

Mitigating Hallucinations in Large Vision-Language Models via Summary-Guided Decoding

Kyungmin Min¹ Minbeom Kim¹ Kang-il Lee²

Dongryeol Lee² Kyomin Jung^{1,2*}

¹IPAI, Seoul National University

²Dept. of ECE, Seoul National University

{kyungmin97,minbeomkim,4bkang,dr1123,kjung}@snu.ac.kr

Abstract

Large Vision-Language Models (LVLMs) demonstrate impressive capabilities in generating detailed and coherent responses from visual inputs. However, they are prone to generate hallucinations due to an over-reliance on language priors. To address this issue, we investigate the language priors in LVLMs and make two key observations: (1) Even when predicting the tokens associated with image-related part-of-speech (POS), models increasingly rely on linguistic priors as the token sequences grow, thereby amplifying hallucinations. (2) Methods that directly calibrate LVLM’s output distribution to mitigate language priors can lead to a degradation in text quality or even exacerbate hallucinations. Based on these findings, we propose a novel method, **Summary-Guided Decoding (SumGD)**. This method naturally encourages the model to focus more on image information by reducing the text context through summaries, while controlling only the image-related POS tokens to maintain text quality. Through experiments, we demonstrate that SumGD achieves state-of-the-art performance on object hallucination benchmarks. Furthermore, in terms of the trade-off between precision and recall, SumGD achieves Pareto optimality among the existing methods. Lastly, we observe that although existing methods struggle to balance the reduction of object hallucinations with maintaining text quality, SumGD demonstrates robustness in handling this challenge.

1 Introduction

Large Vision-Language Models (LVLMs) have shown remarkable advancements by integrating the reasoning capabilities of Large Language Models (LLMs) to interpret visual knowledge (Zhu et al., 2023; Dai et al., 2023; Liu et al., 2024c; Li et al., 2023a). Despite their significant utility, they suffer

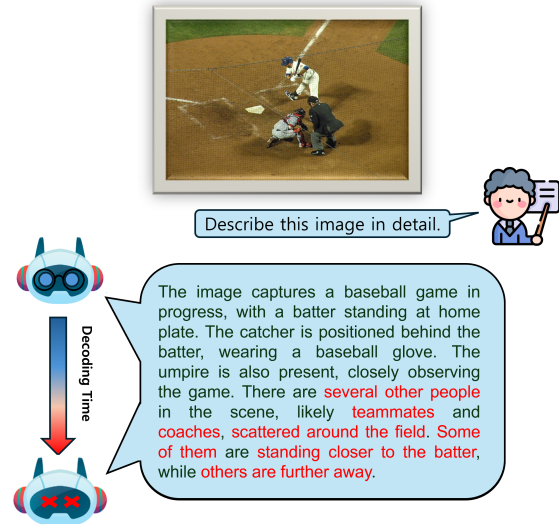


Figure 1: An example of LVLMs’ hallucination. LVLMs hallucinate due to their over-reliance on previously generated text. The red fonts represent the hallucinatory content.

from a critical drawback known as *object hallucination*, where the model generate responses that contradict the visual input (Li et al., 2023c; Liu et al., 2024b). Recent studies have shown that this occurs because LVLMs rely too heavily on learned textual patterns, which referred as *language priors* (Zhou et al., 2024; Liu et al., 2024a; Jing et al., 2023; Lee et al., 2024b). This over-reliance on language priors tends to intensify when the model generates longer sequences or detailed descriptions (Favero et al., 2024), leading to frequent hallucinations as shown in Figure 1.

In this paper, we 1) conduct the fundamental analysis of language priors in LVLMs (Section 2), 2) analyze the limitations of existing methods for mitigating language priors and provide insights into potential solutions (Section 5.1), and 3) propose a novel method that effectively reduces object hallucination while preserving text quality (Section 3).

First, we analyze language priors by examining the distance between the next-token probability dis-

*Corresponding authors.

tributions of LVLMs and LLMs, both conditioned on the same text sequence. Breaking this down by part-of-speech (POS) types reveals a significant divergence for image-related POS tokens, such as NOUN (e.g., “tree”) or ADJ (e.g., “green”). Conversely, language-related POS tokens, such as AUX (e.g., “is,” “will”), show nearly identical distributions. These findings suggest that LVLMs still rely heavily on the same linguistic structures as LLMs, except when visual input is particularly relevant — such as when describing specific objects or attributes. In other words, LVLMs incorporate visual information within a linguistic framework that is very similar to that of LLMs.

Problematically, we discover that even for these image-related POS tokens, the distributional distance rapidly decreases as the number of generated tokens increases. In other words, even when visual information is necessary, LVLMs tend to focus more on textual information, leading to frequent occurrences of object hallucination. We identify this phenomenon as an over-reliance on language priors.

Next, we examine the limitations of contrastive decoding, a promising method for mitigating hallucinations (Favero et al., 2024; Wang et al., 2024; Leng et al., 2023; Kim et al., 2024c; Zhu et al., 2024). Our analysis reveals two primary issues: (1) The effort to reduce language priors through contrastive decoding can disrupt the natural distribution of language-related tokens, potentially degrading overall text quality. (2) As token length increases, the model’s reliance on language priors becomes more pronounced, leading the two output distributions being contrasted to become increasingly similar. This similarity reduces the effectiveness of contrastive decoding in steering the model towards an image-aligned distribution. These findings suggest that reducing language priors may be more effectively achieved by integrating visual information naturally, with minimal intervention in the decoding process.

Building on these observations, we propose a novel method called Summary-Guided Decoding (SumGD). Our approach employs a summarization technique that selectively retains essential information from previously generated sentences, encouraging LVLMs to more effectively incorporate image information. To minimize unnecessary intervention for preserving text quality, the summarization is referenced only when predicting image-related POS tokens, which require image-specific

details.

Our experimental results demonstrate that SumGD significantly outperforms all other decoding approaches in object hallucination benchmarks (e.g., up to +16.5% in CHAIR_S and +19% in CHAIR_I) across various models and architecture sizes. Additionally, SumGD demonstrates Pareto optimal performance, effectively balancing the reduction of object hallucinations with the preservation of high object recall. This balance becomes more pronounced as token length increases. Finally, the results confirm that SumGD not only reduces object hallucinations but also preserves the overall text quality of LVLMs.¹

Our contributions are summarized as follows:

- We analyze how LVLMs tend to disregard image information and increasingly rely on language priors, based on the position and POS type of each token.
- Based on these findings, we propose Summary-Guided Decoding (SumGD). SumGD modifies next-token probabilities using summarized contexts, but only for image-related POS tokens. This approach aims to reflect image information while preserving LVLM’s text quality as much as possible.
- SumGD demonstrates state-of-the-art performance in object hallucination benchmarks and achieves Pareto optimal across all methods in terms of the precision-recall trade-off. Additionally, SumGD preserves text quality almost entirely.

2 Language Priors in LVLMs

In this section, we systematically analyze the causes of language priors in LVLMs. Section §2.1 outlines the method for quantifying language priors. Section §2.2 provides an in-depth analysis of how language priors affect LVLMs based on part-of-speech (POS) types. Section §2.3 analyzes the impact of increasing token length on language priors in LVLMs. We conduct this analysis on 5,000 MSCOCO (Lin et al., 2015) image descriptions generated using LLaVA 1.5 7B (Liu et al., 2024c) (see Appendix B for more details).

¹The code will be available at <https://github.com/andy9705/SumGD>

2.1 How to measure language priors in LVLMs

In LVLMs, language priors refer to the model’s over-reliance on learned textual patterns, where responses are generated based on these patterns without fully considering the provided image. From this perspective, if the token distribution of a LVLM, which decodes using both text and images, becomes similar to that of a LLM, which relies solely on text for decoding, this could indicate an over-reliance on language priors. Here, the *LLM* refers to the state of the LVLM where the input image is not provided as a conditioning factor, with both models conditioned on the same text sequence. Therefore, we measure language priors by examining the distributional distance between the next-token probabilities of LVLMs and LLMs, as described in Favero et al. (2024). We employ Jensen-Shannon Divergence (JSD) (Lin, 1991) to quantify this distance.

Formally, at each time step t , the next token y_t is selected as:

$$y_t = \arg \max_{y \in V} \log p_\theta(y \mid I, T, y_{<t}), \quad (1)$$

where θ is the parameters of LVLMs, V is the vocabulary, I denotes the provided image, T represents the textual prompt (e.g., “Please describe this image in detail.”), and $y_{<t}$ denotes the sequence of generated tokens up to the time step ($t - 1$).

We define the distributional distance at each time step t as:

$$\text{dist}_t = \text{JSD}(p_\theta(\cdot \mid I, T, y_{<t}) \parallel p_\theta(\cdot \mid T, y_{<t})). \quad (2)$$

A larger distance dist_t suggests that the LVLM relies more on visual information for predictions, indicating a lower dependence on language priors. Conversely, a smaller distance implies that the model is generating responses primarily based on textual patterns.

2.2 Analysis of language priors by Part-of-Speech (POS) type

We conduct an experiment to investigate whether LVLMs differ in their reliance on language priors based on the need for image information. Specifically, we analyze this tendency by part-of-speech (POS) type, measuring the JSD at each decoding step and averaging the JSD values for each POS type² up to 32 tokens.

²We utilized the Spacy model (en_core_web_sm) for POS tagging

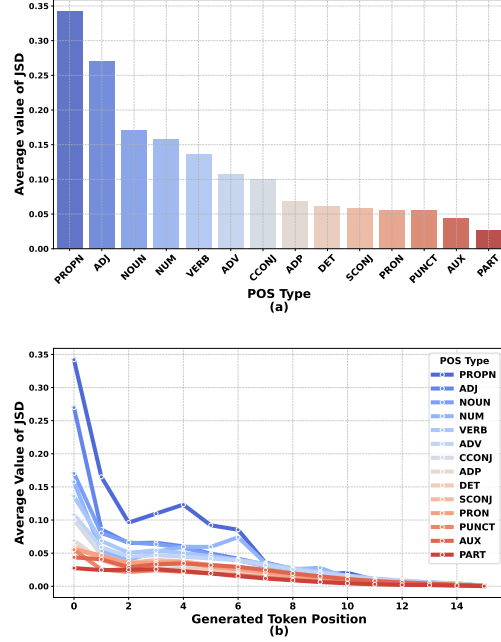


Figure 2: **(Top)** The average JSD between the LVLM and the LLM for each POS category up to 32 tokens. **(Bottom)** The average JSD between the LVLM and the LLM for each POS category across intervals, with each interval consisting of 32 tokens.

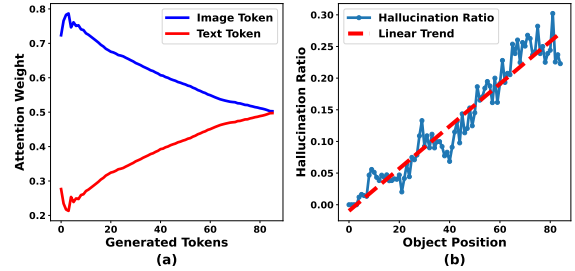


Figure 3: **(Left)** Attention weights of image tokens and text tokens at each decoding step (or token length). **(Right)** Object hallucination ratio at each generated token position.

As shown in Figure 2 (a), we observe significant variation in divergence across different POS categories. POS categories such as PROPN (e.g., “Biden”) and ADJ (e.g., “red”), which related to visual information, exhibit higher divergence. In contrast, language-related POS types, like PART (e.g., “not”, “s”) and AUX (e.g., “are”), show much lower JSD. This indicates that LVLMs integrate visual information within a linguistic framework that is highly aligned with LLMs.

Another important observation, as shown in Figure 2 (b), is that even for image-related POS tokens (e.g., NOUN), the distributional distance decreases significantly as the token length increases.

This suggests that even when image information is required during decoding, models primarily rely

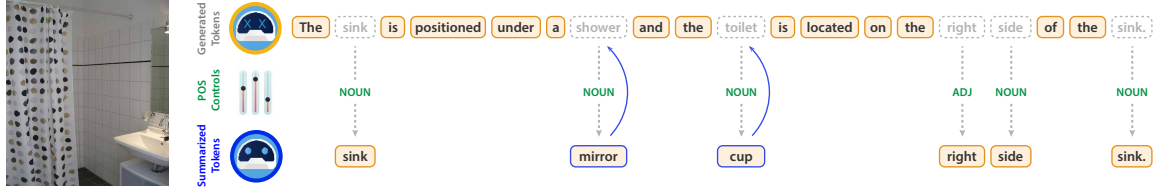


Figure 4: Illustration of our Summary-Guided Decoding.

on textual patterns. In other words, token length (or input length) has a significant influence on how language priors are employed.

2.3 Influence of Token Sequence Length on language priors

We observe that as token sequences grow longer, the model becomes increasingly dependent on language priors in Section 2.2. To explore this effect further, we conduct a detailed analysis of how varying token lengths impact LVLMs, particularly in terms of how attention is distributed between image and text tokens, and the consequent impact on object hallucination.

First, we measure the attention weights assigned to image tokens and text tokens at each decoding step. As shown in Figure 3 (a), initially, LVLMs give sufficient attention to input image tokens when computing the next token. However, as the sentence grows longer, this attention becomes significantly shallower. In other words, when generating long sentences, we can observe that LVLMs tend to rely more on linguistic patterns rather than on visual information. This observation provides additional insight into our earlier findings in Section 2.2, where longer sequences were shown to reinforce the model’s dependence on language priors.

Additionally, to assess the role of input length in hallucination, we evaluate the object hallucination ratio as a function of token length. Figure 3 (b) shows a clear correlation between input length and the likelihood of object hallucinations, indicating that longer text generation increases the chances of hallucination. We hypothesize that this phenomenon is driven by over-reliance on language priors, which amplifies hallucinations in LVLMs.

3 Summary-Guided Decoding

Based on insights from Section 2, we identify that an increase in input length results in greater reliance on language priors, thereby exacerbating

hallucinations in LVLMs. To address this, we present Summary-Guided Decoding (SumGD), a novel method for controlling the length of conditioning input during decoding. In SumGD, we shorten the conditioning input by summarizing the previously generated text after each sentence completion. This process preserves the critical context from earlier outputs while keeping the input concise. The summarized text, combined with the image, serves as part of the conditioned input for generating the next sentence. This approach effectively reduces the input length, allowing the model to stay more focused on the provided image.

Using summarized inputs can reduce contextual information, which may cause discrepancies with the language patterns previously learned by the model. This can result in distributional shifts that weaken the model’s language modeling capabilities, ultimately degrading the quality of the generated text. To address this, we preserve the original distribution for tokens related to language modeling, while using SumGD to control only the image-related POS tokens.³ Our method is illustrated in Figure 4.

We introduce two variations of SumGD for summary model usage. The first approach leverages the instruction-following capabilities inherent in LVLMs. By providing summary instructions directly to the LVM, this method enables the model to perform SumGD without incurring additional memory costs. However, a limitation of this approach is the increased computational burden, as the LVM generates its summaries during the process. To address these challenges, we distill the summarization capability into a smaller, more efficient model, Flan-T5-base (Chung et al., 2022) (see Appendix E for details). This model significantly reduces computational overhead while maintaining the advantage of input length control. We report re-

³As shown in Figure 2, we selected PROP, ADJ, NOUN, and NUM as image-related POS.

sults for both **SumGD with Self-Summarization (SumGD-S)** and **SumGD with the Distilled-Flan-T5 model (SumGD-D)**, highlighting the trade-offs between efficiency and performance.

4 Experiment

4.1 Experiment settings

Datasets and Evaluation Metrics. We generate descriptions for 200 images from the MSCOCO 2014 validation dataset (Lin et al., 2015) prompted with “Please describe this image in detail.” (Huang et al., 2024). We employ the Caption Hallucination Assessment with Image Relevance (CHAIR) (Rohrbach et al., 2019) for evaluating object hallucination. CHAIR consists of two variants: $CHAIR_I$, which calculates the percentage of hallucinated objects out of all objects mentioned in the caption, and $CHAIR_S$, which measures the percentage of captions that contain at least one hallucinated object. Additionally, to complement the precision-based CHAIR metric, we introduce a Recall metric for a more detailed assessment.

$$CHAIR_I = \frac{|\{\text{hallucinated objects}\}|}{|\{\text{all objects mentioned}\}|}$$

$$CHAIR_S = \frac{|\{\text{sentences with hallucinated object}\}|}{|\{\text{all sentences}\}|}$$

$$\text{Recall} = \frac{|\{\text{correct mentioned objects}\}|}{|\{\text{ground truth objects}\}|}$$

To provide a more comprehensive assessment of hallucinations, we use the Sentence-level Hallucination Ratio (SHR) (Zhao et al., 2023), a GPT-4-based evaluation metric. This metric includes hallucinations involving object existence, relationships, and attributes. We generate descriptions for 200 images from the VG dataset (Krishna et al., 2016), using the same prompts as in the CHAIR metric. Specifically, SHR leverages GPT-4⁴ to compare the model’s responses with the manually annotated descriptions from the VG dataset, evaluating each response on a sentence-by-sentence to identify potential hallucinations accurately.

Baseline LVLMS. In LVLMS, two prominent methods for aligning text and vision modalities are the projection layer-based approach and the learnable query-based approach (Li et al., 2023a; Zhu et al., 2023; Chen et al., 2023; Liu et al., 2023). In our experiments, we utilize representative models for each aligning method: LLaVA-1.5 (7B/13B) (Liu

⁴We used GPT-4o (gpt-4o-2024-08-06) for hallucination judgement.

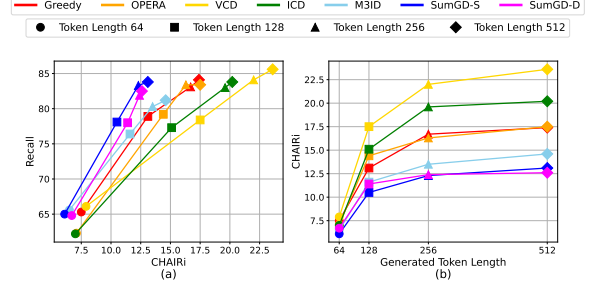


Figure 5: **(Left)** A position closer to the top-left indicates an optimal balance between factuality and recall. **(Right)** Trade-off between generated token length and hallucination (lower is better).

et al., 2024c) and InstructBLIP (7B/13B) (Dai et al., 2023).

Baseline Decoding Methods. We include various decoding methods as baseline approaches in our study, including greedy decoding, nucleus sampling, and beam search for traditional methods. In addition, we incorporate various contrastive decoding techniques, including Visual Contrastive Decoding (VCD) (Leng et al., 2023), which contrasts the original image prompt with a distorted image prompt; Instruction Contrastive Decoding (ICD) (Wang et al., 2024), which contrasts the original instruction prompt with a modified instruction prompt; and Multi-Modal Mutual Information Decoding (M3ID) (Favero et al., 2024), which contrasts the image prompt with a non-image prompt, with the contrast strength progressively increasing as the token length grows. Lastly, we include OPERA (Huang et al., 2024), a beam search-based method designed to counteract the model’s tendency to overemphasize specific anchor tokens.

4.2 Main Results

Results on CHAIR evaluation. As shown in Table 1, SumGD significantly outperforms the baseline methods in the $CHAIR_S$ and $CHAIR_I$ across different model sizes and architectures. Specifically, compared to Greedy decoding, SumGD-S achieves a 16.5% improvement in $CHAIR_S$ and a 19% improvement in $CHAIR_I$ on LLaVA 1.5 7B. On InstructBLIP 7B, the improvements are even more pronounced, with a 23.7% improvement in $CHAIR_I$.

We conduct the CHAIR evaluation by fixing the generated token lengths at 64, 128, 256, and 512, representing a range from short to long text generation to ensure a fair evaluation of object hallucination.

Method	LLAVA-1.5 7B			InstructBLIP 7B			LLAVA-1.5 13B			InstructBLIP 13B			Average		
	$C_S \downarrow$	$C_I \downarrow$	$R \uparrow$	$C_S \downarrow$	$C_I \downarrow$	$R \uparrow$	$C_S \downarrow$	$C_I \downarrow$	$R \uparrow$	$C_S \downarrow$	$C_I \downarrow$	$R \uparrow$	$C_S \downarrow$	$C_I \downarrow$	$R \uparrow$
Greedy	51.5	13.7	<u>79.1</u>	49.0	15.6	72.7	43.5	12.2	<u>78.3</u>	52.0	13.5	69.8	49.0	13.8	<u>75.0</u>
Nucleus	53.0	14.4	76.9	57.0	16.9	72.3	49.5	14.3	74.4	64.5	19.2	68.6	56.0	16.2	73.1
Beam search based ($n=5$)															
Beam Search	47.5	12.5	79.2	45.5	13.1	74.1	43.5	12.0	<u>78.3</u>	58.5	15.0	71.1	48.8	13.2	75.7
OPERA	46.0	13.4	78.3	<u>43.0</u>	13.0	<u>73.8</u>	40.0	12.5	72.2	44.5	12.0	69.5	43.4	12.7	73.5
Contrastive Decoding															
VCD	58.0	16.4	77.8	56.5	16.5	71.6	59.5	16.8	79.5	52.5	13.4	<u>71.2</u>	56.6	15.8	<u>75.0</u>
ICD	45.5	13.4	77.2	60.5	17.8	68.9	47.5	13.0	77.3	66.0	19.3	72.2	54.9	15.9	73.9
M3ID	44.5	12.0	76.1	68.0	18.0	71.6	45.0	11.9	77.8	78.0	20.8	67.8	58.9	15.7	73.3
SumGD-D (Ours)	42.5	<u>11.8</u>	77.8	42.5	<u>12.3</u>	72.7	43.0	10.9	77.7	44.5	<u>11.6</u>	69.2	43.1	<u>11.7</u>	74.4
SumGD-S (Ours)	<u>43.0</u>	11.1	<u>79.1</u>	43.5	11.9	72.2	<u>41.5</u>	<u>11.7</u>	77.3	44.5	10.4	68.8	43.1	11.3	74.4

Table 1: Results on CHAIR Metric ($max\ new\ tokens$ is 512). The best performances are bolded, and the second-best are underlined. Denote $CHAIR_S$ as C_S , $CHAIR_I$ as C_I , and Recall as R . n denotes the number of beams.

Method	LLAVA-1.5 7B		InstructBLIP 7B	
	SHR \downarrow	SPI	SHR \downarrow	SPI
Greedy	43.3	5.00	47.4	5.14
OPERA	42.0	4.74	46.4	4.76
VCD	52.0	5.18	49.5	4.97
ICD	50.2	4.93	57.8	5.93
M3ID	46.4	5.02	59.9	5.51
SumGD-D	<u>41.7</u>	5.08	<u>46.1</u>	5.26
SumGD-S	40.8	5.03	45.7	5.30

Table 2: Results on Sentence-Hallucination Ratio (SHR) and Sentence Per Image (SPI) ($max\ new\ tokens$ is 512). The best performances within each setting are bolded, and the second-best are underlined.

nation across different methods⁵ (see Appendix H for full experimental results). As illustrated in Figure 5 (a) our approach maintains a Pareto optimal position in the factuality-recall trade-off relative to all other methods. Notably, this robustness in managing the trade-off becomes more pronounced as the sequence length increases. Furthermore, Figure 5 (b) shows that, even when considering object hallucination alone, our method exhibits the lowest degree of object hallucination across all variations of generated token lengths. This result is significant, as it suggests that our method can capture both factual accuracy and detailed explanations across short and long generations. This demonstrates the broad applicability of our method.

Results on Sentence-level Hallucination Ratio.

Table 2 shows that SumGD-S achieves the lowest sentence-level hallucination rate on both the LLAVA 1.5 and InstructBLIP models. Additionally, SumGD-D ranks second on both models. Based on these results, our SumGD method demonstrates strong factual accuracy in holistic hallucination evaluations. OPERA performs comparably

to SumGD-D, but since it relies on beam search, it is less efficient than our method in terms of cost. Moreover, an examination of the Sentences Per Image (SPI) reveals that our method does not achieve favorable results simply by generating fewer sentences.

5 Analysis

5.1 Analysis of SumGD and Contrastive Decoding

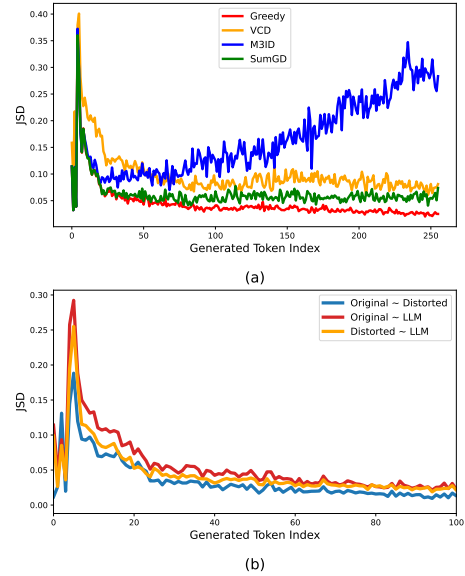


Figure 6: **(Top)** JSD between each method and LLM distributions at each decoding step. **(Bottom)** JSD between the Original Image and LLM, Distorted Image and LLM, Original Image and Distorted Image.

In this section, we analyze SumGD and contrastive decoding, focusing on their relationship with language priors. To explore this, we compute the JSD between each method’s output and LLM

⁵CHAIR is a precision-based metric, which means it can be hacked by generating shorter captions or fewer objects.

Method	Token length 64				Token length 256			
	$C_s \downarrow$	$C_i \downarrow$	$R \uparrow$	$TQ \uparrow$	$C_s \downarrow$	$C_i \downarrow$	$R \uparrow$	$TQ \uparrow$
Greedy	27	7.5	65.3	4.97	67.5	16.7	83.1	4.46
VCD	24.0	7.9	66.1	4.93	82.5	22.0	84.1	4.53
M3ID	20.5	6.5	65.6	4.85	62	13.5	80.3	2.39
SumGD	22.5	6.1	65.0	4.93	54	12.3	83.3	3.75

Table 3: CHAIR metric and Text Quality in various generated token lengths. Denote CHAIR_S as C_S , CHAIR_I as C_I , Recall as R and Text Quality as TQ .

distribution at each decoding step, followed by Section 2.1. For the analysis, we generate descriptions for 200 images from the MSCOCO 2014 validation set using LLAVA 1.5 7B. Factual accuracy is evaluated using the CHAIR metric, while text quality is assessed by GPT-4o (OpenAI, 2024) on a 1 to 5 scale (see details in Appendix G).

Two key questions guide the analysis. **Question1:** *Is significantly deviating from language priors always beneficial?* **Question2:** *Can contrastive decoding reduce hallucinations in LVLMs when language priors heavily influence the two output distributions being contrasted?*

To assess whether significantly deviating from language priors is always beneficial, we examine M3ID, a contrastive decoding method that progressively reduces language priors to focus more on visual information, as shown in Figure 6 (a). However, as presented in Table 3, text quality drops considerably when generating up to 64 tokens compared to 256 tokens. Specifically, it declines from **4.85** to **2.39**, a reduction of about **50.7%**. This suggests that a significant deviation from the language prior disrupts the distribution of language-related tokens, leading to a degradation in text quality.

To investigate the effectiveness of contrastive decoding when language priors significantly influence the original distribution, we investigate VCD. In VCD, the output distribution of the original image prompt is contrasted with that of the distorted image prompt to produce outputs that more align with the original image. A noteworthy observation is that both the output distributions of the original and distorted image prompts progressively converge towards the LLM distribution, as shown in Figure 6 (b). This finding indicates that language priors are influencing both the original output distribution and the output distribution that needs to be compared. Consequently, the two distributions become increasingly similar, diminishing the effectiveness of contrastive decoding. Table 3 demonstrates the reduced effectiveness of contrastive decoding, as VCD results in more instances of object hallucina-

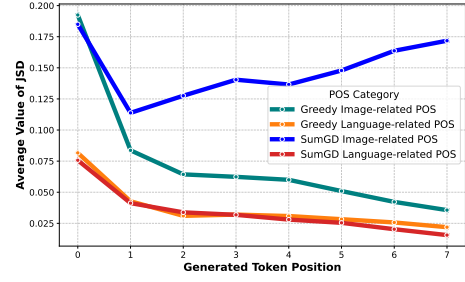


Figure 7: Average JSD values for image-related POS and language-related POS across intervals of 32 tokens, measured in Greedy decoding and SumGD.

tion compared to greedy decoding. Although current contrastive decoding methods focus on distorting the image to create meaningful differences from the original (Leng et al., 2023; Kim et al., 2024c; Wan et al., 2024), the strong influence of language priors may obscure the intended effects of these distortions, undermining the effectiveness of contrastive decoding. This finding is crucial to understanding the limitations of current contrastive decoding approaches.

Unlike contrastive decoding, SumGD excels at reducing object hallucinations while also maintaining a good balance in terms of text quality, as shown in Table 3. To further understand the effectiveness of SumGD, we measure how much SumGD and the Greedy method rely on language priors for each POS type. As seen in Figure 7, SumGD demonstrates a clear reduction in language priors when predicting image-related POS tokens, while preserving the original dependency on language-related POS tokens. These results indicate that our approach effectively mitigates language priors without compromising the core language modeling properties of the LVLm.

5.2 Ablation study of SumGD

In this section, we conduct ablation studies to evaluate the quality of the summary and the effect of POS control in SumGD. For this, we use LLAVA 1.5 7B to generate descriptions for 200 images from the MSCOCO 2014 validation dataset. We employ the CHAIR metric and the text quality metric as described in Section 5.1. Additionally, we include the n-gram fluency indicator (Zhao et al., 2024), represented by $\frac{\text{set}(n\text{grams}(s))}{\text{len}(n\text{gram}(s))}$, where s denotes the description, to measure fluency.

Summary Quality. We conduct an ablation experiment to evaluate the quality of summaries used in SumGD. To achieve this, we employ three distinct summarization models—Distilled-Flan-T5-

	CHAIR _S ↓	CHAIR _I ↓	Recall ↑	Text Quality ↑	1-gram ↑	2-gram ↑
Baseline						
Greedy Decoding	51.5	13.7	79.1	4.9	61.85	92.55
Summary Models						
Distilled-Flan-T5-base(248M)	42.5	11.8	77.8	4.8	59.87	90.57
LLAVA 1.5(7B)	43	11.1	79.1	4.84	60.63	91.41
GPT-4o	43	10.3	78	4.77	59.36	89.77
POS Control in SumGD						
ALL POS	39	10.1	75.8	4.06	52.67	80.66
Image-related POS	43	11.1	79.1	4.84	60.63	91.41

Table 4: Ablation study in terms of Summary Quality and POS Control in SumGD (*max new tokens* is 512).

base, LLAVA 1.5 7B, and GPT-4o. The results, as presented in Table 4, reveal that the effect of summarization quality is consistent across these models in terms of both CHAIR and text quality. This suggests that both SumGD-D and SumGD-S achieve satisfactory levels of summarization quality.

POS Control. We analyze the effect of applying image-related POS control in SumGD. As shown in Table 4, applying SumGD to all POS tokens, as well as selectively to image-related POS tokens, reduces object hallucination compared to the original decoding method. However, when SumGD is applied to all POS tokens, text quality declines compared to the baseline, with the score dropping from **4.9** to **4.06**. This decline is accompanied by a notable decrease in n-gram fluency and object recall, indicating more repetitive generation. In contrast, when SumGD is applied only to image-related POS tokens, the resulting text quality remains almost unchanged, with the score only slightly decreasing from **4.9** to **4.84**. These results demonstrate that applying SumGD selectively to image-related POS tokens effectively preserves the model’s text quality.

6 Related works

Mitigating Language Priors in LVLMS. Large Vision-Language models (LVLMS) extend pre-trained Large Language Models (LLMs) by incorporating visual tokens, enabling them to process visual content (Liu et al., 2023; Dai et al., 2023; Zhu et al., 2023). In LVLMS architectures, the language model is significantly larger than the vision model, creating an imbalanced structure where the language model exerts more significant influence. As a result of this imbalance, the model tends to rely on linguistic patterns rather than adequately considering the visual information provided, a phenomenon known as the language prior problem (Guan et al., 2024; Lee et al., 2024b,c). Several studies have ex-

plored ways to control LLM outputs to better align with desired objectives (Li et al., 2023b; Hallinan et al., 2023; Kim et al., 2023, 2024b). Similarly, research on LVLMS has focused on contrastive decoding techniques to reduce the model’s over-reliance on language priors (Manevich and Tsarfaty, 2024). Visual Contrastive Decoding (VCD) (Leng et al., 2023) works by utilizing distorted images, which amplify the language prior, and Instruction Contrastive Decoding (ICD) (Wang et al., 2024) introduces misleading instructions to achieve a similar effect. Both methods aim to reduce the language prior’s dominance by leveraging these amplified conditions to adjust the model’s behavior. Additionally, Multi-Modal Mutual Information Decoding (M3ID) (Favero et al., 2024) identified that as the token length increases, the model dilutes visual information, leading to a more substantial reliance on language priors. To counter this, M3ID applies more assertive contrastive decoding techniques as the token length grows to calibrate the model’s over-reliance on language priors. However, contrastive decoding can disrupt the distribution of tokens essential for language modeling, leading to a decline in text quality. Additionally, due to the language prior, the two output distributions being contrasted may become more similar, making it less effective in reducing hallucinations. Our method, Summary-Guided Decoding (SumGD), addresses these issues by using summarization techniques to naturally reduce the influence of language priors, allowing the model to focus more on the image. Furthermore, SumGD preserves text quality by controlling only the POS tokens relevant to the image.

7 Conclusion

In this paper, we introduce Summary-Guided Decoding (SumGD) as a novel method to mitigate object hallucinations in LVLMS. Our analysis reveals that as token sequences grow, LVLMS tend to increasingly rely on language priors, reducing the

influence of visual information during the decoding process. To address this, SumGD employs summarization techniques to shorten token length, encouraging the model to incorporate more visual details while controlling only the image-related POS tokens to maintain text quality. Our experimental results demonstrate that SumGD significantly reduces object hallucination and achieves an optimal balance between factual accuracy and recall in both short and long description tasks.

Limitations

In this paper, we propose a Summary-Guided Decoding (SumGD) to mitigate object hallucinations in Large Vision-Language Models (LVLMs). However, this approach comes with some limitations.

First, the generation of summaries during the decoding process incurs additional computational cost, resulting in increased inference time.

Second, while summarization effectively reduces input length and helps mitigate hallucinations, it may also result in the loss of critical contextual information. Additionally, excessively long summaries can increase LVLMs’ dependence on language priors, which may degrade the performance of SumGD. Therefore, it is crucial for future work to train LVLMs in a way that inherently avoids over-reliance on language priors, even when token lengths are extended.

Lastly, we employ part-of-speech (POS) tagging to distinguish between image-related and language-related tokens. However, relying solely on POS tagging for this differentiation can be problematic. The development of more advanced methods for token distinction could enhance the effectiveness of SumGD and create further synergies with this approach.

Ethics Statement

In this paper, we contribute to the future development of a safe and reliable AI community by conducting research focused on reducing hallucinations in Large Vision-Language Models (Xie et al., 2024; Manevich and Tsarfaty, 2024).

Our experiments were conducted by using publicly available datasets, ensuring that no private or sensitive personal data was involved. Furthermore, we utilized publicly accessible models for our experiments, reinforcing the transparency and reproducibility of our approach.

However, the models we used may still exhibit

biases inherent in the underlying datasets and training processes (Howard et al., 2024; Kim et al., 2024a; Fraser and Kiritchenko, 2024; Lee et al., 2024a; Koh et al., 2024). While our focus was on biases related to language priors, we acknowledge the need to address other potential biases as well.

Acknowledgments

This work was partly supported by the Institute of Information & Communications Technology Planning & Evaluation(IITP)-ITRC(Information Technology Research Center) grant funded by the Korea government(MSIT)(IITP-2025-RS-2024-00437633, 60%), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) [NO.RS-2021-II211343, Artificial Intelligence Graduate School Program (Seoul National University)] & RS-2021-II212068, Artificial Intelligence Innovation Hub (Artificial Intelligence Institute, Seoul National University)], and the BK21 FOUR program of the Education and Research Program for Future ICT Pioneers, Seoul National University in 2024. K. Jung is with ASRI, Seoul National University, Korea. The Institute of Engineering Research at Seoul National University provided research facilities for this work.

References

- Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. 2023. [Shikra: Unleashing multimodal llm’s referential dialogue magic](#). *Preprint*, arXiv:2306.15195.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *Preprint*, arXiv:2210.11416.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. [Instructblip: Towards general-purpose vision-language models with instruction tuning](#). *Preprint*, arXiv:2305.06500.
- Alessandro Favero, Luca Zancato, Matthew Trager, Siddharth Choudhary, Pramuditha Perera, Alessandro

- Achille, Ashwin Swaminathan, and Stefano Soatto. 2024. [Multi-modal hallucination control by visual information grounding](#). *Preprint*, arXiv:2403.14003.
- Kathleen Fraser and Svetlana Kiritchenko. 2024. [Examining gender and racial bias in large vision-language models using a novel dataset of parallel images](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 690–713, St. Julian’s, Malta. Association for Computational Linguistics.
- Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. 2024. [Hallusionbench: An advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models](#). *Preprint*, arXiv:2310.14566.
- Skyler Hallinan, Alisa Liu, Yejin Choi, and Maarten Sap. 2023. [Detoxifying text with MaRCO: Controllable revision with experts and anti-experts](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 228–242, Toronto, Canada. Association for Computational Linguistics.
- Phillip Howard, Anahita Bhiwandiwalla, Kathleen C. Fraser, and Svetlana Kiritchenko. 2024. [Uncovering bias in large vision-language models with counterfactuals](#). *Preprint*, arXiv:2404.00166.
- Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming Zhang, and Nenghai Yu. 2024. [Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation](#). *Preprint*, arXiv:2311.17911.
- Liqiang Jing, Ruosen Li, Yunmo Chen, Mengzhao Jia, and Xinya Du. 2023. [Faithscore: Evaluating hallucinations in large vision-language models](#). *Preprint*, arXiv:2311.01477.
- Minbeom Kim, Jahyun Koo, Hwanhee Lee, Joonsuk Park, Hwaran Lee, and Kyomin Jung. 2024a. [Life-Tox: Unveiling implicit toxicity in life advice](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 688–698, Mexico City, Mexico. Association for Computational Linguistics.
- Minbeom Kim, Hwanhee Lee, Kang Min Yoo, Joonsuk Park, Hwaran Lee, and Kyomin Jung. 2023. [Critic-guided decoding for controlled text generation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4598–4612, Toronto, Canada. Association for Computational Linguistics.
- Minbeom Kim, Thibaut Thonet, Jos Rozen, Hwaran Lee, Kyomin Jung, and Marc Dymetman. 2024b. [Guaranteed generation from large language models](#). *Preprint*, arXiv:2410.06716.
- Sihyeon Kim, Boryeong Cho, Sangmin Bae, Sumyeong Ahn, and Se-Young Yun. 2024c. [Vacode: Visual augmented contrastive decoding](#). *Preprint*, arXiv:2408.05337.
- Hyukhun Koh, Dohyung Kim, Minwoo Lee, and Kyomin Jung. 2024. [Can LLMs recognize toxicity? a structured investigation framework and toxicity metric](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 6092–6114, Miami, Florida, USA. Association for Computational Linguistics.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Fei-Fei Li. 2016. [Visual genome: Connecting language and vision using crowdsourced dense image annotations](#). *Preprint*, arXiv:1602.07332.
- Dongryeol Lee, Yerim Hwang, Yongil Kim, Joonsuk Park, and Kyomin Jung. 2024a. [Are llm-judges robust to expressions of uncertainty? investigating the effect of epistemic markers on llm-based evaluation](#). *Preprint*, arXiv:2410.20774.
- Kang-il Lee, Minbeom Kim, Seunghyun Yoon, Minsung Kim, Dongryeol Lee, Hyukhun Koh, and Kyomin Jung. 2024b. [Vlind-bench: Measuring language priors in large vision-language models](#). *Preprint*, arXiv:2406.08702.
- Seongyun Lee, Sue Hyun Park, Yongrae Jo, and Minjoon Seo. 2024c. [Volcano: Mitigating multimodal hallucination through self-feedback guided revision](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 391–404, Mexico City, Mexico. Association for Computational Linguistics.
- Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. 2023. [Mitigating object hallucinations in large vision-language models through visual contrastive decoding](#). *Preprint*, arXiv:2311.16922.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. [Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models](#). *Preprint*, arXiv:2301.12597.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. 2023b. [Contrastive decoding: Open-ended text generation as optimization](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12286–12312, Toronto, Canada. Association for Computational Linguistics.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Xin Zhao, and Ji-Rong Wen. 2023c. [Evaluating object hallucination in large vision-language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in*

- Natural Language Processing*, pages 292–305, Singapore. Association for Computational Linguistics.
- J. Lin. 1991. [Divergence measures based on the shannon entropy](#). *IEEE Transactions on Information Theory*, 37(1):145–151.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. 2015. [Microsoft coco: Common objects in context](#). *Preprint*, arXiv:1405.0312.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. 2024a. [Mitigating hallucination in large multi-modal models via robust instruction tuning](#). *Preprint*, arXiv:2306.14565.
- Hanchao Liu, Wenyuan Xue, Yifei Chen, Dapeng Chen, Xiutian Zhao, Ke Wang, Liping Hou, Rongjun Li, and Wei Peng. 2024b. A survey on hallucination in large vision-language models. *arXiv preprint arXiv:2402.00253*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024c. [Improved baselines with visual instruction tuning](#). *Preprint*, arXiv:2310.03744.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024d. [Llava-next: Improved reasoning, ocr, and world knowledge](#).
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). *Preprint*, arXiv:2304.08485.
- Avshalom Manevich and Reut Tsarfaty. 2024. [Mitigating hallucinations in large vision-language models \(LVLMS\) via language-contrastive decoding \(LCD\)](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 6008–6022, Bangkok, Thailand. Association for Computational Linguistics.
- OpenAI. 2024. [Hello gpt-4o](#).
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2019. [Object hallucination in image captioning](#). *Preprint*, arXiv:1809.02156.
- David Wan, Jaemin Cho, Elias Stengel-Eskin, and Mohit Bansal. 2024. [Contrastive region guidance: Improving grounding in vision-language models without training](#). *Preprint*, arXiv:2403.02325.
- Xintong Wang, Jingheng Pan, Liang Ding, and Chris Biemann. 2024. [Mitigating hallucinations in large vision-language models with instruction contrastive decoding](#). *Preprint*, arXiv:2403.18715.
- Yuxi Xie, Guanzhen Li, Xiao Xu, and Min-Yen Kan. 2024. [V-DPO: Mitigating hallucination in large vision language models via vision-guided direct preference optimization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13258–13273, Miami, Florida, USA. Association for Computational Linguistics.
- Zhiyuan Zhao, Bin Wang, Linke Ouyang, Xiaoyi Dong, Jiaqi Wang, and Conghui He. 2023. [Beyond hallucinations: Enhancing lvlms through hallucination-aware direct preference optimization](#). *Preprint*, arXiv:2311.16839.
- Zhiyuan Zhao, Bin Wang, Linke Ouyang, Xiaoyi Dong, Jiaqi Wang, and Conghui He. 2024. [Beyond hallucinations: Enhancing lvlms through hallucination-aware direct preference optimization](#). *Preprint*, arXiv:2311.16839.
- Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. 2024. [Analyzing and mitigating object hallucination in large vision-language models](#). *Preprint*, arXiv:2310.00754.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. [Minigpt-4: Enhancing vision-language understanding with advanced large language models](#). *Preprint*, arXiv:2304.10592.
- Lanyun Zhu, Deyi Ji, Tianrun Chen, Peng Xu, Jieping Ye, and Jun Liu. 2024. [Ibd: Alleviating hallucinations in large vision-language models via image-biased decoding](#). *Preprint*, arXiv:2402.18476.

A Details of Code, Hyperparameters, and GPU Cost

We conduct our experiments based on the OPERA (Huang et al., 2024) code base which is publicly available. We use the publicly available code provided by the authors for the VCD and OPERA methods. While M3ID and ICD are implemented from scratch due to the lack of public code. For VCD, OPERA, and ICD, we use the hyperparameters as specified in their respective papers. Since only LLaVA 1.5’s hyperparameters were reported in M3ID, we apply these hyperparameters to both LLaVA 1.5 and InstructBLIP for our experiments. Also, we set repetition penalty as 1. All the experiments are conducted using 1 NVIDIA RTX A6000 GPU.

B Experimental Settings for Analyzing Language Priors

We generate descriptions using LLaVA 1.5 7B for 5000 images from the MSCOCO 2014 validation dataset (Lin et al., 2015) and annotate each token to determine whether it represents an object hallucination. An object hallucination is defined as an object not present in the image. We employ the CHAIR metric pipeline (Rohrbach et al., 2019) for evaluation.

C Ablation Study of Inference Time Cost

Method	RIC	C_s	C_i	R
Greedy	1	51.5	13.7	79.1
VCD	2	58.0	16.4	77.8
Beam Search	5	47.5	12.5	79.2
OPERA	5↑	46.0	13.4	78.3
SumGD-S				
+Summarization	2.54	-	-	-
+Summarization + POS Tagging	2.98	43.5	11.0	79.2
SumGD-D				
+Summarization	1.87	-	-	-
+Summarization + POS Tagging	2.3	42.5	11.6	77.7

Table 5: Comparison of Methods with Relative Inference Costs and CHAIR Metrics. Denote $CHAIR_S$ as C_S , $CHAIR_I$ as C_I , Recall as R and Relative Inference Costs as RIC .

We conduct an ablation study to evaluate the impact of the summarization process on inference time. Since the summarization process inherently requires additional token generation, which affects efficiency, we measure inference costs (normalized by relative token generation costs) across baseline methods using the LLaVA 1.5 7B model. Our re-

sults show that the distilled summary model (Flan-T5-base) in SumGD-D requires only half the time per token compared to the LLaVA 1.5 7B model. This finding is incorporated into our inference cost calculations. Additionally, to ensure accurate part-of-speech (POS) tagging, we generate one extra word after each current token and then measure the current token’s POS tag. We perform this ablation study on 200 images. The comparison of inference time costs is presented in Table 5.

D Experiments on LLaVA 1.6

Method	$CHAIR_S$	$CHAIR_I$
Greedy	36.5	10.2
Nucleus	37.0	9.7
Beam-Search (beams=5)	34.5	10.5
SumGD-S	30.5	6.4

Table 6: CHAIR results on LLaVA 1.6 7B (*max new tokens* is 512).

We conduct an evaluation of our SumGD method on the latest model, LLaVA 1.6 7B model (Liu et al., 2024d), to assess its effectiveness. Specifically, we perform the CHAIR evaluation (lower is better) on 200 images. Table 6 shows that SumGD-S effectively reduces object hallucination, demonstrating the applicability of our methodology even with the latest models.

E Distilled Flan-T5-base model

We employ LLaVA 1.5 7B to perform Summary-Guided Decoding with Self-Summarization while generating descriptions for 5,000 images from the MSCOCO dataset. During this process, LLaVA 1.5 iteratively summarizes each previous sentence, and we store each previous sentence along with its corresponding summarized sentence as a pair. This paired dataset is subsequently used to fine-tune the Flan-T5-base model with the prompt “What is a summary of this text?” for training purposes.

F Summarize Prompt for Summary-Guided Decoding

In SumGD-S, we use summary prompt as:

USER: Summarize the following caption in briefly.
 \nCaption: <<caption>> ASSISTANT:

In SumGD-D, we use summary prompt as:
 <<Caption>> \nWhat is a summary of this text?

G GPT-4o Prompt for text quality evaluation

<p>###Task Description:</p> <p>You will be given one caption written for a given image. Your task is to rate the caption on one metric. Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed. The output format should look as follows: Score: [RESULT] (an integer number between 1 and 5). Please do not generate any other opening, closing, and explanations.</p> <p>###Evaluation Criteria:</p> <p>Text Quality (1-5) - Evaluate how well-written the caption is. A high-quality caption is clear, concise, grammatically correct, and well-structured.</p> <p>###Evaluation Steps:</p> <ol style="list-style-type: none"> 1. Read the caption carefully and evaluate its clarity, grammar, and overall readability. 2. Check for any awkward phrasing, grammatical errors, or unnecessary complexity. 3. Assign a score for text quality on a scale of 1 to 5, where 1 is the lowest and 5 is the highest based on the Evaluation Criteria. <p>###Given Caption:</p> <p>{{Caption}}</p> <p>###Score:</p>

Figure 8: GPT-4o prompt for text quality evaluation

Token Length	Method	CHAIRs	CHAIRi	Recall
64	Greedy	27	7.5	65.3
64	Nucleus	31.5	9.8	58.9
64	Beam	20	5.9	62.5
64	VCD	24.0	7.9	66.1
64	ICD	21.5	7.0	62.2
64	M3ID	20.5	6.5	65.6
64	OPERA	22.5	7.1	62.3
64	SumGD-S	22.5	6.1	65.0
64	SumGD-D	24	6.7	64.8
128	Greedy	53	13.1	78.9
128	Nucleus	56.5	16.5	74.2
128	Beam	50.5	13.3	78.3
128	VCD	63.0	17.5	78.4
128	ICD	56.0	15.1	77.3
128	M3ID	46.5	11.6	76.4
128	OPERA	49.5	14.4	79.2
128	SumGD-S	43.5	10.5	78.1
128	SumGD-D	43.5	11.4	78.0
256	Greedy	67.5	16.7	83.1
256	Nucleus	78	20.9	82.8
256	Beam	70	16.2	81.6
256	VCD	82.5	22.0	84.1
256	ICD	71	19.6	83.0
256	M3ID	62	13.5	80.3
256	OPERA	64.5	16.3	83.4
256	SumGD-S	54	12.3	83.3
256	SumGD-D	56.5	12.4	81.9
512	Greedy	69.5	17.4	84.1
512	Nucleus	80	22.0	83.8
512	Beam	71.5	17.4	82.3
512	VCD	83.0	23.6	85.6
512	ICD	73.0	20.2	83.8
512	M3ID	65.5	14.6	81.2
512	OPERA	66.5	17.5	83.4
512	SumGD-S	59	13.1	83.8
512	SumGD-D	61.5	12.6	82.5

Table 7: Performance comparison for CHAIRs, CHAIRi, and Recall

I Case Study

This is the case study of Summary-Guided Decoding (SumGD), Visual Contrastive Decoding (VCD), and Multi-Modal Mutual Information Decoding (M3ID) in generating up to 256 tokens in detailed captioning task. Case study shows that SumGD generated image-related words, while VCD showed a tendency to hallucinate by relying on the word ‘tie’ during decoding. Additionally, M3ID exhibited issues in language modeling.

H CHAIR metric on various token length

In this section, we report CHAIR metric based on various generated token length.



The image features a yellow and green propeller airplane parked on a runway. The airplane is a small, single-engine aircraft, and it appears to be a vintage model. There are two people visible in the scene, one standing near the front of the airplane and another person closer to the middle.

The two people seem to be inspecting the aircraft, possibly discussing its condition or preparing for a flight. The airplane is the main focus of the scene, and the two individuals appear to be the only people present in the image. The airplane is parked on the runway, and the two people are standing close to it, likely examining its features and ensuring it is in good working order. The airplane's wings are clearly visible, and the propeller is spinning, indicating that it is ready for takeoff. The scene captures the essence of aviation and the bond between the people and the aircraft. The runway is the main focus of the image, and the airplane is the central object of interest. The two individuals are the main focus of the image, as they are the ones interacting with the airplane. The runway is the central element of the image

SumGD

Figure 9: SumGD case study.



The scene features a small yellow and green propeller plane on an airport runway. A man is sitting in the cockpit, preparing to take off or having just landed. There is another person standing near the plane, possibly an airport worker or a passenger assisting with the pre-flight or post-flight procedures.

In the background, a field is visible, indicating that the airport is located near a rural or agricultural area. The presence of the runway, plane, and people suggests that the scene is likely taking place at an airfield where small aircraft like this one are used for recreational or commercial purposes. Another notable detail is a **tie** located close to the airplane, possibly belonging to the man in the cockpit. The **tie** is likely a personal accessory and not part of the plane's equipment. A tie can also signify a formal or professional setting, which may indicate the importance or seriousness of the flight. The presence of a **tie** suggests that the man piloting the aircraft may be a professional pilot or someone with significant experience in handling small planes. Overall, the scene portrays a moment of preparation or completion of a flight involving the small yellow and green plane on

VCD

Figure 10: VCD case study. VCD heavily relies on word "tie" to generate descriptions which is not in the provided image.



The image features a yellow and green propeller airplane parked on a runway. The airplane is an old-style yellow and green fighter plane, possibly a World War II model. Two people are visible in the scene, one standing inside the cockpit and another person sitting in the back of the plane. They seem to be enjoying their time on the ground before taking off.

The grassy field surrounding the runway showcases the daytime setting, making it an ideal place for observing and admiring the vintage aircraft. The cloudy sky also adds to the atmosphere of the scene. The J-3 NR number on the green greenwing portion of the plane suggests that it maintains its historical significance and history.

Overall, the image captures a moment of unique American aviation history. The cloudy sky ranges from occupying the center to the cloudy shapes observed in the background. The **cornfield望fields望** lies further distancedfrom the ground strip. The cloudy sky ranges from occupying the center to the cloudy shapes observed in the background. The **cornfield望fields望** lies further distancedfrom the ground strip. The cloudy sky ranges from occupying the center to the cloudy shapes observed

M3ID

Figure 11: M3ID case study. Underline is for repetitive sentences. Red font denotes a degradation of language modeling.