

# Mitigating Object Hallucinations in Large Vision-Language Models through Visual Contrastive Decoding

Sicong Leng<sup>1,2,\*</sup> Hang Zhang<sup>1,3,\*</sup> Guanzheng Chen<sup>1,3</sup> Xin Li<sup>1,3,†</sup>  
 Shijian Lu<sup>2</sup> Chunyan Miao<sup>2</sup> Lidong Bing<sup>1,3</sup>  
<sup>1</sup>DAMO Academy, Alibaba Group <sup>2</sup>Nanyang Technological University  
<sup>3</sup>Hupan Lab, 310023, Hangzhou, China

<https://github.com/DAMO-NLP-SG/VCD>

## Abstract

Large Vision-Language Models (LVLMs) have advanced considerably, intertwining visual recognition and language understanding to generate content that is not only coherent but also contextually attuned. Despite their success, LVLMs still suffer from the issue of object hallucinations, where models generate plausible yet incorrect outputs that include objects that do not exist in the images. To mitigate this issue, we introduce Visual Contrastive Decoding (VCD), a simple and training-free method that contrasts output distributions derived from original and distorted visual inputs. The proposed VCD effectively reduces the over-reliance on statistical bias and unimodal priors, two essential causes of object hallucinations. This adjustment ensures the generated content is closely grounded to visual inputs, resulting in contextually accurate outputs. Our experiments show that VCD, without either additional training or the usage of external tools, significantly mitigates the object hallucination issue across different LVLM families. Beyond mitigating object hallucinations, VCD also excels in general LVLM benchmarks, highlighting its wide-ranging applicability.

## 1. Introduction

Large Vision-Language Models (LVLMs) have become integral in the intersection of computer vision and natural language processing, enabling a range of applications due to their ability to generate contextually relevant textual descriptions from visual inputs. These models are characterized by their effectiveness in capturing and translating complex visual patterns into coherent linguistic representations [5, 12, 18, 33, 45, 49, 71, 73, 78]. The evolution of

\*Equal contribution. Sicong Leng is under the Joint PhD Program between DAMO Academy and Nanyang Technological University.

†Correspondence: xinting.lx@alibaba-inc.com.

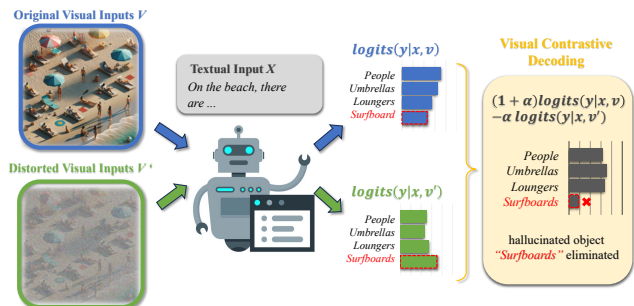


Figure 1. An illustration of Visual Contrastive Decoding. The hallucinated object “Surfboards” is highlighted in red, and it is eliminated during the generative process by contrasting with the output distribution that favors hallucinations.

LVLMs is marked by ongoing improvements in model architecture, training methodologies, and data diversity, leading to enhanced performance and application versatility. Despite these advancements, specific challenges persist, with the issue of object hallucination [20, 38, 43, 48] being a prominent concern that impacts the reliability and applicability of LVLMs across domains.

Object Hallucination in this context refers to the phenomenon where LVLMs generate textual content that is semantically coherent but inconsistent with ground-truth objects in the given image. This challenge not only reveals fundamental issues of LVLMs, such as over-reliance on statistical bias [1, 2, 19, 38] and unimodal priors [21, 22, 51, 68, 70, 75], but also has direct implications for the practical deployment of LVLMs. In applications where precision and reliability of generated content are paramount, object hallucinations can lead to misinformation, misinterpretation, and subsequent erroneous decision-making. In domains like healthcare [26, 66], autonomous systems [8, 69], and robotics [46, 50], such inaccuracies are not just undesirable but could have significant consequences. Addressing the hallucination issue is therefore essential to enhance the integrity,

reliability, and broad applicability of LVLMs in various real-world scenarios.

Various approaches have been explored to curb object hallucinations in VLMs. Early works made attempts on small-scale VLMs by either performing fine-grained modality alignment [6] or reducing the statistical bias of object co-occurrence with data augmentation [30, 54]. However, the behaviors of LVLMs differ significantly from small-scale VLMs, making related methods impractical to generalize and scale up [29, 67]. Several recent studies address this issue by proposing hallucination-targeted datasets for fine-tuning [20, 42], training a post-hoc revisor to reconstruct less hallucinatory outputs [77] or adapting factually augmented Reinforcement Learning from Human Feedback (RLHF) [60]. While existing interventions for object hallucination in LVLMs have shown effectiveness, the incurred human effort and computational cost highlight a pressing need for a simpler but efficient approach.

In this work, we analyze the effect of visual uncertainty on the two primary causes of object hallucinations in LVLMs, namely statistical bias and unimodal priors (i.e., language priors). Building on the analysis above, we introduce Visual Contrastive Decoding (VCD), a training-free technique designed to mitigate object hallucination in LVLMs. As shown in Figure 1, VCD is grounded in the principle of contrasting output distributions from original and distorted visual inputs. Hence, it acts as a corrective mechanism and calibrates the model’s over-reliance on language priors from integrated LLMs and statistical bias of LVLMs’ pretraining corpus. In the realm of efficiency, VCD stands out due to its minimal computational overhead compared with previous studies [20, 42, 60, 77], circumventing the need for additional training or the usage of external tools (e.g., other pretrained models). Our experiments demonstrate VCD’s effectiveness, with consistent improvements on multiple object hallucination benchmarks (e.g., up to +7.4 F1 score boost on POPE [38] and +18% improvement on MME [16]) across different LVLM families, including LLaVA-1.5 [44, 45], InstructBLIP [12], and Qwen-VL [5]. In addition, our method is also beneficial to the general perception capacities of LVLMs as evidenced by benchmarking on MME and LLaVA-Bench<sup>1</sup>, indicating its potential applicability beyond the scope of object hallucination mitigation.

To sum up, our main contributions are as follows:

1. We conduct an in-depth analysis of the effect of visual uncertainty on object hallucinations in LVLMs, particularly from the aspects of statistical bias and unimodal priors.
2. Inspired by the analysis above, we design VCD, a training-free technique that can effectively mitigate object hallucinations in LVLMs. It calibrates the model’s outputs by contrasting output distributions derived from

original and distorted visual inputs, ensuring more consistent content generation.

3. Through comprehensive experiments, we demonstrate the efficacy of the proposed VCD in alleviating object hallucination and enhancing general perception capability. Our method yields notable improvements without the need for additional training or external tools.

## 2. Related Work

### 2.1. Visual-Language Models

The development of Vision-Language Models (VLMs) has transitioned from being rooted in BERT-based language decoders [13, 31, 47] for merging visual and textual data [34, 36, 59, 64], to a notable advancement ushered by the integration of Large Language Models (LLMs) [4, 7, 10, 11, 17, 53, 61–63]. The advent of LLMs heralded the emergence of Large Vision-Language Models (LVLMs) [3, 9, 14, 34], characterized by enhanced capabilities and performance. In this phase, LVLMs, supported by end-to-end training techniques, demonstrated unified decoding of visual and textual tokens, marking a significant enhancement in their performances and adaptability. Recent developments have seen a focus on Visual Instruction Fine-tuning [45], showcasing adaptability to a variety of vision-language tasks. The methodologies adopted, ranging from integrating cross-modal alignment networks to fine-tuning LLaMA models, underscore a trend of diversification and specificity in the approach [5, 12, 33, 71].

### 2.2. Hallucination in VLMs

Prior to the advent of LLMs, the NLP community has primarily defined “hallucination” as the generation of nonsensical content or content that deviates from its sources [28, 32, 39, 57, 74, 76]. In the realm of VLMs, “object hallucination” is also well-documented, referring to models producing plausible outputs that include objects that do not match or are missing from images [6, 38, 54]. Mitigating object hallucination in VLMs has typically involved strategies such as fine-grained contrastive learning [72], ROI feature fusion [6], and the curtailment of co-occurrence patterns via data augmentation [30]. However, with the distinct training paradigms and model architectures that characterize traditional VLMs and contemporary LVLMs, adapting these strategies to the newer auto-regressive approaches in LVLMs poses significant challenges [29, 67].

Recent efforts have sought to navigate these complexities, with studies delving into the evaluation and detection of object hallucinations within the domain of LVLMs [38, 42, 48, 65]. For example, POPE [38] converts the hallucination into a binary classification problem to probe the model’s awareness of whether a specific object exists in the image. Concurrently, there has been a notable push towards

<sup>1</sup><https://huggingface.co/datasets/liuhaotian/llava-bench-in-the-wild>

the development of refined datasets tailored for fine-tuning existing LVLMs [20, 35, 42], training a post-hoc revisor to detect and reconstruct less hallucinatory outputs [77], and adapting factually augmented RLHF [60]. Nevertheless, existing approaches that acquire additional datasets, conduct fine-grained tuning on original or newly introduced models, or utilize other off-the-shell pretrained models can be time-consuming, labor-intensive, and computationally costly. Instead, we propose a conceptually different and training-free approach, VCD, that contrasts the output distributions with original and distorted visual inputs to calibrate the model’s over-reliance on unimodal priors and statistical bias, without utilizing external models.

### 3. Method

#### 3.1. Decoding of Vision-Language Models

We consider an LVLm parametrized by  $\theta$ . The model takes as input a textual query  $x$  and a visual input  $v$ , where  $v$  provides contextual visual information to assist the model in generating a relevant response  $y$  to the textual query. The response  $y$  is sampled auto-regressively from the probability distribution conditioned on the query  $x$  and the visual context  $v$ . Mathematically, this can be formulated as:

$$y_t \sim p_\theta(y_t | v, x, y_{<t}), \quad (1)$$

$$\propto \exp \text{logit}_\theta(y_t | v, x, y_{<t}),$$

where  $y_t$  denotes the token at time step  $t$ , and  $y_{<t}$  represents the sequence of generated tokens up to the time step  $(t - 1)$ . In the decoding phase of LVLms, object hallucinations often emerge when probabilities are erroneously allocated to tokens that do not align with the presented visual input  $v$ . Previous studies have identified two primary causes of this problem: (1) statistical biases inherent in training data (e.g., prevalent but superficial object correlations) [1, 2, 19], and (2) over-reliance on language priors embedded within the powerful LLMs used as decoders [22, 38, 70, 75]. Our approach to mitigate object hallucinations first amplifies these undesirable behaviors with vague inputs and subsequently contrasts with them in the decoding process.

#### 3.2. Visual Uncertainty Amplifies Hallucinations

The fidelity of visual input is pivotal for LVLms to accurately encode visual features and generate outputs faithfully. Yet, the introduction of uncertainty in visual inputs can tilt the equilibrium. This section delves into a comprehensive analysis aiming to validate the assumption that increased visual uncertainty can amplify the language priors and statistical biases in LVLms, thus exacerbating object hallucination.

**Introduction of Visual Uncertainty** In this paper, we propose to adopt the most elementary method—applying a Gaussian noise mask to the original image—to introduce visual

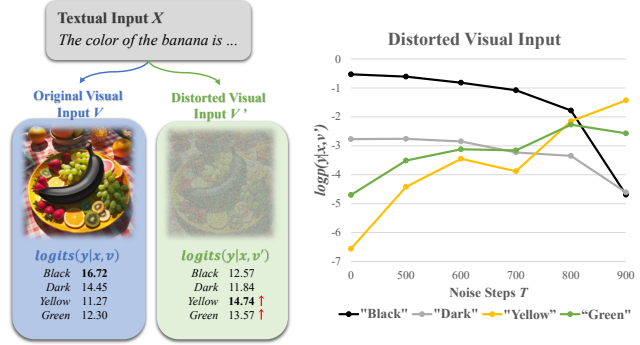


Figure 2. An illustration of visual uncertainty amplifying language priors. Given an image featuring a black banana among other colorful fruits, LVLms favor more conventional banana colors—such as "yellow" and "green", with increasing visual uncertainty. The ground-truth color "black" diminishes in probability ( $\text{log}p(y|x,v')$ ) as the distortion escalates, making LVLms over-reliant on the language priors from LLM pre-training that typically associate bananas with being yellow or green.

uncertainty. This method, although straightforward, provides an initial benchmark to estimate the baseline effects of visual uncertainty on model outputs. Following the forward diffusion process in image generation [24], the distorted image is modeled as follows:

$$q(v_t | v_{t-1}) = \mathcal{N}(v_t; \sqrt{1 - \gamma}v_{t-1}, \gamma\mathbf{I}) \quad (2)$$

$$q(v_T | v_0) = \prod_{t=1}^T q(v_t | v_{t-1}),$$

where  $v_0$  denotes the original visual input (i.e., original image) and  $\mathbf{I}$  refers to an identity matrix. We incrementally add a small amount of Gaussian noise for  $T$  steps, producing a sequence of distorted images  $v_1, \dots, v_T$ . The original image  $v_0$  gradually loses its distinguishable features as step  $t$  goes larger, where the amount of noise added in each step is controlled by  $\gamma$ . Eventually, when  $T \rightarrow \infty$ , visual uncertainty reaches the maximum and  $v_T$  will become indistinguishable from Gaussian noise.

**Visual Uncertainty Amplifies Language Priors** Figure 2 shows that visual uncertainty can compel LVLms to overlook visual evidence and overly exploit language priors for decision-making. However, this tendency is not entirely unexpected, as LLMs are designed to predict next-word probabilities based on vast textual corpora. When confronted with ambiguous visual stimuli, an LVLm might misinterpret these conventional, text-based predictions as a "safety net". These priors, while generally useful, can introduce biases or assumptions that are inconsistent with the actual visual content, particularly when the visual input lacks clarity.

**Visual Uncertainty Amplifies Statistical Bias** The construction of most vision-language pretraining datasets is predominantly based on MSCOCO [40], which inherently suffers

from an unbalanced object distribution and biased object correlations. Previous works [38, 77] point out that LVLMs, trained on such data, may inherit those statistical biases to generate descriptions with hallucinated objects. To further examine the hypothesis that visual uncertainty may amplify statistical biases from pretraining, we designed two targeted experiments to verify (1) if LVLMs hallucinate frequent objects more with distorted visual inputs and (2) if LVLMs are more prone to hallucinate objects that frequently co-occur with ground-truth objects in the image with distorted visual inputs. Figure 3 shows an evident tendency that LVLMs are more prone to hallucinate frequent and co-occurring objects, attributing to the imbalanced object distributions and spurious object correlations inherited from the training data.

### 3.3. Visual Contrastive Decoding

#### 3.3.1 Contrasting the Predictions

Our observations in the previous section reveal that visual uncertainty not only amplifies reliance on language priors but also makes LVLMs more likely to be biased by superficial object correlations present in pretraining datasets, leading to more severe hallucinations. In light of this, we introduce Visual Contrastive Decoding (VCD). VCD is formulated to counteract the statistical biases and language priors in LVLMs by contrasting model outputs generated from original and distorted visual inputs. This is achieved without necessitating additional training or external pretrained models, making VCD a cost-effective and efficient solution.

Specifically, given a textual query  $x$  and a visual input  $v$ , the model generates two distinct output distributions: one conditioned on the original  $v$  and the other on the distorted visual input  $v'$ , which is derived by applying pre-defined distortions (i.e., Gaussian noise mask) to  $v$ . Then, a new contrastive probability distribution is computed by exploiting the differences between the two initially obtained distributions. The new contrastive distribution  $p_{vcd}$  is formulated as:

$$p_{vcd}(y | v, v', x) = \text{softmax}[(1 + \alpha) \text{logit}_\theta(y | v, x) - \alpha \text{logit}_\theta(y | v', x)], \quad (3)$$

where larger  $\alpha$  values indicate a stronger amplification of differences between the two distributions ( $\alpha = 0$  reduces to regular decoding). From the adjusted output distribution  $p_{vcd}$ , we can apply various sampling strategies, such as nucleus sampling [25] and beam search [15].

Essentially, VCD serves as a corrective mechanism, reducing hallucinations by contrasting against a distribution predisposed to favoring them. Alternatively, VCD can also be interpreted as a form of contrastive ensemble that differentiates between the logits of  $p_\theta(y | v, x)$  and  $p_\theta(y | v', x)$ . This method echoes the contrastive objective commonly employed in image generation. For instance, classifier-




Figure 3. The left subfigure shows the correlation between frequent objects in MSCOCO and their propensity to be hallucinated in the validation set. Objects with a higher occurrence rate in the dataset are more likely to be hallucinated by LVLMs under distorted visual scenarios. The right subfigure charts three objects that often appear alongside “dining table”, where they are also more frequently hallucinated when presented with distorted visual inputs.

free diffusion models [23] estimate diffusion noise using  $(1 + \alpha)\epsilon_\theta(x, c) - \alpha\epsilon_\theta(x)$ , where  $c$  serves as a controlling factor. In the realm of text generation, several studies have also exploited contrastive decoding for more faithful generation [37, 41, 52, 56].

#### 3.3.2 Adaptive Plausibility Constraints

According to the formation of the contrastive distribution  $p_{vcd}$  in Equation 3, a challenge may arise as it penalizes the model’s entire output behaviors influenced by distorted visual inputs. However, this is not universally correct – the output distributions with distorted visual inputs can still uphold fundamental linguistic standards and common sense reasoning. Indiscriminate penalization could inaccurately punish these valid outputs and promote the generation of implausible tokens. To address this issue, we follow Li et al. [37] to implement an adaptive plausibility constraint that is contingent upon the confidence level associated with the output distribution with original visual inputs:

$$\begin{aligned} \mathcal{V}_{\text{head}}(y_{<t}) &= \{y_t \in \mathcal{V} : \\ p_\theta(y_t | v, x, y_{<t}) &\geq \beta \max_w p_\theta(w | v, x, y_{<t})\}, \quad (4) \\ p_{vcd}(y_t | v, v', x) &= 0, \text{ if } y_t \notin \mathcal{V}_{\text{head}}(y_{<t}), \end{aligned}$$

where  $\mathcal{V}$  is the output vocabulary of LVLMs and  $\beta$  is a hyperparameter in  $[0, 1]$  for controlling the truncation of the next token distribution. Larger  $\beta$  indicates more aggressive truncation, keeping only high-probability tokens.

Combining the visual contrastive decoding and the adaptive plausibility constraint, we obtain the full formulation:

$$\begin{aligned} y_t &\sim \text{softmax}[(1 + \alpha) \text{logit}_\theta(y_t | v, x, y_{<t}) \\ &\quad - \alpha \text{logit}_\theta(y_t | v', x, y_{<t})], \quad (5) \\ &\text{subject to } y_t \in \mathcal{V}_{\text{head}}(y_{<t}) \end{aligned}$$

Incorporating adaptive plausibility constraints refines the contrastive distribution, bolstering confidence in straightfor-



ward decisions. This ensures that when the model is highly confident in its outputs associated with the original inputs, the candidate pool is streamlined, often retaining a singular token with high probability. Such an approach effectively neutralizes potential adverse effects of VCD, preventing it from inadvertently promoting the generation of implausible tokens and maintaining the integrity of the generated content.

## 4. Experiments

This section details our assessment of the proposed Visual Contrastive Decoding across various LVLMS.

### 4.1. Experimental Settings

#### Datasets & Evaluation Metrics

**POPE**, the Polling-based Object Probing Evaluation [38], presents a streamlined approach to assess object hallucination. Within this benchmark, LVLMS are queried to answer if a specific object exists in the given image. The ratio between queries probing existent objects and non-existent objects is balanced (i.e., 50% vs. 50%). It encompasses three sampling settings: *random*, *popular*, and *adversarial*, each distinct in constructing negative samples. In the *random* setting, objects absent from the image are chosen randomly. The *popular* setting selects missing objects from a high-frequency pool, while in the *adversarial* setting, co-occurring objects not present in the image are prioritized. The POPE benchmark aggregates data from three distinct sources: MSCOCO [40], A-OKVQA [55], and GQA [27]. It involves 500 images from each dataset under each sampling setting and formulates 6 questions per image, culminating in a total of 27,000 query-answer pairs from the development sets of these datasets. The evaluation pivots on four key metrics: Accuracy, Precision, Recall, and the F1 score.

**MME** [16] serves as an extensive benchmark tailored to assess LVLMS across multiple dimensions. It comprises ten perception-related subtasks and four cognition-focused ones. Following Shukang et al. [58], except for adapting the whole dataset, we additionally leverage the existence and count subsets for object-level hallucination evaluation, and the position and color subsets for attribute-level hallucination assessment. Performance is quantified via the combined metric of accuracy and accuracy+ as the official implementation.

**LLaVA-Bench**<sup>2</sup> features a collection of 24 images, accompanying 60 questions that span a range of contexts including indoor and outdoor scenes, memes, paintings, and sketches. This dataset is crafted to assess the capability of LVLMS in tackling more challenging tasks and their adaptability to new domains. We conduct case studies on this dataset to qualitatively demonstrate the effectiveness of our proposed VCD.

<sup>2</sup><https://huggingface.co/datasets/liuhaotian/llava-bench-in-the-wild>

**LVLMS Baselines** We evaluate the effectiveness of our VCD on three state-of-the-art LVLMS. Concretely, we apply our VCD to LLaVA-1.5 and InstructBLIP, which employ Vicuna 7B as language decoder [12, 44], and Qwen-VL, built on top of Qwen 7B backbone [5]. For a more convincing comparison, we report the averaged results as well as the standard deviation over 5 runs on POPE and MME benchmarks.

**Implementation Details** Throughout our experiments, we set  $\alpha = 1$ ,  $\beta = 0.1$ , and  $\gamma = 0.1$  unless explicitly stated otherwise. For a consistent comparative analysis, our baseline decoding strategy employs direct sampling (i.e., denoted as “Regular” in all experimental tables), where the next token is directly sampled from the post-softmax distribution<sup>3</sup>. Conversely, instances labeled as “VCD” in the decoding column of all experimental tables refer to our proposed Visual Contrastive Decoding strategy, which also directly samples from the modified post-softmax distribution after applying VCD. Comprehensive parameter configurations can be found in Supplementary Materials.

### 4.2. Experimental Results

**Results on POPE** Experimental results on POPE under the random, popular, and adversarial settings are summarized in Table 1. A notable observation is the robust effect of our proposed VCD. Specifically, under different sampling settings, the performances of our VCD consistently surpass the baseline results by large margins (up to +5.8 accuracy and +7.4 F1) on all of the LVLMS. This suggests its pivotal role in counteracting statistical biases and language priors in LVLMS, thereby reducing instances of object hallucination. In addition, all LVLMS display a clear performance degradation as we move from the *random* setting to *popular* and experience a further decline while moving to the *adversarial* setting. This trend verifies our hypothesis that statistical biases inherent in LVLMS substantially contribute to the object hallucination problem. In a more detailed model-specific analysis, VCD demonstrates varied effects across different LVLMS. For LLaVA-1.5 and Qwen-VL, the F1 score elevation is predominantly driven by a recall boost (e.g., up to 10 points), showcasing its enhanced ability to accurately detect object presences. Conversely, InstructBLIP’s F1 score improvement is largely due to improved precision, signifying its enhanced capability to accurately filter out false positives. This highlights VCD’s ability to accentuate distinct attributes of various model architectures in binary decision scenarios of POPE.

**Results on MME Hallucination Subset** The MME subset evaluations extend beyond POPE’s scope, encompassing

<sup>3</sup>Optimization of  $\alpha$ ,  $\beta$ ,  $T$ , and applying other sampling strategies as detailed in the ablation studies in Supplementary Materials may yield better results. The current settings serve as a constant baseline to demonstrate the efficacy of our approach.

Dataset	Setting	Model	Decoding	Accuracy $\uparrow$	Precision	Recall	F1 Score $\uparrow$	
MSCOCO	Random	LLaVA1.5	Regular	83.29 $_{(\pm 0.35)}$	92.13 $_{(\pm 0.54)}$	72.80 $_{(\pm 0.57)}$	81.33 $_{(\pm 0.41)}$	
			VCD	<b>87.73</b> $_{(\pm 0.40)}$	91.42 $_{(\pm 0.55)}$	83.28 $_{(\pm 0.42)}$	<b>87.16</b> $_{(\pm 0.41)}$	
		Qwen-VL	Regular	84.73 $_{(\pm 0.36)}$	95.61 $_{(\pm 0.45)}$	72.81 $_{(\pm 0.38)}$	82.67 $_{(\pm 0.41)}$	
			VCD	<b>88.63</b> $_{(\pm 0.10)}$	94.64 $_{(\pm 0.25)}$	81.91 $_{(\pm 0.19)}$	<b>87.81</b> $_{(\pm 0.11)}$	
		InstructBLIP	Regular	80.71 $_{(\pm 0.73)}$	81.67 $_{(\pm 0.67)}$	79.19 $_{(\pm 1.14)}$	80.41 $_{(\pm 0.80)}$	
			VCD	<b>84.53</b> $_{(\pm 0.38)}$	88.55 $_{(\pm 0.54)}$	79.32 $_{(\pm 0.44)}$	<b>83.68</b> $_{(\pm 0.40)}$	
	Popular	LLaVA1.5	Regular	81.88 $_{(\pm 0.48)}$	88.93 $_{(\pm 0.60)}$	72.80 $_{(\pm 0.57)}$	80.06 $_{(\pm 0.05)}$	
			VCD	<b>85.38</b> $_{(\pm 0.38)}$	86.92 $_{(\pm 0.53)}$	83.28 $_{(\pm 0.42)}$	<b>85.06</b> $_{(\pm 0.37)}$	
		Qwen-VL	Regular	84.13 $_{(\pm 0.18)}$	94.31 $_{(\pm 0.43)}$	72.64 $_{(\pm 0.45)}$	82.06 $_{(\pm 0.23)}$	
			VCD	<b>87.12</b> $_{(\pm 0.07)}$	91.49 $_{(\pm 0.10)}$	81.85 $_{(\pm 0.19)}$	<b>86.40</b> $_{(\pm 0.09)}$	
		InstructBLIP	Regular	78.22 $_{(\pm 0.84)}$	77.87 $_{(\pm 1.03)}$	78.85 $_{(\pm 0.52)}$	78.36 $_{(\pm 0.76)}$	
			VCD	<b>81.47</b> $_{(\pm 0.42)}$	82.89 $_{(\pm 0.64)}$	79.32 $_{(\pm 0.44)}$	<b>81.07</b> $_{(\pm 0.39)}$	
	Adversarial	LLaVA1.5	Regular	78.96 $_{(\pm 0.52)}$	83.06 $_{(\pm 0.58)}$	72.75 $_{(\pm 0.59)}$	77.57 $_{(\pm 0.57)}$	
			VCD	<b>80.88</b> $_{(\pm 0.33)}$	79.45 $_{(\pm 0.29)}$	83.29 $_{(\pm 0.43)}$	<b>81.33</b> $_{(\pm 0.34)}$	
		Qwen-VL	Regular	82.26 $_{(\pm 0.30)}$	89.97 $_{(\pm 0.33)}$	72.61 $_{(\pm 0.50)}$	80.37 $_{(\pm 0.37)}$	
			VCD	<b>84.26</b> $_{(\pm 0.39)}$	85.84 $_{(\pm 0.45)}$	82.05 $_{(\pm 0.39)}$	<b>83.90</b> $_{(\pm 0.39)}$	
		InstructBLIP	Regular	75.84 $_{(\pm 0.45)}$	74.30 $_{(\pm 0.63)}$	79.03 $_{(\pm 0.68)}$	76.59 $_{(\pm 0.40)}$	
			VCD	<b>79.56</b> $_{(\pm 0.41)}$	79.67 $_{(\pm 0.59)}$	79.39 $_{(\pm 0.50)}$	<b>79.52</b> $_{(\pm 0.38)}$	
	A-OKVQA	Random	LLaVA1.5	Regular	83.45 $_{(\pm 0.48)}$	87.24 $_{(\pm 0.68)}$	78.36 $_{(\pm 0.54)}$	82.56 $_{(\pm 0.50)}$
				VCD	<b>86.15</b> $_{(\pm 0.23)}$	85.18 $_{(\pm 0.34)}$	87.53 $_{(\pm 0.14)}$	<b>86.34</b> $_{(\pm 0.21)}$
			Qwen-VL	Regular	86.67 $_{(\pm 0.48)}$	93.16 $_{(\pm 0.55)}$	79.16 $_{(\pm 0.59)}$	85.59 $_{(\pm 0.53)}$
				VCD	<b>89.22</b> $_{(\pm 0.14)}$	90.77 $_{(\pm 0.04)}$	87.32 $_{(\pm 0.34)}$	<b>89.01</b> $_{(\pm 0.16)}$
			InstructBLIP	Regular	80.91 $_{(\pm 0.34)}$	77.97 $_{(\pm 0.59)}$	86.16 $_{(\pm 0.88)}$	81.86 $_{(\pm 0.32)}$
				VCD	<b>84.11</b> $_{(\pm 0.27)}$	82.21 $_{(\pm 0.35)}$	87.05 $_{(\pm 0.53)}$	<b>84.56</b> $_{(\pm 0.28)}$
Popular		LLaVA1.5	Regular	79.90 $_{(\pm 0.33)}$	80.85 $_{(\pm 0.31)}$	78.36 $_{(\pm 0.54)}$	79.59 $_{(\pm 0.37)}$	
			VCD	<b>81.85</b> $_{(\pm 0.44)}$	78.60 $_{(\pm 0.58)}$	87.53 $_{(\pm 0.14)}$	<b>82.82</b> $_{(\pm 0.36)}$	
		Qwen-VL	Regular	85.56 $_{(\pm 0.35)}$	90.44 $_{(\pm 0.56)}$	79.53 $_{(\pm 0.84)}$	84.63 $_{(\pm 0.42)}$	
			VCD	<b>87.85</b> $_{(\pm 0.30)}$	88.10 $_{(\pm 0.36)}$	87.53 $_{(\pm 0.47)}$	<b>87.81</b> $_{(\pm 0.31)}$	
		InstructBLIP	Regular	76.19 $_{(\pm 0.80)}$	72.16 $_{(\pm 0.69)}$	85.28 $_{(\pm 0.79)}$	78.17 $_{(\pm 0.73)}$	
			VCD	<b>79.78</b> $_{(\pm 0.47)}$	76.00 $_{(\pm 0.52)}$	87.05 $_{(\pm 0.53)}$	<b>81.15</b> $_{(\pm 0.42)}$	
Adversarial		LLaVA1.5	Regular	74.04 $_{(\pm 0.34)}$	72.08 $_{(\pm 0.53)}$	78.49 $_{(\pm 0.38)}$	75.15 $_{(\pm 0.23)}$	
			VCD	<b>74.97</b> $_{(\pm 0.39)}$	70.01 $_{(\pm 0.40)}$	87.36 $_{(\pm 0.15)}$	<b>77.73</b> $_{(\pm 0.29)}$	
		Qwen-VL	Regular	79.57 $_{(\pm 0.31)}$	79.77 $_{(\pm 0.34)}$	79.23 $_{(\pm 0.73)}$	79.50 $_{(\pm 0.38)}$	
			VCD	<b>81.27</b> $_{(\pm 0.09)}$	77.79 $_{(\pm 0.20)}$	87.53 $_{(\pm 0.34)}$	<b>82.38</b> $_{(\pm 0.10)}$	
		InstructBLIP	Regular	70.71 $_{(\pm 0.76)}$	65.91 $_{(\pm 0.74)}$	85.83 $_{(\pm 0.80)}$	75.56 $_{(\pm 0.57)}$	
			VCD	<b>74.33</b> $_{(\pm 0.67)}$	69.46 $_{(\pm 0.73)}$	86.87 $_{(\pm 0.27)}$	<b>77.19</b> $_{(\pm 0.47)}$	
GQA		Random	LLaVA1.5	Regular	83.73 $_{(\pm 0.27)}$	87.16 $_{(\pm 0.39)}$	79.12 $_{(\pm 0.35)}$	82.95 $_{(\pm 0.28)}$
				VCD	<b>86.65</b> $_{(\pm 0.45)}$	84.85 $_{(\pm 0.59)}$	89.24 $_{(\pm 0.34)}$	<b>86.99</b> $_{(\pm 0.41)}$
			Qwen-VL	Regular	80.97 $_{(\pm 0.32)}$	88.07 $_{(\pm 0.34)}$	71.64 $_{(\pm 0.57)}$	79.01 $_{(\pm 0.40)}$
				VCD	<b>85.59</b> $_{(\pm 0.38)}$	86.88 $_{(\pm 0.44)}$	83.84 $_{(\pm 0.36)}$	<b>85.33</b> $_{(\pm 0.38)}$
			InstructBLIP	Regular	79.65 $_{(\pm 0.24)}$	77.14 $_{(\pm 0.43)}$	84.29 $_{(\pm 0.36)}$	80.56 $_{(\pm 0.18)}$
				VCD	<b>83.69</b> $_{(\pm 0.11)}$	81.84 $_{(\pm 0.42)}$	86.61 $_{(\pm 0.48)}$	<b>84.16</b> $_{(\pm 0.01)}$
	Popular	LLaVA1.5	Regular	78.17 $_{(\pm 0.17)}$	77.64 $_{(\pm 0.26)}$	79.12 $_{(\pm 0.35)}$	78.37 $_{(\pm 0.18)}$	
			VCD	<b>80.73</b> $_{(\pm 0.47)}$	76.26 $_{(\pm 0.68)}$	89.24 $_{(\pm 0.34)}$	<b>82.24</b> $_{(\pm 0.35)}$	
		Qwen-VL	Regular	75.99 $_{(\pm 0.33)}$	78.62 $_{(\pm 0.41)}$	71.40 $_{(\pm 0.38)}$	74.84 $_{(\pm 0.34)}$	
			VCD	<b>81.83</b> $_{(\pm 0.27)}$	80.45 $_{(\pm 0.47)}$	84.09 $_{(\pm 0.32)}$	<b>82.23</b> $_{(\pm 0.22)}$	
		InstructBLIP	Regular	73.87 $_{(\pm 0.58)}$	69.63 $_{(\pm 0.54)}$	84.69 $_{(\pm 0.68)}$	76.42 $_{(\pm 0.52)}$	
			VCD	<b>78.57</b> $_{(\pm 0.14)}$	74.62 $_{(\pm 0.22)}$	86.61 $_{(\pm 0.48)}$	<b>80.17</b> $_{(\pm 0.16)}$	
	Adversarial	LLaVA1.5	Regular	75.08 $_{(\pm 0.33)}$	73.19 $_{(\pm 0.49)}$	79.16 $_{(\pm 0.35)}$	76.06 $_{(\pm 0.24)}$	
			VCD	<b>76.09</b> $_{(\pm 0.43)}$	70.83 $_{(\pm 0.45)}$	88.75 $_{(\pm 0.56)}$	<b>78.78</b> $_{(\pm 0.36)}$	
		Qwen-VL	Regular	75.46 $_{(\pm 0.63)}$	77.92 $_{(\pm 0.73)}$	71.07 $_{(\pm 0.97)}$	74.33 $_{(\pm 0.71)}$	
			VCD	<b>80.01</b> $_{(\pm 0.27)}$	77.86 $_{(\pm 0.24)}$	83.85 $_{(\pm 0.35)}$	<b>80.75</b> $_{(\pm 0.27)}$	
		InstructBLIP	Regular	70.56 $_{(\pm 0.53)}$	66.12 $_{(\pm 0.32)}$	84.33 $_{(\pm 1.05)}$	74.12 $_{(\pm 0.58)}$	
			VCD	<b>75.08</b> $_{(\pm 0.13)}$	70.59 $_{(\pm 0.16)}$	85.99 $_{(\pm 0.10)}$	<b>77.53</b> $_{(\pm 0.08)}$	

Table 1. Results on POPE. *Regular* decoding denotes direct sampling, whereas *VCD* refers to sampling from our proposed contrastive distribution  $p_{vcd}$ . The best performances within each setting are **bolded**.

Model	Decoding	Object-level		Attribute-level		Total Scores $\uparrow$
		<i>Existence</i> $\uparrow$	<i>Count</i> $\uparrow$	<i>Position</i> $\uparrow$	<i>Color</i> $\uparrow$	
LLaVA1.5	Regular	175.67 $_{(\pm 7.51)}$	124.67 $_{(\pm 19.59)}$	114.00 $_{(\pm 9.32)}$	151.00 $_{(\pm 10.45)}$	565.33 $_{(\pm 32.92)}$
	VCD	<b>184.66</b> $_{(\pm 6.81)}$	<b>138.33</b> $_{(\pm 15.68)}$	<b>128.67</b> $_{(\pm 7.21)}$	<b>153.00</b> $_{(\pm 7.58)}$	<b>604.66</b> $_{(\pm 18.76)}$
Qwen-VL	Regular	155.00 $_{(\pm 3.54)}$	127.67 $_{(\pm 13.36)}$	<b>131.67</b> $_{(\pm 7.73)}$	173.00 $_{(\pm 9.75)}$	587.33 $_{(\pm 31.06)}$
	VCD	<b>156.00</b> $_{(\pm 6.52)}$	<b>131.00</b> $_{(\pm 6.19)}$	128.00 $_{(\pm 3.61)}$	<b>181.67</b> $_{(\pm 5.14)}$	<b>596.67</b> $_{(\pm 11.61)}$
InstructBLIP	Regular	141.00 $_{(\pm 13.97)}$	75.33 $_{(\pm 14.16)}$	<b>66.67</b> $_{(\pm 3.91)}$	97.33 $_{(\pm 16.94)}$	380.33 $_{(\pm 40.20)}$
	VCD	<b>168.33</b> $_{(\pm 11.55)}$	<b>92.33</b> $_{(\pm 8.47)}$	64.00 $_{(\pm 6.73)}$	<b>123.00</b> $_{(\pm 11.27)}$	<b>447.67</b> $_{(\pm 13.36)}$

Table 2. Results on the hallucination subset of MME. Regular decoding denotes direct sampling, whereas VCD refers to sampling from our proposed contrastive distribution  $p_{vcd}$ . The best performances within each setting are **bolded**.




Figure 4. MME full set results on LLaVA-1.5. VCD leads to consistent enhancement in LVLMs’ perception capacities while preserving their recognition competencies.

both object-level and attribute-level hallucinations. Results in Table 2 show that implementing VCD leads to a uniform enhancement in addressing object-level hallucinations for all models. Additionally, VCD demonstrates an overall positive impact on attribute-level *Color* scores, contributing to substantial overall performance gains. These improvements emphasize VCD’s strength in addressing the embedded statistical bias and language priors of LVLMs, thus bringing a positive impact on a broader range of hallucination challenges. In contrast, the *Position* score is relatively low across four metrics, with minimal uplift from VCD, suggesting the relatively weak ability of LVLMs in position reasoning.

**Results on MME Full Set** As shown in Figure 4, we also include the evaluation of VCD on MME Full Set to assess its impact on the general capability of LVLMs. With all models exhibiting comparable performance trajectories, we present the results of LLaVA-1.5 as a representative<sup>4</sup>. The implementation of VCD leads to a consistent enhancement in perception-based tasks, while the original recognition competencies of the LVLMs are preserved. This may be attributed to VCD’s reduction of statistical bias and language priors, which improves LVLMs’ general perception capacities by ensuring a visually grounded analysis.

<sup>4</sup>Comprehensive results for all three LVLMs on the MME full set are provided in Supplementary Materials.

### 4.3. Further Discussions

**Effect of Visual Uncertainty on Hallucinations** We further study how the object hallucination of LLaVA-1.5 changes along with visual uncertainty. Figure 6 depicts a clear performance drop on the POPE benchmark with the increase of noise steps, suggesting that the object hallucination will become more severe as visual uncertainty goes larger. This observation aligns with our previous findings in Section 3.2 that visual uncertainty will exacerbate object hallucination issues in LVLMs’ generative process. Our proposed VCD emerges as a correction mechanism by contrasting model outputs with original and distorted visual inputs.

#### GPT-4V Aided Evaluation of Open-Ended Generation

Beyond the “Yes-or-No” question format employed in our POPE and MME evaluations, we extend our analysis to open-ended captioning tasks in the LLaVA-Bench using the recently released LLM, GPT-4V<sup>5</sup>, following Shukang et al. [58]<sup>6</sup>. Results in Table 3 show consistent improvements in VCD over regular decoding. The observed enhancement in accuracy points to VCD’s ability to mitigate hallucinations effectively. Simultaneously, VCD’s counteraction of statistical biases and language priors enhances the perceptual capa-

<sup>5</sup><https://openai.com/research/gpt-4v-system-card>

<sup>6</sup>The prompt used for evaluation and an evaluation case is provided in Supplementary Materials.

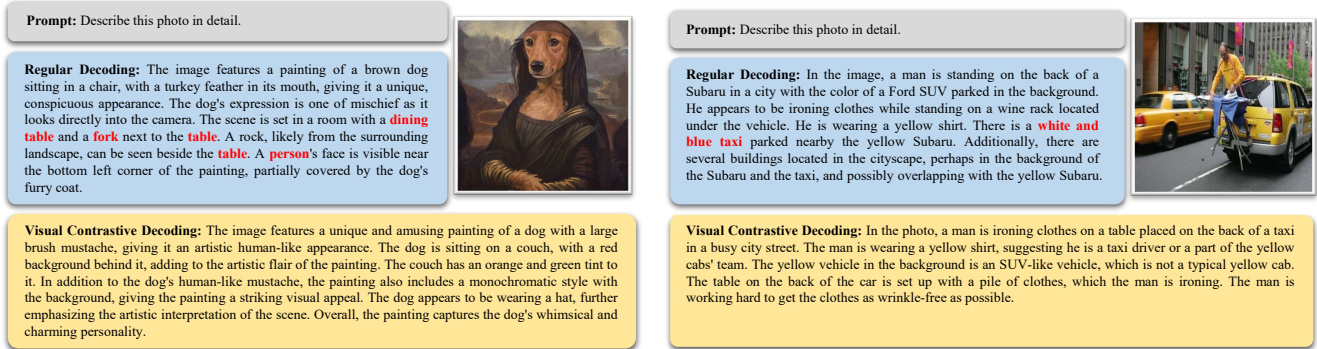


Figure 5. Illustration of hallucination correction by our proposed VCD with two samples from LLaVA-Bench. Hallucinated objects from LLaVA’s regular decoding are highlighted in red.

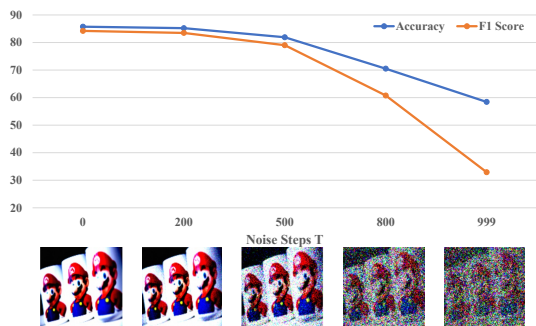


Figure 6. Performance of LLaVA-1.5 on the POPE benchmark across varying noise levels with regular decoding. We visualize the distorted visual inputs subjected to different levels of Gaussian noise at the bottom.

capabilities of LLaVAs, as evidenced by the marked improvement in the detailedness of the responses.

**Case Study on LLaVA-Bench** Figure 5 demonstrates two case studies on how, given identical prompts and images, regular decoding can yield object hallucinations influenced by the statistical bias and language priors inherent during pretraining. For instance, in the displayed examples, objects such as “dining table” and “fork”, which often co-occur with the likely ground-truth object “chair”, are hallucinated. In contrast, the implementation of VCD notably mitigates these hallucination issues and simultaneously preserves the coherence and informativeness of the output text. Due to the page limit, please refer to Supplementary Materials for more cases and ablation studies<sup>7</sup>.

## 5. Conclusion and Limitation

In this paper, we tackle the object hallucination issue in LLaVAs. We conducted an in-depth analysis of how visual uncertainty influences hallucinations, particularly from the aspect of statistical biases and language priors. Our find-

<sup>7</sup>Ablation studies in Supplementary Materials include effects of total noise steps  $T$ , hyper-parameters  $\alpha$ ,  $\beta$ , and effect of VCD on larger LLaVA variants and with other sampling strategies.

Model	Decoding	Accuracy $\uparrow$	Detailedness $\uparrow$
LLaVA-1.5	Regular	3.23	3.54
	VCD	<b>4.15</b>	<b>3.85</b>
InstructBLIP	Regular	3.84	4.07
	VCD	<b>4.23</b>	<b>4.69</b>
Qwen-VL	Regular	4.76	3.46
	VCD	<b>6.69</b>	<b>4.46</b>

Table 3. Results of GPT-4V-aided evaluation on open-ended generation. Accuracy measures the response’s alignment with the image content, and Detailedness gauges the richness of details in the response. Both metrics are on a scale of 10.

ings indicate that visual uncertainty amplifies these factors, contributing to more hallucinations. In light of this, we introduced Visual Contrastive Decoding (VCD), a novel, training-free method that employs contrastive distributions to calibrate the model’s output without the usage of external tools. Our extensive experiments across multiple benchmarks and LLaVA families confirm VCD’s efficacy in reducing hallucinations and also demonstrate its potential to enhance the overall perception capabilities of LLaVAs.

**Limitation** While this study employs a basic Gaussian noise approach to introduce visual uncertainty, more fine-grained techniques, like object-level blurring, hold the potential for improved outcomes. In addition, our focus was limited to LLaVAs processing images and text, not encompassing their emerging applications in video understanding. Future research directions include exploring diverse image distortion methods and extending the Visual Contrastive Decoding (VCD) framework to a broader range of LLaVAs.

**Acknowledgements** We would like to thank the anonymous reviewers and senior area chairs for their constructive comments and support for our work. We would also like to thank Jiayi Li for helpful discussions and paper proofreading. This work was substantially supported by DAMO Academy through DAMO Academy Research Intern Program. This work is also partially funded by the Ministry of Education Singapore under the Tier-1 scheme with project number RG18/22.



## References

- [1] Vedika Agarwal, Rakshith Shetty, and Mario Fritz. Towards causal vqa: Revealing and reducing spurious correlations by invariant and covariant semantic editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9690–9698, 2020. 1, 3
- [2] Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. Analyzing the behavior of visual question answering models. *arXiv preprint arXiv:1606.07356*, 2016. 1, 3
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. 2
- [4] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023. 2
- [5] 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. 1, 2, 5
- [6] Ali Furkan Biten, Lluís Gómez, and Dimosthenis Karatzas. Let there be a clock on the beach: Reducing object hallucination in image captioning. In *Proc. of WACV*, 2022. 2
- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 2
- [8] Long Chen, Oleg Sinavski, Jan Hünemann, Alice Karnsund, Andrew James Willmott, Danny Birch, Daniel Maund, and Jamie Shotton. Driving with llms: Fusing object-level vector modality for explainable autonomous driving. *arXiv preprint arXiv:2310.01957*, 2023. 1
- [9] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022. 2
- [10] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality, 2023. 2
- [11] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022. 2
- [12] 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. *arXiv preprint arXiv:2306.04387*, 2023. 1, 2, 5
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [14] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023. 2
- [15] Markus Freitag and Yaser Al-Onaizan. Beam search strategies for neural machine translation. *arXiv preprint arXiv:1702.01806*, 2017. 4
- [16] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiawu Zheng, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023. 2, 5
- [17] Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. Chatgpt outperforms crowd-workers for text-annotation tasks. *arXiv preprint arXiv:2303.15056*, 2023. 2
- [18] Tao Gong, Chengqi Lyu, Shilong Zhang, Yudong Wang, Miao Zheng, Qian Zhao, Kuikun Liu, Wenwei Zhang, Ping Luo, and Kai Chen. Multimodal-gpt: A vision and language model for dialogue with humans. *arXiv preprint arXiv:2305.04790*, 2023. 1
- [19] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017. 1, 3
- [20] Anisha Gunjal, Jihan Yin, and Erhan Bas. Detecting and preventing hallucinations in large vision language models. *ArXiv*, 2023. 1, 2, 3
- [21] Vipul Gupta, Zhuowan Li, Adam Kortylewski, Chenyu Zhang, Yingwei Li, and Alan Yuille. Swapmix: Diagnosing and regularizing the over-reliance on visual context in visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5078–5088, 2022. 1
- [22] Yudong Han, Liqiang Nie, Jianhua Yin, Jianlong Wu, and Yan Yan. Visual perturbation-aware collaborative learning for overcoming the language prior problem. *arXiv preprint arXiv:2207.11850*, 2022. 1, 3
- [23] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 4
- [24] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3
- [25] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*, 2019. 4
- [26] Mingzhe Hu, Shaoyan Pan, Yuheng Li, and Xiaofeng Yang. Advancing medical imaging with language models: A journey from n-grams to chatgpt. *arXiv preprint arXiv:2304.04920*, 2023. 1
- [27] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional

- question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019. 5
- [28] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38, 2023. 2
- [29] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020. 2
- [30] Jae Myung Kim, A Koepke, Cordelia Schmid, and Zeynep Akata. Exposing and mitigating spurious correlations for cross-modal retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2584–2594, 2023. 2
- [31] MV Koroteev. Bert: a review of applications in natural language processing and understanding. *arXiv preprint arXiv:2103.11943*, 2021. 2
- [32] Katherine Lee, Orhan Firat, Ashish Agarwal, Clara Fannjiang, and David Sussillo. Hallucinations in neural machine translation. *OpenReview*, 2018. 2
- [33] Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Jingkang Yang, and Ziwei Liu. Otter: A multi-modal model with in-context instruction tuning. *arXiv preprint arXiv:2305.03726*, 2023. 1, 2
- [34] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 2
- [35] Lei Li, Yuwei Yin, Shicheng Li, Liang Chen, Peiyi Wang, Shuhuai Ren, Mukai Li, Yazheng Yang, Jingjing Xu, Xu Sun, et al. A large-scale dataset towards multi-modal multilingual instruction tuning. *arXiv preprint arXiv:2306.04387*, 2023. 3
- [36] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 2
- [37] Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. Contrastive decoding: Open-ended text generation as optimization. *arXiv preprint arXiv:2210.15097*, 2022. 4
- [38] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. *ArXiv*, 2023. 1, 2, 3, 4, 5
- [39] Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*, 2021. 2
- [40] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 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, 2014. 3, 5
- [41] Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A Smith, and Yejin Choi. Experts: Decoding-time controlled text generation with experts and anti-experts. *arXiv preprint arXiv:2105.03023*, 2021. 4
- [42] 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*, 2023. 2, 3
- [43] Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. Mitigating hallucination in large multi-modal models via robust instruction tuning. *ArXiv*, 2023. 1
- [44] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*, 2023. 2, 5
- [45] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 1, 2
- [46] Haokun Liu, Yaonan Zhu, Kenji Kato, Izumi Kondo, Tadayoshi Aoyama, and Yasuhisa Hasegawa. Llm-based human-robot collaboration framework for manipulation tasks. *arXiv preprint arXiv:2308.14972*, 2023. 1
- [47] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 2
- [48] Holy Lovenia, Wenliang Dai, Samuel Cahyawijaya, Ziwei Ji, and Pascale Fung. Negative object presence evaluation (nope) to measure object hallucination in vision-language models. *arXiv preprint arXiv:2310.05338*, 2023. 1, 2
- [49] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*, 2023. 1
- [50] Jinjie Mai, Jun Chen, Bing Li, Guocheng Qian, Mohamed Elhoseiny, and Bernard Ghanem. Llm as a robotic brain: Unifying egocentric memory and control. *arXiv preprint arXiv:2304.09349*, 2023. 1
- [51] Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. Counterfactual vqa: A cause-effect look at language bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12700–12710, 2021. 1
- [52] Sean O’Brien and Mike Lewis. Contrastive decoding improves reasoning in large language models. *arXiv preprint arXiv:2309.09117*, 2023. 4
- [53] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020. 2
- [54] Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. Object hallucination in image captioning. *arXiv preprint arXiv:1809.02156*, 2018. 2
- [55] Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. A-okvqa: A benchmark for visual question answering using world knowledge. In *European Conference on Computer Vision*, pages 146–162. Springer, 2022. 5

- [56] Weijia Shi, Xiaochuang Han, Mike Lewis, Yulia Tsvetkov, Luke Zettlemoyer, and Scott Wen-tau Yih. Trusting your evidence: Hallucinate less with context-aware decoding. *arXiv preprint arXiv:2305.14739*, 2023. **4**
- [57] Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. Replug: Retrieval-augmented black-box language models. *arXiv preprint arXiv:2301.12652*, 2023. **2**
- [58] Shukang, 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*, 2023. **5, 7**
- [59] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7464–7473, 2019. **2**
- [60] 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*, 2023. **2, 3**
- [61] 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. **2**
- [62] Yi Tay, Mostafa Dehghani, Vinh Q Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Dara Bahri, Tal Schuster, Steven Zheng, et al. Ul2: Unifying language learning paradigms. In *The Eleventh International Conference on Learning Representations*, 2022.
- [63] 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*, 2023. **2**
- [64] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*, 2022. **2**
- [65] Junyang Wang, Yiyang Zhou, Guohai Xu, Pengcheng Shi, Chenlin Zhao, Haiyang Xu, Qinghao Ye, Ming Yan, Ji Zhang, Jihua Zhu, et al. Evaluation and analysis of hallucination in large vision-language models. *arXiv preprint arXiv:2308.15126*, 2023. **2**
- [66] Sheng Wang, Zihao Zhao, Xi Ouyang, Qian Wang, and Ding-gang Shen. Chatcad: Interactive computer-aided diagnosis on medical image using large language models. *arXiv preprint arXiv:2302.07257*, 2023. **1**
- [67] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022. **2**
- [68] Yike Wu, Yu Zhao, Shivan Zhao, Ying Zhang, Xiaojie Yuan, Guoqing Zhao, and Ning Jiang. Overcoming language priors in visual question answering via distinguishing superficially similar instances. *arXiv preprint arXiv:2209.08529*, 2022. **1**
- [69] Zhenyu Wu, Ziwei Wang, Xiuwei Xu, Jiwen Lu, and Haibin Yan. Embodied task planning with large language models. *arXiv preprint arXiv:2307.01848*, 2023. **1**
- [70] Hong Yan, Lijun Liu, Xupeng Feng, and Qingsong Huang. Overcoming language priors with self-contrastive learning for visual question answering. *Multimedia Tools and Applications*, 82(11):16343–16358, 2023. **1, 3**
- [71] Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, Chenliang Li, Yuanhong Xu, Hehong Chen, Junfeng Tian, Qian Qi, Ji Zhang, and Fei Huang. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023. **1, 2**
- [72] Yan Zeng, Xinsong Zhang, and Hang Li. Multi-grained vision language pre-training: Aligning texts with visual concepts. *arXiv preprint arXiv:2111.08276*, 2021. **2**
- [73] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023. **1**
- [74] Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. Siren’s song in the ai ocean: A survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*, 2023. **2**
- [75] Ren Zhibo, Wang Huizhen, Zhu Muhua, Wang Yichao, Xiao Tong, and Zhu Jingbo. Overcoming language priors with counterfactual inference for visual question answering. In *Proceedings of the 22nd Chinese National Conference on Computational Linguistics*, pages 600–610, 2023. **1, 3**
- [76] Chunting Zhou, Graham Neubig, Jiatao Gu, Mona Diab, Paco Guzman, Luke Zettlemoyer, and Marjan Ghazvininejad. Detecting hallucinated content in conditional neural sequence generation. *arXiv preprint arXiv:2011.02593*, 2020. **2**
- [77] 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*, 2023. **2, 3, 4**
- [78] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. **1**