

# MITIGATING HALLUCINATION IN MULTIMODAL REASONING VIA FUNCTIONAL ATTENTION CONTROL

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Multimodal large reasoning models (MLRMs) are rapidly advancing vision-language reasoning and are emerging as a foundation for cross-modal intelligence. Hallucination remains a persistent failure mode, manifesting itself as erroneous reasoning chains and misinterpretation of visual content. In this study, we observe that attention heads exhibit a staged division: **shallow** heads predominantly serve perception, while **deeper** heads shift toward symbolic reasoning, revealing two major causes of hallucination, namely perceptual bias and reasoning drift. To address these issues, we propose a lightweight and interpretable two-step plugin, Functional Head Identification and Class-conditioned Rescaling, which locates perception- and reasoning-oriented heads and regulates their contributions without retraining. Evaluations on three real-world MLRMs (Kimi-VL, Ocean-R1, R1-Onevision), six benchmarks across three domains, and four baselines show that our plugin achieves an average improvement of **5%** and up to **15%**, with **only <1% additional computation** and **9%** of baseline latency. Our approach is completely model-agnostic and significantly enhances both the reliability and interpretability of the off-the-shelf MLRMs, thereby enabling their safe deployment in high-stakes applications. Our code is available at <https://anonymous.4open.science/r/Functional-Attention-Control>.

## 1 INTRODUCTION

Large language models (LLMs) (Team, 2024; Guo et al., 2025; Team et al., 2025) have rapidly transformed natural language processing, and more recently multimodal large reasoning models (MLRMs) (Xu et al., 2024; Ming et al., 2025; Bai et al., 2025) are extending this revolution to vision-language tasks. By integrating powerful language reasoning with visual understanding, MLRMs are becoming a cornerstone for cross-modal intelligence (Huang et al., 2023), with broad potential for advancing both research (Wang et al., 2024a) and real-world applications (Yin et al., 2024). Despite these advances, a fundamental barrier — hallucinations — still remains (Liu et al., 2024; Wu et al., 2025). Existing models frequently generate conclusions that either conflict with the visual evidence or contradict their own reasoning chains (Bai et al., 2024). Such failures not only diminish the reliability of individual predictions but also erode trust in the reasoning process itself (Ding et al., 2025), posing serious obstacles for deploying MLRMs in high-stakes (Sokol & Vogt, 2023) domains where accountability and interpretability are crucial.

For LLMs, hallucination phenomena are often attributed to limited knowledge coverage (Huang et al., 2025) or decoding bias (Rawte et al., 2023; Ji et al., 2023), and mitigation methods include retrieval-based correction (Lewis et al., 2020), contrastive decoding (Chuang et al., 2023), and self-consistency (Madaan et al., 2023). In multimodal settings, the causes of hallucinations are more complex. The prevailing view is that they mainly stem from *limited use of visual evidence*: MLRMs may overlook details during encoding or lose critical information during cross-modal alignment (Guan et al., 2024; Li et al., 2023b). As a result, most existing designs attempt to “compensate for vision” through stronger supervision (Sun et al., 2024; Liu et al., 2023; Xie et al., 2024), finer-grained alignment (Peng et al., 2023; Chen et al., 2023a), or the use of external visual priors (Zhao et al., 2025b; Kim et al., 2024). While these approaches reduce hallucinations to some extent, they share a common assumption: *the root cause lies mainly in under-utilized visual information*.

Building on this assumption, recent work (Huang et al., 2024; Kang et al., 2025a) reveals that hallucinations also involve **deeper allocation dynamics** inside the model. Interpretability studies (Bi

054  
055  
056  
057  
058  
059  
060  
061  
062  
063  
064  
065  
066  
067  
068  
069  
070  
071  
072  
073  
074  
075  
076  
077  
078  
079  
080  
081  
082  
083  
084  
085  
086  
087  
088  
089  
090  
091  
092  
093  
094  
095  
096  
097  
098  
099  
100  
101  
102  
103  
104  
105  
106  
107

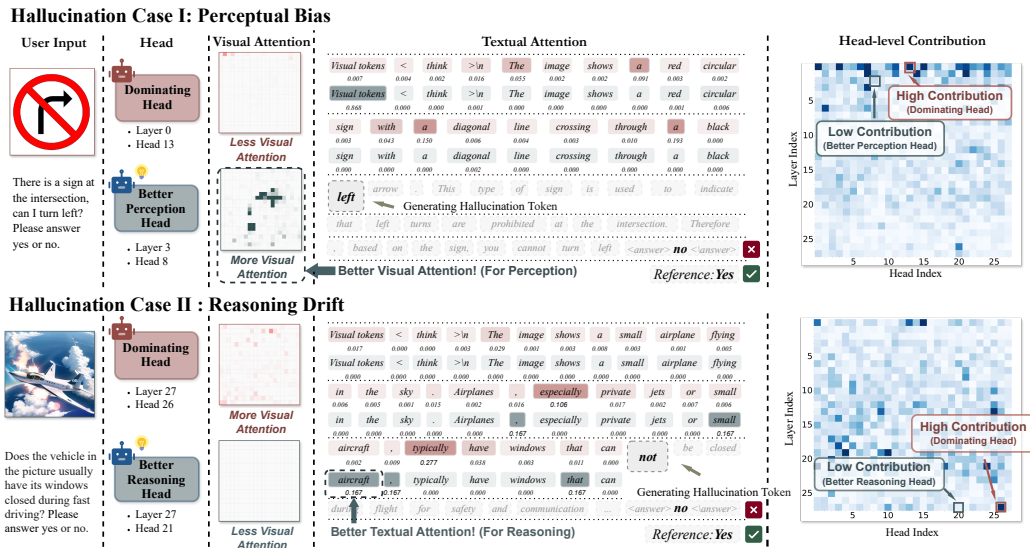


Figure 1: Two hallucination examples corresponding to perception layers (Cause I) and reasoning layers (Cause II). The figure highlights attention patterns over text and image tokens, and the contribution of tokens at the first position where hallucination emerges.

et al., 2025; Li et al., 2023a) show that MLRMs exhibit a staged division of attention: shallow layers rely on visual signals to extract evidence, while deeper layers increasingly shift toward symbolic reasoning over text. This staged allocation suggests that failures may arise not only from limited perception but also from misaligned usage, reflecting an imbalance in how vision and language are allocated and utilized across layers. Naturally, this raises a key research question: *under off-the-shelf architectures, how can we effectively identify and regulate the use of visual and textual information at different stages to mitigate hallucinations in MLRM’s reasoning?*

To tackle this issue, we first decompose hallucinations in MLRMs into two primary failure causes. The first is **Perceptual Bias** (Pan et al., 2024; Zhao et al., 2025a), which occurs in *shallow layers* when attention over visual tokens becomes diffuse and critical evidence is diluted. The second is **Reasoning Drift** (Becker et al., 2025), which arises in *deeper layers* when attention fails to preserve intermediate steps, causing conclusions to deviate from established premises. To address these two causes, we proceed from the premise that models may already contain attention heads with the potential to support perception and reasoning (Jiang et al., 2025), but in their current form, these heads do not play a dominant role (Gong et al., 2024; Zhang et al., 2024a). Consequently, mitigating hallucinations requires identifying such heads and amplifying their contributions, so that perception and reasoning can be more effectively regulated across stages.

In this paper, we propose a lightweight and interpretable two-step plugin: **Functional Head Identification** and **Class-Conditioned Rescaling**. ✦ The goal of **Functional Head Identification** is to isolate heads that naturally specialize in perception or reasoning, so that they can be explicitly leveraged rather than treated uniformly. To this end, we materialize attention weights, compute modality-specific attention ratios, and apply depth-aware boundaries to categorize heads into perception-oriented or reasoning-oriented groups. ✦ The goal of **Class-Conditioned Rescaling** is to amplify the contributions of these functional heads, thereby counteracting perceptual bias and reasoning drift without altering the underlying attention mechanism. This is achieved by assigning targeted multiplicative gains to the identified heads and rescaling their outputs in a structured yet minimally invasive manner.

**Experimental Takeaways.** We conduct systematic experiments on three real-world MLRMs (Kimi-VL, Ocean-R1, R1-Onevision), applying **Identification** and **Rescaling** to analyze their impact on perception and reasoning performance. In the ✦ **Identification Step**, we evaluate over 150 boundary configurations across perception and reasoning stages, finding a 27.4% ↑ between the best and worst settings, an observation that validates the necessity of stage-wise reasoning and highlights clear functional boundaries. During the ✦ **Rescaling Step**, we evaluate over 24 scaling strategies and observe that a moderate factor of 1.14 consistently achieved the best trade-off, under-

scoring that virtually all attention heads may carry latent contributions. **✦ Overall Performance** indicates that across 6 benchmarks spanning three domains and against four baseline methods, our approach delivered an average improvement of 8% over the original MLRM, with gains reaching up to 20% in the most challenging tasks. Moreover, our scheme is *plug-and-play* without retraining, reaching SOTA performance with only  $\sim 4\%$  extra computation, while introducing merely 7% of the latency compared to the best performance baseline.

## 2 MOTIVATION

### 2.1 MODALITY-INDEXED ATTENTION IN TRANSFORMERS

We begin by formalizing how Transformer attention distributes across modalities. Consider a sequence of  $N$  tokens indexed by  $\{1, \dots, N\}$ , partitioned into disjoint sets  $\mathcal{T}_v$  (vision) and  $\mathcal{T}_t$  (text), with  $\mathcal{T}_v \cup \mathcal{T}_t = \{1, \dots, N\}$  and  $\mathcal{T}_v \cap \mathcal{T}_t = \emptyset$ . Unless otherwise noted, we use the entire sequence as the query set, i.e.,  $\mathcal{T}_q = \{1, \dots, N\}$ .

For layer  $\ell$  and head  $h$ , the standard multi-head attention mechanism computes

$$A^{(h,\ell)} = \text{Softmax}_{\text{row}} \left( \frac{(X^{(\ell)} W_Q^{(h,\ell)})(X^{(\ell)} W_K^{(h,\ell)})^\top}{\sqrt{d_k}} \right) \in \mathbb{R}^{N \times N}, \quad (1)$$

where  $X^{(\ell)}$  denotes the input hidden states at layer  $\ell$  ( $N$  tokens, each of dimension  $d$ ),  $W_Q^{(h,\ell)}$ ,  $W_K^{(h,\ell)}$  are the query and key projection matrices, and  $d_k$  is the key dimension used for scaling. The value projection is  $V^{(h,\ell)} = X^{(\ell)} W_V^{(h,\ell)}$ , and the per-head output is  $O^{(h,\ell)} = A^{(h,\ell)} V^{(h,\ell)}$ , which are concatenated across heads and projected to yield the layer output  $Y^{(\ell)}$ .

### 2.2 PERCEPTION AND REASONING OF LMMS

Most large multimodal models (LMMS; e.g., Qwen, LLaVA) and their long-reasoning extensions (MLRMs; e.g., Ocean-R1, Kimi) are architecturally composed of a vision encoder coupled with a language model. This design naturally suggests a division of labor: the encoder **perceives** by compressing raw images into visual tokens, while the LM **reasons** by integrating those tokens with linguistic knowledge to produce answers. If this intuition is mirrored in internal computation, attention dynamics should follow the same workflow: *early layers emphasize visual tokens, whereas deeper layers progressively shift focus toward textual tokens.*

Many empirical studies support this view. An early analysis (Kang et al., 2025b; Zhang et al., 2024c) shows that pruning heads with high visual allocation leads to larger performance drops in vision-conditioned tasks than pruning random heads, with such “visual heads” especially common in shallow and middle layers. Complementarily, a subsequent study (Bi et al., 2025) reports a consistent layerwise trajectory: early layers show a high visual ratio with low concentration (broad scanning), middle layers keep a high ratio while concentration rises (focusing and cross-modal alignment), and late layers reduce the visual ratio as linguistic cues dominate. Together, the results indicate a stratified **perceive–then–reason** pipeline.

Motivated by these findings, we introduce a layerwise abstraction with two boundaries:  $\ell_{\text{perc}}$  denotes the last layer where perception dominates,  $\ell_{\text{reas}}$  denotes the first layer where reasoning dominates.

$$\mathcal{L}_{\text{perc}} = \{1, \dots, \ell_{\text{perc}}\}, \quad \mathcal{L}_{\text{reas}} = \{\ell_{\text{reas}}, \dots, L\}, \quad 1 \leq \ell_{\text{perc}}, \ell_{\text{reas}} \leq L, \quad \ell_{\text{perc}} \leq \ell_{\text{reas}}. \quad (2)$$

In  $\mathcal{L}_{\text{perc}}$  (**perception layers**), attention is biased toward visual tokens and gradually sharpens its focus, enabling the extraction and structuring of visual evidence. In  $\mathcal{L}_{\text{reas}}$  (**reasoning layers**), attention shifts toward textual tokens, allowing linguistic context to dominate inference and generation. Notably,  $\ell_{\text{perc}}$  and  $\ell_{\text{reas}}$  need not coincide: the two sets may overlap (if some layers support both perception and reasoning) or leave a gap (if neither dominates in the intermediate depth).

This flexible abstraction captures the empirical transition from perception to reasoning. It also provides the basis for our hallucination analysis in Section 2.3.

### 2.3 HALLUCINATION IN PERCEPTION AND REASONING

Given the **perception–reasoning** pipeline established in Section 2.2, it is natural to analyze hallucination along these two stages. Figure 1 illustrates how failures in **perception** and **reasoning layers** give rise to two distinct but complementary forms of hallucination, which can be summarized as:

❖ **Cause I (Perceptual Bias).** Perception layers fail to provide sufficient visual evidence because it fails to allocate sufficient attention to the most informative regions.

In  $\mathcal{L}_{\text{perc}}$ , attention is expected to highlight salient cues so that later stages can build on accurate premises. However, we observe that dominant attention frequently under-represent critical regions, which may leaving the perceptual foundation incomplete and weakening subsequent reasoning.

❖ **Cause II (Reasoning Drift).** Reasoning layers fail to retain prior evidence, as attention often diffuses toward image tokens instead of focusing on critical steps in the reasoning chain.

In  $\mathcal{L}_{\text{reas}}$ , effective reasoning requires sustained attention to intermediate inference tokens that support the conclusion. Yet we find that dominating heads are sometimes distracted by residual visual signals, causing them to overlook core textual cues and thereby weakening the consolidation of prior evidence.

Taken together, *perceptual bias* and *reasoning drift* interact synergistically. Perception layers marginalize accurate visual cues before they reach the deeper network, while reasoning layers amplify irrelevant or weakly grounded signals. As illustrated in Figure 1, the compounding effect of these two failures greatly increases the likelihood of hallucination, leading to outputs that contradict the model’s own intermediate evidence.

## 3 METHOD

Illustrated in Figure 2, we directly target the two root causes of hallucination identified in Section 2.3: *perceptual bias* in shallow and *reasoning drift* in deeper layers. The key idea is to reweight attention heads in a *depth-aware* manner, amplifying early heads that correctly capture visual evidence and late heads that preserve reasoning chains, while leaving the standard attention computation intact.

Concretely, the method proceeds in two stages. (I) *Functional Head Identification*: we explicitly materialize attention weights, compute the visual-attention ratio  $S_v^{(\ell)}(h)$ , and use the perception and reasoning boundaries ( $\ell_{\text{perc}}, \ell_{\text{reas}}$ ) to identify perception- and reasoning-oriented heads. (II) *Class-conditioned Rescaling*: we assign multiplicative gains ( $g_{\text{perc}}$  or  $g_{\text{reas}}$ ), lift them to a head-aligned scaling vector, and apply this vector to rescale the concatenated head outputs.

### 3.1 FUNCTIONAL HEAD IDENTIFICATION

The first step of our method is to identify attention heads that function primarily as **perception heads** or **reasoning heads**, following the perceive–then–reason pipeline outlined in Section 2.2. This categorization enables targeted reweighting rather than treating all heads uniformly. For each layer  $\ell$  and head  $h$ , we denote the attention matrix by  $A^{(h,\ell)} \in \mathbb{R}^{N \times N}$ . Its  $(i, j)$ -th entry is written as

$$a_{ij}^{(h,\ell)} = [A^{(h,\ell)}]_{ij}, \quad \sum_{j=1}^N a_{ij}^{(h,\ell)} = 1, \quad \forall i. \quad (3)$$

Here,  $a_{ij}^{(h,\ell)}$  is the normalized attention weight from query token  $i$  to key token  $j$ . In other words, each row of  $A^{(h,\ell)}$  forms a probability distribution over all keys, and  $a_{ij}^{(h,\ell)}$  quantifies how much information query  $i$  draws from token  $j$ .

To measure how head  $h$  distributes attention across modalities, we compute the *modality attention ratio*. Using  $\mathcal{T}_v$  and  $\mathcal{T}_t$  to denote the sets of visual and textual tokens, and  $\mathcal{T}_q$  as the set of query positions, we define the *modality attention ratio*:

$$S_m^{(\ell)}(h) = \frac{1}{|\mathcal{T}_q|} \sum_{i \in \mathcal{T}_q} \sum_{j \in \mathcal{T}_m} a_{ij}^{(h,\ell)}, \quad m \in \{v, t\}. \quad (4)$$

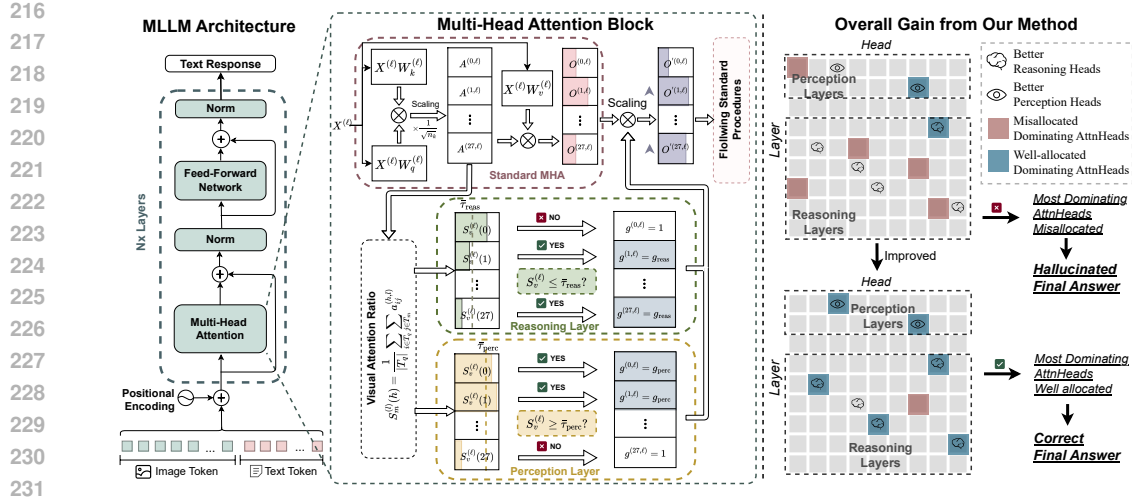


Figure 2: Overall architecture. The attention module is edited by computing the Visual Attention Ratio, which determines the rescaling of specific heads. This process promotes more effective heads to become dominating heads, guiding the output toward correct perception and reasoning. The softmax function is not explicitly shown in the figure.

where  $m = v$  corresponds to visual and  $m = t$  to textual tokens. Intuitively,  $S_v^{(\ell)}(h)$  measures the fraction of attention mass allocated to visual tokens, averaged over all queries, while  $S_t^{(\ell)}(h)$  does the same for textual tokens. Since  $\mathcal{T}_v$  and  $\mathcal{T}_t$  partition the sequence, we have  $S_v^{(\ell)}(h) + S_t^{(\ell)}(h) = 1$ . A larger  $S_v^{(\ell)}(h)$  indicates *stronger visual focus*, while a smaller value indicates *stronger textual focus*.

We then combine these attention ratios with depth information to stratify heads. Specifically, we adopt two ratio thresholds,  $\tau_{\text{perc}}$  and  $\tau_{\text{reas}}$  with  $\tau_{\text{reas}} < \tau_{\text{perc}}$ , which distinguish heads that strongly attend to visual tokens (above  $\tau_{\text{perc}}$ ) from those that focus on text (below  $\tau_{\text{reas}}$ ). In addition, we use the perception and reasoning layer boundaries  $\ell_{\text{perc}}, \ell_{\text{reas}}$  introduced in Section 2.2 to restrict these two head types to shallow and deep layers, respectively. Formally, for each layer  $\ell$  we define

$$\mathcal{H}_{\text{perc}}^{(\ell)} = \{h : \ell \leq \ell_{\text{perc}} \wedge S_v^{(\ell)}(h) \geq \tau_{\text{perc}}\}, \quad \mathcal{H}_{\text{reas}}^{(\ell)} = \{h : \ell \geq \ell_{\text{reas}} \wedge S_v^{(\ell)}(h) \leq \tau_{\text{reas}}\}. \quad (5)$$

Heads that fall between thresholds or outside two-layer boundaries remain unlabeled.

Concretely,  $\mathcal{H}_{\text{perc}}^{(\ell)}$  denotes shallow heads focusing on visual tokens, while  $\mathcal{H}_{\text{reas}}^{(\ell)}$  represents deeper heads attending to textual tokens. At the layer level, the former strengthens visual grounding in perception layers ( $\ell \leq \ell_{\text{perc}}$ ), and the latter reinforces inference consistency in reasoning layers ( $\ell \geq \ell_{\text{reas}}$ ). Thus, enhancing perception heads mitigates **perceptual bias**, and enhancing reasoning heads counters **reasoning drift**, motivating our targeted rescaling strategy.

### 3.2 CLASS-CONDITIONED RESCALING

Building on Section 3.1, we target **perception-oriented heads in shallow layers** and **reasoning-oriented heads in deep layers**, denoted by  $\mathcal{H}_{\text{perc}}^{(\ell)}$  and  $\mathcal{H}_{\text{reas}}^{(\ell)}$ .

Our guiding principle is to intervene on as few heads as possible while maximizing impact on the final outcome. Prior work (Nam et al., 2025; Michel et al., 2019), as well as our case analysis in Figure 2, shows that only a small number of *dominating heads* exert decisive influence on reasoning, while the majority have limited effect. Consequently, the most effective strategy is to increase the likelihood that members of  $\mathcal{H}_{\text{perc}}^{(\ell)}$  and  $\mathcal{H}_{\text{reas}}^{(\ell)}$  enter this dominating subset. Detail in Appx. B.

Based on this insight, we introduce two global gains  $g_{\text{perc}} \geq 1$  and  $g_{\text{reas}} \geq 1$  for perception and reasoning heads, respectively, while leaving all other heads neutral with a factor of 1. Formally, for each layer  $\ell$  and head  $h$ , we define

$$g^{(h,\ell)} = g_{\text{perc}} \cdot \mathbf{1}[h \in \mathcal{H}_{\text{perc}}^{(\ell)}] + g_{\text{reas}} \cdot \mathbf{1}[h \in \mathcal{H}_{\text{reas}}^{(\ell)}] + 1 \cdot \mathbf{1}[h \notin \mathcal{H}_{\text{perc}}^{(\ell)} \cup \mathcal{H}_{\text{reas}}^{(\ell)}], \quad (6)$$

where  $\mathbf{1}[\cdot]$  is the *indicator function*, which returns 1 if true and 0 otherwise. Thus,  $g^{(h,\ell)}$  equals  $g_{\text{perc}}$  for perception heads,  $g_{\text{reas}}$  for reasoning heads, and 1 for all other heads. This compact formulation makes explicit that each head is either amplified according to its functional role or left unchanged.

For each layer  $\ell$ , we apply gains after per-head outputs are computed but before the output projection. Let  $Y^{(\ell)} = \text{Concat}(O^{(1,\ell)}, \dots, O^{(H,\ell)}) W_O^{(\ell)}$  be the standard output, the rescaled output is:

$$Y_{\text{out}}^{(\ell)} = \text{Concat}(g^{(1,\ell)} O^{(1,\ell)}, \dots, g^{(H,\ell)} O^{(H,\ell)}) W_O^{(\ell)}. \quad (7)$$

Conceptually, class-conditioned rescaling enhances perception heads in shallow layers and reasoning heads in deeper layers, thereby addressing the two major failure modes: reinforcing visual grounding to mitigate *Perceptual Bias*, and strengthening logical consistency to reduce *Reasoning Drift*.

Crucially, because the rescaled output  $Y_{\text{out}}^{(\ell)}$  is injected into the residual stream and normalized before passing forward, these adjustments do not remain local. Their influence compounds across subsequent layers, meaning that amplifying the correct functional heads shapes the model’s global reasoning trajectory. In this way, class-conditioned rescaling provides a lightweight yet interpretable mechanism to strengthen the macro-level functions most critical for reducing hallucination.

The design is intentionally minimal: both class gains are *amplifying* and *layer-agnostic*, aiming to provide a lightweight and implementation-friendly means to emphasise screened heads without altering the underlying attention weights or value projections.

## 4 EXPERIMENTS

In this section, we conduct experiments to address the following research questions:

- ❖ **RQ1:** Can the proposed method consistently reduce hallucinations across multimodal reasoning tasks and benchmarks, demonstrating gains over baselines?
- ❖ **RQ2:** How do the amplification of perception and reasoning heads contribute to overall effectiveness?
- ❖ **RQ3:** How do different configurations of layer boundaries ( $\ell_{\text{perc}}$  and  $\ell_{\text{reas}}$ ) affect the overall performance of the model?
- ❖ **RQ4:** How do the visual-attention ratio thresholds ( $\tau_{\text{perc}}$  and  $\tau_{\text{reas}}$ ) impact the head screening process and the method’s overall performance?

**Environment.** Experiments with Ocean-R1 use a 24-core Xeon CPU and an A40 GPU, while R1-Onevision and Kimi-VL use a 16-core Xeon CPU and an A800 GPU. Details in Appx. D.

**Datasets.** We evaluate on **6 benchmarks**: ❶ **Mathematics Reasoning:** MathVista<sub>mini</sub> (Lu et al., 2023), MathVision<sub>mini</sub> (Wang et al., 2024b); ❷ **Visual Reasoning:** CLEVR (Johnson et al., 2017), HallusionBench (Guan et al., 2024); ❸ **Multimodal Integration:** MMStar (Chen et al., 2024), SEED-Bench (Li et al., 2024). Details in Appx. E.

**Baselines.** We compare with **3 inference-time baselines**: VCD (Leng et al., 2024) contrasts original vs. counterfactual image views; AGLA (An et al., 2025) fuses global and local views to improve grounding; CGD (Deng et al., 2024) applies CLIP-guided priors during decoding. Details in Appx. F.

### 4.1 EFFECTIVENESS (RQ1)

To evaluate the effectiveness of our proposed plugin for MLRMs, we compare it against three state-of-the-art hallucination mitigation baselines as well as the vanilla backbone models. The results, summarized in Table 1, cover six benchmarks spanning mathematics reasoning, visual reasoning, and multimodal integration. In *Obs. 1* and *Obs. 2*, we further analyze accuracy and efficiency, confirming that our *plug-and-play* design delivers performance gains with *near-zero* inference overhead.

❖ **Obs. 1: Jointly optimizing perception and reasoning yields broad improvements.** Our method achieves best performance on nearly 95% of the tasks, with an average gain of 4.8% over the vanilla baseline. Unlike prior approaches, which often trade improvements between reasoning (mathematics) and perception (visual)




Figure 3: Efficiency comparison. Our method achieves the best efficiency while simultaneously improving accuracy. The x-axis reports  $(\text{Acc}/\text{AvgBatchTime})^2$  computed over 200 samples from HallusionBench with Kimi-VL.

Table 1: **Overall results.** Performance across six benchmarks covering mathematics, visual reasoning, and multimodal integration. **Bold** indicate the best results, and underlined numbers indicate the second best.

Method	Model	Mathematics Reasoning		Visual Reasoning			Multimodal Integration			
		MathVista <sub>mini</sub>	MathVision <sub>mini</sub>	HallusionBench	CLEVR	MMStar		SEED-Bench		
		Acc	Acc	Acc	F1	Acc	Acc	F1	Acc	F1
Vanilla	Kimi-VL A3B-Thinking	63.48±0.91	56.24±1.42	64.76±0.52	64.97±0.98	76.77±1.03	59.76±0.49	42.75±1.08	66.26±0.41	49.54±1.12
VCD		63.51±0.82	56.14±0.79	<u>65.68±0.62</u>	<u>67.91±0.54</u>	<u>78.01±1.23</u>	59.48±0.61	42.61±1.12	66.52±1.53	49.69±0.82
CGD		53.18±1.55	49.3±0.80	52.13±0.81	55.73±1.06	58.23±1.33	48.73±1.15	31.09±0.75	43.98±1.59	34.88±1.38
AGLA		<u>67.32±1.11</u>	<u>58.88±1.47</u>	61.36±1.11	63.51±1.23	71.73±1.40	<u>63.76±0.85</u>	<u>47.06±1.01</u>	<u>69.27±1.39</u>	<u>51.48±0.99</u>
Ours		<b>69.78±0.82</b>	<b>60.54±0.98</b>	<b>68.19±1.07</b>	<b>68.32±0.80</b>	<b>81.73±1.30</b>	<b>66.49±1.32</b>	<b>49.78±0.63</b>	<b>69.74±1.18</b>	<b>53.62±1.22</b>
Vanilla	Ocean-R1 7B-Instruct	54.58±0.87	20.05±0.80	49.41±0.61	49.45±1.60	38.26±1.50	45.24±0.41	29.38±1.12	59.76±0.92	42.61±0.96
VCD		54.51±1.44	19.92±1.34	<u>50.57±0.52</u>	<u>50.66±0.88</u>	<u>39.52±1.21</u>	45.48±0.96	29.23±0.43	59.98±1.30	42.87±1.01
CGD		46.81±0.88	10.4±0.96	33.72±0.56	33.34±1.06	20.77±1.43	36.24±1.45	21.51±1.25	37.76±0.59	17.55±1.29
AGLA		<u>57.49±1.09</u>	<u>23.69±0.73</u>	45.2±0.92	45.07±0.78	33.51±1.29	<u>48.48±0.53</u>	<u>32.27±0.76</u>	<u>63.01±1.04</u>	<u>45.95±0.61</u>
Ours		<b>59.32±0.88</b>	<b>26.01±0.50</b>	<b>53.64±1.58</b>	<b>53.77±0.49</b>	<b>43.01±1.50</b>	<b>50.77±1.14</b>	<b>34.11±0.99</b>	<b>66.51±1.07</b>	<b>49.66±1.33</b>
Vanilla	R1-Onevision 7B	59.92±1.13	33.54±0.72	58.26±1.59	58.49±1.55	47.98±1.00	56.26±0.53	39.11±0.61	68.48±1.53	52.06±0.67
VCD		59.69±1.34	33.44±1.37	<u>58.96±0.53</u>	<u>59.03±0.54</u>	<u>52.24±0.90</u>	<u>56.74±0.62</u>	<u>39.61±1.25</u>	<u>68.76±0.78</u>	<u>52.31±1.36</u>
CGD		52.58±1.56	28.12±1.19	48.89±1.47	48.82±1.16	34.73±1.48	43.48±1.28	26.87±1.53	63.23±0.60	46.58±0.87
AGLA		<b>60.21±1.48</b>	<u>37.03±1.58</u>	55.69±0.76	55.75±0.44	50.51±1.41	54.49±0.49	37.86±1.42	68.11±1.17	51.65±1.48
Ours		<u>60.09±1.48</u>	<b>39.12±0.44</b>	<b>60.77±0.98</b>	<b>60.88±0.54</b>	<b>62.77±1.02</b>	<b>58.02±0.83</b>	<b>40.94±0.49</b>	<b>69.52±1.52</b>	<b>53.01±1.46</b>

tasks, our patch consistently enhances both sides. For example, VCD yields clear gains on visual benchmarks such as HallusionBench but shows limited benefits on mathematics reasoning tasks, whereas our method provides more balanced improvements across both categories. Cases in Appx. H.

❖ **Obs. 2: Dual-stage gains come with negligible efficiency cost.** As illustrated in Figure 3, our patch introduces only marginal overhead: on HallusionBench (200 samples), vanilla models take  $\sim 101$ s on average, while our method adds merely  $\sim 2$ s (103s total). In contrast, baseline methods (VCD, CGD, AGLA) incur  $1.2\times - 6.6\times$  total inference time. This makes our approach highly pluggable and cost-efficient, substantially improving reliability without sacrificing deployability. A theoretical analysis and more results are provided in Appx. C.

#### 4.2 ABLATION ON HEAD AMPLIFICATION (RQ2)

To isolate the effect of rescaling each functional group, we conduct ablations as summarized in Table 2. Concretely, we evaluate: (i) **w/o reasoning**, which identifies and enhances only *perception* heads; and (ii) **w/o perception**, which identifies and enhances only *reasoning* heads. All other hyperparameters are kept identical to the full plugin to ensure a fair comparison.

❖ **Obs. 3: Perception-only vs. reasoning-only effects are task-skewed and non-additive.** In most settings, **w/o reasoning** (enhancing only perception heads) impacts visual tasks more, while **w/o perception** (enhancing only reasoning heads) drives the gains in mathematics. However, we also observe counterexamples: on R1-OneVision with MathVision<sub>mini</sub>, strengthening a single group alone leads to a  $-3.91\%$  drop, whereas combining both groups yields a  $+5.58\%$  gain. This aligns with our claim in Sec. 2.3 that hallucination is not a single-capability failure but emerges from the *interaction* between perception and reasoning stages.

❖ **Obs. 4: Model-specific reliance on perception vs. reasoning.** On multimodal integration tasks, we observe heterogeneous behaviors across architectures. For instance, on Kimi-VL with MMStar, **w/o reasoning** yields an  $+6.71\%$  improvement, whereas on Ocean-R1 the same setting causes a  $-1.51\%$  drop. This suggests that different MLRM architectures may emphasize distinct functional capacities, highlighting the importance of model-specific balancing between perception and reasoning.

#### 4.3 IMPACT OF BOUNDARY CONFIGURATIONS (RQ3)

Our plugin first identifies perception and reasoning heads, followed by rescaling. A key factor is the choice of layer boundaries ( $\ell_{\text{perc}}$ ,  $\ell_{\text{reas}}$ ), which determines which heads are classified as perception- or reasoning-oriented. Figure 4 reports results on three representative datasets, each drawn from one of the evaluation categories, showing how performance varies under different boundary configurations. Based on these results, we adopt  $\ell_{\text{perc}} = 7$  and  $\ell_{\text{reas}} = 3$  as our final configuration.

Table 2: **Ablation of perception- and reasoning-head rescaling.** We compare the full plugin with variants that disable either perception or reasoning rescaling, alongside the vanilla model.

Method	Model	Mathematics Reasoning		Visual Reasoning			Multimodal Integration			
		MathVista <sub>mini</sub>	MathVision <sub>mini</sub>	HallusionBench	CLEVR	MMStar		SEED-Bench		
		Acc	Acc	Acc	F1	Acc	Acc	F1	Acc	F1
Vanilla	Kimi-VL A3B- Thinking	63.48±0.91	56.24±1.42	64.76±0.52	64.97±0.98	76.77±1.03	59.76±0.49	42.75±1.08	66.26±0.41	49.54±1.12
w/o Reason		69.2±0.59	58.54±0.77	65.78±0.90	65.94±1.15	<b>82.27±0.49</b>	66.47±1.44	49.76±0.58	68.49±1.46	52.45±0.78
w/o Percept		69.37±0.62	<b>62.84±0.61</b>	65.32±0.42	65.49±0.66	81.48±0.89	63.23±0.78	46.31±0.61	69.49±1.14	53.62±0.85
Ours		<b>69.78±0.82</b>	60.54±0.98	<b>68.19±1.07</b>	<b>68.32±0.80</b>	81.73±1.30	<b>66.49±1.32</b>	<b>49.78±0.63</b>	<b>69.74±1.18</b>	<b>53.64±1.22</b>
Vanilla	Ocean-R1 7B-Instruct	54.58±0.87	20.05±0.80	49.41±0.61	49.45±1.60	38.26±1.50	45.24±0.41	29.38±1.12	59.76±0.92	42.61±0.96
w/o Reason		55.11±0.71	26.01±0.58	48.15±1.14	48.16±1.55	40.01±0.86	43.73±1.51	28.18±1.43	60.52±0.70	43.31±0.75
w/o Percept		58.22±0.66	24.35±1.60	<b>54.37±1.34</b>	<b>54.47±0.67</b>	41.01±0.82	50.48±0.73	33.93±1.15	61.99±0.98	44.83±1.26
Ours		<b>59.32±0.88</b>	<b>25.68±0.50</b>	53.64±1.58	53.77±0.49	<b>43.01±1.50</b>	<b>50.77±1.14</b>	<b>34.11±0.99</b>	<b>66.51±1.07</b>	<b>49.66±1.33</b>
Vanilla	R1- Onevision 7B	59.92±1.13	33.54±0.72	58.26±1.59	58.49±1.55	47.98±1.00	56.26±0.53	39.11±0.61	68.48±1.53	52.06±0.67
w/o Reason		59.19±1.18	33.9±1.53	60.66±0.62	60.83±1.23	48.77±0.85	57.77±1.07	40.67±1.38	71.23±0.50	55.41±1.30
w/o Percept		59.29±1.38	29.63±1.12	59.42±0.70	59.51±0.64	55.99±0.42	53.73±0.48	36.82±1.47	68.51±0.93	52.21±0.71
Ours		<b>60.09±1.48</b>	<b>39.12±0.44</b>	<b>60.77±0.98</b>	<b>60.88±0.54</b>	<b>62.77±1.02</b>	<b>58.02±0.83</b>	<b>40.94±0.49</b>	<b>69.52±1.52</b>	<b>53.01±1.46</b>




Figure 4: **Boundary sweep on Ocean-R1.** We fix four hyperparameters ( $g_{\text{reas}}=1.3$ ,  $\tau_{\text{reas}}=0.01$ ,  $g_{\text{perc}}=1.16$ ,  $\tau_{\text{perc}}=0.22$ ) and vary the layer boundaries to assess performance sensitivity.

❖ **Obs. 5: Task-dependent boundary bands rather than a single split.** For *visual-dominant* tasks (CLEVR, SEED-Bench), the best scores concentrate in the lower-left of the boundary grid, pointing to  $l_{\text{perc}} \approx 5-9$  and  $l_{\text{reas}} \approx 2-5$ . Operationally, layers 0-5 act as perception, and  $\gtrsim 9$  as reasoning. By contrast, the *reasoning-dominant* task (MathVision) peaks when  $l_{\text{reas}} \approx 10$ , forming a high-performance region that *does not overlap* with the visual-task region. We hypothesize that this band marks a *perception-to-reasoning transition zone*, where control shifts from visual evidence aggregation to symbolic inference. Hence, the boundaries behave like *task-dependent bands*, with a mid-depth transition zone rather than a single crisp threshold.

❖ **Obs. 6: Deep-layer reasoning exists but is interaction-sensitive.** Strengthening reasoning at deeper layers (e.g.,  $l_{\text{reas}} \approx 24$ ; right-side region) also yields strong performance, suggesting the presence of reasoning-related functional modules. However, their effect is coupled with perception ( $l_{\text{perc}}$ ) and task type, inducing  $\sim 5\%$  performance swings across settings. This supports our choice to keep  $l_{\text{perc}}$  and  $l_{\text{reas}}$  in shallow ( $\leq 10$ ) layers for more consistent performance across diverse tasks. Interestingly, the low Acc observed in the intermediate range ( $l \in [10, 17]$ ), when contrasted with both shallow and deep high-Acc regions, reveals the existence of a transitional zone where perception and reasoning functions are more intricately intertwined rather than cleanly separated.

#### 4.4 EFFECT OF ATTENTION RATIO THRESHOLDS (RQ4)

We further investigate the influence of multiplicative gains ( $g_{\text{reas}}, g_{\text{perc}}$ ) and ratio thresholds ( $\tau_{\text{reas}}, \tau_{\text{perc}}$ ) on model performance. As shown in Figure 5, we report results under the same evaluation setup as Table 1, with default hyperparameters specified in Figure 5.

❖ **Obs. 7: Multiplicative gains show task-dependent yet partially transferable effects.** Peaks appear at similar values (e.g.,  $g_{\text{reas}} = 1.30$ ,  $g_{\text{perc}} = 1.16$ ) for MathVista<sub>mini</sub> and MMStar, yielding  $\sim 15\%$  and  $\sim 6\%$  gains, respectively, and suggesting some cross-task generality. From the overall trend,  $g_{\text{reas}}$  shows a steady pattern, delivering  $\sim 10\%$  improvement already at 1.10, whereas  $g_{\text{perc}}$  fluctuates more noticeably across datasets and settings.



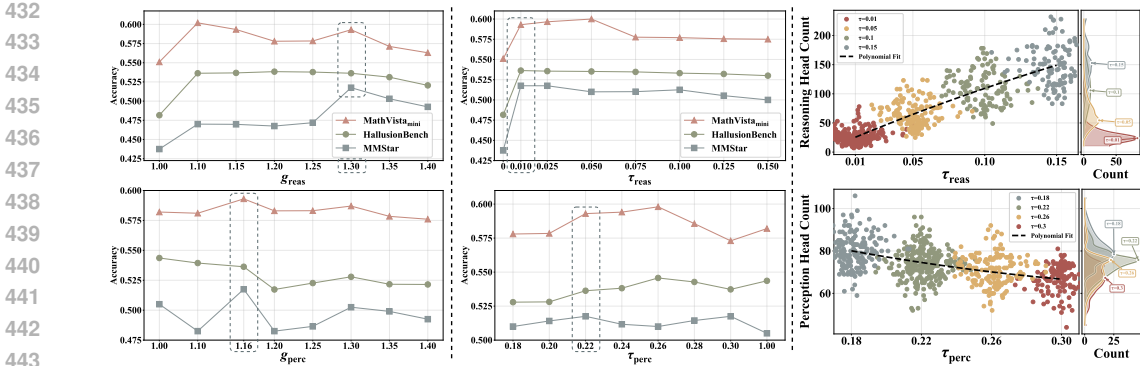


Figure 5: **Analysis of multiplicative gains and ratio thresholds.** On Ocean-R1 with fixed boundaries  $\ell_{reas}=3$ ,  $\ell_{perc}=7$ : *Left*—performance on three datasets under varying multiplicative gains ( $g_{reas}$ ,  $g_{perc}$ ); *Middle*—impact of ratio thresholds ( $\tau_{reas}$ ,  $\tau_{perc}$ ) on performance; *Right*—number of identified heads as a function of ( $\tau_{reas}$ ,  $\tau_{perc}$ ) (small horizontal jitters are added for visual clarity). Default settings are  $g_{reas}=1.30$ ,  $\tau_{reas}=0.01$ ,  $g_{perc}=1.16$ ,  $\tau_{perc}=0.22$ .

❖ **Obs. 8: Ratio thresholds control intervention sparsity and quality.** As shown in the right panel of Figure 5, varying the ratio thresholds directly changes the number of intervened heads (with token-dependent variability). For  $\tau_{reas}$ , strong performance concentrates when  $\sim 50$ – $150$  heads are selected on average (about 6.4% of all heads), indicating a *sparse* intervention. As the target set grows (approximately 6%  $\rightarrow$  18%), performance degrades gradually (e.g., error 0.59  $\rightarrow$  5.7), consistent with dilution from less functional heads.

## 5 RELATED WORK

**Interpretability of Multimodal Reasoning.** Long-chain reasoning is critical for complex multimodal tasks, yet the opacity of current models limits transparency. Existing research follows two main directions. Prompting- and data-driven methods extend Chain-of-Thought reasoning to visual inputs (Wei et al., 2022), either through functionally distinct prompts (Zheng et al., 2023) or curated datasets with annotated reasoning traces and visual highlights (Shao et al., 2024). Architectural and pipeline designs restructure the reasoning process via modular agents (Chen et al., 2023b) or explanation-driven preference optimization (Zhang et al., 2024b). While these approaches expose latent rationales, they often demand heavy supervision or handcrafted prompting. A growing consensus in this line of work is that multimodal reasoning unfolds in stages of perception and symbolic reasoning, which motivates the stage-aware perspective underlying our method (Zheng et al., 2023; Shao et al., 2024).

**Multimodal Hallucination.** Prior studies attribute hallucination in MLRMs to three recurring causes: (i) strong language priors and co-occurrence biases, (ii) insufficient cross-modal alignment, and (iii) training or decoding artifacts (Yue et al., 2024). Analyses across benchmarks such as CHAIR (Rohrbach et al., 2018), POPE (Li et al., 2023b), and MMHal-Bench (Sun et al., 2024) consistently support these views. Mitigation strategies fall into three families. Contrastive decoding perturbs inputs or prompts to down-weight unreliable tokens (e.g., VCD (Leng et al., 2024), ICD (Wang et al., 2024c), LCD (Manevich & Tsarfaty, 2024)). Alignment and preference learning strengthen visual grounding through DPO (Yang et al., 2025b) variants or hallucination-aware training (e.g., HACL (Jiang et al., 2024)). External validation and tool grounding enforce consistency using detectors or CLIP-based guidance (e.g., MARINE (Zhao et al., 2025b)).

## 6 CONCLUSION

We introduce a lightweight, interpretable plugin for MLRMs that identifies perception- and reasoning-oriented heads and selectively enhances them under a minimal-editing principle. Unlike prior decoding-time or training-based approaches, our method is *plug-and-play*, requiring no retraining or architectural changes. Furthermore, the method improves accuracy without introducing any discernible *inference-time overhead*. Across diverse benchmarks, it delivers balanced gains on perception-heavy and reasoning-heavy tasks, confirming that targeted head control can mitigate hallucination without compromising efficiency. We believe this drop-in, cost-free intervention offers a practical path to more reliable and interpretable multimodal reasoning in the deployment stage.

## REFERENCES

- 486  
487  
488 Wenbin An, Feng Tian, Sicong Leng, Jiahao Nie, Haonan Lin, QianYing Wang, Ping Chen, Xiaoqin  
489 Zhang, and Shijian Lu. Mitigating object hallucinations in large vision-language models with  
490 assembly of global and local attention. In *Proceedings of the Computer Vision and Pattern  
491 Recognition Conference*, pp. 29915–29926, 2025.
- 492 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang,  
493 Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*,  
494 2025.
- 495 Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou.  
496 Hallucination of multimodal large language models: A survey. *arXiv e-prints*, pp. arXiv–2404,  
497 2024.
- 499 Jonas Becker, Lars Benedikt Kaesberg, Andreas Stephan, Jan Philip Wahle, Terry Ruas, and Bela  
500 Gipp. Stay focused: Problem drift in multi-agent debate. *arXiv preprint arXiv:2502.19559*, 2025.  
501
- 502 Jing Bi, Junjia Guo, Yunlong Tang, Lianggong Bruce Wen, Zhang Liu, Bingjie Wang, and Chenliang  
503 Xu. Unveiling visual perception in language models: An attention head analysis approach. In  
504 *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 4135–4144, 2025.
- 505 Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing  
506 multimodal llm’s referential dialogue magic. *arXiv e-prints*, pp. arXiv–2306, 2023a.  
507
- 508 Liangyu Chen, Bo Li, Sheng Shen, Jingkan Yang, Chunyuan Li, Kurt Keutzer, Trevor Darrell,  
509 and Ziwei Liu. Large language models are visual reasoning coordinators. *Advances in Neural  
510 Information Processing Systems*, 36:70115–70140, 2023b.
- 511 Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi  
512 Wang, Yu Qiao, Dahua Lin, et al. Are we on the right way for evaluating large vision-language  
513 models? *Advances in Neural Information Processing Systems*, 37:27056–27087, 2024.  
514
- 515 Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James R Glass, and Pengcheng He. Dola:  
516 Decoding by contrasting layers improves factuality in large language models. In *The Twelfth  
517 International Conference on Learning Representations*, 2023.
- 518 Ailin Deng, Zhirui Chen, and Bryan Hooi. Seeing is believing: Mitigating hallucination in large  
519 vision-language models via clip-guided decoding. In *ICLR 2024 Workshop on Reliable and  
520 Responsible Foundation Models*, 2024.  
521
- 522 Yifan Ding, Matthew Facciani, Ellen Joyce, Amrit Poudel, Sanmitra Bhattacharya, Balaji Veeramani,  
523 Sal Aguinaga, and Tim Weninger. Citations and trust in llm generated responses. In *Proceedings  
524 of the AAAI Conference on Artificial Intelligence*, volume 22, pp. 23787–23795, 2025.
- 525 Xuan Gong, Tianshi Ming, Xinpeng Wang, and Zhihua Wei. Damro: Dive into the attention  
526 mechanism of llm to reduce object hallucination. In *Proceedings of the 2024 Conference on  
527 Empirical Methods in Natural Language Processing*, pp. 7696–7712, 2024.  
528
- 529 Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang  
530 Chen, Furong Huang, Yaser Yacoub, et al. Hallusionbench: an advanced diagnostic suite for entan-  
531 gled language hallucination and visual illusion in large vision-language models. In *Proceedings of  
532 the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14375–14385, 2024.
- 533 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,  
534 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms  
535 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.  
536
- 537 Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong  
538 Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language  
539 models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information  
Systems*, 43(2):1–55, 2025.

- 540 Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming  
541 Zhang, and Nenghai Yu. Opera: Alleviating hallucination in multi-modal large language models  
542 via over-trust penalty and retrospection-allocation. In *Proceedings of the IEEE/CVF Conference*  
543 *on Computer Vision and Pattern Recognition*, pp. 13418–13427, 2024.
- 544 Shaohan Huang, Li Dong, Wenhui Wang, Yaru Hao, Saksham Singhal, Shuming Ma, Tengchao Lv,  
545 Lei Cui, Owais Khan Mohammed, Barun Patra, et al. Language is not all you need: Aligning  
546 perception with language models. In *Thirty-seventh Conference on Neural Information Processing*  
547 *Systems*, 2023.
- 549 Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang,  
550 Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM*  
551 *computing surveys*, 55(12):1–38, 2023.
- 552 Chaoya Jiang, Haiyang Xu, Mengfan Dong, Jiaying Chen, Wei Ye, Ming Yan, Qinghao Ye, Ji Zhang,  
553 Fei Huang, and Shikun Zhang. Hallucination augmented contrastive learning for multimodal large  
554 language model. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition*  
555 *(CVPR)*, pp. 27026–27036. IEEE Computer Society, 2024.
- 556 Zhangqi Jiang, Junkai Chen, Beier Zhu, Tingjin Luo, Yankun Shen, and Xu Yang. Devils in middle  
557 layers of large vision-language models: Interpreting, detecting and mitigating object hallucinations  
558 via attention lens. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp.  
559 25004–25014, 2025.
- 561 Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and  
562 Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual  
563 reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.  
564 2901–2910, 2017.
- 565 Seil Kang, Jinyeong Kim, Junhyeok Kim, and Seong Jae Hwang. See what you are told: Visual  
566 attention sink in large multimodal models. In *The Thirteenth International Conference on Learning*  
567 *Representations*, 2025a.
- 568 Seil Kang, Jinyeong Kim, Junhyeok Kim, and Seong Jae Hwang. Your large vision-language model  
569 only needs a few attention heads for visual grounding. In *Proceedings of the Computer Vision and*  
570 *Pattern Recognition Conference*, pp. 9339–9350, 2025b.
- 572 Minchan Kim, Minyeong Kim, Junik Bae, Suhwan Choi, Sungkyung Kim, and Buru Chang. Exploit-  
573 ing semantic reconstruction to mitigate hallucinations in vision-language models. In *European*  
574 *Conference on Computer Vision*, pp. 236–252. Springer, 2024.
- 575 Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing.  
576 Mitigating object hallucinations in large vision-language models through visual contrastive decod-  
577 ing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,  
578 pp. 13872–13882, 2024.
- 579 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal,  
580 Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented genera-  
581 tion for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:  
582 9459–9474, 2020.
- 584 Bohao Li, Yuying Ge, Yixiao Ge, Guangzhi Wang, Rui Wang, Ruimao Zhang, and Ying Shan.  
585 Seed-bench: Benchmarking multimodal large language models. In *Proceedings of the IEEE/CVF*  
586 *Conference on Computer Vision and Pattern Recognition*, pp. 13299–13308, 2024.
- 587 Chong Li, Shaonan Wang, Yunhao Zhang, Jiajun Zhang, and Chengqing Zong. Interpreting and ex-  
588 ploiting functional specialization in multi-head attention under multi-task learning. In *Proceedings*  
589 *of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 16460–16476,  
590 2023a.
- 591 Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating  
592 object hallucination in large vision-language models. In *Proceedings of the 2023 Conference on*  
593 *Empirical Methods in Natural Language Processing*, pp. 292–305, 2023b.

- 594 Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. Mitigating  
595 hallucination in large multi-modal models via robust instruction tuning. In *The Twelfth International*  
596 *Conference on Learning Representations*, 2023.
- 597 Hanchao Liu, Wenyuan Xue, Yifei Chen, Dapeng Chen, Xiutian Zhao, Ke Wang, Liping Hou,  
598 Rongjun Li, and Wei Peng. A survey on hallucination in large vision-language models. *CoRR*,  
599 2024.
- 600 Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng,  
601 Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning  
602 of foundation models in visual contexts. In *The Twelfth International Conference on Learning*  
603 *Representations*, 2023.
- 604 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri  
605 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement  
606 with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594, 2023.
- 607 Avshalom Manevich and Reut Tsarfaty. Mitigating hallucinations in large vision-language models  
608 (lvllms) via language-contrastive decoding (lcd). In *Findings of the Association for Computational*  
609 *Linguistics ACL 2024*, pp. 6008–6022, 2024.
- 610 Paul Michel, Omer Levy, and Graham Neubig. Are sixteen heads really better than one? *Advances in*  
611 *neural information processing systems*, 32, 2019.
- 612 Lingfeng Ming, Yadong Li, Song Chen, Jianhua Xu, Zenan Zhou, and Weipeng Chen. Ocean-r1:  
613 An open and generalizable large vision-language model enhanced by reinforcement learning.  
614 <https://github.com/VLM-RL/Ocean-RL>, 2025. Accessed: 2025-04-03.
- 615 Andrew Nam, Henry Conklin, Yukang Yang, Thomas Griffiths, Jonathan Cohen, and Sarah-Jane  
616 Leslie. Causal head gating: A framework for interpreting roles of attention heads in transformers.  
617 *arXiv preprint arXiv:2505.13737*, 2025.
- 618 Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan,  
619 Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. In-context learning and induction heads.  
620 *arXiv preprint arXiv:2209.11895*, 2022.
- 621 Siyi Pan, Baoliang Chen, Danni Huang, Hanwei Zhu, Zhu Li, Xiangjie Sui, and Shiqi Wang.  
622 Mitigating perception bias: A training-free approach to enhance lmm for image quality assessment.  
623 *CoRR*, 2024.
- 624 Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu  
625 Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint*  
626 *arXiv:2306.14824*, 2023.
- 627 Vipula Rawte, Amit P Sheth, and Amitava Das. A survey of hallucination in large foundation models.  
628 *CoRR*, 2023.
- 629 Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. Object  
630 hallucination in image captioning. In *Proceedings of the 2018 Conference on Empirical Methods*  
631 *in Natural Language Processing*, pp. 4035–4045, 2018.
- 632 Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and  
633 Hongsheng Li. Visual cot: Advancing multi-modal language models with a comprehensive  
634 dataset and benchmark for chain-of-thought reasoning. In *The Thirty-eight Conference on Neural*  
635 *Information Processing Systems Datasets and Benchmarks Track*, 2024.
- 636 Kacper Sokol and Julia E Vogt. (un) reasonable allure of ante-hoc interpretability for high-stakes  
637 domains: Transparency is necessary but insufficient for comprehensibility. In *ICML 3rd Workshop*  
638 *on Interpretable Machine Learning in Healthcare (IMLH)*, 2023.
- 639 Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan,  
640 Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. Aligning large multimodal models with  
641 factually augmented rlhf. In *Annual Meeting of the Association for Computational Linguistics*,  
642 2024.

- 648 Kimi Team, Angang Du, Bohong Yin, Bowei Xing, Bowen Qu, Bowen Wang, Cheng Chen, Chenlin  
649 Zhang, Chenzhuang Du, Chu Wei, et al. Kimi-vl technical report. *arXiv preprint arXiv:2504.07491*,  
650 2025.
- 651 Qwen Team. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.
- 652  
653 Jiayu Wang, Yifei Ming, Zhenmei Shi, Vibhav Vineet, Xin Wang, Yixuan Li, and Neel Joshi. Is a  
654 picture worth a thousand words? delving into spatial reasoning for vision language models. In *The*  
655 *Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024a.
- 656  
657 Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and  
658 Hongsheng Li. Measuring multimodal mathematical reasoning with math-vision dataset. *Advances*  
659 *in Neural Information Processing Systems*, 37:95095–95169, 2024b.
- 660  
661 Xintong Wang, Jingheng Pan, Liang Ding, and Chris Biemann. Mitigating hallucinations in large  
662 vision-language models with instruction contrastive decoding. In *ACL (Findings)*, 2024c.
- 663  
664 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H Chi,  
665 Quoc V Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language  
666 models. In *Proceedings of the 36th International Conference on Neural Information Processing*  
667 *Systems*, pp. 24824–24837, 2022.
- 668  
669 Shengqiong Wu, Hao Fei, Liangming Pan, William Yang Wang, Shuicheng Yan, and Tat-Seng Chua.  
670 Combating multimodal llm hallucination via bottom-up holistic reasoning. In *Proceedings of the*  
671 *AAAI Conference on Artificial Intelligence*, volume 39, pp. 8460–8468, 2025.
- 672  
673 Yuxi Xie, Guanzhen Li, Xiao Xu, and Min-Yen Kan. V-dpo: Mitigating hallucination in large vision  
674 language models via vision-guided direct preference optimization. *CoRR*, 2024.
- 675  
676 Guowei Xu, Peng Jin, Hao Li, Yibing Song, Lichao Sun, and Li Yuan. Llava-cot: Let vision language  
677 models reason step-by-step. *CoRR*, 2024.
- 678  
679 Yi Yang, Xiaoxuan He, Hongkun Pan, Xiyan Jiang, Yan Deng, Xingtao Yang, Haoyu Lu, Dacheng  
680 Yin, Fengyun Rao, Minfeng Zhu, et al. R1-onevision: Advancing generalized multimodal reasoning  
681 through cross-modal formalization. *arXiv preprint arXiv:2503.10615*, 2025a.
- 682  
683 Zhihe Yang, Xufang Luo, Dongqi Han, Yunjian Xu, and Dongsheng Li. Mitigating hallucinations  
684 in large vision-language models via dpo: On-policy data hold the key. In *Proceedings of the*  
685 *Computer Vision and Pattern Recognition Conference*, pp. 10610–10620, 2025b.
- 686  
687 Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. A survey on  
688 multimodal large language models. *National Science Review*, 11(12):nwae403, 2024.
- 689  
690 Zihao Yue, Liang Zhang, and Qin Jin. Less is more: Mitigating multimodal hallucination from  
691 an eos decision perspective. In *Proceedings of the 62nd Annual Meeting of the Association for*  
692 *Computational Linguistics (Volume 1: Long Papers)*, pp. 11766–11781, 2024.
- 693  
694 Xiaofeng Zhang, Yihao Quan, Chaochen Gu, Chen Shen, Xiaosong Yuan, Shaotian Yan, Hao Cheng,  
695 Kaijie Wu, and Jieping Ye. Seeing clearly by layer two: Enhancing attention heads to alleviate  
696 hallucination in vlms. *CoRR*, 2024a.
- 697  
698 Xuan Zhang, Chao Dut, Tianyu Pang, Qian Liu, Wei Gao, and Min Lin. Chain of preference opti-  
699 mization: improving chain-of-thought reasoning in llms. In *Proceedings of the 38th International*  
700 *Conference on Neural Information Processing Systems*, pp. 333–356, 2024b.
- 701  
702 Zeliang Zhang, Phu Pham, Wentian Zhao, Kun Wan, Yu-Jhe Li, Jianing Zhou, Daniel Miranda,  
Ajinkya Kale, and Chenliang Xu. Treat visual tokens as text? but your mllm only needs fewer  
efforts to see. *arXiv preprint arXiv:2410.06169*, 2024c.
- 703  
704 Junru Zhao, Tianqin Li, Dunhan Jiang, Shenghao Wu, Alan Ramirez, and Tai Sing Lee. Perceptual  
inductive bias is what you need before contrastive learning. In *Proceedings of the Computer Vision*  
*and Pattern Recognition Conference*, pp. 9621–9630, 2025a.

702 Linxi Zhao, Yihe Deng, Weitong Zhang, and Quanquan Gu. Mitigating object hallucination in large  
703 vision-language models via image-grounded guidance. In *Forty-second International Conference*  
704 *on Machine Learning*, 2025b.

705  
706 Ge Zheng, Bin Yang, Jiajin Tang, Hong-Yu Zhou, and Sibeï Yang. Ddcot: Duty-distinct chain-of-  
707 thought prompting for multimodal reasoning in language models. In *Thirty-seventh Conference on*  
708 *Neural Information Processing Systems*, 2023.

709  
710  
711  
712  
713  
714  
715  
716  
717  
718  
719  
720  
721  
722  
723  
724  
725  
726  
727  
728  
729  
730  
731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755

756 THE USE OF LARGE LANGUAGE MODELS (LLMs)

757  
758 In preparing this manuscript, we used a large language model (LLM) exclusively for stylistic  
759 refinement, such as polishing grammar, improving clarity, and enhancing readability. The LLM  
760 was not involved in formulating research questions, designing methods, analyzing data, or drawing  
761 conclusions. All scientific content, ideas, and experimental results are entirely the work of the authors.  
762 The LLM served purely as an auxiliary tool for language editing.

763  
764 LIMITATION, DISCUSSION, AND FUTURE WORK

765  
766 **Beyond a single boundary.** Our experiments in Sec. 4.3 and Sec. 4.4 show that the division between  
767 perception and reasoning is not determined by a crisp threshold. Instead, performance peaks appear  
768 in *bands* of layers, and the optimal  $\ell_{\text{perc}}$  and  $\ell_{\text{reas}}$  vary across tasks. This suggests that different  
769 layers may play mixed roles, and that the model undergoes a gradual transition from perception  
770 to reasoning rather than a clean separation. Future work could build on interpretability studies to  
771 map these transitions more precisely, moving from a one-dimensional boundary to richer structural  
772 patterns.

773  
774 **Toward more adaptive selection.** Our method follows the *minimal editing* principle, amplifying  
775 a small set of heads without attenuating others. Formally, the gain in task alignment is  $\Delta =$   
776  $\sum_{h \in \mathcal{H}} (\gamma_h - 1) u_h$ , where  $u_h$  measures the usefulness of head  $h$ . In this work, we identify perception  
777 and reasoning heads using simple ratio thresholds. A natural extension is to design more fine-grained  
778 classifiers that score each head by multiple signals (e.g., depth, modality ratio, consistency). An  
779 adaptive scheme could then select  $S(x)$  depending on the input  $x$ , and set  $\gamma_h > 1$  only for those  
780 heads. While this may yield larger  $\Delta$ , it also risks higher variance and efficiency costs. Thus, the  
781 challenge is to balance stronger, adaptive selection with the stability and efficiency of the current  
782 plug-and-play design.

783 A NOTATIONS AND DEFINITIONS

784  
785  
786 Table 3: Notations and Definitions

788 Notation	788 Definition
789 $\mathcal{T}_v, \mathcal{T}_t, \mathcal{T}_q$	789 Sets of vision, text tokens and query tokens, respectively.
790 $l, h$	790 Indices for layer and attention head in a Transformer, 791 respectively.
792 $X^{(l)}$	792 Input hidden states at layer $l$ .
793 $W_Q^{(h,l)}, W_K^{(h,l)}, W_V^{(h,l)}$	793 Query, key, and value projection matrices for head $h$ in 794 layer $l$ .
795 $A^{(h,l)}$	795 Attention matrix for head $h$ in layer $l$ , $A^{(h,l)} \in \mathbb{R}^{N \times N}$ .
796 $a_{ij}^{(h,l)}$	796 Attention weight from query token $i$ to key token $j$ in 797 matrix $A^{(h,l)}$ .
798 $O^{(h,l)}$	798 Output vector for head $h$ in layer $l$ , $O_h^{(l)} \in \mathbb{R}^{d_{\text{model}}}$ .
799 $V^{(h,l)}$	799 Value tensor for head $h$ in layer $l$ .
800 $Y^{(l)}$	800 Standard multi-head output of the Transformer block at 801 layer $l$ .
802 $\mathcal{L}_{\text{perc}}, \mathcal{L}_{\text{reas}}$	802 Sets of perception layers and reasoning layers, 803 respectively.
804 $l_{\text{perc}}, l_{\text{reas}}$	804 Boundaries for the last layer of the perception stage and 805 the first layer of the reasoning stage.
806 $S_v^{(l)}(h), S_t^{(l)}(h)$	806 Average attention ratio allocated to visual/textual tokens 807 by head $h$ in layer $l$ .
808 $\tau_{\text{perc}}, \tau_{\text{reas}}$	808 Thresholds of visual attention ratio to distinguish 809 perception and reasoning heads.

Continued on next page

Table 3 – Continued from previous page

Notation	Definition
$\mathcal{H}$	The full set of attention heads in a given multi-head attention (MHA) sub-layer.
$\mathcal{H}_{perc}^{(l)}, \mathcal{H}_{reas}^{(l)}$	Sets of perception-oriented and reasoning-oriented heads identified in layer $l$ .
$g_{perc}, g_{reas}$	Global gain factors applied to perception and reasoning heads, respectively.
$\alpha_{perc}, \alpha_{reas}$	Scalar gain assigned to a head based on its attention ratio (used interchangeably with $g_{perc}, g_{reas}$ ).
$g^{(h,l)}$	Specific gain value applied to head $h$ in layer $l$ .
$O_S$	Aggregated output of a subset of heads $S \subseteq \mathcal{H}$ .

## B SELECTIVE ENHANCEMENT AS OUR POLICY CHOICE

**Problem and Objective (Minimal Editing Principle).** When intervening on attention heads, our objective is to *maximize the correction of hallucination* while *minimizing unnecessary edits* to the model’s internal representations. We formalize this idea as the **Minimal Editing Principle**: an effective intervention should (1) amplify signals that are already verified as beneficial, and (2) avoid suppressing or perturbing heads whose functions remain uncertain.

Let  $\mathcal{H}$  denote the full set of attention heads in a given multi-head attention (MHA) sub-layer. Each head  $h \in \mathcal{H}$  produces an output vector  $\mathbf{O}_h^{(l)} \in \mathbb{R}^{d_{\text{model}}}$ , and the aggregated MHA output can be expressed as

$$\tilde{Y}^{(l)} = \sum_{h \in \mathcal{H}} \gamma_h^{(l)} \mathbf{O}_h^{(l)}, \tag{8}$$

where  $\gamma_h^{(l)}$  is a multiplicative gain applied to head  $h$ . In this formulation, any intervention corresponds to editing the gain pattern  $\{\gamma_h^{(l)}\}_{h \in \mathcal{H}}$ . *Minimal editing* requires that such modifications keep the original contributions of most heads, while selectively reinforcing those aligned with perception or reasoning.

Intuitively, this principle reflects a conservative stance. We cannot guarantee that non-target heads are harmful. If they are attenuated indiscriminately, the model may lose useful functions and suffer collateral damage. The safer strategy is therefore to intervene as little as possible: amplify only the identified functional heads, while keeping the rest unchanged to preserve the existing representation.

**Formal Setup and Minimal Assumptions.** To make analysis tractable and verifiable, we adopt two minimal assumptions: (1) interventions operate *after attention computation but before residual and normalization*, ensuring that changes are linear and directly comparable; (2) all enhanced heads share a common amplification factor  $\alpha > 1$ , while all attenuated heads (if any) share a common factor  $\beta < 1$ .

These assumptions preserve the generality of our formulation while guaranteeing that differences between strategies can be derived in closed form. In the next subsection, we instantiate four policies under this setup and analyze their functional differences.

**Four Strategies and a Core Proposition.** Under the above setup, we distinguish four representative policies for modifying head gains:

- **Strategy A (Selective Enhancement).**

$$\gamma_h^{(l)} = \begin{cases} \alpha, & h \in \mathcal{H}_{\text{enhance}}, \\ 1, & h \in \mathcal{H} \setminus \mathcal{H}_{\text{enhance}}. \end{cases} \tag{9}$$

- **Strategy B (Selective Attenuation).**

$$\gamma_h^{(l)} = \begin{cases} \beta, & h \in \mathcal{H}_{\text{attenuate}}, \\ 1, & h \in \mathcal{H} \setminus \mathcal{H}_{\text{attenuate}}. \end{cases} \tag{10}$$



- **Strategy C (Bipolar Scaling).**

$$\gamma_h^{(l)} = \begin{cases} \alpha, & h \in \mathcal{H}_{\text{enhance}}, \\ \beta, & h \in \mathcal{H}_{\text{attenuate}}. \end{cases} \quad (11)$$

- **Strategy D (Mixed Policy).**

$$\gamma_h^{(l)} = \begin{cases} \alpha, & h \in \mathcal{H}_{\text{enhance}}, \\ \beta, & h \in \mathcal{H}_{\text{attenuate}}, \\ 1, & h \in \mathcal{H}_{\text{neutral}}. \end{cases} \quad (12)$$

**Proposition (Difference Equation).** Let  $\mathbf{O}_S = \sum_{h \in S} \mathbf{O}_h^{(l)}$  denote the aggregated output of a subset of heads  $S \subseteq \mathcal{H}$ . Then, relative to Strategy A, any scheme that introduces attenuation (C or D) differs by an explicit subtractive term:

$$\tilde{Y}_C^{(l)} - \tilde{Y}_A^{(l)} = (\beta - 1) \mathbf{O}_{\mathcal{H}_{\text{att}}}, \quad (13)$$

$$\tilde{Y}_D^{(l)} - \tilde{Y}_A^{(l)} = (\beta - 1) \mathbf{O}_{\mathcal{H}_{\text{att}}}, \quad (14)$$

where  $\mathcal{H}_{\text{att}}$  denotes the set of attenuated heads.

*Proof.* Both Strategy A and Strategy C/D share the amplified component  $\alpha \mathbf{O}_{\mathcal{H}_{\text{enh}}}$ . The difference arises in the treatment of non-enhanced heads: Strategy A preserves their outputs unchanged, while C and D rescale them by  $\beta < 1$ . This introduces a correction term  $(\beta - 1) \mathbf{O}_{\mathcal{H}_{\text{att}}}$ , representing the explicit removal of information contributed by attenuated heads.

**Interpretation.** This difference equation makes the design trade-off explicit: ♣ Strategy A amplifies useful heads while keeping all others intact. ♣ Strategies C and D apply an additional subtraction to all attenuated heads, implicitly assuming they are harmful. Since non-target heads may still carry latent but beneficial functions, attenuation risks collateral degradation. Thus, the *minimal editing* choice is Strategy A, which strengthens the identified heads without disturbing the rest of the model.

**Expectation-Level Intuition: Why Selective Enhancement is Safer.** Consider a task-aligned direction  $v \in \mathbb{R}^{d_{\text{model}}}$ , and let each head’s contribution be

$$u_h := \langle \mathbf{O}_h^{(l)}, v \rangle, \quad (15)$$

which measures how much the head  $h$  aligns with the task. For any headset  $\mathcal{S}$ , we then define its total contribution as

$$U(\mathcal{S}) = \sum_{h \in \mathcal{S}} u_h. \quad (16)$$

Ignoring LayerNorm rescaling, the alignment change caused by a gain pattern  $\{\gamma_h^{(l)}\}$  is

$$\Delta(\{\gamma\}) \approx \sum_{h \in \mathcal{H}} (\gamma_h^{(l)} - 1) u_h, \quad (17)$$

which simply adds the weighted shifts of each head relative to the neutral baseline ( $\gamma_h = 1$ ).

**Three Strategies.** Applying this formulation, the alignment increments become:

$$\Delta_A \approx (\alpha - 1) U(\mathcal{H}_{\text{enhance}}), \quad (18)$$

corresponding to *selective enhancement*, where only useful heads are amplified.

$$\Delta_B \approx (\beta - 1) U(\mathcal{H}_{\text{attenuate}}), \quad (19)$$

for *selective attenuation*, where some heads are suppressed while all others remain neutral.

$$\Delta_C \approx (\alpha - 1) U(\mathcal{H}_{\text{enhance}}) + (\beta - 1) U(\mathcal{H}_{\text{attenuate}}), \quad (20)$$

for *bipolar scaling*, where enhancement and attenuation are applied simultaneously.

**Assumption (non-harmful heads).** Prior interpretability studies indicate that most attention heads carry diverse or weakly positive roles rather than being systematically harmful (Olsson et al., 2022). Formally, this suggests

$$\mathbb{E}[U(\mathcal{H}_{\text{attenuate}})] \geq 0, \quad (21)$$

meaning that, on average, attenuated heads are not expected to reduce alignment.

**Implications.** Taking expectations under this assumption yields

$$\mathbb{E}[\Delta_A] > 0, \quad \mathbb{E}[\Delta_B] \leq 0, \quad \mathbb{E}[\Delta_C] \leq \mathbb{E}[\Delta_A]. \quad (22)$$

Hence, Strategy A provides a reliably positive gain by amplifying identified functional heads. Strategies B and C both introduce subtractive terms  $(\beta - 1)U(\mathcal{H}_{\text{attenuate}})$ , which are non-positive in expectation, making them less stable. This formalizes why, under the *minimal-editing* principle, selective enhancement is the safer choice.

**Takeaways.** Strategy A yields a guaranteed positive gain by enhancing only functional heads. Strategies B and C involve attenuation terms  $(\beta - 1)U(\mathcal{H}_{\text{attenuate}})$ , which are non-positive in expectation and risk weakening neutral or beneficial heads. Therefore, under the *minimal editing* principle, Strategy A is the most stable and reliable choice.

## C TIME COMPLEXITY ANALYSIS

In this section, we present a detailed analysis of the algorithmic time complexity of the standard attention mechanism and our proposed method in a single multi-head attention block.

Let the input sequence length be  $N$ , the model hidden dimension be  $d_{\text{model}}$ , and the number of attention heads be  $H$ . For each attention head, the query, key, and value dimensions are denoted by  $d_k$  and  $d_v$ , with  $d_k = d_v = d_{\text{model}}/H$ . For notational simplicity, we refer to the per-head dimension as  $d_h$ . In the following analysis, we focus on the asymptotic complexity with respect to  $N$ , treating  $d_{\text{model}}$ ,  $H$ ,  $d_k$ , and  $d_v$  as constants.

### STANDARD ATTENTION

**Input Projections.** Given the input tensor  $X \in \mathbb{R}^{N \times d_{\text{model}}}$ , three linear projections are applied to generate the query, key, and value tensors:  $Q = XW_Q$ ,  $K = XW_K$ , and  $V = XW_V$ , where  $W_Q, W_K, W_V \in \mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$ . Each projection is a matrix multiplication with complexity  $O(N \cdot d_{\text{model}}^2)$ , resulting in:

$$T_1 = O(3 \cdot N \cdot d_{\text{model}}^2) = O(N \cdot d_{\text{model}}^2) \quad (23)$$

**Attention Score Computation.** For each head  $h \in \{1, \dots, H\}$ , unnormalized attention scores are computed as  $Q_h K_h^T$ , where  $Q_h \in \mathbb{R}^{N \times d_k}$  and  $K_h^T \in \mathbb{R}^{d_k \times N}$ . The cost per head is  $O(N^2 \cdot d_k)$ , and across all heads:

$$T_2 = O(H \cdot N^2 \cdot d_k) \quad (24)$$

**Softmax Normalization.** A row-wise softmax is applied to the  $N \times N$  score matrix for each head. Applying softmax to a vector of length  $N$  costs  $O(N)$ , and thus to  $N$  rows is  $O(N^2)$  per head. Across all heads:

$$T_3 = O(H \cdot N^2) \quad (25)$$

**Value Aggregation.** The normalized attention matrix  $A_h$  is multiplied by  $V_h \in \mathbb{R}^{N \times d_v}$  to produce the head output  $O_h \in \mathbb{R}^{N \times d_v}$ . The cost per head is  $O(N^2 \cdot d_v)$ , yielding:

$$T_4 = O(H \cdot N^2 \cdot d_v) \quad (26)$$

**Output Projection.** The  $H$  head outputs are concatenated into a  $N \times d_{\text{model}}$  matrix and projected via  $W_O \in \mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$  with cost:

$$T_5 = O(N \cdot d_{\text{model}}^2) \quad (27)$$

Summing all terms gives the total time complexity:

$$T_{\text{Standard}} = T_1 + T_2 + T_3 + T_4 + T_5 = O(Nd_{\text{model}}^2 + HN^2d_k + HN^2 + HN^2d_v + Nd_{\text{model}}^2)$$

Table 4: Average per-batch inference time (seconds). Parentheses show relative change vs. Vanilla.

Method	MathVista <sub>mini</sub>	MathVision <sub>mini</sub>	HallusionBench	CLEVR	MMStar	SEED-Bench
Vanilla	98 (+0.0%)	98 (+0.0%)	101 (+0.0%)	74 (+0.0%)	83 (+0.0%)	68 (+0.0%)
VCD	159 (+62.2%)	241 (+145.9%)	170 (+68.3%)	116 (+56.8%)	113 (+36.1%)	98 (+44.1%)
CGD	482 (+391.8%)	709 (+623.5%)	664 (+557.4%)	314 (+324.3%)	297 (+257.8%)	329 (+383.8%)
AGLA	119 (+21.4%)	164 (+67.3%)	123 (+21.7%)	78 (+5.4%)	87 (+4.8%)	88 (+29.4%)
Ours	<b>103 (+5.1%)</b>	<b>101 (+3.1%)</b>	<b>103 (+2.0%)</b>	<b>75 (+1.4%)</b>	<b>83 (+0.0%)</b>	<b>69 (+1.5%)</b>

As  $N$  grows, the  $N^2$  terms dominate. Therefore, the overall asymptotic time complexity is:

$$T_{\text{Standard}} = O(H \cdot N^2 \cdot (d_k + d_v + 1)) = O(N^2) \quad (28)$$

#### OUR METHOD

**Standard Attention Score and Softmax Computation.** Our method first performs the standard attention computation as in the baseline, obtaining all  $H$  normalized attention matrices  $A_h \in \mathbb{R}^{N \times N}$ :

$$T'_1 = O(Nd_{model}^2 + HN^2d_k + HN^2) \quad (29)$$

**Visual Attention Ratio Calculation.** We then compute a visual attention ratio for each head based on Equation 4. This operation traverses all  $N \times N$  attention entries, giving:

$$T'_2 = O(H \cdot N^2) \quad (30)$$

**Head Categorization and Gain Assignment.** Each head is categorized and assigned a scalar gain  $\alpha_h$  based on its  $S_v(h)$ . This involves only  $O(H)$  comparisons and assignments, which are negligible compared to the dominant terms.

**Modulated Value Aggregation.** The values are aggregated as  $O_h = A_h V_h$ , followed by element-wise scaling with  $\alpha_h$ . The aggregation part has cost:

$$T'_{4a} = O(H \cdot N^2 \cdot d_v) \quad (31)$$

and the scaling part adds  $O(H \cdot N \cdot d_v) = O(N \cdot d_{model})$ , which is dominated by the aggregation cost. Thus:

$$T'_4 = O(H \cdot N^2 \cdot d_v) \quad (32)$$

**Final Output Projection.** The modulated outputs are concatenated and projected as in the baseline:

$$T'_5 = O(N \cdot d_{model}^2) \quad (33)$$

Combining all steps, the total complexity of our method is:

$$T_{\text{Ours}} = T'_1 + T'_2 + T'_4 + T'_5 = O(Nd_{model}^2 + HN^2d_k + 2HN^2 + HN^2d_v + Nd_{model}^2)$$

Focusing on the dominant terms with respect to  $N$ , we obtain:

$$T_{\text{Ours}} = O(H \cdot N^2 \cdot (d_k + d_v + 2)) = O(N^2) \quad (34)$$

#### EXPERIMENTS

As reported in Table 4, our method attains per-batch inference times that are essentially indistinguishable from the Vanilla baseline across all six benchmarks, with only marginal fluctuations of a few percentage points. This empirical parity aligns with the theoretical expectation that our additional head-level operations introduce only constant-factor overhead beyond standard multi-head attention. Consequently, the overall time complexity remains unchanged, and in practice, the runtime of our approach matches that of the original model.

## 1026 C.1 COMPARISON AND CONCLUSION

1027  
1028 Although our method introduces an additional  $O(H \cdot N^2)$  step for head-level ratio computation, the  
1029 total complexity remains dominated by the  $O(N^2)$  cost of the standard attention mechanism. As  
1030 a result, both exhibit the same asymptotic time complexity of  $O(N^2)$  with respect to the sequence  
1031 length  $N$ . The additional operations only contribute constant-factor overhead without altering the  
1032 overall computational order.

1033 Consequently, in practice, the runtime of our approach is **essentially indistinguishable** from that of  
1034 standard attention.

## 1036 D IMPLEMENTATION DETAILS

1037  
1038 **Ocean-R1 (Ming et al., 2025)** Ocean-R1 is a large vision-language reasoning model fine-tuned  
1039 from Qwen2.5-VL-Instruct (Bai et al., 2025), designed to enhance cross-modal reasoning and  
1040 visual understanding capabilities through a two-stage rule-based Reinforcement Learning framework.  
1041 The first stage focuses on strengthening the model’s reasoning ability, while the second stage  
1042 improves its visual perception. Experimental results demonstrate that Ocean-R1 achieves substantial  
1043 performance gains, particularly on visual mathematical reasoning benchmarks such as MathVision  
1044 (+2.7/+2.7) and MathVista (+4.9/+4.4), showing strong multimodal reasoning and generalization  
1045 capabilities.

1046 **R1-Onevision (Yang et al., 2025a)** R1-Onevision is a state-of-the-art multimodal reasoning  
1047 model designed to bridge the gap between visual perception and deep reasoning. It is fine-tuned  
1048 from Qwen2.5-VL (Bai et al., 2025), with a focus on cross-modal reasoning that enables precise  
1049 understanding and processing of both visual and textual information. Unlike previous models that  
1050 primarily rely on fixed structures for reasoning, R1-Onevision employs a two-stage post-training  
1051 strategy: Supervised Fine-Tuning (SFT) and Reinforcement Learning (RL), enhancing its ability  
1052 to generalize across diverse tasks. The model leverages a cross-modal reasoning pipeline that  
1053 transforms images into formal text-based representations, which are then processed to generate  
1054 structured reasoning paths. It also incorporates a "role-playing" strategy to iteratively refine visual  
1055 comprehension, ensuring robust multimodal coherence. Experimental evaluations on benchmarks like  
1056 MathVista and MathVerse demonstrate that R1-Onevision outperforms several state-of-the-art  
1057 models, including GPT-4o and Qwen2.5-VL, showcasing its superior reasoning and generalization  
1058 capabilities.

1059 **Kimi-VL (Team et al., 2025)** Kimi-VL is an efficient open-source vision-language model built  
1060 upon a Mixture-of-Experts (MoE) language decoder with only 2.8B activated (16B total) parameters,  
1061 paired with a 400M native-resolution vision encoder (MoonViT). It is designed to provide advanced  
1062 multimodal reasoning, long-context understanding, and strong agent capabilities, while maintaining  
1063 high parameter efficiency. Unlike most dense-architecture VLMs, Kimi-VL achieves competitive  
1064 or superior performance to larger models on diverse tasks, including college-level problem solving,  
1065 OCR, multi-image reasoning, video understanding, and long-document comprehension. Through long  
1066 Chain-of-Thought supervised fine-tuning and reinforcement learning, its enhanced variant Kimi-VL-  
1067 Thinking demonstrates strong long-horizon multimodal reasoning ability, achieving remarkable results  
1068 on benchmarks such as MathVision and MathVista. This demonstrates Kimi-VL’s effectiveness in  
1069 combining parameter efficiency with powerful multimodal reasoning capabilities.

1070 Our method is primarily implemented in the `eager_attention_forward` function within  
1071 the modeling file (`modeling_qwen2_5_vl.py` and `modeling_kimi_vl.py`). Also, we  
1072 incorporate a caching mechanism to store essential information (e.g., the range of visual tokens and  
1073 hyperparameters).

1074 Table 5 summarizes the hyperparameters used in our experiments.

## 1076 E DATASETS

1077  
1078 **MathVista (Lu et al., 2023)** Aggregates 6,141 problems by consolidating 28 existing multimodal  
1079 datasets with three newly curated sources—IQTest (puzzle figures), FunctionQA (function plots),  
and PaperQA (figures from academic papers)—spanning charts, diagrams, textbook geometry, and

Table 5: Selected Hyperparameters

Model	$\ell_{\text{reas}}$	$\tau_{\text{reas}}$	$g_{\text{reas}}$	$\ell_{\text{perc}}$	$\tau_{\text{perc}}$	$g_{\text{perc}}$
Kimi-VL-A3B-Thinking	5	0.01	1.40	10	0.27	1.20
Ocean-R1-7B-Instruct	3	0.01	1.30	7	0.22	1.16
R1-Onevision	3	0.01	1.30	7	0.30	1.20

everyday VQA scenes. It targets cross-domain mathematical reasoning grounded in images, with compositional perception and minimal leakage, and remains challenging for current models.

**MathVision** (Wang et al., 2024b) Curated from real math competitions to provide 3,040 image-based problems across 16 mathematical disciplines (e.g., analytic geometry, topology, graph theory) and five difficulty levels. Designed to stress rigorous multimodal mathematical reasoning beyond earlier benchmarks, it exhibits a large human-model gap and offers *test* and *test-mini* splits for rapid benchmarking.

**MMStar** (Chen et al., 2024) Built by screening items from existing benchmarks and then manually vetting them to ensure vision indispensability and remove data leakage, yielding 1,500 human-selected samples covering six core capabilities along 18 axes. Its goal is a purified, balanced test that measures true multimodal gains—i.e., questions where visual evidence is necessary and difficulty stems from advanced cross-modal reasoning.

**CLEVR** (Johnson et al., 2017) A synthetic diagnostic VQA dataset rendered in Blender with simple 3D objects varying in color, shape, size, and material; questions are programmatically generated and paired with functional programs and scene graphs. The design emphasizes low dataset bias and compositional reasoning difficulty (counting, comparisons, spatial relations) over 100k images and 1M questions.

**HallusionBench** (Guan et al., 2024) Comprises 346 figures from diverse sources and formats (e.g., charts, tables, maps, famous optical illusions, memes) paired with 1,129 expert-authored questions structured into control pairs and “visual-dependent” vs. “visual-supplement” settings. It is purpose-built to disentangle language hallucination from visual illusion and to test consistency under easy/hard conditions, remaining challenging for state-of-the-art LVLMs.

**SEED-Bench** (Li et al., 2024) Uses human-verified multiple-choice questions targeted to capability dimensions; in SEED-Bench-1 (19k), images come from Conceptual Captions and videos from Something-Something V2, Epic-Kitchens, and Breakfast, covering both spatial and temporal understanding. The benchmark’s design goal is dimension-specific difficulty and objective automatic scoring across 12 image/video comprehension dimensions.

## F BASELINES

We adopt a deliberately diverse set of baselines. Methods for mitigating inference-time hallucinations via generation intervention can be grouped into three families: **Contrastive Decoding**, **Guided Decoding**, and **Visual Amplification**. Accordingly, we select VCD, CGD, and AGLA as representative instances. These baselines instantiate three complementary mechanisms—probability calibration, external guidance, and feature enhancement—covering different failure modes while reducing evaluation bias due to methodological homogeneity; together they constitute a representative comparison suite.

**Visual Contrastive Decoding (VCD)** (Leng et al., 2024). VCD is a training-free decoding strategy that explicitly builds a counterfactual view of the image by introducing controlled visual uncertainty and then contrasts the model’s predictions under the original versus the uncertain view. Tokens that are preferred only when the image is ambiguous are down-weighted, which directly combats two root causes of object hallucination: over-reliance on language priors and spurious object co-occurrence. An adaptive plausibility constraint further preserves fluent generation by truncating to high-confidence candidates from the original distribution. The key difference from other baselines is that VCD does not require external models or retraining; it modifies only the sampling distribution and therefore generalizes across LVLm families.

**Assembly of Global and Local Attention (AGLA)** (An et al., 2025). AGLA targets the attention deficiency behind many hallucinations: LVLMs tend to lock onto prompt-irrelevant global patterns while missing fine-grained, prompt-relevant regions. It first performs image–prompt matching (via a Grad-CAM–style analysis) to produce an augmented view that highlights regions tied to the query and suppresses distractors; decoding then fuses evidence from the global view (original image) and the local, discriminative view (augmented image). This design differs from VCD and CGD by enriching the *visual features* themselves rather than only reshaping probabilities, which improves grounding without sacrificing generative context. Empirically, it is especially strong under challenging negative settings, e.g., achieving about 81.36 F1 on the adversarial split of POPE with LLaVA-1.5.

**CLIP-Guided Decoding (CGD)** (Deng et al., 2024). CGD injects an external, vision-language prior—CLIP—directly into decoding, but at the *sentence level*. Instead of judging partial tokens, CGD scores complete sentence candidates by mixing the LVLM’s length-normalized likelihood with CLIP image–text similarity and keeps candidates that are both probable and visually grounded. Two design choices explain its gains: (i) sentence-level guidance avoids the myopia of token-level heuristics and directly penalizes late-caption drift (later sentences are empirically more hallucinatory), and (ii) CLIP similarity provides a stronger, more stable signal of image–text alignment than model likelihood alone. As a result, CGD substantially lowers hallucination while preserving caption quality, e.g., CHAIRs on COCO with LLaVA-1.5 drops from 44.7 to 29.7.

## G METRICS

We adopt the *F1 score* as the primary evaluation metric, with calculation logic adapted to the nature of each dataset. Specifically, two distinct evaluation schemes are applied: (1) a binary judgment scheme for datasets containing binary ground-truth labels (e.g., HallucinationBench), and (2) a multi-class scheme for datasets containing multiple-choice questions (e.g., MMStar and SEED-Bench). In both schemes, the final performance is reported using the *Weighted F1 score*.

**Binary Judgment Scheme** For binary datasets, we use the standard F1 score, defined as:

$$F1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} = \frac{2TP}{2TP + FP + FN}. \quad (35)$$

**Multi-class Scheme** For multi-class datasets, each sample has a ground-truth label  $y \in \{1, \dots, K\}$  and a predicted label  $\hat{y} \in \{1, \dots, K\}$ . For each class  $i \in \{1, \dots, K\}$ , we define:

$$TP_i = |\{\hat{y} = i, y = i\}|, \quad FP_i = |\{\hat{y} = i, y \neq i\}|, \quad (36)$$

$$FN_i = |\{\hat{y} \neq i, y = i\}|, \quad TN_i = |\{\hat{y} \neq i, y \neq i\}|. \quad (37)$$

Then:

$$\text{Precision}(i) = \frac{TP_i}{TP_i + FP_i}, \quad \text{Recall}(i) = \frac{TP_i}{TP_i + FN_i}, \quad (38)$$

$$F1(i) = \frac{2 \times \text{Precision}(i) \times \text{Recall}(i)}{\text{Precision}(i) + \text{Recall}(i)}. \quad (39)$$

**Weighted Aggregation** To account for class imbalance, the weighted F1 score is computed as:

$$\text{Weighted-F1} = \sum_{i=1}^K \frac{\text{Support}(i)}{N} \times F1(i), \quad (40)$$

where  $\text{Support}(i)$  denotes the number of samples with ground-truth label  $i$ , and  $N$  is the total number of samples.

## H CASE ANALYSIS

### H.1 CONTRIBUTION MAP

Following Michel et al (Michel et al., 2019), we quantify the contribution of each attention head to a *specific next-token prediction* by inserting differentiable gates on head outputs and reading out

the loss gradients w.r.t. these gates. Intuitively, if slightly amplifying head  $(l, h)$  would decrease the loss for the current target token, that head is *helpful* for this token; if it would increase the loss, it is *harmful*. We report absolute sensitivities as non-negative importance scores and optionally keep the signed scores to analyze supportive vs. adversarial heads.

### Notation at a glance.

- $B$ : batch size;  $S$ : sequence length seen by the model at the current decoding step (includes prompt and previously generated tokens).
- $L$ : number of transformer layers.
- $H$ : number of attention heads per layer.
- $d_{\text{model}}$ : hidden size;  $d_{\text{head}} = d_{\text{model}}/H$ : per-head width after splitting.
- $\mathbf{x}$ : model inputs (encoded sequence including both image and text);  $\mathbf{y}_{<\tau}$ : generated prefix;  $y_\tau$ : random variable at current position;  $t^*$ : target token for this step.
- $\mathcal{L} = -\log P(y_\tau = t^* | \mathbf{x}, \mathbf{y}_{<\tau})$ : cross-entropy loss defined only for current position  $\tau$ .

#### H.1.1 GATE PARAMETER MECHANISM

We insert lightweight gate parameters onto each attention head to directly modulate their outputs without changing model weights. This setup allows us to treat each gate as a differentiable scalar that controls the influence of its head. With this mechanism in place, we can measure how scaling a head affects the model’s loss for the current token.

For layer  $l$ , let the standard multi-head output (after attention and output projection, before the MLP) be

$$\mathbf{O}^{(l)} \in \mathbb{R}^{B \times S \times d_{\text{model}}}. \quad (41)$$

We reshape to expose heads:

$$\mathbf{O}_{\text{heads}}^{(l)} = \text{reshape} \left( \mathbf{O}^{(l)}, [B, S, H, d_{\text{head}}] \right). \quad (42)$$

We use a trainable-on-the-fly (but *not optimized*) gate vector  $\mathbf{g}^{(l)} \in \mathbb{R}^H$  initialized as  $\mathbf{1}_H$  and broadcast it multiplicatively:

$$\forall b \in [B], s \in [S], h \in [H] : \quad \mathbf{O}_{\text{gated}}^{(l)}[b, s, h, :] = g_h^{(l)} \mathbf{O}_{\text{heads}}^{(l)}[b, s, h, :]. \quad (43)$$

Finally we fold heads back to  $d_{\text{model}}$  and continue the standard forward. Implementation uses lightweight forward hooks; model weights remain frozen. Gates are created with `requires_grad=True` so that backprop can produce  $\frac{\partial \mathcal{L}}{\partial g_h^{(l)}}$ .

**Shape sanity check.**  $\mathbf{O}_{\text{heads}}^{(l)}[b, s, h, :] \in \mathbb{R}^{d_{\text{head}}}$  is the contribution of head  $h$  at token position  $s$ . Multiplying by  $g_h^{(l)}$  scales the entire vector produced by that head uniformly at all positions in the current pass; this matches the "masking/scale" abstraction in Michel et al.

#### H.1.2 GRADIENT-BASED IMPORTANCE

By backpropagating the cross-entropy loss through these gates, we obtain per-head gradients that quantify how loss would change if a head were amplified. Negative gradients indicate supportive heads, while positive gradients reveal harmful ones. Taking the absolute values yields stable importance scores for both visualization and analysis.

At decoding step  $\tau$  with target  $t^*$ ,

$$\mathcal{L} = -\log P(y_\tau = t^* | \mathbf{x}, \mathbf{y}_{<\tau}). \quad (44)$$

We define the *signed sensitivity* and the *importance* for head  $(l, h)$ :

$$S^{(l,h)} = \frac{\partial \mathcal{L}}{\partial g_h^{(l)}}, \quad I^{(l,h)} = \left| S^{(l,h)} \right|. \quad (45)$$

Interpretation: a negative  $S^{(l,h)}$  means "if we scale up this head, loss goes down" (supportive head); a positive  $S^{(l,h)}$  suggests the opposite (potentially harmful). Reporting  $|S|$  yields non-negative "importance heatmaps", while retaining the sign of  $S$  allows for more fine-grained adversarial analysis.

**Connection to first-order Taylor pruning.** Expanding  $\mathcal{L}$  at  $\mathbf{g} = \mathbf{1}$ , a small perturbation  $\Delta g_h^{(l)}$  changes the loss by  $\approx S^{(l,h)} \Delta g_h^{(l)}$ ; therefore  $|S^{(l,h)}|$  can directly serve as a "influence measure for current sample/step", which is the per-sample instantiation of the expected sensitivity  $I_h = \mathbb{E} |\partial \mathcal{L} / \partial \xi_h|$  proposed by Michel et al. under our gate parameterization.

### H.1.3 FROM PER-STEP SCORES TO A RANKING

Finally, we aggregate these per-head importance scores and normalize them either layer-wise or globally. Sorting the normalized scores produces a ranking that highlights which heads are most influential for the given token (or across tokens when averaged). This ranking provides a principled basis for head attribution and comparison.

Collect raw scores

$$\mathcal{I} = \{I^{(l,h)} \mid l = 1, \dots, L; h = 1, \dots, H\}. \quad (46)$$

We consider two normalizations:

**Layer-wise** (robust to layer-scale differences):

$$\tilde{I}_{\text{layer}}^{(l,h)} = \frac{I^{(l,h)}}{\sum_{h'=1}^H I^{(l,h')}}. \quad (47)$$

**Global** (for global ranking across the whole stack):


$$\tilde{I}_{\text{global}}^{(l,h)} = \frac{I^{(l,h)}}{\sum_{l'=1}^L \sum_{h'=1}^H I^{(l',h')}}. \quad (48)$$

We then sort

$$\pi = \text{argsort}(\tilde{I}, \text{descending}), \quad (49)$$

obtaining a head ranking tailored to the current token  $t^*$  (or to an aggregate of tokens; see below).

## H.2 CASE I: REASONING DRIFT



Question - Image (a)

Is the text in this figure trying to talk to Gorbachev?



1296  
1297  
1298  
1299  
1300  
1301  
1302  
1303  
1304  
1305  
1306  
1307  
1308  
1309  
1310  
1311  
1312  
1313  
1314  
1315  
1316  
1317  
1318  
1319  
1320  
1321  
1322  
1323  
1324  
1325  
1326  
1327  
1328  
1329  
1330  
1331  
1332  
1333  
1334  
1335  
1336  
1337  
1338  
1339  
1340  
1341  
1342  
1343  
1344  
1345  
1346  
1347  
1348  
1349

### Vanilla - Image (b)

<think>

Okay, so I need to figure out whether the text in the image is trying to talk to Gorbachev. Let's start out by understanding the context. The image shows a text message or a letter. The text is talking about change and openness and mentions General Secretary Brezhnev. So, Brezhnev was the leader of the Soviet Union during the Cold War. Gorbachev came later, so he might have been the General Secretary later in the 80s.

Looking at the text, it says things like "Come here to this gate!" and "Tear down this wall!". These phrases sound like they're addressing someone specific, probably someone in power authority, which would be the General Secretary. The image also has uses the term "General Secretary Brezhnev", which is a typo for Brezhnev. This reinforces the idea that the speaker is talking to someone in authority.

Also, the message is urgent to to the cause of freedom and peace and economic prosperity in Eastern Europe. This kind of plea would be directed at someone who is in power, like a leader. **Gorbachev** was famous for bringing Soviet policies and for making reforms like glasvez, which aimed at openness and change in the Soviet Union. So, it makes sense that **this message is addressed to him.**

The image shows a letter by someone urging change and openness, asking Brezhnev to do something, and even mentioning tearing down a wall. All these points to in message intended at **Gorbachev**. The reference to Brezhnev being the leader, even if the name spelling is incorrect, points to the main target of the message. So, **yes, the text is trying to talk to Gorbachev.**

</think>

The text in the image is clearly directed at Gorbachev. The mention of General Secretary Brezhnev, the reference to tearing down a wall, and the plea for change and openness all point to in address to Gorbachev, who a leader known for the reforms aimed at opening in in Eastern Europe. Therefore, the message is addressed to Gorbachev.

**\*\*Answer:\*\* Yes, the text in the image is trying to talk to Gorbachev.**

### Ours - Image (c)

<think>

Okay, so I'm trying to determine if the text is trying to talk to Gorbachev. Let's start with the image directly.

The image shows a request to open the gates. However, the text in the image is addressed to General Secretary Brezhnev. The image itself depicts a scene where Gorbachev might be the leader of the Soviet Union. The image also a reference to " Gorbachev, which gives a hint that the leader might Gorbachev's the leader of the Soviet Union.

The image shows a request to open the gates. The text in the image is addressed to General Secretary Brezhnev. The image itself shows a scene where Gorbachev might be the leader of the Soviet Union. The image also refers to Gorbachev, but the text is addressed to Brezhnev.

Therefore, the text is not directly addressing Gorbachev; it is instead addressed to Brezhnev, the leader of the Soviet Union at the time.


</think>

The image shows a request to open the gates. Therefore, the text is not trying to talk to Gorbachev; it is trying to talk to Brezhnev, the leader of the Soviet Union at the time.


**\*\*Answer:\*\* No**, the text in the image is not trying to talk to Gorbachev; it is referring to Brezhnev, the leader of the Soviet Union at the time.

**Analysis.** From the contribution maps, our method markedly amplifies the mid-to-deep *reasoning heads* in Layer 10–23, thereby substantially improving the model’s reasoning capacity. In contrast, the Vanilla model shows no obvious perceptual failure but suffers during reasoning: because the question mentions “Gorbachev”, Vanilla model erroneously binds this cue to the entire message, accumulates evidence under the hypothesis “target = Gorbachev,” and drifts away from the correct logical chain, ultimately misclassifying an appeal addressed to Brezhnev as speaking to Gorbachev. By strengthening these mid-to-deep reasoning heads, our method maintains stable entity grounding and a consistent causal narrative, **preventing reasoning drift** and producing the correct negative answer.


### H.3 CASE II: PERCEPTUAL BIAS



(a) Question



(b) Contribution Map (Vanilla)



(c) Contribution Map (Ours)

Question - Image (a)

Based on the map, did the Republic Party win Texas in the 2020 elections?

Vanilla - Image (b)

<think>

Okay, so I’m trying to determine whether the Republican Party won Texas in the 2020 elections based on the provided map. Let me break this down step.

First, the image shows a map of the United States indicating the election results by state for the 2020 U.S. presidential election. Each state is colored either **\*\*red\*\*** or blue, representing Republican and Democratic wins, respectively. The image shows Texas colored in **\*\*red\*\***.

To confirm, the image legend has the **\*\*red\*\*** color for Republican (Trump) and blue for Democratic (Biden). Therefore, since Texas is colored **\*\*red\*\***, it indicates that the **Republican Party won Texas in the 2020 election**.

There are no other states mentioned that would affect this conclusion. The other states might be colored differently, but Texas being **red** is the key point of this question.

So, putting it all, the image directly shows Texas as **\*\*red\*\***, which means the **Republican Party won Texas in the 2020 election**.

</think>

The image shows Texas colored **red**, which indicates that the Republican Party won Texas in the 2020 election.

**Answer:** Yes, the Republican Party won Texas in the 2020 election.

Ours - Image (c)

<think>

Okay, so I'm trying to determine whether the Republican Party won Texas in the 2020 elections based on the provided map. Let me go through this step by step.

First, I need to understand what the map is showing. It is an electoral map from the 2020 U.S. Presidential Election, where each state is colored either **red** or blue. **Red** represents Republican candidates, and blue represents Democratic candidates. Each state also have a number 2020, so this is the 2020 election results.


Next, I notice that Texas is colored **blue**. This indicates that Texas was won by the Democrats. Therefore, the Republican Party did not win Texas in the 2020 elections.

</think>


The answer is that the Republican Party did not win Texas in the 2020 elections. Texas was won by the Democrats.

**Analysis.** In this case, the contribution maps reveal that our method effectively mitigates the dispersed and unfocused attention observed in the Vanilla model by strengthening both shallow perception heads and deep reasoning heads. The primary issue with the Vanilla model arises in the second paragraph of its reasoning, where it misperceives the color of Texas as red. This initial perceptual bias propagates through its logical reasoning process, ultimately leading to an incorrect conclusion that the Republican Party won Texas. In contrast, our method, supported by **stronger perceptual attention** in shallow layers and more robust reasoning capacity in deeper layers, accurately perceives Texas as blue and correctly infers that the state was won by the Democrats. This demonstrates how our approach enhances the interplay between perception and reasoning, enabling the model to avoid compounding errors and to arrive at the correct answer.


#### H.4 CASE III: COMPREHENSIVE HALLUCINATION



(a) Question



(b) Contribution Map (Vanilla)



(c) Contribution Map (Ours)

Question - Image (a)

According to the image, does the value of the square root of 2 range from 1.40 to 1.41?

1458  
1459  
1460  
1461  
1462  
1463  
1464  
1465  
1466  
1467  
1468  
1469  
1470  
1471  
1472  
1473  
1474  
1475  
1476  
1477  
1478  
1479  
1480  
1481  
1482  
1483  
1484  
1485  
1486  
1487  
1488  
1489  
1490  
1491  
1492  
1493  
1494  
1495  
1496  
1497  
1498  
1499  
1500  
1501  
1502  
1503  
1504  
1505  
1506  
1507  
1508  
1509  
1510  
1511

#### Vanilla - Image (b)

<think>

Okay, so I need to determine whether the square root of 2 ranges from 1.40 to 1.41 based on the image provided. The image shows **several** square roots, and I need to identify the square root of 2.

1. As the image shows **several** square roots, and the first one 2 is approximately **1.44**.
2. The image also shows the square root of 2 as approximately **1.44**.
3. Therefore, the image shows the square root of 2 as approximately **1.44**.

</think>

Answer: The square root of 2 is not **between 1.0 and 1.1**; it is approximately 1.4.

#### Ours - Image (c)

<think>

Okay, so I need to figure out whether the value of the square root of 2, denoted as  $\sqrt{2}$ , falls within the range of 1.40 to 1.41. Let's start directly based on the image and the information provided.

The image shows the square root of 2 as approximately **1.41421**. The image also shows various mathematical constants and numbers sequences, including the first 150 decimal places, and it provides many of the first decimal places as 1.4 ...

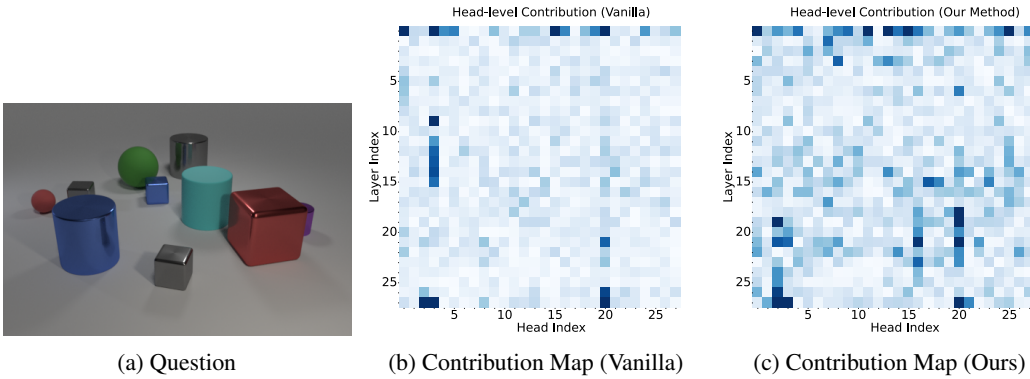
</think>

Answer: No. It's about 1.414, which is slightly above 1.41.

**Analysis.** In this case, we can clearly observe the differences in perception and reasoning abilities between the two methods. The Vanilla model, when tasked with identifying the value of the square root of 2, incorrectly perceived it as 1.44 rather than the correct value of 1.414. This perceptual error indicates insufficient contribution from the shallow perception heads, leading to poor sensitivity to numerical details. Furthermore, the subsequent reasoning process in the Vanilla model was unclear, with its step-by-step logic lacking coherence, and the final answer deviated from the original question. This combination of perceptual inaccuracy and weak reasoning aligns with the contribution map, which shows limited activation in both shallow perception heads and mid-to-deep reasoning heads. In contrast, our method demonstrates significant enhancements in both shallow perception heads and mid-to-deep reasoning heads, resulting in a more balanced improvement across perception and reasoning capabilities. Specifically, in this example, our method not only correctly identified the square root of 2 as 1.41421 but also precisely reasoned that the value is slightly above 1.41, thereby providing an accurate and contextually relevant answer. This highlights the superiority of our approach in achieving **improvements in both perceptual accuracy and logical reasoning**.

#### H.5 CASE IV: COMPREHENSIVE HALLUCINATION

1512  
1513  
1514  
1515  
1516  
1517  
1518  
1519  
1520  
1521  
1522  
1523  
1524  
1525  
1526  
1527  
1528  
1529  
1530  
1531  
1532  
1533  
1534  
1535  
1536  
1537  
1538  
1539  
1540  
1541  
1542  
1543  
1544  
1545  
1546  
1547  
1548  
1549  
1550  
1551  
1552  
1553  
1554  
1555  
1556  
1557  
1558  
1559  
1560  
1561  
1562  
1563  
1564  
1565



#### Question - Image (a)

There is a big thing that is the same color as the tiny ball; what shape is it?

#### Vanilla - Image (b)

<think>

Okay, so I'm trying to solve this problem where I need to figure out the shape of a big object that's the same color as a tiny ball. Let's start with understanding the question carefully. The word "same color" probably refers to the same color group, meaning, or category, but I think it might be a bit too complicated.

The image shows a 3D geometric composition with a large red sphere, which is the "same color" of the small red sphere. Hmm, so the question is asking for the shape of a large object that's the same color as the tiny ball. So the tiny ball is a sphere, and the large one the same color as the tiny ball would also be a sphere as well, but much bigger in size.

The image shows a large red sphere within other shapes, so the large red sphere stands out the same color of another 3D object. The only other 3D object that could be described as the same color would also be another sphere. The image shows a large red sphere, where another 3D object can be a sphere. Therefore, the large red sphere stands out the same color of another sphere, which would also be a sphere as well. So, the large red sphere is a sphere, and another 3D object sharing the same color can be a sphere as well.

So the answer should be a sphere.

</think>

The image shows a large red sphere, which is the same color of another sphere. Therefore, the large red sphere is a sphere, and the one same color would be a sphere as well. The image shows a large red sphere, with another shape being a large red sphere. So, the large red sphere is a sphere, and another 3D object can be a sphere as well.

**\*\*Answer:\*\* Sphere**

#### Ours - Image (c)

<think>

Okay, so I'm trying to figure out the answer to this question: There is a big thing that is the same color as the tiny ball; what shape is it?

The image shows a big red cube next to a tiny red sphere. The tiny ball is a sphere, and the big thing mentioned is a red cube. Since the image shows a big red cube next to a tiny red sphere, the big thing must be a red cube.

1566  
1567  
1568  
1569  
1570  
1571  
1572  
1573  
1574  
1575  
1576  
1577  
1578  
1579  
1580  
1581  
1582  
1583  
1584  
1585  
1586  
1587  
1588  
1589  
1590  
1591  
1592  
1593  
1594  
1595  
1596  
1597  
1598  
1599  
1600  
1601  
1602  
1603  
1604  
1605  
1606  
1607  
1608  
1609  
1610  
1611  
1612  
1613  
1614  
1615  
1616  
1617  
1618  
1619

Therefore, the answer is a red cube.  
</think>

The image shows a big red cube next to a tiny red sphere. The tiny ball is a sphere, and the big thing mentioned is a red cube.

Answer: The shape is a red cube.

**Analysis.** The Vanilla model’s reasoning exhibits systemic failure: it is convoluted, opaque, and internally inconsistent. Contribution maps reveal diffuse attention without stable anchoring, producing weak and unreliable signals across both perceptual and relational heads. More critically, attribute cross-talk collapses the distinction between color and shape, preventing the model from executing the necessary reasoning pipeline of anchoring → same-color filtering → applying the “big” constraint. As a result, it defaults to a spurious shortcut and outputs the incorrect label “sphere.” In contrast, our method imposes structured coordination: shallow heads focus sharply on color and size cues, while deep heads strengthen cross-object reasoning that jointly enforces the predicates “same color as the tiny ball” and “big,” leading to the correct identification of the large red cube. These results highlight that without principled alignment between perception and reasoning, models are prone to brittle and misleading inference.