

# ALIGNING MODALITIES IN VISION LARGE LANGUAGE MODELS VIA PREFERENCE FINE-TUNING

**Anonymous authors**

Paper under double-blind review

## 1 INTRODUCTION

Vision Large Language Models (VLLMs) have achieved significant success in various vision understanding tasks, such as image captioning Vinyals et al. (2015); Li et al. (2022; 2023c) and vision question answering Ye et al. (2023); Antol et al. (2015). These VLLM models fuse larger-scale pre-trained vision models into the representation space of a large language models (LLM), allowing the LLM access to the visual representations. However, such VLLMs are not perfect and even suffer from “hallucinations”, a phenomenon in which the language model generates content that is not grounded in the image, such as imagined objects and even scenes, wrong spatial relationships or categories, etc. Such artifacts are present even when both the vision backbone produces high-quality visual features and the language model itself is factual and accurate. These issues can pose significant risks when VLLMs are deployed in high-stakes scenarios, such as medical domains Li et al. (2023b) or autonomous driving Dewangan et al. (2023).

As discussed by Cui et al. (2023), the potential reason for hallucinations in VLLMs lies in their tendency to prioritize common sense or stereotypes present in the training language data, often disregarding the actual visual input information. In this paper, we attribute this issue to the lack of alignment between the image and text modalities, resulting in a reduced focus on input image information. Recent research efforts have sought to enhance the alignment between modalities through preference fine-tuning techniques, such as reinforcement learning from human feedback (RLHF) Sun et al. (2023). Concurrent works Li et al. (2023d); Zhao et al. (2023) also use the Direct Preference Optimization (DPO) framework, but they rely on the traditional preference data generation process in LLMs, where both preferred and dispreferred responses may potentially be incorrect. However, in VLLMs, the produced responses are centered around the image data rather than being generated freely like in LLMs. When comparing two responses, both of which may be incorrect for the given task, the model may struggle to accurately align the image with the correct generated response. In Yu et al. (2023a) the authors propose to solve this issue by collection corrective feedback, which shows strong results, but relies on costly human data gathering.

Unlike prior works that generate both preferred and dispreferred data, we propose **Preference Optimization in VLLM with AI-Generated Dispreferences (POVID)** framework, aiming to exclusively generate dispreferred feedback data using AI models. In POVID we employ a high-quality ground truth multi-modal instruction as the preferred answer and employ two strategies to generate dispreferred responses. *First*, we utilize GPT-4V to introduce plausible hallucinations into the answer, which we then use as the dispreferred response. *Second*, we aim to provoke inherent hallucination patterns and subsequently correct them within the target VLLM that requires fine-tuning. We achieve this goal by introducing noise, triggering inherent hallucination patterns within the VLLMs. The introduction of noise disrupts the VLLM’s comprehension of the image, leading it to generate uncertain responses that rely more on textual context or the knowledge it has acquired from the training data. Given that the inherent hallucination patterns of the target VLLM evolve during the training process, the response generation with the noisy image occurs in real-time during training, and this is treated as dispreference. Finally, we integrate both forms of dispreference into the DPO optimization framework, specifically targeting the alignment of language generation with the image.

The primary contribution of this paper is POVID, which utilizes AI-generated dispreference to align the image and text modalities in VLLMs. This approach explicitly contrasts a hallucinatory answer with a truthful one, eliminating the need for gathering human feedback and making it easily deployable at scale. Our empirical results demonstrate the promise of our framework in reducing hallucinations and enhancing other VLLM-related tasks. In particular, we visualized the experimental

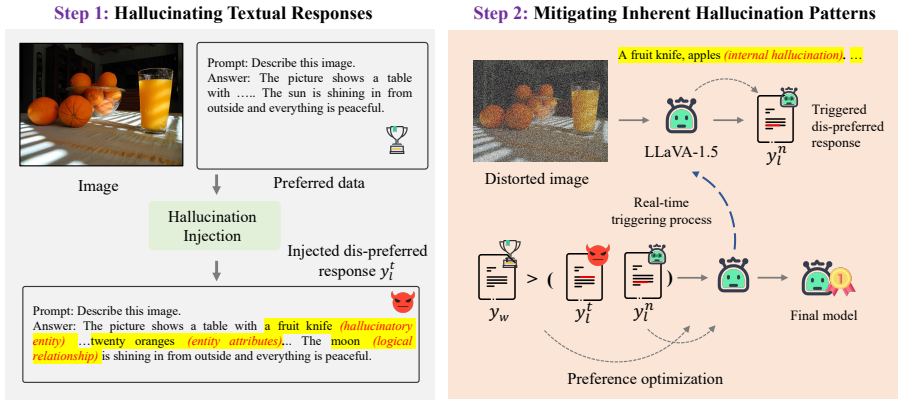


Figure 1: The framework of POVID. The preference generation process is divided into two steps: hallucinating textual responses and trigger dispreference during training. Here, different types of triggered hallucinations are labeled in (*types of hallucinations*).

results in the appendix as shown in Figure 4, our approach significantly improves performance compared to other preference tuning methods in VLLMs, achieving an average improvement of 12.4% improvements on average. Additionally, we demonstrate that POVID can redirect the attention of VLLMs towards the image modality, resulting in better modality alignment.

## 2 CONSTRUCTING PREFERENCES TO ALIGNING MODALITIES IN VLLMS

While preference learning approaches (e.g., DPO) facilitate the lightweight and stable training of VLLMs, they require data in the form of preferences. In contrast to LLMs, which support more freestyle generation in many scenarios, VLLMs used in various applications, such as VQA or image captioning, produce responses linked to input images. This inherent image-centricity presents distinct challenges in the preference data generation process for VLLMs, setting it apart from the process in LLMs. Specifically, in VLLMs, when comparing two responses, neither of which is correct for the required task (e.g., image captioning), the model may not be able to accurately align the image with the response.

To address this challenge, we propose **Preference Optimization in VLLM with AI-Generated Dispreferences (POVID)**, a novel approach aimed at better aligning image and text modalities. As illustrated in Figure 1, POVID leverages AI models to generate dispreferred responses without the need for human labeling efforts. These generated dispreferred responses, when combined with groundtruth image descriptions (treated as preferred responses), form the preference data pairs. Specifically, we employ two strategies to generate the dispreferred response: (1) Firstly, we manipulate the groundtruth text response by transforming the groundtruth response into hallucinated response, which serves as the dispreferred response; (2) Secondly, we introduce distortion to the image input during the training process, intending to trigger inherent hallucination patterns within the VLLMs. These patterns are then formalized as the dispreferred response, motivating the model to correct its inherent dispreferred patterns. In the remainder of this section, we will provide detailed explanations of both strategies and demonstrate how to integrate them into the preference training framework.

### 2.1 HALLUCINATING TEXTUAL RESPONSES

In our first strategy, we aim to generate dispreferred hallucinatory responses by hallucinating the groundtruth correct response. We construct the hallucinatory response based on a subset with 17K examples that are randomly sampled from LLaVA-Instruct-150K Liu et al. (2023b) dataset. Here, the LLaVA-Instruct-150K datasets is used to train LLaVA LLaVA with supervised fine-tuning. The 17K examples includes various task types, including image captioning, simple VQA and logical reasoning.

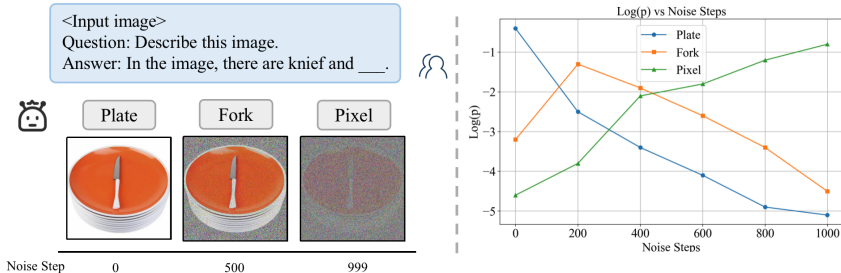


Figure 2: Illustration of logits for the next token generation with “In the image, there are knife and \_”. This figure shows the predictive uncertainty in token generation, emphasizing the influence of visual cues from objects identified as “knife” and “plate”.

To construct the preferences, we treat the original answers in the 17K examples as preferred responses. In terms of constructing dispreferred responses, we hallucinate the original answers using GPT-4V OpenAI (2023). Here, we adopt two hallucinating approaches tailored to different tasks:

**I. Hallucinating Image Captioning Tasks.** First, we hallucinate the image captioning tasks by considering three fundamental causes of hallucination in VLLMs: (1) *Object Co-occurrence*: This phenomenon arises when the training data contains spurious co-occurring patterns between objects, leading VLLMs to generate objects based on these learned spurious correlations. In this context, we aim to leverage GPT-4V to deduce object co-occurrence within the given image and subsequently revise the original responses accordingly; (2) *Logical Relationships Between Entities*: This involves using GPT-4V to modify the relationships between the original objects; (3) *Incorrect Attributes*: In this case, we employ GPT-4V to alter the attributes of various objects, such as changing their colors. We illustrate these three distinct hallucination scenarios with an example provided in Figure 5(a) in appendix. In addition, the prompt we used to generate the dispreferred response is in Appendix D.

**II. Hallucinating Reasoning Tasks.** Secondly, when dealing with tasks involving reasoning, such as VQA and logical reasoning, we task GPT-4V with modifying the reasoning process. This entails introducing errors related to logical relationships, entity information, entity attributes, and more. Additionally, we recommend that GPT-4V attempts to make subtle changes to the reasoning process, ensuring it remains independent of factual reasoning results, meaning that an incorrect reasoning process may still yield correct results. However, if the introduction of errors necessitates alterations to the reasoning results, we instruct GPT-4V to adjust the results accordingly. Likewise, in Figure 5(b) in appendix, we provide an example to demonstrate both the original and the generated dispreferred responses. The prompt we used is detailed in Appendix D.

## 2.2 MITIGATING INHERENT HALLUCINATION PATTERNS

In addition to generating the dispreferred response using powerful external AI models like GPT-4V, we also aim to provoke inherent hallucination patterns within the VLLM to be finetuned. Our second strategy introduces noise into the image to trigger inherent hallucination patterns in the VLLMs. This noise disrupts the VLLM’s understanding of the image, leading it to produce uncertain responses that rely more on textual context or acquired knowledge from the training data. This occurs because, in the presence of noisy images, the model tends to prioritize inherent object associations over visual information. Notably, the noise step should remain within a reasonable range, ensuring that the image remains easily recognizable by humans. For example, as depicted in Figure 2, when presented with the context “There are a knife and \_”, under specific noisy conditions, the likelihood of “fork” surpasses that of “plate” (ground truth). This may occur because “plate” is more likely to co-occur with “fork” in the training data. With an increase in noise steps, the term “pixel” becomes predominant, owing to the noticeable noise patterns within the image. Consequently, establishing an appropriate noise step to activate inherent hallucination patterns is a reasonable approach.

To achieve this goal, we introduce diffusion noise into the original image. We define the noise step as  $k$ , and the noised image with step  $k$  can be expressed as follows:

$$x(k) = \sqrt{\xi_k} \cdot x + \sqrt{1 - \xi_k} \cdot \epsilon, \tag{1}$$

Table 1: Comparison between POVID and other preferences construction approaches in both hallucination and comprehensive evaluation benchmarks. We bold the best results and underline the second best results.

Method	Hallucination Benchmark				Comprehensive Benchmark			
	CHAIR <sub>S</sub> ↓	CHAIR <sub>i</sub> ↓	POPE ↑	MMHal ↑	SciQA-IMG ↑	MM-Vet ↑	MMBench ↑	LLaVA-Bench ↑
LLaVA-1.5	66.8	12.7	85.90	2.42	66.8	30.5	63.0	63.4
+ Vfeedback	56.3	11.4	83.72	<u>2.62</u>	66.2	<u>31.2</u>	<u>63.9</u>	62.1
+ Human-Preference	54.0	9.3	81.50	2.53	65.8	31.1	60.4	63.7
+ RLHF-V	<u>44.6</u>	<u>7.9</u>	<u>86.20</u>	2.59	<u>67.1</u>	30.9	63.6	<u>65.4</u>
<b>POVID (ours)</b>	<b>31.8</b>	<b>5.4</b>	<b>86.90</b>	<b>2.69</b>	<b>68.8</b>	<b>31.8</b>	<b>64.9</b>	<b>68.7</b>

where  $\bar{\xi}_t = \prod_{i=0}^k \xi_i$  and  $\xi_k \in (0, 1)$  is a hyperparameter chosen prior to model training. Detailed settings can be found in Appendix C. After obtaining the noised image, in order to more effectively capture changes in inherent hallucination patterns during the fine-tuning process of the VLLM, we integrate the image noising process into the DPO fine-tuning process. Specifically, for each input prompt  $x$ , we take into account the dispreferred responses from both the hallucinated text responses discussed in Section 2.1 and the responses triggered by distorted images. We then reformulate the DPO loss as follows:

$$\mathcal{L}_{POVID}(\pi_{\theta}; \pi_{\text{ref}}) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[ \log \sigma \left( \alpha \log \frac{\pi_{\theta}(y_w|x)}{\pi_{\text{ref}}(y_w|x)} - \left( \beta_1 \log \frac{\pi_{\theta}(y_l^t|x)}{\pi_{\text{ref}}(y_l^t|x)} + \beta_2 \log \frac{\pi_{\theta}(y_l^n|x^n)}{\pi_{\text{ref}}(y_l^n|x^n)} \right) \right) \right], \quad (2)$$

where  $\alpha$ ,  $\beta_1$  and  $\beta_2$  are coefficients that balance preferred and dispreferred terms.  $y_l^g$  represents the dispreferred response generated using the approach outlined in Section 2.1. Additionally,  $x^n$  represents the noisy image, which triggers the generation of the dispreferred response  $y_l^n$ . It’s important to note that for each token  $i$  in the sequence  $y_l^n$ , the value of  $y_{l,i}^n$  is determined by selecting the maximum probability from the set  $\pi_{\theta}(\cdot | x^n, y_{w, < i})$ . Here, each generated token in the dispreferred response  $y_{l,i}^n$  is conditioned on the prior tokens from the preferred response  $y_{w, < i}$ . This conditioning allows us to control the reliability of the triggered dispreferred response. As a result, we aim to capture the most significant changes between the preferred and dispreferred responses, since a substantial portion of dispreferred response overlaps with preferred response.

### 3 EXPERIMENT

In this section, we empirically investigate the effectiveness of POVID in aligning image and text modalities in VLLMs and reducing hallucination. We aim to answer the following questions: (1) Can POVID effectively reduce hallucination in VLLMs compared to other preference fine-tuning strategies? (2) Does POVID improve performance compared to other benchmarks and tasks like VQA? (3) Can hallucinating textual responses and image distortion benefit performance? (4) How does POVID change attention weights to align image and text modalities? For the experimental setting and comparison model and benchmark, we have introduced it in detail in Appendix C, and we have also supplemented the analytical experiments such as modality alignment and benchmark fine-grained in Appendix B.

#### 3.1 RESULTS

**Comparison with Different Preferences in VLLMs.** In Table 1, we present the results of a comparison between various VLLM preferences, evaluating both hallucination and comprehensive benchmarks. Firstly, in the hallucination benchmarks, POVID effectively enhances performance by creating dispreferred preferences through textual data manipulation and image distortion. We achieve a significant improvement of 31.78% across all hallucination benchmarks, effectively reducing hallucinations in the generated responses. This outcome aligns with our expectations, as constructing dispreferences from the ground-truth correct responses maximally enables the model to discern differences between correct and incorrect responses while optimizing alignment between the image

Table 2: Comparison between POVID and other state-of-the-art VLLMs across both hallucination and comprehensive evaluation benchmarks. We bold the best results and underline the second best results.

Method	Vision Encoder	Language Model	CHAIR <sub>S</sub> ↓	CHAIR <sub>i</sub> ↓	POPE ↑	MMHal ↑	Avg. ranking ↓
InstructBLIP	ViT-g (1.3B)	Vicuna (7B)	<u>40.0</u>	<u>8.0</u>	77.83	2.10	3.00
Qwen-VL-Chat	ViT-G (1.9B)	Qwen (7B)	48.2	9.1	<b>87.07</b>	<b>2.89</b>	<u>2.50</u>
mPLUG-Owl2	ViT-L (0.3B)	LLaMA (7B)	54.4	12.0	86.20	2.17	4.00
<b>POVID (ours)</b>	ViT-L (0.3B)	Vicuna (7B)	<b>31.8</b>	<b>5.4</b>	<u>86.29</u>	<u>2.69</u>	<b>1.50</b>

Method	Vision Encoder	Language Model	SciQA-IMG ↑	MM-Vet ↑	MMBench ↑	LLaVA-Bench ↑	Avg. ranking ↓
InstructBLIP	ViT-g (1.3B)	Vicuna (7B)	60.5	26.2	36.0	60.9	4.00
Qwen-VL-Chat	ViT-G (1.9B)	Qwen (7B)	<u>68.2</u>	<b>41.2</b>	60.6	<u>67.7</u>	<u>2.25</u>
mPLUG-Owl2	ViT-L (0.3B)	LLaMA (7B)	64.5	<u>36.2</u>	<u>64.5</u>	59.9	3.00
<b>POVID (ours)</b>	ViT-L (0.3B)	Vicuna (7B)	<b>68.8</b>	31.8	<b>64.9</b>	<b>68.7</b>	<b>1.75</b>

Table 3: Results of ablation study. Text disprefer means we only using hallucinated textual response to train DPO. Image distortion means that we use distorted images to trigger inherent hallucination patterns.

		Hallucination Benchmarks				Comprehensive Benchmarks			
Text disprefer	Image distortion	CHAIR <sub>S</sub> ↓	CHAIR <sub>i</sub> ↓	POPE ↑	MMHal ↑	SciQA-IMG ↑	MM-Vet ↑	MMBench ↑	LLaVA-Bench ↑
×	×	66.8	12.7	85.90	2.42	66.2	<u>31.2</u>	63.9	62.1
✓	×	<u>39.6</u>	<u>6.3</u>	<u>86.04</u>	<u>2.65</u>	<u>67.2</u>	30.9	<u>64.7</u>	<u>67.5</u>
×	✓	50.4	9.6	85.19	2.54	66.9	30.7	64.3	66.9
✓	✓	<b>31.8</b>	<b>5.4</b>	<b>86.90</b>	<b>2.69</b>	<b>68.8</b>	<b>31.8</b>	<b>64.9</b>	<b>68.7</b>

and text modalities within the model. Moreover, in more comprehensive evaluation benchmarks, which encompass not only factuality and hallucination assessment but also other aspects, POVID continues to demonstrate superior performance when compared to other preference data collection methods. This further indicates our model’s capacity to enhance VLLM performance through improved modality alignment.

**Comparison with Open-Sourced VLLMs Models.** We present a comparison between POVID and other open-sourced VLLMs in Table 2. Although various approaches utilize different image and text encoders, POVID outperforms other popular VLLMs in five out of eight benchmarks. In contrast, the second-best baseline, Qwen-VL-Chat, achieves the best performance in only two out of eight benchmarks. This underscores the superiority of POVID and further corroborates its effectiveness in aligning image and text modalities to improve the performance of VLLMs.

**Ablation Studies.** To further demonstrate the essential role of the key components of POVID in contributing to performance improvement, we conducted ablation experiments on both hallucination and comprehensive benchmarks, and present the results in Table 3. In this ablation study, we evaluate the effectiveness of two aspects: (1) hallucinating groundtruth responses and (2) image distortion. According to the results, we initially observe that image distortion can enhance performance across all benchmarks. This indicates its effectiveness in aligning multimodalities by compelling the model to rectify inherent hallucination patterns. Additionally, generating dispreference from groundtruth responses significantly enhances performance, underscoring the effectiveness of the AI-generated dispreference strategy. Finally, when combining both strategies, POVID achieves the best performance, further affirming its effectiveness in enhancing VLLMs through improved modality alignment.

## 4 CONCLUSION

In this work, we introduce a novel approach, Preference Optimization in VLLM with AI-Generated Dispreferences (POVID) to address the challenges in modality alignment for large vision-language models. In POVID, we adopt two strategies to generate dispreferred responses: first, we use synthetic data from GPT-4V to inject plausible hallucinations into the correct answer. Second, we use distorted images to trigger the inherent hallucination behavior of the VLLM. Then both of these answers are integrated into an RLHF framework via Direct Preference Optimization. Empirical evaluations across multiple benchmarks reveal that POVID not only mitigates hallucination effectively but boosts the overall performance of model.

## REFERENCES

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pp. 2425–2433, 2015.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- Junbum Cha, Wooyoung Kang, Jonghwan Mun, and Byungseok Roh. Honeybee: Locality-enhanced projector for multimodal llm. *arXiv preprint arXiv:2312.06742*, 2023.
- Gongwei Chen, Leyang Shen, Rui Shao, Xiang Deng, and Liqiang Nie. Lion: Empowering multi-modal large language model with dual-level visual knowledge. *arXiv preprint arXiv:2311.11860*, 2023.
- Ching-Yao Chuang, Varun Jampani, Yuanzhen Li, Antonio Torralba, and Stefanie Jegelka. Debiasing vision-language models via biased prompts. *arXiv preprint arXiv:2302.00070*, 2023.
- Chenhang Cui, Yiyang Zhou, Xinyu Yang, Shirley Wu, Linjun Zhang, James Zou, and Huaxiu Yao. Holistic analysis of hallucination in gpt-4v (ision): Bias and interference challenges. *arXiv preprint arXiv:2311.03287*, 2023.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023.
- Vikrant Dewangan, Tushar Choudhary, Shivam Chandhok, Shubham Priyadarshan, Anushka Jain, Arun K Singh, Siddharth Srivastava, Krishna Murthy Jatavallabhula, and K Madhava Krishna. Talk2bev: Language-enhanced bird’s-eye view maps for autonomous driving. *arXiv preprint arXiv:2310.02251*, 2023.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023.
- Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming Zhang, and Nenghai Yu. Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation. *arXiv preprint arXiv:2311.17911*, 2023.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pp. 4904–4916. PMLR, 2021.
- Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*, 2023.
- Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. *arXiv preprint arXiv:2311.16922*, 2023.
- Chen Li, Yixiao Ge, Dian Li, and Ying Shan. Vision-language instruction tuning: A review and analysis. 2023a.
- Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *arXiv preprint arXiv:2306.00890*, 2023b.

- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023c.
- Lei Li, Zhihui Xie, Mukai Li, Shunian Chen, Peiyi Wang, Liang Chen, Yazheng Yang, Benyou Wang, and Lingpeng Kong. Silkie: Preference distillation for large visual language models. *arXiv preprint arXiv:2312.10665*, 2023d.
- Lei Li, Yuwei Yin, Shicheng Li, Liang Chen, Peiyi Wang, Shuhuai Ren, Mukai Li, Yazheng Yang, Jingjing Xu, Xu Sun, et al. M<sup>3</sup>it: A large-scale dataset towards multi-modal multilingual instruction tuning. *arXiv preprint arXiv:2306.04387*, 2023e.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*, 2023f.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. Aligning large multi-modal model with robust instruction tuning. *arXiv preprint arXiv:2306.14565*, 2023a.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023b.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*, 2023c.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *The 36th Conference on Neural Information Processing Systems (NeurIPS)*, 2022.
- OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023. URL <https://arxiv.org/abs/2303.08774>.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://arxiv.org/abs/2305.18290>.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. Object hallucination in image captioning. *arXiv preprint arXiv:1809.02156*, 2018.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. Aligning large multimodal models with factually augmented rlhf. *arXiv preprint arXiv:2309.14525*, 2023.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca), 2023.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023a.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023b.
- Haoqin Tu, Chenhang Cui, Zijun Wang, Yiyang Zhou, Bingchen Zhao, Junlin Han, Wangchunshu Zhou, Huaxiu Yao, and Cihang Xie. How many unicorns are in this image? a safety evaluation benchmark for vision llms. *arXiv preprint arXiv:2311.16101*, 2023.

- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3156–3164, 2015.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023.
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Tong Xu, Hao Wang, Dianbo Sui, Yunhang Shen, Ke Li, Xing Sun, and Enhong Chen. Woodpecker: Hallucination correction for multimodal large language models. *arXiv preprint arXiv:2310.16045*, 2023.
- Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwen He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, and Tat-Seng Chua. Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. *arxiv*, 2023a.
- Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwen He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, et al. Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. *arXiv preprint arXiv:2312.00849*, 2023b.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*, 2023c.
- Zhiyuan Zhao, Bin Wang, Linke Ouyang, Xiaoyi Dong, Jiaqi Wang, and Conghui He. Beyond hallucinations: Enhancing lvlms through hallucination-aware direct preference optimization. *arXiv preprint arXiv:2311.16839*, 2023.
- Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. Analyzing and mitigating object hallucination in large vision-language models. *arXiv preprint arXiv:2310.00754*, 2023.

## A RELATED WORK

**VLLMs and VLLM Hallucination.** The advent of autoregressive large-scale language models (LLMs), highlighted in works by Touvron et al. (2023a;b); Taori et al. (2023), has led to the development of Vision-Large Language Models (VLLMs). To align the image and text modalities, recent research has concentrated on instruction tuning Li et al. (2023a), scaling up training dataset Jia et al. (2021), and better alignment between image and text with local feature enhancement Cha et al. (2023). These advancements have successfully combined LLMs with image inputs and excel in image comprehension.

However, such VLLMs are not perfect and even suffer from “hallucinations”, generating outputs that may not accurately or faithfully represent the content of a user-provided image. There are various sources of hallucinations in VLLMs, including biased data Chuang et al. (2023); Tu et al. (2023), insufficient training Chen et al. (2023), and imperfect inference Huang et al. (2023). Recently, addressing hallucination in LVLMS is primarily achieved through various techniques such as decoding approaches Leng et al. (2023); Huang et al. (2023), post-processing Zhou et al. (2023); Yin et al. (2023) and the construction of higher-quality dataset Liu et al. (2023a); Li et al. (2023e). While these approaches can mitigate hallucination to some extent, they often fail to directly guide VLLMs to align image and text modalities.

**Preference Alignment** Aligning with human preferences for large models has emerged as a critical issue due to the limitations imposed by safety and ethical considerations in real-world applications. Preference alignment can be broadly categorized into two main approaches: alignment through feedback, which encompasses both human Bai et al. (2022); Rafailov et al. (2023) and AI-generated feedback Lee et al. (2023) and alignment via prompt guidance Wei et al. (2022). Initial investigations



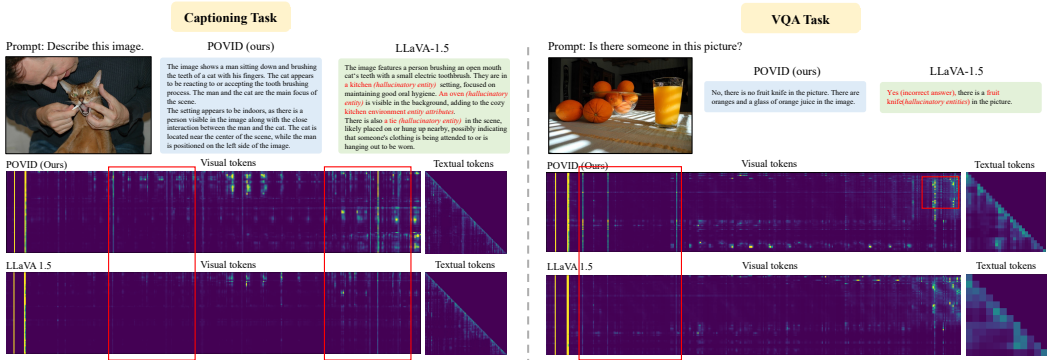


Figure 3: Comparison of attention map between POVID and LLaVA-1.5 at different tasks. The red box region is labeled with the image attentions that can be significantly improved by POVID.

into preference alignment for VLLMs have recently been conducted. Sun et al. (2023) introduced LLaVA-RLHF, which utilizes a preference dataset annotated by humans to decrease hallucinations in LLaVA. Li et al. (2023d) proposed a method for distilling preferences into VLLMs to enhance their ability to generate relevant and accurate responses based on visual context. Yu et al. (2023b) collected human preferences in the form of segment-level corrections to hallucinatory content and optimizing the model’s behavior based on dense, direct feedback. While these initial results are promising, these works heavily rely on the traditional preference data generation process in LLMs, which generate both preferred and dispreferred responses, but none of them are guaranteed to be correct. In VLLMs, when both responses prove incorrect for the given task, accurately aligning the image with the correct generated response becomes challenging. In contrast, POVID directly generates dispreferred responses, effectively addressing this challenge.

## B SUPPLEMENTARY EXPERIMENT & VISUALIZATION

### B.1 ANALYSIS

In this section, we provide a comprehensive analysis to demonstrate how different components contribute to the performance of POVID and illustrate how POVID enhances overall performance.

Table 4: Fine-grained performance comparison of various models on LLaVA-Bench, where we adopt the following abbreviation: Convo for Conversation, Captioning for Detail description, Reasoning for Complex reasoning.

Method	Convo	Captioning	Reasoning	Overall
LLaVA-1.5	53.3	53.4	79.6	63.4
+ Vlfedback	51.3	49.3	78.5	62.1
+ Human-Preference	49.6	43.3	81.3	63.7
+ RLHF-V	55.8	56.1	80.3	65.4
<b>POVID (ours)</b>	<b>55.9</b>	<b>60.1</b>	<b>81.5</b>	<b>68.7</b>

**Fine-grained Performance Analysis.** Table 4 presents a fine-grained performance analysis of different preference collection strategies on the LLaVA-Bench benchmark. This analysis encompasses a spectrum of multi-modal reasoning and perception dimensions, such as Conversation, Detail Description, and Complex Reasoning. According to Table 4, it is evident that, when compared with other preference data collection approaches, POVID excels in image captioning and providing detailed descriptions for a given image. This outcome aligns with our expectations, as our training data includes various long-form captions, and such comprehensive preference comparisons result in improved alignment and stronger image captioning results. Moreover, across other categories, our approach consistently outperforms alternative preference collection strategies, underscoring its effectiveness across various tasks.

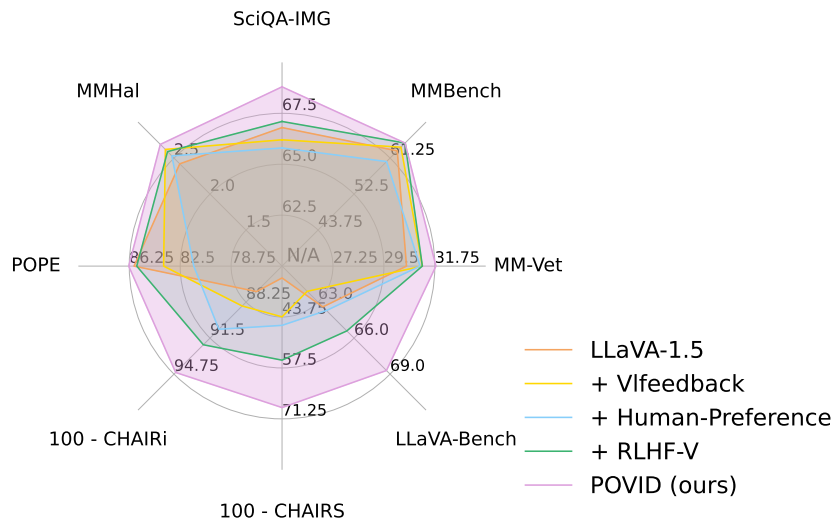


Figure 4: An overall performance comparison between POVID and other VLLM preference learning approaches.

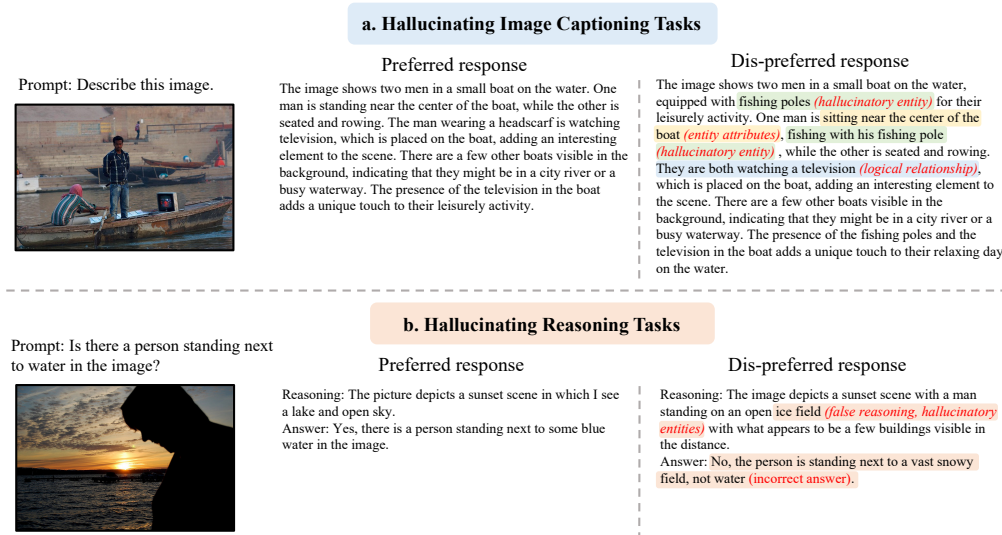


Figure 5: Two examples extracted from hallucinated image captioning tasks and reasoning tasks. Different types of hallucinations are labeled in (*types of hallucinations*).

**Modality Alignment Analysis.** We assess the impact of POVID on modality alignment by comparing the attention maps generated by POVID with those of the original LLaVA-1.5 model, with a specific focus on image captioning and VQA tasks. We illustrate two cases in Figure 3, where these attention maps reveal the distribution of attention scores assigned to generated textual tokens within the input image-text sequence throughout the VLLM’s output generation phase. Our findings reveal that the original LLaVA-1.5 model tends to overemphasize the context of the text, which can result in hallucinations. In contrast, POVID increasingly prioritizes attention towards the image, indicating a strong alignment between image and text modalities. One potential explanation for this phenomenon is that, through a comparison between the ground truth and the generated dispreferred data, along with the mitigation of internal hallucination patterns, POVID redirects the VLLM’s attention, leading to a greater focus on the image tokens.

## C TRAINING AND EXPERIMENTS SETUP

Training hyperparameters are shown in Table 5. For the first phase, we trained for 3 epochs, and for the second phase, the training was conducted for 1 epoch. Training for 20 hours on one A100 80G GPU. For the second phase, we adjust the diffusion noise level, symbolized by  $\xi$  through a specific formula:  $\xi = \text{Sigmoid}(l_t) \times (0.5 \times 10^{-2} - 10^{-5}) + 10^{-5}$ , where  $\epsilon$  is drawn from a normal distribution.

**Implementation Details.** Following concurrent VLLM preference tuning studies Yu et al. (2023b); Li et al. (2023d), we have chosen LLaVA-1.5 (7B) as our backbone model for all experiments and have applied POVID to fine-tune LLaVA-1.5 (7B). The overall training process is divided into two stages. In the first stage, we exclusively utilize the preferences generated through the hallucinating textual responses, as discussed in Section 2.1, to fine-tune LLaVA-1.5 using DPO. In the second stage, we employ image distortion to rectify the model’s inherent hallucinatory behaviors using the loss defined in Eq. 2. The first stage involves training for 3 epochs, and the second stage for 1 epoch. The entire training process requires a single A100 80GB GPU and takes approximately 6 hours.

**Baseline Approaches.** We first compare the proposed approach with other VLLM preference tuning methods, which include Silkie Li et al. (2023d), LLaVA-RLHF Sun et al. (2023), and RLHF-V Yu et al. (2023b). These methods share a common goal of enhancing model performance by creating curated datasets and subsequently applying preference tuning techniques to fine-tune the model based on these datasets. To ensure a fair and equitable comparison, we utilize the same curated datasets employed by these approaches and apply DPO to fine-tune LLaVA-1.5 (7B).

Furthermore, we compare the performance with other open source VLLMs, including InstructBLIP Dai et al. (2023), Qwen-VL-Chat Bai et al. (2023) and mPLUG-Owl2 Ye et al. (2023).

**Evaluation Benchmark.** To evaluate the performance of POVID and other baselines, we first adopt VLLM hallucination evaluation benchmarks, including CHAIR Rohrbach et al. (2018), POPE Li et al. (2023f), and MMHal Sun et al. (2023). Here, CHAIR, including CHAIR<sub>S</sub> and CHAIR<sub>I</sub>, is a metric used in image captioning tasks to evaluate the accuracy of object descriptions in captions. It compares the objects mentioned in a caption with those present in the image. MMHal Sun et al. (2023) assesses hallucinations and response informativeness by utilizing GPT-4V to compare model output with human responses and various object labels, determining the scores accordingly. POPE Li et al. (2023f) uses a set of binary classification tasks, prompting VLLMs with simple Yes-or-No questions about the existence of certain objects in images.

We further evaluate all approaches on comprehensive VLLM evaluation benchmarks, including SciQA-IMG Lu et al. (2022), MME Fu et al. (2023), MMbench Liu et al. (2023c), MM-Vet Yu et al. (2023c) and LLaVA-bench Liu et al. (2023b). Each benchmark contains tasks to evaluate perception, cognition, and reasoning abilities of VLLMs.

## D CONSTRUCTION OF THE DISPREFERENCE DATASET

This section details the prompts utilized to compile the dataset focusing on dispreferences, specifically within the realms of image captioning and reasoning tasks. The prompts are designed to elicit responses that reveal dispreference patterns, categorized into two main types: image captioning tasks intended to provoke imaginative descriptions, and reasoning tasks aimed at stimulating inferential thought processes. These prompts, central to our methodology, are enumerated in Table 6, offering a comprehensive view of the data generation framework.

Table 5: Training hyperparameters.

Hyperparameters	
<code>lora_r</code>	128
<code>lora_alpha</code>	256
<code>lora_target</code>	all
<code>mm_projector_lr</code>	2e-5
Batch size	1
Learning rate	1e-7
<code>model_max_length</code>	1024
<code>noise_step</code> (only for internal preference optimization)	500

Table 6: Two types of prompts to GPT4V (The format of the obtained data is {image, prefer data, disprefer data}).

---

**Prompts for hallucinating image captioning tasks:**

Help me generate one highly confusing response based on the image and the standard caption in the Question-Answer Pair.

\*\*\*\*\*

Question-answer Pair:

Q: {question}

A: {answer}

Requirements:

- (1) The generated caption is generally similar to the given A, with the same main meaning;
- (2) You can refer to the following errors to generate the wrong caption (1. The wrong caption can contain some co-occurring objects, which are prone to appear in such scenarios but do not appear in the image; 2. The wrong caption can be an error in the number of entities or the logical relationships between entities; 3. The attributes of entities in the caption can also be modified, such as color, appearance, etc.) (3) Compared to the original caption A, the caption you modified is incorrect based on the provided image.

\*\*\*\*\*

Output Format:

Answer: your answer

---

**Prompts for hallucinating reasoning tasks:**

Now, please help me generate new answers with hallucination errors based on the image, question, and answer provided. There are two cases now:

1. If the given question and answer are short and do not require logical reasoning, then modify the answer to a hallucination error answer, such as some quantity errors or entity and property errors.
2. If the entire question requires logical reasoning, then help me reorganize the answers based on the given image, questions, and answers into the format “Reason: xxx, Result: xxx” (Answer 1). Modify the reasons by introducing errors related to logical relationships, entity information, entity attributes, etc. If the error in the reason would lead to a new result, modify the result accordingly. If the error does not lead to a new result, keep the original result. Similarly, organize it in the format “Reason: xxx, Result: xxx” (Answer 2).

\*\*\*\*\*

Question-answer Pair:

Q: {question}

A: {answer}

Requirements:

- (1) The generated wrong answer and reasoning process should be combined with the image and be misleading..

\*\*\*\*\*

Output Format:

Answer: your answer

---