionally, the *Len* metric in Table 1 and the qualitative results in Figure 7 also indicate that VHR has only a minor impact on language generation quality, despite its enhancement of vision-aware attention heads.

Model	Method	CHAIR _S ↓	CHAIR _I ↓
InstructBLIP	VHR	37.76	9.75
	fixed VHR	45.40	13.57
	outlier VHR	37.76	10.18
LLaVA-1.5	VHR	33.32	9.71
	fixed VHR	44.72	13.81
	outlier VHR	36.88	10.36
LLaVA-NeXT	VHR	24.96	6.80
	fixed VHR	36.96	9.80
	outlier VHR	24.64	6.37

Table 4: Ablation study on adaptively determining key heads per sample and removing outlier VHD scores.

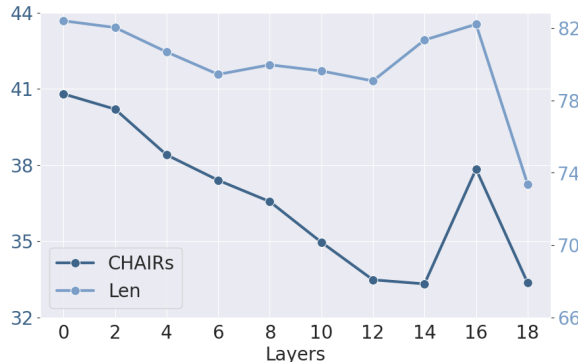


Figure 5: Results of VHR applied across different numbers of the last few layers in LLaVA-1.5.

4.6 Ablation Study

Impact of Adaptive Attention Head Selection.

Since VHR reinforces different attention heads for each sample, we conduct an ablation study to validate the necessity of this strategy. As shown in Table 4, fixing a set of attention heads identified by one sample for reinforcement across all samples leads to a significant performance drop.

Impact of Outlier VHD Score Removal. VHR relies on VHD scores to identify critical attention heads, but these scores measure absolute influence, conflating both positive and negative vision sensitivity. To selectively enhance heads that improve visual context awareness, we propose excluding outlier VHD scores during selection (Equation 6). As shown in Table 4, this refinement consistently matches or outperforms the baseline, demonstrating that pruning unhelpful sensitivity is essential for optimal head selection.

Impact of Reinforced Layers. Figure 5 shows the ablation study results on the number of the last few layers for reinforcement. Increasing the number of reinforced layers continuously alleviates hallucinations, with optimal performance reached

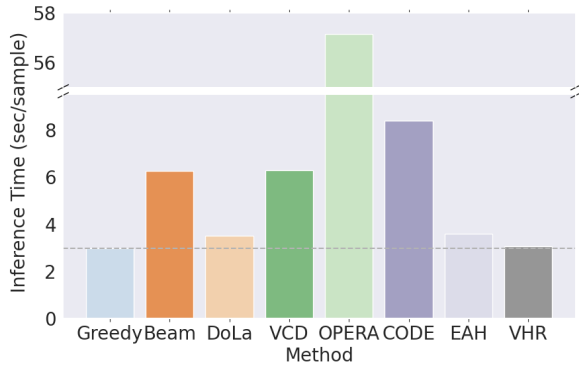


Figure 6: Comparison of inference time for different methods.

at the last 14 layers. However, further reinforcement degrades the model’s generation quality and fails to effectively mitigate hallucinations. More discussion on the choice of reinforced layers and scale factor can be found in Appendix C and D.

4.7 Further Analysis

Additional Time Analysis. For the first generation step, VHR requires an additional forward pass, removing the image context to calculate the VHD scores. In subsequent generation steps, only the scaling operation is needed. As a result, the extra computation introduced by VHR is negligible throughout the entire generation process. A detailed inference time comparison between VHR and baseline methods is presented in Figure 6.

Qualitative Results. To clearly demonstrate the effect of VHR in reducing hallucinations, we provide a concrete example in Figure 7. Without VHR, the LVLM generates content that is absent from the image, such as mentioning people watching the game in the background. This could stem from inherent language bias in the training data. When VHR is applied, the outputs are more accurately aligned with the actual content of the image.

5 Related Work

5.1 Hallucination in LVLMs

Hallucination in LVLMs refers to discrepancies between the generated text and the actual content present in the corresponding image. This issue stems from multiple factors, including data bias, insufficient image grounding by vision encoders, and misalignment between modalities (Liu et al., 2024a). Current approaches to mitigating hallucinations in LVLMs can be broadly classified into three categories: training alignment (Zhao et al.,

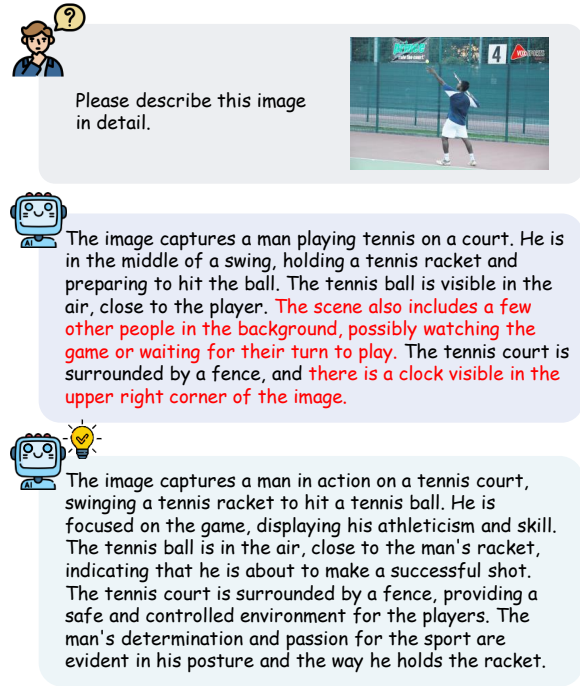


Figure 7: An example of VHR eliminating the hallucinated content.

2023; Yu et al., 2024), post-processing (Zhou et al., 2023; Yin et al., 2023; Fang et al., 2024b), and decoding strategies (Leng et al., 2024; Zhu et al., 2024; Huang et al., 2024; Kim et al., 2024; Gong et al., 2024). The first two categories often rely on external resources or models to improve performance, whereas our proposed VHR method operates solely based on the model’s internal states, without requiring additional external information. In contrast to existing decoding strategies, VHR intervenes proactively within the model to address hallucinations before they occur, rather than modifying the logits distribution at the output stage. A recent approach, EAH (Zhang et al., 2024), also targets the attention mechanism to reduce hallucinations, but it focuses specifically on the attention sink phenomenon. In comparison, VHR reduces language bias in LVLMs by exploiting cross-modal contrast and re-adjusting the contributions of attention heads across multiple layers. Consequently, the underlying motivations and methodologies of the two works differ significantly.

5.2 Language Bias in LVLMs

Language bias refers to the tendency of models to prioritize language patterns or prior knowledge over the actual visual context presented in the input. This issue has a long-standing research history pre-

dating the emergence of LVLMs and was studied through methods like balanced multimodal training (Goyal et al., 2017) and causal inference (Niu et al., 2021). With the advent of LVLMs, language bias has become an even more pressing concern. These models are typically pre-trained on massive corpora of text data, further exacerbating the potential for language overfitting. Studies have identified that LVLMs, despite their capabilities in visual recognition, struggle to fully integrate visual context (Parcalabescu and Frank, 2024), leading to hallucinations, particularly in reasoning tasks (Ghosh et al., 2024). To address this growing challenge, several recent methods have proposed solutions such as contrastive decoding (Leng et al., 2024; Zhu et al., 2024) and visual description grounding decoding (Ghosh et al., 2024). However, these approaches directly manipulate the output logits, which introduces instability during generation and lacks a thorough analysis of the model’s internal mechanisms. In contrast, our approach intervenes directly within the model, providing a more interpretable and effective way to address the internal factors driving language bias in LVLMs while complementing existing methods.

6 Conclusion

This work investigates the connection between hallucination in LVLMs and the multi-head attention mechanism. We introduce the VHD metric, which quantifies the sensitivity of attention head outputs to visual context, revealing that language bias can contribute to hallucinations in LVLMs. Building on these findings, we propose VHR, a training-free approach that strengthens the role of vision-aware attention heads to mitigate hallucinations. Extensive experiments demonstrate that VHR outperforms existing methods, significantly improving the alignment of LVLMs with visual information.

Limitations

Our analysis and mitigation strategy primarily focus on the multi-head attention mechanism of LVLMs. While this is a critical component influencing hallucinations, there may be other architectural factors—such as those in the vision encoder and the FFN module in the LLMs—that contribute to hallucinations but were not directly addressed in this study. Future work could focus on more comprehensive interventions that span the entire model, going beyond attention head manipulation.

Acknowledgements

This work was supported by National Key R&D Program of China under Grant No.2024YFC3015501, also partly supported by National Natural Science Foundation of China under Grant No.62276260, Beijing Municipal Science and Technology Project under Grant Z231100007423004, Beijing Natural Science Foundation under Grant 4244099, Aeronautical Science Foundation of China under Grant 2024M0710M0002.

References

- Liang Chen, Haozhe Zhao, Tianyu Liu, Shuai Bai, Junyang Lin, Chang Zhou, and Baobao Chang. 2025. An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models. In *European Conference on Computer Vision*, pages 19–35. Springer.
- Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. 2023. Dola: Decoding by contrasting layers improves factuality in large language models. *arXiv preprint arXiv:2309.03883*.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. *Instructblip: Towards general-purpose vision-language models with instruction tuning*. *Preprint*, arXiv:2305.06500.
- Junfeng Fang, Zac Bi, Ruipeng Wang, Houcheng Jiang, Yuan Gao, Kun Wang, An Zhang, Jie Shi, Xiang Wang, and Tat-Seng Chua. 2024a. Towards neuron attributions in multi-modal large language models. *Advances in Neural Information Processing Systems*, 37:122867–122890.
- Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Shi Jie, Xiang Wang, Xiangnan He, and Tat-Seng Chua. 2024b. Alphaedit: Null-space constrained knowledge editing for language models. *arXiv preprint arXiv:2410.02355*.
- Sreyan Ghosh, Chandra Kiran Reddy Evuru, Sonal Kumar, Utkarsh Tyagi, Oriol Nieto, Zeyu Jin, and Dinesh Manocha. 2024. Visual description grounding reduces hallucinations and boosts reasoning in lvlms. *arXiv preprint arXiv:2405.15683*.
- Xuan Gong, Tianshi Ming, Xinpeng Wang, and Zhihua Wei. 2024. Damro: Dive into the attention mechanism of lvlm to reduce object hallucination. *arXiv preprint arXiv:2410.04514*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding

- in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913.
- Jinghan He, Haiyun Guo, Kuan Zhu, Zihan Zhao, Ming Tang, and Jinqiao Wang. 2024. Seekr: Selective attention-guided knowledge retention for continual learning of large language models. *arXiv preprint arXiv:2411.06171*.
- Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming Zhang, and Nenghai Yu. 2024. Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13418–13427.
- Junho Kim, Hyunjun Kim, Yeonju Kim, and Yong Man Ro. 2024. Code: Contrasting self-generated description to combat hallucination in large multi-modal models. *arXiv preprint arXiv:2406.01920*.
- 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.
- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2024. Inference-time intervention: Eliciting truthful answers from a language model. *Advances in Neural Information Processing Systems*, 36.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*.
- 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 *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer.
- Hanchao Liu, Wenyuan Xue, Yifei Chen, Dapeng Chen, Xiutian Zhao, Ke Wang, Liping Hou, Rongjun Li, and Wei Peng. 2024a. A survey on hallucination in large vision-language models. *arXiv preprint arXiv:2402.00253*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024b. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024c. [Llava-next: Improved reasoning, ocr, and world knowledge](#).
- Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. 2021. Counterfactual vqa: A cause-effect look at language bias. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12700–12710.
- Letitia Parcalabescu and Anette Frank. 2024. Do vision & language decoders use images and text equally? how self-consistent are their explanations? *arXiv preprint arXiv:2404.18624*.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. Object hallucination in image captioning. *arXiv preprint arXiv:1809.02156*.
- John Tabak. 2014. *Geometry: the language of space and form*. Infobase Publishing.
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Tong Xu, Hao Wang, Dianbo Sui, Yunhang Shen, Ke Li, Xing Sun, and Enhong Chen. 2023. Woodpecker: Hallucination correction for multimodal large language models. *arXiv preprint arXiv:2310.16045*.
- Junwei You, Haotian Shi, Zhuoyu Jiang, Zilin Huang, Rui Gan, Keshu Wu, Xi Cheng, Xiaopeng Li, and Bin Ran. 2024. V2x-vlm: End-to-end v2x cooperative autonomous driving through large vision-language models. *arXiv preprint arXiv:2408.09251*.
- Qinan Yu, Jack Merullo, and Ellie Pavlick. 2023. Characterizing mechanisms for factual recall in language models. *arXiv preprint arXiv:2310.15910*.
- Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwen He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, et al. 2024. Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13807–13816.
- Xiaofeng Zhang, Yihao Quan, Chaochen Gu, Chen Shen, Xiaosong Yuan, Shaotian Yan, Hao Cheng, Kaijie Wu, and Jieping Ye. 2024. Seeing clearly by layer two: Enhancing attention heads to alleviate hallucination in lvlms. *arXiv preprint arXiv:2411.09968*.
- Zhiyuan Zhao, Bin Wang, Linke Ouyang, Xiaoyi Dong, Jiaqi Wang, and Conghui He. 2023. Beyond hallucinations: Enhancing lvlms through hallucination-aware direct preference optimization. *arXiv preprint arXiv:2311.16839*.
- Yiyang Zhou, Chenhong Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. 2023. Analyzing and mitigating object hallucination in large vision-language models. *arXiv preprint arXiv:2310.00754*.
- Zhenhong Zhou, Haiyang Yu, Xinghua Zhang, Rongwu Xu, Fei Huang, Kun Wang, Yang Liu, Junfeng Fang, and Yongbin Li. 2024. On the role of attention

heads in large language model safety. *arXiv preprint arXiv:2410.13708*.

Lanyun Zhu, Deyi Ji, Tianrun Chen, Peng Xu, Jieping Ye, and Jun Liu. 2024. *Ibd: Alleviating hallucinations in large vision-language models via image-biased decoding. arXiv preprint arXiv:2402.18476*.

A A Case Study of Language Bias

Figure 8 illustrates an example that reflects language bias in LVLMs. We first prompted the model to describe the image, resulting in a complete description. We observed that the latter part of this description contained hallucinated content. To investigate the role of language bias in this case, we removed the image input and re-prompted the model using only the original text prompt and the non-hallucinated part of the initial response. This allowed us to observe how the model would continue the text based solely on its internal knowledge. We found that the continuation closely resembled the previously hallucinated content, suggesting it was generated based on internal language priors rather than visual evidence. This example clearly demonstrates the connection between language bias and hallucination in LVLMs.

In the lower part of Figure 8, we visualize the T-VHD scores for each word in the generated description. The word color intensity reflects the T-VHD scores, with darker shades indicating higher sensitivity to visual input. The results show that words within fixed phrases (e.g., *surrounded by*, *on the right/left side*) generally have lower T-VHD scores, suggesting stronger reliance on language priors. In contrast, object terms (e.g. *wooden dining table*, *vase*, *red roses*) typically exhibit higher T-VHD scores upon first mention, indicating greater dependence on visual context. Notably, among all object terms, hallucinated items (e.g., *cup*, *bowl*, *chair*) tend to have lower T-VHD scores, highlighting reduced visual grounding. This example demonstrates how the proposed T-VHD metric effectively captures the model’s reliance on visual information versus language priors at the token level.

B Attention Output Reorientation

Proposition 1 Consider a layer l within an LVLm, and let h be the index of the attention head to be reinforced. Let \tilde{Z}_l be the input to the FFN module obtained with $\tilde{A}_{l,h} = \alpha \cdot A_{l,h}$ ($\alpha > 1$), Z_l be the original input obtained with $A_{l,h}$, $Z_{l,h}$ be the pseudo-input obtained with only the $A_{l,h}$ component. Then it holds that $\cos(\tilde{Z}_l, Z_{l,h}) > \cos(Z_l, Z_{l,h})$.



Figure 8: An example indicating the connection between hallucination in LVLMs and language bias. The proposed T-VHD metric reflects the model’s reliance on visual content versus language priors at the token level.

Proof. By partitioning the projection matrix W_l^O , we can further express the output of MHA as the sum of the contributions from each attention head:

$$\begin{aligned} \text{MHA}_l(X_l) &= [A_{l,1}(X_{l,1}), \dots, A_{l,n_h}(X_{l,n_h})]W_l^O \\ &= A_{l,1}(X_{l,1})W_{l,1}^O + \dots + A_{l,n_h}(X_{l,n_h})W_{l,h}^O. \end{aligned} \quad (10)$$

To simplify the notation, we define \mathbf{x} and \mathbf{y} as follows:

$$\begin{aligned} \mathbf{x} &= \hat{X}_l + \text{MHA}_l(X_l), \\ \mathbf{y} &= A_{l,h}(X_{l,h})W_{l,h}^O. \end{aligned} \quad (11)$$

We then substitute \mathbf{x} and \mathbf{y} for the variables in Equation 8 to derive $Z_l, \tilde{Z}_l, Z_{l,h}$:

$$\begin{aligned} Z_l &= \hat{g}_t \frac{\mathbf{x}}{\|\mathbf{x}\|}, \\ \tilde{Z}_l &= \hat{g}_t \frac{\mathbf{x} + (\alpha - 1)\mathbf{y}}{\|\mathbf{x} + (\alpha - 1)\mathbf{y}\|}, \\ Z_{l,h} &= \hat{g}_t \frac{\mathbf{y}}{\|\mathbf{y}\|}. \end{aligned} \quad (12)$$

Lastly, we prove that $\cos(\tilde{Z}_l, Z_{l,h})$ is greater than $\cos(Z_l, Z_{l,h})$:

$$\begin{aligned}
& \cos(\tilde{Z}_l, Z_{l,h}) - \cos(Z_l, Z_{l,h}) \\
&= \frac{\langle \mathbf{x} + (\alpha - 1)\mathbf{y}, \mathbf{y} \rangle}{\|\mathbf{x} + (\alpha - 1)\mathbf{y}\| \|\mathbf{y}\|} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{x}\| \|\mathbf{y}\|} \\
&= \frac{\langle \mathbf{x} + (\alpha - 1)\mathbf{y}, (\alpha - 1)\mathbf{y} \rangle}{\|\mathbf{x} + (\alpha - 1)\mathbf{y}\| \|(\alpha - 1)\mathbf{y}\|} - \frac{\langle \mathbf{x}, (\alpha - 1)\mathbf{y} \rangle}{\|\mathbf{x}\| \|(\alpha - 1)\mathbf{y}\|} \\
&= \frac{\langle \mathbf{x}, \hat{\mathbf{y}} \rangle + \|\hat{\mathbf{y}}\|^2}{\|\mathbf{x} + \hat{\mathbf{y}}\| \|\hat{\mathbf{y}}\|} - \frac{\langle \mathbf{x}, \hat{\mathbf{y}} \rangle}{\|\mathbf{x}\| \|\hat{\mathbf{y}}\|} \\
&> \frac{-\|\mathbf{x}\| + \|\hat{\mathbf{y}}\|}{\|\mathbf{x} + \hat{\mathbf{y}}\|} + 1 \\
&= \frac{\|\mathbf{x} + \hat{\mathbf{y}}\| + \|\hat{\mathbf{y}}\| - \|\mathbf{x}\|}{\|\mathbf{x} + \hat{\mathbf{y}}\|} \\
&> 0,
\end{aligned} \tag{13}$$

which concludes the proof.

C Choice of Reinforced Layers

Since Chen et al., 2025 and Zhang et al., 2024 have highlighted the unique role of the second layer in integrating visual information through attention map analysis, we include this layer for VHR and further validated its significance in mitigating hallucinations. Table 5 presents the results of ablation experiments conducted on this layer and the last few layers on LLaVA-1.5. The results show that enhancing layer1 and the deeper layers both significantly alleviate hallucinations, with the combination of both yielding even better results. This suggests that VHR in the shallow and deep layers alleviates hallucinations through distinct mechanisms. However, we note that applying VHR to the second layer does not universally improve performance; in some challenging benchmarks, its benefits diminish or even introduce instability. Further analysis of the layer-specific mechanisms in LVLMs remains an important direction for future work.

D Choice of Scale Factor

Table 6 shows the results of ablation experiments on the scale factor α in VHR. When α is set to 2 or 3, hallucinations are effectively alleviated. However, as α increases to 4, excessive intervention disrupts the model’s behavior, causing anomalies in the hallucination metric. Conversely, when $\alpha < 1$, which weakens the contribution of attention heads sensitive to visual information, hallucinations become significantly more pronounced. This further confirms the crucial role of the attention heads identified based on the VHD scores in mitigating hallucinations in LVLMs.

	CHAIR _S ↓	CHAIR _I ↓
LLaVA-1.5	49.68	14.32
w/ VHR on layer1	40.80	12.00
w/ VHR on last 14 layers	41.96	12.56
w/ VHR on both	33.32	9.71

Table 5: Ablation study on the reinforced layers.

	CHAIR _S ↓	CHAIR _I ↓	Len
LLaVA-1.5	49.68	14.32	83.06
w/ VHR $\alpha = 0.2$	63.28	21.04	86.42
w/ VHR $\alpha = 0.5$	55.80	17.32	84.47
w/ VHR $\alpha = 2$	33.32	9.71	81.33
w/ VHR $\alpha = 3$	27.04	8.68	88.31
w/ VHR $\alpha = 4$	3.64	2.01	144.54

Table 6: Ablation study on the scale factor.

E Examples of VHD Scores

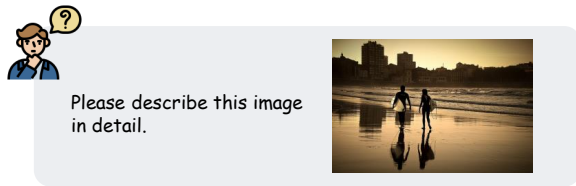
As shown in Figure 10, the first row presents the VHD scores during the first generation step across different samples, while the second row shows the VHD scores for different object terms generated within the same sample. It can be observed that VHD scores vary across different samples and generation steps; however, significant differences between the VHD scores of the attention heads within the model are consistently present.

F More Qualitative Results

As shown in Figure 9, we present more examples that illustrate the effect of VHR in eliminating hallucinated objects. After incorporating VHR, the descriptions generated by the LVLMs faithfully align with the content of the images, while preserving the richness of the descriptions.

G Details on the GPT-4V Evaluation

To evaluate the performance of LVLMs on LLaVA-Bench (In-the-Wild), we use GPT-4o as the evaluator. The prompt template adapted from (Gong et al., 2024) is shown in Table 7, with an additional metric, *Naturalness*, introduced to assess the fluency and coherence of the generated language. For each sample, GPT-4o is provided with the original image, the baseline LVLM output, and the output from the VHR-enhanced model. The evaluation focuses on three key aspects: accuracy, detailedness, and naturalness, with particular emphasis on the reduction of hallucinations in the VHR-enhanced responses compared to the baseline.



Please describe this image in detail.



Please describe this image in detail.

The image captures a beautiful beach scene with two people walking along the shoreline. Both individuals are carrying surfboards, likely preparing to enjoy a day of surfing. The surfboards are visible in their hands, and the people are positioned close to the water. **In the background, there are several cars parked near the beach, indicating that this is a popular spot for visitors. The cars are scattered around the area, with some closer to the water and others further back.**

The image features a red and yellow train traveling down the tracks. The train is quite large, occupying a significant portion of the scene. The train is passing under a bridge, which can be seen in the background. **There are several people visible in the image, likely passengers or bystanders, scattered around the train and the surrounding area. Some of them are closer to the train, while others are further away, possibly waiting for the train to come to a stop.**

The image captures a beautiful sunset over a beach, with two people walking along the water's edge. Both individuals are holding surfboards, likely preparing for a surfing session. The surfboards are positioned horizontally, with one person's surfboard on the left side and the other person's surfboard on the right side. In the background, there are several buildings, including a tall one on the left side and another one on the right side. The combination of the sunset, the beach, and the surfboards creates a serene and picturesque scene.

The image features a red and gray train traveling down the tracks, with a large red train car positioned in the middle of the scene. The train is moving through a city, with buildings visible in the background. The train appears to be a DB Schenker train, which is a German rail company. The train is traveling on a steel track, providing a sense of speed and movement.

Figure 9: More examples of VHR eliminating the hallucinated content.




Figure 10: More examples of the VHD scores during different generation steps of different samples.

GPT-4V Prompt

You are required to score the performance of two AI assistants in describing a given image. You should pay extra attention to the hallucination, which refers to the part of descriptions that are inconsistent with the image content, such as claiming the existence of something not present in the image or describing incorrectly in terms of the counts, positions, or colors of objects in the image. Please rate the responses of the assistants on a scale of 1 to 10, where a higher score indicates better performance, according to the following criteria:

1: Accuracy: whether the response is accurate with respect to the image content. Responses with fewer hallucinations should be given higher scores.

2: Detailedness: whether the response is rich in necessary details. Note that hallucinated descriptions should not count as necessary details.

3: Naturalness: assess the language quality, focusing on: fluency of sentence structure, appropriateness of word choice, smoothness of language flow, absence of awkward or unnatural phrasing.

Please output the scores for each criterion, containing only two values indicating the scores for Assistant 1 and 2, respectively. The two scores are separated by a space. Following the scores, please provide an explanation of your evaluation, avoiding any potential bias and ensuring that the order in which the responses were presented does not affect your judgment.

[Assistant 1]

{ }

[End of Assistant 1]

[Assistant 2]

{ }

[End of Assistant 2]

Output format:

Accuracy:

Reason:

Detailedness:

Reason:

Naturalness:

Reason:

Table 7: The prompt used for GPT-4V evaluation.