

# Anatomical Region-Guided Contrastive Decoding: A Plug-and-Play Strategy for Mitigating Hallucinations in Medical VLMs

Xiao Liang,<sup>1</sup> Chenxi Liu,<sup>1</sup> Zhi Ma,<sup>\*1</sup> Di Wang,<sup>\*1</sup> Bin Jing,<sup>2</sup> Quan Wang,<sup>1</sup> Yuanyuan Shi<sup>3</sup>

<sup>1</sup>School of Computer Science and Technology, Xidian University, China

<sup>2</sup>School of Biomedical Engineering, Capital Medical University, China

<sup>3</sup>Department of Ophthalmology, the Ninth Medical Center of the Chinese PLA General Hospital, China  
ecoxial2012@outlook.com

## Abstract

Medical Vision-Language Models (MedVLMs) show immense promise in clinical applicability. However, their reliability is hindered by hallucinations, where models often fail to derive answers from visual evidence, instead relying on learned textual priors. Existing mitigation strategies for MedVLMs have distinct limitations: training-based methods rely on costly expert annotations, limiting scalability, while training-free interventions like contrastive decoding, though data-efficient, apply a global, untargeted correction whose effects in complex real-world clinical settings can be unreliable. To address these challenges, we introduce **Anatomical Region-Guided Contrastive Decoding (ARCD)**, a plug-and-play strategy that mitigates hallucinations by providing targeted, region-specific guidance. Our module leverages an anatomical mask to direct a three-tiered contrastive decoding process. By dynamically re-weighting at the token, attention, and logits levels, it verifiably steers the model's focus onto specified regions, reinforcing anatomical understanding and suppressing factually incorrect outputs. Extensive experiments across diverse datasets, including chest X-ray, CT, brain MRI, and ocular ultrasound, demonstrate our method's effectiveness in improving regional understanding, reducing hallucinations, and enhancing overall diagnostic accuracy.

**Code** — <https://github.com/ecoxial2007/ARegionCD>

## Introduction

In recent years, Medical Vision-Language Models (MedVLMs) have demonstrated impressive capabilities across diverse modalities such as chest radiography (Wu et al. 2023), pathology (Seyfioglu et al. 2023), and dermatology (Zeng et al. 2025), showing great promise in advancing medical intelligence for tasks like automated reporting and visual question answering. However, the clinical applicability of these models is critically undermined by their propensity for hallucination (Bai et al. 2024)—generating plausible yet factually incorrect statements that contradict the visual evidence. This unreliability stems from a fundamental lack of visual grounding, where models often prioritize statistical biases and unimodal priors over actual image content (Leng et al. 2023). For instance, a model may misidentify *ECG leads*

\*Corresponding authors.

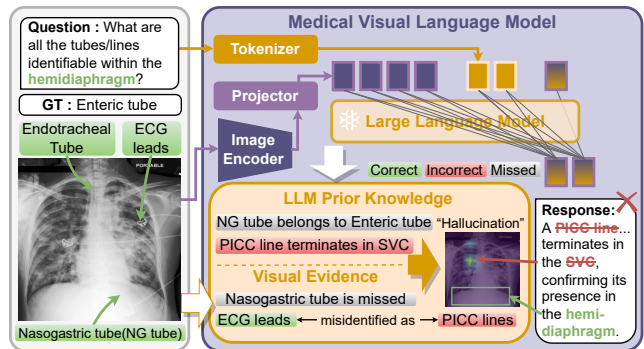


Figure 1: An example of hallucination driven by a statistical bias. The model misidentifies visually apparent *ECG leads* as a *PICC line* because the latter is far more common in training corpora reports. This flawed prior-visual association leads to a factually incorrect response and demonstrates a critical failure in visual grounding.

as a *PICC line*, as shown in Figure 1. This confusion often arises because textual reports in its training data frequently describe *PICC lines* but rarely mention the visually similar *ECG leads*, creating a strong statistical prior that overrides the visual evidence. The opaque nature of this failure mechanism raises a critical question for clinical trust: *How can we ensure a VLM's attention is verifiably focused on the specific, diagnostically relevant regions of an image, mirroring the analytical process of a physician?*

To mitigate hallucinations and enhance the visual understanding of VLMs, researchers have pursued two main avenues. The first involves training-based approaches. For instance, methods like MMedPO (Zhu et al. 2024) leverage preference optimization to compel the model to ground its responses in visual evidence rather than relying on textual priors. The second avenue is training-free inference-time intervention. Techniques such as visual contrastive decoding (Leng et al. 2023) work by contrasting outputs from original and distorted images to penalize the model's over-reliance on learned data biases. While these approaches are valuable, they present a difficult trade-off. Training-based methods are constrained by the prohibitively high cost of annotating medical data, which limits their broad applicability. On

the other hand, training-free methods, though data-efficient, apply a global, untargeted correction whose effects on clinically complex medical data can be unreliable.

To address these challenges, we propose **Anatomical Region-Guided Contrastive Decoding (ARCD)**, a plug-and-play decoding strategy that provides fine-grained, multi-level guidance on the model’s visual understanding, enhancing its reliability on clinically complex data without requiring additional training. We first build upon a strong MedVLM baseline established by fine-tuning Phi-3.5V on high-quality medical visual instruction data. The proposed ARCD strategy then performs **Dynamic Attention Mask Generation** to convert the segmentation annotation of a given anatomical region into a token-level attentional mask, which in turn directs our **Mask-Guided Conditional Token Weighting** module. This module steers the decoding process using a three-tiered strategy: it applies dynamic, contrastive re-weighting between guided and unguided generation branches sequentially at the token, attention, and logits levels, which ultimately reinforces anatomical region-specific understanding at inference time, thereby mitigating hallucinations. Our primary contributions are:

- We propose **Anatomical Region-Guided Contrastive Decoding**, a novel decoding strategy that improves visual grounding by guiding language generation with a three-tiered, contrastive mechanism based on specified anatomical regions.
- The proposed **ARCD** method is entirely training-free and plug-and-play, allowing for seamless integration with various VLMs and segmentation models without requiring any parameter updates.
- We conduct extensive experiments across four diverse medical modalities (chest X-ray, CT, brain MRI, and ocular B-ultrasound) to validate the effectiveness and generalization of our method.

## Related work

### Medical Visual Language Models

Recent advancements in Medical Vision-Language Models (Med-VLMs) have primarily focused on adapting general-domain models to the specialized needs of healthcare. Foundational works established this paradigm, with LLaVA-Med using instruction-tuning to align biomedical visual features with language embeddings for cost-effective adaptation (Li et al. 2023a), while HuatuoVision leveraged massive, high-quality VQA datasets to further enhance performance (Chen et al. 2024a). Following this, a trend towards domain specialization emerged. This includes models purpose-built for specific modalities, such as RadFM (Wu et al. 2023) for chest radiography with 2D/3D capabilities, Quilt-LLaVA (Seyfioglu et al. 2023) for pathology, which creates spatially-grounded data by tracking narrator cursors in educational videos, and MMSkin (Zeng et al. 2025) for dermatology, which uses high-fidelity image-text pairs from professional textbooks. To enhance reliability and mitigate hallucinations, advanced techniques have emerged. These range from using preference optimization for clinical

accuracy (Zhu et al. 2024) to integrating weak region-of-interest annotations for better localization (Chen et al. 2024c). Concurrently, other methods like rationale-based explanations (Gai et al. 2024), latent prompts (Gu et al. 2024), feature fusion (Ha et al. 2024), and self-training pipelines (Sun et al. 2024) have collectively improved model interpretability and robustness. However, despite progress in textual accuracy, a key challenge remains in fine-grained visual grounding, where models often fail to explicitly link generated statements to specific, diagnostically relevant image regions (Chen et al. 2024c; Wu, Kim, and Wu 2024).

### Hallucination-Mitigated VLMs

Object hallucination in general Vision-Language Models (VLMs), where generated text misaligns with visual facts, is a well-studied issue (Liu et al. 2024a; Lan et al. 2024). Its causes include knowledge scarcity from limited data (Chen et al. 2024b), an over-reliance on statistical and language priors (Favero et al. 2024; Leng et al. 2023), and perceptual failures from poor image quality (Liu et al. 2024b). Mitigation strategies are either training-based or training-free. Training-based methods aim for intrinsic correction by enriching data (e.g., ShareGPT4V (Chen et al. 2023)), using preference optimization like DPO to favor factual responses (Zhao 2023; Yang et al. 2025), fusing visual reasoning (Park et al. 2025), or applying adversarial training (Chen et al. 2025). Training-free methods intervene at inference, using techniques like Visual and Instructional Contrastive Decoding (VCD (Leng et al. 2023), ICD (Wang et al. 2024)), retrieval-augmented generation (Feng et al. 2024), or post-processing correction frameworks such as Woodpecker (Yin et al. 2024). Despite their effectiveness in general domain, applying these methods to the medical domain presents unique challenges (Yan et al. 2025). Training-based approaches are often hindered by the prohibitive cost and expertise required for medical data annotation. Meanwhile, training-free methods may lack the high degree of interpretability and the rigorous, verifiable accuracy essential for high-stakes clinical applications.

## Methodology

Our proposed Anatomical Region-Guided Contrastive Decoding is a plug-and-play strategy that builds upon a strong medically-adapted VLM baseline (Phi3.5V-Med). This module first generates a token-level mask from a spatial annotation that delineates a specific anatomical region (*Dynamic Attention Mask Generation*) and subsequently guides the contrastive decoding process via a three-tiered strategy (*Mask-Guided Conditional Token Weighting*). An overview of our method is illustrated in Figure 2.

### Task Formulation and Baseline

A Vision-Language Model (VLM) aims to generate a coherent and contextually relevant text sequence  $\mathcal{Y} = \{y_1, y_2, \dots, y_T\}$  in response to a given image  $\mathcal{V}$  and a textual prompt or question  $\mathcal{Q}$ . The generation process is typically auto-regressive, where the objective is to maximize the

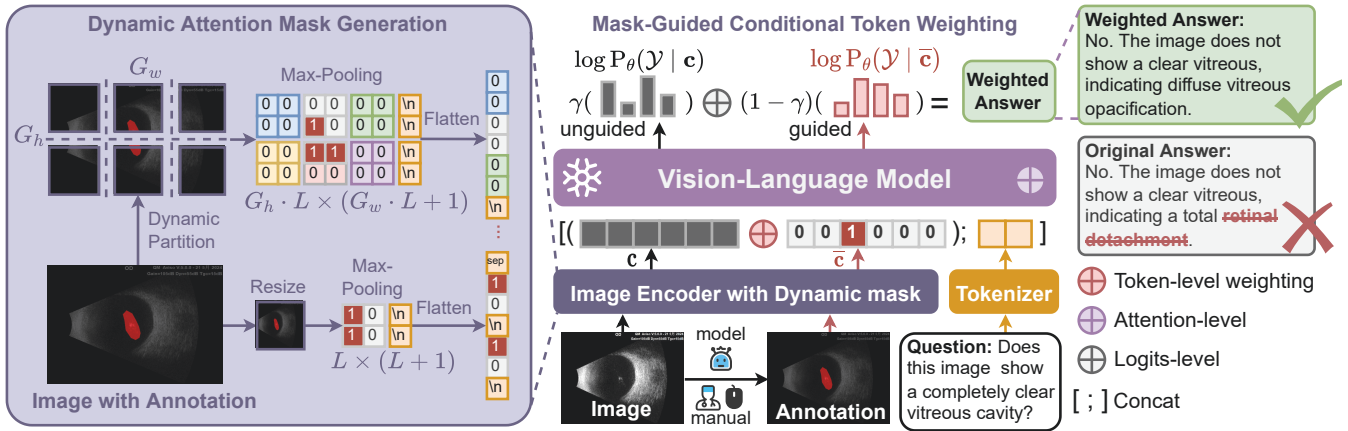


Figure 2: Overview of our proposed Anatomical Region-Guided Contrastive Decoding strategy. **Left:** *Dynamic Attention Mask Generation* module converts a specified anatomical region (e.g., a segmentation annotation) into a multi-scale token-level mask. **Right:** *Mask-Guided Conditional Token Weighting* module then uses this mask to steer the decoding process via a strategy that applies contrastive re-weighting at the token level, attention level, and logits level, ensuring the generated answer is grounded in the specified visual region.

guided log-probability of the target sequence. This is formally expressed as:

$$\log P_{\theta}(\mathcal{Y}|\mathcal{V}, \mathcal{Q}) = \sum_{t=1}^T \log P_{\theta}(y_t|\mathcal{V}, \mathcal{Q}, y_{<t}), \quad (1)$$

where  $y_t$  is the token at timestep  $t$ , and  $y_{<t}$  represents all previously generated tokens. While powerful, lightweight models like Phi-3.5 Vision (Abdin et al. 2024) possess strong general visual understanding capabilities, they lack the specialized domain knowledge required for medical applications. Consequently, they struggle to answer medical questions accurately and often fail to identify specific pathologies, even when they are visually prominent. To create a robust baseline, we must inject medical expertise into the model. We perform parameter-efficient fine-tuning on Phi-3.5 Vision using Low-Rank Adaptation (LoRA) (Hu et al. 2021) with the large-scale PubMedVision (Chen et al. 2024a) dataset. This process enriches the model with essential medical knowledge and aligns its outputs with clinical context. The resulting model, which we term Phi3.5V-Med, serves as our baseline for all subsequent experiments.

### Dynamic Attention Mask Generation

Despite infusing medical knowledge into VLMs through visual instruction fine-tuning, the data-driven nature of Supervised Fine-Tuning (SFT) makes it challenging to guarantee that the VLM’s responses are genuinely grounded in the provided visual evidence, as highlighted in Section 1. To address this, we propose our Anatomical Region-Guided Contrastive Decoding strategy, aiming to explicitly direct the model’s focus towards clinically relevant anatomical regions. Specifically, an expert or a pre-trained segmentation model (e.g., PSPNet or MedSAM) first annotates or generates a segmentation mask  $\mathcal{S}$  of the same spatial dimensions as the input image  $\mathcal{V} \in \mathbb{R}^{H \times W \times C}$ . This mask  $\mathcal{S} \in \{0, 1\}^{H \times W}$  defines the anatomical region of interest,

with pixel values of 1 indicating the relevant region. Subsequently, this region-specific information must be effectively injected into the MedVLM’s attention mechanism by transforming it into a token-level mask that aligns with the image token embeddings produced by the VLM’s vision encoder.

To achieve this, we introduce **Dynamic Attention Mask Generation**, a module that creates attention masks structurally aligned with the VLM’s visual tokens. This approach allows for handling dynamic input resolutions and capturing fine-grained details by generating two distinct masks: a global mask  $\mathcal{M}_g$  and a local mask  $\mathcal{M}_l$ . Specifically, the ViT-L/14 vision encoder from Phi-3.5V processes a  $336 \times 336$  input image by dividing it into a  $24 \times 24$  grid of  $14 \times 14$  pixel patches. This grid is subsequently reshaped into a  $12 \times 12$  feature grid, and we denote this dimension as  $L = 12$ . The global mask  $\mathcal{M}_g$  is thus derived by downsampling the high-resolution mask  $\mathcal{S}$  to a feature-level grid of size  $L \times L$ . A newline separator  $m_{sep}$  (mask value 0) is then appended to each of its  $L$  rows, resulting in a flattened one-dimensional mask  $\mathcal{M}_g$  of length  $L \times (L + 1)$ . Similarly, the local mask  $\mathcal{M}_l$  is generated by downsampling  $\mathcal{S}$  to a larger composite grid of size  $(G_h \cdot L) \times (G_w \cdot L + 1)$ . Here, the user-defined grid dimension  $G = (G_h, G_w)$  determines the granularity of the local analysis. A larger  $G$  partitions the image into a finer-grained  $G_h \times G_w$  grid of local views, allowing the model to focus on smaller details. Subsequently, a newline separator is appended to each of the  $G_h \cdot L$  rows of this composite grid before it is flattened into the final sequence  $\mathcal{M}_l$ . Finally, the complete mask  $\mathcal{M}$  is assembled by concatenating the local mask, a single separator token  $m_{sep} = (0)$ , and the global mask:  $\mathcal{M} = [\mathcal{M}_l; m_{sep}; \mathcal{M}_g] \in \mathbb{R}^N$ . This produces a composite and structured attention mask with a total length  $N$ :

$$N = G_h \cdot L \cdot (G_w \cdot L + 1) + 1 + L \cdot (L + 1). \quad (2)$$

This mask  $\mathcal{M}$  then takes effect during the model’s self-attention computations, thereby prioritizing information

from the specified anatomical regions during response generation.

### Mask-Guided Conditional Token Weighting

While VCD (Leng et al. 2023) mitigates language priors by merging an “*original*” and a noise-augmented “*distorted*” branch at the logits level, its original branch still contains redundant visual information that can lead to hallucinations. Additionally, this post-hoc intervention at the logits level fails to correct the decoder’s erroneous focus on irrelevant visual regions during the attention computation stage.

To address these limitations, we extend VCD with a multi-level guidance mechanism, fusing a “*original (unguided)*” branch  $\bar{c}$  and a “*guided*” branch  $c$  via weighted fusion at the token, attention, and logits levels. The conditional guidance for the branch  $c$  is provided by our previously generated visual token mask  $\mathcal{M} = \{m_1, m_2, \dots\}$ , where the binary indicator  $m_i \in \{0, 1\}$  is set to 1 for tokens corresponding to regions requiring enhanced focus. First, the input image  $\mathcal{V}$  and question  $\mathcal{Q}$  are processed by a tokenizer and encoder  $f(\cdot)$  to produce token embeddings. For clarity, we denote the input tokens collectively as  $\mathbf{x}$  and the resulting embeddings as  $\mathbf{c} = f(\mathbf{x})$ . At the token-level, a small weight  $\alpha$  (e.g., 0.01) is used to suppress the embeddings of the guided tokens within the unguided branch  $\bar{c}$ :

$$\bar{c}_i = \begin{cases} \alpha \cdot c_i & \text{if } m_i = 1 \\ c_i & \text{if } m_i = 0 \end{cases} \quad (3)$$

Subsequently, at the attention-level, a weight  $\beta$  (e.g., 3, where a larger value implies stronger guidance) is applied to amplify the attention probability  $p_i$  corresponding to the visual regions of interest:

$$p_i = \frac{\beta^{m_i} \cdot \exp(e_i)}{\sum_{j=1}^N \beta^{m_j} \cdot \exp(e_j)}. \quad (4)$$

Here,  $e_i$  represents the  $i$ -th element of the pre-softmax query-key attention score matrix within a single attention head, and  $N$  is the number of tokens in the input sequence. This step is designed to counteract the tendency for textual information to dominate, ensuring that visual cues are not diminished before being processed by the Large Language Model. Finally, at the logits-level, the parameter  $\gamma$  (e.g., 1.5, where a larger value implies stronger guidance) controls the intensity of the guidance towards the guided branch:

$$\mathbb{P}_{\mathcal{Y}} = (1 - \gamma) \log P_{\theta}(\mathcal{Y}|\bar{c}) + \gamma \log P_{\theta}(\mathcal{Y}|c). \quad (5)$$

Here,  $\mathbb{P}_{\mathcal{Y}}$  represents the final, guided log-probability distribution for the next token, from which the guided response  $\mathcal{Y}$  is generated auto-regressively at each decoding step. Through this three-tiered weighted fusion, we achieve a comprehensive and controllable mechanism for anatomical region-guided generation. Beyond steering focus towards visual regions, this framework can also be adapted for textual alignment by modifying the mask  $\mathcal{M}$ , ensuring that generated content more closely adheres to specific topics (Zhang et al. 2023). The optimal hyperparameter settings are detailed in the implementation section.

## Experiments

### Datasets

For our experimental setup, we leveraged **PubMedVision** (Chen et al. 2024a), a large dataset containing 1.3 million MedVQA instances for visual instruction tuning. We then used this to fine-tune the Phi-3.5V model, which served as our zero-shot baseline for subsequent evaluations. Our downstream tasks encompassed three distinct types of MedVQA challenges. These included **MIMIC-Ext-VQA** (Bae et al. 2023) for chest X-rays and **SLAKE** (Liu et al. 2021), a comprehensive radiology dataset covering abdominal CT, chest X-rays, and brain MRI. Additionally, we incorporated a specialized Ocular B-ultrasound dataset, named **OBScan**, specifically constructed by us. OBScan is composed of ocular B-scan ultrasound images from patients with various eye diseases, each paired with an expert-authored diagnostic report. These reports are structured into three sections: a description of the left eye, a description of the right eye, and a final impression. To provide granular grounding for the diagnoses, experts manually segmented key anatomical and pathological regions in each image, including the retina, lens, vitreous opacity, and vitreous detachment. Following the methodology of MIMIC-Ext-VQA, we then formulated a diverse set of QA pairs featuring three question formats (*choose*, *verify*, and *query*) and two answer types (*closed* and *open-ended*) for our experiments. The fine-tuning baseline underwent separate fine-tuning processes on each of these three diverse downstream datasets. Further details for all datasets are provided in the Appendix.

### Implementation Details

Our experiments were conducted using 24GB NVIDIA 3090 GPUs, leveraging bfloat16 precision for all training processes. For visual instruction fine-tuning, we initialized the model with pre-trained weights from Phi-3.5-Vision. We employed the AdamW optimizer with a cosine learning rate scheduler, setting the learning rate to  $2e-4$ , the batch size to 256, and training for 1 epoch. For downstream task fine-tuning, we adapted the learning rate to  $1e-4$ , reduced the batch size to 64, and trained for 10 epochs. Only the LoRA layers (rank 64) were updated throughout both phases.

During inference, our proposed Anatomical Region-Guided Contrastive Decoding is employed as a decoding strategy for both our zero-shot and fine-tuned Phi-3.5V-Med baselines. The core of this strategy is to modulate the influence of anatomical region-guided embeddings relative to normal embeddings at three key stages of generation. This is achieved using parameters  $\alpha$  (in Eq. 3),  $\beta$  (in Eq. 4), and  $\gamma$  (in Eq. 5), which control the weighting at the token-level, attention-level, and logits-level, respectively. By adjusting these weights, we can precisely steer the model’s final output. A detailed ablation study on the effects of these parameters is discussed in the next section. Following the methodology of LLaVA-Med (Li et al. 2023b), we assessed performance using accuracy for closed-set questions and recall, defined as the ratio of ground-truth tokens present in the generated response, for open-set questions.

Method	MIMIC-Ext-VQA*			SLAKE			OBScan		
	Open	Closed	Overall	Open	Closed	Overall	Open	Closed	Overall
<i>General Visual-Language Model</i>									
LLaVA-1.6-7B	12.70	57.07	46.06	32.28	53.54	42.91	47.78	62.80	57.48
Qwen-VL-7B	15.87	56.54	46.46	39.37	61.42	50.39	48.89	71.95	63.78
Phi3.5V-4.2B	17.46	56.02	46.46	33.07	53.54	43.31	40.00	56.10	50.39
GPT-4o	28.57	74.87	63.39	38.58	74.80	56.69	72.22	70.73	71.26
<i>Medical Visual-Language Model (Zero-Shot)</i>									
LLaVA-Med-7B	14.29	55.50	45.28	37.80	55.12	46.46	52.22	60.37	57.48
Huatuov-34B	17.46	65.45	53.54	41.73	74.02	57.87	60.00	67.68	64.96
Phi3.5V-Med	14.29	58.12	47.24	38.58	61.42	50.00	54.44	65.85	61.81
w/ VCD	15.87 (+1.58)	59.16 (+1.04)	48.43 (+1.19)	40.16 (+1.58)	61.42 (+0.00)	50.79 (+0.79)	<b>61.11 (+6.67)</b>	64.02 (-1.83)	62.99 (+1.18)
w/ DoLA	<b>19.05 (+4.76)</b>	56.54 (-1.58)	47.24 (+0.00)	<u>47.24 (+8.66)</u>	59.84 (-1.58)	<u>53.54 (+3.54)</u>	57.78 (+3.34)	62.80 (-3.05)	61.02 (-0.79)
w/ OPERA	14.29 (+0.00)	<u>61.78 (+3.66)</u>	<u>50.00 (+2.76)</u>	36.22 (-2.36)	<b>65.35 (+3.93)</b>	50.79 (+0.79)	<u>60.00 (+5.56)</u>	64.02 (-1.83)	62.06 (+0.25)
w/ ARCD	14.29 (+0.00)	<b>62.83 (+4.71)</b>	<b>50.79 (+3.55)</b>	<b>48.82 (+10.24)</b>	<u>61.42 (+0.00)</u>	<b>55.11 (+5.11)</b>	<b>61.11 (+6.67)</b>	<b>68.29 (+2.44)</b>	<b>65.75 (+3.94)</b>
<i>Medical Visual-Language Model (Fine-Tuning)</i>									
Phi3.5V-Med	28.57	82.72	69.29	76.38	88.19	82.28	74.44	98.17	89.76
w/ VCD	33.33 (+4.76)	81.68 (-1.04)	69.69 (+0.40)	74.02 (-2.36)	<b>88.19 (+0.00)</b>	81.10 (-1.18)	<u>75.56 (+1.12)</u>	97.56 (-0.61)	89.76 (+0.00)
w/ DoLA	<u>47.62 (+19.05)</u>	80.10 (-2.62)	<u>72.05 (+2.76)</u>	<u>77.17 (+0.79)</u>	86.61 (-1.58)	81.89 (-0.39)	73.33 (-1.11)	98.17 (+0.00)	89.37 (-0.39)
w/ OPERA	<b>50.79 (+22.22)</b>	78.53 (-4.19)	71.65 (+2.36)	<u>77.17 (+0.79)</u>	<u>87.40 (-0.79)</u>	<u>82.28 (+0.00)</u>	65.56 (-8.88)	<b>99.39 (+1.22)</b>	87.40 (-2.36)
w/ ARCD	42.86 (+14.29)	<b>89.53 (+6.81)</b>	<b>77.95 (+8.66)</b>	<b>77.95 (+1.57)</b>	<b>88.19 (+0.00)</b>	<b>83.07 (+0.79)</b>	<b>81.11 (+6.67)</b>	98.17 (+0.00)	<b>92.13 (+2.37)</b>

Table 1: Performance comparison of MedVQA methods on three datasets. The proposed method (w/ARCD) provides attentional guidance via a segmentation mask and is evaluated against strong baselines and decoding strategies. \*On the MIMIC dataset, evaluation is restricted to queries with an identifiable bounding box, where the mask is generated by filling the specified region.

## Baselines

For our experiments, we use greedy decoding as the standard methodology to assess the VLMs’ inherent performance. We then apply and compare various advanced decoding strategies, including our proposed ARCD, to this baseline. Our baseline models are two variants of Phi3.5V. The first, named *Phi3.5V-Med Zero-shot*, is the model fine-tuned on the general medical vision-language dataset, PubMedVision. The second, *Phi3.5V-Med Fine-Tuning*, is the model further fine-tuned specifically on each downstream task dataset. For a comprehensive performance benchmark, we compare our models against several state-of-the-art methods, which are categorized as follows:

- **General VLMs:** We select representative models from the general domain, including LLaVA-1.6 (Li et al. 2023a), Qwen-VL (Bai et al. 2023), Phi-3.5V (Abdin et al. 2024), and GPT-4o (OpenAI 2023).
- **Medical VLMs:** We include models specifically tailored for the medical domain, such as LLaVA-Med-7B (Li et al. 2023b) and Huatuov-34B (Chen et al. 2024a).
- **Decoding Strategies:** We apply several plug-and-play decoding strategies to our baseline, including VCD (Leng et al. 2023), DoLa (Chuang et al. 2023), and OPERA (Huang et al. 2023).

## Results

### Comparisons with Baselines

As detailed in Table 1, our proposed ARCD strategy, which uses anatomical region masks for guidance, consistently demonstrates accuracy improvements across all three MedVQA datasets. In the zero-shot scenario, our approach achieves pronounced accuracy gains on SLAKE (+5.11%) and OBScan (+3.94%). The improvement on the more fine-grained MIMIC dataset is more modest (+3.55%). This suggests that while the anatomical mask provides correct spatial

guidance, the zero-shot model’s performance is ultimately constrained by its lack of specialized diagnostic knowledge. In contrast, the MIMIC fine-tuned model, which already possesses this knowledge, fully leverages the anatomical guidance for a substantial +8.66% improvement. Furthermore, our technique consistently outperforms other decoding strategies like VCD and DoLa. We attribute this to the strong anatomical prior introduced by the segmentation mask, which effectively constrains the model’s attention to clinically relevant anatomical structures and reduces hallucinations.

### Ablation Study

In this section, we conduct ablation studies to: 1) Analyze the causes of hallucinations in MedVLMs; 2) Investigate the impact of different guidance information and parameters; 3) Evaluate the influence of alternative visual prompts.

**MedVLM Hallucination Evaluation.** While prior work in the general domain has analyzed and investigated the sources of VLM hallucinations, a unified evaluation standard tailored for the medical domain is lacking. To address this, we qualitatively analyze model outputs, which GPT-4o classifies into five categories: *Correct*, *Correct but Overly General*, *Correct with Detail Nuances*, *Irrelevant/Refusal*, and *Hallucination/Factual Contradiction*, as shown in Figure 3. Our results indicate that fine-tuning on both general (*Phi3.5V-Med-ZS*) and task-specific (*Phi3.5V-Med-FT*) medical data significantly reduces hallucinations and refusals. Notably, our proposed w/ARCD method consistently increases the proportion of fully *Correct* responses over its corresponding baseline in every setting. However, the complexity of the MIMIC dataset, where reports often detail numerous co-occurring diseases, can lead the fine-tuned model to generate more plausible but unverified information. Consequently, such outputs must be subjected to expert review for clinical validation.




Figure 3: GPT-4o evaluation of model responses under different settings on 250 samples uniformly drawn from three datasets. *Phi-3.5V-ZS* and *Phi-3.5V-Med-ZS* represent the zero-shot results for the base model and the model adapted with PubMedVision. *Phi-3.5V-Med-FT* is the model fine-tuned on the three MedVQA datasets, while *w/ARCD* denotes our proposed method with attentional masking.

Method	MIMIC-Ext-VQA	SLAKE*	OBScan
Phi3.5V-Med-ZS	47.24	50.00	61.81
w/ Label Prompt	49.60	51.97	53.14
w/ Bbox Prompt	<b>50.79</b>	<b>55.11</b>	64.96
w/ Mask Prompt	-	<b>55.11</b>	<b>65.74</b>
w/ Label+Bbox	<b>50.79</b>	54.33	59.44
w/ Label+Mask	-	54.33	62.59

Table 2: Prompt type ablation of Phi3.5V-Med-ZS on medical VQA benchmarks. In the SLAKE\*, only a portion of the images have mask annotations. **Bolded** results are the default setting.




Figure 4: Parameter ablation of  $\beta$  and  $\gamma$  using the Phi-3.5V-Med zero-shot model, with  $\alpha = 0.01$  fixed.

**Impact of Guidance Information.** While our primary approach uses region mask guidance, other information sources like bounding box or class labels could also influence performance. We thus ablate these components and their combinations, with results in Table 2. Across all datasets, mask-based guidance provides the most substantial performance boost. Bbox guidance, though less precise, also yields consistent gains, confirming that strengthening the model’s focus on specific spatial regions helps mitigate hallucinations. For the MIMIC and SLAKE datasets, all guidance components improved performance. However, on OBScan, using only the class label was detrimental (-8.67%). We hypothesize this is because our labels are often clinical abbreviations (e.g., “VD” as “Vitreous Hemorrhage”), which may introduce ambiguity or conflict with the visual evidence without precise spatial grounding.

**Impact of Control Parameters.** We present a detailed




Figure 5: Ablation Study on Visual Prompting: Impact of Visual ROI and Attention on MedVLM Performance.

ablation study in Figure 4 on the control parameters  $\alpha$  (in Eq. 3),  $\beta$  (in Eq. 4), and  $\gamma$  (in Eq. 5). Setting  $\alpha = 1$  disables our ARCD module, meaning the mask  $m$  does not affect the input token embedding  $c$ . Setting  $\beta = 1$  or  $\gamma = 1$  disables the corresponding attention-level and logits-level guidance, respectively. For attention-level guidance, performance peaks when  $\beta$  is 3 or 5 across datasets, suggesting an optimal  $\beta$  balances enhanced focus with preserving broader context. For logits-level guidance,  $\gamma = 1.3$  is consistently optimal or near-optimal. Since model performance is insensitive to  $\alpha$  within the  $[0, 0.1]$  range, we defer detailed results to the appendix. These findings validate that our balanced, multi-level guidance is superior to single-level interventions.

**Impact of Visual Prompting.** To further investigate the effectiveness of our proposed method, we conduct an ablation study on various forms of input visual information, as shown in Figure 5. As the results show, simple cropping or masking methods are clearly outperformed by our carefully designed ARCD, although performance varies slightly across different datasets. On the MIMIC dataset, a simple center crop yields a notable improvement, while other methods show only marginal gains or even a decrease in performance. For the SLAKE dataset, all visual prompting strategies provide a consistent and modest performance boost over the baseline. On the OBS dataset, methods guiding attention to the ROI, particularly using a bounding box, achieve the best results; interestingly, any cropping method harms

Correct	Overly General	Detail Nuances	Irrelevant	Factual Contradiction
				
<b>SLAKE Abdominal CT</b>				
				
<b>SLAKE - Head MRI</b>				
				
<b>MIMIC - Chest X-Ray</b>				
				
<b>Ocular B-scan Ultrasound</b>				
<b>Question:</b> What diseases are included in the picture? <b>Answer:</b> Liver cancer				
<b>Phi-3.5V-Med-ZS:</b> The image shows multiple lesions in the liver, which are likely metastatic lesions from a primary tumor elsewhere in the body. The reference information indicates that the patient has a history of a primary tumor, and the CT scan is being used to evaluate the extent of the disease and plan appropriate treatment.				
<b>Phi-3.5V-Med-ZS w/ ARCD:</b> The image shows liver findings suggestive of multiple diseases, potentially including cirrhosis, hepatocellular carcinoma, and metastatic lesions. The liver appears to have a nodular, irregular surface, indicating the presence of cirrhosis. Additionally, there are multiple hypodense lesions within the liver, which are likely metastatic lesions from a primary tumor elsewhere in the body.				
<b>Q:</b> Does the picture contain lung? <b>A:</b> No				
<b>Phi-3.5V-Med-FT:</b> Yes. The image includes the lungs, which are located in the thoracic cavity above the diaphragm and are visible in the transverse plane view of the abdomen.				
<b>Phi-3.5V-Med-FT w/ ARCD:</b> No, the image does not contain lung. The focus of the CT scan is on the abdomen, which includes organs like the liver, kidneys, and spleen, but does not include the lungs, as they are located in the thoracic cavity above the abdomen.				
<b>Question:</b> Does the picture contain brain? <b>Answer:</b> Yes				
<b>Phi-3.5V-Med-ZS:</b> Yes, the image contains the brain, which is the central part of the human body. The brain is a complex organ that is responsible for various functions, including cognition, memory, and motor control. The image shows the brain in a cross-sectional view, allowing for the visualization of its internal structures.				
<b>Phi-3.5V-Med-ZS w/ ARCD:</b> Yes, the image contains the brain. It shows a cross-sectional view of the brain, with the cerebral cortex, ventricles, and other anatomical structures clearly visible.				
<b>Question:</b> Where is/are the abnormality located? <b>A:</b> Lower right				
<b>Phi-3.5V-Med-FT:</b> Left lobe. The MRI image reveals an abnormality situated in the left lobe of the brain, indicating potential issues such as edema or a tumor affecting that specific region.				
<b>Phi-3.5V-Med-FT w/ ARCD:</b> Lower Right. The MRI image reveals brain edema, specifically located in the lower right region of the brain, indicating swelling due to excess fluid accumulation.				
<b>Question:</b> What are all the tubes/lines located within the right apical zone? <b>Answer:</b> Chest tube				
<b>Phi-3.5V-Med-ZS:</b> The image shows the presence of multiple tubes and lines within the right apical zone, which are likely medical devices or equipment used for monitoring or treatment purposes. These tubes and lines are visible as radiopaque structures on the chest X-ray. The specific purpose and location of these tubes and lines would require further clinical information to determine their exact nature and function.				
<b>Phi-3.5V-Med-ZS w/ ARCD:</b> The image shows several tubes and lines within the right apical zone. These include a central venous catheter, which is a tube inserted into a large vein, likely the superior vena cava. Additionally, there are two chest tubes, which are used to drain fluid or air from the pleural space around the lungs. These tubes and lines are essential for monitoring and managing the patient's condition.				
<b>Phi-3.5V-Med-FT:</b> The only tube located within the right apical zone is the endotracheal tube, which is noted to terminate within this zone in the image.				
<b>Phi-3.5V-Med-FT w/ ARCD:</b> Chest tube. The image shows a new right-sided chest tube terminating at the apex, confirming its presence in the right apical zone.				
<b>Question:</b> What anatomical finding, lens opacification or vitreous opacification, is associated with the left eye? <b>Answer:</b> Vitreous opacification				
<b>Phi-3.5V-Med-ZS:</b> The image shows a lens opacification in the left eye, as indicated by the reference information provided.				
<b>Phi-3.5V-Med-ZS w/ ARCD:</b> Based on the information provided in the image, the left eye shows lens opacification rather than vitreous opacification. The image clearly depicts the opacification of the lens, which is a characteristic finding of cataracts or other lens-related pathologies.				
<b>Phi-3.5V-Med-FT:</b> The presence of lens opacification is evident in the left eye, as indicated by the observed echo enhancement within the lens structure. This finding suggests a significant degree of opacification affecting visual clarity.				
<b>Phi-3.5V-Med-FT w/ ARCD:</b> The presence of weak echoes in the vitreous suggests vitreous opacification. The image indicates that there are scattered punctate and band-like weak echoes, which are consistent with this condition.				

Figure 6: Qualitative comparison of model-generated responses across four medical imaging modalities: Abdominal CT, Head MRI, Chest X-ray, and Ocular B-scan ultrasound. We compare our proposed method against *Phi-3.5V-Med* zero-shot and fine-tuned settings. Since the global and local masks coincide, only the local mask is visualized as small red boxes.

performance, and we speculate this is because unprocessed textual information, which are lost during cropping, provide a helpful prior. Notably, simply overlaying a bounding box on the relevant region consistently provides a stable performance improvement, as noted in (Shtedritski, Rupprecht, and Vedaldi 2023), which indirectly highlights the importance of attention priors for enhancing VLM accuracy and reliability and demonstrates the effectiveness of ARCD.

## Case Study

Figure 6 showcases a qualitative comparison of the responses generated by our proposed ARCD against those from the Zero-shot Med-Phi3.5V and the Fine-tuned Med-Phi3.5V baselines. The cases are selected from a diverse range of medical imaging modalities, including chest X-ray, abdominal CT, brain MRI, and ocular B-scan ultrasound, with expert-annotated masks provided as a ground-truth reference for the clinically relevant anatomical regions. As observed, the baseline models exhibit distinct failure modes across various tasks. The fine-tuned model, for instance, often contradicts facts by incorrectly identifying an abnormality in the *Left lobe* of the brain instead of the correct *Lower right* and hallucinates the presence of lungs in an abdominal CT scan. Similarly, both baseline models misidentify the condition in an ocular ultrasound as *lens opacification* when it is *vitreous opacification*. In contrast, our ARCD strategy

consistently provides the correct answers in these challenging cases. This demonstrates that by leveraging anatomical masks to ground the model’s attention, our approach effectively mitigates factual errors and hallucinations, leading to more accurate and reliable diagnostic responses.

## Conclusion

In conclusion, hallucination remains a critical barrier to the clinical adoption of Medical Vision-Language Models (MedVLMs). To address this, we introduced Anatomical Region-Guided Contrastive Decoding (ARCD), a novel, training-free strategy that steers language generation using specified anatomical regions. By applying a three-tiered, contrastive guidance mechanism, our method provides fine-grained control over the model’s visual attention, directly tackling the core problem of poor visual grounding. This plug-and-play approach avoids expensive retraining while offering more reliable correction than existing inference-time methods. Extensive experiments across diverse datasets validate its effectiveness and generalization. Our work offers a new approach for mitigating hallucinations in medical vision-language models, driving progress in the field.

## Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grants 62192730, 62192734, 62577041, the National Science and Technology Major Project under Grant 2022ZD0117103, the Outstanding Youth Science Foundation of Shaanxi Province under Grant 2025JC-JCQN-083, the Key Research and Development Program of Shaanxi Province under Grants 2025CY-YBXM-047 and 2024GX-YBXM-140, and the CCF-Huawei Populus Grove Fund under Grant CCF-HuaweiFM202507.

## References

- Abdin, M.; Jacobs, S. A.; Awan, A. A.; et al. 2024. Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone. *ArXiv*, abs/2404.14219.
- Bae, S.; Kyung, D.; Ryu, J.; Cho, E.; Lee, G.; Kweon, S.; Oh, J.; Ji, L.; Chang, E. I.-C.; Kim, T.; and Choi, E. 2023. EHRXQA: A Multi-Modal Question Answering Dataset for Electronic Health Records with Chest X-ray Images. *ArXiv*, abs/2310.18652.
- Bai, J.; Bai, S.; Yang, S.; Wang, S.; and et al., S. T. 2023. Qwen-VL: A Frontier Large Vision-Language Model with Versatile Abilities. volume abs/2308.12966.
- Bai, Z.; Wang, P.; Xiao, T.; He, T.; Han, Z.; Zhang, Z.; and Shou, M. Z. 2024. Hallucination of Multimodal Large Language Models: A Survey. *ArXiv*, abs/2404.18930.
- Chen, C.; Liu, M.; Jing, C.; Zhou, Y.; Rao, F.; Chen, H.; Zhang, B.; and Shen, C. 2025. PerturboLLaVA: Reducing Multimodal Hallucinations with Perturbative Visual Training. *arXiv preprint arXiv:2503.06486*.
- Chen, J.; Ouyang, R.; Gao, A.; Chen, S.; and et al., G. H. C. 2024a. HuatuoGPT-Vision, Towards Injecting Medical Visual Knowledge into Multimodal LLMs at Scale. *ArXiv*, abs/2406.19280.
- Chen, L.; Li, J.; Dong, X.; Zhang, P.; He, C.; Wang, J.; Zhao, F.; and Lin, D. 2023. ShareGPT4V: Improving Large Multi-Modal Models with Better Captions. *arXiv preprint arXiv:2311.12793*.
- Chen, L.; et al. 2024b. Low-hallucination Synthetic Captions for Large-Scale Vision-Language Model Pre-training. *arXiv preprint arXiv:2504.13123*.
- Chen, X.; Lai, Z.; Ruan, K.; Chen, S.; Liu, J.; and Liu, Z. 2024c. R-LLaVA: Improving Med-VQA Understanding through Visual Region of Interest. *arXiv preprint arXiv:2410.20327*.
- Chuang, Y.; Xie, Y.; Luo, H.; Kim, Y.; Glass, J.; and He, P. 2023. DoLa: Decoding by Contrasting Layers Improves Factuality in Large Language Models. *CoRR*, abs/2309.03883.
- Favero, A.; Zancato, L.; Trager, M.; Choudhary, S.; Perera, P.; Achille, A.; Swaminathan, A.; and Soatto, S. 2024. Multi-modal hallucination control by visual information grounding. *arXiv preprint arXiv:2403.14003*.
- Feng, Y.; Hu, H.; Hou, X.; Liu, S.; Ying, S.; Du, S.; Hu, H.; and Gao, Y. 2024. Hyper-RAG: Combating LLM Hallucinations using Hypergraph-Driven Retrieval-Augmented Generation. *arXiv preprint arXiv:2504.08758*.
- Gai, X.; Zhou, C.; Liu, J.; Feng, Y.; Wu, J.; and Liu, Z. 2024. Medthink: Explaining medical visual question answering via multimodal decision-making rationale. *arXiv preprint arXiv:2404.12372*.
- Gu, T.; Yang, K.; Liu, D.; and Cai, W. 2024. Lapa: Latent prompt assist model for medical visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4971–4980.
- Ha, C. N.; Asaadi, S.; Karn, S. K.; Farri, O.; Heimann, T.; and Runkler, T. 2024. Fusion of Domain-Adapted Vision and Language Models for Medical Visual Question Answering. *arXiv preprint arXiv:2404.16192*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; and Chen, W. 2021. LoRA: Low-Rank Adaptation of Large Language Models. *ArXiv*, abs/2106.09685.
- Huang, Q.; Dong, X.; Zhang, P.; et al. 2023. OPERA: Alluviating Hallucination in Multi-Modal Large Language Models via Over-Trust Penalty and Retrospection-Allocation. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13418–13427.
- Lan, W.; Chen, W.; Chen, Q.; Pan, S.; Zhou, H.; and Pan, Y. 2024. A Survey of Hallucination in Large Visual Language Models. *arXiv preprint arXiv:2410.15359*.
- Leng, S.; Zhang, H.; Chen, G.; Li, X.; Lu, S.; Miao, C.; and Bing, L. 2023. Mitigating Object Hallucinations in Large Vision-Language Models through Visual Contrastive Decoding. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13872–13882.
- Li, C.; Wong, C.; Zhang, S.; Usuyama, N.; Liu, H.; Yang, J.; Naumann, T.; Poon, H.; and Gao, J. 2023a. Llavamed: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36: 28541–28564.
- Li, C.; Wong, C.; Zhang, S.; Usuyama, N.; Liu, H.; Yang, J.; Naumann, T.; Poon, H.; and Gao, J. 2023b. LLaVA-Med: Training a Large Language-and-Vision Assistant for Biomedicine in One Day. *ArXiv*, abs/2306.00890.
- Liu, B.; Zhan, L.-M.; Xu, L.; Ma, L.; Yang, Y. F.; and Wu, X.-M. 2021. Slake: A Semantically-Labeled Knowledge-Enhanced Dataset For Medical Visual Question Answering. *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 1650–1654.
- Liu, H.; Xue, W.; Chen, Y.; Chen, D.; Zhao, X.; Wang, K.; Hou, L.; Li, R.; and Peng, W. 2024a. A Survey on Hallucination in Large Vision-Language Models. *arXiv preprint arXiv:2402.00253*.
- Liu, Y.; et al. 2024b. Re-Balancing Contrastive Decoding: A Unified Framework for VLM Hallucination Mitigation. *arXiv preprint arXiv:2409.06485*.
- OpenAI. 2023. GPT-4 Technical Report. *ArXiv*, abs/2303.08774.

Park, W.; Kim, W.; Kim, J.; and Do, J. 2025. SECOND: Mitigating Perceptual Hallucination in Vision-Language Models via Selective and Contrastive Decoding. *arXiv preprint arXiv:2506.08391*.

Seyfioglu, M. S.; Ikezogwo, W. O.; Ghezloo, F.; Krishna, R.; and Shapiro, L. G. 2023. Quilt-LLaVA: Visual Instruction Tuning by Extracting Localized Narratives from Open-Source Histopathology Videos. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13183–13192.

Shtedritski, A.; Rupprecht, C.; and Vedaldi, A. 2023. What does CLIP know about a red circle? Visual prompt engineering for VLMs. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 11953–11963.

Sun, G.; Qin, C.; Fu, H.; Wang, L.; and Tao, Z. 2024. Self-Training Large Language and Vision Assistant for Medical Question Answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 20052–20060.

Wang, X.; Pan, J.; Ding, L.; and Biemann, C. 2024. Mitigating Hallucinations in Large Vision-Language Models with Instruction Contrastive Decoding. In *Findings of the Association for Computational Linguistics: ACL 2024*.

Wu, C.; Zhang, X.; Zhang, Y.; Wang, Y.; and Xie, W. 2023. Towards Generalist Foundation Model for Radiology. *ArXiv*, abs/2308.02463.

Wu, J.; Kim, Y.; and Wu, H. 2024. Hallucination benchmark in medical visual question answering. *arXiv preprint arXiv:2401.05827*.

Yan, Q.; Yuan, Y.; Hu, X.; Wang, Y.; Xu, J.; Li, J.; Fu, C.-W.; and Heng, P.-A. 2025. MedHallTune: An Instruction-Tuning Benchmark for Mitigating Medical Hallucination in Vision-Language Models. *arXiv preprint arXiv:2502.20780*.

Yang, Z.; Luo, X.; Han, D.; Xu, Y.; and Li, D. 2025. Mitigating Hallucinations in Large Vision-Language Models via DPO: On-Policy Data Hold the Key. *arXiv preprint arXiv:2501.09695*.

Yin, S.; Fu, C.; Zhao, S.; Xu, T.; Wang, H.; Sui, D.; Shen, Y.; Li, K.; Sun, X.; and Chen, E. 2024. Woodpecker: Hallucination correction for multimodal large language models. *Science China Information Sciences*.

Zeng, W.; Sun, Y.; Ma, C.; Tan, W.; and Yan, B. 2025. MM-Skin: Enhancing Dermatology Vision-Language Model with an Image-Text Dataset Derived from Textbooks. *ArXiv*, abs/2505.06152.

Zhang, Y.; Qian, S.; Peng, B.; Liu, S.; and Jia, J. 2023. Prompt Highlighter: Interactive Control for Multi-Modal LLMs. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13215–13224.

Zhao, e. a. 2023. Hallucination-Aware Direct Preference Optimization for Large Multimodal Models. *arXiv preprint arXiv:2311.16839*.

Zhu, K.; Xia, P.; Li, Y.; Zhu, H.; Wang, S.; and Yao, H. 2024. MMedPO: Aligning Medical Vision-Language Models with Clinical-Aware Multimodal Preference Optimization. *arXiv preprint arXiv:2412.06141*.