

Mitigating Object Hallucinations in MLLMs via Multi-Frequency Perturbations

Shuo Li^{*,1}, Jiajun Sun^{*,1}, Guodong Zheng¹, Xiaoran Fan¹, Yujiong Shen¹,
Yi Lu¹, Zhiheng Xi¹, Wenming Tan², Tao Ji^{†,1}, Tao Gui^{†,1,3,4}, Qi Zhang¹, Xuanjing Huang¹
Fudan University¹ Hikvision Research Institute²
Shanghai Key Lab of Intelligent Information Processing³
Pengcheng Laboratory⁴
lis23@m.fudan.edu.cn, {taoji, tgui}@fudan.edu.cn

Abstract

Recently, multimodal large language models (MLLMs) have demonstrated remarkable performance in visual-language tasks. However, the authenticity of the responses generated by MLLMs is often compromised by object hallucinations. We identify that a key cause of these hallucinations is the model’s over-susceptibility to image frequency features in detecting objects. In this paper, we introduce Multi-Frequency Perturbations (MFP), a simple, cost-effective, and pluggable adversarial training method that leverages both low-frequency and high-frequency features of images to perturb visual feature representations and explicitly suppress redundant frequency-domain features during inference, thereby mitigating hallucinations. Experimental results demonstrate that our method significantly mitigates object hallucinations across various model architectures. Furthermore, as a training-time method, MFP can be combined with inference-time methods to achieve state-of-the-art performance on the CHAIR benchmark.

1 Introduction

Large language models (LLMs), exemplified by ChatGPT (OpenAI, 2022), demonstrate remarkable performance across diverse text-based tasks. By integrating visual encoders like CLIP (Radford et al., 2021), multimodal large language models (MLLMs) extend these capabilities to visual domains. These models excel in a range of vision-language tasks, including image caption (Wang et al., 2020), visual question answering (Antol et al., 2015), and visual dialogue (Das et al., 2017).

However, MLLMs always face the challenge of object hallucinations (Rohrbach et al., 2018; Li et al., 2023d), where the model’s outputs fail to accurately correspond to the objects in the real

image. This issue undermines model performance and diminishes its credibility.

To tackle this issue, previous studies propose various methods (Bai et al., 2024; Lan et al., 2024), which can be broadly categorized into training-time and inference-time methods. However, these approaches overlook a critical weakness in MLLMs. Specifically, the MLLMs tend to be over-susceptible to frequency features in the image—namely, abrupt changes in pixel values—when detecting objects, often overlooking the actual features of real images. For instance, as illustrated in Figure 1, MLLMs can still identify objects in images, even when blurred images containing certain frequencies are used. This characteristic may contribute to the occurrence of object hallucination. This phenomenon has also been widely observed in studies on image frequency-domain attacks targeting vision models (Wang et al., 2021; Long et al., 2022; Feng et al., 2022) in other models.

Based on these insights, we introduce Multi-Frequency Perturbations (MFP), a simple, cost-effective, and pluggable method that leverages both low-frequency and high-frequency features of images to perturb visual feature representations and explicitly suppress redundant frequency-domain features during inference. MFP partitions an image into high-frequency and low-frequency components, extracts their respective features, and performs a fine-grained fusion of these features with the original image features at the visual token level. This process results in more robust image features. The module is composed of two main components: (1) Multi-Frequency Feature Extraction. In this step, Gaussian high-pass and low-pass filters (Young and Van Vliet, 1995) are applied to the original image to obtain its raw high-frequency and low-frequency features. These raw features are then fed into the visual encoder, producing the corresponding high-frequency and low-frequency

* Equal contributions.

† Corresponding author.

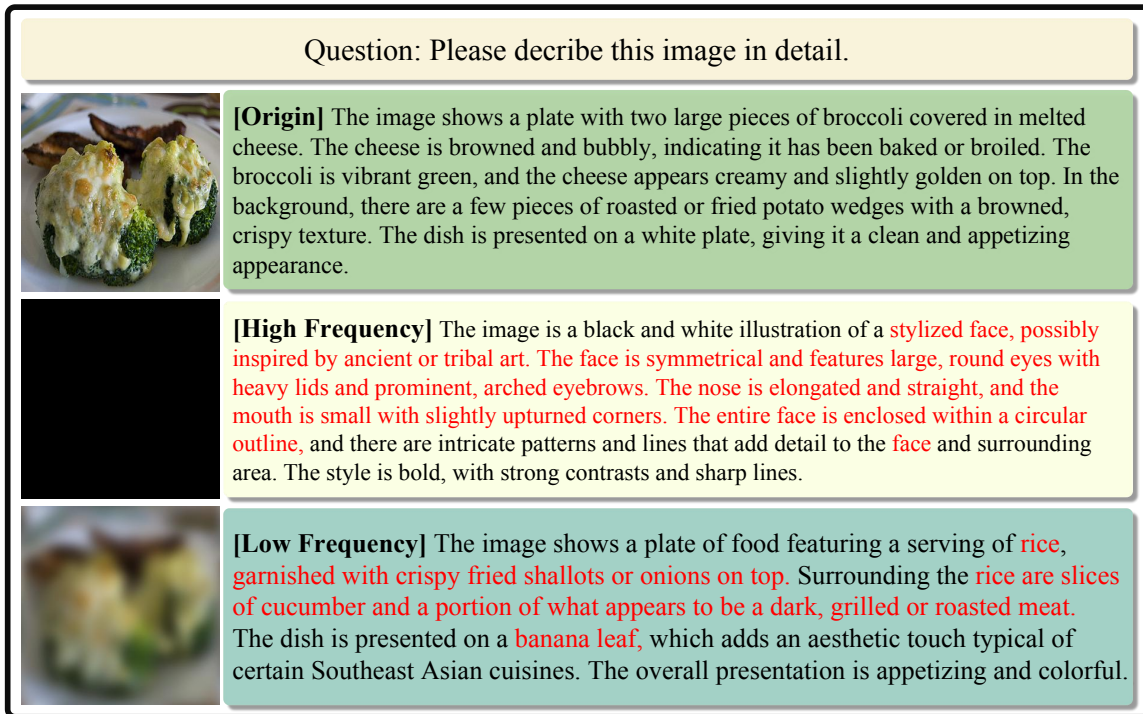


Figure 1: An example of GPT-4o. Unlike humans, the model is over-susceptible to limited high- and low-frequency image features to detect objects, leading to incorrect object detection and erroneous image caption.

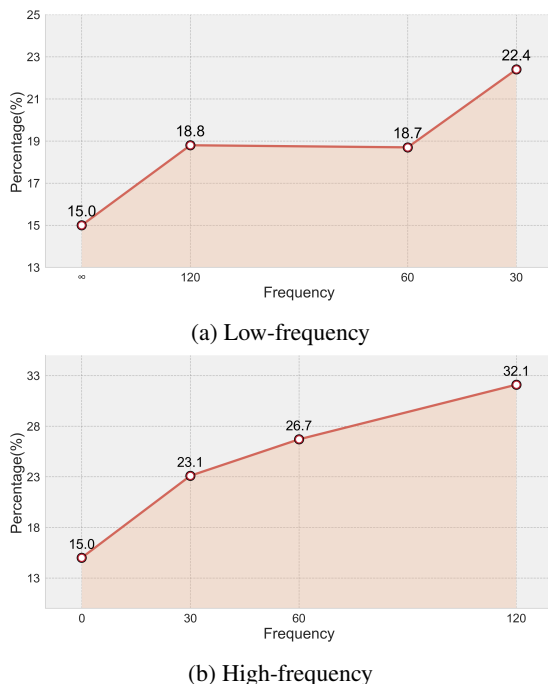


Figure 2: Instance-level hallucination rate when using only low or high frequency features. The x-axis represents the cutoff frequency. Features with frequencies higher than the cutoff are retained as high-frequency features, while those below the cutoff are selected as low-frequency features.

features, represented as visual token sequences. (2) Fine-Grained Frequency Feature Fusion. We use the original visual token sequence derived from the image encoder along with the high-frequency and low-frequency token sequences. These are fused at the token level using a cross-attention mechanism to generate the final perturbed visual token sequences. During inference, we apply a decay to both high-frequency and low-frequency feature perturbations, thereby reducing redundant high-frequency and low-frequency features. Experimental results demonstrate that our method achieves remarkable effectiveness in MLLM object hallucination benchmarks, regardless of visual encoders, LLM backbones, resolutions, or size. Additionally, our training-time approach can be combined with inference-time methods to achieve better performance and even achieve SOTA results on CHAIR (Rohrbach et al., 2018).

In this paper, our main contributions are:

- we are the first to address the cause of object hallucinations from the frequency domain perspective and introduce MFP, a simple, efficient, and pluggable method that effectively mitigates object hallucination in MLLMs;
- we demonstrate the effectiveness of our proposed method across models of various architectures,

3 Over-Susceptibility to Frequency in MLLM

Unlike most previous work that attributes hallucinations to text tokens’ interference or insufficient attention to visual tokens, we find that the models are over-susceptible to the frequency of the image (that is, how drastically the pixel values of the image change) when recognizing objects. We explore two experimental settings: (1) natural images and (2) isolated frequency features.

3.1 Analysis in Natural Images

Logits Lens To exclude the interference of text tokens and focus solely on image features, we **use images as the only input** and decode each visual token into text representations using the logits lens method (Nostalgebraist, 2020; Wendler et al., 2024; Neo et al., 2025) in LLaVA-1.5 (Liu et al., 2024a). Formally, given model parameters $\bar{\theta}$, an input image I and vocabulary VB , we obtain the last layer’s hidden states for each i -th visual token from the forward pass $h_i = H(\bar{\theta}, I_i)$ and project it onto the vocabulary probability distribution using the unembedding matrix W_u :

$$t_i = \arg \max(h_i W_u), t_i \in VB, \quad (1)$$

where t_i is the text token corresponding to the i -th visual token. Figure 3 presents an example of applying the logits-lens method.

Using the logits-lens method, we probe 500 randomly selected images from the MSCOCO dataset (Lin et al., 2014) and find that in **at least 34% of hallucinating cases**, the hallucinated objects mentioned in the generated captions also appear in the textual representations of the corresponding visual tokens extracted via logits lens. This suggests that **the model hallucinates objects as early as the visual perception**.

Slight Perturbations To reveal the model’s over-susceptibility to frequency changes in images, we apply slight perturbations to the image’s frequency content—specifically, mild blurring and sharpening using Gaussian filtering. These alterations are subtle and typically inconspicuous to the human. We then compare the logits-lens probing results before and after the perturbation. Our observations, as illustrated in Figure 4, show that the objects recognized by the model through logits-lens change significantly, and this shift also influences the generated image captions.

3.2 Analysis in Isolated Frequency Features

Considering that the experimental results on natural images may be influenced by various intrinsic factors—such as absolute pixel values and objects’ layouts—we further explore the model’s outputs using isolated frequency features.

As shown in Figure 5, we apply filters to retain only the high- or low-frequency features of an image, causing significant distortion that makes the objects barely distinguishable. We then prompt the MLLMs with the instruction: “Please describe this image in detail.” However, we observe a common phenomenon: almost all of the MLLMs, including GPT-4o (Hurst et al., 2024)—one of the most advanced MLLMs—fail to recognize images as meaningless or devoid of objects. Instead, they often generate incorrect captions and hallucinate objects.

Next, we use quantitative methods to validate the existence of this phenomenon. As shown in Figure 2, we randomly selected 200 images from the MSCOCO dataset, retaining only a portion of the high- or low-frequency features. We then prompted the LLaVA-1.5 to describe these images and calculated the proportion of instances containing object hallucinations (Recognizing images as meaningless or devoid of objects will not be classified as hallucination). We observed a significant increase in the proportion of hallucinatory sentences as less high-frequency feature was preserved (with the cut-off frequency ranging from 0 to 120), rising from 15.0 to 32.1. A similar trend was seen with low-frequency features (cutoff frequency from infinity to 30), where the proportion increased from 15.0 to 24.4.

Based on these results, we observe that MLLMs tend to be over-susceptible to frequency features in the image for detecting objects. However, this tendency makes MLLMs vulnerable and overly sensitive, which can lead to hallucinations in the models.

4 Multi-Frequency Perturbations

In Section 3, we demonstrate that MLLMs are overly sensitive to frequency features during object recognition, which leads to object hallucinations. In this section, we present our proposed robust training method based on multi-frequency perturbations to address this issue.

Model Architecture Our model architecture, similar to most MLLMs, as shown in Figure 5,

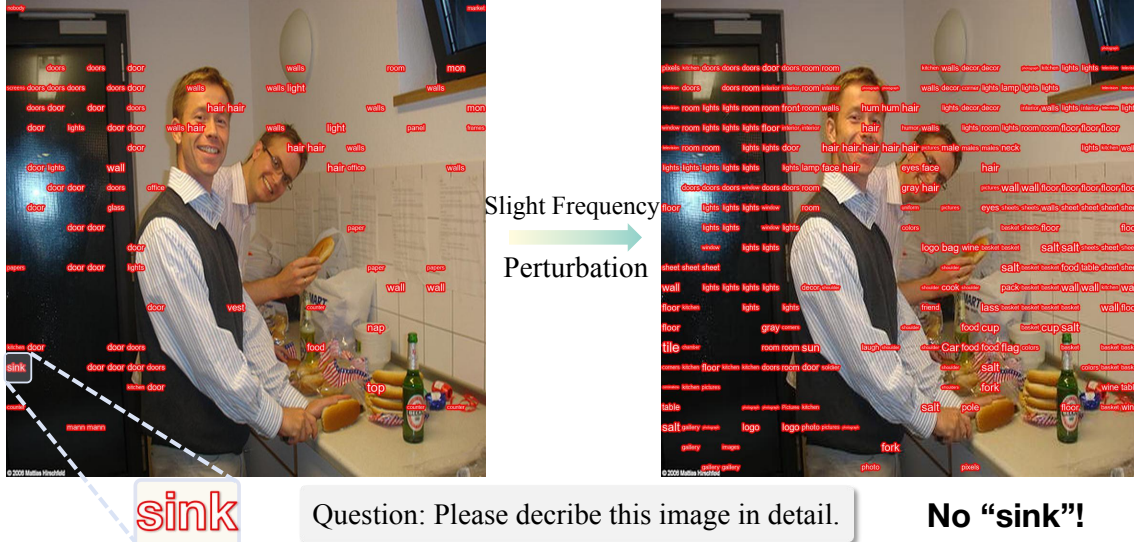


Figure 4: An example of the comparison of results before and after slight frequency perturbation.

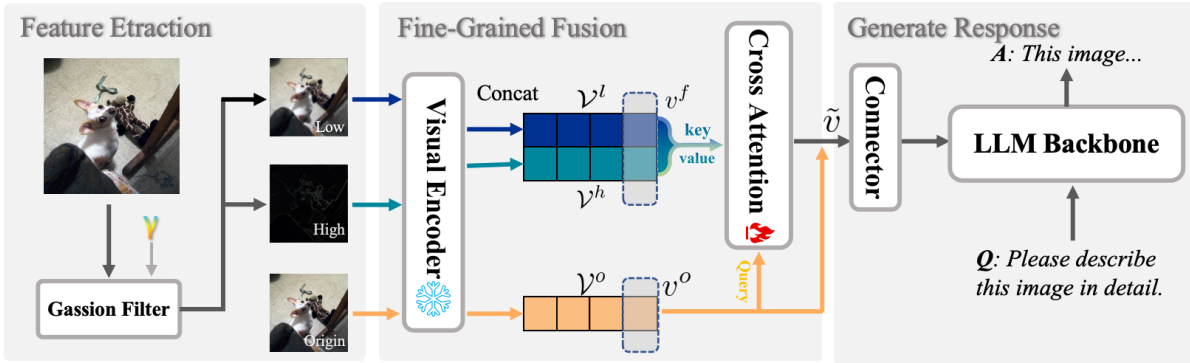


Figure 5: The model architecture of our proposed method. Where γ is only employed at inference time.

consists of three main components: a CLIP-like visual encoder, connectors, and a large Language Model (LLM). During training, the Gaussian filter is used to extract both high- and low-frequency features from the image. These features are then fine-grained and fused with the original image features. The resulting fused features as visual tokens are aligned with the text tokens via the connector. The concatenated tokens are then fed into the LLM backbone to generate the response.

Multi-Frequency Feature Extraction We use a Gaussian filter to extract the high-frequency and low-frequency features from the image. Specif-

ically, let $\mathcal{I}^o(u, v) \in \mathbb{R}^{h \times w \times 3}$ denote the original image, an RGB three-channel image of height $u \in [0, h]$ and width $v \in [0, w]$. First, we apply the Fourier transform to each channel separately:

$$\mathcal{F}_c(u, v) = FFT(\mathcal{I}_c(u, v)), c \in \{R, G, B\}, \quad (2)$$

then we define the Gaussian low-frequency filter $\mathcal{H}_c^l(u, v)$ and high-frequency filter $\mathcal{H}_c^h(u, v)$ as follows:

$$\begin{cases} \mathcal{H}_c^l(u, v) = \exp\left(-\frac{\mathcal{D}^2(u, v)}{2\mathcal{D}_0^2}\right) \\ \mathcal{H}_c^h(u, v) = 1 - \exp\left(-\frac{\mathcal{D}^2(u, v)}{2\mathcal{D}_0^2}\right), \end{cases} \quad (3)$$

where $\mathcal{D}^2(u, v)$ represents the distance from the pixel point (u, v) to the center of the frequency rectangle, and \mathcal{D}_0 denotes the cutoff frequency. These filters are then applied to filter the low and high frequency of the original image:

$$\begin{cases} \mathcal{F}_c^l(u, v) = \mathcal{F}_c(u, v) \cdot \mathcal{H}_c^l(u, v) \\ \mathcal{F}_c^h(u, v) = \mathcal{F}_c(u, v) \cdot \mathcal{H}_c^h(u, v). \end{cases} \quad (4)$$

Finally, we apply the inverse Fourier transform to these frequencies:

$$\begin{cases} \mathcal{I}_c^l(u, v) = FFT^{-1}(\mathcal{F}_c^l(u, v)) \\ \mathcal{I}_c^h(u, v) = FFT^{-1}(\mathcal{F}_c^h(u, v)), \end{cases} \quad (5)$$

and combine them into the RGB channels to obtain the low-frequency features $\mathcal{I}^l(u, v)$ and high-frequency features $\mathcal{I}^h(u, v)$ of the image.

Fine-Grained Frequency Feature Fusion After obtaining the low-frequency features $\mathcal{I}^l(u, v)$ and high-frequency features $\mathcal{I}^h(u, v)$ of the image, we fuse them to the origin image features using a cost-effective method. We then encode these features with the visual encoder VG (e.g., CLIP) to obtain a sequence of visual tokens $V \in \mathbb{R}^{L \times dim}$:

$$\mathcal{V} = VG(\mathcal{I}), \quad (6)$$

where L depends on the settings of the visual encoder, and dim refers to the dimension of the visual encoder's hidden layer.

For each token v in the visual token sequence we obtained, we use cross-attention to fuse the high-frequency token v^h and low-frequency visual token v^l into the visual token v^o of the original image at the same position in the sequence. This process can be expressed as follows:

$$\begin{aligned} v^o &\in \mathcal{V}^o, v^l \in \mathcal{V}^l, v^h \in \mathcal{V}^h, \\ v^f &= v^l \oplus v^h, \\ \tilde{v} &= \text{softmax} \left(\frac{v^o W^q (v^f W^k)^\top}{\sqrt{d_k}} \right) v^f W^v + v^o, \end{aligned} \quad (7)$$

where \oplus represents the concatenation operation, $v_f \in \mathbb{R}^{2 \times dim}$ is the multi-frequency token obtained by concatenating v^l and v^h . W^q , W^k , and $W^v \in \mathbb{R}^{dim \times dim}$ are the projection matrices. This means that the original image token is used as the query in the attention calculation, while the multi-frequency token serves as the key and value. Since our cross attention only applies to tokens in the same position

of the sequence, we only have a small attention score matrix, which means a small computational cost.

Finally, we stack \tilde{v} to obtain the final visual token sequence \tilde{V} , which is aligned with the text tokens through the connector. These tokens are then concatenated and fed into the LLM backbone to generate the response.

Training & Inference Our training method follows the setting used in LLaVA (Liu et al., 2024c), consisting of two stages: the pre-training (PT) stage and the supervised-finetuning (SFT) stage. During the PT stage, only the connector and the three projection matrices (W_q, W_k, W_v) are trained. In the SFT stage, the model connector, the three projection matrices in Equation 7, and the LLM backbone are trained. Compared to LLaVA, we introduce very few additional training parameters (only $W_q, W_k, W_v, \approx 3M$).

During model inference, similar to practices in adversarial training, we introduce a decay factor γ in the multi-frequency feature extraction process to modulate the influence of perturbed frequency features, considering that these features are not pure noise but still contain informative signals. Specifically, we rewrite the process in the Equation 4 as follows:

$$\begin{cases} \mathcal{F}_c^l(u, v) = \mathcal{F}_c(u, v) \cdot \mathcal{H}_c^l(u, v) \cdot G(\gamma) \\ \mathcal{F}_c^h(u, v) = \mathcal{F}_c(u, v) \cdot \mathcal{H}_c^h(u, v) \cdot G(\gamma), \end{cases} \quad (8)$$

where $G(\gamma)$ is a matrix whose values are obtained by randomly sampling from the uniform distribution $U(0, \gamma)$, where $\gamma \leq 1$.

5 Experiment

5.1 Setup

Implementation Detail Similar to most work in the field of MLLM hallucination, we apply our proposed method on LLaVA-1.5-7B (Liu et al., 2024a) for our experiments. The LCS-558k dataset (Liu et al., 2024a) is used during the pre-training phase, while the LLaVA-mixed-665k dataset (Liu et al., 2024a) is used during the SFT stage. As Appendix C, We keep the training parameters consistent with LLaVA-1.5. We set $\mathcal{D}_0=30$. At inference time, we employ sampling decoding with temperature=0.2. The codebase framework is PyTorch (Paszke et al., 2019), and experiments are conducted with $8 \times H100$ GPUs.

Baseline In this paper, we select several established methods as baselines, including DoLa (Chuang et al., 2023), ITI (Li et al., 2023c), VCD (Leng et al., 2024), ICD (Wang et al., 2024), SID (Huo et al., 2024), AGLA (An et al., 2024), OPERA (Huang et al., 2024), DOPRA (Wei and Zhang, 2024), HALC (Chen et al., 2024a), and CCA-LLaVA (Xing et al., 2025). Further details are provided in the Appendix A.1.

Benchmark We adopt widely used benchmarks from prior work, including CHAIR (Rohrbach et al., 2018), POPE (Li et al., 2023d), MME (Fu et al., 2024), and MMBench (Liu et al., 2024f), to ensure comprehensive and consistent evaluation. Further details are provided in the Appendix A.2.

5.2 Main Results

Comparison of MFP with existing methods To comprehensively evaluate the effectiveness of our proposed method, we compare its performance against various existing approaches across multiple evaluation metrics, as presented in Table 1. Our method consistently outperforms the baseline and other competing methods in key aspects.

First, in terms of POPE, our method achieves the highest F1 score of 86.2, surpassing the baseline (85.9) and all other methods, including DOPRA (85.6), SID (85.1), and CCA-LLaVA (85.5). Regarding CHAIR_s, our method attains a significantly lower score of 41.2, indicating mitigated sentence-level object hallucinations compared to the baseline (50.2) and most other methods, such as VCD (51.0) and DoLa (57.0). Notably, our method also outperforms AGLA (43.0) and CCA-LLaVA (43.0). For CHAIR_i, our method achieves a score of 11.7, which is competitive with the best-performing method (CCA-LLaVA, 11.5). This suggests that our method effectively mitigates instance-level hallucinations. Furthermore, the image captions generated by our model also maintain completeness relative to the baseline, as evidenced by the CHAIR F1 score (77.6 vs. 76.8). On MME, our method achieves the highest scores in Existence (195.0), Count (150.0), and Position (138.3), demonstrating strong performance in reducing hallucinations related to object presence, quantity, and spatial arrangement. With an overall MME score of 643.3, our approach ranks first, outperforming strong baselines such as VCD (604.6) and OPERA (592.3), further validating its effectiveness in hallucination mitigation.

Second, in terms of MMBench, our method achieves the highest accuracy of 68.2, outperforming all other approaches, including SID (65.1), OPERA (64.4), and DoLa (63.8).

In summary, our method demonstrates competitive performance in mitigating object hallucinations while maintaining and improving the baseline model’s general capabilities.

Generalization of MFP across Architectures

Table 2 demonstrates that MFP consistently improves performance across various model architectures, including different LLMs, visual encoders, model sizes, and input resolutions. For LLMs such as Vicuna1.5-7B (Zheng et al., 2023) and LLaMA2-7B (Touvron et al., 2023), MFP slightly increases POPE scores while significantly reducing CHAIR_s and CHAIR_i, indicating reduced hallucination. Similar improvements are observed across model sizes (e.g., CHAIR_s drops from 42.3 to 37.6 in Vicuna1.5-13B), and across visual encoders like CLIP (Radford et al., 2021) and SigLIP (Zhai et al., 2023). MFP also generalizes well across input resolutions, maintaining or improving POPE scores (e.g., 86.4 to 86.8 in LLaVA-Next (Liu et al., 2024b)) while reducing hallucination rates. These results confirm that MFP is architecture-agnostic and effective across diverse model settings.

Compatibility of MFP with existing SOTA method

As our method mitigates hallucinations from a novel perspective, it can be seamlessly integrated with existing method in an orthogonal manner. To our knowledge, PAI (Liu et al., 2024d) is the current SOTA method without additional data, achieving the best performance on the CHAIR benchmark. PAI enhances image token attention and adjusts logits to reduce text bias. However, due to its operation at the decoding and attention layers, PAI is not easily compatible with many methods (e.g., VCD, OPERA). Nonetheless, As a training-free inference-time method, PAI can be effectively combined with our training-time MFP.

As shown in Table 3, both methods perform well independently—MFP reduces CHAIR_s by 9.0 and CHAIR_i by 3.3, while PAI achieves reductions of 25.6 and 7.2, respectively. When combined, CHAIR_s and CHAIR_i are further reduced by 6.6 and 2.0 over PAI alone, with the F1 score maintained at 74.4, establishing a new SOTA on the CHAIR benchmark. These results demonstrate the strong compatibility of our method.

Method	POPE		CHAIR			MME					MMB
	F1	CHAIR _s [↓]	CHAIR _i [↓]	F1	Avg. Len	Existence	Count	Position	Color	Overall	Acc.
Baseline	85.9	50.2	15.0	76.8	99.2	175.7	124.7	114.0	51.0	565.3	63.0
DoLa (2023)	80.2	57.0	15.2	-	97.5	175.0	108.3	90.0	138.3	511.7	63.8
ITI (2023c)	83.7	48.2	13.9	-	98.6	-	-	-	-	-	-
VCD (2024)	83.2	51.0	14.9	-	101.9	184.7	138.3	<u>128.7</u>	53.0	604.7	63.9
ICD (2024)	-	47.4	13.9	-	-	185.0	117.9	117.5	162.1	582.5	63.1
SID (2024)	85.1	45.0	11.7	-	-	190.0	148.3	128.3	<u>175.0</u>	641.7	65.1
AGLA (2024)	84.6	<u>43.0</u>	14.1	78.9	98.8	-	-	-	-	-	-
OPERA (2024)	85.2	47.0	14.6	-	95.3	180.7	133.3	123.3	155.0	592.3	64.4
DOPRA (2024)	<u>85.6</u>	46.3	13.8	-	96.1	-	-	-	-	-	-
HALC (2024a)	83.9	50.2	12.4	-	97.2	-	-	-	-	-	-
CCA-LLaVA (2025)	85.5	43.0	11.5	-	96.6	<u>190.0</u>	<u>148.3</u>	128.3	175.0	<u>641.7</u>	<u>65.4</u>
Ours	86.2	41.2	<u>11.7</u>	77.6	94.4	195.0	150.0	138.3	160.0	643.3	68.2

Table 1: Compare results of MFP with other SOTA methods on POPE and CHAIR datasets. We set $\gamma = 0.23$. The evaluation results of the compared methods are from published papers. The best performances within each metric are **bolded**. The second best performances are underlined.

LLaVA	V. E.	LLM	Res.	POPE	CHAIR _s [↓]	CHAIR _i [↓]
V1.5	CLIP	vicuna1.5 ^{7B}	336	85.9	50.2	15.0
w/ MFP	CLIP	vicuna1.5 ^{7B}	336	86.2	41.2	11.7
V1.5	CLIP	vicuna1.5 ^{13B}	336	85.9	53.0	14.6
w/ MFP	CLIP	vicuna1.5 ^{13B}	336	86.4	37.6	10.4
V1.5	SigLIP	vicuna1.5 ^{7B}	384	86.4	47.0	12.0
w/ MFP	SigLIP	vicuna1.5 ^{7B}	384	85.4	42.0	11.9
V1.5	CLIP	llama2 ^{7B}	336	85.4	48.0	14.8
w/ MFP	CLIP	llama2 ^{7B}	336	86.2	46.0	14.1
Next	CLIP	vicuna1.5 ^{7B}	672	86.4	51.0	12.8
w/ MFP	CLIP	vicuna1.5 ^{7B}	672	86.8	45.4	11.7

Table 2: Results of MFP cross different architectures. The best performances within each setting are **bolded**. V. E. refer to visual encoder and Res. refer to resolution. From top to bottom we set γ to 0.23, 0.1, 0.1, 0.3 and 0.4. Considering that the data of LLaVA-Next is not publicly available, we train LLaVA-Next using the data of LLaVA-1.5.

Training	Inference	CHAIR _s [↓]	CHAIR _i [↓]	F1	Avg. Len
Vanilla	Vanilla	50.2	15.0	76.8	99.2
MFP	Vanilla	41.2	11.7	77.6	94.4
Vanilla	PAI	24.6	7.2	74.4	87.6
MFP	PAI	18.0	5.2	74.4	68.8

Table 3: Results of compatibility of MFP with existing SOTA method on CHAIR benchmark. $\gamma = 0.23$. The best performances within each setting are **bolded**. The experiments are conducted on the LLaVA-1.5-7B model.

5.3 More Results

Due to space limitations, we present additional experimental results and detailed analyses in the Appendix B. Specifically, we demonstrate that our proposed method preserves the general capabilities of MLLMs across several widely used benchmarks. To further validate its effectiveness, we evaluate the method on a broader set of hallucination benchmarks, where consistent improvements are observed. Additionally, we conduct a sensitivity analysis to better understand the underlying mechanism of MFP, and perform comprehensive ablation studies to evaluate the contribution of each component in the method.

6 Conclusion

We present the first study of object hallucination in MLLMs from a frequency-domain perspective, showing that models are overly sensitive to frequency features during object recognition. To address this, we propose MFP—a simple and efficient training-time method that effectively mitigates hallucinations. MFP achieves strong results across multiple benchmarks, generalizes well across architectures, and can be seamlessly combined with inference-time methods like PAI to achieve state-of-the-art performance on the CHAIR benchmark. These findings underscore the potential of frequency-domain techniques for hallucination mitigation and highlight the importance of further investigating the relationship between image features and object recognition in MLLMs.

Limitations

Due to time and computational constraints, our evaluation is conducted on a limited set of model architectures, which may impact the generalizability of our findings to a broader range of architectures. Additionally, our experiments do not include a comprehensive analysis of key hyperparameters, such as the number of feature fusion layers and the specific frequency ranges utilized. A more systematic exploration of these factors could provide deeper insights into their influence on performance and potentially lead to further improvements.

Acknowledgments

The authors wish to thank the anonymous reviewers for their helpful comments. This work was partially funded by Guangdong S&T Program 2024B0101050003, National Natural Science Foundation of China (No.62206057, 62076069, 61976056), Shanghai Rising-Star Program (23QA1400200), and Natural Science Foundation of Shanghai (23ZR1403500).

References

- Wenbin An, Feng Tian, Sicong Leng, Jiahao Nie, Haonan Lin, QianYing Wang, Guang Dai, Ping Chen, and Shijian Lu. 2024. Agla: Mitigating object hallucinations in large vision-language models with assembly of global and local attention. *arXiv preprint arXiv:2406.12718*.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433.
- Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou. 2024. Hallucination of multimodal large language models: A survey. *arXiv preprint arXiv:2404.18930*.
- Parichehr Behjati, Pau Rodriguez, Carles Fernández Tena, Armin Mehri, F Xavier Roca, Seiichi Ozawa, and Jordi Gonzalez. 2022. Frequency-based enhancement network for efficient super-resolution. *IEEE Access*, 10:57383–57397.
- Assaf Ben-Kish, Moran Yanuka, Morris Alper, Raja Giryes, and Hadar Averbuch-Elor. 2023. Mitigating open-vocabulary caption hallucinations. *arXiv preprint arXiv:2312.03631*.
- Ronald N Bracewell. 1989. The fourier transform. *Scientific American*, 260(6):86–95.
- Zhaorun Chen, Zhuokai Zhao, Hongyin Luo, Huaxiu Yao, Bo Li, and Jiawei Zhou. 2024a. Halc: Object hallucination reduction via adaptive focal-contrast decoding. *arXiv preprint arXiv:2403.00425*.
- Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, and 1 others. 2024b. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. 2023a. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv preprint arXiv:2312.14238*.
- Zhiyang Chen, Yousong Zhu, Yufei Zhan, Zhaowen Li, Chaoyang Zhao, Jinqiao Wang, and Ming Tang. 2023b. Mitigating hallucination in visual language models with visual supervision. *arXiv preprint arXiv:2311.16479*.
- Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. 2023. Dola: Decoding by contrasting layers improves factuality in large language models. *arXiv preprint arXiv:2309.03883*.
- Runmin Cong, Mengyao Sun, Sanyi Zhang, Xiaofei Zhou, Wei Zhang, and Yao Zhao. 2023. Frequency perception network for camouflaged object detection. In *Proceedings of the 31st ACM international conference on multimedia*, pages 1179–1189.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Preprint*, arXiv:2305.06500.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. 2017. Visual dialog. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 326–335.
- Xiaoyi Dong, Pan Zhang, Yuhang Zang, Yuhang Cao, Bin Wang, Linke Ouyang, Xilin Wei, Songyang Zhang, Haodong Duan, Maosong Cao, Wenwei Zhang, Yining Li, Hang Yan, Yang Gao, Xinyue Zhang, Wei Li, Jingwen Li, Kai Chen, Conghui He, and 4 others. 2024. Internlm-xcomposer2: Mastering free-form text-image composition and comprehension in vision-language large model. *arXiv preprint arXiv:2401.16420*.
- Xiaoran Fan, Tao Ji, Changhao Jiang, Shuo Li, Senjie Jin, Sirui Song, Junke Wang, Boyang Hong, Lu Chen, Guodong Zheng, and 1 others. 2024. Mousi: Poly-visual-expert vision-language models. *arXiv preprint arXiv:2401.17221*.

- Yu Feng, Benteng Ma, Jing Zhang, Shanshan Zhao, Yong Xia, and Dacheng Tao. 2022. Fiba: Frequency-injection based backdoor attack in medical image analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20876–20885.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, Yunsheng Wu, and Rongrong Ji. 2024. *Mme: A comprehensive evaluation benchmark for multimodal large language models*. *Preprint*, arXiv:2306.13394.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. *Making the v in vqa matter: Elevating the role of image understanding in visual question answering*. *Preprint*, arXiv:1612.00837.
- Wenxue Guan, Haobo Li, Dawei Xu, Jiabin Liu, Shenghua Gong, and Jun Liu. 2024. Frequency generation for real-world image super-resolution. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. 2018. *Vizwiz grand challenge: Answering visual questions from blind people*. *Preprint*, arXiv:1802.08218.
- Zongbo Han, Zechen Bai, Haiyang Mei, Qianli Xu, Changqing Zhang, and Mike Zheng Shou. 2024. Skip\n: A simple method to reduce hallucination in large vision-language models. *arXiv preprint arXiv:2402.01345*.
- 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13418–13427.
- Fushuo Huo, Wenchao Xu, Zhong Zhang, Haozhao Wang, Zhicheng Chen, and Peilin Zhao. 2024. Self-introspective decoding: Alleviating hallucinations for large vision-language models. *arXiv preprint arXiv:2408.02032*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Chaoya Jiang, Haiyang Xu, Mengfan Dong, Jiaying Chen, Wei Ye, Ming Yan, Qinghao Ye, Ji Zhang, Fei Huang, and Shikun Zhang. 2024. Hallucination augmented contrastive learning for multimodal large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27036–27046.
- Wei Lan, Wenyi Chen, Qingfeng Chen, Shirui Pan, Huiyu Zhou, and Yi Pan. 2024. A survey of hallucination in large visual language models. *arXiv preprint arXiv:2410.15359*.
- Seongyun Lee, Sue Hyun Park, Yongrae Jo, and Minjoon Seo. 2023. Volcano: mitigating multimodal hallucination through self-feedback guided revision. *arXiv preprint arXiv:2311.07362*.
- Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. 2024. 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.
- Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. 2023a. *Seed-bench: Benchmarking multimodal llms with generative comprehension*. *Preprint*, arXiv:2307.16125.
- 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.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR.
- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023c. Inference-time intervention: Eliciting truthful answers from a language model. *Advances in Neural Information Processing Systems*, 36:41451–41530.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023d. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*.
- Jiaying Lin, Xin Tan, Ke Xu, Lizhuang Ma, and Rynson WH Lau. 2023. Frequency-aware camouflaged object detection. *ACM Transactions on Multimedia Computing, Communications and Applications*, 19(2):1–16.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6-12, 2014, proceedings, part v 13*, pages 740–755. Springer.

- 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, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024b. *Llava-next: Improved reasoning, ocr, and world knowledge*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024c. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Shi Liu, Kecheng Zheng, and Wei Chen. 2024d. Paying more attention to image: A training-free method for alleviating hallucination in vlms. In *European Conference on Computer Vision*, pages 125–140. Springer.
- Xinyu Liu, Yingqing He, Lanqing Guo, Xiang Li, Bu Jin, Peng Li, Yan Li, Chi-Min Chan, Qifeng Chen, Wei Xue, and 1 others. 2024e. Hiprompt: Tuning-free higher-resolution generation with hierarchical mllm prompts. *arXiv preprint arXiv:2409.02919*.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, and 1 others. 2024f. Mmbench: Is your multi-modal model an all-around player? In *European conference on computer vision*, pages 216–233. Springer.
- Yuyang Long, Qilong Zhang, Boheng Zeng, Lianli Gao, Xianglong Liu, Jian Zhang, and Jingkuan Song. 2022. Frequency domain model augmentation for adversarial attack. In *European conference on computer vision*, pages 549–566. Springer.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. *Learn to explain: Multimodal reasoning via thought chains for science question answering*. *Preprint*, arXiv:2209.09513.
- Clement Neo, Luke Ong, Philip Torr, Mor Geva, David Krueger, and Fazl Barez. 2025. *Towards interpreting visual information processing in vision-language models*. *Preprint*, arXiv:2410.07149.
- Nostalgebraist. 2020. Interpreting gpt: The logit lens. <https://www.alignmentforum.org/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens>. Accessed: 2025-05-17.
- OpenAI. 2022. ChatGPT: Optimizing language models for dialogue. <https://openai.com/blog/chatgpt/>.
- OpenAI. 2024. *Gpt-4 technical report*. *Preprint*, arXiv:2303.08774.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, and 1 others. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. *Learning transferable visual models from natural language supervision*. *Preprint*, arXiv:2103.00020.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. Object hallucination in image captioning. *arXiv preprint arXiv:1809.02156*.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. *Towards vqa models that can read*. *Preprint*, arXiv:1904.08920.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, Kurt Keutzer, and Trevor Darrell. 2023a. *Aligning large multimodal models with factually augmented rlhf*. *Preprint*, arXiv:2309.14525.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, and 1 others. 2023b. *Aligning large multimodal models with factually augmented rlhf*. *arXiv preprint arXiv:2309.14525*.
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, Austin Wang, Rob Fergus, Yann LeCun, and Saining Xie. 2024a. *Cambrian-1: A fully open, vision-centric exploration of multimodal llms*. *Preprint*, arXiv:2406.16860.
- Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. 2024b. *Eyes wide shut? exploring the visual shortcomings of multimodal llms*. *Preprint*, arXiv:2401.06209.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, and 1 others. 2023. *Llama 2: Open foundation and fine-tuned chat models*. *arXiv preprint arXiv:2307.09288*.
- Haoran Wang, Yue Zhang, and Xiaosheng Yu. 2020. An overview of image caption generation methods. *Computational intelligence and neuroscience*, 2020(1):3062706.
- Tong Wang, Yuan Yao, Feng Xu, Shengwei An, Hanghang Tong, and Ting Wang. 2021. *Backdoor attack through frequency domain*. *arXiv preprint arXiv:2111.10991*.

- Xintong Wang, Jingheng Pan, Liang Ding, and Chris Biemann. 2024. [Mitigating hallucinations in large vision-language models with instruction contrastive decoding](#). *Preprint*, arXiv:2403.18715.
- Jinfeng Wei and Xiaofeng Zhang. 2024. Dopro: Decoding over-accumulation penalization and re-allocation in specific weighting layer. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 7065–7074.
- Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2024. [Do llamas work in english? on the latent language of multilingual transformers](#). *Preprint*, arXiv:2402.10588.
- Chenxi Xie, Changqun Xia, Tianshu Yu, and Jia Li. 2023. Frequency representation integration for camouflaged object detection. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 1789–1797.
- Yun Xing, Yiheng Li, Ivan Laptev, and Shijian Lu. 2025. Mitigating object hallucination via concentric causal attention. *Advances in Neural Information Processing Systems*, 37:92012–92035.
- Kai Xu, Minghai Qin, Fei Sun, Yuhao Wang, Yen-Kuang Chen, and Fengbo Ren. 2020. Learning in the frequency domain. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1740–1749.
- 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. *Science China Information Sciences*, 67(12):220105.
- Ian T Young and Lucas J Van Vliet. 1995. Recursive implementation of the gaussian filter. *Signal processing*, 44(2):139–151.
- Qifan Yu, Juncheng Li, Longhui Wei, Liang Pang, Wentao Ye, Bosheng Qin, Siliang Tang, Qi Tian, and Yueting Zhuang. 2024a. Hallucidoctor: Mitigating hallucinatory toxicity in visual instruction data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12944–12953.
- Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwen He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, and 1 others. 2024b. Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13807–13816.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2024c. [Mm-vet: Evaluating large multimodal models for integrated capabilities](#). *Preprint*, arXiv:2308.02490.
- Zihao Yue, Liang Zhang, and Qin Jin. 2024. Less is more: Mitigating multimodal hallucination from an eos decision perspective. *arXiv preprint arXiv:2402.14545*.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11975–11986.
- Dengsheng Zhang. 2019. Wavelet transform. In *Fundamentals of image data mining: Analysis, Features, Classification and Retrieval*, pages 35–44. Springer.
- Pan Zhang, Xiaoyi Dong, Bin Wang, Yuhang Cao, Chao Xu, Linke Ouyang, Zhiyuan Zhao, Shuangrui Ding, Songyang Zhang, Haodong Duan, Wenwei Zhang, Hang Yan, Xinyue Zhang, Wei Li, Jingwen Li, Kai Chen, Conghui He, Xingcheng Zhang, Yu Qiao, and 2 others. 2023. Internlm-xcomposer: A vision-language large model for advanced text-image comprehension and composition. *arXiv preprint arXiv:2309.15112*.
- Linxi Zhao, Yihe Deng, Weitong Zhang, and Quanquan Gu. 2024. Mitigating object hallucination in large vision-language models via classifier-free guidance. *arXiv preprint arXiv:2402.08680*.
- Zhiyuan Zhao, Bin Wang, Linke Ouyang, Xiaoyi Dong, Jiaqi Wang, and Conghui He. 2023. Beyond hallucinations: Enhancing lvlms through hallucination-aware direct preference optimization. *arXiv preprint arXiv:2311.16839*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *Preprint*, arXiv:2306.05685.
- Yiyang Zhou, Chenhang Cui, Rafael Rafailov, Chelsea Finn, and Huaxiu Yao. 2024. Aligning modalities in vision large language models via preference fine-tuning. *arXiv preprint arXiv:2402.11411*.
- Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. 2023. Analyzing and mitigating object hallucination in large vision-language models. *arXiv preprint arXiv:2310.00754*.
- Lanyun Zhu, Deyi Ji, Tianrun Chen, Peng Xu, Jieping Ye, and Jun Liu. 2024. Ibd: Alleviating hallucinations in large vision-language models via image-biased decoding. *arXiv preprint arXiv:2402.18476*.

A More Details About the Experiments Setup

A.1 More Details About the Baselines

DoLa (Chuang et al., 2023) is a simple decoding strategy that mitigates hallucinations in LLMs by

contrasting logits from different layers, without relying on external knowledge or fine-tuning. ITI (Li et al., 2023c) improves LLM truthfulness by modulating activations along specific attention head directions during inference. VCD (Leng et al., 2024) is a training-free method that contrasts output distributions from original and perturbed visual inputs to minimize bias and unimodal priors. ICD (Wang et al., 2024) contrasts distributions from standard and instruction disturbance, thereby increasing alignment uncertainty and effectively subtracting hallucinated concepts from the original distribution. SID (Huo et al., 2024) reduces vision-and-text association hallucinations by selectively preserving only the least important vision tokens early in decoding. AGLA (An et al., 2024) is a training-free method that mitigates object hallucinations by leveraging global features for response generation and local features for visual discrimination. OPERA (Huang et al., 2024) introduces a novel MLLM decoding strategy that reduces hallucinations through an over-trust penalty and a retrospection-allocation mechanism. DOPRA (Wei and Zhang, 2024) addresses hallucinations by applying weighted layer penalties and redistribution during decoding. HALC (Chen et al., 2024a) enhances vision-language tasks by incorporating fine-grained visual information and integrating both local and global contexts. CCA-LLaVA (Xing et al., 2025) proposes Concentric Causal Attention (CCA), a positional alignment strategy that mitigates RoPE’s long-term decay by reducing the relative distance between visual and instruction tokens in MLLMs.

A.2 More Detail About the Benchmarks

CHAIR (Rohrbach et al., 2018) Caption Hallucination Assessment with Image Relevance (CHAIR) is a widely used benchmark for assessing hallucinations in image captioning tasks. It works by creating a set of ground-truth object labels for each image, where any object mentioned in the caption that is not present in the label set is considered a hallucinated object. CHAIR consists of two variants: CHAIR_i (instance-level) and CHAIR_s (sentence-level), which are calculated as follows:

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

$$\text{CHAIR}_s = \frac{|\{\text{captions w/ hal. objects}\}|}{|\{\text{all captions}\}|}. \quad (10)$$

Consistent with the evaluation settings of previous work (Leng et al., 2024), we randomly sample 500 images from the MSCOCO 2014 validation set (Lin et al., 2014), set the max-tokens to 512, and use the prompt “Please describe this image in detail.” In addition to the CHAIR metric, we also report F1 scores to assess the completeness of the generated image descriptions.

POPE (Li et al., 2023d) The Polling-based Object Probing Evaluation (POPE) is a benchmark created to evaluate object hallucinations within the VQA framework. It works by asking MLLMs targeted questions like “Is there a {object} in the image?”, where {object} refers to a ground-truth item selected from three distinct splits. In the “random” split, objects are randomly chosen from the entire dataset. The “popular” split includes the most frequently occurring objects, and the “adversarial” split presents objects that are closely related to those in the image. For our analysis, we use the COCO dataset (Lin et al., 2014), selecting 500 images and posing six questions per image for each POPE split. We use the average F1 score of the three split sets as the evaluation metric.

MME (Fu et al., 2024) In line with prior methods (Leng et al., 2024; Chuang et al., 2023; Huo et al., 2024), we employ the hallucination subset of MME to systematically assess the model’s performance. This subset enables a detailed evaluation of both object-level and attribute-level hallucinations. For object-level hallucination, we employ the existence and count subsets, which enable the evaluation of a model’s ability to correctly identify the presence and quantity of objects within an image. Meanwhile, attribute-level hallucinations are examined using the position and color subsets, allowing us to assess how well MLLMs capture spatial relationships and color attributes.

MMBench (Liu et al., 2024f) To assess whether the overall capabilities of MLLMs are well preserved, we employ MMBench, a widely recognized benchmark specifically designed for evaluating the comprehensive abilities of MLLMs. MMBench serves as an objective and standardized evaluation framework, enabling a rigorous assessment of model performance across diverse multimodal tasks. It encompasses a meticulously curated set of over 3,000 multiple-choice questions, systematically categorized into 20 distinct ability dimensions, including but not limited to object localiza-

tion, social reasoning, spatial understanding, and commonsense inference. By covering a broad spectrum of cognitive and perceptual skills, MMBench provides a holistic measure of an MLLM’s capacity, ensuring a robust and fair comparison across different models.

B More Experimental Results

B.1 Comparison on General Benchmarks

Table 4 illustrates the performance comparison between our proposed MFP method and the baseline model across 9 general benchmarks. Our evaluation is conducted on a diverse set of benchmarks, including TextVQA (Singh et al., 2019), VizWiz (Gurari et al., 2018), MMBench (English & Chinese), MME, MM-Vet (Yu et al., 2024c), ScienceQA (Lu et al., 2022), SEED-Image (Li et al., 2023a), and VQAv2 (Goyal et al., 2017), covering various aspects of multimodal understanding and reasoning.

Overall, MFP achieves a comparable performance to the baseline, with an average score of 61.2, slightly surpassing the baseline’s 61.0. While MFP demonstrates improvements on certain benchmarks, such as MMBench (EN & CN) and ScienceQA, it exhibits minor performance drops on others, including VQAv2 and MME. According to the evaluation results, our model maintains the same general capability as the baseline.

B.2 Evaluation Results on More Hallucination Benchmarks

GPT-4o Assistant Evaluation In keeping with PAI, we also choose to use GPT-assisted evaluation. Since GPT-4V has been removed from the market, we choose the more advanced GPT-4o to evaluate our model. We construct prompts and input both the images and the description responses from two assistants into GPT-4o as follow:

{image}

You are required to score the performance of two AI assistants in describing a given image. You should pay extra attention to the hallucination, which refers to the part of descriptions that are inconsistent with the image content, such as claiming the existence of something not present in the image or describing incorrectly in terms of the counts, positions, or colors of objects in the image. Please rate the responses of the assistants

on a scale of 1 to 10, where a higher score indicates better performance, according to the following criteria:

1: Accuracy: whether the response is accurate with respect to the image content. Responses with fewer hallucinations should be given higher scores.

2: Detailedness: whether the response is rich in necessary details. Note that hallucinated descriptions should not count as necessary details.

Please output the scores for each criterion, containing only two values indicating the scores for Assistant 1 and 2, respectively. The two scores are separated by a space. Following the scores, please provide an explanation of your evaluation, avoiding any potential bias and ensuring that the order in which the responses were presented does not affect your judgment.

[Assistant 1]

{Response of Assistant 1}

[Assistant 2]

{Response of Assistant 2}

Output format:

Accuracy: {Scores}

Reason:

Detailedness: {Scores}

Reason:

The evaluation considers two key dimensions: Accuracy and Detailedness. Our evaluation Settings are consistent with PAI, and the results are shown in Table 5. Compared with the baseline and VCD, our method has better accuracy and detailedness.

Evaluation Results on MMHal-Bench We adopt MMHal-Bench (Sun et al., 2023a), a benchmark specifically designed to assess hallucination in multimodal large language models. MMHal-Bench comprises 96 image-question pairs, structured across 8 question categories and 12 object topics, encompassing a diverse set of visual and reasoning challenges. The question types include object attributes, adversarial object presence, object comparisons, counting, spatial relations, environmental context, holistic scene descriptions, and other complex reasoning forms. Since the original evaluation was conducted using GPT-4-0314, which is no

Method	VQA ^T	VizWiz	MMB ^{EN}	MMB ^{CN}	MME	MM-Vet	SQA ^I	SEED ^I	VQA ^{v2}	Overall
Baseline	58.6	50.0	64.3	58.3	75.5	30.5	66.8	66.1	78.5	61.0
MFP	56.5	53.4	68.2	59.0	73.3	30.4	67.7	65.1	76.9	61.2

Table 4: Result of General Benchmarks. The best performance is **bolded**.

Method	Accuracy \uparrow	Detailedness \uparrow
Baseline	5.38	5.88
VCD	5.83	5.93
MFP	6.29	6.52

Table 5: Result of GPT-4o Assistant Evaluation. The best performance is **bolded**.

Method	Avg. Score \uparrow	Hal. Rate \downarrow
Baseline	1.55	0.74
MFP	1.82	0.68

Table 6: Result of MMHal-Bench Evaluation. The best performance is **bolded**.

longer accessible, we employ the latest available version, GPT-4o, for scoring the model responses.

Table 6 presents the evaluation results on MMHal-Bench. Our proposed method, MFP, achieves a higher average score of 1.82 compared to the baseline’s 1.55, indicating improved response quality across diverse hallucination-related question types. Furthermore, MFP significantly reduces the hallucination rate from 0.74 to 0.68, demonstrating its effectiveness in mitigating object hallucinations under complex and high-difficulty scenarios. These results further validate the effectiveness of MFP in enhancing model performance and mitigating hallucinations.

B.3 Evaluation Results on More Models

We additionally evaluate MFP on the InternVL2 and Qwen2.5-VL model. Considering that MFP incorporates additional modules, both the baseline and MFP are fine-tuned on the same post-SFT

Model	V. E.	LLM	Res.	POPE	CHAIR \downarrow	CHAIR \downarrow
InternVL2 ^{1B}	InternViT	qwen2.5 ^{0.5B}	anyres	87.1	45.6	12.0
w/ MFP	InternViT	qwen2.5 ^{0.5B}	anyres	88.1	42.2	10.9
QwenVL2.5 ^{3B}	ViT	qwen2.5 ^{3B}	anyres	84.3	42.4	14.6
w/ MFP	ViT	qwen2.5 ^{3B}	anyres	83.4	30.8	9.9

Table 7: Results of MFP cross different models. From top to bottom we set γ to 0.23 and 0.33.

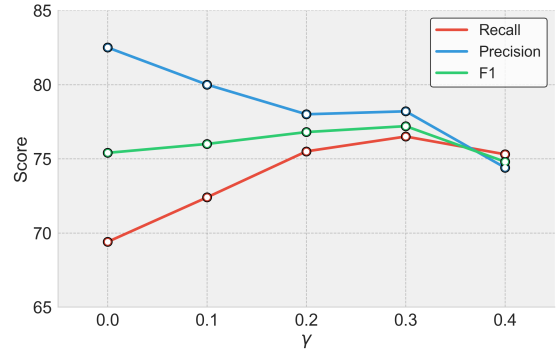


Figure 6: Results of sensitivity analysis on CHAIR benchmark for the parameter γ . The experiments are conducted on the LLaVA-1.5-7B model.

dataset, LLaVA-665K, to ensure a fair and controlled comparison. The results are shown in Table 7.

According to the results on POPE and CHAIR benchmarks, the effectiveness of MFP on different MLLM architectures is further demonstrated.

B.4 Efficiency Comparison Between MFP and Baseline

Considering that MFP introduces additional components, we compare its training and inference time with that of the baseline to assess the impact on computational efficiency.

For training time, we report the average time per step across 100 steps after the training process has stabilized. Additionally, we measure the time spent specifically on image processing (i.e., frequency filtering and image encoding) during each of these steps. For inference time, we compute the average time per sample based on evaluation conducted on the CHAIR benchmark.

B.5 Sensitivity Analysis

We choose the parameter γ at inference time for sensitivity analysis, using recall, precision, and F1 score from the CHAIR benchmark as evaluation metrics. A lower γ means fewer low- and high-frequency features are injected, and vice versa. As shown in Figure 6, when γ decreases from 0.4 to 0.3, recall, precision, and F1 score all in-

Method	Stage	POPE	CHAIR _s ↓	CHAIR _i ↓
Baseline	-	85.9	50.2	15.0
MFP	-	86.2	41.2	11.7
w/o low	Training	86.7	49.8	14.8
w/o high	Training	86.5	51.4	14.4
w/o γ	Inference	86.5	51.2	14.3

Table 8: Result of ablating different parts of MFP. The best performance is **bolded**. $\gamma=0.23$ is set for all experiments. The experiments are conducted on the LLaVA1.5-7B model.

crease, reaching their peak values of 76.5, 78.2, and 77.2, respectively. This improvement suggests that our method effectively suppresses redundant high- and low-frequency features, thereby enhancing the model’s object detection capacity. However, when γ is further reduced from 0.3 to 0.0, recall exhibits a decreasing trend (from 76.5 to 69.4), while precision increases (from 78.2 to 82.5). The F1 score initially improves slightly at $\gamma = 0.2$ before declining from 77.2 to 75.4 at $\gamma = 0.0$. This indicates that while reducing multi-frequency features alleviates object hallucination and improves precision, it may also lead to the loss of essential frequency-domain information, negatively impacting object detection and reducing recall. The observed trade-off between recall and precision highlights the importance of selecting an optimal γ to balance feature suppression and retention. This result supports the conclusion that the core mechanism of our proposed method is to reduce redundant frequency-domain features, thereby improving overall model robustness.

B.6 Ablation Study

To further investigate the effectiveness of our proposed Multi-Frequency Processing (MFP) method, we conduct an ablation study by removing different frequency features and the inference-time parameter γ , as shown in Table 8. Compared to the baseline, MFP significantly reduces CHAIR_s and CHAIR_i scores from 50.2 to 41.2 and from 15.0 to 11.7, respectively, demonstrating its effectiveness in mitigating object hallucinations. When low-frequency features are removed during training (w/o low), CHAIR_s and CHAIR_i increase to 49.8 and 14.8, respectively, suggesting that low-frequency features play a crucial role in suppressing spurious detections. Removing low-frequency features (w/o low) or high-frequency features (w/o

high) degrades performance, indicating that both components are critical—though low-frequency features contribute slightly more. Furthermore, when γ is removed at inference time (w/o γ), CHAIR_s and CHAIR_i degrade to 51.2 and 14.3, respectively, underscoring the importance of decaying frequency features during inference. Importantly, across all ablation settings, the POPE score remains consistently higher than the baseline, ranging from 86.2 to 86.7, indicating that our method robustly increases POPE score regardless of specific frequency feature being removed. These results validate that while both low- and high-frequency features contribute to hallucination mitigating, their joint utilization within MFP leads to the most effective mitigating.

Method	Train (s / step)		Test (s / item)
	Img. Pro.	Total	Total
Baseline	0.1	3.4	2.1
MFP	0.3	4.0	2.1

Table 9: Efficiency comparison between MFP and the baseline.

C More Training Parameters Details

Hyperparameter	PT	SFT
batch size	256	128
lr	1e-3	2e-5
lr schedule	cosine decay	
lr warmup ratio	0.03	
weight decay	0	
epoch	1	
optimizer	AdamW	
DeepSpeed stage	2	3

Table 10: Hyperparameters of our model’s pretrain and finetune.

For main results, we keep all training hyperparameters roughly the same as the LLaVA series. Table 10 presents a detailed training setup during PT and SFT stage.

Table 9 presents the training and inference efficiency comparison between MFP and the baseline. During training, MFP introduces a slight increase in image processing time (0.3s vs. 0.1s) due to the additional frequency processing. As a result, the

total training time per step rises modestly from 3.4s to 4.0s. However, this overhead remains relatively small considering the performance gains achieved. At inference time, the average duration per item remains almost unchanged (2.1s for both), indicating that MFP introduces negligible overhead during inference. These results confirm that MFP offers enhanced hallucination mitigation with minimal impact on computational efficiency.

D Case Study

Figure 7 presents a case-by-case comparison between our proposed MFP method and the original output. The results demonstrate that our method significantly mitigates hallucinations.



Figure 7: Comparison between our proposed MFP method and the original output in some cases. The hallucinating responses are highlighted in red.