



Don't Miss the Forest for the Trees: Attentional Vision Calibration for Large Vision Language Models

Sangmin Woo*

Donguk Kim*

Jaehyuk Jang*

Yubin Choi

Changick Kim

KAIST

{smwoo95, kdu3613, jhyuk, choibinbin, changick}@kaist.ac.kr

Project: <https://sangminwoo.github.io/AvisC/>

Abstract

Large Vision Language Models (LVLMs) demonstrate strong capabilities in visual understanding and description, yet often suffer from hallucinations—attributing incorrect or misleading features to images. We observe that LVLMs disproportionately focus on a small subset of image tokens—termed *blind tokens*—which are typically irrelevant to the query (e.g., background or non-object regions). We hypothesize that such attention misalignment plays a key role in generating hallucinated responses. To mitigate this issue, we propose **Attentional Vision Calibration (AVISC)**, a test-time approach that dynamically recalibrates the influence of blind tokens without modifying the underlying attention mechanism. AVISC first identifies blind tokens by analyzing layer-wise attention distributions over image tokens, then employs a contrastive decoding strategy to balance the influence of original and blind-token-biased logits. Experiments on standard benchmarks, including POPE, MME, and AMBER, demonstrate that AVISC effectively reduces hallucinations in LVLMs.

1 Introduction

Large Vision Language Models (LVLMs) (Dai et al., 2024; Zhu et al., 2023; Liu et al., 2023c,b; Bai et al., 2023; Tong et al., 2024a) have recently shown remarkable capabilities in generating coherent and contextually relevant descriptions of visual inputs. Yet, these models are prone to "hallucinations", producing responses that do not accurately reflect the underlying image. Such hallucinations pose a critical challenge for applications, demanding reliability, precision, and trustworthy visual interpretation.

In this work, we hypothesize that a key contributing factor to hallucinations in LVLMs is an attention misalignment during inference. Specifically,

we observe that LVLMs often allocate excessive attention to a small subset of image tokens—referred to as *blind tokens*—which appear to be uninformative or irrelevant to the query (e.g., background or repetitive regions; see Fig. 1). Our preliminary experiments support this observation (Fig. 2): despite receiving high attention weights, these tokens do not typically carry the query-relevant information necessary for precise interpretation.

We posit that this skewed focus may cause LVLMs to rely disproportionately on such blind tokens during text generation, potentially leading to hallucinated or inaccurate responses. Rather than grounding their outputs in the most semantically relevant visual details, the models may be influenced by peripheral or misleading cues, undermining the factual integrity of their descriptions.

To investigate this hypothesis, we propose AVISC, a novel, training-free decoding method that dynamically recalibrates the influence of blind tokens at inference time. AVISC operates in two key stages: first, it analyzes the layer-wise attention distributions to identify image tokens that receive excessive attention and flags them as blind tokens; then, it applies a contrastive decoding strategy (Leng et al., 2023; Favero et al., 2024). By comparing the prediction logits generated using all image tokens with those obtained when blind tokens are selectively zeroed out, AVISC adjusts the final token probabilities to mitigate the undue influence of irrelevant tokens.

Notably, AVISC does not directly manipulate model's attention mechanism or require any retraining. Instead, it operates solely at the decoding stage by modifying the final token probabilities, reducing the influence of blind tokens while amplifying the impact of non-blind tokens.

We validate our approach on multiple benchmarks—including POPE (Rohrbach et al., 2018), MME (Fu et al., 2024), and AMBER (Wang et al., 2023)—demonstrating that AvisC not

*Equal contribution

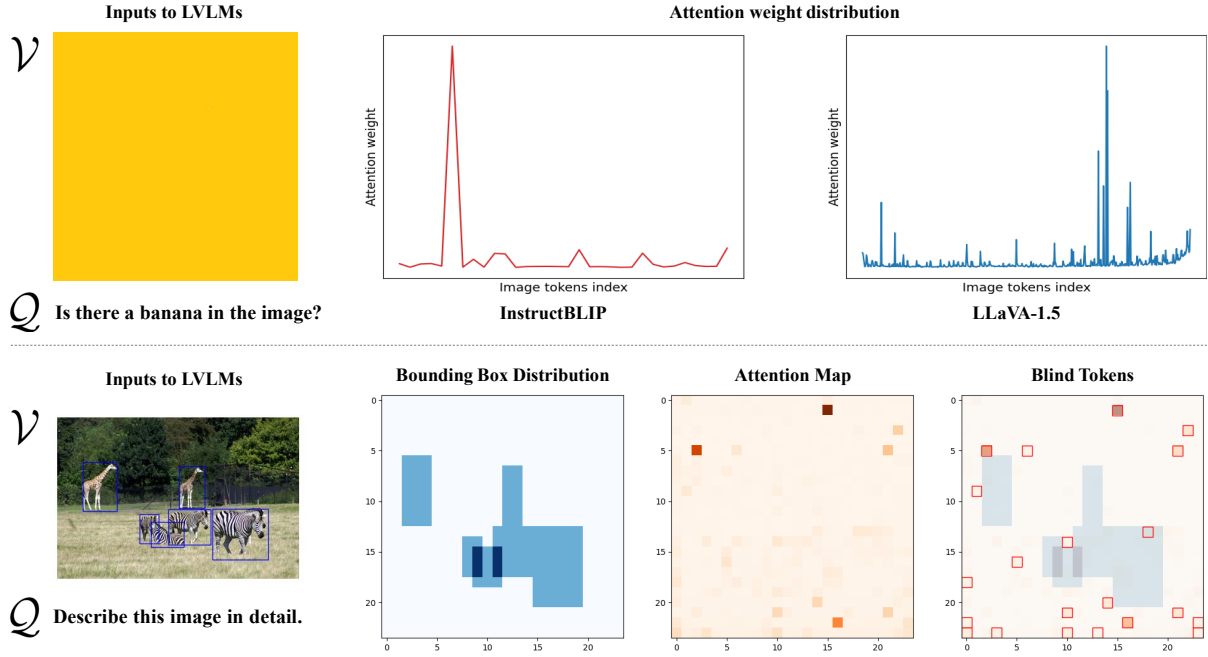


Figure 1: **Blind tokens in LVLMs.** (Top) Even when the image (‘V’) lacks meaningful content for the textual query (‘Q’), modern LVLMs (Dai et al., 2024; Liu et al., 2023c) still assign disproportionate attention to a few image tokens (*i.e.*, *blind tokens*). Despite having identical, featureless yellow patches, these tokens dominate the attention distribution. (Bottom) In a real image, overlaying bounding boxes and LLaVA 1.5’s attention map highlights a clear mismatch between blind tokens (red boxes) and genuinely informative regions. *Note*: attention weights are averaged across all layers for the first generated token. See Appendix A for more examples.

only significantly reduces hallucinations but also enhances both the factual accuracy and descriptive richness of LVLM outputs. Importantly, our method is model-agnostic and requires no additional training data or external modules, making it a practical plug-and-play solution for improving existing LVLM systems.

In summary, our contribution is threefold: (1) We uncover and analyze the phenomenon of *blind tokens*—image tokens with disproportionately high attention that contribute to hallucinations in LVLMs. (2) We introduce AVISC, a training-free approach that dynamically recalibrates the influence of blind tokens without modifying the underlying model; and (3) We demonstrate, through comprehensive experiments on standard benchmarks, that AvisC effectively reduces hallucinations and enhances the performance of diverse LVLMs.

2 Observations

Modern LVLMs (Dai et al., 2024; Liu et al., 2023c) build upon the transformer architecture (Vaswani et al., 2017), where attention weights are intended to highlight the most relevant tokens for generating the next output token. Intuitively, tokens receiving higher attention should correspond to key elements

in the image—an idea that has proven effective in both vision- and text-based transformers (Caron et al., 2021; Ilharco et al., 2021; Vaswani et al., 2017). However, we find that this principle does not always hold in current LVLMs.

Blind tokens in uniform images. A striking illustration of this issue arises when an image has no meaningful content for the query—such as a uniformly colored background. As shown in Fig. 1 (top), even in a plain yellow image with no discernible objects, LVLMs often concentrate most of their attention on a few patches. We refer to these excessively attended yet semantically uninformative patches as *blind tokens*. This phenomenon echoes findings in vision transformers (Darcet et al., 2023)¹, where certain background regions disproportionately attract high attention, possibly serving as global information “reservoirs” at the expense of local, detail-rich areas.

Mismatch between blind tokens and actual objects. Beyond artificially simple images, we also analyze real images from COCO2014 (Lin et al., 2014). We ask LVLMs to describe the given image

¹We discuss our findings in relation to (Darcet et al., 2023) in Appendix D.

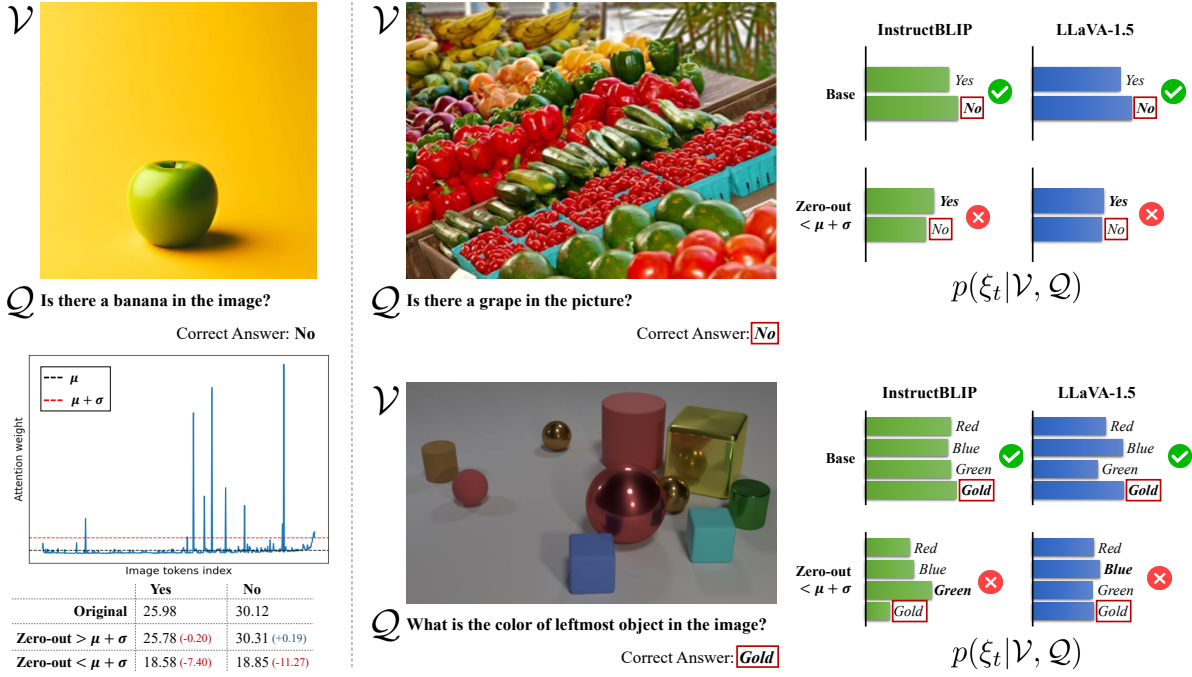


Figure 2: **Blind tokens contribute little to actual predictions.** (a) We perform zero-out experiments to measure the impact of blind vs. non-blind tokens. Zeroing out blind tokens ($\text{Zero-out} > \mu + \sigma$), where attention weights are above mean + standard deviation, leaves the model’s predicted probabilities nearly unchanged, suggesting that these tokens carry minimal object-discriminative information. In contrast, zeroing out *non-blind tokens* yields near 50:50 probabilities, underscoring their critical role in correct prediction. (b) When non-blind tokens are zeroed out, the models fails to correctly predict previously well-classified instances.

and measure how much attention goes to patches corresponding to object bounding boxes. As shown in Fig. 1 (bottom), many tokens receiving disproportionately high attention (blind tokens) have little overlap with genuine object regions, while actual objects receive comparatively less attention. Specifically, only 3.7% of blind tokens overlap with these regions, and merely 23.2% of total attention weight is goes to them.² This mismatch indicates that, despite carrying little or no query-relevant information, blind tokens consume a large share of the model’s attention. Consequently, truly informative tokens that capture critical visual details are under-emphasized, potentially compromising the model’s descriptive quality and reliability.

Zero-out experiments. To better understand the functional role of blind tokens, we conduct a zero-out analysis on LLaVA-1.5 (Liu et al., 2023c), shown in Fig. 2. Specifically, we either zero out the blind tokens or the non-blind tokens, then observe changes in the model’s predicted logits. When we remove blind tokens, the logits remain almost identical to those of the original model—indicating that these tokens contribute little to the final prediction.

By contrast, removing non-blind tokens causes the logits to collapse to near-uniform probabilities, revealing that the essential, object-discriminative information resides in those less-attended tokens.³

Attention bias and hallucinations. These findings suggest that LVLMs systematically overemphasize certain patches that do not meaningfully aid the prediction process. Consequently, truly informative tokens—often corresponding to the actual objects or key details—receive insufficient attention. We hypothesize that this imbalance predisposes the model to hallucinate, as the generation process leans on blind tokens that fail to encode crucial visual details. In the next sections, we propose a simple yet effective decoding method to mitigate this problem by recalibrating the model’s attention usage at decoding stage, thereby reducing its reliance on blind tokens and improving visual grounding.

Hypothesis. Our hypothesis is that blind tokens arises as a structural byproduct of the deep, layered architecture of these models, similar to the “high-norm outlier tokens” observed in vision trans-

²Detailed statistics can be found in Fig. 9.

³For dataset-level zero-out experiment details, see Tab. 7.

formers (Darcet et al., 2023) (see Appendix D). As information is propagated through layers, global representations from earlier layers are progressively compressed. However, instead of being allocated to semantically meaningful tokens, this global information often becomes concentrated in structurally convenient but semantically irrelevant tokens—frequently those in repetitive or low-information regions. These tokens consequently accumulate disproportionate representational weight and attract excessive attention during decoding. Despite their lack of semantic relevance, they misguide the model’s focus and contribute to the generation of hallucinated content. While establishing a definitive causal link is inherently challenging, our qualitative and quantitative evidence suggest that blind tokens are a recurring and impactful phenomenon in LVLMs. Visualizations and token distribution analyses (see Figs. 6 to 11) demonstrate that blind tokens frequently emerge in spatially uninformative regions and exhibit anomalously high attention scores.

3 Approach: AVISC

We propose AVISC, a test-time approach to enhance visual object understanding in LVLMs during decoding. AVISC dynamically recalibrates the model’s attention at every token generation step by reducing the over-emphasis on *blind tokens*—image tokens that receive disproportionate attention despite lacking task-relevant information. An overview of AVISC is shown in Fig. 3. Our approach modifies the decoding process in three key steps: (1) **Layer selection**: identify layers that exhibit a high proportion of image-related attention; (2) **Blind token identification**: detect tokens that capture an unusually high share of attention; and (3) **Contrastive decoding**: adjust output logits to mitigate the influence of these blind tokens.

3.1 LVLM Framework

Uni-modal encoding. LVLMs process visual inputs and textual queries into compact representations. Pre-trained encoders like CLIP (Radford et al., 2021) are commonly used for processing visual data. The text data is tokenized, turning it into a sequence of manageable pieces for further processing.

Cross-modal alignment. Since LLMs natively process only text, a learnable cross-modal alignment module—such as the Q-Former (Li et al., 2023a) or a linear projection layer (Liu et al.,

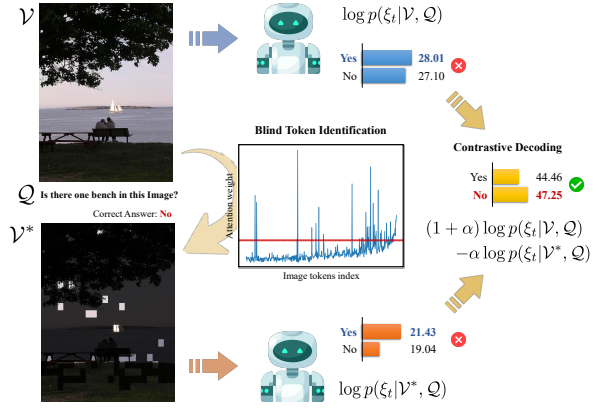


Figure 3: An overview of AVISC.

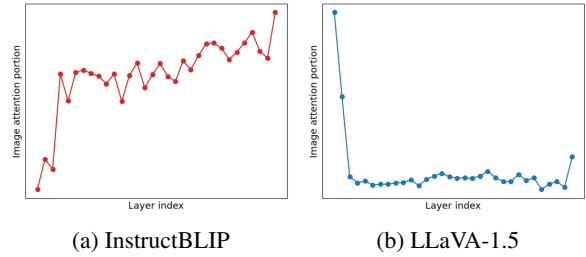


Figure 4: **Layer-wise image attention proportion in LVLMs** (Liu et al., 2023b; Dai et al., 2024). This shows the proportion of attention given to image tokens at each layer relative to total attention. Different layers exhibit distinct attention patterns, which vary across models. Attention weights are averaged over 60 questions from the LLaVA-Bench (Liu et al., 2023c).

2023c)—transform visual features into tokens compatible with the LLM’s input space. This yields a set of visual tokens, $\mathcal{V} = \{v_0, v_1, \dots, v_{N-1}\}$, which are then concatenated with text tokens, $Q = \{\sigma_N, \sigma_{N+1}, \dots, \sigma_{N+M-1}\}$, forming a unified input sequence of length $N + M$.

Next token prediction via LLM. The concatenated token sequence is processed by LVLM (parametrized by θ) in an auto-regressive manner. The model computes logits ℓ_t for each potential next token:

$$\ell_t = \log p(\xi_t | \mathcal{V}, Q, \xi_{<t}), \quad (1)$$

where ξ_t is the token predicted at timestep t , and $\xi_{<t}$ denotes the sequence generated up to timestep $(t - 1)$. These logits are converted into a probability distribution via the softmax function:

$$p(\xi_t) = \text{Softmax}(\ell_t). \quad (2)$$

from which the next token is sampled.

3.2 Attentional Vision Calibration

Visual hallucinations in LVLMs often arise during decoding when the model’s token selection

is guided by skewed probability distributions that do not faithfully reflect the underlying visual input. Our observations (see Sec. 2) indicate that this problem is linked to an attentional bias toward certain non-relevant tokens, which we term *blind tokens*. Our approach aims to recalibrate these attention patterns to mitigate hallucinations.

Layer selection. Different layers in LVLMs contribute variably to processing visual information. As illustrated in Fig. 4, models such as InstructBLIP (Dai et al., 2024) and LLaVA-1.5 (Liu et al., 2023c) exhibit different attention distributions across layers. To accommodate these differences, we first select layers that exhibit a high proportion of attention on image tokens. Formally, for the i th layer, we define the attention weight matrix:

$$\mathbf{A}_i = \left[\mathbf{a}_{h,q,k}^i \right]_{(h,q,k)=(1,1,1)}^{(H,N+M,N+M)}, \quad (3)$$

where $\mathbf{a}_{h,q,k}^i$ represents the attention weight from head h , for query q to key l_k in layer i . With image tokens $\mathcal{V} \in \mathbb{R}^{N \times D}$ and query tokens $\mathcal{Q} \in \mathbb{R}^{M \times D}$, we compute the proportion of attention dedicated to image tokens in layer i :

$$AP_i^{\text{layer}} = \frac{\sum_h \sum_{k=1}^N \mathbf{a}_{h,(N+M),k}^i}{\sum_{i,h} \sum_{k=1}^N \mathbf{a}_{h,(N+M),k}^i}, \quad (4)$$

where H is the total number of attention heads, N is the number of image tokens, and M is the number of query tokens. We then sort the layers by this proportion and select layers using top-P sampling with threshold γ :

$$\{\text{Selected Layers}\} = \text{top-P}(\{AP_i^{\text{layer}}\}_{i=1}^L, \gamma). \quad (5)$$

These selected layers provide the basis for further token-level analysis.

Blind token identification. Within the selected layers, we compute the attention proportion for each image token by averaging over all heads:

$$AP^{\text{image}} = \frac{\sum_{i \in \{\text{Selected Layers}\}} \sum_{h=1}^H \mathbf{a}_{h,(N+M),[1:N]}^i}{|\{\text{Selected Layers}\}| \times H}. \quad (6)$$

To identify tokens that disproportionately capture attention, *i.e.*, *blind tokens*, we calculate the mean (μ) and standard deviation (σ) of the image attention weights. Tokens with an attention proportion exceeding $\mu + \lambda\sigma$ (where λ is a hyperparameter) are classified as blind tokens:

$$\{\text{Blind Token Indices}\} = \{j | AP_j^{\text{image}} > \mu + \lambda\sigma\}. \quad (7)$$

Contrastive decoding. To mitigate hallucinations, we reduce the influence of blind tokens during decoding via contrastive decoding (Leng et al., 2023; Favero et al., 2024). We construct a biased set of visual tokens by zeroing out non-blind tokens:

$$\mathcal{V}^* = \bigcup_{j=1}^N \mathbb{1}_{\{j \in \text{Blind Token Indices}\}}(j) \mathbf{v}_j. \quad (8)$$

We then compute the logits for the next token using both the original visual tokens (\mathcal{V}) and the biased tokens (\mathcal{V}^*):

$$\begin{aligned} \ell_t &= \log p(\xi_t | \mathcal{V}, \mathcal{Q}, \xi_{<t}; \theta), \\ \ell_t^* &= \log p(\xi_t | \mathcal{V}^*, \mathcal{Q}, \xi_{<t}; \theta), \end{aligned} \quad (9)$$

Finally, we adjust the logits using a contrastive scheme and sample the next token from:

$$\xi_t \sim \text{Softmax}((1 + \alpha)\ell_t - \alpha\ell_t^*). \quad (10)$$

Here, α controls the strength of the contrast. This adjustment effectively down-weights the contribution of blind tokens and promotes a more balanced attention distribution, thereby reducing hallucinations in the final output.

4 Experiments

Additional experimental results are provided in Appendix C.

4.1 Evaluation Setup

We deliberately avoid constraining LVLMs to provide one-word responses (*e.g.*, “Yes” or “No”) for discriminative tasks, as our analysis (see Tab. 13) shows that such constraints can significantly alter performance. For our experiments, we set $P = 0.5$ in Eq. (5), $\lambda = 1$ in Eq. (7), and $\alpha = 3$ for InstructBLIP (Dai et al., 2024) and $\alpha = 2.5$ for LLaVA-1.5 (Liu et al., 2023c) in Eq. (10).⁴

LVLMs. We evaluate AVISC on two state-of-the-art LVLMs: **InstructBLIP** (Dai et al., 2024) and **LLaVA-1.5** (Liu et al., 2023c), both of which use Vicuna 7B (Chiang et al., 2023) as the LLM backbone. InstructBLIP employs the Q-Former (Li et al., 2023a) to efficiently fuse visual and textual features using a fixed number of tokens (*e.g.*, 32 tokens), while LLaVA-1.5 aligns image and text modalities via linear projection layers. Notably, AVISC is model-agnostic and can be integrated into various LVLM architectures.

⁴Further experimental details are provided in Appendix B.

Setup	Method	InstructBLIP (Dai et al., 2024)				LLaVA-1.5 (Liu et al., 2023c)			
		Acc. ↑	Prec. ↑	Rec. ↑	F1 ↑	Acc. ↑	Prec. ↑	Rec. ↑	F1 ↑
MS-COCO	Random	<i>base</i>	82.27	82.84	81.40	82.11	84.47	83.35	84.72
		VCD	83.37	83.39	82.60	83.24	84.80	83.00	85.20
		M3ID	84.37	84.62	84.00	84.31	86.00	85.11	86.18
		AVISC	88.73	93.88	82.87	88.03	87.93	88.24	87.88
	Popular	<i>base</i>	77.77	74.81	83.73	79.02	82.23	79.72	82.95
		VCD	78.00	75.12	83.73	79.19	82.27	79.19	87.53
		M3ID	77.30	74.10	83.93	78.71	82.83	79.62	83.72
		AVISC	83.90	81.33	88.00	84.53	84.33	81.71	88.47
	Adversarial	<i>base</i>	73.13	69.41	82.60	75.46	77.10	72.57	87.13
		VCD	75.87	72.85	82.47	77.36	76.10	71.50	86.80
		M3ID	76.03	72.47	83.93	77.79	77.70	73.23	87.33
		AVISC	81.57	80.37	83.53	81.92	77.53	72.82	87.87
A-OKVQA	Random	<i>base</i>	81.00	77.71	86.93	82.06	82.73	77.43	92.40
		VCD	81.73	78.67	87.07	82.66	81.30	75.45	92.80
		M3ID	82.33	77.81	90.47	83.66	83.57	77.86	93.80
		AVISC	88.47	87.66	89.53	88.59	84.60	79.29	93.67
	Popular	<i>base</i>	75.00	70.14	87.07	77.69	76.10	69.86	91.80
		VCD	75.33	70.52	87.07	77.92	75.43	68.58	93.87
		M3ID	75.60	70.40	88.33	78.36	76.80	70.20	93.13
		AVISC	81.77	77.82	88.87	82.98	78.83	72.10	94.07
	Adversarial	<i>base</i>	68.80	63.57	88.07	73.84	67.90	62.11	91.80
		VCD	69.70	64.54	87.47	74.27	67.43	61.50	93.20
		M3ID	69.57	64.21	88.40	74.39	68.10	61.99	93.60
		AVISC	72.53	67.12	88.33	76.28	68.97	62.70	93.67
GQA	Random	<i>base</i>	80.00	77.08	85.40	81.02	82.40	77.03	92.33
		VCD	81.73	79.35	85.80	82.45	82.27	75.85	94.67
		M3ID	80.57	76.77	87.67	81.85	82.83	76.64	94.47
		AVISC	86.47	85.89	87.27	86.57	85.00	78.81	95.73
	Popular	<i>base</i>	73.53	68.80	86.13	76.49	72.03	65.57	92.80
		VCD	74.10	69.45	86.07	76.87	71.77	64.90	94.80
		M3ID	74.57	69.45	87.83	77.53	72.83	66.04	94.00
		AVISC	78.00	73.68	87.13	79.84	74.80	67.46	95.80
	Adversarial	<i>base</i>	68.00	63.49	84.73	72.59	68.73	62.54	93.40
		VCD	70.27	65.43	85.93	74.29	68.27	62.00	94.40
		M3ID	68.90	64.06	86.13	73.47	68.13	61.88	94.47
		AVISC	73.07	67.80	87.87	76.54	69.20	62.61	95.33

Table 1: **POPE benchmark results.** AVISC consistently outperforms *base* decoding and other methods: VCD (Leng et al., 2023) and M3ID (Favero et al., 2024). We reimplemented VCD and M3ID in our evaluation setup.

Benchmarks. (1) **POPE** (Li et al., 2023b) views hallucination evaluation as a binary classification task (yes/no) with questions on object presence (e.g., "Is there a cat in the image?"). It evaluates both visible and imaginary objects across across three setups: random, popular, and adversarial. (2) **MME** (Fu et al., 2024) evaluates 14 subtasks—including object existence, count, position, and color—via binary questions. (3) **AMBER** (Wang et al., 2023) combines generative and discriminative tasks, focusing on hallucinations in object existence, attributes, and relationships. Generative performance is measured by CHAIR, discriminative by F1, with the overall AMBER score computed as $((100 - \text{CHAIR}) + \text{F1})/2$.

Baselines. We compare AVISC against recent contrastive decoding methods, notably **VCD** (Leng

et al., 2023) and **M3ID** (Favero et al., 2024), which mitigate hallucinations by enhancing the reference image’s influence relative to the language model’s prior through contrasting output distributions from original and altered visual inputs. Our reimplementations of VCD and M3ID serve as baselines, as they too avoid external models, costly self-feedback, and additional training.

4.2 Benchmark Results

POPE. Table 1 summarizes performance on the POPE benchmark (Li et al., 2023b) across MS-COCO (Lin et al., 2014), A-OKVQA (Schwenk et al., 2022), and GQA (Hudson and Manning, 2019) datasets under the Random, Popular, and Adversarial setups. Overall, AVISC consistently outperforms the baseline (*base*) and other decoding methods (VCD, M3ID) in most cases, achieving the

Model	Method	Object-level		Attribute-level		Total Score
		Existence \uparrow	Count \uparrow	Position \uparrow	Color \uparrow	
InstructBLIP	<i>base</i>	170.19(± 11.12)	89.52(± 11.04)	67.62(± 14.04)	114.76(± 9.60)	442.09(± 31.51)
	VCD	172.62(± 8.92)	98.33(± 15.99)	71.90(± 13.42)	117.14(± 10.70)	459.99(± 16.56)
	M3ID	173.89(± 10.52)	89.72(± 13.44)	72.72(± 14.77)	110.56(± 7.20)	446.88(± 28.54)
	AVISC	184.76(± 5.56)	82.85(± 12.16)	74.76(± 6.19)	131.43(± 4.76)	473.80(± 19.67)
LLaVA 1.5	<i>base</i>	173.57(± 8.16)	110.00(± 15.82)	100.47(± 18.78)	125.24(± 15.91)	509.28(± 30.57)
	VCD	172.14(± 8.09)	117.14(± 8.76)	103.33(± 20.56)	119.52(± 8.58)	512.14(± 31.82)
	M3ID	178.33(± 6.83)	107.22(± 14.78)	96.39(± 5.52)	127.50(± 8.28)	509.44(± 22.52)
	AVISC	189.29(± 1.82)	104.76(± 11.66)	106.19(± 13.93)	127.86(± 9.13)	528.09(± 24.70)

Table 2: **MME-Hallucination (Fu et al., 2024) benchmark results.** Our method effectively reduces hallucinations at both object and attribute levels, surpassing VCD (Leng et al., 2023) and M3ID (Favero et al., 2024) in Total Score.

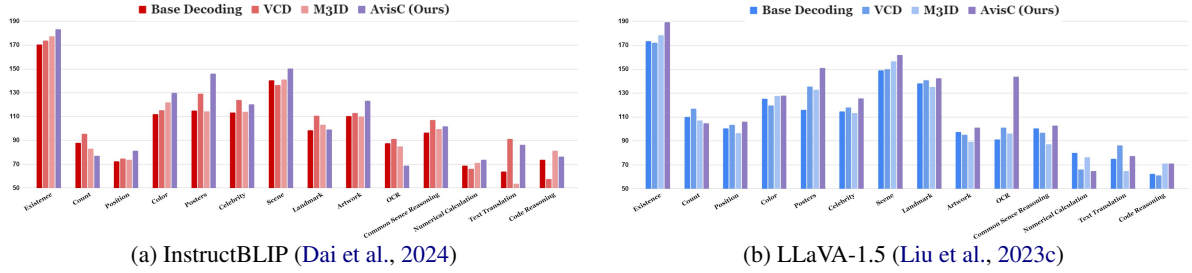


Figure 5: **Performance comparison on MME-Fullset.** AVISC achieves top performance in 7 of 14 categories with InstructBLIP (Dai et al., 2024) and in 11 categories with LLaVA-1.5 (Liu et al., 2023c). Beyond minimizing hallucinations, AVISC also boosts the general functionality of LVLMs.

highest Accuracy and F1 scores. Additionally, balanced improvements in Precision and Recall suggest reduced errors and better information capture. For InstructBLIP, AVISC yields a significant performance boost—especially in mitigating hallucinations related to object existence—while LLaVA-1.5 shows somewhat lower gains in the more challenging Popular and Adversarial setups. Nonetheless, AVISC proves robust across different datasets and query configurations.

MME-Hallucination. Table 2 presents performance results for InstructBLIP and LLaVA-1.5 on the MME-Hallucination benchmark (Fu et al., 2024). We evaluate both object-level metrics (Existence, Count) and attribute-level metrics (Position, Color). Both models show marked improvements in the Existence category when using AVISC, achieving the highest scores. While VCD slightly outperforms in the Count metric, AVISC excels in Position and Color, leading to superior Total Scores overall. These results affirm that AVISC effectively reduces hallucinations and improves accuracy across multiple dimensions.

MME-Fullset. Figure 5 compares various decoding methods on the MME-Fullset benchmark (Fu et al., 2024) across 14 categories. AVISC achieves

top results in 7 categories for InstructBLIP and 11 for LLaVA-1.5. This indicates that AVISC enhances the model’s ability to extract and utilize informative visual features through attention calibration. Although both models experience a slight decline in the Count category with AVISC—and InstructBLIP shows lower performance on OCR tasks—LLaVA-1.5 sees significant OCR improvements, demonstrating that the impact of AVISC can vary across different models. Overall, AVISC delivers superior results across most tasks compared to the baselines.

AMBER. Table 3 shows results on the AMBER benchmark (Wang et al., 2023), which includes both generative and discriminative tasks. AVISC significantly improves discriminative performance (Accuracy and F1) for both InstructBLIP and LLaVA-1.5, outperforming Base, VCD, and M3ID. In generative tasks, it also achieves substantial gains, particularly in the Existence metric, indicating better object detection. Overall, AVISC enables both models to achieve the highest scores across most AMBER metrics.

4.3 Ablation Study

Ablations on α and λ . In our approach, λ is the threshold for identifying blind tokens that receive

Metric		InstructBLIP (Dai et al., 2024)				LLaVA 1.5 (Liu et al., 2023c)			
		base	VCD	M3ID	AVIS	base	VCD	M3ID	AVIS
Generative	CHAIR↓	8.40(±0.57)	7.60(±0.42)	6.85(±0.07)	6.70(±0.28)	7.95(±0.64)	6.70(±0.42)	6.00(±0.14)	6.25(±0.07)
	Cover↑	46.40(±1.27)	47.65(±0.35)	47.20(±0.71)	46.65(±1.48)	44.45(±0.21)	46.50(±0.28)	48.90(±0.28)	46.55(±0.64)
	Hal↓	31.10(±0.64)	29.90(±0.99)	27.50(±0.71)	28.00(±0.28)	31.00(±2.83)	27.80(±1.70)	26.00(±0.28)	25.60(±1.70)
	Cog↓	2.60(±0.05)	2.20(±0.14)	2.20(±0.14)	2.55(±0.35)	2.15(±0.35)	1.95(±0.35)	1.45(±0.07)	2.00(±0.04)
Discriminative	Acc.↑	68.20(±0.14)	69.65(±0.35)	69.05(±0.35)	72.60(±0.42)	67.00(±0.71)	67.30(±1.41)	67.25(±0.21)	70.70(±0.57)
	Prec.↑	79.00(±0.14)	80.70(±0.42)	79.70(±0.28)	72.60(±0.42)	85.45(±0.49)	86.10(±1.70)	86.50(±0.57)	85.45(±0.21)
	Rec.↑	70.70(±0.42)	71.60(±0.42)	71.25(±0.35)	76.10(±0.05)	60.95(±1.20)	60.55(±1.34)	60.05(±0.07)	67.55(±0.92)
	F1↑	74.60(±0.14)	75.90(±0.42)	75.25(±0.07)	78.60(±0.28)	71.10(±0.99)	71.10(±1.56)	70.90(±0.14)	75.45(±0.64)
	AMBER↑	83.10(±0.35)	84.15(±0.05)	84.20(±0.07)	85.95(±0.05)	81.58(±0.18)	82.20(±0.99)	82.45(±0.14)	84.60(±0.35)

Table 3: **AMBER (Wang et al., 2023) benchmark results.** AVISC outperforms contrastive decoding baselines (Leng et al., 2023; Favero et al., 2024) in both generative and discriminative tasks, achieving the highest AMBER score.

(a) InstructBLIP (Dai et al., 2024) ($\lambda = 1$)						(b) InstructBLIP (Dai et al., 2024) ($\alpha = 3$)						(c) LLaVA-1.5 (Liu et al., 2023b) ($\lambda = 1$)					
α	Object		Attribute		Total Score	λ	Object		Attribute		Total Score	α	Object		Attribute		Total Score
	Exist.	Count	Position	Color			Exist.	Count	Position	Color			Exist.	Count	Position	Color	
0.5	180	83.33	80.00	130	473.33	0.0	180	75.00	60.00	115.00	430.00	0.5	185	111.66	103.33	115.00	514.99
2.0	180	86.66	75.00	135	476.66	0.1	185	60.00	65.00	123.33	433.33	2.0	180	103.33	101.66	120.00	504.99
2.5	180	85.00	71.66	135	471.66	1.0	195	75.00	73.33	135.00	478.33	2.5	180	105.00	111.66	120.00	516.66
3.0	195	75.00	73.33	135	478.33	1.5	195	75.00	73.33	135.00	478.33	3.0	180	105.00	111.66	120.00	516.66

Table 4: α and λ ablations on MME-Hallucination (Fu et al., 2024). We set $\alpha = 3$, $\lambda = 1$ for InstructBLIP (Dai et al., 2024) and $\alpha = 2.5$, $\lambda = 1$ for LLaVA-1.5 (Liu et al., 2023b).

excessive attention (see Eq. (7)), and α controls the strength of contrastive decoding (see Eq. (10)). We conducted ablation experiments on the MME-Hallucination benchmark (Liu et al., 2023d) to study their effects. Tab. 4 (a) and (c) show results using InstructBLIP (Dai et al., 2024) and LLaVA-1.5 (Liu et al., 2023c), respectively, with λ fixed at 1 and α varied from 0.5 to 3. Overall, performance consistently improves with higher α values, with InstructBLIP achieving the highest total score at $\alpha=3$ and LLaVA-1.5 at $\alpha=2.5$. These findings suggest that a stronger contrastive signal can better mitigate hallucinations. Additionally, Tab. 4 (b) shows that performance for InstructBLIP improves as λ increases, indicating that restricting the application of our method to a smaller set of highly attended tokens yields better results.

Ablations on γ . We further evaluated the sensitivity of our approach to the parameter γ , which determines the cumulative threshold for selecting layers based on image attention (see Eq. (5)). Using LLaVA-1.5 with $\lambda=1.0$ and $\alpha=2.5$, our experiments (shown in Tab. 5) reveal that performance remains robust across a range of γ values, except for extreme settings (e.g., $\gamma=0.1$). Our default value of $\gamma=0.5$ yields high accuracy and balanced metrics on the POPE-COCO-Random benchmark,

and achieves the highest total score on the MME-Hallucination benchmark. Overall, these results indicate that our approach is not highly sensitive to γ , thereby reducing the need for extensive parameter tuning. Consequently, we fixed $\lambda=1.0$ and $\gamma=0.5$ in our experiments.

Layer Selection. We performed an ablation study to evaluate the effectiveness of our targeted layer selection strategy, which selects layers with the highest proportion of image-related attention for blind-token localization and calibration. To assess the contribution of this mechanism, we compared our method with three variants that manually fixed the selected layers (early, middle, or late five layers in the network). As shown in Tab. 6, our approach consistently outperforms all alternatives across both the POPE-COCO-Random and MME-Hallucination benchmarks. These results demonstrate that our targeted layer selection yields tangible improvements, albeit with modest margins, indicating that selecting layers with stronger image-related attention leads to more accurate blind-token localization and reduces hallucination.

5 Related Work

To mitigate hallucinations in LVLMs, researchers have developed strategies across three levels:

(a) POPE-COCO-Random					
γ	Acc. \uparrow	Prec. \uparrow	Rec. \uparrow	F1 \uparrow	
0.5 (Ours)	87.93	88.24	87.53	87.88	
0.1	86.77	83.98	90.87	87.29	
0.3	87.47	85.35	90.47	87.83	
1.0	88.27	88.06	88.53	88.30	

(b) MME-Hallucination					
γ	Existence \uparrow	Count \uparrow	Position \uparrow	Color \uparrow	Total Score \uparrow
0.5 (Ours)	189.29	104.76	106.19	127.86	528.10
0.1	167.50	101.80	103.33	117.50	490.13
0.3	180.00	98.33	114.16	125.00	517.49
1.0	182.50	108.33	109.99	117.50	518.32

Table 5: γ ablations on (a) POPE-COCO-Random and (b) MME-Hallucination benchmarks with LLaVA-1.5 ($\lambda = 1$, $\alpha = 2.5$).

(a) POPE-COCO-Random					
Method	Acc. \uparrow	Prec. \uparrow	Rec. \uparrow	F1 \uparrow	
w/ layer selection (Ours)	87.93	88.24	87.53	87.88	
Early 5 layers	87.50	87.99	86.27	87.12	
Mid 5 layers	87.53	88.20	86.07	87.13	
Last 5 layers	87.46	88.56	85.47	87.00	

(b) MME-Hallucination					
Method	Existence \uparrow	Count \uparrow	Position \uparrow	Color \uparrow	Total Score \uparrow
w/ layer selection (Ours)	189.29	104.76	106.19	127.86	528.10
Early 5 layers	180.00	108.33	106.66	120.00	514.99
Mid 5 layers	180.00	108.33	105.00	120.00	513.33
Last 5 layers	185.00	103.33	105.00	120.00	513.33

Table 6: **Ablation of layer selection** on (a) POPE-COCO-Random and (b) MME-Hallucination benchmarks. Our targeted layer selection outperforms manual alternatives in all key metrics.

Input-level. These methods improve data quality and diversity by incorporating negative (Liu et al., 2023a) and counterfactual data (Yu et al., 2023a) or through dataset cleansing (Wang et al., 2024; Yue et al., 2024), thereby fostering more robust visual-text alignments during training.

Model-level. Approaches at this level enhance visual representations by increasing image processing resolution (Chen et al., 2023; Liu et al., 2023b; Zhai et al., 2023) or by leveraging advanced vision encoders (He et al., 2024; Jain et al., 2023; Tong et al., 2024b). Typically, these methods involve additional training with auxiliary supervision or reinforcement learning (Zhao et al., 2023; Gunjal et al., 2024; Sun et al., 2023; Yu et al., 2023b).

Output-level. Output-level methods directly refine the generated outputs. Contrastive decoding techniques (Leng et al., 2023; Favero et al., 2024) mitigate hallucinations by contrasting outputs from original and modified visual inputs, while guided decoding leverages external models like CLIP (Radford et al., 2021) or DETR (Carion et al., 2020) to steer generation. Other approaches include training-free methods (Wan et al., 2024; Zhang et al., 2024; Huang et al., 2023) and post-hoc corrections (Lee et al., 2023; Wu et al., 2024).

Our work falls within the output-level category.

Unlike prior contrastive decoding methods that contrast whole-image representations, AVISC analyzes the internal attention patterns of LVLMs to identify *blind tokens*—tokens that attract excessive attention but contribute little to the final output—and applies a contrastive decoding strategy to recalibrate their influence.

6 Conclusion

We identify and characterize blind tokens in LVLMs—image tokens that receive excessive attention while conveying little task-relevant information. These tokens misdirect the model’s focus, increasing the likelihood of hallucinated responses. To address this, we propose Attentional Vision Calibration (AVISC), a novel, training-free decoding technique that dynamically detects and mitigates the effect of blind tokens using image-wise attention analysis and contrastive decoding. Extensive evaluations on hallucination benchmarks demonstrate that AVISC improves both visual grounding and response accuracy, surpassing existing decoding strategies.

Limitations

While AVISC reduces hallucinations, its effectiveness declines in tasks requiring precise object counting (e.g., the “Count” category in MME and “Number” in AMBER; see Tabs. 14 and 15). This suggests that blind tokens may sometimes carry essential information for quantification. AVISC adds some overhead due to dynamic test-time recalibration but maintains competitive tokens-per-second throughput compared to other contrastive decoding methods. Unlike high-latency beam search approaches (e.g., OPERA), AVISC offers a better efficiency–accuracy trade-off. Though currently sequential, recalibration can be parallelized to reduce wall-clock time at the expense of memory. Performance on MME and AMBER varies with dataset scope and evaluation protocols, which are sensitive to token usage. Still, AVISC consistently lowers hallucination rates and yields statistically significant gains.

Future work. Building on insights from (Darcet et al., 2023), we hypothesize that blind token phenomenon may be intrinsic to large-scale transformer architectures, not limited to LVLMs. Future research will further explore these blind tokens and develop strategies to address them while balancing computational efficiency.

References

- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Zhong Muyan, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2023. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv preprint arXiv:2312.14238*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. 2024. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Advances in Neural Information Processing Systems*, 36.
- Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. 2023. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*.
- Alessandro Favero, Luca Zancato, Matthew Trager, Siddharth Choudhary, Pramuditha Perera, Alessandro Achille, Ashwin Swaminathan, and Stefano Soatto. 2024. Multi-modal hallucination control by visual information grounding. *arXiv preprint arXiv:2403.14003*.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, Yunsheng Wu, and Rongrong Ji. 2024. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*.
- Anisha Gunjal, Jihan Yin, and Erhan Bas. 2024. Detecting and preventing hallucinations in large vision language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18135–18143.
- Xin He, Longhui Wei, Lingxi Xie, and Qi Tian. 2024. Incorporating visual experts to resolve the information loss in multimodal large language models. *arXiv preprint arXiv:2401.03105*.
- Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming Zhang, and Nenghai Yu. 2023. Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation. *arXiv preprint arXiv:2311.17911*.
- Drew A Hudson and Christopher D Manning. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709.
- Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. 2021. **Openclip**. If you use this software, please cite it as below.
- Jitesh Jain, Jianwei Yang, and Humphrey Shi. 2023. Vcoder: Versatile vision encoders for multimodal large language models. *arXiv preprint arXiv:2312.14233*.
- Seongyun Lee, Sue Hyun Park, Yongrae Jo, and Minjoon Seo. 2023. Volcano: mitigating multimodal hallucination through self-feedback guided revision. *arXiv preprint arXiv:2311.07362*.
- Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. 2023. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. *arXiv preprint arXiv:2311.16922*.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023b. 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.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. 2023a. Mitigating hallucination in large multi-modal models via robust instruction tuning. In *The Twelfth International Conference on Learning Representations*.

- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023b. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023c. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. 2023d. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. Object hallucination in image captioning. *arXiv preprint arXiv:1809.02156*.
- Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. 2022. A-okvqa: A benchmark for visual question answering using world knowledge. In *European Conference on Computer Vision*, pages 146–162. Springer.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. 2023. Aligning large multimodal models with factually augmented rlhf. *arXiv preprint arXiv:2309.14525*.
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. 2024a. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860*.
- Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. 2024b. Eyes wide shut? exploring the visual shortcomings of multimodal llms. *arXiv preprint arXiv:2401.06209*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- David Wan, Jaemin Cho, Elias Stengel-Eskin, and Mohit Bansal. 2024. Contrastive region guidance: Improving grounding in vision-language models without training. *arXiv preprint arXiv:2403.02325*.
- Junyang Wang, Yuhang Wang, Guohai Xu, Jing Zhang, Yukai Gu, Haitao Jia, Ming Yan, Ji Zhang, and Jitao Sang. 2023. Amber: An llm-free multi-dimensional benchmark for mllms hallucination evaluation. *arXiv preprint arXiv:2311.07397*.
- Xintong Wang, Jingheng Pan, Liang Ding, and Chris Biemann. 2024. Mitigating hallucinations in large vision-language models with instruction contrastive decoding. *arXiv preprint arXiv:2403.18715*.
- Junfei Wu, Qiang Liu, Ding Wang, Jinghao Zhang, Shu Wu, Liang Wang, and Tieniu Tan. 2024. Logical closed loop: Uncovering object hallucinations in large vision-language models. *arXiv preprint arXiv:2402.11622*.
- Qifan Yu, Juncheng Li, Longhui Wei, Liang Pang, Wentao Ye, Bosheng Qin, Siliang Tang, Qi Tian, and Yuet-ing Zhuang. 2023a. Hallucidoctor: Mitigating hallucinatory toxicity in visual instruction data. *arXiv preprint arXiv:2311.13614*.
- Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwan He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, et al. 2023b. Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. *arXiv preprint arXiv:2312.00849*.
- Zihao Yue, Liang Zhang, and Qin Jin. 2024. Less is more: Mitigating multimodal hallucination from an eos decision perspective. *arXiv preprint arXiv:2402.14545*.
- Bohan Zhai, Shijia Yang, Chenfeng Xu, Sheng Shen, Kurt Keutzer, and Manling Li. 2023. Halle-switch: Controlling object hallucination in large vision language models. *arXiv e-prints*, pages arXiv–2310.
- Yi-Fan Zhang, Weichen Yu, Qingsong Wen, Xue Wang, Zhang Zhang, Liang Wang, Rong Jin, and Tieniu Tan. 2024. Debiasing large visual language models. *arXiv preprint arXiv:2403.05262*.
- Zhiyuan Zhao, Bin Wang, Linke Ouyang, Xiaoyi Dong, Jiaqi Wang, and Conghui He. 2023. Beyond hallucinations: Enhancing llms through hallucination-aware direct preference optimization. *arXiv preprint arXiv:2311.16839*.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.

Appendix

Contents

A Visualizations & Analysis on Blind Tokens	12
B More Experimental Details	13
B.1 Further Implementation details . . .	13
B.2 Evaluation Benchmarks	15
B.3 Metrics	16
C Additional Experiments	16
C.1 Zero-Out Experiments on POPE-COCO Benchmark	16
C.2 Inference Time and OPERA . . .	17
C.3 Alternatives to Zero-Out	17
C.4 Results of Larger LVLM	18
C.5 Additional Evaluation on off-the-shelf LVLMs	18
C.6 Additional Evaluation on Generative Benchmarks	18
C.7 POPE (Li et al., 2023b) with Single-Word Constraint	18
C.8 Detailed Results on MME-Fullset	19
C.9 Detailed Results on AMBER Discriminative Tasks	19
D Comparison with "Vision Transformers Need Registers" (Darcet et al., 2023)	19
E Qualitative Results	25

A Visualizations & Analysis on Blind Tokens

In this section, we provide a comprehensive analysis of the attention biases observed in LVLMs through extensive visualizations. Our findings reveal that LVLMs tend to allocate excessive attention to certain image tokens—termed *blind tokens*—which, despite receiving high attention weights, contribute little to the final prediction logits.

More examples of attention bias. To illustrate the phenomenon, we present several examples using LLaVA-1.5-7B. As shown in Fig. 6, the figure depicts three images—black, Gaussian noise, and white—lacking any semantic or query-related information. Despite the question “Is there a carrot in the image?” and the absence of any meaningful objects or features, the model still concentrates its attention on certain regions. This highlights how blind tokens can dominate the attention mechanism, even when there are no informative cues present. Such behavior demonstrates the tendency of LVLMs to latch onto seemingly random patches in the absence of salient visual details. We define blind tokens as image tokens that draw excessive attention while contributing little to the final prediction logits.

Visualization of blind tokens and target objects.

In Fig. 7, we overlay the locations of blind tokens with bounding boxes of target objects on images from the POPE-COCO benchmark. This visualization supports our claim that there is a mismatch between the highly attended blind tokens and the regions containing query-relevant information.

Distributions of bounding boxes and blind tokens. Heatmaps in Fig. 8 illustrate that while object bounding boxes tend to be centered, blind tokens are predominantly located along the image edges, revealing a significant spatial disparity.

Visualization and statistics of blind tokens. We conducted a correlation analysis on 3,000 images from the COCO2014 validation dataset. The results are in Fig. 9. LVLMs were tasked with describing images, and we analyzed the attention distribution over 24×24 patches. Our results indicate that, on average, only 3.7% of blind tokens overlap with actual object regions, with merely 23.3% of the total attention weight allocated to these regions—highlighting the disconnect between blind tokens and task-relevant information.

Histogram of blind tokens. Fig. 10 presents a histogram of the number of blind tokens and their

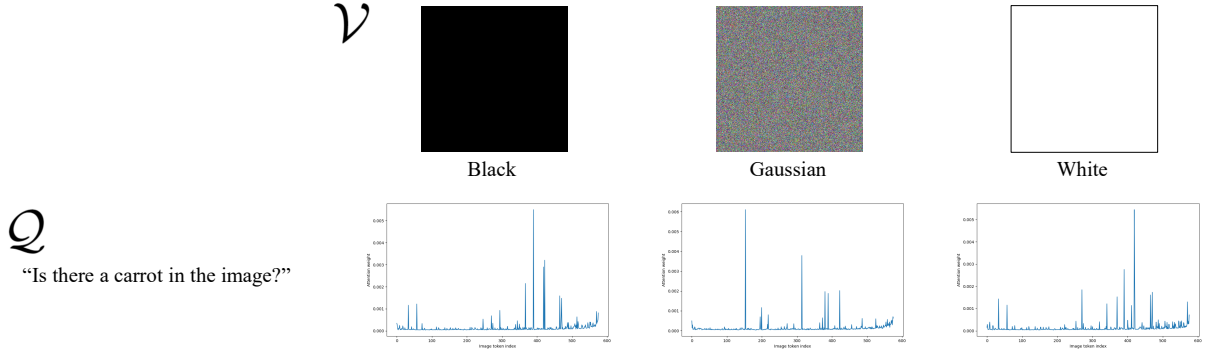


Figure 6: **Attention distribution for images that lack semantic or query-relevant information.** Despite the absence of meaningful content, the model still focuses on certain regions, illustrating how blind tokens can dominate attention even in non-informative scenarios.

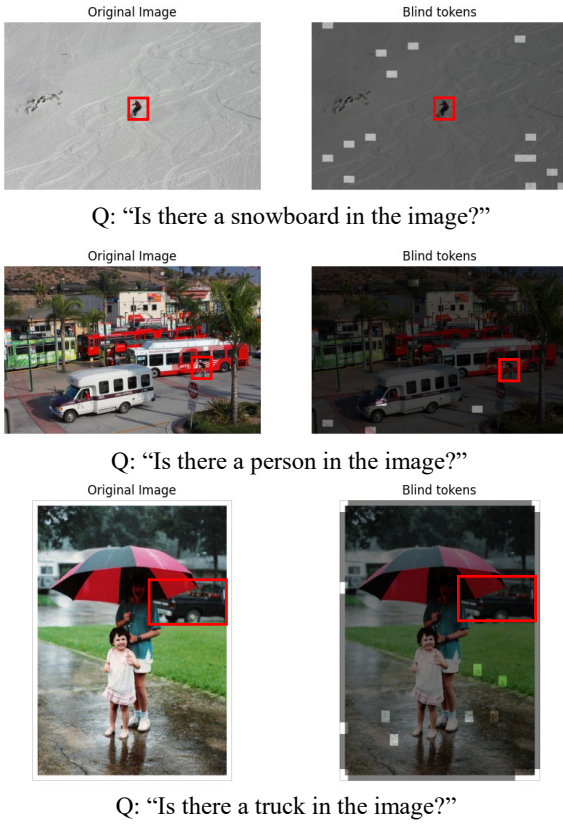


Figure 7: **Visualization of blind tokens in real images from the POPE-COCO benchmark.** Each row displays (left) the original image and (right) the same image with blind tokens highlighted. The red boxes indicate areas where query-related objects are located.

corresponding attention weights. In our evaluation with LLaVA-1.5-7B on the POPE-COCO-Random benchmark, we identified an average of 12.95 blind tokens, which accounted for 33.23% of the total image token attention weight.

Blind tokens and token probability distribution.

Fig. 11 visualizes the location of blind tokens for a given image and query, and presents the token logit values of both the baseline model and AVISC.

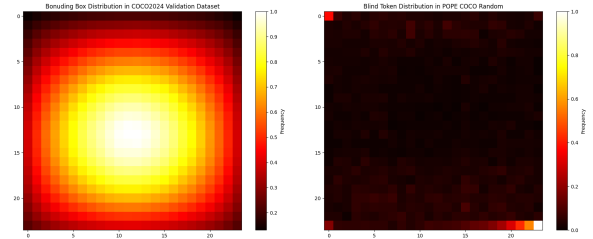


Figure 8: **Comparison of (left) bounding box distribution with (right) the distribution of blind tokens in the COCO dataset.** Warmer colors indicate higher density, revealing that bounding boxes cluster near the center while blind tokens are more prevalent around the periphery. This highlights a spatial mismatch between regions containing genuine objects and areas receiving disproportionately high attention.

For example, in the first problem, which asked whether there is a banana in the image, the original probability distribution was: 'No' at 89.62%, 'Yes' at 8.46%, and 'There' at 1.56%. After applying AVISC, the logit distribution shifted to: 'No' at 98.00%, 'There' at 1.35%, and 'Yes' at 0.61%.

B More Experimental Details

B.1 Further Implementation details

Our decoding process employs cut-off sampling following VCD (Leng et al., 2023). Tokens with probability below β times the maximum probability at each generation step are masked. Formally, we consider text tokens $\xi_t \in \mathcal{H}$ satisfying:

$$\begin{aligned} \mathcal{H}(\xi_{<t}) &= \{\xi_t \in \mathcal{H} \mid p(\xi_t \mid \mathcal{V}, Q, \xi_{<t}; \theta)\} \\ &\geq \beta \max_w p(w \mid \mathcal{V}, Q, \xi_{<t}; \theta). \end{aligned} \quad (11)$$

We set $\beta=0.1$ and limit generation to a maximum of 64 tokens per task. For LLaVA-1.5 (Liu et al., 2023c) experiments, we used the llava_{v1} conversation template.

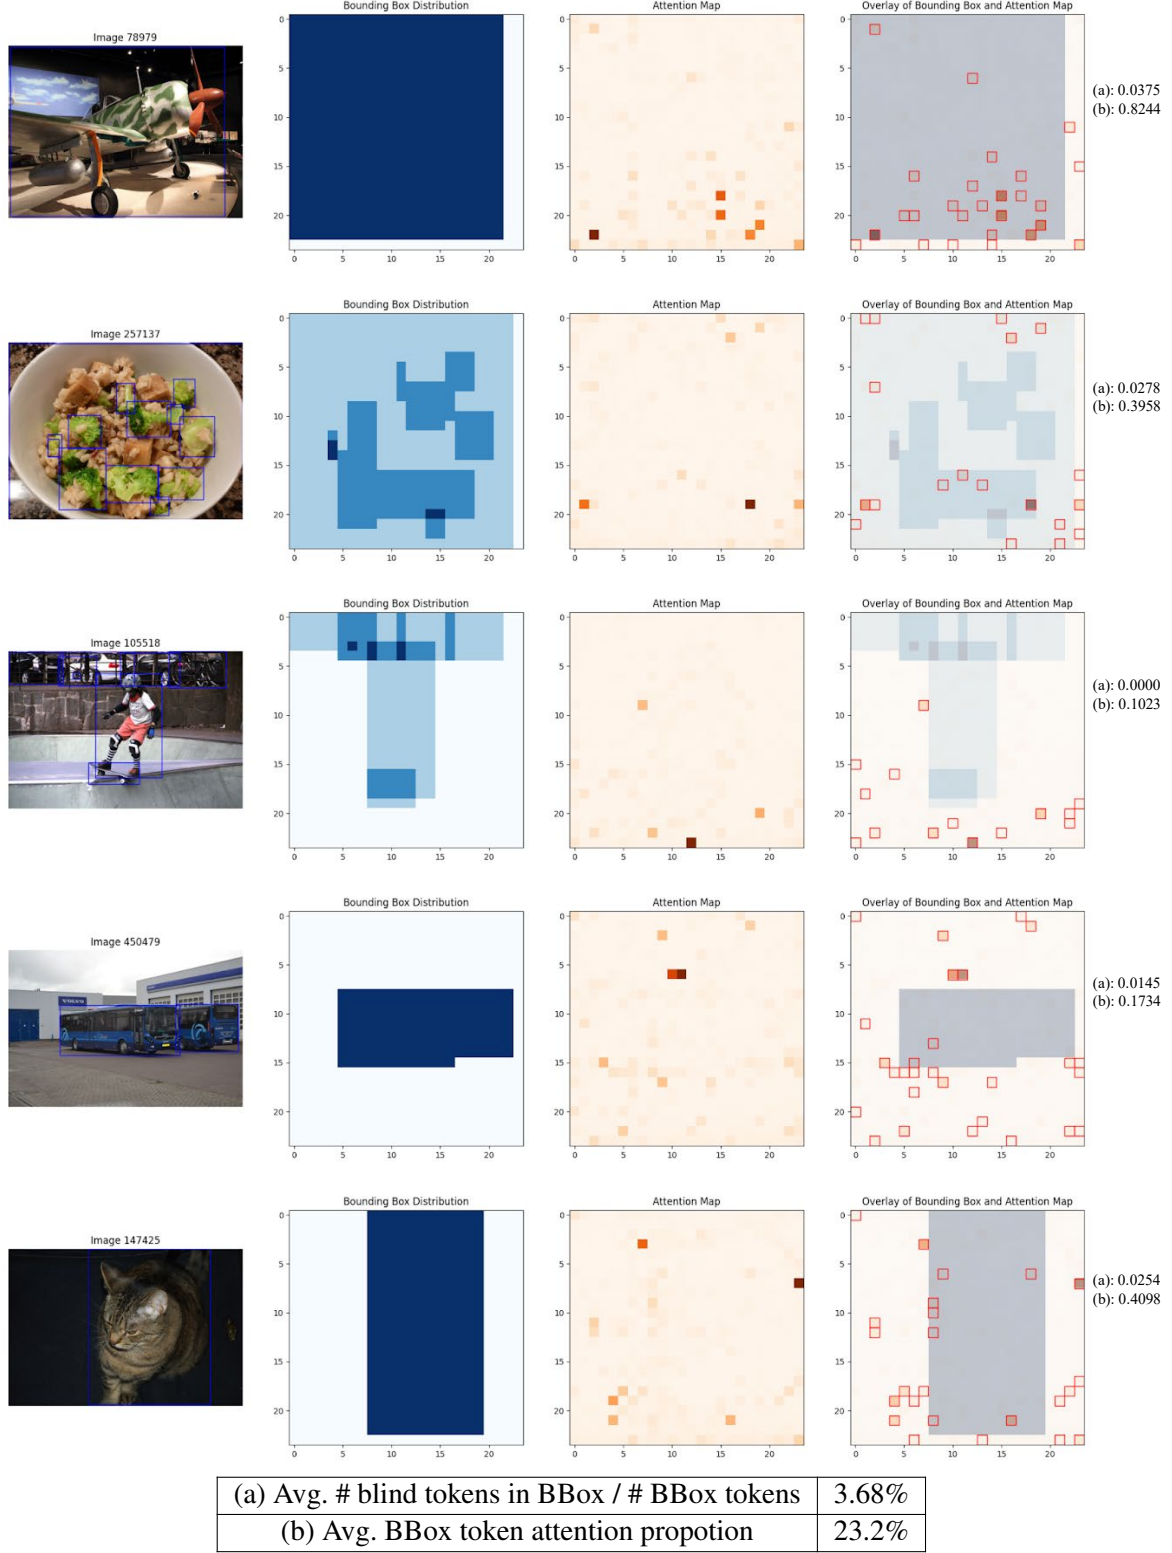


Figure 9: **Visualization and statistics of object bounding boxes and blind tokens in the COCO2014 dataset.** Each row displays (from left to right) the original image, bounding box distribution, attention map, and blind tokens (in red boxes). On average, only 3.68% of blind tokens overlap with bounding boxes, while bounding box regions receive just 23.2% of the total attention. This highlights a clear mismatch between regions containing genuine objects and those receiving high attention.

For reproducing VCD (Leng et al., 2023), we followed the official code with $\alpha = 1.0$, $\beta = 0.1$, and a diffusion noise step $T = 500$. In our M3ID (Favero

et al., 2024) reproduction, we set $\lambda = 0.2$. These settings ensure fair comparisons across methods.

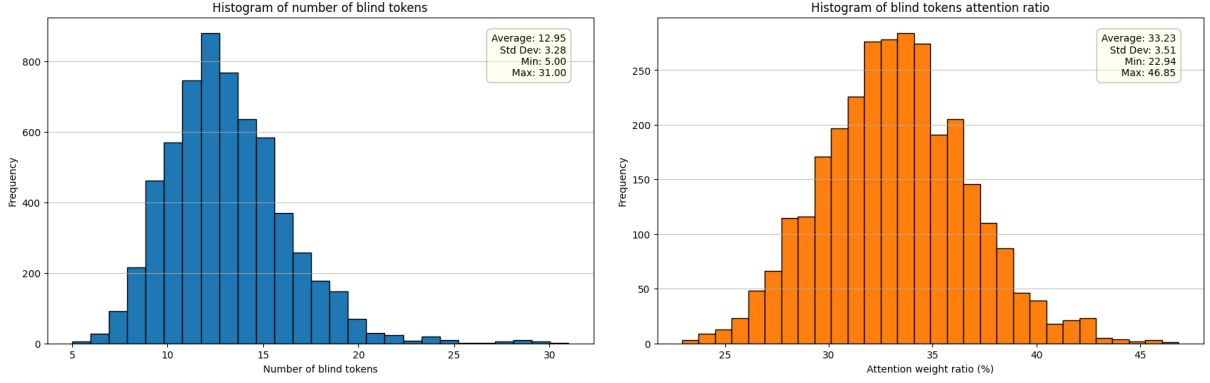


Figure 10: **Histograms illustrating the distribution of blind tokens in the POPE-COCO-Random benchmark.** The left histogram shows the average number of blind tokens per image (about 13), while the right histogram indicates that these tokens account for roughly 33% of the total attention weight. This highlights the disproportionate influence blind tokens exert on the model’s attention.

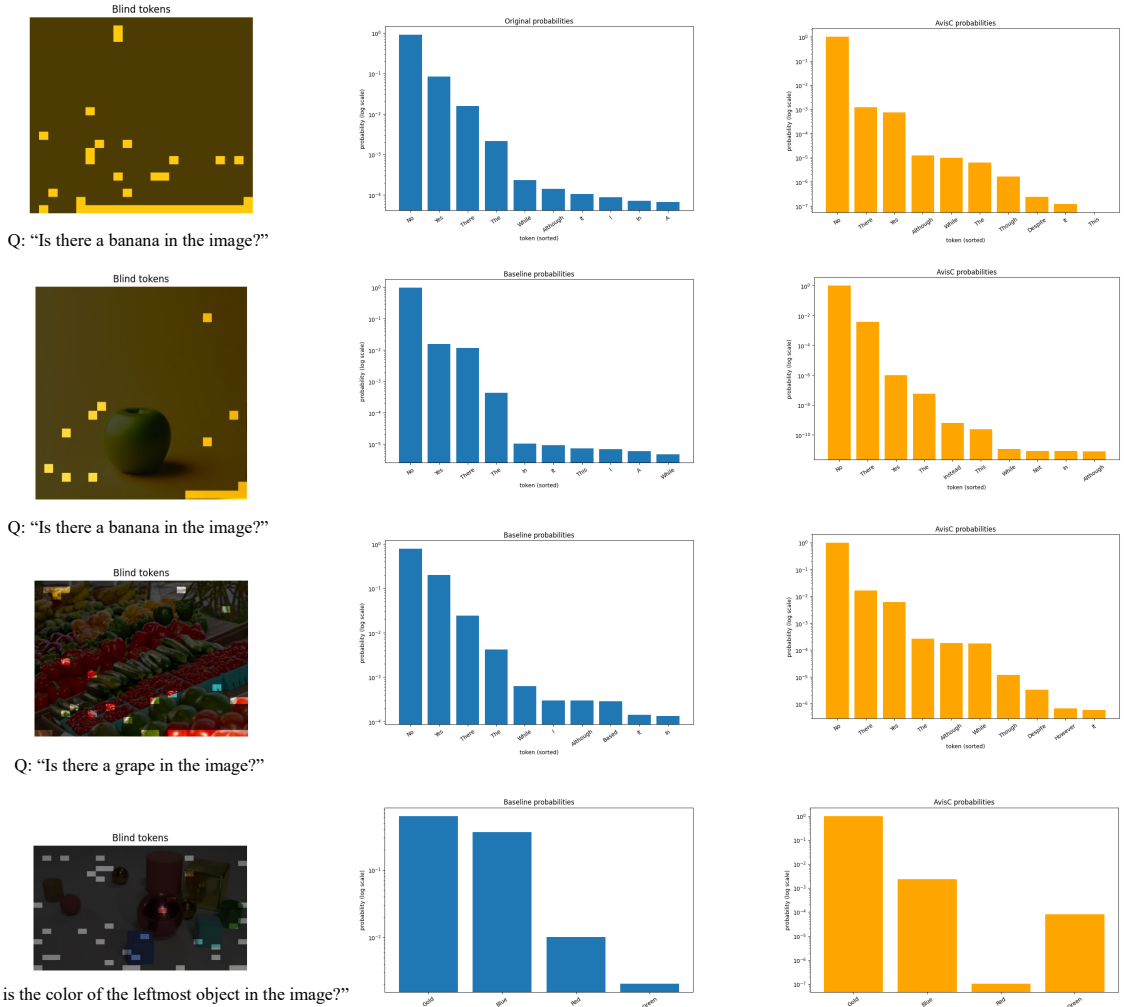


Figure 11: **Visualization of blind tokens and the logit probability distributions before and after AVISC.** Each row shows (left) the image with blind tokens highlighted, (center) the model’s original prediction logits, and (right) the logits adjusted by AVISC. By recalibrating attention away from blind tokens, AVISC increases the accuracy and confidence of the model’s responses to queries.

B.2 Evaluation Benchmarks

POPE. We utilize the official POPE benchmark (Li et al., 2023b), which includes 3,000

question-answer pairs (across random, popular, and adversarial setups) with queries of the form “Is there a [object] in the image?”. Performance is

measured by accuracy, precision, recall, and mean F1-score.⁵

MME. The MME dataset (Fu et al., 2024) is divided into 10 perceptual categories (existence, count, position, color, posters, celebrity, scene, landmark, artwork, OCR) and 4 cognitive categories (commonsense reasoning, numerical calculation, text translation, code reasoning). We use the official dataset but remove the one-word response constraint to allow natural responses.⁶

AMBER. AMBER (Wang et al., 2023) comprises 1004 images with both generative (e.g., “Describe this image.”) and discriminative (existence, attribute, relation) tasks. We randomly sample 500 questions for generative and 5000 for discriminative tasks, following official protocols.⁷

LLaVA-Bench. LLaVA-Bench (Liu et al., 2023c) features 24 images and 60 questions covering diverse contexts (e.g., indoor, outdoor, paintings, sketches) to test LVLM adaptability.⁸

B.3 Metrics

Metrics on the MME. For each visual input \mathcal{V} and its discriminative questions $\{q_1, q_2\}$, we define accuracy (ACC) as:

$$ACC(\mathcal{V}, q_i) = \begin{cases} 1 & \text{if LVLMs}(\mathcal{V}, q_i) \\ & = \text{Answer}(\mathcal{V}, q_i), \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

An additional metric, $ACC+$ (Fu et al., 2024), is 1 if both answers for an image are correct, and 0 otherwise.

$$ACC+(\mathcal{V}) = \begin{cases} 1 & \text{if LVLMs}(\mathcal{V}, q_i) \\ & = \text{Answer}(\mathcal{V}, q_i) \text{ for any } i, \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

The overall MME score is the sum of ACC and $ACC+$.

Metrics on the generative tasks. Let R denote the response generated for a visual input V . We employ:

(1) **CHAIR** (Rohrbach et al., 2018; Wang et al., 2023) evaluates the occurrence of hallucinatory

objects in responses to LVLMs. *CHAIR* uses an annotated list of objects $A=\{a_{obj}^1, a_{obj}^2, \dots, a_{obj}^n\}$ to calculate how often hallucinated objects appear in the responses. Let $R=\{r_{obj}^1, r_{obj}^2, \dots, r_{obj}^m\}$ be the list of objects mentioned in the response of LVLMs, the formula for *CHAIR* is given as:

$$CHAIR = 1 - \frac{\text{len}(R \cap A)}{\text{len}(R)}. \quad (14)$$

(2) **Cover** (Wang et al., 2023) The *Cover* metric measures how completely the objects in the response cover the identified objects in the image. *Cover* calculates the ratio of objects mentioned in the response to the total objects listed. The formula for *Cover* is:

$$Cover = \frac{\text{len}(R \cap A)}{\text{len}(A)}. \quad (15)$$

(3) **Hal** (Wang et al., 2023) The *Hal* metric quantifies the presence of hallucinations by checking if the *CHAIR* value is not zero, indicating the presence of hallucinations. The *Hal* is presented by the following formula:

$$Hal = \begin{cases} 1 & \text{if } CHAIR \neq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

(4) **Cog** (Wang et al., 2023) The *Cog* metric evaluates whether the hallucinations in LVLMs responses resemble human cognition. The *Cog* calculates the ratio of the human hallucinatory object targets, denoted as $H=\{h_{obj}^1, h_{obj}^2, \dots, h_{obj}^n\}$ to the objects mentioned in the response. The formula for *Cog* is:

$$Cog = \frac{\text{len}(R \cap H)}{\text{len}(R)}. \quad (17)$$

(5) **AMBER Score** (Wang et al., 2023) The *AMBER Score* metric evaluates the comprehensive performance of LVLMs for generative tasks and discriminative tasks. This score combines the *CHAIR* metric for generative tasks with the F1 metric for discriminative tasks. The formula representing the *AMBER Score* is as follows:

$$AMBER\ Score = \frac{1}{2} \times (1 - CHAIR + F1). \quad (18)$$

C Additional Experiments

C.1 Zero-Out Experiments on POPE-COCO Benchmark

In Tab. 7, we compare how logits change under two strategies on the POPE-COCO benchmark (Rohrbach et al., 2018): (1) zeroing out blind

⁵<https://github.com/RUCAIBox/POPE>

⁶<https://github.com/BradyFU/Awesome-Multimodal-Large-Language-Models/tree/Evaluation>

⁷<https://github.com/junyangwang0410/AMBER.git>

⁸<https://huggingface.co/datasets/liuhaotian/llava-bench-in-the-wild>

Case	# Case (Base)	# Case (AVISC)	Logit Yype	"Yes" Logit	"No" Logit	GT Logit - Wrong Logit \uparrow
TP (GT = Yes) \uparrow	3952	3958 (+6)	Baseline	30.34	25.68	4.67
			Zero-out $> \mu + \sigma$	28.78	25.56	3.22
			Zero-out $< \mu + \sigma$	20.14	19.40	0.74
TN (GT = No) \uparrow	3317	3536 (+219)	Baseline	26.60	28.88	2.28
			Zero-out $> \mu + \sigma$	25.82	28.45	2.63
			Zero-out $< \mu + \sigma$	18.78	19.21	0.43
FP (GT = No) \downarrow	1183	964 (-219)	Baseline	28.01	27.61	-0.40
			Zero-out $> \mu + \sigma$	26.75	27.48	0.74
			Zero-out $< \mu + \sigma$	19.33	19.29	-0.04
FN (GT = Yes) \downarrow	548	542 (-6)	Baseline	27.42	28.05	-0.63
			Zero-out $> \mu + \sigma$	26.41	27.76	-1.36
			Zero-out $< \mu + \sigma$	19.08	19.26	-0.18

Table 7: **Zero-out experiments on the POPE-COCO benchmark (Rohrbach et al., 2018).** We compare how logits change under two strategies: (1) zeroing out blind tokens (Zero-out $> \mu + \sigma$) and (2) zeroing out non-blind tokens (Zero-out $< \mu + \sigma$). Rows denote true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). For each case, we show the average “Yes” logit, “No” logit, and the difference between the ground-truth (GT) logit and the wrong logit. All results are obtained with LLaVA-1.5-7B (Liu et al., 2023b) on 3,000 MS-COCO (Lin et al., 2014) images.

Method	Acc. \uparrow	Prec. \uparrow	Rec. \uparrow	F1 \uparrow	tokens/sec \uparrow
base	84.47	83.35	86.13	84.72	24.44
VCD	84.80	83.00	87.53	85.20	11.53
M3ID	86.00	85.11	87.27	86.18	13.14
AvisC	87.93	88.24	87.53	87.88	12.28
OPERA (Beam=2)	89.35	90.37	88.80	89.58	0.17

Table 8: **Comparison of inference time and performance on the POPE-COCO-Random benchmark for LLaVA-1.5.** While OPERA achieves the highest performance metrics, it operates at a substantially slower speed compared to the other methods.

tokens (*i.e.*, tokens with attention $> \mu + \sigma$) and (2) zeroing out non-blind tokens (*i.e.*, tokens with attention $< \mu + \sigma$). Here, removing blind tokens minimally alters the model’s predictions, indicating that they hold little object-discriminative information. In contrast, removing non-blind tokens drastically shifts the logits, underscoring their critical importance. This indicates that blind tokens have a smaller impact on prediction logits than non-blind tokens. Compared to base decoding, AVISC effectively reduces over-emphasis on blind tokens, improving performance, particularly for TN and FP cases.

C.2 Inference Time and OPERA

Tab. 8 presents an efficiency and performance comparison between contrastive decoding methods (AVISC, M3ID, OPERA, and VCD) and AVISC. Inference speed is measured with a TiTAN RTX GPU on the POPE-COCO-Random benchmark. OPERA introduces the concept of an “anchor token” and uses this token to guide sentence generation and rollback, thereby mitigating hallucinations. OPERA is implemented on the beam search decod-

Case	Acc. \uparrow	Prec. \uparrow	Rec. \uparrow	F1 \uparrow
InstructBLIP	Zeros	88.50	93.00	83.27
	Ones	82.50	75.48	96.27
	Noise	86.77	84.71	89.73
	Mask	88.53	90.14	86.53
LLaVA 1.5	Zeros	87.87	88.12	87.53
	Ones	79.97	72.22	97.40
	Noise	88.47	93.19	83.00
	Mask	84.77	86.29	82.67

Table 9: **Design choices for non-blind image token deactivation.** Each row presents a different method for handling non-blind tokens (Zeros, Ones, Noise, or Mask), and shows the resulting performance.

ing method of LLMs, so a fair comparison with AVISC is not possible. However, OPERA showed the best performance overall. However, its inference speed was approximately $\times 72.23$ slower than AVISC.

C.3 Alternatives to Zero-Out

Table 9 shows ablation results for various deactivation schemes applied to non-blind image tokens on the POPE-COCO-random benchmark (Li et al., 2023b), using both InstructBLIP (Dai et al., 2024) and LLaVA 1.5 (Liu et al., 2023c). We compare four methods: setting tokens to zero (Zeros), to ones (Ones), replacing tokens with noise (Noise), and masking tokens out in the attention mechanism (Mask). For InstructBLIP, the Mask approach achieves the highest Accuracy and F1 score, while the Zeros method excels in Precision; Ones yields the best Recall, and Noise offers balanced performance across Precision and Recall. For LLaVA 1.5, Noise achieves the highest Accuracy and Precision, whereas Zeros demonstrates consistent, balanced

Setup	Method	LLaVA-1.5 (13B)			
		Acc. ↑	Prec. ↑	Rec. ↑	F1 ↑
Random	<i>base</i>	83.17	79.49	89.40	84.15
	VCD	82.97	78.90	90.00	84.09
	M3ID	83.43	79.31	90.47	84.52
	AVIS	88.40	86.05	91.67	88.77
Popular	<i>base</i>	80.93	76.45	89.40	82.42
	VCD	79.67	74.59	90.00	81.57
	M3ID	80.90	75.94	90.47	82.57
	AVIS	85.73	81.94	91.67	86.53
Adversarial	<i>base</i>	76.03	70.74	88.80	78.75
	VCD	75.57	69.86	89.93	78.64
	M3ID	75.80	69.97	90.40	78.88
	AVIS	79.27	73.65	91.13	81.47

Table 10: **Results of LLaVA-1.5-13B on POPE-COCO benchmark.**

Setup	Method	LLaVA-OneVision (Qwen2-7B)		Qwen2.5-VL-7B	
		Acc.	F1	Acc.	F1
Random	<i>base</i>	88.60	87.34	88.84	87.82
	VCD	90.57	89.71	90.47	89.37
	M3ID	89.87	88.91	90.77	88.57
	AVIS	91.46	90.84	92.36	91.50
Popular	<i>base</i>	84.60	83.79	85.05	83.72
	VCD	87.20	86.50	87.10	86.21
	M3ID	87.60	85.65	88.50	85.36
	AVIS	89.86	89.45	90.76	90.16
Adversarial	<i>base</i>	83.00	82.40	84.07	82.73
	VCD	86.00	85.42	86.50	85.52
	M3ID	86.80	84.86	87.30	84.96
	AVIS	86.12	85.83	87.62	86.93

Table 11: **POPE (MS-COCO) results on LLaVA-OneVision-Qwen2-7B and Qwen2.5-VL-7B.**

performance across all metrics. Overall, the Zeros approach proved most effective in calibrating attention to image tokens and improving model performance.

C.4 Results of Larger LVL

Tab. 10 presents the performance of each method on the POPE benchmark using the COCO dataset based on the LLaVA-1.5v-13B model. In this experiment setup, compared to the 7B small model shown in Tab. 1, the performance improvement of AVISC is even more pronounced. For other methods (*i.e.*, VCD, M3ID), the performance increase is slight or, in some cases, decreases depending on the metric. However, AVISC demonstrates robust performance improvement, remaining resilient to changes in the size of LVLs.

C.5 Additional Evaluation on off-the-shelf LVLs

To further validate the model-agnostic robustness and generalizability of AVISC, we extended our

Model	Method	MMHal-Bench		Object-Hallucination	
		Score ↑	HalRate ↓	CHAIR _s ↓	CHAIR _t ↓
InstructBLIP	<i>base</i>	1.84	0.64	0.70	9.1
	VCD	1.75	0.64	0.80	8.9
	M3ID	1.70	0.65	0.90	7.6
	OPERA (Beam)	—	—	16.6	6.8
	AVIS	2.03	0.59	0.70	8.3
LLaVA-1.5	<i>base</i>	1.59	0.72	25.0	9.2
	VCD	1.96	0.64	23.6	8.4
	M3ID	2.14	0.61	23.2	7.3
	OPERA (Beam)	2.15	0.54	45.1	22.3
	AVIS	2.19	0.59	22.1	7.8

Table 12: **Results on free-form generative benchmarks.**

evaluation to include recent and diverse LVLs with varying attention mechanisms. Specifically, we tested AVISC on LLaVA-OneVision-Qwen2-7B, Qwen2.5-VL-7B. As shown in Tab. 11, AVISC consistently outperforms prior methods across all setups on the POPE benchmark, achieving significant reductions in hallucination without compromising the original model capabilities. These results support AVISC’s effectiveness as a test-time, plug-and-play strategy applicable across off-the-shelf LVLs.

C.6 Additional Evaluation on Generative Benchmarks

To more comprehensively evaluate the impact of AVISC on generative capabilities, we conducted additional experiments on two free-form generation benchmarks: MMHal-Bench and Object-Hallucination. The results, summarized in Tab. 12, show that AVISC not only preserves the generative quality of the base models (*e.g.*, maintaining or improving MMHal scores), but also consistently reduces hallucination metrics compared to existing approaches. These findings highlight that AVISC serves as a reliable and effective test-time method that retains the strengths of pretrained LVLs in generative settings.

C.7 POPE (Li et al., 2023b) with Single-Word Constraint

As shown in Tab. 13, we see that imposing a one-word response constraint on LVLs leads to notable changes in performance compared to Tab. 1. Despite the change in query setup, AVISC shows the best performance on the POPE benchmark. Specifically, precision and recall vary significantly in the COCO random setup comparing scenarios with and without the instruction, "Please answer this question with one word." To mitigate these impacts and better evaluate discriminative capabilities,

Setup	Method	InstructBLIP (Dai et al., 2024)				LLaVA 1.5 (Liu et al., 2023c)				
		Acc. ↑	Prec. ↑	Rec. ↑	F1 ↑	Acc. ↑	Prec. ↑	Rec. ↑	F1 ↑	
MS-COCO	Random	base	81.53	82.71	79.73	81.19	83.77	92.31	73.67	81.94
		VCD	82.03	83.77	79.47	81.56	85.43	93.25	76.40	83.99
		AVIS-C	86.03	95.53	75.60	84.41	84.67	97.88	70.87	82.21
	Popular	base	78.47	77.73	79.80	78.75	82.57	89.62	73.67	80.86
		VCD	79.13	78.94	79.47	79.20	83.17	88.36	76.40	81.94
		AVIS-C	84.27	91.45	75.60	82.77	83.67	95.25	70.87	81.27
	Adversarial	base	77.43	76.09	80.00	78.00	79.77	83.85	73.73	78.47
		VCD	77.23	76.10	79.40	77.72	80.27	82.76	76.47	79.49
		AVIS-C	81.83	86.20	75.80	80.67	81.83	90.99	70.67	79.55
A-OKVQA	Random	base	81.33	78.52	86.27	82.21	84.93	89.16	79.53	84.07
		VCD	81.57	78.78	86.40	82.42	85.53	87.64	82.73	85.12
		AVIS-C	87.10	89.95	83.53	86.62	87.33	95.09	78.73	86.14
	Popular	base	76.87	72.69	86.07	78.82	80.90	81.77	79.53	80.64
		VCD	77.30	73.10	86.40	79.19	81.17	80.22	82.73	81.46
		AVIS-C	82.47	81.79	83.53	82.65	85.03	90.08	78.73	84.03
	Adversarial	base	71.40	66.67	85.60	74.96	74.80	72.63	79.60	75.95
		VCD	72.47	67.39	87.07	75.97	75.03	71.87	82.27	76.72
		AVIS-C	76.47	73.16	83.60	78.03	79.27	79.58	78.73	79.16
GQA	Random	base	80.57	77.47	86.20	81.60	84.80	87.88	80.73	84.16
		VCD	81.73	79.02	86.40	82.55	85.63	86.89	83.93	85.38
		AVIS-C	85.30	88.57	81.07	84.65	87.40	95.17	78.80	86.21
	Popular	base	74.67	70.17	85.80	77.20	79.37	78.59	80.73	79.64
		VCD	74.63	69.94	86.40	77.30	78.73	76.03	83.93	79.78
		AVIS-C	80.63	80.37	81.07	80.72	83.33	86.66	78.80	82.54
	Adversarial	base	72.63	67.78	86.27	75.92	76.00	74.13	79.87	76.89
		VCD	71.93	67.21	85.67	75.32	76.40	72.76	84.40	78.15
		AVIS-C	77.60	75.91	80.87	78.31	80.37	81.52	78.53	80.00

Table 13: **POPE (Li et al., 2023b) results with one-word constraint.** We use the instruction "Please answer in one word." at the end of the query text.

we designed experiments that allow the LVLMs to freely make judgments and provide explanations for these judgments rather than restricting them to answers in one word.

C.8 Detailed Results on MME-Fullset

The detailed results on MME-Fullset are provided in Tab. 14. AVISC demonstrates substantial improvements in both LLaVA-1.5 and InstructBLIP across a wide range of perception and recognition tasks. These findings highlight the capability of AVISC to effectively handle diverse tasks, extending beyond hallucination mitigation, and suggest its potential to enhance the ability of LVLMs to accurately interpret and analyze visual information and query text appropriately.

C.9 Detailed Results on AMBER Discriminative Tasks

Tab. 15 presents the performance of the discriminative task on the AMBER benchmark across different categories. The discriminative task in the AMBER benchmark is divided into six categories: 'Existence', 'Attribute', 'State', 'Number', 'Action', and 'Relation', to evaluate the model's performance. For most categories, except for a few, both the LLaVA-1.5 and InstructBLIP models show performance improvements due to the applied AVISC.

D Comparison with "Vision Transformers Need Registers" (Darcet et al., 2023)

Summary of (Darcet et al., 2023). (Darcet et al., 2023) identify artifacts in vision transformer fea-

Task	Category	LLaVA 1.5 (Liu et al., 2023c)				InstructBLIP (Dai et al., 2024)			
		<i>base</i>	VCD	M3ID	AVISC	<i>base</i>	VCD	M3ID	AVISC
Perception	Existence	173.57 (±8.16)	172.14 (±8.09)	178.33 (±6.83)	189.29 (±1.89)	170.19 (±11.12)	172.62 (±8.92)	173.89 (±10.52)	184.76 (±5.56)
	Count	110.00 (±15.82)	117.14 (±8.76)	107.22 (±14.78)	104.76 (±11.66)	89.52 (±11.04)	98.33 (±15.99)	89.72 (±13.44)	82.85 (±12.16)
	Position	100.47 (±18.78)	103.33 (±20.56)	96.39 (±5.52)	106.19 (±13.93)	67.62 (±14.04)	71.90 (±13.42)	72.72 (±14.77)	74.76 (±6.19)
	Color	125.24 (±15.91)	119.52 (±8.58)	127.50 (±8.28)	127.86 (±9.13)	114.76 (±9.60)	117.14 (±10.70)	110.56 (±7.20)	131.43 (±4.76)
	Posters	132.31 (±6.73)	135.54 (±3.61)	132.82 (±7.94)	150.85 (±6.49)	114.97 (±6.25)	129.08 (±6.97)	114.46 (±6.97)	145.92 (±2.41)
	Celebrity	114.56 (±6.45)	118.09 (±7.69)	113.38 (±0.21)	125.59 (±2.50)	113.38 (±3.95)	123.82 (±4.99)	114.12 (±2.91)	120.29 (±7.90)
	Scene	149.13 (±0.53)	150.00 (±3.54)	156.63 (±1.59)	162.00 (±1.06)	140.50 (±0.71)	136.50 (±10.25)	141.00 (±1.06)	150.38 (±3.36)
	Landmark	138.25 (±4.95)	140.75 (±4.95)	135.13 (±4.77)	142.38 (±0.53)	98.50 (±0.35)	110.75 (±4.24)	103.25 (±6.72)	99.25 (±0.35)
	Artwork	97.50 (±2.83)	95.25 (±4.24)	89.38 (±3.36)	101.00 (±7.42)	110.38 (±4.42)	113.00 (±3.54)	110.13 (±6.89)	123.38 (±2.30)
	OCR	91.25 (±19.45)	101.25 (±1.77)	96.25 (±15.91)	143.75 (±5.3)	87.50 (±21.21)	91.25 (±8.84)	85.00 (±10.61)	68.75 (±5.3)
Recognition	Commonsense Reasoning	100.36 (±2.53)	96.79 (±5.56)	87.14 (±12.12)	102.86 (±7.07)	96.43 (±1.01)	107.14 (±8.08)	99.64 (±2.53)	101.79 (±6.57)
	Numerical Calculation	80.00 (±7.07)	66.25 (±8.84)	76.25 (±12.37)	65.00 (±14.14)	68.75 (±1.77)	66.25 (±15.91)	71.25 (±22.98)	73.75 (±5.30)
	Text Translation	75.00 (±3.54)	86.25 (±22.98)	65.00 (±14.14)	77.50 (±17.68)	63.75 (±5.3)	91.25 (±1.77)	53.75 (±5.3)	86.25 (±1.77)
	Code Reasoning	62.50 (±10.61)	61.25 (±1.77)	71.25 (±15.91)	71.25 (±5.30)	73.75 (±5.30)	57.50 (±0.00)	81.25 (±1.77)	76.25 (±5.3)

Table 14: Results on MME-Fullset (Fu et al., 2024).

Category	LLaVA 1.5 (Liu et al., 2023c)				InstructBLIP (Dai et al., 2024)			
	<i>base</i>	VCD	M3ID	AVISC	<i>base</i>	VCD	M3ID	AVISC
Existence	68.55 (±0.21)	67.15 (±1.91)	68.50 (±0.14)	75.35 (±0.21)	72.05 (±0.49)	73.20 (±1.27)	72.95 (±0.21)	81.35 (±0.07)
Attribute	67.85 (±0.49)	69.50 (±1.27)	68.20 (±0.42)	69.80 (±0.85)	68.40 (±0.14)	69.90 (±0.14)	69.15 (±0.92)	70.80 (±1.56)
State	65.55 (±0.35)	67.80 (±0.28)	65.75 (±0.64)	68.40 (±1.70)	70.55 (±0.64)	72.40 (±0.00)	70.70 (±0.85)	72.85 (±1.77)
Number	69.05 (±0.78)	68.50 (±2.40)	68.95 (±0.92)	67.10 (±1.84)	60.90 (±0.00)	60.70 (±0.85)	61.80 (±0.71)	60.85 (±0.49)
Action	78.50 (±3.96)	81.90 (±3.39)	81.50 (±1.84)	84.50 (±3.25)	74.95 (±2.05)	79.05 (±2.62)	78.70 (±1.27)	85.20 (±2.40)
Relation	58.80 (±4.10)	57.75 (±0.07)	59.70 (±3.39)	60.50 (±0.14)	56.05 (±1.63)	58.00 (±1.41)	57.00 (±1.98)	54.65 (±2.76)

Table 15: Results on AMBER discriminative tasks (Wang et al., 2023).

ture maps—specifically, “high-norm outlier tokens” that concentrate attention in redundant background areas. These tokens capture significant global information despite lacking local details, leading to poor performance in tasks requiring spatial precision. Notably, when additional memory (register tokens) is introduced, these artifacts vanish.

Differences from blind tokens. While both high-norm outlier tokens and our *blind tokens* exhibit unusually high attention weights in seemingly irrelevant regions, key differences exist:

- **Source of Attention:** High-norm tokens are computed within vision transformer layers, whereas our blind tokens are derived from the LLM’s attention in LVLMs (e.g., Vicuna-7B



Query: Is there a chair in the image?



GT: No

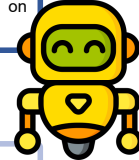
InstructBLIP

Base: Yes, there is a wooden chair in the image, used primarily by the children.

VCD: There is a chair in the image, with kids sitting on it.

M3ID: Yes, there is a chair in the image.

AvisC: There are several children sitting at tables on mats, with no visible chairs in the image.



LLaVA-1.5

Base: Yes, there is a tv in the image. It sits atop a large modern dresser in the dining and living room area.

VCD: Yes, there is a flat screen television in the living room.

M3ID: Yes, there is a TV in the image, located in the living room area of the house.

AvisC: No, there is no TV in the image. The image features a dining room table with chairs, but it does not show any TV present in the image.

Query: Is there a tv in the image?



GT: No

InstructBLIP

Base: No, there isn't a chair in the image.

VCD: There is no chair in the image.

M3ID: No, there is no chair in the image.

AvisC: There is a step stool or a chair in the image, which the little girl is standing on while brushing her teeth.



Query: Is there a chair in the image?



GT: Yes

LLaVA-1.5

Base: No, there is no chair in the image. The image features a boy standing by a wall, surrounded by a bed with a comforter and a pillow.

VCD: No, there is no chair in the image. The scene takes place in a bedroom.

M3ID: No, there is no chair in the image. The image features a young boy standing on a bed with a flowered blanket.

AvisC: Yes, there is a chair in the image, located next to the boy standing on the pillows.

Query: Is there a chair in the image?



GT: Yes

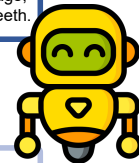


Figure 12: Qualitative examples on POPE (Li et al., 2023b).



Query: Is there a black couch in the image?



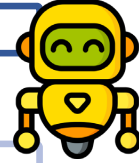
GT: No

InstructBLIP

Base: Yes, there is a black couch in the image.

VCD: Yes, there is a red couch with a black dog sitting on top of it in the image.

AvisC: No, there is a red couch in the image.



LLaVA-1.5

Query: Is there a brown scarf in the image?



GT: No

Base: Yes, there is a red and brown scarf wrapped around the teddy bear's neck in the image.

VCD: Yes, there is a brown scarf in the image, wrapped around the teddy bear.

AvisC: No, there is no brown scarf in the image. The teddy bear in the picture is wrapped in a red bandana or a red scarf.



Query: Is the motorcycle on the right side of the bus?



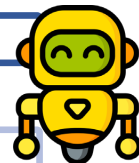
GT: Yes

InstructBLIP

Base: No.

VCD: No.

AvisC: Yes, the motorcycle is on the right side of the bus.



LLaVA-1.5

Query: Is there a blue and yellow fire hydrant in the image?



GT: Yes

Base: No, there is no blue and yellow fire hydrant in the image.

VCD: Yes, there is a blue and yellow fire hydrant in the image.

AvisC: Yes, there is a blue and yellow fire hydrant in the image, located on the side of the road.

Figure 13: Qualitative examples on MME (Fu et al., 2024).

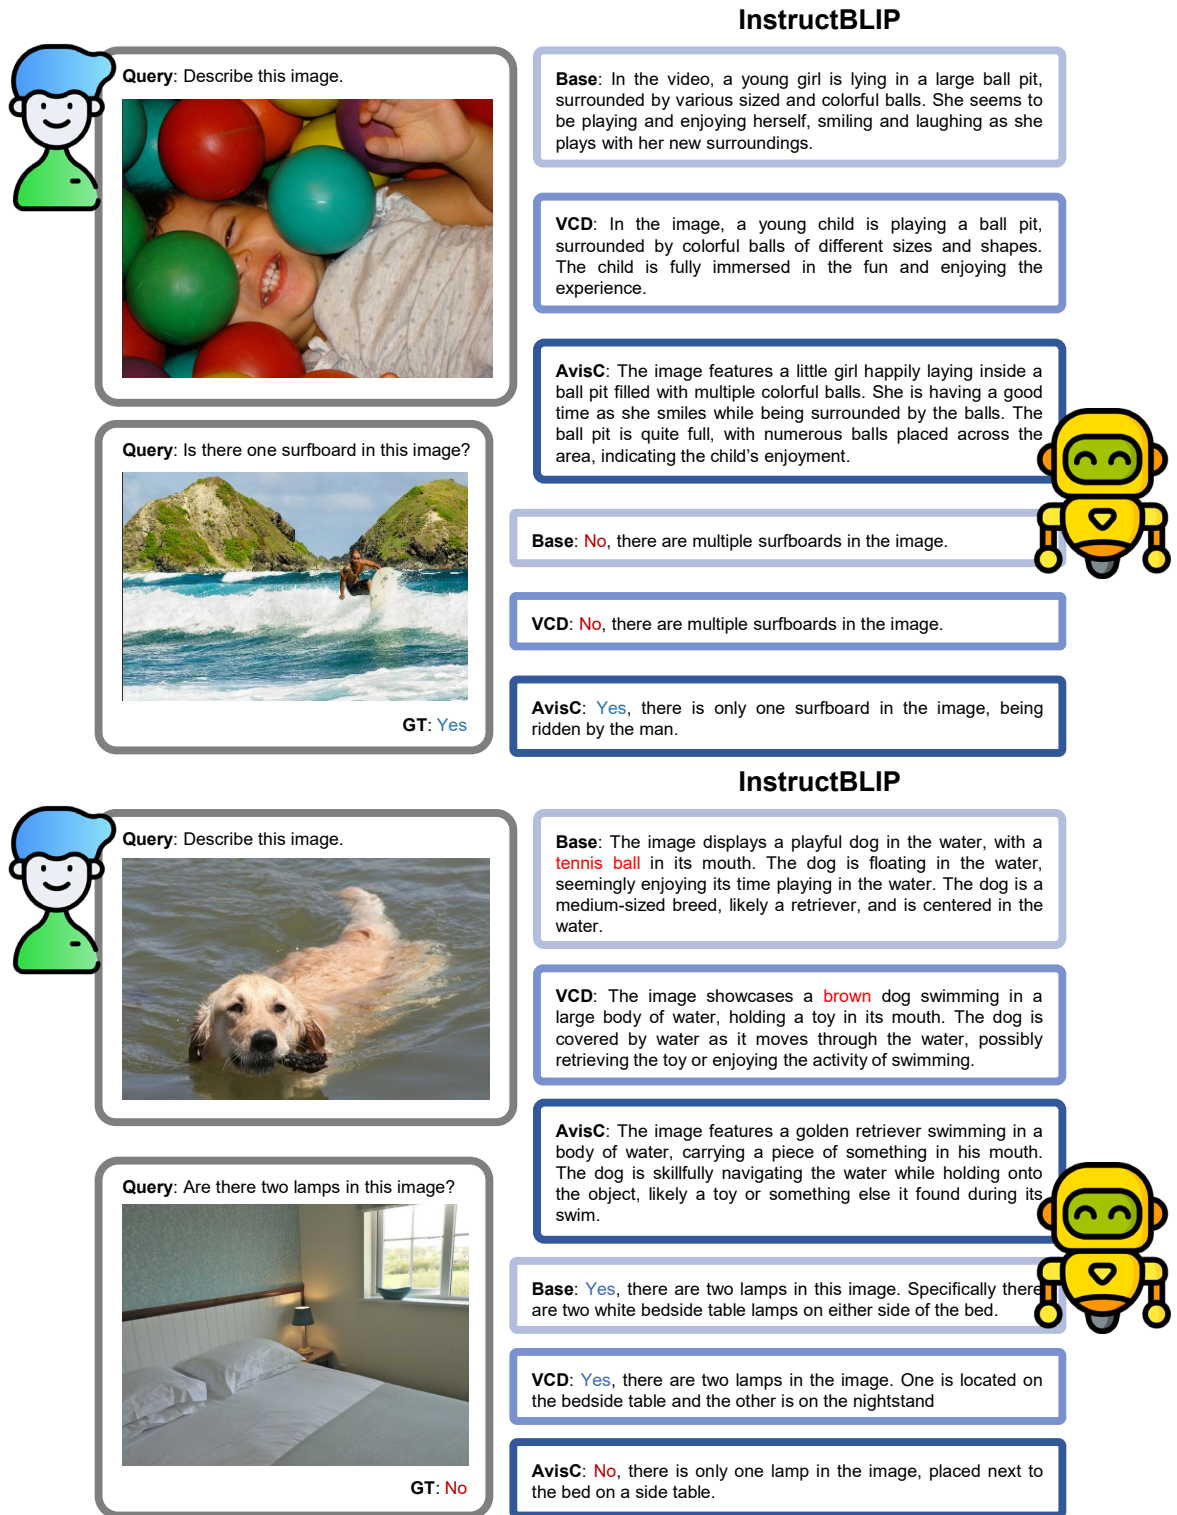


Figure 14: Qualitative examples of InstructBLIP (Dai et al., 2024) on AMBER (Wang et al., 2023).

in LLaVA-1.5-7B), with differences in mask-
ing strategies.

- **Task and Architecture:** Vision transformers are optimized for dense prediction tasks, and the emergence of high-norm tokens is sensitive to the training regime (e.g., DINOv1

vs. DINOv2). In contrast, LVLMs integrate a visual encoder with an LLM via an auto-regressive prediction scheme for image-based Q&A tasks.

- **Domain:** LVLMs project image tokens into an LLM space, altering attention dynamics



Figure 15: Qualitative examples of LLaVA-1.5 (Liu et al., 2023b) on AMBER (Wang et al., 2023).

compared to pure vision transformers.

Our experiments show a moderate correlation between high-norm and blind tokens (with $P(\text{blind token} \mid \text{high-norm token}) = 40.38\%$ and

$P(\text{high-norm token} \mid \text{blind token}) = 31.27\%$), suggesting shared underlying properties despite their differences. Additionally, blind tokens tend to appear at the beginning and end of the image token sequence—a pattern not clearly observed for



Figure 16: **Response comparison on LLaVA-Bench (Liu et al., 2023c)**. Hallucinations are colored in **red**. AVISC demonstrates a robust understanding of images and reduces hallucinations in responses.

high-norm tokens in vision transformers.

On reducing dependency on blind tokens. Although high-norm tokens in (Darcet et al., 2023) encode global information, our findings indicate that blind tokens in LVLMs often lack query-relevant details. As shown in Fig. 2, the essential information is typically captured by non-blind tokens. Therefore, reducing the influence of blind tokens via our contrastive decoding scheme—while enhancing the role of non-blind tokens—effectively mitigates hallucinations without sacrificing critical information.

E Qualitative Results

We provide qualitative results on all benchmarks (POPE (Li et al., 2023b), MME (Fu et al., 2024), AMBER (Wang et al., 2023), and LLaVA-Bench (Liu et al., 2023c)) in Figs. 12 to 15. These highlight the differences between sentences generated by standard decoding (Base), VCD (Leng et al., 2023), and those produced by AVISC. The results demonstrate the effectiveness of AVISC in dealing with a variety of challenging visual contexts. Base and VCD often generate descriptions that include errors or hallucinations where elements

not present in the image are described. In contrast, AVISC helps counteract these hallucinations, generating sentences that reflect a more accurate comprehension of the image.