

Revealing and Enhancing Core Visual Regions: Harnessing Internal Attention Dynamics for Hallucination Mitigation in LVLMs

Guangtao Lyu¹, Qi Liu¹, Chenghao Xu³, Jiexi Yan², Muli Yang⁴, Xueting Li¹, Fen Fang⁴, Cheng Deng^{1*}

¹School of Electronic Engineering, Xidian University, Xi'an, China

²School of Computer Science and Technology, Xidian University, Xi'an, China

³College of Information Science and Engineering, Hohai University, Nanjing, China

⁴Institute for Infocomm Research, A*STAR, Singapore

{guangtaolyu, qiliu, lxt}@stu.xidian.edu.cn, {jxyan1995, chdeng.xd}@gmail.com, chx@hhu.edu.cn, {yangml, fang_fen}@astar.edu.sg

Abstract

LVLMs have achieved strong multimodal reasoning capabilities but remain prone to hallucinations, producing outputs inconsistent with visual inputs or user instructions. Existing training-free methods, including contrastive decoding and auxiliary expert models, which incur several times more computational overhead and may introduce potential interference, as well as static internal signal enhancement, are often vulnerable to the attention sink phenomenon. We find that internal Positive Attention Dynamics (PAD) in LVLMs naturally reveal semantically core visual regions under the distortions of attention sinks. Based on this, we propose Positive Attention Dynamics Enhancement (PADE), a training-free attention intervention that constructs a PAD map to identify semantically core visual regions, applies per-head Median Absolute Deviation Scaling to adaptively control the intervention strength, and leverages System-Token Compensation to maintain attention to complex user instructions and support long-term output consistency. Experiments on multiple LVLMs and benchmarks show that PADE improves visual grounding and reduces hallucinations, validating the effectiveness of leveraging internal attention dynamics for reliable multimodal reasoning.

1 Introduction

Large Vision Language Models (LVLMs) have achieved remarkable progress in multimodal reasoning and demonstrate strong performance across a wide range of tasks (Liu et al., 2023; Achiam et al., 2023; Bai et al., 2023). Despite these advances, LVLMs remain prone to hallucinations (Liu et al., 2024b; Ji et al., 2023; Bi et al., 2025b), generating content inconsistent with visual inputs or user instructions, which undermines their reliability, particularly in safety-critical applications such as medical analysis (Sun et al., 2024a; Chen et al., 2024a;

*Corresponding author

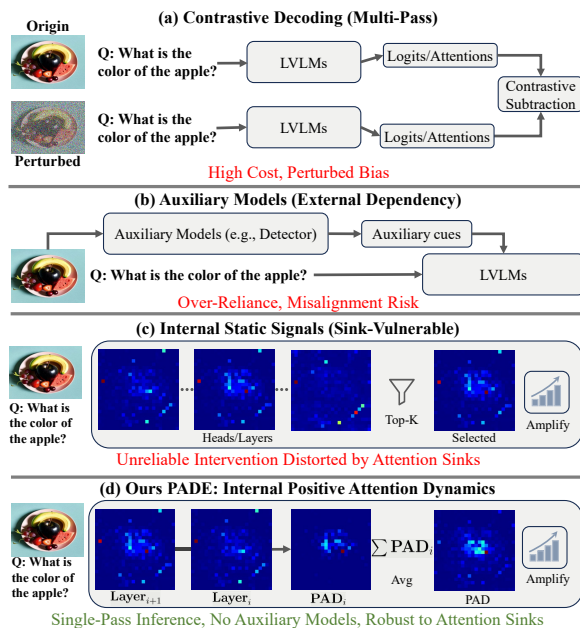


Figure 1: Comparison of hallucination mitigation paradigms. (a) Contrastive decoding methods. (b) Auxiliary expert methods. (c) Static internal signal methods. (d) Ours PADE: internal positive attention dynamics.

Hu et al., 2024b) and autonomous driving (Jiang et al., 2024; Sun et al., 2025a; Shao et al., 2024).

Recent studies have shown that hallucinations in LVLMs primarily arise from an over-reliance on linguistic priors and insufficient utilization of visual inputs (Bai et al., 2024; Ji et al., 2023; Liu et al., 2024b; Bi et al., 2025a). To address this issue, many inference-time intervention methods have been proposed to enhance visual grounding by amplifying vision-related signals, such as visual tokens, attention activations, or output logits. Existing strategies mainly fall into three types (Figure 1): (1) **contrastive decoding**, which amplifies visual contributions by contrasting outputs generated under different visual conditions (e.g., PAI (Liu et al., 2024d), IBD (Zhu et al., 2024), VCD (Leng et al., 2024)). These methods require

multiple forward passes and may introduce additional bias from the contrastive perturbed signal. (2) **auxiliary expert models**, which leverage external models to provide auxiliary cues or highlight salient regions (e.g., HALC (Chen et al., 2024b), AGLA (An et al., 2025)), at the cost of introducing external dependencies and potential semantic misalignment with the target LVLM. and (3) **internal static signals**, which enhance vision-related heads, layers, or tokens by selecting top-ranked elements based on attention values or other heuristic scores (e.g., VHR (He et al., 2025), VAF (Yin et al., 2025)). Relying on static criteria such as top- k selection or thresholding, these methods are highly vulnerable to the attention sink phenomenon (Xiao et al., 2023; Kang et al., 2025), where dominant but semantically irrelevant sink tokens are repeatedly amplified, biasing attention away from truly informative visual regions.

To address these limitations, we revisit a key question: *how can semantically core visual regions be reliably identified and enhanced in the presence of attention sink distortions, without relying on external models or inputs?* Our key finding is that positive attention dynamics across layers reveal semantically core visual regions (Figure 2). Core regions exhibit stronger positive inter-layer attention changes, while irrelevant regions remain weakly attended, and attention sinks show irregular fluctuations. Leveraging these internal Positive Attention Dynamics (PAD) allows us to reveal visual evidence that emerges coherently through the model’s internal understanding process, thereby enabling more reliable identification of semantically core visual regions without external auxiliary models.

Motivated by this observation, we propose Positive Attention Dynamics Enhancement (PADE), a training-free attention intervention that selectively reinforces semantically core visual regions to improve visual grounding and mitigate hallucinations. PADE identifies core regions by leveraging the model’s internal positive attention dynamics and enhances them in the target layer, without relying on external models or multiple forward passes. Specifically, PADE constructs a Positive Attention Dynamics (PAD) map from positive inter-layer attention deltas to highlight semantically core regions. To adaptively control the intervention strength, PAD is scaled per attention head using the Median Absolute Deviation (MAD), ensuring robustness to extreme values while preserving proportionality to the underlying signal. Finally, we pro-

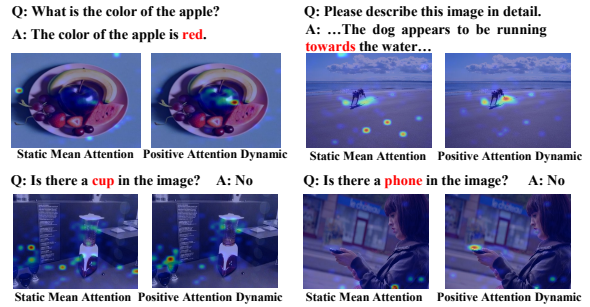


Figure 2: Static versus dynamic internal attention signals. Static mean attention is dominated by attention sinks, while Positive Attention Dynamics (PAD) more reliably highlight semantically core visual regions.

pose System-Token Compensation (STC), which leverages system tokens with high attention ratios but limited semantic relevance as a compensation source, thereby preserving attention to complex instructions and maintaining long-term generation.

In summary, our contributions are as follows:

- We demonstrate that internal Positive Attention Dynamics (PAD) provide a more reliable signal for identifying semantically core visual regions than static signal-based metrics, especially under the distortions of attention sinks.
- We propose Positive Attention Dynamics Enhancement (PADE), a training-free attention intervention that leverages PAD to identify and selectively reinforce semantically core visual regions during inference.
- Extensive experiments on both hallucination-focused and general-purpose benchmarks show that PADE effectively improves visual grounding and reduces hallucinations, while preserving overall multimodal understanding.

2 Related Work

LVLMs. LVLMs (Dai et al., 2023; Liu et al., 2023; Bai et al., 2023; Anthropic, 2024; Achiam et al., 2023; Ye et al., 2024; Hu et al., 2024a; Liu et al., 2024c; Chen et al., 2024c; Yang et al., 2025a; Touvron et al., 2023; Chiang et al., 2023; Lu et al., 2024; DeepSeek-AI et al., 2025) achieve strong performance across a wide range of vision-language tasks (Jiang et al., 2024; Hu et al., 2024b; Xu et al., 2026b; Zhu et al., 2023; Lyu et al., 2025b; Xu et al., 2024, 2026a; Lin et al., 2024). Despite this progress, hallucination remains a fundamental challenge for LVLMs (Lee et al., 2018; Gunjal et al., 2024; Liu et al., 2024b; Woo et al., 2025).

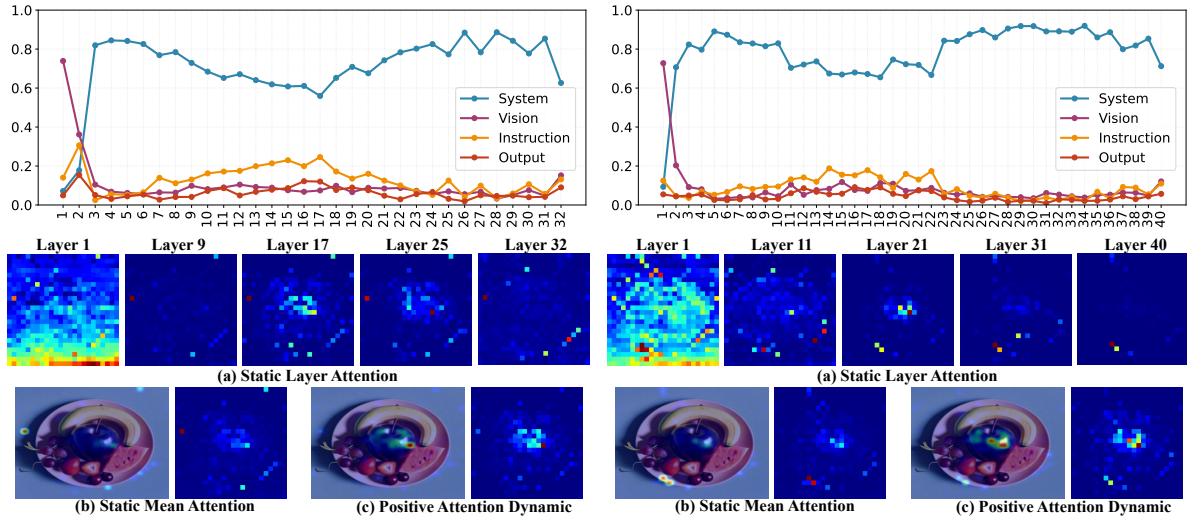


Figure 3: Attention analysis of LLaVA-1.5-7B (left) and 13B (right). Top: the attention ratio of different token types (System, Vision, Instruction, Output). Bottom: heatmap visualizations of attention distributions, including (a) static attention from uniformly sampled layers, (b) layer-averaged static attention, and (c) positive attention dynamics.

Hallucinations in LVLMs. Existing mitigation approaches can be broadly categorized into two classes (Yan et al., 2026; Gong et al., 2026; Khaki et al., 2025; Khaki et al.; Yang et al., 2026; Liu et al., 2026). **Training-based methods** aim to reduce hallucinations by strengthening modality alignment and robustness, typically through improved data curation, alignment objectives, retrieval-augmented generation, or reinforcement learning (Liu et al., 2024a; Yu et al., 2024a,b; Chen et al., 2025; Ouali et al., 2024). These approaches require substantial computational resources and retraining costs, limiting their practicality and flexibility. **Training-free methods** intervene directly in the decoding process by manipulating logits, attention, or hidden states, and can be grouped into three categories (Yang et al., 2025b; Zhang et al., 2024; Park et al., 2025; Kim et al., 2024; Huo et al., 2024; Zhou et al., 2024; Lyu et al., 2025a; Li et al., 2025; Zhang et al., 2025; Huangyw et al., 2025; Huang et al., 2025; Tang et al., 2025; Jiang et al., 2025b,a; Darcet et al., 2024; Liu et al., 2025b,a; Wang et al., 2025, 2024a; Xing et al., 2024; Jiang et al.). (1) *Contrastive decoding* contrasts outputs under different visual conditions, such as PAI (Liu et al., 2024d), IBD (Zhu et al., 2024), and VCD (Leng et al., 2024). These methods require multiple forward passes and may introduce bias from the contrastive perturbed signal. (2) *Auxiliary expert models* leverage external models to provide auxiliary cues (e.g., HALC (Chen et al., 2024b), AGLA (An et al., 2025), Wood-

pecker (Yin et al., 2023)), which over-rely on external models and may not be aligned with the target LVLM. (3) *Static internal signal* methods select and amplify vision-related heads, layers, or tokens based on heuristic scores or internal signals (e.g., VAR (Kang et al., 2025), VAF (Yin et al., 2025), OPERA (Huang et al., 2023), MemVR (Zou et al., 2025)). These methods are vulnerable to the *attention sink or massive activations* (Sun et al., 2024b; Kang et al., 2025), where semantically irrelevant but dominant tokens are repeatedly amplified. Concurrent work GIFT (Qi et al., 2025) also leverages attention variations, but it relies on part-of-speech analysis to identify key words, which may fail for generic prompts without explicit key objects (e.g., “Please describe this image in detail”), and selects specific visual heads for intervention. In contrast, PADE extracts semantic core regions directly from the LVLM’s internal attention dynamics, uses MAD to adaptively control intervention strength, and employs STC to preserve both complex user instruction and historical long-term outputs.

3 Analysis: Internal Attention Dynamics

In this section, we analyze the internal attention behaviors of LVLMs to understand how attention is distributed and evolves across layers during multimodal understanding and reasoning. We systematically study different token types attention distributions and visual attention evolution in LLaVA-1.5 models at two scales (7B and 13B), as shown in Figure 3, with additional qualitative examples

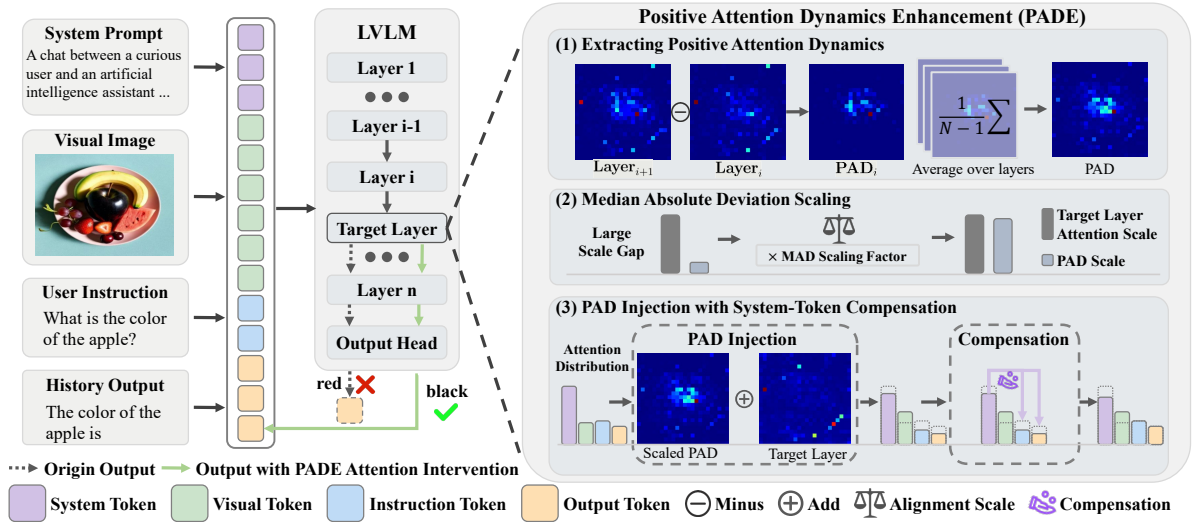


Figure 4: Overview of our PADE. PADE identifies semantically core visual regions via Positive Attention Dynamics (PAD) and selectively enhances them in the target layer, with Median Absolute Deviation Scaling for adaptively controlling the intervention strength and System-Token Compensation to preserve attention for understanding complex instructions and ensuring consistent long-term generation.

provided in Section A under diverse visual contexts and prompt variations.

3.1 Attention Sink Dominance in LLMs

As shown in Figure 3, both LLaVA-1.5-7B and 13B exhibit a highly imbalanced attention distribution, where a small set of tokens consistently absorb a disproportionate amount of attention mass. These tokens, commonly referred to as *attention sinks* (Xiao et al., 2023; Kang et al., 2025; Sun et al., 2024b), persist across layers despite limited semantic relevance to the visual content.

3.2 Static Metrics Are Vulnerable to Sinks

Attention sinks manifest as isolated extreme activations with large magnitudes, absolute attention maps become heavily skewed. Semantically irrelevant sink tokens are often assigned high attention scores, while genuinely informative visual regions are comparatively suppressed. This distortion fundamentally limits prior attention intervention methods (Yin et al., 2025; Wan et al., 2025) that rely on internal static signals. Since sink tokens frequently rank among the highest-attended elements, they are repeatedly selected and further amplified, biasing intervention toward spurious structures and degrading visual grounding reliability.

3.3 Attention Dynamics Reveal Core Regions

As illustrated in Figure 3, we observe distinguishable attention dynamics among semantically core

regions, attention sinks, and irrelevant background areas. Semantically meaningful regions exhibit multiple instances of pronounced positive attention changes across layers, even though their attention may occasionally decrease. In contrast, most irrelevant regions maintain consistently low attention with only minor fluctuations. Attention sink tokens can display large or sporadic spikes, but these changes are irregular and do not align consistently with semantic understanding and reasoning.

Motivated by these observations, we quantify and highlight semantically core visual regions using *Positive Attention Dynamics (PAD)*, computed as the *positive inter-layer attention deltas* between consecutive layers. Semantically core regions often exhibit repeated substantial increases in attention, which typically correspond to decreases in other regions. By retaining only positive deltas and differential formulation, PAD naturally suppresses attention in less relevant areas, and inherently mitigates the influence of most attention sinks. As a result, PAD provides a more reliable signal for identifying semantically core visual regions than static attention maps, by leveraging the model’s internal attention evolution, without requiring external models or auxiliary cues (Figures 2 and 3).

4 Method: PADE

Building on our analysis, we propose Positive Attention Dynamics Enhancement (PADE), a training-free attention intervention that selectively rein-

forces semantically core visual regions to improve visual grounding and mitigate hallucinations in LVLMs. PADE identifies core regions by leveraging the model’s internal *positive attention dynamics* and enhances them in the target layer, without relying on external models or multiple forward passes. As illustrated in Figure 4, PADE consists of three key steps: (1) extracting *Positive Attention Dynamics (PAD)* to identify semantically core visual regions, (2) applying per-head *Median Absolute Deviation (MAD)* scaling to adaptively control intervention strength, and (3) introducing *System-Token Compensation (STC)* to preserve attention to complex instructions and long-term generation.

4.1 Positive Attention Dynamics

To enable effective attention intervention, we first identify semantically core visual regions that should be selectively reinforced. Following our analysis in Section 3, such core visual regions tend to receive positive attention increases as the model refines its internal understanding, whereas irrelevant regions rarely exhibit small changes and attention sinks show irregular spikes. Motivated by this observation, we extract Positive Attention Dynamics (PAD), which aggregates positive inter-layer attention deltas to reveal semantically core visual regions. By retaining only positive deltas, PAD emphasizes regions whose importance increases during reasoning, while naturally suppressing noisy fluctuations and attention sinks.

Let \mathbf{A}_l denote the visual attention map at layer l , averaged over all attention heads. The positive inter-layer attention delta is defined as

$$\Delta^+ \mathbf{A}_l = \max(0, \mathbf{A}_l - \mathbf{A}_{l-1}), \quad l = 2, \dots, L. \quad (1)$$

We aggregate these deltas across layers to obtain the PAD:

$$\mathbf{P} = \frac{1}{L-1} \sum_{l=2}^L \Delta^+ \mathbf{A}_l. \quad (2)$$

4.2 Per-Head MAD Scaling

The PAD is injected into attention logits rather than post-softmax attention, preserving the inherent properties of attention and avoiding an additional softmax operation. However, attention logits often contain extreme outliers induced by sink tokens and operate at a much larger scale than PAD, with significant variation across samples. Without proper scaling, the same intervention coefficient λ

can produce inconsistent perturbation magnitudes, leading to poorly calibrated interventions.

We adaptively control the intervention strength by scaling each attention head with the median absolute deviation (MAD), which uses the median instead of the mean to provide robust per-head, per-sample calibration and reduce the influence of extreme attention sink values.

Let $\mathbf{Z}_{l,h}^v$ denote the visual attention logits of head h at layer l . The MAD is computed as

$$\text{MAD}(\mathbf{Z}_{l,h}^v) = \text{median}(|\mathbf{Z}_{l,h}^v - \text{median}(\mathbf{Z}_{l,h}^v)|), \quad (3)$$

and the PAD is scaled accordingly:

$$\hat{\mathbf{P}}_{l,h} = \text{MAD}(\mathbf{Z}_{l,h}^v) \cdot \tilde{\mathbf{P}}. \quad (4)$$

4.3 System-Token Compensation (STC)

After scaling PAD to match the magnitude of attention logits, we inject it into the visual attention logits of the target layer:

$$\hat{\mathbf{Z}}_{l,h}^v \leftarrow \mathbf{Z}_{l,h}^v + \lambda \cdot \hat{\mathbf{P}}_{l,h}, \quad (5)$$

where λ controls the intervention strength.

Directly increasing visual attention will reduce attention to user instructions or previously generated outputs, potentially impairing instruction following and output coherence for complex or long-form tasks. Formally, attention is computed over different token groups as

$$\mathbf{A} = \text{softmax}([\mathbf{Z}^s, \mathbf{Z}^v, \mathbf{Z}^i, \mathbf{Z}^o]), \quad (6)$$

where \mathbf{Z}^s , \mathbf{Z}^v , \mathbf{Z}^i , and \mathbf{Z}^o denote the logits of system tokens, user instruction tokens, visual tokens, and output tokens, respectively.

As shown in Figure 3, system tokens consistently receive a large proportion of attention and remain largely unrelated to the semantic content of user instructions or visual inputs. Based on this observation, we introduce System-Token Compensation (STC), which leverages the high-attention but semantically irrelevant system tokens to compensate for the increased visual attention, rather than affecting instruction or history tokens. Specifically, system-token logits are adjusted as

$$\check{\mathbf{Z}}^s \leftarrow \mathbf{Z}^s - \text{mean}(\lambda \cdot \hat{\mathbf{P}}_{l,h}), \quad (7)$$

allowing selective enhancement of semantically core visual regions while preserving attention to user instructions and previously generated outputs. The final attention weights are computed as

$$\hat{\mathbf{A}} = \text{softmax}([\check{\mathbf{Z}}^s, \hat{\mathbf{Z}}^v, \mathbf{Z}^i, \mathbf{Z}^o]). \quad (8)$$

Table 1: Results on the POPE (Accuracy and F1). \uparrow indicates that higher is better. Best results are **bolded**.

Setup	Method	LLaVA-1.5		InstructBLIP		Qwen-VL		LLaVA-1.5-13B		LLaVA-Next	
		Accuracy \uparrow	F1 \uparrow	Accuracy \uparrow	F1 \uparrow	Accuracy \uparrow	F1 \uparrow	Accuracy \uparrow	F1 \uparrow	Accuracy \uparrow	F1 \uparrow
Random	Vanilla	84.63	84.99	83.33	83.57	85.17	83.00	83.27	84.27	84.23	82.14
	VCD	84.57	85.02	84.60	84.49	84.69	82.91	83.47	84.52	83.48	81.36
	PAI	85.12	85.64	83.82	83.98	85.63	83.74	84.06	84.98	83.76	81.52
	VAF	85.64	85.82	85.12	85.28	85.93	84.08	84.32	85.28	84.84	82.64
	VAR	86.12	86.56	85.62	85.71	86.54	84.66	84.82	85.76	85.12	83.26
	PADE (ours)	86.96	87.42	86.52	86.82	87.14	85.68	86.08	87.16	86.24	84.74
Popular	Vanilla	81.33	82.33	76.00	77.94	84.50	82.50	80.57	82.19	82.33	80.44
	VCD	81.57	83.02	76.68	77.82	84.37	82.46	80.96	82.47	82.46	80.68
	PAI	81.82	83.46	76.82	78.56	84.83	82.76	80.78	82.74	82.76	80.82
	VAF	82.62	84.18	77.36	79.02	86.12	84.28	81.62	83.46	83.57	81.68
	VAR	83.16	85.52	78.12	79.96	86.89	84.92	82.74	84.62	84.62	82.74
	PADE (ours)	84.56	86.28	78.74	80.76	87.72	86.12	83.82	85.28	85.56	84.12
Adversarial	Vanilla	75.87	78.27	74.17	76.58	82.53	80.56	75.12	78.35	79.37	77.88
	VCD	75.76	78.12	74.62	76.72	82.68	80.78	75.46	78.42	79.56	78.12
	PAI	76.12	78.53	74.82	76.96	82.83	80.82	75.72	78.64	79.68	78.14
	VAF	76.72	79.14	75.48	77.73	83.52	81.38	76.63	79.12	80.23	78.74
	VAR	77.43	80.16	76.22	78.16	83.96	81.82	77.52	79.46	80.92	79.56
	PADE (ours)	78.47	81.12	77.26	79.13	85.12	83.08	78.62	80.62	81.08	80.92

Table 2: Results on CHAIR (Max Token 128). \downarrow denotes lower is better. – denotes unavailable results.

Method	LLaVA-1.5		InstructBLIP		Qwen-VL		LLaVA-1.5-13B		LLaVA-Next	
	CHAIR _S \downarrow	CHAIR _I \downarrow	CHAIR _S \downarrow	CHAIR _I \downarrow	CHAIR _S \downarrow	CHAIR _I \downarrow	CHAIR _S \downarrow	CHAIR _I \downarrow	CHAIR _S \downarrow	CHAIR _I \downarrow
Vanilla	55.1	16.4	57.4	17.6	52.1	16.7	50.4	14.7	30.2	10.9
VCD	54.4	16.6	60.7	18.0	50.4	17.2	49.6	14.3	29.8	10.6
M3ID	56.6	15.8	62.3	18.2	49.8	17.4	—	—	—	—
Woodpecker	57.6	16.7	60.8	17.6	51.8	16.3	—	—	—	—
HALC	51.0	14.8	53.8	15.7	49.6	15.4	—	—	—	—
ONLY	49.8	14.3	52.2	15.5	48.0	14.3	48.9	13.6	28.6	9.7
AGLA	52.4	14.6	54.8	16.2	49.8	15.6	—	—	—	—
OPERA	51.6	14.2	54.2	14.8	48.6	14.6	—	—	—	—
PAI	53.1	15.1	54.2	15.6	49.1	15.6	49.4	14.1	29.4	10.4
VAF	50.1	14.2	53.4	15.1	48.7	14.4	48.6	13.4	28.4	9.6
VAR	49.6	14.1	52.3	14.7	48.4	13.9	48.2	13.2	28.2	9.2
PADE (ours)	48.6	13.7	51.8	14.2	47.8	13.4	47.3	12.6	27.8	8.9

5 Experiments

Benchmarks. To comprehensively evaluate PADE, we conduct extensive experiments on two categories of benchmarks: hallucination-focused and general-purpose multimodal benchmarks. The hallucination-focused benchmarks include POPE (Li et al., 2023) (binary hallucination classification), CHAIR (Rohrbach et al., 2018) (object hallucination in open-ended captioning), HallusionBench (Guan et al., 2024) (fine-grained visual consistency), and AMBER (Wang et al., 2023) (visually grounded reasoning and generation). The general-purpose benchmarks include VizWiz (Gurari et al., 2018), MME (Fu et al., 2023), LLaVA-Wild (Liu et al., 2023), and MM-Vet (Yu et al., 2023), covering diverse tasks in visual understanding, reasoning, and real-world multimodal scenarios.

Evaluated LLMs. We evaluate our PADE on sev-

eral representative open-source LLMs, including LLaVA-1.5 (Liu et al., 2023), InstructBLIP (Dai et al., 2023), and Qwen-VL (Bai et al., 2023). We additionally consider a larger scale variant, LLaVA-1.5-13B, as well as the stronger and newer LLaVA-NeXT (Liu et al., 2024c). Together, these models cover diverse backbone architectures and model scales, providing a comprehensive evaluation setting. Following prior works (Leng et al., 2024; Wan et al., 2025), we apply sampling-based decoding in default. Unless otherwise specified, LLaVA-1.5 is used as the default model.

Baselines. We compare PADE with various training-free hallucination mitigation methods: (1) *contrastive decoding* methods (VCD (Leng et al., 2024), PAI (Liu et al., 2024d), M3ID (Favero et al., 2024), ICD (Wang et al., 2024b), DoLA (Chuang et al., 2023)); (2) *auxiliary expert model* methods (HALC (Chen et al., 2024b), AGLA (An et al.,

Table 3: Results on multiple general vision language benchmarks. \uparrow indicates that higher is better.

Methods	VizWiz	MME			LLaVA-Wild	MM-Vet
	Accuracy \uparrow	Perception \uparrow	Cognition \uparrow	Overall \uparrow	Average \uparrow	Total \uparrow
Vanilla	50.00	1508.97	355.71	1864.68	64.80	31.1
VCD	44.90	1515.01	357.86	1872.87	63.21	30.2
ICD	37.62	1306.91	287.86	1594.77	56.90	25.9
OPERA	50.76	1473.62	310.71	1784.34	64.31	32.0
INTER	48.77	1502.35	336.18	1838.53	61.70	30.9
V-ITI	51.72	1518.32	369.03	1887.35	65.44	31.7
PADE (ours)	52.08	1520.68	371.44	1892.12	65.92	32.4

Table 4: Results on the HallusionBench.

Methods	fACC \uparrow	qACC \uparrow	easyA \uparrow	hardA \uparrow
Vanilla	17.9	8.13	36.0	36.7
VCD	13.9	11.4	33.0	34.7
ICD	13.9	8.35	36.9	33.5
OPERA	16.2	5.49	37.6	35.4
INTER	15.8	8.21	36.9	34.5
V-ITI	17.9	10.27	36.6	37.0
PADE (ours)	18.1	11.56	37.9	37.4

Table 5: Results on AMBER Generative Subset.

Method	CHAIR (\downarrow)	Cover (\uparrow)	Hall (\downarrow)	Cog (\downarrow)
Vanilla	7.8	51.0	36.4	4.2
VCD	7.5	50.8	36.2	4.1
OPERA	7.3	49.6	32.0	3.5
DoLA	7.6	51.6	36.0	4.0
PAI	7.4	49.9	33.2	3.7
PADE (ours)	7.1	51.8	31.4	3.4

2025), Woodpecker (Yin et al., 2023), V-ITI (Sun et al., 2025b)); and (3) *static internal signal* methods (OPERA (Huang et al., 2023), VAF (Yin et al., 2025), INTER (Dong et al., 2025), ONLY (Wan et al., 2025), VAR (Kang et al., 2025)).

Implementation Details. All experiments are conducted on a single NVIDIA RTX A6000 GPU (48GB). Unless otherwise specified, PADE is applied to the final layer with an intervention strength of $\lambda = 0.1$. PADE introduces negligible computational and memory overhead. The method maintains only a single additional attention map and relies on lightweight operations, including inter-layer differencing and MAD scaling. It requires neither auxiliary models nor multiple forward passes, achieving inference speed comparable to vanilla decoding while improving visual grounding and mitigating hallucinations.

Table 6: Ablation results on the proposed components.

Model	Components	CHAIR _S \downarrow	CHAIR _T \downarrow
LLaVA-1.5-7B	PADE	48.6	13.7
	w/o MAD	54.9	16.3
	w/o STC	49.2	14.0
LLaVA-1.5-13B	PADE	47.9	12.8
	w/o MAD	50.3	14.5
	w/o STC	48.6	13.3
LLaVA-1.5-NeXT	PADE	27.8	8.9
	w/o MAD	30.1	10.7
	w/o STC	28.3	9.2
Qwen-VL	PADE	47.3	12.6
	w/o MAD	51.8	16.5
	w/o STC	48.1	13.7
InstructBLIP	PADE	51.8	14.2
	w/o MAD	57.2	17.3
	w/o STC	52.2	14.5

5.1 Main Experimental Results

Results on Hallucination Benchmarks. As shown in Tables 1, 2, 4 and 5, PADE consistently achieves superior performance across different model architectures and scales on diverse comprehension and generation hallucination benchmarks, demonstrating its effectiveness in mitigating hallucinations. Compared with contrastive decoding methods (PAI, VCD, IBD, ICD), which introduce perturbed visual inputs that may disrupt semantic alignment, and auxiliary expert approaches (AGLA, HALC, Woodpecker), which rely on external models or conditions not necessarily aligned with the target LLM, PADE directly leverages the model’s internal positive attention dynamics to identify and reinforce semantically core visual regions. In contrast to static internal signal methods (VAF, OPERA, VAR), which are either sensitive to attention sinks or merely reallocate attention, PADE simultaneously enhances the global visual attention ratio while emphasizing semantically core regions, leading to more effective hallucination mitigation.

Results on General Benchmarks. As shown in Table 3, many existing hallucination mitigation methods rely on contrastive decoding or auxiliary expert models, which use perturbed images, instructions, or external experts to guide outputs. While effective for hallucination reduction, these approaches often compromise general multimodal understanding and reasoning capabilities, as the interventions are not naturally aligned with the LVLM’s internal reasoning. In contrast, PADE leverages the LVLM’s own internal attention dynamics, selectively reinforcing semantically core visual regions. This allows it to reduce hallucinations while preserving the model’s inherent broad visual reasoning and multimodal understanding, demonstrating a more reliable attention intervention.

5.2 Ablation Studies

Effect of Proposed Components. We evaluate the contribution of each component in PADE on the CHAIR benchmark across multiple LVLMs. Specifically, we consider two variants: (i) *w/o MAD* and (ii) *w/o STC*. As shown in Table 6, removing either component consistently degrades performance, while the full PADE achieves the best results across all models. Notably, removing MAD leads to a substantial degradation in performance, as the raw PAD signal is several orders of magnitude smaller than the original attention logits. Without proper scaling, the intervention becomes largely ineffective and the model behavior closely resembles the vanilla baseline, underscoring the necessity of robust and adaptive scale calibration. These results confirm that MAD and STC play roles in ensuring PADE’s effectiveness and stability.

Intervention Layer. We study the impact of the intervention layer using LLaVA-1.5 7b and 13b on the CHAIR benchmark. As shown in Figure 5, the effectiveness of PADE generally improves when applied to later layers, with the best performance achieved at the final layer. This trend can be attributed to the evolution of visual attention across layers: the model typically attends to semantically core visual regions in intermediate layers, while attention becomes more diffuse and spreads to less relevant regions in later layers. Although the overall visual attention ratio is already high at the final layer, it is often dominated by non-core regions. Injecting PADE at this stage effectively re-emphasizes semantically core regions while suppressing dispersed attention, leading to the most pronounced improvement. Moreover, the




Figure 5: Ablation results on the intervention layer and strength λ of LLaVA-1.5-7B (top) and 13B (bottom).

final layer naturally allocates less attention to system tokens and more to visual tokens and historical outputs, making them more receptive to targeted visual enhancement.

Intervention Strength λ . We study the effect of the intervention strength λ on LLaVA-1.5 7b and 13b using the CHAIR benchmark. As shown in Figure 5, PADE achieves better performance with relatively small values of λ (e.g., 0.1 and 0.3), with $\lambda=0.1$ performing best overall. Moderate intervention strengths effectively enhance semantically core visual regions while preserving the original attention distribution, whereas excessively large λ introduces overly strong interventions and perturbations that deviate from the model’s learned attention dynamics and degrade performance.

6 Conclusion

In this paper, we investigate the role of internal positive attention dynamics in LVLMs and show that they can reveal semantically core visual regions under attention sink distortions. Based on this finding, we propose Positive Attention Dynamics Enhancement (PADE), a lightweight, training-free intervention that constructs a PAD map to identify core visual regions, applies per-head Median Absolute Deviation (MAD) scaling to adaptively adjust the intervention strength, and leverages a System-Token Compensation (STC) module to maintain attention to complex user instructions and support long-term output consistency. Extensive experiments demonstrate that PADE improves visual grounding and reduces hallucinations, validating the effectiveness of exploiting internal signal dynamics.

7 Limitations

PADE emphasizes internal dynamic signals rather than static magnitude-based criteria, leveraging the evolution of attention across layers to identify semantically core regions. While our analysis demonstrates the effectiveness of attention dynamics for inference-time intervention, the current study focuses exclusively on attention-based mechanisms. Other internal representations, such as the dynamics of hidden states, activation patterns in feed-forward networks, or variations in output logits during decoding, are not explored. These components may also encode complementary signals related to semantic grounding and hallucination behavior. Extending the dynamic analysis framework beyond attention to encompass a broader range of internal model signals remains an important direction for future work.

8 Acknowledgments

Our work is supported in part by the National Key R&D Program of China (No. 2023YFC3305600) and the National Natural Science Foundation of China (U25B2048, 62132016).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv:2303.08774*.
- Wenbin An, Feng Tian, Sicong Leng, Jiahao Nie, Haonan Lin, QianYing Wang, Ping Chen, Xiaoqin Zhang, and Shijian Lu. 2025. Mitigating object hallucinations in large vision-language models with assembly of global and local attention. In *CVPR*.
- Anthropic. 2024. [Introducing Claude 3.5 Sonnet](#). Announcement of Claude 3.5 Sonnet model release, featuring improved intelligence, vision capabilities, and new Artifacts feature.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv:2308.12966*.
- Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou. 2024. Hallucination of multimodal large language models: A survey. *arXiv:2404.18930*.
- Jinhe Bi, Yifan Wang, Danqi Yan, Xun Xiao, Artur Hecker, Volker Tresp, and Yunpu Ma. 2025a. [Prism: Self-pruning intrinsic selection method for training-free multimodal data selection](#). *Preprint*, arXiv:2502.12119.
- Jinhe Bi, Yujun Wang, Haokun Chen, Xun Xiao, Artur Hecker, Volker Tresp, and Yunpu Ma. 2025b. Llava steering: Visual instruction tuning with 500x fewer parameters through modality linear representation-steering. In *ACL*.
- Cong Chen, Mingyu Liu, Chenchen Jing, Yizhou Zhou, Fengyun Rao, Hao Chen, Bo Zhang, and Chunhua Shen. 2025. Perturbollava: Reducing multimodal hallucinations with perturbative visual training. In *ICLR*.
- Junying Chen, Chi Gui, Ruyi Ouyang, Anningzhe Gao, Shunian Chen, Guiming Hardy Chen, Xidong Wang, Zhenyang Cai, Ke Ji, Xiang Wan, and 1 others. 2024a. Towards injecting medical visual knowledge into multimodal llms at scale. In *NeurIPS*.
- Zhaorun Chen, Zhuokai Zhao, Hongyin Luo, Huaxiu Yao, Bo Li, and Jiawei Zhou. 2024b. Halc: Object hallucination reduction via adaptive focal-contrast decoding. *arXiv:2403.00425*.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, and 1 others. 2024c. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *CVPR*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality](#).
- 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. In *ICLR*.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *NeurIPS*.
- Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. 2024. Vision transformers need registers. In *ICLR*.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, and 1 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- Xin Dong, Shichao Dong, Jin Wang, Jing Huang, Li Zhou, Zenghui Sun, Lihua Jing, Jinsong Lan, Xiaoyong Zhu, and Bo Zheng. 2025. Inter: Mitigating hallucination in large vision-language models by interaction guidance sampling. In *CVPR*, pages 2534–2544.

- Alessandro Favero, Luca Zancato, Matthew Trager, Sidharth Choudhary, Pramuditha Perera, Alessandro Achille, Ashwin Swaminathan, and Stefano Soatto. 2024. Multi-modal hallucination control by visual information grounding. In *CVPR*.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiawu Zheng, and 1 others. 2023. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv:2306.13394*.
- Boyang Gong, Yu Zheng, Fanye Kong, Jie Zhou, and Jiwen Lu. 2026. Attention at rest stays at rest: Breaking visual inertia for cognitive hallucination mitigation. *arXiv preprint arXiv:2604.01989*.
- Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, and 1 others. 2024. Hallusionbench: an advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. In *CVPR*, pages 14375–14385.
- Anisha Gunjal, Jihan Yin, and Erhan Bas. 2024. Detecting and preventing hallucinations in large vision language models. In *AAAI*, 16.
- Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. 2018. Vizwiz grand challenge: Answering visual questions from blind people. In *CVPR*, pages 3608–3617.
- Jinghan He, Kuan Zhu, Haiyun Guo, Junfeng Fang, Zhenglin Hua, Yuheng Jia, Ming Tang, Tat-Seng Chua, and Jinqiao Wang. 2025. Cracking the code of hallucination in vlms with vision-aware head divergence. In *ACL*.
- Shengding Hu, Yuge Tu, Xu Han, Chaoqun He, Ganqu Cui, Xiang Long, Zhi Zheng, Yewei Fang, Yuxiang Huang, Weilin Zhao, and 1 others. 2024a. Minicpm: Unveiling the potential of small language models with scalable training strategies. *arXiv:2404.06395*.
- Yutao Hu, Tianbin Li, Quanfeng Lu, Wenqi Shao, Junjun He, Yu Qiao, and Ping Luo. 2024b. Omnimedvqa: A new large-scale comprehensive evaluation benchmark for medical lvm. In *CVPR*.
- Lei Huang, Xiaocheng Feng, Weitao Ma, Yuchun Fan, Xiachong Feng, Yuxuan Gu, Yangfan Ye, Liang Zhao, Weihong Zhong, Baoxin Wang, and 1 others. 2025. Alleviating hallucinations from knowledge misalignment in large language models via selective abstention learning. In *ACL*, pages 24564–24579.
- 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:2311.17911*.
- Huangyw Huangyw, Yong Zhang, Ning Cheng, Zhitao Li, Shaojun Wang, and Jing Xiao. 2025. Dynamic attention-guided context decoding for mitigating context faithfulness hallucinations in large language models. In *ACL Findings*, pages 5174–5193.
- Fushuo Huo, Wenchao Xu, Zhong Zhang, Haozhao Wang, Zhicheng Chen, and Peilin Zhao. 2024. Self-introspective decoding: Alleviating hallucinations for large vision-language models. *arXiv:2408.02032*.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*.
- Bo Jiang, Shaoyu Chen, Bencheng Liao, Xingyu Zhang, Wei Yin, Qian Zhang, Chang Huang, Wenyu Liu, and Xinggong Wang. 2024. Senna: Bridging large vision-language models and end-to-end autonomous driving. *arXiv:2410.22313*.
- Nicholas Jiang, Anish Kachinthaya, Suzanne Petryk, and Yossi Gandelsman. Interpreting and editing vision-language representations to mitigate hallucinations. In *ICLR*.
- Nick Jiang, Amil Dravid, Alexei Efros, and Yossi Gandelsman. 2025a. Vision transformers don’t need trained registers. In *NeurIPS*.
- Zhangqi Jiang, Junkai Chen, Beier Zhu, Tingjin Luo, Yankun Shen, and Xu Yang. 2025b. Devils in middle layers of large vision-language models: Interpreting, detecting and mitigating object hallucinations via attention lens. In *CVPR*.
- Seil Kang, Jinyeong Kim, Junhyeok Kim, and Seong Jae Hwang. 2025. See what you are told: Visual attention sink in large multimodal models. *arXiv:2503.03321*.
- Samir Khaki, Junxian Guo, Jiaming Tang, Shang Yang, Yukang Chen, Konstantinos N Platanotis, Yao Lu, Song Han, and Zhijian Liu. Sparsevila-r1: Decoupling visual sparsity for efficient vlm reasoning. In *NeurIPS 2025 Workshop on Efficient Reasoning*.
- Samir Khaki, Junxian Guo, Jiaming Tang, Shang Yang, Yukang Chen, Konstantinos N Platanotis, Yao Lu, Song Han, and Zhijian Liu. 2025. Sparsevila: Decoupling visual sparsity for efficient vlm inference. In *ICCV*, pages 23784–23794.
- Junho Kim, Hyunjun Kim, Kim Yeonju, and Yong Man Ro. 2024. Code: Contrasting self-generated description to combat hallucination in large multi-modal models. *NeurIPS*, 37:133571–133599.
- Katherine Lee, Orhan Firat, Ashish Agarwal, Clara Fan-jiang, and David Sussillo. 2018. Hallucinations in neural machine translation.
- 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 *CVPR*.

- Wei Li, Zhen Huang, Houqiang Li, Le Lu, Yang Lu, Xinmei Tian, Xu Shen, and Jieping Ye. 2025. Visual evidence prompting mitigates hallucinations in large vision-language models. In *ACL*, pages 4048–4080.
- 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:2305.10355*.
- Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, and Li Yuan. 2024. Video-llava: Learning united visual representation by alignment before projection. In *EMNLP*.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. 2024a. Mitigating hallucination in large multi-modal models via robust instruction tuning. In *ICLR*.
- Hanchao Liu, Wenyuan Xue, Yifei Chen, Dapeng Chen, Xiutian Zhao, Ke Wang, Liping Hou, Rongjun Li, and Wei Peng. 2024b. A survey on hallucination in large vision-language models. *arXiv:2402.00253*.
- 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.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. In *NeurIPS*.
- Shi Liu, Kecheng Zheng, and Wei Chen. 2024d. Paying more attention to image: A training-free method for alleviating hallucination in vlms. *arXiv:2407.21771*.
- Shiyu Liu, Xinyi Wen, Zhibin Lan, Ante Wang, and Jinsong Su. 2026. Countering the over-reliance trap: Mitigating object hallucination for vlms via a self-validation framework. *arXiv preprint arXiv:2601.22451*.
- Yexin Liu, Zhengyang Liang, Yueze Wang, Xianfeng Wu, Feilong Tang, Muyang He, Jian Li, Zheng Liu, Harry Yang, Sernam Lim, and 1 others. 2025a. Unveiling the ignorance of mllms: Seeing clearly, answering incorrectly. In *CVPR*.
- Zhining Liu, Ziyi Chen, Hui Liu, Chen Luo, Xianfeng Tang, Suhang Wang, Joy Zeng, Zhenwei Dai, Zhan Shi, Tianxin Wei, and 1 others. 2025b. Seeing but not believing: Probing the disconnect between visual attention and answer correctness in vlms. *arXiv:2510.17771*.
- Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Yaofeng Sun, and 1 others. 2024. Deepseek-vl: towards real-world vision-language understanding. *arXiv:2403.05525*.
- Guangtao Lyu, Xinyi Cheng, Chenghao Xu, Qi Liu, Muli Yang, Fen Fang, Huilin Chen, Jiexi Yan, Xu Yang, and Cheng Deng. 2025a. Revealing perception and generation dynamics in vlms: Mitigating hallucinations via validated dominance correction. *arXiv:2512.18813*.
- Guangtao Lyu, Chenghao Xu, Jiexi Yan, Muli Yang, and Cheng Deng. 2025b. Towards unified human motion-language understanding via sparse interpretable characterization. In *ICLR*.
- Yassine Ouali, Adrian Bulat, Brais Martinez, and Georgios Tzimiropoulos. 2024. Clip-dpo: Vision-language models as a source of preference for fixing hallucinations in vlms. In *ECCV*, pages 395–413.
- Woohyeon Park, Woojin Kim, Jaek Kim, and Jaeyoung Do. 2025. Second: Mitigating perceptual hallucination in vision-language models via selective and contrastive decoding. In *ICML*.
- Zheng Qi, Chao Shang, Evangelia Spiliopoulou, and Nikolaos Pappas. 2025. Capturing gaze shifts for guidance: Cross-modal fusion enhancement for vlm hallucination mitigation. *arXiv:2510.22067*.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. Object hallucination in image captioning. *arXiv:1809.02156*.
- Hao Shao, Yuxuan Hu, Letian Wang, Guanglu Song, Steven L Waslander, Yu Liu, and Hongsheng Li. 2024. Lmdrive: Closed-loop end-to-end driving with large language models. In *CVPR*.
- Fengzhao Sun, Jun Yu, Yunxiang Zhang, Jiaming Hou, Xilong Lu, Heng Song, and Fang Gao. 2025a. Towards robust autonomous driving: Conditional multimodal large language models for fine-grained perception. In *ICRA*.
- Guohao Sun, Can Qin, Huazhu Fu, Linwei Wang, and Zhiqiang Tao. 2024a. Self-training large language and vision assistant for medical question answering. *NeurIPS*.
- Mingjie Sun, Xinlei Chen, J Zico Kolter, and Zhuang Liu. 2024b. Massive activations in large language models. *arXiv:2402.17762*.
- Nan Sun, Zhenyu Zhang, Xixun Lin, Kun Wang, Yanmin Shang, Naibin Gu, Shuohuan Wang, Yu Sun, Hua Wu, Haifeng Wang, and 1 others. 2025b. V-iti: Mitigating hallucinations in multimodal large language models via visual inference-time intervention. *arXiv:2512.03542*.
- Feilong Tang, Chengzhi Liu, Zhongxing Xu, Ming Hu, Zile Huang, Haochen Xue, Ziyang Chen, Zelin Peng, Zhiwei Yang, Sijin Zhou, and 1 others. 2025. Seeing far and clearly: Mitigating hallucinations in mllms with attention causal decoding. In *CVPR*, pages 26147–26159.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. Llama: Open and efficient foundation language models. *arXiv:2302.13971*.

- Zifu Wan, Ce Zhang, Silong Yong, Martin Q Ma, Simon Stepputtis, Louis-Philippe Morency, Deva Ramanan, Katia Sycara, and Yaqi Xie. 2025. Only: One-layer intervention sufficiently mitigates hallucinations in large vision-language models. *arXiv:2507.00898*.
- Chenxi Wang, Xiang Chen, Ningyu Zhang, Bozhong Tian, Haoming Xu, Shumin Deng, and Huajun Chen. 2024a. Mllm can see? dynamic correction decoding for hallucination mitigation. *arXiv:2410.11779*.
- Fei Wang, Xingchen Wan, Ruoxi Sun, Jiefeng Chen, and Sercan O Arik. 2025. Astute rag: Overcoming imperfect retrieval augmentation and knowledge conflicts for large language models. In *ACL*.
- Junyang Wang, Yuhang Wang, Guohai Xu, Jing Zhang, Yukai Gu, Haitao Jia, Jiaqi Wang, Haiyang Xu, Ming Yan, Ji Zhang, and 1 others. 2023. Amber: An llm-free multi-dimensional benchmark for mllms hallucination evaluation. *arXiv:2311.07397*.
- Xintong Wang, Jingheng Pan, Liang Ding, and Chris Biemann. 2024b. Mitigating hallucinations in large vision-language models with instruction contrastive decoding. In *ACL*.
- Sangmin Woo, Donguk Kim, Jaehyuk Jang, Yubin Choi, and Changick Kim. 2025. Don't miss the forest for the trees: Attentional vision calibration for large vision language models. In *ACL Findings*.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2023. Efficient streaming language models with attention sinks. *arXiv:2309.17453*.
- Yun Xing, Yiheng Li, Ivan Laptev, and Shijian Lu. 2024. Mitigating object hallucination via concentric causal attention. *NeurIPS*, 37:92012–92035.
- Chenghao Xu, Lyu Guangtao, Yan Jiexi, Yang Muli, and Cheng Deng. 2024. Llm knows body language, too: Translating speech voices into human gestures. In *ACL*.
- Chenghao Xu, Jiexi Yan, Muli Yang, Fen Fang, Huilin Chen, and Cheng Deng. 2026a. Editing is a bargaining game: Balanced knowledge editing in large language models. In *AAAI*, volume 40, pages 34097–34105.
- Shijia Xu, Zhou Wu, Xiaolong Jia, Yu Wang, Kai Liu, and April Xiaowen Dong. 2026b. Self-correcting rag: Enhancing faithfulness via mmkp context selection and nli-guided mcts. *arXiv:2604.10734*.
- Bei Yan, Yuecong Min, Jie Zhang, Shiguang Shan, and Xilin Chen. 2026. Act now: Preempting lvlm hallucinations via adaptive context integration. *arXiv preprint arXiv:2604.00983*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025a. Qwen3 technical report. *arXiv:2505.09388*.
- Junqi Yang, Yuecong Min, Jie Zhang, Shiguang Shan, and Xilin Chen. 2026. Infact: A diagnostic benchmark for induced faithfulness and factuality hallucinations in video-llms. *arXiv preprint arXiv:2603.11481*.
- Tianyun Yang, Ziniu Li, Juan Cao, and Chang Xu. 2025b. Mitigating hallucination in large vision-language models via modular attribution and intervention. In *Adaptive Foundation Models: Evolving AI for Personalized and Efficient Learning*.
- Qinghao Ye, Haiyang Xu, Jiabo Ye, Ming Yan, Anwen Hu, Haowei Liu, Qi Qian, Ji Zhang, and Fei Huang. 2024. mplug-owl2: Revolutionizing multi-modal large language model with modality collaboration. In *CVPR*.
- Hao Yin, Guangzong Si, and Zilei Wang. 2025. Clear-sight: Visual signal enhancement for object hallucination mitigation in multimodal large language models. In *CVPR*, pages 14625–14634.
- 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:2310.16045*.
- Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwen He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, and 1 others. 2024a. Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. In *CVPR*, pages 13807–13816.
- Tianyu Yu, Haoye Zhang, Yuan Yao, Yunkai Dang, Da Chen, Xiaoman Lu, Ganqu Cui, Taiwen He, Zhiyuan Liu, Tat-Seng Chua, and 1 others. 2024b. Rlaif-v: Aligning mllms through open-source ai feedback for super gpt-4v trustworthiness.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2023. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv:2308.02490*.
- 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:2411.09968*.
- Yongheng Zhang, Xu Liu, Ruoxi Zhou, Qiguang Chen, Hao Fei, Wenpeng Lu, and Libo Qin. 2025. Cchallenge: A novel benchmark for joint cross-lingual and cross-modal hallucinations detection in large language models. *ACL*.
- Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. 2024. Analyzing and mitigating object hallucination in large vision-language models. In *ICLR*.

Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv:2304.10592*.

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:2402.18476*.

Xin Zou, Yizhou Wang, Yibo Yan, Sirui Huang, Ken-ting Zheng, Junkai Chen, Chang Tang, and Xuming Hu. 2025. Look twice before you answer: Memory-space visual retracing for hallucination mitigation in multimodal large language models. In *ICML*.

A Additional Qualitative Examples

We present additional qualitative examples in Figures 6 to 10 to illustrate how *Positive Attention Dynamics (PAD)* reveal semantically core visual regions under diverse and challenging conditions. These examples cover direction hallucinations in open-ended long-term generation, object existence hallucinations, atypical or counterintuitive attribute reasoning (e.g., color), as well as small-object and occlusion scenarios across different user prompts.

Across all cases, a common attention pattern emerges in LVLMs: while the model may initially attend to relevant regions in early or intermediate layers, attention progressively diffuses in later layers and becomes dominated by semantically irrelevant sink tokens, which often maintain abnormally high activations across layers. This behavior underscores the unreliability of static attention maps for identifying core visual regions.

Unlike static attention, which is often skewed by attention sinks and fails to highlight small or occluded objects (e.g., Figures 7 and 8), PAD captures positive inter-layer attention deltas that reliably reflect semantically core regions. By applying PADE at later layers, the model selectively re-emphasizes these core regions (see Figures 6, 9 and 10), thereby improving visual grounding and mitigating hallucinations in multimodal understanding.

Together, these examples demonstrate that PAD highlights semantically core visual regions across varying object scales, scene complexities, and user instructions. In contrast, static attention is often dominated by attention sinks, particularly in later layers, which can obscure meaningful regions.

Q: Please describe this image in detail.

A: ...The dog appears to be running **towards** the water...

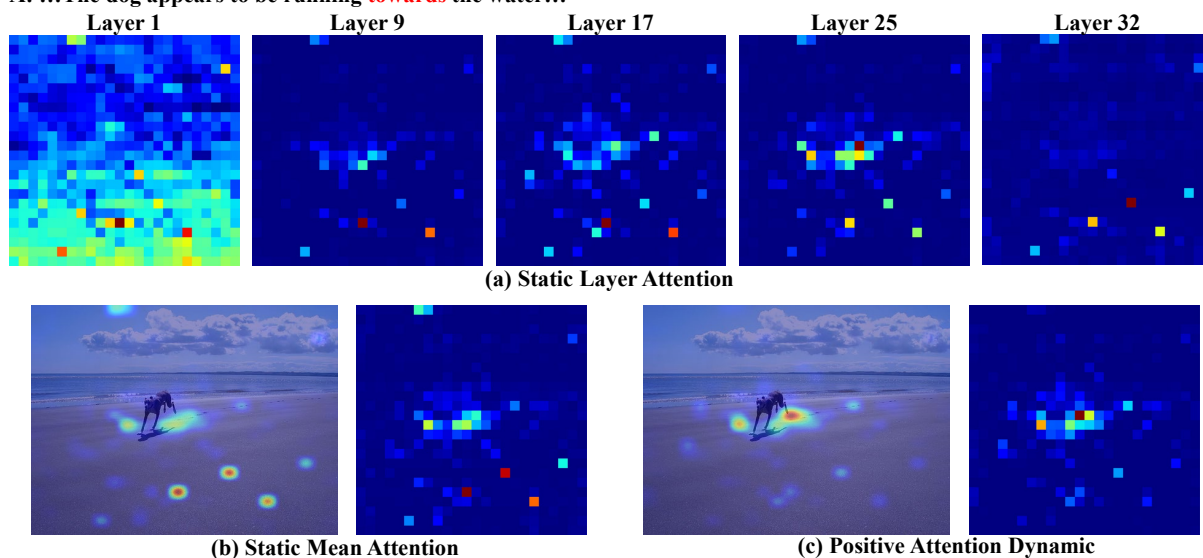


Figure 6: Direction hallucination in open-ended long-term generation. Although the model attends to the dog throughout generation, it incorrectly describes the running direction as *towards* the water, while the dog is in fact moving away from it. A substantial portion of attention is absorbed by several semantically irrelevant sink tokens, especially in later layers, diluting the focus on the dog region. Positive Attention Dynamics (PAD) captures consistent attention changes concentrated around the dog region across layers. When applied at later layers, PADE selectively reinforces this core region after attention diffusion, correcting the directional hallucination.

Q: Is there a **cup** in the image?

A: No.

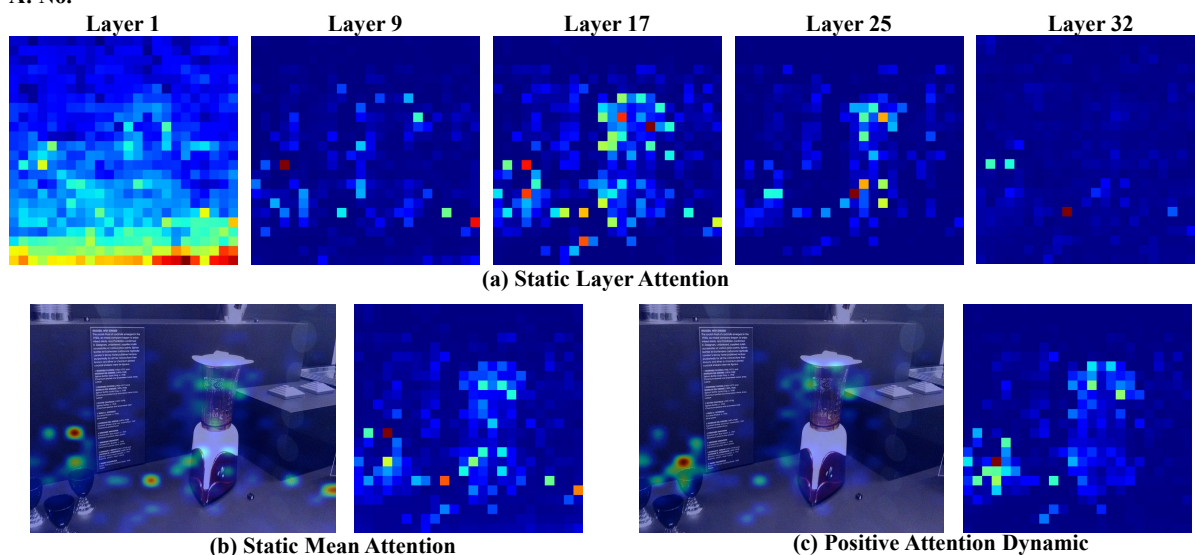


Figure 7: Object existence hallucination in a cluttered scene. The image contains multiple objects, including several cups and a visually dominant bottle. Due to the presence of multiple salient objects, visual attention becomes fragmented, and smaller target objects such as the cups receive substantially less attention than the large bottle. Although both the cups and the bottle are attended to at intermediate layers, attention gradually diffuses away from semantically relevant regions, and sink tokens increasingly dominate the final-layer attention map, causing the cup regions to become inconspicuous. PAD captures the transient attention concentration on the cup regions before diffusion occurs, enabling PADE to re-emphasize these small but semantically critical regions in later layers and correctly recognize the existence of cups, thereby mitigating the hallucinated judgment.

Q: Is there a **phone** in the image?
A: No.

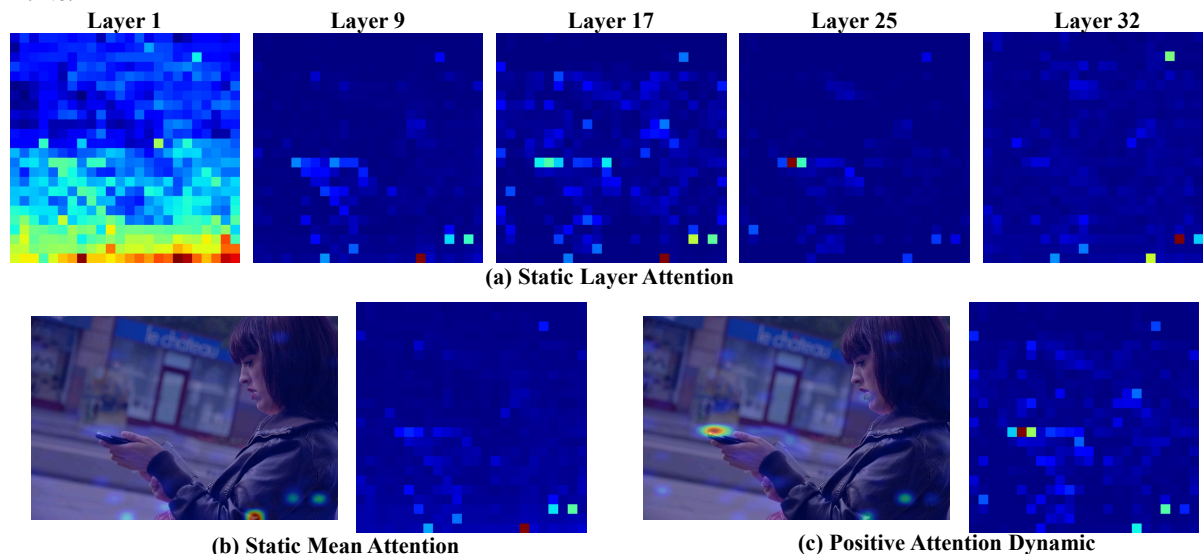


Figure 8: Object existence hallucination for a small target. The phone occupies a relatively small visual region, resulting in weak and sparse attention signals. In later layers, attention disperses from the phone region and becomes highly imbalanced, with sink tokens dominating the distribution. Despite the absence of strong static attention, PAD identifies consistent attention dynamics localized around the phone region. By reinforcing this region at later layers, PADE restores visual grounding and alleviates the hallucination.

Q: Is there a **bag** in the image?
A: No.

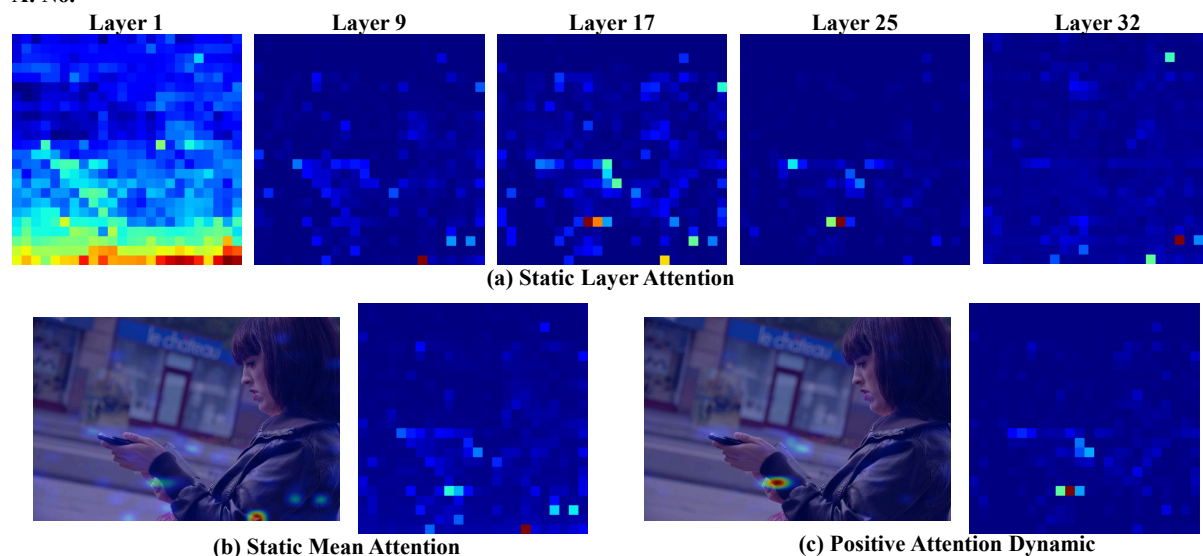


Figure 9: Object existence hallucination under occlusion and prompt variation. The bag is partially occluded and occupies a small visible area, leading to weak static attention signals. As attention diffuses in later layers, sink tokens dominate the distribution, resulting in hallucinated responses. PAD captures coherent attention dynamics around the hand and bag regions across layers and enables PADE to selectively enhance these regions in the target layer. This figure shares the same image as Figure 8 but uses a different user instruction (bag vs. phone). Both prompts exhibit similar sink-dominated attention patterns and hallucination behaviors, indicating that such failures arise from intrinsic LVLM attention dynamics rather than prompt-specific artifacts. PADE consistently leverages attention evolution to recover semantically core regions and mitigate hallucinations across prompts.

(a) Old Output (b) Static Mean Attention (c) Positive Attention Dynamic (d) New Output




Figure 10: Comparison of static mean attention and positive attention dynamics during the generation of different tokens. (a) Original model output. (b) Static mean attention. (c) Positive Attention Dynamics (PAD). (d) Output after applying PADE, where semantically core regions are enhanced via PAD. Across all tokens, PAD more reliably highlights semantically core visual regions than static attention. In particular, when generating tokens corresponding to key objects (e.g., “apple”) and attributes (e.g., “red”), attention changes are most strongly concentrated on the apple region. These results demonstrate that PAD effectively captures semantically relevant visual areas and validate the effectiveness of PADE in reinforcing core visual regions for reliable multimodal understanding.

B Details of Benchmarks

To comprehensively evaluate PADE, we conduct experiments on two categories of benchmarks: *hallucination-focused* and *general-purpose* multimodal benchmarks. Following prior works (Liu et al., 2023; Leng et al., 2024; Sun et al., 2025b), we adopt standard evaluation protocols for each benchmark to measure both hallucination mitigation and overall multimodal understanding performance. Below, we provide a brief overview of each benchmark and its corresponding evaluation focus.

B.1 Hallucination Benchmarks

We use the following benchmarks to evaluate the visual hallucination performance:

- **CHAIR** (Rohrbach et al., 2018): CHAIR evaluates object hallucination in open-ended captioning and generation by comparing objects mentioned in model outputs with ground-truth objects present in the image. We report both instance-level and sentence-level metrics:

$$\text{CHAIR}_I = \frac{|\{\text{hallucinated objects}\}|}{|\{\text{all mentioned objects}\}|}, \quad (9)$$

$$\text{CHAIR}_S = \frac{|\{\text{hallucinated captions}\}|}{|\{\text{generated captions}\}|}. \quad (10)$$

Here, CHAIR_I measures the proportion of hallucinated objects among all mentioned objects, while CHAIR_S measures the fraction of generated outputs that contain at least one hallucinated object. Lower values indicate fewer hallucinations.

- **POPE** (Li et al., 2023): POPE (Polling-based Object Probing Evaluation) assesses object existence hallucination via binary (yes/no) questions. Queries are evenly split between existent and non-existent objects under random, popular, and adversarial settings. We report Accuracy, Precision, Recall, and F1 score, where higher values indicate better object existence discrimination.
- **HallusionBench** (Guan et al., 2024): HallusionBench evaluates fine-grained visual hallucinations in open-ended multimodal responses. It reports question-level accuracy (qACC) and full-response accuracy (fACC), measuring whether answers are fully supported by visual evidence. In addition, easyA

and hardA separately evaluate performance on visually simple and challenging cases, respectively. Higher scores indicate better visual grounding and hallucination robustness.

- **AMBER** (Wang et al., 2023): AMBER focuses on open-ended hallucinations in multimodal reasoning and generation. It reports multiple metrics, including CHAIR-style object hallucination rates, Coverage (Cover), and overall Hallucination rate. It further distinguishes Cognitive Hallucination (Cog.), which measures hallucinations arising from incorrect reasoning rather than missing visual evidence. Lower hallucination-related scores indicate better performance.

B.2 General Benchmarks

To examine whether mitigating hallucinations comes at the cost of general multimodal understanding, we further evaluate PADE on a suite of general-purpose benchmarks, following prior work (Liu et al., 2023; Sun et al., 2025b):

- **VizWiz** (Gurari et al., 2018): VizWiz is a visual question answering benchmark consisting of images captured by blind or low-vision users. The dataset contains diverse real-world challenges such as poor lighting, blur, and occlusion.
- **MME** (Fu et al., 2023): MME is a comprehensive evaluation benchmark for multimodal large language models, covering both perception and reasoning abilities. It includes tasks such as object existence, counting, spatial relations, color recognition, and commonsense reasoning.
- **LLaVA-Bench (In-the-Wild)** (Liu et al., 2023): LLaVA-Wild consists of diverse real-world images paired with open-ended questions spanning various domains. It is designed to assess robustness and generalization under unconstrained scenarios.
- **MM-Vet** (Yu et al., 2023): MM-Vet evaluates a model’s ability to conduct visually grounded conversations across multiple reasoning skills. It covers tasks such as object recognition, spatial reasoning, and commonsense inference.