
Revisit What You See: Disclose Language Prior in Vision Tokens for Efficient Guided Decoding of LVLMs

Beomsik Cho¹ Jaehyung Kim¹

Abstract

Large Vision-Language Models (LVLMs) have demonstrated remarkable performance across various multimodal tasks by integrating visual perception with language understanding. However, conventional decoding strategies of LVLMs often fail to successfully utilize visual information, leading to visually ungrounded responses. While various approaches have been proposed to address this limitation, they typically require additional training, multi-step inference procedures, or external model dependencies. This paper introduces ReVisiT, a simple yet effective decoding method that references vision tokens to guide the text generation process in LVLMs. Our approach leverages the semantic information embedded within vision tokens by projecting them into the text token distribution space, and dynamically selecting the most relevant vision token at each decoding step through constrained divergence minimization. This selected vision token is then used to refine the output distribution to better incorporate visual semantics. Experiments on three LVLM hallucination benchmarks with two recent LVLMs demonstrate that ReVisiT consistently enhances visual grounding with minimal computational overhead. Moreover, our method achieves competitive or superior results relative to state-of-the-art baselines while reducing computational costs for up to 2×. The code is available at <https://github.com/bscho333/ReVisiT>.

1. Introduction

With the recent success of Large Language Models (LLMs; (Touvron et al., 2023a;b; Achiam et al., 2023; Team et al., 2023)), Large Vision-Language Models (LVLMs; (Dai et al.,

2023; Liu et al., 2023; Ye et al., 2023; Zhu et al., 2024; Chen et al., 2023; Bai et al., 2023; Chen et al., 2024a; Liu et al., 2024a; Bai et al., 2025)) have emerged as powerful multimodal architecture that integrate visual perception with language understanding and text generation. LVLMs typically incorporate visual information by encoding images into embeddings using a vision encoder, then mapping them to *vision tokens* via a cross-modality projector aligned with LLM decoder’s linguistic embedding space. Then, LVLMs utilize these vision tokens alongside text tokens during the text generation in autoregressive manner. These models have demonstrated impressive performance on foundational vision-language tasks such as Visual Question Answering (VQA; (Antol et al., 2015; Li et al., 2023a; Selvaraju et al., 2017)) and image captioning (Vinyals et al., 2015; Xu et al., 2015; You et al., 2016), while also showing promising results in more complex reasoning tasks, *e.g.*, multimodal conversational agents (Sundar & Heck, 2022). Despite these advances, the decoding process of LVLMs largely follows conventional LLM decoding strategies without specific consideration for vision tokens. This insufficient extraction and utilization of visual information often result in visually ungrounded responses (Zhou et al., 2023; Yin et al., 2024; Leng et al., 2024; Favero et al., 2024; Huo et al., 2025; Kang et al., 2025; Rohrbach et al., 2018; Li et al., 2023c; Wang et al., 2023).

Prior efforts to address this limitation can be broadly categorized based on *when* and *how* additional visual information is incorporated relative to the decoding process. Specifically, we distinguish between (1) *pre-alignment*, (2) *post-alignment*, and (3) *intra-alignment*. Pre-alignment methods aim to enhance visual information before the decoding phase, typically by re-designing model architectures (Dai et al., 2024) or developing specialized visual tokenization techniques (You et al., 2024; Lu et al., 2024), but they generally require additional training and model modification. Post-alignment strategies revise the generated outputs through external models, such as object detectors or proprietary models (Zhou et al., 2023; Yin et al., 2024), inevitably introducing complexity and external dependencies. Due to these inherent overheads of pre- and post-alignment methods, researchers have increasingly focused on intra-alignment approaches, which modify attention mechanisms

¹Yonsei University. Correspondence to: Jaehyung Kim <jae-hyun@yonsei.ac.kr>.

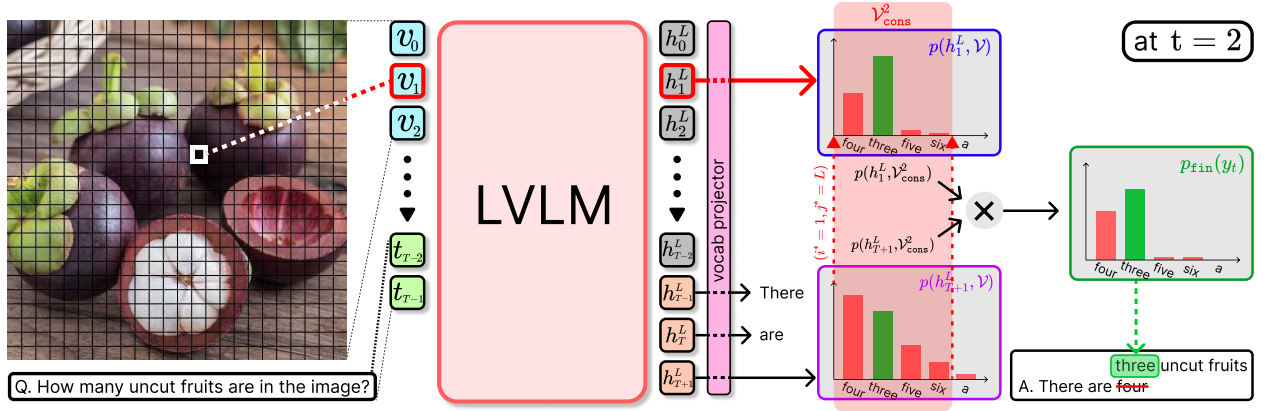


Figure 1. **An overview of ReVisiT.** Given an input image and text prompt, the LVLM first encodes the image into vision tokens through a vision encoder and a cross-modal projector. ReVisiT re-purposes these vision tokens as reference informers to guide the text generation process. At each decoding step, ReVisiT (1) constrains the vocabulary \mathcal{V} to $\mathcal{V}_{\text{cons}}^t$, (2) projects vision token embeddings into $\mathcal{V}_{\text{cons}}^t$ and selects most relevant token, and (3) refines the final output distribution. Specifically, ReVisiT adoptively leverages vision tokens to serve as reference signals for decoding, enhancing visual grounding without additional forward passes or external supervision.

or output distributions during inference (Leng et al., 2024; Favero et al., 2024; Huo et al., 2025) to improve visual grounding without altering model parameters. However, these methods still incur notable computational overhead due to multi-step procedures such as additional inference under input perturbation or recalibration of attention scores. Consequently, these limitations highlight the need for a streamlined approach that can better utilize visual information during LVLM decoding without additional training, external models, or complex multi-step inference while maintaining computational efficiency.

Contribution. In this work, we propose a simple yet effective decoding method by **Referencing Vision** tokens to guide the **Text** generation process of LVLMs (**ReVisiT**). Vision tokens inherently contain semantic information aligned with textual representations, as they are mapped to the linguistic embedding space through the cross-modality projector. Despite this, they are typically only utilized as the additional input context as same as text tokens, which might be insufficient to use their rich semantic information. Motivated by this, our key insight is to explicitly extract this semantics embedded in vision tokens and leverage them as internal reference signals during decoding. However, directly leveraging these embeddings could be difficult due to their inherent noise and complexity.

To address this challenge, ReVisiT identifies the vision token that best aligns with the current output distribution by minimizing a constrained divergence at each decoding step. Then, ReVisiT uses it as a reference logit to refine the output distribution to better incorporate visual semantics (see Figure 1 for the overview). Specifically, ReVisiT first projects the embeddings of all vision tokens into the text token distri-

bution space. Since vision tokens remain fixed throughout the decoding process, this projection only needs to be computed once before the decoding. The projected probability distributions are constrained to top- K text tokens in the current output distribution, to focus on semantically meaningful ones. ReVisiT then selects the vision token which minimize the Jensen-Shannon Divergence (JSD) between its constrained distribution and the original output distribution, as the most relevant visual guide for each decoding step. Finally, ReVisiT refines the output token distribution by re-weighting it using the constrained distribution of the selected vision token.

We validate the effectiveness of ReVisiT on multiple hallucination benchmarks, including CHAIR (Rohrbach et al., 2018), POPE (Li et al., 2023c), and AMBER (Wang et al., 2023), using LLaVA-1.5-7B (Liu et al., 2024a) and Qwen2.5-VL-7B (Bai et al., 2025) as the base LVLMs. Across all benchmarks, ReVisiT consistently improves performance over the vanilla greedy decoding baseline, demonstrating enhanced visual grounding capabilities with only a **1-3%** increase in inference time. Furthermore, our method shows competitive or superior results relative to state-of-the-art intra-alignment methods (Leng et al., 2024; Favero et al., 2024; Huo et al., 2025; Chuang et al., 2024) with substantially reduced computational costs as these baselines typically incur up to 100% additional computation overhead. These results highlight that appropriate selection of vision tokens enables them to serve as effective informers, enhancing the visual grounding capabilities of LVLMs while offering a practical free lunch.

2. Preliminary and Motivation

We first describe our problem setup formalizing the conventional decoding process of LVLMs in Section 2.1. Then, in Section 2.2, we motivate our approach by identifying the limitations of existing intra-alignment and early exit methods, and propose leveraging vision token representations as internal reference signals for visual grounding. Finally, Section 2.3 provides qualitative analyses demonstrating that vision token hidden states indeed encode meaningful visual semantics, supporting their use during decoding. Detailed discussions of related prior works are provided in Appendix A.1.

2.1. Problem setup

Let’s denote the given input text tokens (e.g., query) and vision tokens as \mathbf{x} and \mathbf{v} , respectively. Specifically, the vision tokens \mathbf{v} are generated from the given input image, by forwarding it into vision encoder and cross-modality projector sequentially (Liu et al., 2024a; Bai et al., 2025). Then, our goal is to generate output text tokens \mathbf{y} using LVLM \mathcal{M} ; following the auto-regressive nature of LLM, output tokens \mathbf{y} (e.g., response) are sequentially generated, i.e., $y_t \sim \mathcal{M}(\cdot | \mathbf{v}, \mathbf{x}, \mathbf{y}_{<t-1})$. Let assume that LVLM \mathcal{M} has L decoding layers and text token vocabulary \mathcal{V} , and denote its hidden state of l -th layer as h^l . Then, at time step t , the output probability over \mathcal{V} is calculated as below:

$$\begin{aligned} p(h_{T+t-1}^L, \mathcal{V}) &= \text{softmax}_{\mathcal{V}}(\phi(h_{T+t-1}^L)), \\ \sum_{w \in \mathcal{V}} p(w | h_{T+t-1}^L, \mathcal{V}) &= 1, \end{aligned} \quad (1)$$

where $T = |\mathbf{v}| + |\mathbf{x}|$, ϕ is the language modeling head which maps final hidden state of LVLM into logits over vocabulary \mathcal{V} , and $\text{softmax}_{\mathcal{V}}$ indicates a softmax operation over \mathcal{V} . Then, the output token y_t is sampled from the obtained distribution:

$$y_t \sim p(h_{T+t-1}^L, \mathcal{V}). \quad (2)$$

2.2. Motivation from prior works

To better align LVLM’s responses with visual semantics in given image, prior works (Leng et al., 2024; Favero et al., 2024; Huo et al., 2025) have explored *intra-alignment methods*. They adjust the model’s output logits (Eq. 2) using the reference logits during decoding. Formally, the decoding process of these methods can be described as:

$$y_t \sim f(p(h_{T+t-1}^L, \mathcal{V}), p(\tilde{h}_{T+t-1}^L, \mathcal{V})), \quad (3)$$

where \tilde{h} denotes the hidden state to obtain the reference logit, and f denotes the logit adjustment operation, such as logit addition, subtraction, or interpolation. For example, VCD (Leng et al., 2024) adds random noise into input

image and forward it to obtain a distorted output logits. M3ID (Favero et al., 2024) generates vision-agnostic output logits to identify language priors derived from LLM decoder, and SID (Huo et al., 2025) constructs the reference output logits from low-attention token subsets. While these approaches are effective at aligning visual semantics, they commonly require extra computational cost as reference information is obtained through additional inference passes or recalibration. These observations lead us to the following research question: (RQ) *How can we find reference information that is relevant to the visual input to guide decoding and improve visual grounding efficiently?*

Interestingly, recent studies in the NLP domain offer a compelling alternative. Rather than relying on externally induced reference signals, the prior works about *early exit* (Teerapittayanon et al., 2016; Elbayad et al., 2020; Schuster et al., 2022) revealed that intermediate hidden states already encode semantically meaningful information. Specifically, early exit projects intermediate hidden states h^j into logits over vocabulary \mathcal{V} using the model’s language modeling head ϕ , which can be formally expressed as $p(h_{T+t-1}^j | \mathcal{V})$; here, j denotes the index of an intermediate layer. Motivated by this, DoLa (Chuang et al., 2024) have demonstrated that adjusting logits based on this early exit can effectively resolve output uncertainty during decoding. Roughly, the decoding with early exit can be generalized in the following formula:

$$y_t \sim f(p(h_{T+t-1}^L, \mathcal{V}), p(h_{T+t-1}^j, \mathcal{V})), \quad (4)$$

where the intermediate layer j is selected by maximizing the Jensen-Shannon Divergence (JSD) between the intermediate and final layer logits (Chuang et al., 2024), to identify the layer with the greatest predictive uncertainty. Despite its original design for language-only tasks, DoLa (Chuang et al., 2024) also proves effective for improving visual grounding in LVLMs without any model retraining or architectural changes. These observations indicate that visually relevant auxiliary signals may reside within the model itself, without requiring external augmentation like aforementioned intra-alignment methods.

Prior early exit works have primarily focused on adjusting output distributions based on intermediate representations corresponding to the language modality. Namely, they largely overlook the vision tokens attended by the decoder. However, as LVLMs encode visual information through vision tokens alongside text tokens in the same decoding sequence, intermediate representations of vision tokens may also contain rich visual semantics that are highly relevant for grounding. Moreover, leveraging vision token embeddings offers a distinct advantage: they directly originate from the image input, providing an immediate and explicit pathway to extract visually grounded signals. Therefore, we shift our focus towards vision tokens, investigating whether their

intermediate hidden representations can serve as effective internal reference signals for visual grounding by projecting them into logits.

2.3. Deriving vision tokens as reference signals

While early exit (Teerapittayanon et al., 2016; Elbayad et al., 2020; Schuster et al., 2022) has initially explored to accelerate inference, *logit-lens* (nostalgebraist, 2020) employed it to analyze the semantic information encoded in text token hidden states. In this work, we refer logit-lens as the projection of vision tokens’ intermediate representation into the text token distribution space for convenience, although our use case differs from its original intent. Specifically, given a vision token hidden state h_i^j from layer j and vision token index i , we project it into the text token distribution over vocabulary \mathcal{V} , which yields $p(h_i^j, \mathcal{V})$.

Under this definition, we first conducted analysis by applying logit-lens to the vision tokens \mathbf{v} at the final decoder layer L . As shown in Figure 2 (left), directly applying logit lens to vision tokens over full vocabulary \mathcal{V} results in diffuse and uninformative token distributions. However, manually constraining the vocabulary \mathcal{V} to the subset \mathcal{V}_{env} in Figure 2 effectively reveals the relevant visual semantics, as shown by the top-ranked text tokens in the distribution $p(h_i^j, \mathcal{V}_{\text{env}})$. This observation suggests that visual semantics are encoded in visual tokens and can provide reference signals based on that to guide decoding, if one can restrict logit-lens to the proper subset in vocabulary.

While manually defined semantic subsets is considerable as demonstrated in Figure 2 (left), it is not scalable nor adaptive across diverse images and tasks. To resolve this limitation, we draw inspiration by analyzing the inherent decoding dynamics of LVLM with hallucination cases. As illustrated in Figure 2 (right), we observe that hallucinations typically occur when the greedy top-1 prediction diverges from the visual content. However, the top-5 candidate tokens at hallucinatory starting point often contain visually grounded alternatives such as "mountain", "landscape" or "painting". This result indicates that the LVLM inherently contains grounded candidates and has a chance to avoid the hallucination, even if the top-1 prediction is incorrect. Thus, at each decoding time step, we construct the constrained vocabulary by selecting the high probable candidates from the vanilla output distribution $p(h_{T+t-1}^L, \mathcal{V})$. Then, by projecting the vision token representation onto these dynamically selected subsets, we can extract visually grounded reference signals without external inference or heavy computational overhead.

3. ReVisiT: Referencing Vision Tokens to Guide Text Generation

Based on insights in Section 2.3, we propose ReVisiT, a decoding method that adaptively leverages intermediate vision token representations as internal reference signals to guide text generation, enhancing visual grounding with minimal computational overhead. At each decoding step, ReVisiT focuses the decoding process onto a context-relevant vocabulary subset and refines the output by referencing the most relevant vision token. To be specific, ReVisiT consists of the following three main steps: (1) adaptively constraining the vocabulary based on the vanilla output distribution, (2) projecting and selecting vision token hidden states within the subset, and (3) refining the output distribution through element-wise combination and normalization to generate the next token.

3.1. Adaptive context-relevant vocabulary subset construction

At each decoding timestep t , LVLM \mathcal{M} generate output distributions $p(h_{T+t-1}^L, \mathcal{V})$ over the entire vocabulary space \mathcal{V} . However, as observed in Figure 2, the original distribution is often diffusely spread across many irrelevant tokens. Therefore, we dynamically constrain a vocabulary \mathcal{V} into subset $\mathcal{V}_{\text{cons}}^t$ more suitable to capture the plausible semantic space of the output distribution. By restricting the output space to a smaller, contextually focused subset, this constraint reduce semantic distraction and facilitate more precise integration of visual information. To be specific, following (Li et al., 2023b), we adaptively define the vocabulary subset $\mathcal{V}_{\text{cons}}^t$ as below:

$$\mathcal{V}_{\text{cons}}^t = \{w \in \mathcal{V} : p(w|h_{T+t-1}^L, \mathcal{V}) \geq \alpha \cdot \max_{w'} p(w'|h_{T+t-1}^L, \mathcal{V})\},$$

where $\alpha \in (0, 1)$ controls the sparsity of the selected candidates. This vocabulary subset $\mathcal{V}_{\text{cons}}^t$ serves as a dynamically adapted output space, which all subsequent vision token projections and refinements are performed to enhance visual grounding without introducing additional computation.

3.2. Vision token projection and selection

Given the dynamically constrained vocabulary subset $\mathcal{V}_{\text{cons}}^t$ at each decoding timestep t , we next describe how vision token hidden states h_i^j are projected and selected within this constrained space. Let \mathcal{J} as the set of candidate decoder layers where vision token hidden states are selected. Specifically, for each h_i^j , where i denotes the index of the vision tokens and j denotes the candidate layer, we obtain the projected distribution over the constrained vocabulary $\mathcal{V}_{\text{cons}}^t$:

$$p(h_i^j, \mathcal{V}_{\text{cons}}^t), \quad \text{for } i \in \{0, \dots, |\mathbf{v}| - 1\}, j \in \mathcal{J}. \quad (5)$$

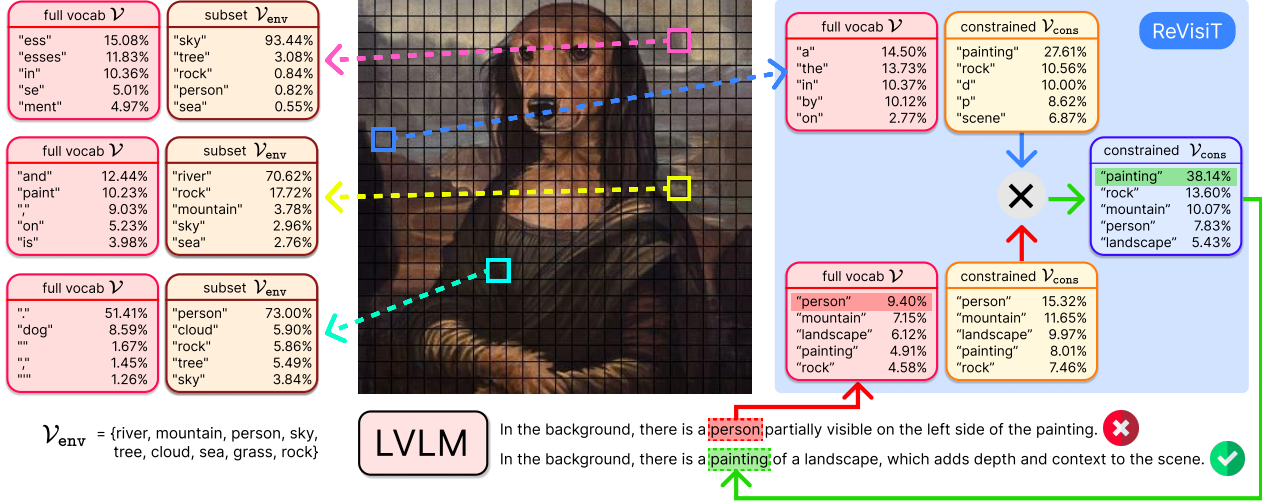


Figure 2. Motivation of ReVisiT. We qualitatively analyzed various vision tokens. Dotted arrows represent vision token projection over specified vocabulary set. For each box, representing text token distribution, we annotated top-5 probable text tokens. Left part illustrate the effectiveness of vocabulary constraint, whereas right part shows the distribution shift during ReVisiT. See Appendix C.1 for a detailed discussion of the underlying values and analysis.

Intuitively, this distribution represents the likelihood of each candidate token within \mathcal{V}_{cons}^t conditioned on the vision token embedding h_i^j . Next, to identify the vision token most relevant to the current decoding context, we compare each vision token distribution $p(h_i^j, \mathcal{V}_{cons}^t)$ against the vanilla output distribution $p(h_{T+t-1}^L, \mathcal{V}_{cons}^t)$ by computing the Jensen-Shannon Divergence (JSD):

$$(i^*, j^*) = \arg \min_{i,j} \text{JSD} \left(p(h_{T+t-1}^L, \mathcal{V}_{cons}^t) \parallel p(h_i^j, \mathcal{V}_{cons}^t) \right).$$

Here, i^* and j^* denote the indices of the selected vision token and decoder layer, respectively. We adopt JSD as the divergence metric due to its symmetric properties and bounded output range.

Remark. For the efficient implementation, the vision token projections over the full vocabulary \mathcal{V} , $p(h_i^j, \mathcal{V})$, are precomputed once before decoding begins. Then, at each timestep t , we apply slicing and masking to adapt the cached projections to the dynamically constrained subset \mathcal{V}_{cons}^t , enabling efficient selection without additional forward computation.

3.3. Output logit refinement using distribution from selected vision token

After identifying the most contextually relevant vision token at each decoding step, we then refine the vanilla output distribution guided by the reference visual grounding signal from the selected vision token. Specifically, given the vanilla output distribution $p(h_{T+t-1}^L, \mathcal{V}_{cons}^t)$ and the selected vision token distribution $p(h_{i^*}^{j^*}, \mathcal{V}_{cons}^t)$, we combine them through

element-wise multiplication:

$$p_{fin}(y_t) \propto \begin{cases} p_{y_t}(h_{T+t-1}^L, \mathcal{V}_{cons}^t) \times p_{y_t}(h_{i^*}^{j^*}, \mathcal{V}_{cons}^t), & \text{if } y_t \in \mathcal{V}_{cons}^t \\ 0, & \text{otherwise.} \end{cases}$$

The combined distribution p_{fin} is subsequently normalized over \mathcal{V}_{cons}^t to form a valid probability distribution, and the next output token y_t is sampled from the normalized p_{fin} . This refinement mechanism allows the decoding process to dynamically integrate visual grounding signals extracted from vision tokens, without introducing additional inference passes or architectural modifications. The overall algorithm of ReVisiT is presented in Algorithm 1.

4. Experiments

4.1. Setups

Datasets and evaluation metrics. To comprehensively assess the visual grounding capability of our method, we consider both generative and discriminative aspects of hallucination mitigation. Specifically, we evaluate our approach on three different benchmarks used to measure robustness and generalizability across vision-language tasks. 1) **CHAIR** (Rohrbach et al., 2018) measures hallucination rates in image captioning by comparing generated captions to ground-truth object annotations from MS COCO (Lin et al., 2014) dataset. We report sentence-level ($CHAIR_S$) and instance-level ($CHAIR_I$) hallucination scores on 500 sampled images. In addition, we also report *Recall*, computed as the proportion of ground-truth objects correctly mentioned in the generated captions. 2) **POPE** (Li et al., 2023c) assesses visual grounding by posing binary ob-

Table 1. **Results on CHAIR benchmark.** Lower scores (\downarrow) on CHAIR_S, CHAIR_I and higher (\uparrow) recall indicate better performance. The best results in each setting are **bolded**, and the second-best are underlined.

METHOD	LLAVA-1.5			QWEN2.5-VL		
	CHAIR _S \downarrow	CHAIR _I \downarrow	RECALL \uparrow	CHAIR _S \downarrow	CHAIR _I \downarrow	RECALL \uparrow
GREEDY	53.8	14.66	80.81	<u>32.2</u>	8.51	<u>71.79</u>
DoLA	53.6	14.45	80.23	<u>35.2</u>	<u>8.14</u>	<u>69.80</u>
VCD	<u>52.8</u>	15.61	80.17	37.8	9.33	<u>71.79</u>
M3ID	57.0	16.57	81.96	32.8	9.40	68.78
SID	50.6	<u>13.58</u>	81.00	36.3	9.99	72.48
OURS	50.6	13.43	<u>81.45</u>	31.6	7.56	70.83

Table 2. **Average results on POPE benchmark.** Average for 3 datasets and 3 types. Higher scores (\uparrow) on Accuracy, Precision, and F1 indicate better performance. The best results are **bolded**, and the second-best are underlined. Full results are presented in Table 5.

METHOD	LLAVA-1.5			QWEN2.5-VL		
	ACCURACY \uparrow	PRECISION \uparrow	F1 \uparrow	ACCURACY \uparrow	PRECISION \uparrow	F1 \uparrow
GREEDY	79.47	74.48	82.36	84.49	93.34	82.82
DoLA	79.58	74.67	82.41	81.40	93.71	78.45
VCD	77.74	72.51	80.97	84.69	92.85	83.17
M3ID	79.48	74.48	82.37	83.93	92.32	82.31
SID	<u>80.24</u>	<u>75.61</u>	<u>82.84</u>	<u>84.76</u>	92.73	83.30
OURS	81.80	78.03	83.45	84.77	<u>93.41</u>	<u>83.18</u>

ject presence questions. We evaluate on MS COCO, A-OKVQA (Schwenk et al., 2022), and GQA (Hudson & Manning, 2019) datasets constructing 3,000 samples per dataset and query type (random, popular, adversarial), with an equal split between positive and negative instances. We report *accuracy*, *precision*, and *F1 score*. 3) **AMBER** (Wang et al., 2023) provides fine-grained object-level datasets for both generative and discriminative evaluations of hallucination mitigation. We focus on the generative setting, reporting CHAIR, Cover, Hal, and Cog metrics over 1,004 samples to further analyze fine-grained. 4) **LLaVA-Bench-In-the-Wild** (Liu et al., 2023) consists of 24 diverse real-world images paired with 60 open-ended questions. We use this dataset with LLaVA-1.5-7B to qualitatively assess the ReVisiT’s robustness in complex, uncontrolled visual environments. Further dataset details and evaluation setups are provided in Appendix A.2.

Models and baselines. To validate the generality of our method across different architectures and capabilities, we conduct experiments using two representative LVLMs: LLaVA-1.5-7B (Liu et al., 2024a) as a widely adopted conventional baseline, and Qwen2.5-VL-7B (Bai et al., 2025) as a representative of state-of-the-art. As baselines, we adopt various intra-alignment decoding strategies to compare the effectiveness of our method in enhancing visual grounding. (a) *Greedy decoding* is a standard decoding baseline, and (b) *DoLa* (Chuang et al., 2024) improves factuality by contrasting output logits from earlier and later decoder layers at each timestep, and has been widely adopted for hallucina-

tion mitigation in LVLMs despite being originally proposed for LLMs. For intra-alignment baselines, (c) *VCD* (Leng et al., 2024) mitigates hallucination by perturbing visual inputs with Gaussian noise and enforcing consistency between the original and perturbed outputs; (d) *M3ID* (Favero et al., 2024) contrasts outputs with and without visual input to alleviate language prior bias and enhance visual grounding; (e) *SID* (Huo et al., 2025) applies contrastive decoding on low-attention vision patches to mitigate hallucination and refine visual grounding.

Implementation details. In all experiments, we apply greedy decoding (*i.e.*, temperature is 0) with a maximum output length of 512 tokens for all decoding methods. Regarding baselines, we follow the default hyperparameters provided by their original implementations. For ReVisiT, we set the hyperparameters as follows: for LLaVA-1.5-7B, we select vision tokens mostly from the last decoder layer (“last”). For Qwen2.5-VL-7B, we select vision tokens from all even-numbered decoder layers (“all”). The APC threshold α is set for each task. More details are presented in Appendix A.3.

4.2. Main results

We first present the experimental results on CHAIR benchmark in Table 1. Notably, all baseline methods, including strong intra-alignment baselines (SID, M3ID, VCD) that explicitly attempt to ground visual semantics during decoding, underperform relative to vanilla greedy decoding

Table 3. **Results on AMBER benchmark.** Lower scores (\downarrow) on CHAIR, Hal, Cog, and higher scores (\uparrow) on Cover indicate better performance. The best result for each metric is **bolded**, and the second-best is underlined.

METHOD	LLAVA-1.5				QWEN2.5-VL			
	CHAIR \downarrow	COVER \uparrow	HAL \downarrow	COG \downarrow	CHAIR \downarrow	COVER \uparrow	HAL \downarrow	COG \downarrow
GREEDY	7.2	50.6	32.7	3.8	<u>4.8</u>	64.0	27.8	1.6
DoLA	7.2	50.9	32.9	3.9	<u>5.2</u>	65.4	31.3	1.9
VCD	8.7	51.7	40.5	4.3	5.8	<u>64.3</u>	31.3	1.5
M3ID	7.2	57.0	40.7	<u>3.0</u>	4.9	53.4	21.8	0.9
SID	<u>6.3</u>	51.9	<u>32.8</u>	3.3	4.9	62.6	28.3	<u>1.1</u>
OURS	6.2	<u>55.2</u>	33.0	2.7	4.4	64.0	<u>25.2</u>	1.4

on Qwen2.5-VL, highlighting the challenge of hallucination mitigation on stronger LVLMS. In contrast, ReVisiT consistently reduces the hallucination rates and achieves the best performance across both sentence-level (CHAIR_S) and instance-level (CHAIR_I) metrics for LLaVA-1.5 (Liu et al., 2024a) and Qwen2.5-VL (Bai et al., 2025). To be specific, compared to vanilla greedy decoding, ReVisiT reduces CHAIR_S by 5.95% and CHAIR_I by 8.39% on LLaVA-1.5, and achieves 1.86% and 11.16% reductions respectively on Qwen2.5-VL. Moreover, ReVisiT also attains competitive Recall scores, achieving the second-highest object coverage among all methods while effectively suppressing hallucinations.

Next, we present the results on POPE benchmark. Since POPE evaluates hallucination in a discriminative manner covering three datasets and three object types (random, popular, and adversarial), it poses a significant challenge for models to achieve robust and consistent performance across diverse settings. Table 2 summarizes the average accuracy across all 18 evaluation settings. Here, ReVisiT consistently achieves superior results, further confirming its effectiveness in mitigating hallucination across diverse object types and datasets. Particularly on GQA dataset under the popular and adversarial object settings, ReVisiT demonstrates its effectiveness even under more visually and semantically challenging scenarios, achieving absolute accuracy gains of over 5 percentage points compared to the second-best methods when using LLaVA-1.5. In addition, as shown in Table 5, ReVisiT outperforms vanilla greedy decoding across all 18 evaluation settings, demonstrating its strong robustness. In contrast, all existing baselines occasionally underperform vanilla greedy sampling which indicates achieving consistent improvements under POPE’s discriminative evaluation setting is non-trivial.

Furthermore, we evaluate ReVisiT and baseline approaches using the AMBER benchmark that provides a more recent and diverse evaluation of hallucination phenomena. Table 3 summarizes the evaluation results. Here, ReVisiT achieves an 8.33% reduction in the CHAIR score compared to vanilla greedy decoding on Qwen2.5-VL, while all other baselines

exhibit higher CHAIR scores than the vanilla greedy. Given that CHAIR directly measures hallucination frequency and serves as the main evaluation criterion in the generative setting, we believe that it is the most important measurement for visual grounding quality. Across other metrics, ReVisiT demonstrates competitive or superior performance in general, which highlight the robustness and stability of our approach across diverse evaluation conditions. Overall, these results demonstrate that adaptively leveraging the most informative vision token at each decoding step, as proposed in ReVisiT, enables more faithful and visually grounded generation while maintaining simplicity.

4.3. Additional analyses

In this section, we conduct additional analyses to further investigate the properties of ReVisiT.

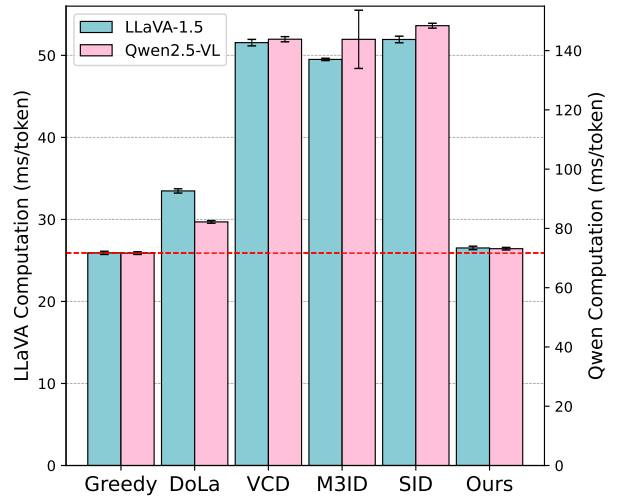


Figure 3. **Inference speed.** Comparison of per-token inference latency across different decoding strategies for LLaVA-1.5-7B (left y-axis) and Qwen2.5-VL-7B (right y-axis), with standard deviations visualized as error bars.

Inference speed improvement. To evaluate the inference efficiency of ReVisiT compared to baseline decoding strategies, we measure the per-token computational time. All

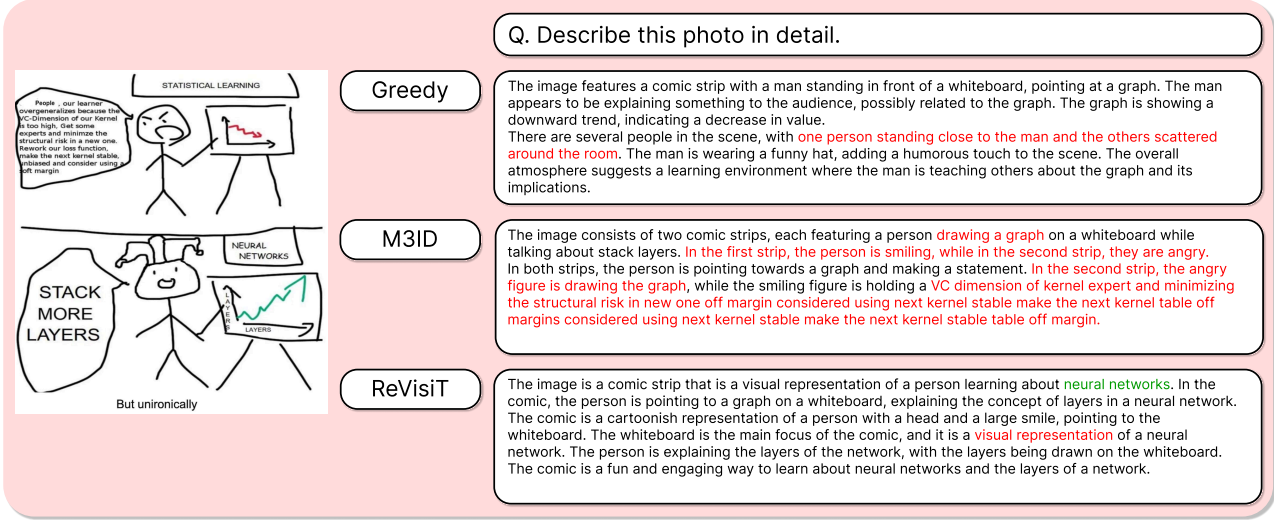


Figure 4. **Qualitative example.** The input image is a cartoon-style illustration contrasting classical statistical learning and neural network reasoning via a visual metaphor, emphasizing the shift from theoretical rigor to the heuristic of “stacking more layers.” We compare the generated responses of vanilla greedy decoding, M3ID, and ReVisiT, highlighting how ReVisiT better captures the intended visual analogy compared to the baseline methods.

measurements are conducted on an NVIDIA RTX A6000 GPU with 48GB memory and an AMD EPYC 9354 32-Core Processor running at 3.25 GHz. To eliminate initial overhead, we perform 6 warmup iterations before measurement. During evaluation, each of the 300 samples is decoded three times, and the results are averaged to reduce variance. We report the average decoding time per generated token, measured in milliseconds (ms/token). Other evaluation configurations, including preprocessing pipelines and evaluation scripts, follow the official CHAIR benchmark settings. Exact numerical values are provided in the Appendix B.2.

As shown in Figure 3, ReVisiT achieves inference efficiency comparable to vanilla greedy decoding, exhibiting only 2.3% and 3.5% increases in computational time on LLaVA-1.5 and Qwen2.5-VL, respectively. Unlike ReVisiT, the intra-alignment baselines generally incur nearly $2\times$ higher computational cost. In addition, DoLa, which requires an additional forward pass at every decoding step to project intermediate layers, also introduces 29.34% and 14.64% additional computation on LLaVA-1.5 and Qwen2.5-VL, respectively. These results highlight the key advantage of ReVisiT: it enhances visual grounding with negligible computational overhead, unlike prior methods that significantly compromise inference efficiency.

Ablation study. To investigate the impact of each design choice, we conduct an ablation study on Qwen2.5-VL, evaluated under the CHAIR benchmark. Table 4 summarizes the results, varying three factors: (1) w/ or w/o vocabulary subset constraint, (2) layer selection strategy (all layers

Table 4. **Ablation study.** We evaluate the impact of vocabulary subset constraint, layer selection, and threshold α variation on hallucination mitigation under the CHAIR benchmark. Main table settings are in gray.

SUBSET LAYER α			CHAIR _S ↓	CHAIR _I ↓	RECALL ↑
w/o	ALL	1e−5	0.2	7.14	0.83
		1e−4	31.8	7.93	72.17
		1e−5	31.6	7.56	70.83
w/	ALL	1e−6	32.4	9.38	70.44
		1e−4	34.6	7.89	71.27
		1e−5	33.8	7.92	70.31
	LAST	1e−6	33.2	9.07	70.83

vs. last layer), and (3) threshold α for vocabulary selection. First, it is observed that selecting vision tokens from all layers achieves better performance than relying solely on the last layer on Qwen2.5-VL, suggesting that aggregating multi-layer information can be beneficial for visual grounding. We also observe that increasing α results in a smaller and sparser vocabulary subset, which improves the Recall score by promoting the generation of more plausible ground-truth objects, as observed with $\alpha = 1e-4$. However, as enlarging subset can include noisy tokens and result in the reduced precision, which indicates the importance of proper selection. As shown in Table 4, when vision token projection is performed over the full vocabulary without constraining the subset (w/o subset), the CHAIR_S score is abnormally low (0.2) and the CHAIR_I score also appears improved compared to the best setting. However, the Recall

score drops sharply to 0.83, indicating that the model fails to generate not only hallucinated objects but also a significant portion of ground-truth objects. This highlights the necessity of constraining the candidate set to context-relevant tokens. Examples of the generation failures are provided in Appendix C.2.

Qualitative analysis. To illustrate the qualitative improvements, we present examples from the LLaVA-Bench-In-The-Wild dataset, which consists of diverse real-world images requiring complex reasoning beyond conventional benchmarks. As shown in Figure 4, the input image is a cartoon illustration requiring contextual and conceptual understanding. While baseline methods such as vanilla greedy decoding fail to capture the main theme and generate overly generic or literal descriptions, or even getting stuck in repetitive phrasing. In contrast, ReVisiT successfully identifies the underlying concept: about neural network layers. More qualitative samples are provided in Appendix C.3

5. Conclusion

In this paper, we proposed ReVisiT, a simple yet effective LVLM decoding method that references vision tokens to guide the text generation process. Unlike prior approaches that introduce extra computational cost to obtain the reference logit at every decoding step, ReVisiT leverages the intrinsic vision token representations with negligible overhead. By dynamically constraining the output vocabulary based on output context, ReVisiT enhances visual grounding and mitigates object hallucinations across both generative and discriminative tasks.

References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C. L., and Parikh, D. Vqa: Visual question answering. In *European Conference on Computer Vision (ECCV)*, 2015.
- Bai, J., Bai, S., Yang, S., Wang, S., Tan, S., Wang, P., Lin, J., Zhou, C., and Zhou, J. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023.
- Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- Chen, K., Zhang, Z., Zeng, W., Zhang, R., Zhu, F., and Zhao, R. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023.
- Chen, Z., Wu, J., Wang, W., Su, W., Chen, G., Xing, S., Zhong, M., Zhang, Q., Zhu, X., Lu, L., et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024a.
- Chen, Z., Zhao, Z., Luo, H., Yao, H., Li, B., and Zhou, J. Halc: Object hallucination reduction via adaptive focal-contrast decoding. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2024b.
- Chuang, Y.-S., Xie, Y., Luo, H., Kim, Y., Glass, J., and He, P. Dola: Decoding by contrasting layers improves factuality in large language models. In *International Conference on Learning Representations (ICLR)*, 2024.
- Dai, W., Li, J., Li, D., Tiong, A. M. H., Zhao, J., Wang, W., Li, B., Fung, P., and Hoi, S. Instructblip: Towards general-purpose vision-language models with instruction tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- Dai, W., Lee, N., Wang, B., Yang, Z., Liu, Z., Barker, J., Rintamaki, T., Shoeybi, M., Catanzaro, B., and Ping, W. Nvlm: Open frontier-class multimodal llms. *arXiv preprint arXiv:2409.11402*, 2024.
- Elbayad, M., Gu, J., Grave, E., and Auli, M. Depth-adaptive transformer. In *International Conference on Learning Representations (ICLR)*, 2020.
- Favero, A., Zancato, L., Trager, M., Choudhary, S., Perera, P., Achille, A., Swaminathan, A., and Soatto, S. Multi-modal hallucination control by visual information grounding. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- Gong, X., Ming, T., Wang, X., and Wei, Z. Damro: Dive into the attention mechanism of lvlm to reduce object hallucination. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
- Huang, Q., Dong, X., Zhang, P., Wang, B., He, C., Wang, J., Lin, D., Zhang, W., and Yu, N. Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- Hudson, D. A. and Manning, C. D. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

- Huo, F., Xu, W., Zhang, Z., Wang, H., Chen, Z., and Zhao, P. Self-introspective decoding: Alleviating hallucinations for large vision-language models. In *International Conference on Learning Representations (ICLR)*, 2025.
- Kang, S., Kim, J., Kim, J., and Hwang, S. J. See what you are told: Visual attention sink in large multimodal models. In *International Conference on Learning Representations (ICLR)*, 2025.
- Leng, S., Zhang, H., Chen, G., Li, X., Lu, S., Miao, C., and Bing, L. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- Li, J., Li, D., Savarese, S., and Hoi, S. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2023a.
- Li, X. L., Holtzman, A., Fried, D., Liang, P., Eisner, J., Hashimoto, T., Zettlemoyer, L., and Lewis, M. Contrastive decoding: Open-ended text generation as optimization. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023b.
- Li, Y., Du, Y., Zhou, K., Wang, J., Zhao, W. X., and Wen, J.-R. Evaluating object hallucination in large vision-language models. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2023c.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *International Conference on Computer Vision (ICCV)*, 2014.
- Liu, H., Li, C., Wu, Q., and Lee, Y. J. Visual instruction tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- Liu, H., Li, C., Li, Y., and Lee, Y. J. Improved baselines with visual instruction tuning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024a.
- Liu, S., Zheng, K., and Chen, W. Paying more attention to image: A training-free method for alleviating hallucination in lvlms. In *International Conference on Computer Vision (ICCV)*, 2024b.
- Lu, S., Li, Y., Chen, Q.-G., Xu, Z., Luo, W., Zhang, K., and Ye, H.-J. Ovis: Structural embedding alignment for multimodal large language model. *arXiv preprint arXiv:2405.20797*, 2024.
- nostalgebraist. interpreting gpt: the logit lens. <https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens>, 2020.
- Rohrbach, A., Hendricks, L. A., Burns, K., Darrell, T., and Saenko, K. Object hallucination in image captioning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.
- Schuster, T., Fisch, A., Gupta, J., Dehghani, M., Bahri, D., Tran, V., Tay, Y., and Metzler, D. Confident adaptive language modeling. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Schwenk, D., Khandelwal, A., Clark, C., Marino, K., and Mottaghi, R. A-okvqa: A benchmark for visual question answering using world knowledge. In *International Conference on Computer Vision (ICCV)*, 2022.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *European Conference on Computer Vision (ECCV)*, 2017.
- Sundar, A. and Heck, L. Multimodal conversational ai: A survey of datasets and approaches. *arXiv preprint arXiv:2205.06907*, 2022.
- Team, G., Anil, R., Borgeaud, S., Alayrac, J.-B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A. M., Hauth, A., Millican, K., et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Teerapittayanon, S., McDanel, B., and Kung, H.-T. Branchynet: Fast inference via early exiting from deep neural networks. In *International Conference on Pattern Recognition (ICPR)*, 2016.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023a.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023b.
- Vinyals, O., Toshev, A., Bengio, S., and Erhan, D. Show and tell: A neural image caption generator. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- Wang, C., Chen, X., Zhang, N., Tian, B., Xu, H., Deng, S., and Chen, H. Mllm can see? dynamic correction decoding for hallucination mitigation. In *International Conference on Learning Representations (ICLR)*, 2025.

- Wang, J., Wang, Y., Xu, G., Zhang, J., Gu, Y., Jia, H., Wang, J., Xu, H., Yan, M., Zhang, J., et al. Amber: An llm-free multi-dimensional benchmark for mllms hallucination evaluation. *arXiv preprint arXiv:2311.07397*, 2023.
- Wang, X., Pan, J., Ding, L., and Biemann, C. Mitigating hallucinations in large vision-language models with instruction contrastive decoding. *arXiv preprint arXiv:2403.18715*, 2024.
- Woo, S., Kim, D., Jang, J., Choi, Y., and Kim, C. Don’t miss the forest for the trees: Attentional vision calibration for large vision language models. *arXiv preprint arXiv:2405.17820*, 2024.
- Wu, Z., Chen, X., Pan, Z., Liu, X., Liu, W., Dai, D., Gao, H., Ma, Y., Wu, C., Wang, B., et al. Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding. *arXiv preprint arXiv:2412.10302*, 2024.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., and Bengio, Y. Show, attend and tell: Neural image caption generation with visual attention. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- Ye, Q., Xu, H., Xu, G., Ye, J., Yan, M., Zhou, Y., Wang, J., Hu, A., Shi, P., Shi, Y., et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023.
- Yin, S., Fu, C., Zhao, S., Xu, T., Wang, H., Sui, D., Shen, Y., Li, K., Sun, X., and Chen, E. Woodpecker: Hallucination correction for multimodal large language models. In *Science China Information Sciences (SCIS)*, 2024.
- You, H., Zhang, H., Gan, Z., Du, X., Zhang, B., Wang, Z., Cao, L., Chang, S.-F., and Yang, Y. Ferret: Refer and ground anything anywhere at any granularity. In *International Conference on Learning Representations (ICLR)*, 2024.
- You, Q., Jin, H., Wang, Z., Fang, C., and Luo, J. Image captioning with semantic attention. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Zhou, Y., Cui, C., Yoon, J., Zhang, L., Deng, Z., Finn, C., Bansal, M., and Yao, H. Analyzing and mitigating object hallucination in large vision-language models. *arXiv preprint arXiv:2310.00754*, 2023.
- Zhu, D., Chen, J., Shen, X., Li, X., and Elhoseiny, M. Minigt-4: Enhancing vision-language understanding with advanced large language models. In *International Conference on Learning Representations (ICLR)*, 2024.

A. Detailed Discussions

A.1. Prior works

In this section, we provide a detailed discussion of prior works that aim to align visual inputs with language generation. We organize these approaches into three categories—*pre-alignment*, *intra-alignment*, and *post-alignment*—based on the stage where the alignment is applied with respect to the inference process. While this taxonomy follows the temporal flow of inference, the categories are largely orthogonal in methodology and can be composed to enhance multimodal grounding.

A.1.1. PRE-ALIGNMENT

Recent advances in LVLMs have explored a range of *pre-alignment* strategies, in which visual and textual modalities are aligned during training to enhance multimodal understanding. Most existing approaches adopt a combination of instruction tuning and architectural refinements to achieve improved visual grounding. For instance, Qwen-VL (Bai et al., 2023) and InternVL (Chen et al., 2024a) retain decoder-only LLM backbones while incorporating high-resolution vision encoders and cross-modal projection modules. These models leverage large-scale curated image-text data and modality-specific tuning objectives, leading to strong general-purpose performance across a variety of vision-language tasks. Building upon these designs, Qwen2.5-VL (Bai et al., 2025) introduces a Dynamic High-Resolution (DHR) processing pipeline and rotary position encodings tailored for vision-language alignment, further enhancing the model’s ability to process diverse image resolutions efficiently. NVLM (Dai et al., 2024) explores a broader architectural design space by implementing decoder-only, cross-attention, and hybrid variants, along with a DHR input pipeline and tile tagging mechanisms. On the other hand, there are another line of works designed to mitigate the architectural challenges within the pre-alignment space. Ferret (You et al., 2024) focuses on fine-grained spatial referring and grounding by introducing hybrid region representations and a spatial-aware visual sampler, enabling the model to handle points, boxes, and free-form regions. In a different direction, Ovis (Lu et al., 2024) addresses the structural mismatch between continuous visual embeddings and discrete textual embeddings. Specifically, Ovis introduces a learnable visual embedding table and a probabilistic vision tokenization scheme that mimics the structure of language token embeddings, thereby aligning the two modalities at the representation level. While these pre-alignment strategies have proven effective, they typically require extensive retraining and are not designed to enable inference-time intervention based on visual context.

A.1.2. POST-ALIGNMENT

Post-alignment refers to strategies that operate on the generated output sequence, aiming to correct hallucinations or factual inconsistencies after the generation process is complete. Unlike pre- or intra-alignment methods that intervene during training or decoding, post-alignment introduces no changes to the model’s parameters or inference-time logic. A key distinction among post-alignment methods lies in how the correction is applied. Some approaches train an auxiliary revision module, while others rely on external vision models to validate and refine the generated content. For example, LURE (Zhou et al., 2023) learns a dedicated revisor that takes the original image and the LVLM-generated caption as input, and outputs a hallucination-corrected caption. The training signal is derived from statistical indicators such as co-occurrence bias and positional uncertainty observed in model outputs. In contrast, Woodpecker (Yin et al., 2024) does not require finetuning, but instead performs post-hoc correction through a multi-stage pipeline. Starting from the generated text, it extracts visual concepts, formulates corresponding questions, and verifies them using a pretrained VQA model. Each stage of this process relies on existing pretrained components including question generation, visual verification, and text editing modules to systematically identify and correct hallucinated content. While post-alignment strategies offer a practical way to improve output consistency without altering the base model, they typically rely on additional training or auxiliary models.

A.1.3. INTRA-ALIGNMENT

Improving visual grounding (e.g., mitigating object hallucination) of LVLMs during decoding has evolved through distinct methodological paradigms, each addressing unique aspects of visual-textual misalignment. The research trajectory begins with Contrastive Decoding (Li et al., 2023b) strategies, pioneered by VCD (Leng et al., 2024), which introduced input distortions to contrast original and perturbed visual distributions, thereby reducing reliance on language priors within LVLMs. This foundation expanded into instruction-aware variants such as ICD (Wang et al., 2024), which amplifies alignment uncertainty through textual prompt disturbances, and M3ID (Favero et al., 2024), which identifies language priors in the LLM decoder via vision-agnostic input perturbations. Concurrently, attention-centric approaches emerged (Gong et al., 2024; Liu et al., 2024b; Woo et al., 2024; Huo et al., 2025; Kang et al., 2025), aiming to mitigate attention misallocation by

recalibrating cross-modal attention or introducing contrastive signals for token-level grounding. Other paradigms (Chen et al., 2024b; Huang et al., 2024; Wang et al., 2025) have also been explored in recent literature, but all these intra-alignment approaches require additional computation to obtain vision-related reference information.

A.2. Dataset and evaluation setup

CHAIR evaluation. CHAIR (Rohrbach et al., 2018) is a generative benchmark designed to measure object hallucinations in image captioning. A hallucinated object refers to any entity mentioned in the generated caption that is not present in the corresponding image. This metric is widely adopted to evaluate the visual grounding capability of vision-language models. We follow the standard CHAIR evaluation procedure based on the object annotations from the MS COCO dataset (Lin et al., 2014). Each image is associated with a set of ground-truth objects defined by the 80 category COCO detection vocabulary. During evaluation, noun phrases extracted from generated captions are matched against this predefined object list. Mentions of out-of-vocabulary objects are ignored, while in-vocabulary objects absent from the ground-truth annotations are considered hallucinated. We report two standard metrics provided by the CHAIR benchmark and one additional metric:

$$\begin{aligned}\text{CHAIR}_S &= \frac{|\{\text{hallucinated captions}\}|}{|\{\text{total captions}\}|}, \\ \text{CHAIR}_I &= \frac{|\{\text{hallucinated object mentions}\}|}{|\{\text{total object mentions}\}|}, \\ \text{Recall} &= \frac{|\{\text{correctly predicted objects}\}|}{|\{\text{total ground-truth objects}\}|}\end{aligned}$$

We use the official CHAIR evaluation code¹ and randomly sample 500 images from the MS COCO validation set using a fixed seed (seed=42) for reproducibility. For each image, we prompt the model with "Please describe this image in detail." to generate captions.

POPE evaluation. POPE (Li et al., 2023c) is a discriminative benchmark designed to evaluate the object-level visual grounding capabilities of LVLMs. It formulates the task as a binary object presence question-answering problem, where the model is asked to determine whether a specific object is present in a given image. Evaluation is conducted using standard classification metrics derived from the confusion matrix, including accuracy, precision, and F1 score. Accuracy measures the overall proportion of correct predictions, while precision captures the fraction of predicted positive instances that are indeed correct. F1 score, defined as the harmonic mean of precision and recall, provides a balanced assessment of model performance on this binary classification task. The benchmark covers three datasets: MS COCO (Lin et al., 2014), A-OKVQA (Schwenk et al., 2022), and GQA (Hudson & Manning, 2019). For each dataset, POPE defines three types of query strategies—*random*, *popular*, and *adversarial*—yielding a total of nine evaluation scenarios. Binary questions are constructed in the form "Is there a <object> in the image?", with the queried object selected according to the designated strategy. Random queries are uniformly sampled from the dataset’s object vocabulary. Popular queries focus on frequently occurring objects, while adversarial queries target objects that are semantically plausible in the given context but are absent from the image, thus probing the model’s reliance on prior co-occurrence statistics rather than visual evidence. Ground-truth labels indicating object presence are provided by the official POPE benchmark². For MS COCO, labels are derived from object detection annotations, whereas for A-OKVQA and GQA, they are obtained via SEEM-based segmentation annotations. Each query type consists of 3,000 binary QA instances, evenly split between positive and negative samples. As a result, we evaluate across a total of nine scenarios comprising 27,000 QA examples.

AMBER evaluation. AMBER (Wang et al., 2023) is a comprehensive benchmark designed to assess hallucination phenomena in LVLMs across multiple dimensions. Unlike CHAIR and POPE that focus on object existence, AMBER evaluates three core hallucination types: existence (fabricating non-existent entities), attribute (misrepresenting object properties), and relation (incorrectly describing inter-object interactions). The benchmark supports both generative tasks and discriminative tasks. We focused on generative setting, where AMBER calculates hallucination rates through automated object matching against ground-truth annotations, similar to CHAIR but extended to attribute and relation verification. We

¹<https://github.com/LisaAnne/Hallucination>

²<https://github.com/AoiDragon/POPE>

report four metrics provided in the AMBER benchmark to quantify hallucination phenomena in the generative setting:

- **CHAIR** measures the proportion of hallucinated object mentions in the response. Given a generated object set R'_{obj} and the ground-truth annotated object set A_{obj} , it is defined as:

$$\text{CHAIR}(R) = 1 - \frac{|R'_{\text{obj}} \cap A_{\text{obj}}|}{|R'_{\text{obj}}|}.$$

- **Cover** evaluates the object coverage of the model output, measuring how well the response mentions the ground-truth objects:

$$\text{Cover}(R) = \frac{|R'_{\text{obj}} \cap A_{\text{obj}}|}{|A_{\text{obj}}|}.$$

- **Hal** is a binary indicator of hallucination presence. A response is considered hallucinatory if any hallucinated object is detected:

$$\text{Hal}(R) = \begin{cases} 1 & \text{if } \text{CHAIR}(R) \neq 0 \\ 0 & \text{otherwise} \end{cases}.$$

- **Cog** captures the proportion of hallucinated objects that belong to a predefined cognitively plausible hallucination set H_{obj} :

$$\text{Cog}(R) = \frac{|R'_{\text{obj}} \cap H_{\text{obj}}|}{|R'_{\text{obj}}|}.$$

All metrics are averaged over the full validation split comprising 1,004 generative samples. Lower values of CHAIR and Hal indicate fewer hallucinations, while higher Cover and Cog scores reflect better object recall and alignment with cognitively plausible patterns. We use the official AMBER benchmark³ without modification. For Qwen2.5-VL experiments, input images exceeding a resolution of 1400×1400 were resized to fit GPU memory constraints, due to its dynamic high-resolution encoding.

LLaVA-Bench-In-the-Wild dataset. We additionally employ the LLaVA-Bench-In-the-Wild (Liu et al., 2023) dataset to qualitatively evaluate the visual grounding ability of LVLMs in open-ended image understanding. The dataset consists of 24 diverse images, including both real-world photographs and abstract illustrations, and 60 textual prompts designed to elicit complex visual reasoning and language generation. Following prior work (Leng et al., 2024), we present representative qualitative examples to illustrate how different decoding strategies affect grounding fidelity and response quality in challenging scenarios.

³<https://github.com/junyangwang0410/AMBER>

Algorithm 1 ReVisiT decoding algorithm

Input: LVLMM \mathcal{M} with final decoder layer L , projection (language modeling) head W_{proj} , and vocabulary \mathcal{V} , input image I , set of candidate decoder layers \mathcal{J} , input text prompt x_{raw}

Tokenize input text: $\mathbf{x} = \{x_0, \dots, x_{|\mathbf{x}|-1}\} \leftarrow \text{Tokenizer}(x_{\text{raw}})$

Encode image: $\mathbf{v} = \{v_0, \dots, v_{|\mathbf{v}|-1}\} \leftarrow \text{VisionEncoder}(I)$

Construct input sequence: $\mathbf{z} = \mathbf{v} \parallel \mathbf{x}$

$\triangleright \mathbf{z} = \{z_0, \dots, z_{T-1}\}$ where $T = |\mathbf{v}| + |\mathbf{x}|$

Initialize output sequence $\mathbf{y} \leftarrow \emptyset$

$\triangleright t \leftarrow 0$

while not end-of-sequence **do**

 Get decoder hidden states h_{T+t-1}^L from \mathcal{M} with context (\mathbf{z}, \mathbf{y})

if first-timestep **then**

 Given $\{h_i^j\}$ for $i = 0, \dots, T-1$ and $j \in \mathcal{J}$ from initial forward pass

 Initialize cache $\mathbf{l}_{\text{vision}} \leftarrow \emptyset$

for each layer $j \in \mathcal{J}$ **do**

for $i = 0$ to $|\mathbf{v}| - 1$ **do**

$l_i^j = W_{\text{proj}}^\top h_i^j$

 Append l_i^j to $\mathbf{l}_{\text{vision}}$

end for

end for

 Cache matrix $\mathbf{l}_{\text{vision}} \in \mathbb{R}^{|\mathcal{J}| \cdot |\mathbf{v}| \times |\mathcal{V}|}$

 first-timestep $\leftarrow \text{False}$

end if

 Compute vanilla distribution $p_{\text{base}} = \text{softmax}_{\mathcal{V}}(W_{\text{proj}}^\top h_{T+t-1}^L)$ (Eq. 1)

 Select constrained vocabulary $\mathcal{V}_{\text{cons}}^t \subseteq \mathcal{V}$ based on APC of p_{base} (Eq. 3.1)

 Let $p_{\text{base}}^{\text{cons}} = \text{softmax}_{\mathcal{V}_{\text{cons}}^t}(W_{\text{proj}}^\top h_{T+t-1}^L)$

$(j^*, i^*) \leftarrow \arg \min_{(j,i)} \text{JSD}(p_{\text{base}}^{\text{cons}} \parallel \text{softmax}_{\mathcal{V}_{\text{cons}}^t}(\mathbf{l}_{\text{vision}}[j, i]))$

 Let $p_{\text{ref}}^{\text{cons}} = \text{softmax}_{\mathcal{V}_{\text{cons}}^t}(\mathbf{l}_{\text{vision}}[j^*, i^*])$

 Logit adjustment via element-wise multiplication: $l_{\text{ReVisiT}} = p_{\text{base}}^{\text{cons}} \odot p_{\text{ref}}^{\text{cons}}$

 Normalize over constrained vocabulary: $p_{\text{ReVisiT}} = \text{softmax}_{\mathcal{V}_{\text{cons}}^t}(l_{\text{ReVisiT}})$

 Sample next token: $y_t \sim p_{\text{ReVisiT}}$

 Append y_t to \mathbf{y}

$t \leftarrow t + 1$

end while

return \mathbf{y}

A.3. ReVisiT

In this section, we first provide the pseudo-code of ReVisiT in Algorithm 1, which illustrates how vision tokens are leveraged to refine the output distribution at each decoding step. We highlight that our implementation is memory-efficient, model-agnostic, and requires no additional training or modification to the base model; the decoding function alone is modified. ReVisiT involves two hyperparameters: (1) the pool of decoder layer(s) where vision tokens are selected, and (2) the alignment weight α that determines the size of constrained vocabulary subset. For LLaVA-1.5-7B, we empirically found that using vision tokens from the final decoder layer (“last”) mostly yielded the best grounding fidelity. In contrast, for Qwen2.5-VL-7B, aggregating vision tokens from all even-numbered decoder layers (“all”) was more effective. These selections were fixed per model across all benchmarks. The threshold α controls the size of the constrained vocabulary subset. Lower values (resulting in larger subsets) are used for generative tasks where broader linguistic diversity is desired, while higher values (yielding smaller subsets) are preferred in discriminative tasks to enhance precision. We select α from a logarithmic range between 10^{-1} and 10^{-6} , with 10^{-5} serving as the default when task-specific adaptation is not required.

B. Detailed Quantitative Results

B.1. Comprehensive POPE results

Table 5. **Comprehensive Results on POPE benchmark.** Higher scores (\uparrow) on Accuracy, Precision, and F1 indicate better performance.

	SETUP	METHOD	LLAVA-1.5			QWEN2.5-VL		
			ACC. \uparrow	PREC. \uparrow	F1 \uparrow	ACC. \uparrow	PREC. \uparrow	F1 \uparrow
MS-COCO	RANDOM	GREEDY	89.07	89.54	89.00	84.33	99.71	81.47
		DoLA	89.00	89.58	88.92	81.40	99.89	77.17
		VCD	87.23	86.62	87.34	84.50	98.96	81.81
		M3ID	89.07	89.54	89.00	84.27	99.42	81.42
		SID	89.73	91.22	89.55	86.00	99.09	83.85
		OURS	89.10	89.82	89.00	84.67	99.62	81.95
	POPULAR	GREEDY	85.63	83.72	86.03	83.77	98.10	80.92
		DoLA	85.63	83.85	86.00	81.10	98.95	76.89
		VCD	83.53	80.75	84.25	83.93	97.39	81.27
		M3ID	85.63	83.72	86.03	83.77	98.01	80.94
		SID	86.03	84.71	86.29	85.33	97.32	83.21
		OURS	85.80	84.16	86.13	84.17	98.21	81.47
	ADVERSARIAL	GREEDY	79.27	74.72	81.01	83.43	97.18	80.61
		DoLA	79.40	74.97	81.08	80.83	98.03	76.65
		VCD	77.67	72.90	79.77	83.50	96.31	80.85
		M3ID	79.27	74.72	81.01	83.40	97.00	80.59
		SID	80.20	76.22	81.60	84.33	94.86	82.25
		OURS	79.90	76.21	81.22	83.70	97.11	80.99
A-OKVQA	RANDOM	GREEDY	86.33	81.00	87.42	87.87	96.86	86.58
		DoLA	86.43	81.25	87.47	84.97	97.73	82.65
		VCD	83.70	77.91	85.23	88.43	96.91	87.28
		M3ID	86.33	81.00	87.42	87.30	96.05	85.97
		SID	87.07	82.55	87.91	87.87	96.18	86.67
		OURS	87.37	84.75	87.83	88.20	97.36	86.94
	POPULAR	GREEDY	78.77	71.74	81.72	86.07	92.73	84.89
		DoLA	79.00	72.06	81.85	83.07	92.91	80.87
		VCD	76.47	69.58	79.99	86.20	92.36	85.12
		M3ID	78.77	71.74	81.72	85.80	92.22	84.63
		SID	80.33	73.82	82.70	86.33	92.71	85.23
		OURS	81.97	77.54	83.31	86.17	92.42	85.07
	ADVERSARIAL	GREEDY	68.20	61.86	74.91	81.27	83.26	80.69
		DoLA	68.33	62.00	74.95	79.37	84.77	77.63
		VCD	67.80	61.69	74.47	81.53	83.12	81.08
		M3ID	68.20	61.86	74.91	80.23	81.51	79.82
		SID	70.90	64.30	76.36	81.57	83.37	81.06
		OURS	72.83	67.00	76.81	81.40	83.31	80.85
GQA	RANDOM	GREEDY	85.97	79.72	87.30	87.50	97.15	86.08
		DoLA	86.07	79.92	87.36	83.47	97.36	80.62
		VCD	83.77	77.16	85.53	87.87	96.86	86.58
		M3ID	86.00	79.74	87.33	87.27	96.82	85.82
		SID	85.97	79.72	87.30	86.80	96.94	85.20
		OURS	87.13	83.47	87.80	87.83	97.17	86.50
	POPULAR	GREEDY	73.80	66.38	78.64	84.40	90.12	83.20
		DoLA	73.97	66.56	78.73	79.60	87.76	77.13
		VCD	71.87	64.76	77.32	84.33	89.07	83.32
		M3ID	73.83	66.39	78.67	82.20	85.88	81.24
		SID	73.80	66.38	78.64	83.47	89.34	82.13
		OURS	78.67	72.42	81.28	84.63	90.37	83.46
	ADVERSARIAL	GREEDY	68.17	61.60	75.19	81.80	84.97	80.94
		DoLA	68.40	61.80	75.30	78.80	86.00	76.44
		VCD	67.63	61.22	74.82	81.93	84.66	81.19
		M3ID	68.20	61.62	75.22	81.17	83.95	80.36
		SID	68.17	61.60	75.19	81.17	84.76	80.14
		OURS	73.40	66.91	77.68	82.17	85.14	81.38

B.2. Token-Level computation speed

Table 6. **Per-token inference latency** (ms/token) of different decoding strategies. We report the mean \pm standard deviation over 300 samples.

METHOD	LLAVA-1.5-7B	QWEN2.5-VL-7B
GREEDY	25.9 \pm 0.2	71.7 \pm 0.5
DoLA	33.5 \pm 0.3	82.2 \pm 0.4
VCD	51.6 \pm 0.4	143.8 \pm 0.9
M3ID	49.5 \pm 0.1	143.8 \pm 9.8
SID	51.9 \pm 0.4	148.4 \pm 0.8
OURS	<u>26.5 \pm 0.2</u>	<u>73.2 \pm 0.4</u>

B.3. Additional evaluation on DeepSeekVL2

To further demonstrate the model-agnostic applicability of our decoding strategy, we evaluate ReVisiT on the recent **DeepSeekVL2** model (Wu et al., 2024), a 27B-parameter LVLM with a Mixture-of-Experts (MoE) decoder backbone based on DeepSeek-MoE. Despite its large total parameter count, only 4.1B parameters are activated per inference due to expert sparsity. This architecture poses a distinct challenge compared to dense decoder models such as LLaVA or Qwen.

During implementation, we observed frequent degeneration in the form of repetition even under vanilla greedy decoding. To ensure fair comparison, we apply a repetition penalty of 1.1 across all methods for this experiment, following practices from prior work on decoding calibration.

As shown in Table 7, ReVisiT consistently reduces hallucinations in the CHAIR benchmark: CHAIR_S drops from 20.4 to 14.2, and CHAIR_I from 7.04 to 2.1. These results reaffirm the effectiveness of ReVisiT in mitigating hallucinated object mentions across diverse LVLM architectures, including sparsely activated expert-based models.

Table 7. **DeepSeekVL2 results on CHAIR benchmark.** ReVisiT effectively reduces hallucinations while maintaining comparable recall.

METHOD	CHAIR _S ↓	CHAIR _I ↓	RECALL ↑
GREEDY	20.40	7.04	59.51
REVISIT	14.20	2.10	57.49

C. Detailed Qualitative Results

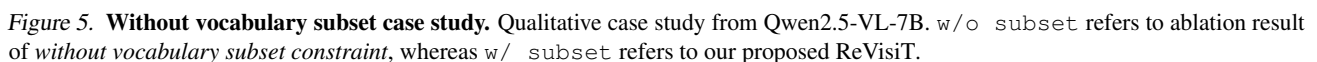
C.1. Case studies for vision token analysis

To support the qualitative visualization presented in Figures 1 and 2, we provide the detailed numerical values and positional metadata used in our vision token analysis. All examples are drawn from the LLaVA-Bench-In-the-Wild dataset, using the LLaVA-1.5-7B. Vision token embeddings are extracted from decoder layer $L = 32$, and the threshold for constrained vocabulary selection is set to $\alpha = 10^{-5}$.


First, Figure 1 presents a case study from the LLaVA-Bench-In-the-Wild dataset (image ID: 002 . jpg, query: “*How many uncut fruits are in the image?*”). At decoding step $t = 2$, ReVisiT selects vision token index 229 as the most relevant reference based on the constrained vocabulary distribution. Table 8 reports the log-probability values from three sources: (1) the base model output (Eq. 1), (2) the vision token projection at index 229 (Eq. 5), and (3) the final combined distribution after log-probability fusion (Eq. 3.3). These values are computed over the constrained subset $\mathcal{V}_{\text{cons}}^2$. The token ‘three’ is ultimately selected, as the vision token significantly boosts its relevance, increasing its probability from 35.45% to 58.27%, while suppressing ‘four’ from 50.78% to 38.21%.

Next, Figure 2 visualizes a comprehensive case from the same dataset (image ID: 003 . jpg, prompt: “*Describe this photo in detail.*”). Four vision tokens are selected to highlight different aspects of visual grounding. The three tokens shown on the **left** of the figure are used to examine interpretability and semantic alignment of vision token projections. For each token, we compute the projected distribution over (1) the full vocabulary \mathcal{V} and (2) a manually defined constrained subset $\mathcal{V}_{\text{env}} = \{\text{river, mountain, person, sky, tree, cloud, sea, grass, rock}\}$. This comparison reveals how visually grounded semantics become more prominent when irrelevant lexical candidates are excluded. The full distribution shifts are detailed in Table 9.

The token on the **right** (vision token index 193) is selected to demonstrate the actual decoding behavior of ReVisiT. At generation step $t = 72$, this token serves as a reference informer to refine the output distribution through log-probability fusion. The final token ‘painting’ is selected, with its probability increasing from 8.08% (base) to 39.15% (final), while alternatives such as hallucinatory token ‘person’ are suppressed. Detailed values are presented in Table 10.



C.3. Additional qualitative examples



Q. What are the animals in the painting and what are they doing?

Greedy

The painting features a bear, a cat, and a rabbit sitting around a table, eating food. They are enjoying a meal together, with the bear eating a donut, the cat drinking from a cup, and the rabbit eating a cookie.

DoLa

The painting features a bear, a cat, and a rabbit sitting around a table, eating food. They are enjoying a meal together, with the bear eating a donut, the cat drinking from a cup, and the rabbit eating a cookie.

VCD

The painting features a bear, a bunny, and a cat sitting around a table, eating cookies. They are all gathered around the table, enjoying their treats together.

M3ID

The painting features a bear, a cat, and a rabbit sitting around a wooden table, enjoying a meal together. They are eating cookies and drinking pink cups.

SID

The painting features a bear, a cat, and a rabbit sitting around a table, eating cookies and drinking from cups.

ReVisiT

The painting features a bear, a cat, and a rabbit. The bear is eating a donut, while the cat and the rabbit are sitting at a table with a plate of donuts.

Figure 6. Additional qualitative example. The input image is a illustration showing a bear, a cat, and a rabbit seated around a table with a plate of donuts. We compare the responses of vanilla greedy decoding and ReVisiT to the question, “*What are the animals in the painting and what are they doing?*” While the greedy output introduces a hallucinated detail (“*cookie*”) and assigns actions not visually supported (e.g., “*the cat drinking from a cup*”), ReVisiT provides a more faithful description aligned with the visual content, accurately identifying the animals and their activities.

20

D. Broader Impacts and Limitations

Broader impacts. Like most approaches built upon LLMs, the proposed method inherits the limitations and potential societal risks of the underlying models, including biases in training data, reinforcement of stereotypes, and the generation of factually incorrect or harmful content. While our decoding strategy improves visual grounding by constraining generation based on vision token signals, it does not explicitly eliminate such biases and may inadvertently propagate them through vision-conditioned outputs. On the other hand, by incorporating grounded visual evidence during generation, our method may contribute to mitigating hallucinations and increasing the interpretability of multimodal outputs. We believe that this is an important step toward trustworthy and reliable deployment of LVLMs in the real-world.

Limitations. A potential limitation of our method is its strong reliance on the semantics encoded in vision tokens, which may lead to overfitting to visually salient cues. While this behavior improves visual grounding in most cases, it can result in an overemphasis on perceptually grounded elements at the expense of implicit, commonsense, or non-visual contextual information. To mitigate such overfitting issue, future work could explore adaptive interpolation strategies between reference logits and vanilla logits, or incorporate auxiliary confidence calibration mechanisms to balance multimodal importance dynamically. Another limitation of our approach is its exclusive reliance on the model’s internal representations without access to external knowledge sources. As a result, it cannot correct factual errors or hallucinations originating from pretraining data, particularly in cases where visual evidence alone is insufficient for disambiguation or grounding. While our method improves the alignment between visual inputs and generated text, it does not incorporate retrieval mechanisms that could supplement the model with up-to-date or context-specific information. Nevertheless, due to its modular and decoding-time nature, our method is compatible with retrieval-augmented generation (RAG) pipelines and could be integrated with external knowledge modules in future work.

Rank	Token	Base Logit		(229) Vision Token Logit		Final Logit	
		Log-Prob	Prob (%)	Log-Prob	Prob (%)	Log-Prob	Prob (%)
1	three	-1.04	35.45	-1.47	22.98	-2.51	58.27
2	four	-0.68	50.78	-2.25	10.52	-2.93	38.21
3	two	-4.57	1.04	-1.19	30.44	-5.76	2.26
4	five	-2.54	7.91	-4.45	1.16	-6.99	0.66
5	a	-4.33	1.31	-3.08	4.58	-7.42	0.43
6	six	-3.94	1.94	-5.15	0.58	-9.10	0.08
7		-5.10	0.61	-4.22	1.47	-9.32	0.06
8	several	-5.75	0.32	-5.39	0.45	-11.15	0.01
9	seven	-5.79	0.31	-6.12	0.22	-11.90	0.00
10	un	-8.65	0.02	-3.80	2.24	-12.45	0.00
11	at	-8.28	0.03	-4.77	0.85	-13.05	0.00
12	one	-10.05	0.00	-3.69	2.50	-13.74	0.00
13	in	-11.15	0.00	-2.99	5.01	-14.14	0.00
14	eight	-7.38	0.06	-6.78	0.11	-14.16	0.00
15	f	-9.52	0.01	-4.76	0.86	-14.28	0.00
16	multiple	-7.66	0.05	-6.66	0.13	-14.32	0.00
17	only	-8.68	0.02	-5.89	0.28	-14.57	0.00
18	nine	-7.71	0.04	-6.97	0.09	-14.68	0.00
19	half	-8.62	0.02	-6.22	0.20	-14.85	0.00
20	many	-8.70	0.02	-6.21	0.20	-14.91	0.00
21	all	-10.38	0.00	-4.85	0.78	-15.23	0.00
22	still	-11.11	0.00	-4.15	1.57	-15.26	0.00
23	an	-10.33	0.00	-5.10	0.61	-15.43	0.00
24	some	-10.05	0.00	-5.40	0.45	-15.46	0.00
25	more	-11.15	0.00	-4.68	0.92	-15.83	0.00
26	about	-10.05	0.00	-5.90	0.27	-15.95	0.00
27	both	-10.74	0.00	-5.23	0.53	-15.98	0.00
28	the	-11.98	0.00	-4.07	1.71	-16.05	0.00
29	not	-11.58	0.00	-4.59	1.02	-16.16	0.00
30	over	-11.99	0.00	-4.25	1.42	-16.25	0.00
31	different	-11.93	0.00	-4.73	0.88	-16.66	0.00
32	ten	-8.56	0.02	-8.15	0.03	-16.71	0.00
33	just	-10.97	0.00	-5.78	0.31	-16.75	0.00
34	fruit	-11.81	0.00	-5.04	0.65	-16.85	0.00
35	as	-12.06	0.00	-4.82	0.81	-16.88	0.00
36	small	-12.11	0.00	-4.86	0.78	-16.97	0.00
37	various	-10.39	0.00	-6.71	0.12	-17.10	0.00
38	also	-11.94	0.00	-5.21	0.55	-17.15	0.00
39	very	-11.80	0.00	-5.54	0.39	-17.35	0.00
40	around	-11.28	0.00	-6.07	0.23	-17.35	0.00
41	numerous	-10.02	0.00	-7.51	0.05	-17.54	0.00
42	no	-10.70	0.00	-6.93	0.10	-17.64	0.00
43	total	-11.17	0.00	-6.80	0.11	-17.97	0.00
44	currently	-11.17	0.00	-6.89	0.10	-18.06	0.00
45	cut	-11.78	0.00	-6.46	0.16	-18.24	0.00
46	quite	-11.62	0.00	-6.71	0.12	-18.33	0.00
47	few	-11.91	0.00	-6.64	0.13	-18.55	0.00
48	twelve	-10.07	0.00	-8.53	0.02	-18.60	0.00
49	lots	-11.52	0.00	-7.37	0.06	-18.89	0.00
50	th	-11.34	0.00	-7.63	0.05	-18.98	0.00
51	actually	-11.84	0.00	-7.40	0.06	-19.24	0.00
52	approximately	-11.05	0.00	-8.45	0.02	-19.50	0.00
53	exactly	-12.14	0.00	-7.48	0.06	-19.62	0.00
54	plenty	-11.78	0.00	-7.95	0.04	-19.73	0.00
55	eleven	-10.82	0.00	-9.25	0.01	-20.07	0.00

Table 8. Detailed numerical values for Figure 1. Token scores over the constrained vocabulary $\mathcal{V}_{\text{cons}}^2$ at decoding step $t = 2$ with ReVisiT are presented. The selected vision token (index 229) amplifies the relevance of \texttt{three}, increasing its probability from 35.45% (base) to 58.27% (final), while suppressing \texttt{four} from 50.78% to 38.21%.

Vision Token	Top (Pink)				Middle (Yellow)				Bottom (Cyan)			
Vocab	\mathcal{V}		\mathcal{V}_{env}		\mathcal{V}		\mathcal{V}_{env}		\mathcal{V}		\mathcal{V}_{env}	
Token	Log-Prob	%	Log-Prob	%	Log-Prob	%	Log-Prob	%	Log-Prob	%	Log-Prob	%
river	-14.74	0.00	-6.16	0.21	-3.61	2.70	-0.36	69.75	-14.58	0.00	-5.02	0.66
mountain	-14.07	0.00	-5.48	0.42	-6.47	0.15	-3.22	3.99	-15.02	0.00	-5.46	0.43
person	-13.39	0.00	-4.80	0.82	-12.27	0.00	-9.01	0.01	-9.86	0.01	-0.30	73.81
sky	-8.66	0.02	-0.07	93.44	-6.76	0.12	-3.51	2.99	-12.82	0.00	-3.26	3.83
tree	-12.06	0.00	-3.48	3.09	-11.71	0.00	-8.46	0.02	-12.47	0.00	-2.92	5.42
cloud	-13.88	0.00	-5.30	0.50	-9.52	0.01	-6.26	0.19	-12.46	0.00	-2.91	5.47
sea	-13.80	0.00	-5.21	0.55	-6.82	0.11	-3.57	2.81	-13.24	0.00	-3.68	2.52
grass	-15.41	0.00	-6.82	0.11	-7.19	0.08	-3.94	1.95	-13.41	0.00	-3.86	2.11
rock	-13.35	0.00	-4.77	0.85	-4.95	0.71	-1.70	18.33	-12.41	0.00	-2.85	5.77

Table 9. Detailed numerical values for Figure 2 (left). Log-probabilities and probabilities (%) of selected constrained vocabulary tokens as projected by three vision tokens (pink, yellow, cyan) under the full vocabulary \mathcal{V} and the constrained vocabulary \mathcal{V}_{env} .

Rank	Token	Base Logit		(193) Vision Token Logit		Final Logit	
		Log-Prob	Prob (%)	Log-Prob	Prob (%)	Log-Prob	Prob (%)
1	painting	-2.52	8.08	-1.27	27.99	-3.79	39.15
2	rock	-2.60	7.42	-2.28	10.22	-4.88	13.11
3	mountain	-2.14	11.76	-3.03	4.83	-5.17	9.82
4	person	-1.89	15.10	-3.59	2.75	-5.48	7.19
5	landscape	-2.27	10.30	-3.42	3.27	-5.70	5.82
6	d	-3.45	3.16	-2.29	10.14	-5.74	5.55
7	p	-3.91	2.01	-2.45	8.60	-6.36	3.00
8	scene	-3.80	2.23	-2.70	6.75	-6.50	2.60
9	small	-2.83	5.91	-3.70	2.48	-6.52	2.54
10	beautiful	-3.67	2.54	-3.03	4.84	-6.70	2.13
11	large	-2.93	5.34	-3.78	2.28	-6.71	2.11
12	scen	-3.57	2.81	-3.17	4.19	-6.74	2.04
13	hill	-4.01	1.82	-2.83	5.91	-6.84	1.86
14	chair	-2.85	5.78	-4.38	1.26	-7.23	1.26
15	bow	-3.87	2.09	-3.68	2.53	-7.54	0.92
16	second	-3.80	2.23	-4.52	1.09	-8.33	0.42
17	smaller	-3.02	4.90	-6.04	0.24	-9.06	0.20
18	boat	-3.00	4.98	-6.34	0.18	-9.34	0.15
19	distant	-4.14	1.59	-5.43	0.44	-9.57	0.12

Table 10. Detailed numerical values for Figure 2 (right). Token scores over the constrained vocabulary \mathcal{V}_{cons}^{72} at decoding step $t = 72$ with ReVisiT. The selected vision token (index 193) amplifies the relevance of \texttt{painting}, increasing its probability from 8.08% (base) to 39.15% (final) while suppressing hallucinatory token \texttt{person} from 15.10% to 7.19%.