LOTUS: A Leaderboard for Detailed Image Captioning from Quality to Societal Bias and User Preferences

Yusuke Hirota[1,2*] **Boyi Li**[1] **Ryo Hachiuma**[1] **Yueh-Hua Wu**[1] **Boris Ivanovic**[1,3]
Yuta Nakashima[2] **Marco Pavone**[1,3] **Yejin Choi**[1] **Yu-Chiang Frank Wang**[1] **Huck Yang**[1]
[1]NVIDIA Research [2]Osaka University [3]Stanford University

Abstract

Large Vision-Language Models (LVLMs) have transformed image captioning, shifting from concise captions to detailed descriptions. We introduce LOTUS, a leaderboard for evaluating detailed captions, addressing three main gaps in existing evaluations: lack of standardized criteria, bias-aware assessments, and user preference considerations. LOTUS comprehensively evaluates various aspects, including caption quality (*e.g.*, alignment, descriptiveness), risks (*e.g.*, hallucination), and societal biases (*e.g.*, gender bias) while enabling preference-oriented evaluations by tailoring criteria to diverse user preferences. Our analysis of recent LVLMs reveals no single model excels across all criteria, while correlations emerge between caption detail and bias risks. Preference-oriented evaluations demonstrate that optimal model selection depends on user priorities. [1]

1 Introduction

Image captioning has evolved with Large Vision-Language Models (LVLMs) such as LLaVA (Liu et al., 2024), moving from generating concise captions (Chen et al., 2015) to more *detailed* descriptions (Chen et al., 2024; Liu et al., 2024). This transition, driven by LVLMs' improved ability to follow instructions, enhances visual-semantic understanding and strengthens vision-language applications, including pre-training (Zheng et al., 2024; Liu et al., 2023b).

A crucial challenge in detailed image captioning lies in effectively evaluating the generated captions. Traditional n-gram-based metrics, such as BLEU (Papineni et al., 2002), which are well-suited for concise captions, prove inadequate for assessing detailed descriptions (Chan et al., 2023). This limitation has spurred the development of new evaluations tailored to detailed captions.

However, we argue that current approaches to evaluating detailed captions face challenges:

Lack of a unified evaluation framework. While existing studies tend to target specific dimensions like descriptiveness, alignment, or hallucination detection, there is no overarching, standardized evaluation framework. This fragmentation leads to inconsistent performance assessments across studies, hindering comparability in the field.

Absence of side-effect evaluation. Despite recent findings (Zhang et al., 2024b) showing that LVLMs often exhibit societal *biases* (e.g., gender bias), current evaluation methods largely overlook these biases, raising the risk of perpetuating harmful stereotypes in generated captions.

User preference-agnostic evaluation. The quality of detailed captions is highly subjective, as system preferences vary significantly. While some users favor highly descriptive captions, others prioritize minimizing risks such as hallucinations. This variability poses a challenge for designing a universal metric that accommodates diverse needs.

In this paper, we contribute to establishing a unified leaderboard, LOTUS (unified LeaderbOard to socieTal bias and USer preferences), that overcomes the challenges in existing evaluations. Specifically, LOTUS 1) **comprehensively evaluates various aspects** of detailed captions (Figure 1 (a)), including caption quality-related criteria (*e.g.*, descriptiveness (Chan et al., 2023), alignment (Li et al., 2024)), potential risks (*e.g.*, hallucinations (Jing et al., 2024)), and societal bias (*e.g.*, gender bias (Buolamwini and Gebru, 2018)), enabling diverse, unified model assessments; 2) **supports preference-oriented evaluation** by tailoring criteria to different user preferences (Figure 1 (b)), allowing for customized assessments that better align with diverse user needs.

Leveraging LOTUS's multifaceted and adaptable framework, we evaluate recent LVLMs (Liu et al., 2024; Dai et al., 2023; Chen et al., 2023; Ye

*Work done as an intern at NVIDIA Research.
[1]Leaderboard: https://lotus-vlm.github.io/

Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track), pages 295–309

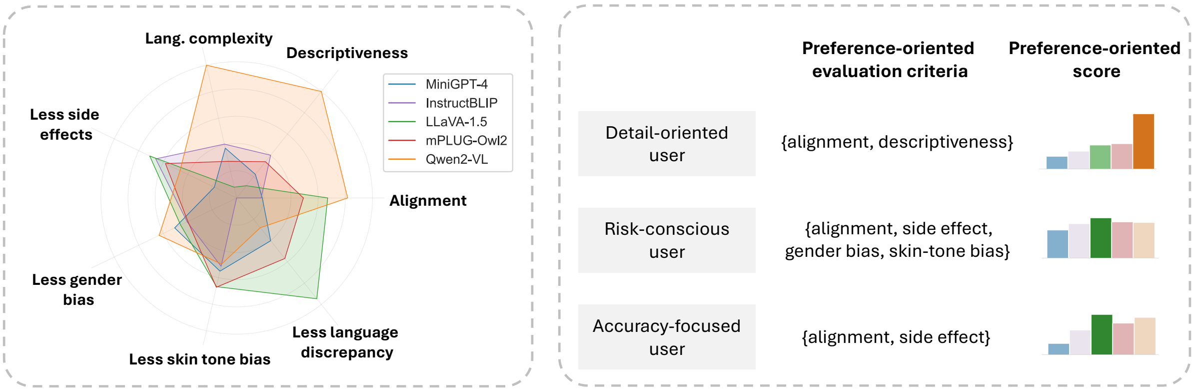

(a) Unified evaluation leaderboard

(b) Preference-oriented evaluation

Figure 1: Overview of the LOTUS leaderboard. LOTUS enables (a) unified evaluation of various aspects of detailed captions, including societal bias, and (b) preference-oriented assessment tailored to different user preferences.

et al., 2024; Wang et al., 2024), uncovering various insights:

- Different models exhibit distinct strengths and weaknesses across various aspects, with no single model consistently performing well across all criteria. For instance, Qwen2-VL (Wang et al., 2024) generates high-quality captions but shows higher risks of hallucination and skin tone bias (Figure 1). This observation highlights the importance of LOTUS's comprehensive evaluation characteristic.
- We discover correlations among evaluation criteria, revealing that models producing more detailed captions tend to have higher risks of specific biases (*e.g.*, skin tone bias) and hallucinations (Figure 2). This finding suggests a potential trade-off between descriptiveness and risk mitigation in caption generation.
- By selecting evaluation criteria based on user preferences, we can accurately reflect what different users value in captions (Figure 1 (b)). For instance, while Qwen2-VL is the best option for users who prioritize caption quality, it is not suitable for those who prefer captions with minimal risks of side effects and social bias. This finding highlights the importance of customized evaluation criteria in addressing the specific needs of diverse users.

2 LOTUS: A Unified Leaderboard for Detailed Captions

As discussed in Section 1, prior work on evaluating detailed captions faces several challenges: 1) lack of a unified evaluation framework, 2) absence of bias-aware evaluation, and 3) user preference-agnostic evaluation. Here, we introduce our proposed leaderboard, LOTUS, which unifies various evaluation criteria (Section 2.1), including societal bias (Section 2.2) and enables preference-oriented evaluation (Section 2.3).

Preliminaries. Let $\mathcal{D} = \{(I, y, a)\}$ denote a test set of the captioning dataset, where I is an image, y is its corresponding ground-truth detailed caption, and a is an optional protected attribute label of the person in the image (*e.g.*, `woman` or `man` for binary gender). Our target task is detailed image captioning: given a prompt[2] p and an image, we use an LVLM M to generate a detailed caption y', *i.e.*, $y' = M(I, p)$.

2.1 Unified and Comprehensive Evaluation

For a comprehensive, multifaceted assessment, LOTUS unifies four main criteria for detailed caption evaluation that have been previously assessed separately: alignment, descriptiveness, language complexity, and side effects. LOTUS incorporates multiple metrics for each criterion to enhance reliability (Naidu et al., 2023). Otherwise stated, the average is computed over $\mathcal{D}$ for each metric. We summarize each **criterion** and its ***metrics***:[3]

Alignment measures how well a caption matches the image content using two metrics: ***CLIPScore*** (Hessel et al., 2021) quantifies the semantic similarity between the image and caption using CLIP embeddings:

$$\text{CLIPScore} = \max(0, \cos(\phi_I(I), \phi_T(y'))) \quad (1)$$

[2] We use "Describe this image in detail." as the prompt.
[3] Detailed metric descriptions are in Appendix E.

where ϕ_I and ϕ_T are CLIP image and text encoders,[4] and $\cos(\cdot, \cdot)$ denotes cosine similarity. ***CapScore*** (Li et al., 2024) prompts GPT-4 to rate a caption based on its similarity to the ground truth (CapScore_S) and alignment (CapScore_A), both ranging from 0 to 1.

Descriptiveness evaluates how detailed a caption is in describing image elements using two metrics: ***CLIP recall*** (Chan et al., 2023) evaluates whether a caption is specific enough to identify its corresponding image. Specifically, CLIPScore is computed between the image I and all generated captions, and Recall@k determines if the correct caption y' appears in the top-k most similar captions. ***Noun and verb coverage*** (Chan et al., 2023) assesses how well a caption y' covers key objects (nouns) and actions (verbs) present in an image by comparing it to the ground-truth y. Noun coverage is calculated as:

$$\text{Noun Coverage} = \frac{|N(y) \cap N(y')|}{|N(y')|} \tag{2}$$

where $N(y')$ is the set of all nouns in y'. Verb coverage is calculated for verbs likewise.

Language complexity (Onoe et al., 2024) evaluates the structural complexity of the sentences and language used in captions. We use the following metrics: ***Syntactic complexity*** measures the maximum depth of the dependency tree (Ohta and Sakai, 2017) of y'. A greater depth indicates a more complex sentence structure. ***Semantic complexity*** is indicated by the number of nodes in a scene graph derived from y' (Spacy, 2024). A higher number of nodes suggests a more detailed representation of objects and attributes within the scene.

Side effects identify negative aspects in captions. We consider two issues: hallucination and harmfulness (*i.e.*, existence of NSFW (Not safe for work) words) for this criterion. We assess hallucination through two methods: ***CHAIR$_s$*** (Rohrbach et al., 2018) quantifies object hallucination by computing the fraction of objects in y' that are not present in the image I:

$$\text{CHAIR}_s = \frac{O_H}{O_T}, \tag{3}$$

where O_H is the number of hallucinated objects, and O_T is the total number of annotated objects. ***FaithScore*** (Jing et al., 2024) evaluates the faithfulness of long captions by breaking down each caption into atomic *facts* that represent specific, verifiable statements about the image content. Let V denote an indicator function of visual entailment (Wang et al., 2022), giving 1 if f is entailed by I, and 0 otherwise. Each atomic fact f_k (*e.g.*, "A man playing baseball") is checked with V to compute FaithScore as:

$$\text{FaithScore} = \frac{1}{K} \sum_{k=1}^{K} V(f_k, I) \tag{4}$$

where K is the total number of facts. Additionally, we employ a sentence-level FaithScore, FaithScore_S, which measures the proportion of sentences in y' that are free from hallucinations.

To evaluate the harmfulness of captions, we examine the ***existence of NSFW words***[5] in y'. Specifically, if y' contains an NSFW word, this metric gives 1 (which is averaged over $\mathcal{D}$).

2.2 Bias-Aware Evaluation

LOTUS not only unifies various criteria but also addresses a critical aspect often overlooked in prior work: societal bias. Following previous works (Zhao et al., 2021; Tang et al., 2021), we examine binary **gender and skin tone biases**.

To measure societal bias, we use a popular and standard way of quantifying bias, ***performance disparity*** (Buolamwini and Gebru, 2018), comparing the performance of the captioning model across different demographic groups. In the case of binary gender bias (*i.e.*, $a \in \{\texttt{woman}, \texttt{man}\}$), we first prepare two separate sets of woman and man images, $\mathcal{D}_{\texttt{woman}}$ and $\mathcal{D}_{\texttt{man}}$:

$$\mathcal{D}_g = \{(I, y, a) \in \mathcal{D} | a = g\}, \tag{5}$$

where $g \in \{\texttt{woman}, \texttt{man}\}$. For each set, we generate detailed captions, obtaining $\mathcal{D}'_g = \{(I, y', a) \mid y' = M(I, p)\}$. The performance disparity is defined as the absolute difference in performance between $\mathcal{D}'_{\texttt{woman}}$ and $\mathcal{D}'_{\texttt{man}}$.[6] We compute performance disparity for each metric in Section 2.1. For skin tone bias, we conduct the same process based on the binary skin tone class (*i.e.*, $a \in \{\texttt{darker-skin}, \texttt{lighter-skin}\}$).

Beyond societal bias, we also investigate **language discrepancy**. We examine how the choice of prompt language affects the model's performance

[4] To handle detailed input captions, we utilize the CLIP variant (Zhang et al., 2024a) capable of processing long text.

[5] We adopt the NSFW word list in (LDNOOBW, 2024).

[6] Note that the average is computed over $\mathcal{D}'_g$.

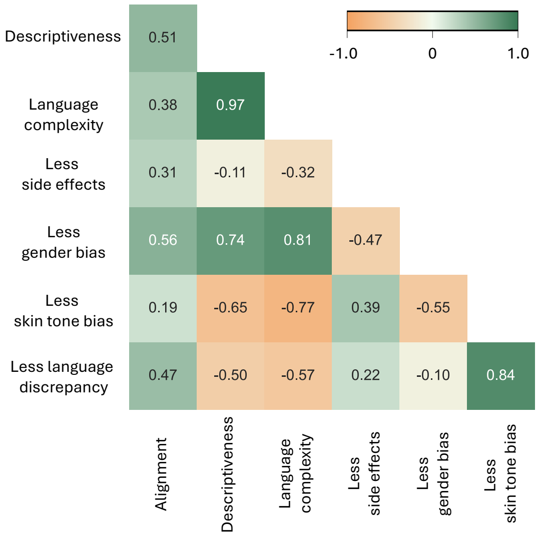

Figure 2: Correlation matrix of evaluation criteria.

across different languages. Let $\mathcal{L}$ be a set of languages. For each language $l \in \mathcal{L}$, we use a prompt[7] p_l in that language to generate captions and evaluate their performance using the same metrics as in Section 2.1. In our experiments, we consider three languages $\mathcal{L} = \{\texttt{English}, \texttt{Japanese}, \texttt{Chinese}\}$. As in societal bias, we define language discrepancy as the performance disparity between the best- and worst-performing languages.

2.3 User Preference-Oriented Evaluation

While our unified criteria offer diverse model evaluations, another benefit is the ability to tailor evaluations to specific user preferences. To achieve this, we categorize user types based on different priorities in captioning as shown Figure 1 (b). For example, a *detail-oriented user* may prioritize metrics that assess descriptiveness, whereas a *risk-conscious user* might emphasize minimizing side effects and societal bias. By selecting criteria that align with these user profiles, our framework provides a prioritized assessment of model performance (*e.g.*, selecting "alignment" and "descriptiveness" for *detail-oriented user*). This preference-oriented approach allows for a more specific evaluation of model performance, demonstrating that tailored criteria can effectively capture the preferences of each user type (Section 3.2).

3 Experiments

Dataset. We evaluate captioning models on the COCO Karpathy test set ($5,000$ images) (Karpathy and Fei-Fei, 2015). For societal bias analysis, we use binary gender and skin tone annotations from (Zhao et al., 2021), sampling images to balance demographic groups (*e.g.*, 6,628 for gender, 2,192 for skin tone). Ground-truth detailed captions are sourced from the Localized Narratives dataset (Pont-Tuset et al., 2020).

Evaluation metrics. We use the evaluation metrics summarized in Section 2.1 and compute the **normalized average score** (N-avg) to summarize each criterion. For each criterion, scores are Min-Max normalized to [0, 1], with inversion applied for metrics where lower is better (*e.g.*, CHAIR). N-avg is then calculated as the mean of normalized scores per criterion, such as CLIPScore and CapScores for alignment. For gender and skin tone biases and language discrepancy, the N-avg score is the mean of normalized performance disparity scores across all metrics.

Captioning models. We evaluate detailed captions from five representative LVLMs: MiniGPT-4 (Chen et al., 2023), InstructBLIP (Dai et al., 2023), LLaVA-1.5 (Liu et al., 2024), mPLUG-Owl2 (Ye et al., 2024), and Qwen2-VL (Wang et al., 2024). To ensure a fair comparison, we use the 7B parameter variant for all models, as this size is commonly available across these models.

3.1 Results on LOTUS

Tables 1 and 2 present the results of the four criteria in Section 2.1 and bias-aware evaluation. Additionally, we visualize the normalized average scores (N-avg in the tables) in Figure 1 (a). The visual examples of the generated captions are shown in Figure 9. The key findings are summarized below:

Models show varied performance across criteria, with no model excelling in all areas. The N-avg scores for each criterion and Figure 1 (a) indicate that models have distinct strengths and weaknesses. For example, Qwen2-VL performs the best on criteria related to caption quality (*i.e.*, alignment, descriptiveness, complexity) but scores relatively lower on side effects (0.46). Also, it shows a strong skin bias tone and language discrepancy, showing the lowest scores for both criteria. Conversely, LLaVA-1.5, while weaker in descriptiveness and complexity, has minimal side effects and skin tone bias, complementing Qwen2-VL. This underscores the value of unified evaluation criteria to reveal each model's unique strengths and weaknesses.

Unexpected trade-offs emerge from criteria cor-

[7] For each language $l \neq \texttt{English}$, we use the prompt "Describe this image in detail in English" translated into l.

Table 1: Unified evaluation of LVLM captioners on LOTUS with CLIPScore (CLIP-S), CapScore (CapS_S, CapS_A), CLIP recall (recall), noun/verb coverage (noun, verb), syntactic and semantic complexities (syn, sem), CHAIR_s (CH_s), FaithScore (FS, FS_s), and existence of NSFW words (harm). Values in **bold** and underline indicate the best and second-best, respectively. All metrics are scaled by 100.

Model	Alignment ↑				Descriptiveness ↑				Complexity ↑			Side effects				
	CLIP-S	CapS_S	CapS_A	N-avg	Recall	Noun	Verb	N-avg	Syn	Sem	N-avg	CH_s ↓	FS ↑	FS_s ↑	Harm ↓	N-avg ↑
MiniGPT-4	60.8	33.0	35.9	0.19	75.3	33.0	34.7	0.22	8.0	32.6	0.38	37.8	55.0	37.6	0.31	0.18
InstructBLIP	59.9	36.0	35.5	0.18	82.1	34.2	34.7	0.40	7.7	46.0	0.41	58.5	62.4	**43.3**	0.10	0.66
LLaVA-1.5	60.1	38.5	**45.0**	0.67	80.5	32.5	31.0	0.11	7.1	39.6	0.08	49.0	**65.7**	41.6	0.12	**0.71**
mPLUG-Owl2	59.7	**39.7**	40.0	0.49	83.3	35.0	32.8	0.34	7.4	45.6	0.28	59.1	62.0	41.3	**0.08**	0.58
Qwen2-VL	**61.8**	37.3	43.2	**0.82**	**90.4**	**45.9**	**36.9**	**1.00**	**8.3**	**75.7**	**1.00**	**26.8**	54.2	41.7	0.28	0.46

Table 2: Bias-aware evaluation of LVLM captioners on LOTUS. Language discrepancy evaluation cannot be applicable to InstructBLIP due to a lack of Japanese support. **Bold** and underline indicate the best and second-best, respectively. All metrics are scaled by 100.

Model	Alignment			Descriptiveness			Complexity		Side effects				
	CLIP-S	CapS_S	CapS_A	Recall	Noun	Verb	Syn	Sem	CH_s	FS	FS_S	Harm	N-avg↑
Gender bias													
MiniGPT-4	0.3	0.9	1.1	7.8	1.7	**2.6**	6.3	3.2	4.8	6.3	4.0	1.64	0.51
InstructBLIP	0.8	2.7	1.2	8.4	1.9	3.3	**1.0**	0.1	6.8	3.8	5.0	0.72	0.40
LLaVA-1.5	0.7	2.2	0.7	9.5	2.2	4.1	1.5	0.2	7.6	3.8	**3.7**	0.39	0.46
mPLUG-Owl2	0.6	2.2	1.2	9.1	2.3	3.5	1.6	**0.0**	7.2	3.1	5.8	**0.33**	0.40
Qwen2-VL	**0.2**	**0.7**	**0.5**	**6.3**	**0.1**	3.6	13.5	2.5	**4.4**	**0.9**	5.7	1.77	**0.63**
Skin tone bias													
MiniGPT-4	0.8	1.5	0.8	4.8	**0.2**	2.3	19.4	0.2	2.0	0.9	0.5	0.09	0.55
InstructBLIP	0.5	1.4	**0.2**	8.4	1.9	1.1	6.8	**0.1**	4.0	2.4	1.1	0.09	0.51
LLaVA-1.5	0.4	1.3	0.7	4.0	**0.2**	**1.0**	**5.3**	0.6	2.7	1.4	1.3	0.18	**0.67**
mPLUG-Owl2	0.6	1.9	0.5	5.1	0.8	2.2	7.6	0.4	**1.7**	**0.1**	**0.4**	**0.00**	**0.67**
Qwen2-VL	**0.2**	**1.1**	1.5	**2.3**	0.5	1.3	14.9	2.3	2.7	3.1	1.8	0.09	0.50
Language discrepancy													
MiniGPT-4	0.8	1.5	3.9	2.3	4.3	5.2	52.2	5.0	5.4	5.6	3.4	0.10	0.40
InstructBLIP	-	-	-	-	-	-	-	-	-	-	-	-	-
LLaVA-1.5	0.4	**0.8**	**2.0**	**1.1**	**1.1**	**1.8**	**11.4**	**1.8**	**4.7**	**2.0**	1.6	0.06	**0.95**
mPLUG-Owl2	1.4	1.6	4.9	1.5	**1.1**	3.7	37.5	8.4	17.0	6.3	**1.3**	**0.02**	0.57
Qwen2-VL	**0.2**	3.6	6.7	1.9	3.9	3.8	90.8	26.2	6.4	7.5	2.1	0.14	0.28

relations. The correlation analysis of our evaluation criteria in Figure 2 reveals several intriguing patterns in LVLM captioner performance:

1. Models with better descriptiveness tend to give less gender bias but more skin tone bias (0.74 and -0.65, respectively). This suggests a potential trade-off between information richness and different aspects of fairness.
2. Side effects have only weak to moderate correlations with other criteria (ranging from -0.47 to 0.39), implying that hallucinations or NSFW content might not significantly impact caption quality or societal bias.
3. Gender bias and skin tone bias show a moderate negative correlation (-0.55), indicating an inverse relationship between these two biases. This highlights the complexity of addressing multiple aspects of fairness simultaneously.
4. Alignment correlates positively with all other criteria, suggesting that improvements in one area often enhance image-caption alignment, though to varying extents.

These findings underscore the intricate interplay between different performance aspects in LVLM captioners, emphasizing the need for a holistic approach to model improvement that considers multiple criteria simultaneously.

Descriptiveness amplifies societal bias trade-offs. To further explore why higher descriptiveness reduces gender bias but amplifies skin tone bias (observations 1 and 3 above), we analyze gender and skin tone representation in captions. For gender bias, we calculate the difference ($|\Delta|$) between the ratio of captions mentioning female-related terms[8] for woman images (recall_F) and male-related terms for man images (recall_M). For skin tone bias, we compare the ratio of captions containing race-related terms[9] for images of individuals with darker skin tones (recall_D) versus lighter skin tones (recall_L). We then examine correlations between $|\Delta|$ values and our descriptiveness and bias scores from Tables 1 and 2 (N-avg).

Table 3 presents the recall values (%) and $|\Delta|$

[8]We use the gender word list in (Hirota et al., 2023).

[9]We use race-related terms defined in (Hirota et al., 2025).

Table 3: Gender and skin tone representations in generated captions. $Rec_{F/M}$ denotes recall of gender terms for woman/man images. $Rec_{D/L}$ represents recall of racial terms for darker/lighter skin. $|\Delta|$ is recall disparities.

Model	Gender images			Skin images		
	Rec_F	Rec_M	$\|\Delta\|$	Rec_D	Rec_L	$\|\Delta\|$
MiniGPT.	68.0	71.2	3.2	3.0	2.3	0.7
Instruct.	75.3	78.7	3.4	1.1	0.7	0.4
LLaVA.	74.0	80.1	6.1	0.3	0.4	0.1
mPLUG.	77.9	82.0	4.1	0.6	0.6	0.0
Qwen2.	41.0	40.7	0.3	7.0	2.9	4.1

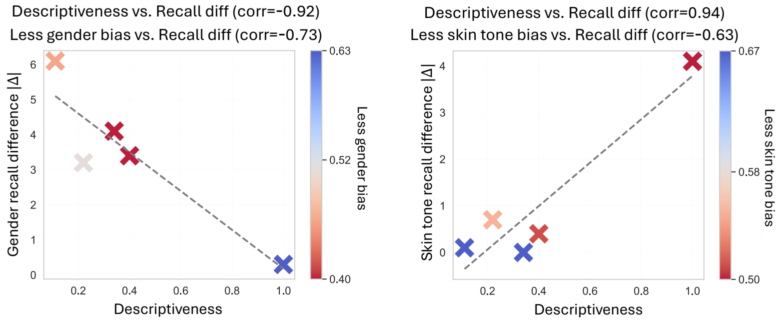

Figure 3: Correlations between descriptiveness, gender/skin tone bias, and Δ. Descriptiveness and gender/skin tone bias are the normalized average scores in Tables 1 and 2 (N-avg).

for gender and skin tone biases, while Figure 3 visualizes the correlations between descriptiveness, gender/skin tone bias scores, and the $|\Delta|$ values. The results indicate that more descriptive models tend to have smaller gender representation disparities (corr $= -0.92$) but larger differences in racial word usage based on skin tone (corr $= 0.94$). We observe strong correlations between these disparities and less gender and skin tone biases (corr $= -0.73$ and -0.63, respectively).

This suggests that as captions become more descriptive, the gender term usage gap between woman and man images narrows, likely because gender tends to be described for both genders (Hirota et al., 2023). Consequently, with increased descriptiveness, models tend to include gender terms regardless of specific gender. For racial attributes, while captioning models generally avoid racial terms, they more frequently describe minoritized groups, such as people of color, than White individuals (Zhao et al., 2021). As descriptiveness rises, racial term usage increases, and due to inherent skin tone bias, this leads to greater disparities in racial term usage for darker-skinned individuals.

3.2 Results for Preference-Oriented Evaluation

As introduced in Section 2.3, our evaluation framework supports assessments tailored to user preferences. To demonstrate this, we consider three user types: 1) **Detail-oriented users** prioritize comprehensive descriptions that cover detailed contents in images (selected criteria: {alignment, descriptiveness}), **Risk-conscious users** seek to minimize risks like hallucinations and biases (selected criteria: {alignment, side effects, gender bias, skin-tone bias}), and 3) **Accuracy-focused users** value fact-based, error-free captions (selected criteria: {alignment, side effects}).

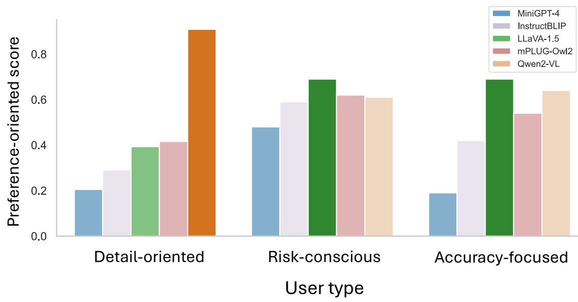

Figure 4: Preference-oriented scores for detail-oriented user (left), risk-conscious user (middle), and accuracy-focused user (right). The best models for each user type are highlighted in darker colors.

In Figure 4, we show the *preference-oriented scores* for each user type, computed by taking the average of the N-avg scores of the selected criteria. The figure demonstrates that the performance of models greatly varies depending on user preferences. For *detail-oriented user*, Qwen2-VL can be the best option, presenting much higher scores than the other models. However, for the users who focus on the risks (*i.e.*, *risk-conscious user*), LLaVA-1.5 might be the most suitable to reduce the risks of generating captions with hallucinations, NSFW words, and societal bias. Similarly, LLaVA-1.5 also performs best for the *accuracy-focused user*, indicating its strength in producing reliable and precise captions. These results highlight that LVLM captioning models should be chosen based on specific user needs, not a universal approach. [10]

4 Related Work

Detailed image captioning. Recent advancements in LVLMs have significantly enhanced mul-

[10] In Appendix B, we validate whether our preference-oriented evaluation accurately reflects real users' preferences through LLM agent-simulated analysis.

timodal understanding (Liu et al., 2024; Ye et al., 2024). Techniques like visual instruction tuning (Liu et al., 2023a), which combines visual inputs with textual guidance during training, enable LVLMs to effectively follow user instructions. Leveraging these advancements, recent research (Chen et al., 2024; Lai et al., 2023) has explored generating detailed image descriptions to improve alignment and utility for downstream tasks. For instance, Zheng et al. (2024) proposed a pipeline using detailed captions from LVLMs (*i.e.*, LLaVA-1.5 (Liu et al., 2024)) for pre-training, boosting the performance of CLIP (Radford et al., 2021).

Evaluation for detailed captions. A critical challenge in detailed image captioning is evaluating generated captions. Conventional metrics like CIDEr (Vedantam et al., 2015) are inadequate for assessing detailed captions (Chan et al., 2023), prompting researchers to develop new methods. For example, Chan et al. (2023) proposed measuring noun and verb coverage by comparing these elements in generated and ground-truth captions.

However, as discussed in Section 1, existing works lack a unified evaluation framework and often overlook societal biases. To address these limitations, we propose LOTUS, a unified evaluation leaderboard for detailed captions. LOTUS provides a comprehensive assessment across multiple dimensions, including previously underexplored areas such as gender and skin tone bias.

5 Conclusion

We introduced LOTUS, a unified leaderboard for evaluating detailed captions from LVLMs. Our analysis uncovered insights unexplored in the existing literature: a trade-off between caption descriptiveness and bias risks, and the impact of user preferences on optimal model selection. LOTUS paves the way for detailed captioning models that holistically optimize performance, mitigate societal biases, and adapt to diverse user preferences.

Ethical Considerations

LOTUS integrates the evaluation of societal biases, including gender, skin tone, and language bias, emphasizing the ethical considerations central to LVLM development. However, it is important to recognize that LOTUS does not capture all potential societal biases, and its scores should not be viewed as a comprehensive measure of a model's bias.

For instance, researchers and practitioners must exercise caution when interpreting LOTUS scores. A favorable score does not imply that a model is free of bias. LOTUS should be seen as one of several tools for evaluating LVLMs, rather than a definitive measure of ethical integrity.

The definition and assessment of bias can vary significantly depending on the context. While LOTUS provides a standardized approach, it may not be universally applicable. We encourage users to critically assess its relevance to their specific use cases and to complement LOTUS with additional bias evaluation methods when appropriate. In sum, by acknowledging these limitations, we advocate for a more nuanced and holistic approach to addressing societal biases in LVLMs, fostering the responsible and ethical development of these technologies.

Fairness recommendations. While we categorized different user types and validated that our user-oriented evaluation can meet the user needs for each type in Section 3.2, we recommend that fairness criteria (*i.e.*, gender and skin tone biases) be considered for all users. Recent works (Zhao et al., 2021; Hirota et al., 2023; Burns et al., 2018; Garcia et al., 2023; Hirota et al., 2022) have demonstrated that image captioning models can perpetuate or amplify societal bias in training datasets, resulting in harmful descriptions for minoritized demographic groups. To mitigate such risks, we emphasize the importance of incorporating fairness criteria into caption evaluation.

Use of binary gender and skin tone categories. In our study, we employed a binary approach to evaluate gender and skin tone biases, classifying gender as female/male and skin tone as darker/lighter, in line with prior work (Zhao et al., 2017; Burns et al., 2018; Wang et al., 2019; Zhao et al., 2023, 2021; Hirota et al., 2024). While this approach addresses bias to some extent, we acknowledge its limitations in reflecting the complexity of real-world diversity. As more comprehensive data becomes available, future research will aim to incorporate non-binary gender categories and more nuanced skin tone classifications.

References

Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *FAccT*.

Kaylee Burns, Lisa Anne Hendricks, Kate Saenko, Trevor Darrell, and Anna Rohrbach. 2018. Women also snowboard: Overcoming bias in captioning models. In *ECCV*.

David Chan, Austin Myers, Sudheendra Vijayanarasimhan, David A Ross, and John Canny. 2023. Ic3: Image captioning by committee consensus. In *EMNLP*.

Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. 2023. Minigpt-v2: large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*.

Lin Chen, Jisong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. 2024. Sharegpt4v: Improving large multi-modal models with better captions. In *ECCV*.

Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.

Cheng-Han Chiang and Hung-yi Lee. 2023. Can large language models be an alternative to human evaluations? In *ACL*.

Wenliang Dai, Junnan Li, AMH Li, Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. In *NeurIPS*.

Noa Garcia, Yusuke Hirota, Yankun Wu, and Yuta Nakashima. 2023. Uncurated image-text datasets: Shedding light on demographic bias. In *CVPR*.

Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. Clipscore: A reference-free evaluation metric for image captioning. In *EMNLP*.

Yusuke Hirota, Jerone TA Andrew, Dora Zhao, Orestis Papakyriakopoulos, Apostolos Modas, Yuta Nakashima, and Alice Xiang. 2024. Resampled datasets are not enough: Mitigating societal bias beyond single attributes. In *EMNLP*.

Yusuke Hirota, Min-Hung Chen, Chien-Yi Wang, Yuta Nakashima, Yu-Chiang Frank Wang, and Ryo Hachiuma. 2025. Saner: Annotation-free societal attribute neutralizer for debiasing clip. In *ICLR*.

Yusuke Hirota, Yuta Nakashima, and Noa Garcia. 2022. Quantifying societal bias amplification in image captioning. In *CVPR*.

Yusuke Hirota, Yuta Nakashima, and Noa Garcia. 2023. Model-agnostic gender debiased image captioning. In *CVPR*.

Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming Zhang, and Nenghai Yu. 2024. Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation. In *CVPR*.

Liqiang Jing, Ruosen Li, Yunmo Chen, Mengzhao Jia, and Xinya Du. 2024. Faithscore: Evaluating hallucinations in large vision-language models. In *EMNLP*.

Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *CVPR*.

Zhengfeng Lai, Haotian Zhang, Wentao Wu, Haoping Bai, Aleksei Timofeev, Xianzhi Du, Zhe Gan, Jiulong Shan, Chen-Nee Chuah, Yinfei Yang, et al. 2023. Veclip: Improving clip training via visual-enriched captions. *arXiv preprint arXiv:2310.07699*.

LDNOOBW. 2024. List-of-dirty-naughty-obscene-and-otherwise-bad-words. Accessed: 2024-10-13.

Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. 2024. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In *CVPR*.

Boyi Li, Ligeng Zhu, Ran Tian, Shuhan Tan, Yuxiao Chen, Yao Lu, Yin Cui, Sushant Veer, Max Ehrlich, Jonah Philion, et al. 2024. Wolf: Captioning everything with a world summarization framework. *arXiv preprint arXiv:2407.18908*.

Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. Improved baselines with visual instruction tuning. In *CVPR*.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. Visual instruction tuning. *NeurIPS*.

Yanqing Liu, Kai Wang, Wenqi Shao, Ping Luo, Yu Qiao, Mike Zheng Shou, Kaipeng Zhang, and Yang You. 2023b. Mllms-augmented visual-language representation learning. *arXiv preprint arXiv:2311.18765*.

Gireen Naidu, Tranos Zuva, and Elias Mmbongeni Sibanda. 2023. A review of evaluation metrics in machine learning algorithms. In *Computer Science On-line Conference*.

Shinri Ohta and Kuniyoshi L Sakai. 2017. Computational principles of syntax in the regions specialized for language: integrating theoretical linguistics and functional neuroimaging. In *Merge in the Mind-Brain*.

Yasumasa Onoe, Sunayana Rane, Zachary Berger, Yonatan Bitton, Jaemin Cho, Roopal Garg, Alexander Ku, Zarana Parekh, Jordi Pont-Tuset, Garrett Tanzer, et al. 2024. Docci: Descriptions of connected and contrasting images. In *ECCV*.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*.

Jordi Pont-Tuset, Jasper Uijlings, Soravit Changpinyo, Radu Soricut, and Vittorio Ferrari. 2020. Connecting vision and language with localized narratives. In *ECCV*.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.

Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. Object hallucination in image captioning. In *EMNLP*.

Spacy. 2024. Spacy: linguistic features. Accessed: 2024-10-13.

Ruixiang Tang, Mengnan Du, Yuening Li, Zirui Liu, Na Zou, and Xia Hu. 2021. Mitigating gender bias in captioning systems. In *WWW*.

Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *CVPR*.

Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. 2024. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.

Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *ICML*, pages 23318–23340.

Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. 2019. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *ICCV*.

Qinghao Ye, Haiyang Xu, Jiabo Ye, Ming Yan, Anwen Hu, Haowei Liu, Qi Qian, Ji Zhang, and Fei Huang. 2024. mplug-owl2: Revolutionizing multi-modal large language model with modality collaboration. In *CVPR*.

Beichen Zhang, Pan Zhang, Xiaoyi Dong, Yuhang Zang, and Jiaqi Wang. 2024a. Long-clip: Unlocking the long-text capability of clip. In *ECCV*.

Jie Zhang, Sibo Wang, Xiangkui Cao, Zheng Yuan, Shiguang Shan, Xilin Chen, and Wen Gao. 2024b. Vlbiasbench: A comprehensive benchmark for evaluating bias in large vision-language model. *arXiv preprint arXiv:2406.14194*.

Dora Zhao, Jerone TA Andrews, and Alice Xiang. 2023. Men also do laundry: Multi-attribute bias amplification. In *ICML*.

Dora Zhao, Angelina Wang, and Olga Russakovsky. 2021. Understanding and evaluating racial biases in image captioning. In *ICCV*.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *EMNLP*.

Kecheng Zheng, Yifei Zhang, Wei Wu, Fan Lu, Shuailei Ma, Xin Jin, Wei Chen, and Yujun Shen. 2024. Dreamlip: Language-image pre-training with long captions. In *ECCV*.

A Detailed Experimental Settings

In this section, we provide the details of the experiments.

A.1 LLM-agent based evaluation

In Section 3.2, we conduct an experiment to validate whether our preference-oriented scores accurately reflect real users' preferences. For this experiment, we rely on GPT-4o instead of human workers, simulating humans. Specifically, we give an instruction prompt to simulate a specific user type and rate the generated caption based on the specified user type. The simulated prompts for each user type are shown in Figure 5. Using these prompts, we compute the simulated user scores (*i.e.*, answers to the question "How well does this caption meet your expectations for describing the image?", rating from 1 to 10). Then, we take an average over the dataset.

A.2 Instruct prompts for LVLMs

The prompts to generate detailed captions, including the ones written in English, Japanese, and Chinese, are presented in Figure 6.

B User-Simulation

How well do our preference-oriented scores match real users' preferences? While our preference-oriented evaluation offers valuable insights, it is essential to validate whether our scoring system accurately reflects real users' preferences. To this end, we use GPT-4o to simulate real user feedback based on recent findings on language models' ability to simulate human responses (Chiang and Lee, 2023), addressing the challenges of recruiting a large, diverse user base.

Figure 8 depicts our evaluation pipeline. We first instruct GPT-4o to simulate specific user types using prompts that reflect each user type's preferences, then rate captions on a 1-10 scale (refer to the simulated user prompt in Figure 8). For example, a prompt for the risk-conscious user focuses on minimizing potential risks in captions. We compare these simulated user scores with our preference-oriented scores to assess the alignment between our framework and simulated user preferences.

Figure 7 presents high correlations between the simulated user scores and our preference-oriented scores (*e.g.*, for risk-conscious users, corr $= 0.84$ between simulated scores and preference-oriented scores). These results indicate that tailored sets of criteria are well-aligned with what actual users would likely prioritize in generated captions.

C Visual examples

Figure 9 shows examples of the generated captions by MiniGPT-4, InstructBLIP, LLaVA-1.5, mPLUG-Owl2, and Qwen2-VL. This figure demonstrates the characteristics of each model. For example, Qwen2-VL gives more detailed and informative sentences compared to the other models, which is consistent with the results in LOTUS (*i.e.*, Qwen2-VL has the best scores for *descriptiveness*). However, only Qwen2-VL contains a race-related word "India" in the first sentence, which cannot be guessed from this image. Based on our evaluation of the relationship between skin tone bias and the existence of race-related terms, this observation can further validate the experimental results on LOTUS, where Qwen2-VL shows the worst score for skin tone bias.

D LOTUS leaderboard

In Figures 10 and 11, we show the actual pages of our LOTUS leaderboard for the unified evaluation (Figure 10) and bias-aware evaluation (Figure 11). The link to the leaderboard is `https://lotus-vlm.github.io/` (anonymized).

E Detailed Explanation of the Evaluation Metrics

In this section, we provide detailed explanations of the metrics used in LOTUS.

CapScore. (Li et al., 2024) Leveraging the ability of GPT-4 to understand and comprehend the long input texts, CapScore utilizes GPT-4 to rate a generated caption. We show the prompt to give GPT-4, evaluating the generated captions in the two criteria (CapScore$_S$ and CapSocre$_A$):

> Can you evaluate the following generated caption based on two metrics:
>
> 1. Similarity to the ground truth caption: How closely does the generated caption match the ground truth in content and meaning? Provide a score between 0 and 1 (two decimal places).
>
> 2. Absence of hallucinations and misalignments: Does the generated caption avoid incorrect information not present in the ground truth?

Detail-oriented user

You are a Detail-oriented user. You value captions that provide accurate and comprehensive descriptions, covering all aspects of an image thoroughly. You appreciate detailed captions that include as many relevant elements as possible, such as objects, people, actions, and settings.
Please evaluate the following Generated caption based on the following question, rating from 1 to 10:
How well does this caption meet your expectations for describing the image?
You can refer to the Ground truth caption to check the contents of the image. No additional text in the output.
Ground truth caption: {ground-truth caption}
Generated caption: {generated caption}

Risk-conscious user

You are a Risk-conscious user. You focus on minimizing potential risks in captions, such as NSFW content, information not directly observable in the image, and descriptions with societal bias. You value captions that strictly align with what can be verified in the image and avoid adding any speculative details. Captions should be factually accurate, and free from additional assumptions or unnecessary elaboration.
Please evaluate the following Generated caption based on the following question, rating from 1 to 10:
How well does this caption meet your expectations for describing the image?
You can refer to the Ground truth caption to check the contents of the image. No additional text in the output.
Ground truth caption: {ground-truth caption}
Generated caption: {generated caption}

Accuracy-focused user

You are a Accuracy-focused user. You prioritize the factual correctness and quality of the captions and expect them to strictly align with the actual content of the image. You value precise and reliable descriptions, with minimal errors or assumptions. Captions should be high-quality, factually accurate, and free from additional assumptions or unnecessary elaboration.
Please evaluate the following Generated caption based on the following question, rating from 1 to 10:
How well does this caption meet your expectations for describing the image?
You can refer to the Ground truth caption to check the contents of the image. No additional text in the output.
Ground truth caption: {ground-truth caption}
Generated caption: {generated caption}

Figure 5: Simulated user prompts for each user type.

> Provide a score between 0 and 1 (two decimal places). Please output only the two scores separated by a semicolon in the format 'similarity score;hallucination score'. No additional text in the output.
>
> Ground truth caption: {ground-truth caption}
>
> Generated caption: {generated caption}

We compute the average of the scores from the two questions across the test set, obtaining the final CapScore.

CLIP Recall (Chan et al., 2023) is a metric that evaluates how well a generated caption uniquely identifies its corresponding image by checking if the correct caption is within the top 5 closest matches when comparing the image embedding to the caption embeddings. This metric helps determine if the caption includes enough distinctive details to set its image apart from others.

For each image I_i, we use CLIP to generate an embedding $\mathbf{I}_i$ that represents the image. We also generate embeddings for the generated captions associated with this image and other images. Then, we check whether the caption embedding $\mathbf{Y}_i$ of the correct caption appears in the top-5 closest caption embeddings based on similarity to $\mathbf{I}_i$. The Recall@5 over a dataset of n images is CLIP Recall:

$$\text{CLIP Recall} = \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}\left(\mathbf{Y}_i \in \text{Top } 5(\mathbf{I}_i)\right), \quad (6)$$

where Top $5(\mathbf{I}_i)$ represents the set of the top 5 closest caption embeddings to the image embedding $\mathbf{I}_i$, and $\mathbb{1}$ is an indicator function that returns 1 if $\mathbf{Y}_i$ is among the top 5 closest captions to $\mathbf{I}_i$ and 0 otherwise.

A higher CLIP Recall score implies that the captions effectively reflect image content in a way that allows accurate identification, which is particularly useful for tasks requiring captions that are detailed and distinct.

Noun/verb coverage (Chan et al., 2023) is a metric used to evaluate how thoroughly a generated caption describes an image by focusing on the nouns and verbs present in the text. The coverage is determined by comparing the nouns and verbs in

- English: “Describe this image in detail.”
- Japanese: “この画像を英語で詳しく説明してください。”
- Chinese: “请用英语详细描述这张图片。”

Figure 6: The prompts to generate detailed captions, written in English, Japanese, and Chinese. The prompts written in Japanese and Chinese mean “Describe this image in detail in English.”, and are used for the *language discrepancy* evaluation.

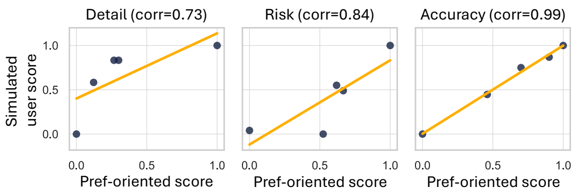

Figure 7: Preference-oriented score vs. simulated user score.

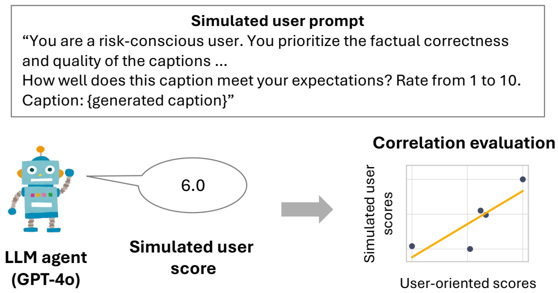

Figure 8: Correlation analysis between preference-oriented scores and user-simulated scores. Full prompts are provided in Appendix A.

the caption with those in reference captions, assessing whether the caption captures essential objects and actions depicted in the image.

Noun coverage counts the nouns in a caption that match exactly with those in the set of reference captions (*i.e.*, we use COCO captions for the reference captions) for the same image. This is done as follows:

$$\text{Noun Coverage} = \frac{1}{\left|\bigcup_{j=1}^{n} N(R_j^i)\right|} \times \sum_{k \in N(C_i)} \mathbb{1}\left(k \in \bigcup_{j=1}^{n} N(R_j^i)\right) \tag{7}$$

where $N(y_i')$ is the set of nouns in the generated caption y_i', and $N(r_j^i)$ represents the set of nouns in the j-th reference caption for image I_i, and $\mathbb{1}$ is an indicator function that returns 1 if the noun k is present in any reference caption’s noun set $\bigcup_{j=1}^{n} N(R_j^i)$, and 0 otherwise.

Verb coverage is calculated similarly, using verbs instead of nouns. The exact match method strictly requires the same words to appear in both the generated and reference captions.

Syntactic complexity (Onoe et al., 2024) measures the structural depth of sentences within the descriptions, specifically by examining the maximum depth of the dependency tree for each sentence (Ohta and Sakai, 2017). The deeper the tree, the more complex the sentence structure. Formally, syntactic complexity can be defined as:

$$\text{Syntactic comp.} = \frac{1}{n}\sum_{i=1}^{n} \tag{8}$$

$$\max(\text{depth of dependency tree})_i \tag{9}$$

Semantic complexity (Onoe et al., 2024) evaluates the richness of content by looking at the number of elements, or nodes, described within a scene graph from y'. The scene graph represents objects and their relationships within the image. A higher number of nodes indicates a more detailed and conceptually rich description. Semantic complexity is expressed as:

$$\text{Semantic comp.} = \frac{1}{n}\sum_{i=1}^{n} \tag{10}$$

$$\text{num. of nodes in scene graph}_i \tag{11}$$

To extract scene graphs from the generated captions, we use the tool in spacy (Spacy, 2024).

FaithScore (Jing et al., 2024) In the context of detailed captioning, FaithScore is used to evaluate

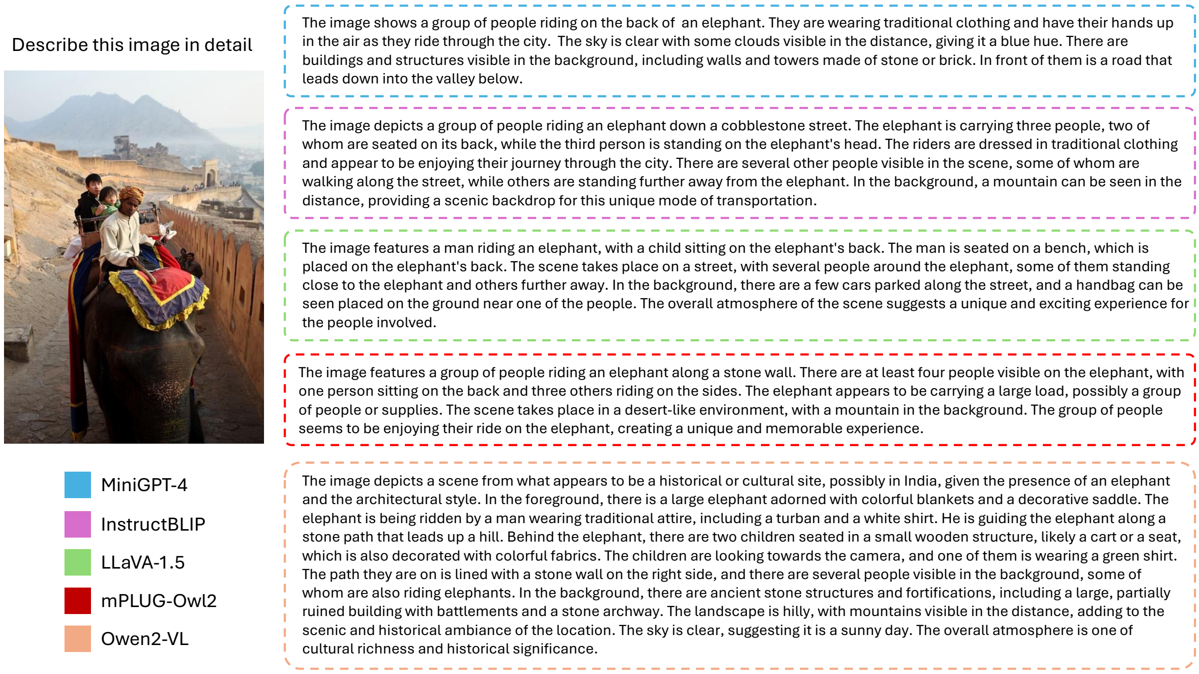

Figure 9: Examples of the generated captions.

how accurately a generated caption y' aligns with the content of an image I. To achieve this, the caption y' is first broken down into atomic facts by a large language model (LLM). The LLM identifies and isolates specific elements such as entities (e.g., objects or people), attributes (descriptive traits), and relationships (interactions or connections between entities). By separating these components, the model produces discrete fact-based units, allowing for a more detailed examination of how each part of the image is represented in the caption.

To evaluate how faithfully a generated caption y' aligns with the visual content of an image I, the caption is first decomposed into atomic facts, denoted as f. Each fact f is then verified against the image I by a verification function V, which utilizes a visual entailment model (VEM). The verification function checks whether each fact is supported by the image. Specifically, the verification function V is defined as:

$$V(f, I) = \begin{cases} 1 & \text{if VEM}(f, I) > 0 \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

(13)

In this formulation, the VEM determines the likelihood that the fact f aligns with the image I. If the entailment score for f in the context of I is greater than 0, the fact is considered supported by the image, and $V(f, I)$ returns 1; otherwise, it returns 0.

The overall FaithScore for the caption y' with K atomic facts is calculated by averaging the verification results for all facts:

$$\text{FaithScore} = \frac{1}{K} \sum_{k=1}^{K} V(f_k, I) \tag{14}$$

where K is the total number of atomic facts in the caption y', and $V(f, I)$ indicates whether each fact is consistent with the image. This metric provides an averaged score reflecting the proportion of facts within y' that are verified to be accurate representations of the content in I. For a dataset with n samples, the overall average FaithScore S can be computed as:

$$S = \frac{1}{n} \sum_{i=1}^{n} \text{FaithScore}_i \tag{15}$$

where FaithScore_i represents the FaithScore for the i-th caption in the dataset. This dataset-level average offers a comprehensive measure of the model's ability to generate captions that faithfully describe images across all samples consistently.

Additionally, we employ a sentence-level FaithScore, which measures the proportion of sentences in y' that are free from hallucinations.

LOTUS Leaderboard: Unified Evaluation of LVLM Captioners

Model	Alignment ↑				Descriptiveness ↑				Complexity ↑			Side effects				
	CLIP-S	CapS_S	CapS_A	N-avg	Recall	Noun	Verb	N-avg	Syn	Sem	N-avg	CH_s ↓	FS ↑	FS_s ↑	Harm ↓	N-avg ↑
MiniGPT-4	60.8	33.0	35.9	0.19	75.3	33.0	34.7	0.22	8.0	32.6	0.38	37.8	55.0	37.6	0.31	0.18
InstructBLIP	59.9	36.0	35.5	0.18	82.1	34.2	34.7	0.40	7.7	46.0	0.41	58.5	62.4	**43.3**	0.10	0.66
LLaVA-1.5	60.1	38.5	**45.0**	0.67	80.5	32.5	31.0	0.11	7.1	39.6	0.08	49.0	**65.7**	41.6	0.12	**0.71**
mPLUG-Owl2	59.7	**39.7**	40.0	0.49	83.3	35.0	32.8	0.34	7.4	45.6	0.28	59.1	62.0	41.3	**0.08**	0.58
Qwen2-VL	**61.8**	37.3	43.2	**0.82**	**90.4**	**45.9**	**36.9**	**1.00**	**8.3**	**75.7**	**1.00**	**26.8**	54.2	41.7	0.28	0.46

Note: All metrics are scaled by 100. Green (bold) indicates the best performance, and blue (underline) indicates the second-best performance for each metric. For CH_s and Harm, lower values are better (↓), while for other metrics, higher values are better (↑).

Figure 10: LOTUS leaderboard for the unified evaluation.

Existence of NSFW words. To estimate the harmfulness of the generated captions, we measure the ratio of captions with NSFW words. Given a function H to check if one or more NSFW words exist in y', we define the harmfulness as follows:

$$\text{Harmfulness} = \frac{1}{n}\sum_{i=1}^{n} H(y'_i) \tag{16}$$

$$H(y') = \begin{cases} 1 & \text{if a NSFW word exists in } y' \\ 0 & \text{otherwise} \end{cases}$$

(17)

(18)

E.1 Evaluation on Hallucination Mitigation Methods

Having established a unified evaluation leaderboard, we use it to assess the impact of hallucination mitigation techniques. Specifically, we analyze the two prominent methods, VCD (Leng et al., 2024) and OPERA (Huang et al., 2024), when applied to LLaVA-1.5 on LOTUS. Both approaches aim to increase the model's reliance on visual evidence when decoding. Tables 4 and 5 show the results of LLaVA-1.5 and its variants with VCD and OPERA applied, driving the following insights:

Mitigating hallucinations in reduced gender bias. The results on hallucination metrics (CH$_s$, FS, FS$_S$ in Table 4) and gender bias in Table 5 demonstrate that applying mitigation methods not only reduces hallucinations but results in gender bias mitigation. In Table 5, applying VCD and OPERA leads to lessening gender disparity (*e.g.*, 10 out of 12 metrics for VCD). A possible hypothesis on this observation is that hallucination mitigation methods, which encourage the model to rely more heavily on visual evidence, may reduce the influence of gender stereotypes present in the training data, leading to decreased gender bias.

Mitigation methods increase the performance disparity among different languages. While reducing gender bias, the results of language discrepancy in Table 5 exhibit performance disparity among the languages is amplified after applying the mitigation methods (*e.g.*, 11 out of 12 metrics worsen for OPERA). This observation may result from the methods' increased reliance on visual evidence and factual accuracy, potentially exposing or exacerbating existing disparities in the model's visual recognition and linguistic representation across different cultures and languages.

LOTUS Leaderboard: Bias-aware Evaluation of LVLM Captioners

Model	Alignment			Descriptiveness			Complexity		Side effects				N-avg↑
	CLIP-S	CapS_S	CapS_A	Recall	Noun	Verb	Syn	Sem	CH_s	FS	FS_S	Harm	
Gender bias													
MiniGPT-4	0.3	0.9	1.1	7.8	1.7	**2.6**	6.3	3.2	4.8	6.3	4.0	1.64	0.51
InstructBLIP	0.8	2.7	1.2	8.4	1.9	3.3	**1.0**	0.1	6.8	3.8	5.0	0.72	0.40
LLaVA-1.5	0.7	2.2	0.7	9.5	2.2	4.1	1.5	0.2	7.6	3.8	**3.7**	0.39	0.46
mPLUG-Owl2	0.6	2.2	1.2	9.1	2.3	3.5	1.6	**0.0**	7.2	3.1	5.8	**0.33**	0.40
Qwen2-VL	**0.2**	**0.7**	**0.5**	**6.3**	**0.1**	3.6	13.5	2.5	**4.4**	**0.9**	5.7	1.77	**0.63**
Skin tone bias													
MiniGPT-4	0.8	1.5	0.8	4.8	**0.2**	2.3	19.4	0.2	2.0	0.9	0.5	0.09	0.55
InstructBLIP	0.5	1.4	**0.2**	8.4	1.9	1.1	6.8	**0.1**	4.0	2.4	1.1	0.09	0.51
LLaVA-1.5	0.4	1.3	0.7	4.0	**0.2**	**1.0**	**5.3**	0.6	2.7	1.4	1.3	0.18	**0.67**
mPLUG-Owl2	0.6	1.9	0.5	5.1	0.8	2.2	7.6	0.4	**1.7**	**0.1**	**0.4**	**0.00**	**0.67**
Qwen2-VL	**0.2**	**1.1**	1.5	**2.3**	0.5	1.3	14.9	2.3	2.7	3.1	1.8	0.09	0.50
Language discrepancy													
MiniGPT-4	0.8	1.5	3.9	2.3	4.3	5.2	52.2	5.0	5.4	5.6	3.4	0.10	0.40
InstructBLIP	-	-	-	-	-	-	-	-	-	-	-	-	-
LLaVA-1.5	0.4	**0.8**	**2.0**	**1.1**	**1.1**	**1.8**	**11.4**	**1.8**	**4.7**	2.0	1.6	0.06	**0.95**
mPLUG-Owl2	1.4	1.6	4.9	1.5	**1.1**	3.7	37.5	8.4	17.0	6.3	**1.3**	**0.02**	0.57
Qwen2-VL	**0.2**	3.6	6.7	1.9	3.9	3.8	90.8	26.2	6.4	7.5	2.1	0.14	0.28

Note: All metrics are scaled by 100. Green (bold) indicates the best performance, and blue (underline) indicates the second-best performance for each metric. Language discrepancy evaluation is not applicable to InstructBLIP due to a lack of Japanese support.

Figure 11: LOTUS leaderboard for bias-aware evaluation.

Table 4: Unified evaluation of hallucination mitigation methods on LOTUS. All metrics are scaled by 100.

Model	Alignment ↑			Descriptiveness ↑			Complexity ↑		Side effect			
	CLIP-S	CapS_S	CapS_A	Recall	Noun	Verb	Syn	Sem	CH_s↓	FS ↑	FS_S ↑	Harm ↓
LLaVA-1.5	60.8	38.5	45.0	80.5	32.5	31.0	7.1	39.6	49.0	65.7	41.6	0.12
+ VCD	60.1	36.3	41.8	82.4	32.7	28.8	7.5	43.0	48.4	64.8	42.4	0.08
+ OPERA	60.6	37.3	44.2	82.9	33.2	30.9	7.3	40.6	47.7	66.1	42.6	0.12

Table 5: Bias-aware evaluation of hallucination mitigation methods on LOTUS. All metrics are scaled by 100.

Model	Alignment			Descriptiveness			Complexity		Side effect			
	CLIP-S	CapS_S	CapS_A	Recall	Noun	Verb	Syn	Sem	CH_s	FS	FS_S	Harm
Gender bias												
LLaVA-1.5	0.7	2.2	0.7	9.5	2.2	4.1	1.5	0.2	7.6	3.8	3.7	0.39
+ VCD	0.6	1.1	0.2	9.0	2.0	3.1	6.2	0.1	4.6	4.3	3.2	0.33
+ OPERA	0.6	2.8	0.2	8.1	2.0	0.9	8.5	0.3	7.2	2.9	3.5	0.54
Skin tone bias												
LLaVA-1.5	0.4	1.3	0.7	4.0	0.2	1.0	5.3	0.6	2.7	1.4	1.3	0.18
+ VCD	0.6	0.6	0.6	5.7	0.3	1.2	6.3	1.1	1.2	1.2	2.1	0.27
+ OPERA	0.3	0.2	0.1	3.8	0.2	0.6	20.9	0.7	0.0	0.1	1.3	0.00
Language discrepancy												
LLaVA-1.5	0.4	0.8	2.0	1.1	1.1	1.8	11.4	1.8	4.7	2.0	1.6	0.06
+ VCD	0.6	1.1	4.0	2.7	1.5	1.9	21.2	4.7	4.0	1.5	2.3	0.10
+ OPERA	0.8	2.2	4.7	2.1	1.5	2.9	23.6	5.1	11.7	2.9	3.9	0.02