

VIB-Probe: Detecting and Mitigating Hallucinations in Vision-Language Models via Variational Information Bottleneck

Feiran Zhang^{*}, Yixin Wu^{*}, Zhenghua Wang, Xiaohua Wang,
Changze Lv, Xuanjing Huang, Xiaoqing Zheng[†]

College of Computer Science and Artificial Intelligence, Fudan University, Shanghai, China
Shanghai Key Laboratory of Intelligent Information Processing

{frzhang25,yixinwu23}@m.fudan.edu.cn {xjhuang,zhengxq}@fudan.edu.cn

Abstract

Vision-Language Models (VLMs) have demonstrated remarkable progress in multimodal tasks, but remain susceptible to hallucinations, where generated text deviates from the underlying visual content. Existing hallucination detection methods primarily rely on output logits or external verification tools, often overlooking their internal mechanisms. In this work, we investigate the outputs of internal attention heads, postulating that specific heads carry the primary signals for truthful generation. However, directly probing these high-dimensional states is challenging due to the entanglement of visual-linguistic syntax and noise. To address this, we propose VIB-Probe, a novel hallucination detection and mitigation framework leveraging the Variational Information Bottleneck (VIB) theory. Our method extracts discriminative patterns across layers and heads while filtering out semantic nuisances through the information bottleneck principle. Furthermore, by leveraging the gradients of our VIB probe, we identify attention heads with strong causal influence on hallucinations and introduce an inference-time intervention strategy for hallucination mitigation. Extensive experiments across diverse benchmarks demonstrate that VIB-Probe significantly outperforms existing baselines in both settings. Our code will be made publicly available.

1 Introduction

Vision-Language Models (VLMs) have emerged as an influential force in multimodal artificial intelligence, demonstrating a sophisticated ability to generate contextually rich natural language descriptions grounded in visual patterns (Ye et al., 2023; Bai et al., 2025; Li et al., 2023b). By integrating visual encoders with large language models, VLMs have shown impressive performance across




Figure 1: Hallucination detection performance comparison across 6 benchmarks, based on the AUPRC metric. Our proposed VIB-Probe consistently achieves state-of-the-art overall results.

diverse vision-language tasks, including image captioning, visual question answering, and multimodal machine translation (Liu et al., 2023; Zhu et al., 2024a; Chen et al., 2023; Lee et al., 2024). Despite these advancements, VLMs remains prone to hallucinations, where generated descriptions are unfaithful to the objects or relations present in the source image (Yin et al., 2024; He et al., 2025). This lack of visual fidelity undermines the reliability and applicability of VLMs, particularly in high-stakes domains that demand precise multimodal reasoning and factual accuracy.

Existing approaches to hallucination detection primarily rely on surface-level confidence indicators, such as logit-based entropy or divergence (Fieback et al., 2025b,a; Zollicoffer et al., 2025; Hendrycks and Gimpel, 2017). These heuristic-based classifiers typically exploit only a narrow slice of the model’s internal dynamics and depend on manually engineered features that may fail to generalize across diverse architectures. Consequently, developing robust and efficient mechanisms for detecting hallucinations in VLM outputs remains a significant open challenge.

^{*}Equal contribution.

[†]Corresponding Author.

Recent research in interpretability suggests that VLM hallucinations are often rooted in fragile attention dynamics introduced by the visual modality (Jiang et al., 2025b; Tang et al., 2025; Jiang et al., 2025a; Yang et al., 2025). Specifically, a model may attend to irrelevant regions, infer non-existent objects, or over-rely on linguistic priors at the expense of visual grounding (Zheng et al., 2025). Crucially, this informational drift is not confined to the final output layer, but rather emerges progressively across internal layers (Zhang et al., 2025a; He et al., 2025). Hallucination-related signals are encoded within the outputs of specific attention heads across layers, while these signals are often entangled with task-irrelevant syntactic noise.

Motivated by these insights, we propose **VIB-Probe**, a novel framework grounded in **Variational Information Bottleneck (VIB)** theory (Tishby et al., 2000; Alemi et al., 2017). As illustrated in Figure 2, VIB-Probe distills a compact latent representation from the high-dimensional attention head outputs across all Transformer layers, retaining information predictive of hallucinations while suppressing noise and spurious correlations. We employ a multi-layer encoder to capture the bottleneck features for robust detection. Furthermore, we extend our approach to hallucination mitigation by applying gradient-based attribution from the probe’s logits to the attention heads. By these means, we identify specific “hallucination-sensitive” heads that exert strong causal influence on unfaithful generation. We then introduce an inference-time mitigation strategy that selectively suppresses these heads during decoding when the detected hallucination risk exceeds a predefined threshold. Extensive experiments across multiple benchmarks demonstrate that our approach yields consistent gains in both detection and mitigation across diverse VLM architectures.

The contributions of this study can be summarized as follows:

- We introduce VIB-Probe, a novel framework for hallucination detection that exploits the information of multi-layer, multi-head attention outputs in VLMs. By grounding our approach in Variational Information Bottleneck theory, we distill a compact yet highly predictive latent representation, enabling robust detection across both open-ended generation and closed-form QA settings.
- We propose a training-free, inference-time

mitigation strategy that bridges the gap between detection and control. By employing probe-based attribution, we identify hallucination-sensitive attention heads and dynamically suppress their outputs upon high risks of hallucination.

- We conduct comprehensive experiments across both discriminative and generative hallucinatory benchmarks, demonstrating that VIB-Probe achieves state-of-the-art performance in hallucination detection and mitigation, while further highlighting its robustness and generalizability across diverse perturbations and architectures.

2 Related Work

2.1 Hallucinations in VLMs

Vision-Language Models (VLMs) integrate visual encoders with Large Language Models (LLMs) via projection layers to enable multimodal reasoning (Liu et al., 2023). Compared to factual errors in text-only LLMs, VLM hallucinations mainly arise from failures in visual grounding. They are commonly grouped into **object**, **attribute**, and **relational** hallucinations (Zhou et al., 2024).

Hallucination Detection Early detectors relied on shallow output statistics (e.g., token confidence or entropy), which often generalize poorly under complex reasoning. Reference-free methods aim to verify outputs without external evidence (Li et al., 2024; Prabhakaran et al., 2025). Recent work probes mechanistic signals in attention, e.g., Lookback Lens (Chuang et al., 2024) and OPERA (Huang et al., 2024), by analyzing aggregated attention patterns during decoding, improving discrimination between grounded and hallucinatory outputs. Building on this direction, we move beyond raw attention weights and apply VIB to attention head outputs to better isolate hallucination-relevant signals from high-dimensional noise.

Hallucination Mitigation Mitigation methods are typically categorized into **training-based**, **post-generation**, and **inference-time** approaches. Training-based methods enhance robustness via instruction-tuning on curated data (Liu et al., 2023; Zhang et al., 2024; Zhou et al., 2024), while post-generation methods employ external verifiers for iterative refinement (Yin et al., 2024). Inference-time

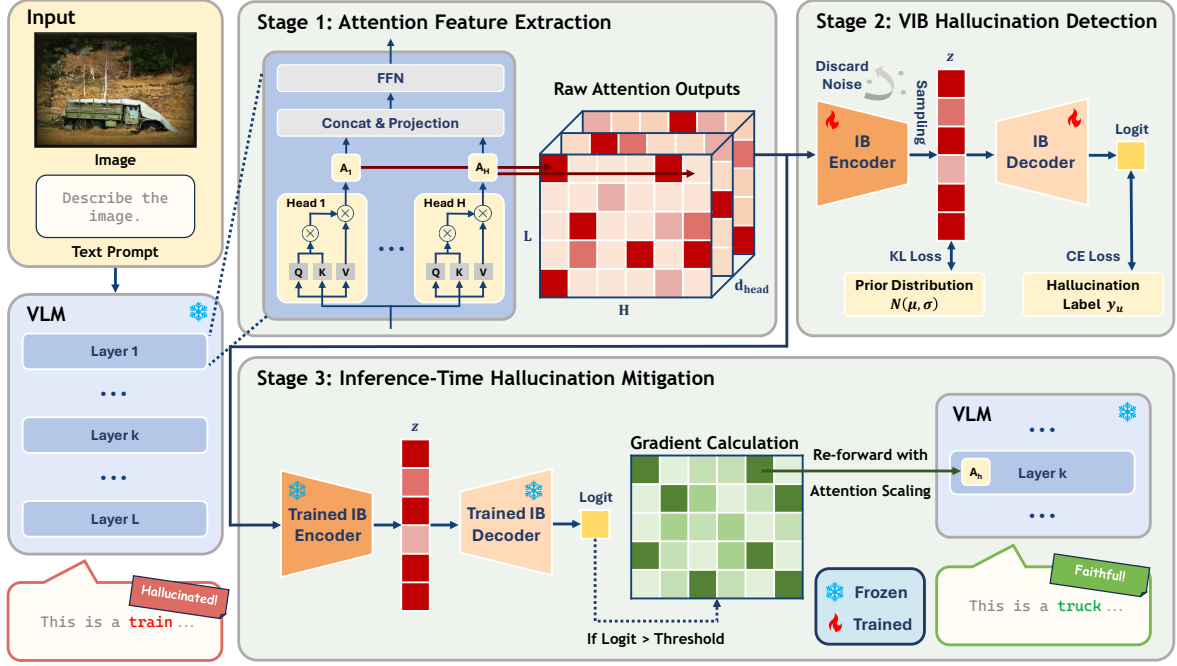


Figure 2: Overview of the VIB-Probe framework. The Information Bottleneck (IB) theory is leveraged to detect and mitigate hallucinations by probing internal attention features. **Stage 1:** We extract raw output vectors from all attention heads across all Transformer layers ($L \times H$) during VLM decoding. **Stage 2:** The extracted features are fed into an IB Encoder, which compresses the high-dimensional input into a compact latent representation z . This process filters out task-irrelevant noise while retaining minimal sufficient statistics for prediction. **Stage 3:** Leveraging the trained VIB modules, inference-time mitigation is achieved by suppressing the attention heads with high risks of hallucination for each token, producing a more faithful output generation.

interventions have gained attention for avoiding re-training costs: VCD (Leng et al., 2024a) reduces reliance on linguistic priors via visual perturbation, and PAI (Liu et al., 2024b) and IBD (Zhu et al., 2025) strengthen visual grounding by adjusting attention to image tokens. Our method follows this paradigm but introduces gradient-based attribution to target hallucination-sensitive heads for training-free intervention.

2.2 Information Bottleneck Theory

The Information Bottleneck (IB) principle (Tishby et al., 2000) serves as a robust information-theoretic framework for regularizing internal representations. By compressing model input to minimize mutual information, IB encourages the model to discard irrelevant features while retaining essential semantic content, thereby enhancing generalization capabilities. This principle has been extensively adopted across various machine learning paradigms, including image generation (Jeon et al., 2025), generative classification (Ardizzone et al., 2020), explanation regeneration (Li et al., 2023c), and retrieval-augmented generation (Zhu

et al., 2024b). To operationalize the IB objective in deep neural networks, Alemi et al. (2017) introduced the Variational Information Bottleneck (VIB). Inspired by the architecture of Variational Autoencoders (VAEs) (Kingma and Welling, 2014), VIB employs a variational approach to approximate the IB trade-off and has demonstrated significant efficacy in parsing (Li and Eisner, 2019a), low-resource fine-tuning (Mahabadi et al., 2021), and graph structure learning (Sun et al., 2022) domains.

3 Method

3.1 Preliminaries

Attention Head Output For vision-language models of the most prevalent LLaVA-style (Liu et al., 2023) architecture, a vision encoder is coupled with a decoder-only large language model via a projection layer. An input image is encoded into a sequence of visual tokens, which are projected into the LLM’s embedding space and concatenated with the textual prompt tokens. This multimodal sequence is processed by L Transformer decoder layers, each containing H attention heads. During each autoregressive decoding step t , the model pre-

dicts the next token conditioned on the image, the prompt, and the previously generated tokens. For a given layer $l \in [1, L]$ and head $h \in [1, H]$, the input hidden states \mathbf{X}^l are transformed into query, key, and value matrices:

$$Q_{l,h} = X^l W_{l,h}^Q, K_{l,h} = X^l W_{l,h}^K, V_{l,h} = X^l W_{l,h}^V \quad (1)$$

where $W^Q, W^K, W^V \in \mathbb{R}^{d_{model} \times d_h}$ are the projection weights. The attention weights $\mathbf{A}_{l,h}$ are then computed via the scaled dot-product:

$$\mathbf{A}_{l,h} = \text{softmax} \left(\frac{\mathbf{Q}_{l,h} (\mathbf{K}_{l,h})^\top}{\sqrt{d_h}} \right) \quad (2)$$

To capture the raw, disentangled information flow prior to the final head-mixing, we extract the **pre-projection attention head output** $\mathbf{O}_{l,h}$:

$$\mathbf{O}_{l,h} = \mathbf{A}_{l,h} \mathbf{V}_{l,h} \quad (3)$$

For each token generated at step t , we aggregate $\mathbf{O}_{l,h}$ across all layers and heads to construct a representation tensor $\mathcal{T} \in \mathbb{R}^{L \times H \times d_h}$. This tensor provides a comprehensive ‘‘snapshot’’ of the model’s internal multimodal processing and serves as the primary input for our VIB-Probe framework.

Information Bottleneck The Information Bottleneck principle defines an optimal representation \mathbf{z} of an input signal \mathbf{v} that maximizes its predictive power regarding a target \mathbf{y} while minimizing the information retained from \mathbf{v} . Formally, the IB objective is formulated as the following constrained optimization:

$$\min \mathcal{L}_{IB} = \beta I(\mathbf{v}; \mathbf{z}) - I(\mathbf{z}; \mathbf{y}), \quad (4)$$

where $I(\cdot; \cdot)$ denotes mutual information and $\beta > 0$ is a Lagrange multiplier controlling the trade-off between *compression* (minimizing $I(\mathbf{v}; \mathbf{z})$) and *prediction* (maximizing $I(\mathbf{z}; \mathbf{y})$). By penalizing $I(\mathbf{v}; \mathbf{z})$, the model is forced to discard ‘‘semantic nuisances’’ features that are irrelevant to the grounding of visual content.

Directly optimizing Eq. (4) is generally intractable, as computing mutual information requires knowledge of the underlying data distributions. The **Variational Information Bottleneck** addresses this by introducing a variational upper bound on the compression term $I(\mathbf{v}; \mathbf{z})$ and replaces the predictive term with a tractable likelihood model. Specifically, VIB parameterizes an

encoder $p_\theta(\mathbf{z} | \mathbf{v})$ and a decoder $p_\phi(\mathbf{y} | \mathbf{z})$, and uses a prior $r(\mathbf{z})$. The resulting objective is:

$$\min \mathcal{L}_{VIB} = \beta \mathbb{E}_{\mathbf{v}} [\text{KL}(p_\theta(\mathbf{z} | \mathbf{v}) \| r(\mathbf{z}))] + \mathbb{E}_{\mathbf{v}} \mathbb{E}_{\mathbf{z} \sim p_\theta} [-\log p_\phi(\mathbf{y} | \mathbf{z})], \quad (5)$$

where the first term acts as a *compression* regularizer and the second term represents the negative log-likelihood loss of *prediction*. In practice, for binary labels, the prediction loss is implemented as a binary cross-entropy (BCE) loss.

3.2 Hallucination Detection via VIB on Attention Head Outputs

Building on prior observations, we propose VIB-Probe, a lightweight detector based on Information Bottleneck theory. VIB-Probe is designed to aggregate the internal holistic attention information of VLMs for hallucination detection.

Problem Setup Given an input image and a text prompt, a VLM generates a total of N tokens autoregressively. At each decoding step u , we extract the pre-projection attention head outputs from all layers and heads (Eq. (3)), stacking them into a tensor $\mathcal{T} \in \mathbb{R}^{L \times H \times d_h}$. Our goal is to predict a binary hallucination label $\mathbf{y}_u \in \{0, 1\}$, where $\mathbf{y}_u = 1$ denotes a hallucination and $\mathbf{y}_u = 0$ signifies those visually-grounded. The resulting training set is defined as $\mathcal{D} = \{(\mathcal{T}_u, \mathbf{y}_u)\}_{u=1}^N$.

VIB Detector Architecture We treat the tensor \mathcal{T} as the raw internal signal \mathbf{v}_u and feed it into a lightweight convolutional or multi-layer perceptron encoder $f_\psi(\cdot)$ to extract a high-level feature representation $\mathbf{h}_u \in \mathbb{R}^{d_f}$:

$$\mathbf{v}_u := \mathcal{T}, \quad \mathbf{h}_u = f_\psi(\mathbf{v}_u), \quad (6)$$

A variational bottleneck then parameterizes an approximate posterior $q_\psi(\mathbf{z}_u | \mathbf{v}_u)$ as a multivariate diagonal Gaussian:

$$q_\psi(\mathbf{z}_u | \mathbf{v}_u) = \mathcal{N}(\boldsymbol{\mu}_u, \text{diag}(\boldsymbol{\sigma}_u^2)), \quad (7)$$

$$[\boldsymbol{\mu}_u, \log \boldsymbol{\sigma}_u^2] = g_\psi(\mathbf{h}_u),$$

where the decoder g_ψ is a single or multiple fully-connected layers. Similar to (Li and Eisner, 2019b), we sample \mathbf{z}_u during training using the reparameterization trick (Kingma and Welling, 2014) to ensure end-to-end differentiability:

$$\mathbf{z}_u = \boldsymbol{\mu}_u + \boldsymbol{\sigma}_u \odot \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (8)$$

At inference time, we adopt a deterministic approach for stability, utilizing the mean representation $\mathbf{z}_u = \boldsymbol{\mu}_u$ for prediction. Finally, a linear classification layer computes the **hallucination risk logit** s_u and the corresponding probability \hat{p}_u through the sigmoid function $\sigma(\cdot)$:

$$s_u = w^\top \mathbf{z}_u + b, \quad \hat{p}_u = \sigma(s_u), \quad (9)$$

Training Objective Following the Variational Information Bottleneck principle, we optimize the latent representation \mathbf{z}_u to be maximally informative about the label \mathbf{y}_u while remaining minimally sufficient with respect to the input \mathbf{v}_u . We regularize the information flow by penalizing the KL divergence between the approximate posterior $q_\psi(\mathbf{z}_u | \mathbf{v}_u)$ and a standard normal prior $r(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$. The token-level detection loss is formulated as:

$$\mathcal{L}_{\text{det}} = \mathbb{E}_{(\mathbf{v}_u, \mathbf{y}_u) \sim \mathcal{D}} \left[\underbrace{\text{BCE}(\mathbf{y}_u, \hat{p}_u)}_{\text{prediction}} + \beta \underbrace{\text{KL}(q_\psi(\mathbf{z}_u | \mathbf{v}_u) \| r(\mathbf{z}))}_{\text{compression}} \right], \quad (10)$$

where $\beta > 0$ is a Lagrange multiplier that controls the trade-off between prediction accuracy and representation compression. The first term here is the standard binary cross-entropy (BCE) loss:

$$\text{BCE}(\mathbf{y}_u, \hat{p}_u) = -\mathbf{y}_u \log \hat{p}_u - (1 - \mathbf{y}_u) \log(1 - \hat{p}_u), \quad (11)$$

Given our choice of a diagonal Gaussian posterior (Eq. (7)), the KL term has a closed form:

$$\text{KL}(\mathcal{N}(\boldsymbol{\mu}_u, \text{diag}(\boldsymbol{\sigma}_u^2)) \| \mathcal{N}(\mathbf{0}, \mathbf{I})) = \frac{1}{2} \sum_{i=1}^{d_z} (\mu_{u,i}^2 + \sigma_{u,i}^2 - \log \sigma_{u,i}^2 - 1), \quad (12)$$

where d_z denotes the dimension of the bottleneck latent space. During training, we minimize the objective function L_{det} with respect to the parameters of the encoder f_ψ and decoder g_ψ . At inference time, the raw logit s_u is utilized to assess hallucination risk and further mitigation.

3.3 Hallucination Mitigation

Building upon the trained VIB detector, we propose an inference-time mitigation strategy that translates detection signals into actionable model control. By

attributing the predicted hallucination risk to specific internal components, we can dynamically suppress the most influential attention heads that leads to hallucinations.

At each decoding step u , we perform a VLM forward pass to extract the attention head outputs $\mathcal{T} \in \mathbb{R}^{L \times H \times d_h}$ and compute the VIB hallucination risk logit s_u . If $s_u \leq \tau$ (where τ is a risk threshold), the model samples the next token normally. If $s_u > \tau$, an intervention is triggered to rectify the potential hallucination, by modifying attention heads and regenerating the token.

Gradient-based Attribution and Head Selection

To identify which heads contribute most to hallucination risks, we perform a backward pass through the **frozen** VIB detector. We compute the gradient of the risk logit by each attention head at the current step: $g^{l,h} = \nabla_{o^{l,h}} s_u$. Since our intervention involves scaling the head outputs by a coefficient $\alpha^{l,h}$, such that the modified output becomes $\tilde{o}^{l,h} = \alpha^{l,h} o^{l,h}$. The sensitivity of the risk logit to this scaling is:

$$\nabla_{\alpha^{l,h}} s_u = \langle g^{l,h}, o^{l,h} \rangle. \quad (13)$$

We define the **head importance score** as the magnitude of this sensitivity: $I^{l,h} = |\langle g^{l,h}, o^{l,h} \rangle|$. We then select the set of most influential heads $\mathcal{K} = \text{TopK}(\{I^{l,h}\})$ for targeted suppression.

Inference-Time Single-Step Head Suppression

We initialize all the output scaling coefficients as $\alpha^{l,h} = 1$. For the heads identified in \mathcal{K} , we apply a single-step suppression update to reduce hallucinatory risk:

$$\alpha^{l,h} \leftarrow 1 - \lambda \cdot \text{ReLU}(\langle g^{l,h}, o^{l,h} \rangle), \quad (l, h) \in \mathcal{K}, \quad (14)$$

where λ is a hyperparameter for controlling the suppression strength. We keep $\alpha_i^{l,h} = 1$ unmodified for $(l, h) \notin \mathcal{K}$. Finally, we rerun the VLM decoding step using the modified head outputs $\tilde{o}^{l,h} = \alpha^{l,h} o^{l,h}$ to obtain the edited logits and then sample the regenerated token.

4 Experiments

4.1 Benchmarks

We evaluate VIB-Probe across a diverse suite of hallucination detection benchmarks covering both discriminative and generative datasets.

Benchmark	Method	MiniGPT-4		LLaVA-v1.5		LLaVA-v1.6		Qwen2.5-VL		Average	
		A-ROC	A-PR	A-ROC	A-PR	A-ROC	A-PR	A-ROC	A-PR	A-ROC	A-PR
Discriminative Benchmarks											
POPE	AvgEnt	76.27	68.64	77.43	67.52	79.51	70.66	78.99	70.20	78.05	69.26
	AvgProb	61.56	63.39	64.25	63.90	63.06	66.44	68.28	64.55	64.29	64.57
	RepProbing	91.18	92.30	94.68	94.50	93.01	93.87	96.82	95.89	93.92	94.14
	MetaToken	89.69	90.07	93.07	92.22	94.21	94.33	94.10	94.84	92.77	92.87
	DHCP	93.80	91.76	94.87	94.53	94.92	94.20	96.80	96.52	95.10	94.25
	VIB-Probe (ours)	94.19	93.37	96.52	96.96	95.99	95.51	96.98	96.40	95.92	95.56
AMBER	AvgEnt	61.25	58.53	62.05	62.43	65.48	62.20	66.42	66.80	63.80	62.49
	AvgProb	59.68	55.90	64.74	60.28	64.81	63.79	64.51	63.33	63.44	60.83
	RepProbing	72.25	71.11	77.53	76.82	75.35	74.74	74.61	74.52	74.94	74.30
	MetaToken	74.18	73.39	74.60	75.20	74.46	74.81	75.59	75.10	74.71	74.63
	DHCP	83.18	82.27	82.07	81.89	84.77	83.64	84.77	83.98	83.70	82.95
	VIB-Probe (ours)	83.40	82.94	82.95	82.43	85.99	85.91	85.51	85.82	84.46	84.28
Generative Benchmarks											
M-HalDetect	AvgEnt	54.90	38.22	53.27	36.87	55.90	37.52	63.52	41.39	56.90	38.50
	AvgProb	54.00	38.93	59.01	39.54	60.21	40.36	66.47	42.71	59.92	40.39
	RepProbing	78.21	70.04	77.18	69.80	77.38	71.20	80.92	71.13	78.42	70.54
	MetaToken	77.28	69.13	82.02	71.14	81.23	73.56	75.19	69.30	78.93	70.78
	DHCP	79.58	74.62	88.13	80.20	86.51	78.87	84.82	80.40	84.76	78.52
	VIB-Probe (ours)	83.33	77.26	89.98	82.35	88.36	81.23	85.17	80.79	86.71	80.41
COCO-Caption	AvgEnt	52.08	30.81	58.93	32.01	55.89	34.72	60.21	35.26	56.78	33.20
	AvgProb	55.36	32.67	54.45	33.92	58.88	36.95	59.05	34.18	56.94	34.43
	RepProbing	65.96	56.88	72.33	62.56	71.92	64.99	77.11	66.14	71.83	62.64
	MetaToken	65.70	55.34	67.28	58.30	67.23	59.35	70.89	61.20	67.78	58.55
	DHCP	69.52	58.13	74.06	64.99	74.20	68.17	74.14	67.64	72.98	64.73
	VIB-Probe (ours)	72.55	62.82	75.24	66.51	75.16	69.32	76.83	70.52	74.95	67.29

Table 1: Results of hallucination detection across multiple baselines on discriminative and generative benchmarks. We report AUROC (A-ROC) and AUPRC (A-PR) as metrics and compare our method with baselines across four base VLMs (MiniGPT-4, LLaVA-v1.5-7B, LLaVA-v1.6-Mistral-7B, and Qwen2.5-VL-7B-Instruct).

POPE POPE (Li et al., 2023d) is a standard diagnostic for VLM object hallucinations. For each image, the dataset provides three positive questions regarding existing objects and three negative questions. The negative samples are selected based on random sampling (*Random*), global frequency (*Popular*), or co-occurrences with present objects (*Adversarial*). Throughout our experiments, we utilize the official POPE dataset, which comprises a total of 9,000 questions across 1,500 images.

AMBER AMBER (Wang et al., 2023) extends the scope of evaluation beyond POPE’s objects to include *attribute* and *relation* hallucinations. The origin dataset contains 14,216 discriminative queries. We randomly sampled 5,000 queries from the original dataset for the experiments.

M-HalDetect M-HalDetect (Gunjal et al., 2024) provides a more granular assessment of hallucinations in detailed responses. Based on the MS COCO (Lin et al., 2014) 2014 validation set, it includes 12,800 training and 3,200 validation samples. Responses are segmented and expert-

annotated into four categories: *Accurate*, *Inaccurate*, *Analysis*, and *Unsure*. Approximately 25% of segments are labeled as hallucinatory, presenting a challenge for fine-grained description tasks.

COCO-Caption To evaluate generative hallucinations in open-ended captioning, We randomly sampled 2,000 images from the MS COCO 2014 validation set, splitting them into training and validation subsets by an 80:20 ratio. We identify hallucinations from the image captions generated.

4.2 Hallucination Detection

4.2.1 Experimental Setup

Base Models and Datasets We evaluate the efficacy of VIB-Probe on four representative VLMs: MiniGPT-4 (Zhu et al.), LLaVA-v1.5-7B (Liu et al., 2023), LLaVA-v1.6-Mistral-7B (Liu et al., 2024a), and Qwen2.5-VL-7B-Instruct (Bai et al., 2025). Experiments cover two extensively adopted discriminative benchmarks, POPE and AMBER (subset averages reported), alongside two generative datasets M-HalDetect and COCO-Caption. To as-

ness detection performance, we report AUPRC and AUROC (Davis and Goadrich, 2006). Detailed configurations are provided in Appendix B.1.

Baselines We compare our method with classic methods based on model uncertainty and probing classifiers, as well as two strong baselines. Meta-Token (Fieback et al.) trains a lightweight classifier by ensembling multiple statistical features derived from object token generation. Meanwhile, DHCP (Zhang et al., 2025b) detects hallucinations by training a lightweight prober that leverages cross-modal attention patterns during decoding. Implementation details are included in Appendix A.2.

Implementation Details VIB-Probe is implemented as a multi-layer MLP encoder for dimensionality reduction, followed by a simple linear decoder. The latent distribution is constrained by a standard Gaussian prior $\mathcal{N}(0, I)$. We set the bottleneck dimension $d = 256$. We optimize the framework using AdamW with a learning rate of 2×10^{-5} and a linear warm-up for the KL-divergence coefficient β , capped at 3×10^{-4} . For discriminative tasks, we extract representations from the last answer token; for generative tasks, we utilize the internal states corresponding to the final token of each sentence or annotated span.

4.2.2 Result Analysis

Table 1 presents the hallucination detection performance of baselines and our VIB-Probe across both discriminative and generative benchmarks. Our VIB-Probe consistently outperforms existing state-of-the-art methods across the four evaluated VLMs. While achieving competitive results on the discriminative benchmarks (+1.20%), our method also demonstrates a pronounced advantage on the challenging generative tasks (+2.84%). This underscores its superior capability in detecting hallucinations within complex, free-form text.

Among the baselines, uncertainty-based heuristics like AvgEnt and AvgProb perform reasonably on closed-set tasks but falter in generative settings. Conversely, RepProbing significantly outperforms these metrics, confirming that hidden states serve as effective indicators of visual fidelity. While MetaToken excels at object-level detection, its performance degrades on generative benchmarks, likely because its heuristic features are too specialized for object tokens to capture span-level or sentence-level relational errors. DHCP emerges as the strongest baseline, validating the utility of

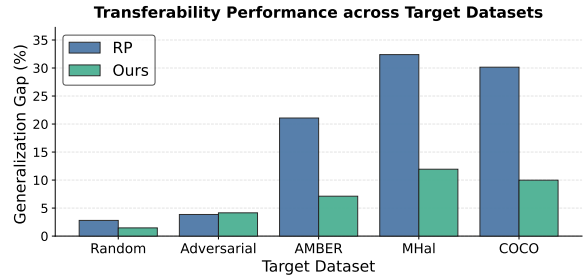


Figure 3: Generalization gap from POPE-Popular to other test sets. A lower generalization gap indicates stronger transferability performance. Results are compared based on LLaVA-v1.5-7B.

Method	POPE		COCO	
	A-ROC	A-PR	A-ROC	A-PR
AvgEnt	53.14	52.73	50.98	31.77
RepProbing	76.92	77.46	63.23	54.10
DHCP	84.77	83.63	66.40	58.58
VIB-Probe	88.78	87.30	73.76	64.81

Table 2: Robustness performance of hallucination detection on input images with random perturbations. Methods are compared based on LLaVA-v1.5-7B.

attention-based hallucination detection.

Transferability Performance To evaluate the ability of VIB-Probe to extract representations highly-correlated with hallucinations that remain invariant to shifts in data distribution and task format, we conducted a series of cross-distribution and cross-task generalization experiments. We first assessed cross-distribution generalization by training on the POPE-Popular subset and evaluating it across all discriminative benchmarks. Subsequently, we evaluated cross-task generalization by evaluating the POPE-Popular detector directly on generative tasks. As illustrated in Figure 3, while baseline methods like RepProbing experiences significant performance degradation under domain shift (e.g., a 32.4% decline on M-HalDetect), our VIB-Probe exhibits stability and stronger transferability. This indicates that the Information Bottleneck successfully distills domain-invariant hallucination signals from the internal attention dynamics, effectively filtering out dataset-specific biases.

Robustness Performance To verify that VIB-Probe isolates compact representations specifically aligned with hallucination signals rather than low-level visual noise, we further designed a robustness experiment to evaluate its performance under varying image quality conditions. Specifi-

Base Model	Method	POPE		COCO	
		ACC \uparrow	F1 \uparrow	C _i \downarrow	C _s \downarrow
LLaVA-v1.5-7B	Vanilla	82.6	83.3	18.2	59.3
	BeamSearch	82.2	84.1	19.5	60.6
	PAI	84.0	84.6	14.4	46.7
	VCD	83.6	83.9	15.8	52.2
	VIB-Probe	83.7	85.2	14.1	44.9
LLaVA-v1.6-7B	Vanilla	84.1	85.1	11.8	40.7
	BeamSearch	84.3	85.6	10.9	39.2
	PAI	87.9	88.4	9.2	35.3
	VCD	86.3	87.8	9.0	36.4
	VIB-Probe	88.2	89.5	8.7	32.1

Table 3: Performance of hallucination mitigation on the validation sets of POPE and COCO. Methods are compared based on LLaVA-v1.5-7B.

cally, we introduced random perturbations to the input images from the POPE and COCO-Caption datasets for evaluation only. These perturbations include rotation, Gaussian blur, and brightness adjustments, while ensuring that the ground-truth labels remained valid. Results in Table 2 demonstrate that VIB-Probe maintains high detection accuracy despite these image perturbations. This resilience indicates that our framework effectively extracts the core internal states associated with unfaithful generation, even when the model’s representations are subjected to external visual noise.

4.3 Hallucination Mitigation

To validate our mitigation capabilities, we performed experiments on the POPE benchmark and a randomly selected 500-image subset of COCO val 2014. For generative evaluation, we utilized the CHAIR (Rohrbach et al., 2018) metric, which quantifies object-level hallucinations by cross-referencing generated entities against ground-truth object lists. For POPE, we reported the Accuracy and F1 score metrics. Experimental results in Table 3 indicate that while contrastive decoding-based VCD (Leng et al., 2024b) provide a viable baseline for hallucination mitigation, inference-time attention intervention strategies such as PAI (Liu et al., 2024c) generally delivers stronger performance. VIB-Probe attains the best performance across most metrics as compared to baselines, demonstrating the effectiveness of intervention on hallucination-related attention heads.

4.4 Ablation Studies

Information Bottleneck Constraint To verify the effectiveness of the Information Bottleneck con-

Base Model	Setting	POPE	M-Hal
LLaVA-v1.5-7B	VIB-Probe	96.96	82.35
	– KL Loss	88.32	71.91
Qwen2.5-VL-7B-Instruct	VIB-Probe	96.40	80.79
	– KL Loss	92.11	67.34

Table 4: Impact of removing the Information Bottleneck constraint (KL loss) on detection performance. The AUPRC metric is reported.

Base Model	Layers	POPE	M-Hal
LLaVA-v1.5-7B	<i>All</i>	96.96	82.35
	<i>1–8</i>	68.71	49.66
	<i>1–16</i>	73.80	52.39
	<i>9–24</i>	91.45	69.18
	<i>17–32</i>	93.22	65.94
	<i>25–32</i>	89.68	59.44

Table 5: Impact of layers selected for the extraction of attention head outputs on detection performance. The AUPRC metric is reported.

straint, we test a variant that retains the VIB-Probe encoder-decoder structure but removes the KL loss, solely optimizing the BCE loss. Experimental results in Table 4 indicate that removing the KL loss degrades performance to a level comparable to the RepProbing baseline. This further demonstrates that explicitly introducing the Information Bottleneck KL divergence constraint is crucial to our gains, making our approach more effective than a simple probing classifier.

Layer Feature Selection We evaluate the impact of extracting features from a specific layers to train the VIB, rather than utilizing attention heads from all VLM layers. For LLaVA-v1.5-7B with 32 layers, results on POPE and M-HalDetect are presented in Table 5. Using information from only a small subset of layers results in performance degradation, particularly on the more challenging M-HalDetect. Notably, employing only deeper layers yields better performance than using shallower layers, likely due to the fact that cross-modal information is not yet fully fused in shallow layers.

5 Conclusion

Hallucinations remain a formidable challenge for the deployment of Vision-Language Models in reliability-critical environments. Unfaithful generations often emerge progressively from internal attention dynamics, rather than solely from the final output. To address this, we introduce VIB-Probe, a framework that leverages high-dimensional multi-

head attention outputs across all layers. By grounding our approach in the Variational Information Bottleneck theory, we effectively distill a compact latent representation that isolates hallucination-related signals from task-irrelevant noise. Beyond detection, we further demonstrate that VIB-Probe supports lightweight inference-time mitigation by identifying and down-weighting a small set of hallucination-sensitive heads upon high risks. Extensive experiments across diverse architectures and benchmarks demonstrate state-of-the-art performance in detection and mitigation, highlighting the robustness and practicality of our framework.

Limitations

Our study primarily focuses on transformer-based vision–language models with standard attention mechanisms. While these architectures cover most widely used VLMs, the applicability of VIB-Probe to alternative multimodal architectures or models that do not rely on explicit attention structures has not been explored and remains an interesting direction for future work. In addition, our method requires access to the model’s internal representations and attention outputs, which restricts it to a white-box setting and may be a potential limitation.

References

- Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. 2017. [Deep variational information bottleneck](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Lynton Ardizzone, Radek Mackowiak, Carsten Rother, and Ullrich Köthe. 2020. [Training normalizing flows with the information bottleneck for competitive generative classification](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yanzhi Zhu, Ming-Hsuan Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025. [Qwen2.5-vl technical report](#). *CoRR*, abs/2502.13923.
- Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. 2023. [Shikra: Unleashing multimodal llm’s referential dialogue magic](#). *arXiv preprint*.
- Yung-Sung Chuang, Linlu Qiu, Cheng-Yu Hsieh, Ranjay Krishna, Yoon Kim, and James Glass. 2024. [Lookback lens: Detecting and mitigating contextual hallucinations in large language models using only attention maps](#). *Preprint*, arXiv:2407.07071.
- Jesse Davis and Mark Goadrich. 2006. [The relationship between precision-recall and roc curves](#). In *Proceedings of the 23rd international conference on Machine learning - ICML '06*.
- Laura Fieback, Nishilkumar Balar, Jakob Spiegelberg, and Hanno Gottschalk. 2025a. [Efficient contrastive decoding with probabilistic hallucination detection - mitigating hallucinations in large vision language models -](#). *arXiv preprint*.
- Laura Fieback, Jakob Spiegelberg, and Hanno Gottschalk. 2025b. [Metatoken: Detecting hallucination in image descriptions by meta classification](#). *Proceedings Copyright*, 126:137.
- Laura Fieback, Jakob Spiegelberg, and Hanno Gottschalk. 2025b. [Metatoken: Detecting hallucination in image descriptions by meta classification](#). In *Proceedings of the 20th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, pages 126–137. SCITEPRESS - Science and Technology Publications.
- Anisha Gunjal, Jihan Yin, and Erhan Bas. 2024. [Detecting and preventing hallucinations in large vision language models](#). In *Proceedings of the AAI Conference on Artificial Intelligence*, volume 38, pages 18135–18143.
- Jinghan He, Kuan Zhu, Haiyun Guo, Junfeng Fang, Zhenglin Hua, Yuheng Jia, Ming Tang, Tat-Seng Chua, and Jinqiao Wang. 2025. [Cracking the code of hallucination in llms with vision-aware head divergence](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 3488–3501. Association for Computational Linguistics.
- Dan Hendrycks and Kevin Gimpel. 2017. [A baseline for detecting misclassified and out-of-distribution examples in neural networks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming Zhang, and Nenghai Yu. 2024. [Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation](#). In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13418–13427. IEEE.
- Insu Jeon, Wonkwang Lee, Myeongjang Pyeon, and Gunhee Kim. 2025. [IB-GAN: disentangled representation learning with information bottleneck generative adversarial networks](#). *CoRR*, abs/2510.20165.

- Nick Jiang, Anish Kachinthaya, Suzanne Petryk, and Yossi Gandelsman. 2025a. [Interpreting and editing vision-language representations to mitigate hallucinations](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Zhangqi Jiang, Junkai Chen, Beier Zhu, Tingjin Luo, Yankun Shen, and Xu Yang. 2025b. [Devils in middle layers of large vision-language models: Interpreting, detecting and mitigating object hallucinations via attention lens](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, pages 25004–25014. Computer Vision Foundation / IEEE.
- Diederik P. Kingma and Max Welling. 2014. [Auto-encoding variational bayes](#). In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- Jusung Lee, Sungguk Cha, Younghyun Lee, and Cheoljong Yang. 2024. [Visual question answering instruction: Unlocking multimodal large language model to domain-specific visual multitasks](#). *CoRR*, abs/2402.08360.
- Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. 2024a. [Mitigating object hallucinations in large vision-language models through visual contrastive decoding](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 13872–13882. IEEE.
- Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. 2024b. [Mitigating object hallucinations in large vision-language models through visual contrastive decoding](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13872–13882.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. [Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models](#). In *International conference on machine learning*, pages 19730–19742. PMLR.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. 2023b. [BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 19730–19742. PMLR.
- Qing Li, Jiahui Geng, Chenyang Lyu, Derui Zhu, Maxim Panov, and Fakhri Karray. 2024. [Reference-free hallucination detection for large vision-language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4542–4551. Association for Computational Linguistics.
- Qintong Li, Zhiyong Wu, Lingpeng Kong, and Wei Bi. 2023c. [Explanation regeneration via information bottleneck](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 12081–12102. Association for Computational Linguistics.
- Xiang Lisa Li and Jason Eisner. 2019a. [Specializing word embeddings \(for parsing\) by information bottleneck](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics.
- Xiang Lisa Li and Jason Eisner. 2019b. [Specializing word embeddings \(for parsing\) by information bottleneck](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023d. [Evaluating object hallucination in large vision-language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 292–305.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. [Microsoft coco: Common objects in context](#). In *European conference on computer vision*, pages 740–755. Springer.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024a. [Improved baselines with visual instruction tuning](#). In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). *CoRR*, abs/2304.08485.
- Shi Liu, Kecheng Zheng, and Wei Chen. 2024b. [Paying more attention to image: A training-free method for alleviating hallucination in vlms](#). In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part LXXXIII*, volume 15141 of *Lecture Notes in Computer Science*, pages 125–140. Springer.
- Shi Liu, Kecheng Zheng, and Wei Chen. 2024c. [Paying more attention to image: A training-free method for alleviating hallucination in vlms](#). In *European Conference on Computer Vision*, pages 125–140. Springer.
- Rabeeh Karimi Mahabadi, Yonatan Belinkov, and James Henderson. 2021. [Variational information bottleneck for effective low-resource fine-tuning](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

- Vishnu Prabhakaran, Purav Aggarwal, Vinay Kumar Verma, Gokul Swamy, and Anoop Saladi. 2025. [VADE: Visual attention guided hallucination detection and elimination](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 14949–14965, Vienna, Austria. Association for Computational Linguistics.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. [Object hallucination in image captioning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Qingyun Sun, Jianxin Li, Hao Peng, Jia Wu, Xingcheng Fu, Cheng Ji, and Philip S. Yu. 2022. [Graph structure learning with variational information bottleneck](#). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 4165–4174. AAAI Press.
- Feilong Tang, Chengzhi Liu, Zhongxing Xu, Ming Hu, Zile Huang, Haochen Xue, Ziyang Chen, Zelin Peng, Zhiwei Yang, Sijin Zhou, Wenxue Li, Yulong Li, Wenxuan Song, Shiyang Su, Wei Feng, Jionglong Su, Mingquan Lin, Yifan Peng, Xuelian Cheng, and 2 others. 2025. [Seeing far and clearly: Mitigating hallucinations in mllms with attention causal decoding](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, pages 26147–26159. Computer Vision Foundation / IEEE.
- Naftali Tishby, Fernando C. N. Pereira, and William Bialek. 2000. [The information bottleneck method](#). *CoRR*, physics/0004057.
- Junyang Wang, Yuhang Wang, Guohai Xu, Jing Zhang, Yukai Gu, Haitao Jia, Jiaqi Wang, Haiyang Xu, Ming Yan, Ji Zhang, and 1 others. 2023. [Amber: An llm-free multi-dimensional benchmark for mllms hallucination evaluation](#). *arXiv preprint arXiv:2311.07397*.
- Tianyun Yang, Ziniu Li, Juan Cao, and Chang Xu. 2025. [Understanding and mitigating hallucination in large vision-language models via modular attribution and intervention](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, Chenliang Li, Yuanhong Xu, Hehong Chen, Junfeng Tian, Qian Qi, Ji Zhang, and Fei Huang. 2023. [mplug-owl: Modularization empowers large language models with multimodality](#). *CoRR*, abs/2304.14178.
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Tong Xu, Hao Wang, Dianbo Sui, Yunhang Shen, Ke Li, Xing Sun, and Enhong Chen. 2024. [Woodpecker: hallucination correction for multimodal large language models](#). *Sci. China Inf. Sci.*, 67(12).
- Jinrui Zhang, Teng Wang, Haigang Zhang, Ping Lu, and Feng Zheng. 2024. [Reflective instruction tuning: Mitigating hallucinations in large vision-language models](#). In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part LXVIII*, volume 15126 of *Lecture Notes in Computer Science*, pages 196–213. Springer.
- Xiaofeng Zhang, Yihao Quan, Chen Shen, Chaochen Gu, Xiaosong Yuan, Shaotian Yan, Jiawei Cao, Hao Cheng, Kaijie Wu, and Jieping Ye. 2025a. [Shallow focus, deep fixes: Enhancing shallow layers vision attention sinks to alleviate hallucination in vlms](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 3512–3534. Association for Computational Linguistics.
- Yudong Zhang, Ruobing Xie, Xingwu Sun, Yiqing Huang, Jiansheng Chen, Zhanhui Kang, Di Wang, and Yu Wang. 2025b. [Dhcp: Detecting hallucinations by cross-modal attention pattern in large vision-language models](#). In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 3555–3564.
- Ge Zheng, Jiaye Qian, Jiajin Tang, and Sibe Yang. 2025. [Why vlms are more prone to hallucinations in longer responses: The role of context](#). *CoRR*, abs/2510.20229.
- Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. 2024. [Analyzing and mitigating object hallucination in large vision-language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2024. [Minigt-4: Enhancing vision-language understanding with advanced large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2024a. [Minigt-4: Enhancing vision-language understanding with advanced large language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Kun Zhu, Xiaocheng Feng, Xiyuan Du, Yuxuan Gu, Weijiang Yu, Haotian Wang, Qianglong Chen, Zheng Chu, Jingchang Chen, and Bing Qin. 2024b. [An information bottleneck perspective for effective noise filtering on retrieval-augmented generation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1044–1069. Association for Computational Linguistics.

Lanyun Zhu, Deyi Ji, Tianrun Chen, Peng Xu, Jieping Ye, and Jun Liu. 2025. **IBD: alleviating hallucinations in large vision-language models via image-biased decoding**. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2025, Nashville, TN, USA, June 11-15, 2025*, pages 1624–1633. Computer Vision Foundation / IEEE.

Geigh Zollicoffer, Minh Vu, and Manish Bhattarai. 2025. **Mtre: Multi-token reliability estimation for hallucination detection in vlms**. *arXiv preprint*.

A Models and Baselines

A.1 Vision Language Models

MiniGPT-4 MiniGPT-4 citezhuminigpt connects visual and textual modalities using a single linear projection layer. It utilizes a frozen BLIP-2 (Li et al., 2023a) visual encoder, which consists of ViT-G/14 (EVA-CLIP) and a Q-Former. The language backbone is Vicuna-7B (based on LLaMA-1), comprising 32 transformer layers and 32 attention heads.

LLaVA-v1.5-7B LLaVA-v1.5 (Liu et al., 2023) employs a two-layer MLP projector to align visual features with the language model. Its visual encoder is CLIP-ViT-L-336px. The language backbone is Vicuna-7B-v1.5 (based on Llama-2), which contains 32 layers and 32 attention heads.

LLaVA-v1.6-Mistral-7B LLaVA-v1.6 (LLaVA-NeXT) (Liu et al., 2024a) introduces an "AnyRes" technique that splits high-resolution images into grids to overcome resolution limits, while still using the CLIP-ViT-L-336px visual encoder. The backbone is Mistral-7B-Instruct-v0.2, featuring 32 layers and 32 attention heads.

Qwen2.5-VL-7B-Instruct Qwen2.5-VL (Bai et al., 2025) utilizes Naive Dynamic Resolution and M-RoPE to handle variable image sizes naturally without fixed patching. It uses a customized SigLIP-based visual encoder (approx. 600M params) with a C-Abstractor for feature compression. The backbone is Qwen2.5-7B, consisting of 28 layers and 28 attention heads.

A.2 Hallucination Detection Baselines

AvgProb Given a generated sentence (or sequence) indexed by i with J_i tokens, let p_{ij} denote the model-assigned conditional probability of the *actually generated* token at position j . AvgProb quantifies sentence-level uncertainty by the mean

negative log-probability over all positions:

$$\text{AvgProb}(i) = -\frac{1}{J_i} \sum_{j=1}^{J_i} \log p_{ij}.$$

A larger AvgProb(i) indicates that the model tends to assign lower likelihood to the produced tokens, reflecting higher uncertainty for the whole sentence.

AvgEnt AvgEnt computes uncertainty using the full predictive distribution at each position. Let $\mathbf{p}_{ij}(\cdot)$ be the predicted distribution over the vocabulary \mathcal{V} at position j in sentence i , and define the token-level predictive entropy as

$$H_{ij} = -\sum_{v \in \mathcal{V}} \mathbf{p}_{ij}(v) \log \mathbf{p}_{ij}(v).$$

We then aggregate token entropies into a sentence-level score via averaging:

$$\text{AvgEnt}(i) = \frac{1}{J_i} \sum_{j=1}^{J_i} H_{ij}.$$

Higher AvgEnt(i) suggests more diffuse (less confident) predictive distributions across tokens, hence greater sentence-level uncertainty.

RepProbing RepProbing includes a lightweight classifier trained on the VLM decoder’s last-layer hidden states to estimate hallucination risk. Let $z_t^L \in \mathbb{R}^d$ be the hidden state at token position t from the top decoder layer L . The probe outputs a hallucination score (or probability) as

$$\hat{y}_t^h = f_\theta(z_t^L), \quad (15)$$

where f_θ is typically a linear head or a shallow MLP.

A.3 Hallucination Mitigation Baselines

BeamSearch Beam search is a deterministic decoding strategy that approximates the most likely output sequence by maintaining the top- B partial hypotheses (“beams”) at each step. Starting from the prompt, it repeatedly expands each beam with candidate next tokens and keeps only the B sequences with the highest cumulative log-probability (often with length normalization), continuing until an end-of-sequence token is produced.

PAI PAI (Liu et al., 2024c) is a training-free method that mitigates text inertia in LVLMs—when the LLM dominates so outputs rely more on text context than visual evidence. It boosts attention to image tokens and subtracts text-only logits from multimodal logits to suppress language-only bias, encouraging stronger visual grounding and reducing hallucinations.

VCD VCD (Visual Contrastive Decoding) (Leng et al., 2024b) is a simple, training-free decoding method that contrasts the output distributions produced from an original image and a distorted version of the same image. By using this contrast to suppress statistical biases and unimodal language priors, it encourages stronger visual grounding, substantially reducing object hallucinations across LVLM families while also performing well on general LVLM benchmarks.

B Implementation Details

B.1 Hallucination Detection

In the hallucination detection experiments, for the discriminative benchmarks POPE and AMBER, we follow the work of (Li et al., 2024) to extract images, questions, and ground truths (GT) from the original datasets. For each sample, we construct responses that either contain or do not contain hallucinations; specifically, for samples where the GT is “Yes”, we generate “Yes” (containing hallucination) and “No” (free from hallucination) responses.

For the POPE benchmark, we construct training and validation splits across its three subsets (popular, random, and adversarial) and report the average metrics over these subsets. For the AMBER benchmark, we conduct experiments using a curated subset of 5,000 samples. We manually partition the datasets to ensure that different samples associated with the same image do not overlap between the training and validation sets.

For the M-HalDetect benchmark, we further divide the official validation set into training and validation subsets using an 80 : 20 ratio and report span-based hallucination detection results. Regarding the COCO-Caption task, we employ the LLaVA-v1.5-7B model to generate responses for images from the COCO 2014 Val set. We annotate hallucinated objects in the responses using the official COCO 2014 Val annotations and report sentence-based hallucination detection results.

B.2 Model Architecture

Regarding the VIB-Probe encoder, we utilize a 3-layer MLP network with dimensions (1024, 512, 256) to reduce the dimensionality of the original attention output feature vectors, followed by processing with two residual blocks. For the decoder, we employ a simple single linear layer. Throughout the network, we apply the GELU activation function and LayerNorm.

B.3 Hallucination Mitigation

We evaluate object hallucinations in VLM’s generation with the CHAIR (Captioning Hallucination Assessment with Image Relevance) metrics, which compare model-generated captions against ground-truth object annotations to quantify objects mentioned in text but not present in the image. Specifically, CHAIR_i reports the proportion of hallucinated object mentions among all generated object mentions, while CHAIR_s reports the percentage of captions that contain at least one hallucinated object.

$$\text{CHAIR}_i = \frac{|\{\text{hallucinated objects}\}|}{|\{\text{all objects mentioned}\}|} \quad (16)$$

$$\text{CHAIR}_s = \frac{|\{\text{sentences with hallucinated objects}\}|}{|\{\text{all sentences}\}|} \quad (17)$$

In the hallucination mitigation experiments, we intervene on the attention heads that rank in the top 5% of head importance scores. The threshold for triggering this intervention is determined based on the average logit values from the training set used in the hallucination detection experiments. We set the suppression strength λ to 0.001.