

BEAF: Observing BEfore-AFter Changes to Evaluate Hallucination in Vision-language Models

Moon Ye-Bin^{1*}  Nam Hyeon-Woo^{1*}  Wonseok Choi²  Tae-Hyun Oh^{1,2,3} 

¹Dept. of EE, POSTECH, Korea ²Grad. School of AI, POSTECH, Korea

³Institute for Convergence Research and Education in Advanced Technology, Yonsei University, Korea.

{ybmooon, hyeonw.nam, wonseok.c, taehyun}@postech.ac.kr

Abstract. Vision language models (VLMs) perceive the world through a combination of a visual encoder and a large language model (LLM). The visual encoder, pre-trained on large-scale vision-text datasets, provides zero-shot generalization to visual data, and the LLM endows its high reasoning ability to VLMs. It leads VLMs to achieve high performance on wide benchmarks without fine-tuning, exhibiting zero or few-shot capability. However, recent studies show that VLMs are vulnerable to hallucination. This undesirable behavior degrades reliability and credibility, thereby making users unable to fully trust the output from VLMs. To enhance trustworthiness and better tackle the hallucination of VLMs, we curate a new evaluation dataset, called the BEfore-AFter hallucination dataset (BEAF), and introduce new metrics: True Understanding (TU), IGnorance (IG), StuBbornness (SB), and InDecision (ID). Unlike prior works that focus only on constructing questions and answers, the key idea of our benchmark is to manipulate visual scene information by image editing models and to design the metrics based on scene changes. This allows us to clearly assess whether VLMs correctly understand a given scene by observing the ability to perceive changes. We also visualize image-wise object relationship by virtue of our two-axis view: vision and text. Upon evaluating VLMs with our dataset, we observed that our metrics reveal different aspects of VLM hallucination that have not been reported before. Project page: <https://beafbench.github.io/>

1 Introduction

Vision language models (VLMs) [1, 5, 8, 23, 24] have recently emerged by its multi-modal reasoning capability. It is indeed promising research toward making intelligent agents perceive the world. VLMs comprise a vision encoder as a visual perception module and a large language model (LLM) as a reasoning module. Recent LLMs [6, 32, 33] have exhibited powerful reasoning capabilities across various tasks, including commonsense reasoning [3, 16] and math [7], in zero-/few-shot ways. This success motivates researchers to integrate LLMs with vision encoders to open LLMs’ eyes to see the world. Leveraging the reasoning

* Authors contributed equally to this work.

capabilities of LLMs alongside visual perception enables VLMs to excel in various visual tasks, such as image captioning [2, 38] and visual question answering [11, 15] without task-specific fine-tuning.

Despite this encouraging demonstration, VLMs are susceptible to the hallucination [12, 14, 19, 34, 39, 42] wherein their outputs do not reflect facts presented in input images. For example, although a given input image does not present an object, VLMs often incorrectly answer for the presence of the object. This tendency is problematic, akin to spreading falsehoods in human communication. Hallucination prevents genuine communication, consequently diminishing the reliability of interactions between agents or humans. No matter how good VLM’s answer is, if hallucination is included in the answer, it can undermine trust in the system’s reliability and credibility. Therefore, it is crucial to address hallucinations to improve the system’s responsibility and trustworthiness.

Understanding the impact of hallucinations on model performance requires the development of an assessment framework, which requires a deep exploration of the underlying causes. The prior works [14, 41, 42] have been proposed to unveil the causation of hallucinations in terms of data and models by reporting their performance on existing hallucination benchmarks. As a widely used hallucination benchmark, POPE [19] follows the question-and-answer (QnA) format. For the convenience of measuring, POPE uses a discriminative question, *e.g.*, “Is there {object} in the image?”, and gets a YES or NO answer. The following hallucination benchmarks [14, 19, 34] commonly leverage LLMs or human annotators to generate question prompts with available scene information such as bounding boxes and captions. This text-axis evaluation may assess certain aspects of the hallucination performance well, but it is hard to disentangle sources of hallucination, particularly in VLMs that handle both visual and text modalities. Thus, to thoroughly analyze and understand the behavior of VLMs, it is both natural and necessary to consider both modalities.

In this work, we propose a BEfore-AFTer (BEAF) hallucination evaluation benchmark. This benchmark comprises a dataset enhanced along the vision-axis and metrics that are aware of changes. We manipulate scenes by selectively removing objects on the vision-axis while augmenting QnA pairs along the text-axis. We can observe how the models’ answers change as the image is manipulated for the same question. For example, as illustrated in Fig. 1, if we remove an apple from the given image and then ask “Is there an apple?”, the answer from the model should change from “Yes” to “No”, between the original and modified images, if it truly relies on what it sees. Therefore, for a more thorough evaluation and analysis, we need change-aware metrics beyond just using standard accuracy [14, 19, 34].

We propose four new metrics: True Understanding (TU), IGNorance (IG), StUBbornness (SB), and InDecision (ID), which are enabled by our vision and text augmentation. TU measures whether the model answers correctly before and after the manipulation. IG evaluates whether the models perceive the images consistently both before and after manipulation. SB is related to the phenomenon of giving the same answer instead of correctly recognizing the image. ID measures

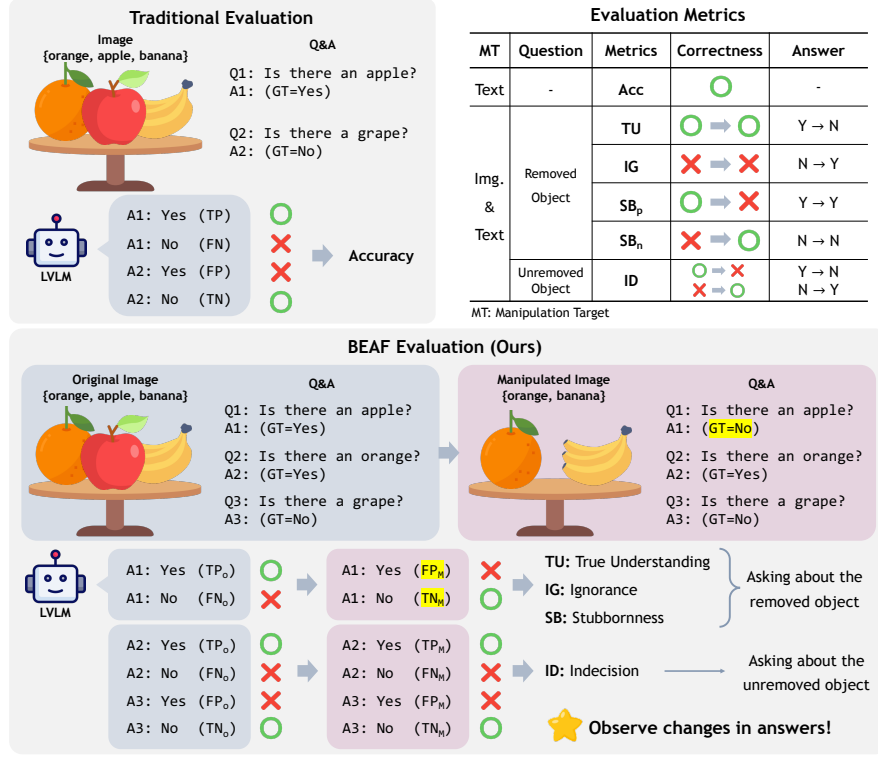


Fig. 1: Illustration of Before-After (BEAF) benchmark. We present a comparison between traditional evaluation benchmarks and our BEAF benchmark for assessing the hallucination behavior in VLMs. Traditional evaluation methods solely manipulate questions based on the existence of an object and measure accuracy or F1 score. In contrast, our BEAF benchmark not only constructs questions but also manipulates images and tracks changes in answers as the images undergo manipulation. The BEAF benchmark introduces novel metrics including True Understanding (TU), IGnorance (IG), StuBbornness (SB), and InDecision (ID), which consider the changes in answers.

cases where the answer changes even though the target object does not change or does not exist in the image. The first three metrics address questions regarding target objects that have been removed, while the last one belongs to questions about other objects not related to the targets.

Based on our BEAF dataset and change-aware metrics, we can evaluate and understand the models, even the parts we did not know through the previous methods. The evaluation results imply that answers assumed as non-hallucination from the existing hallucination benchmarks could be a hallucination. We also analyze the influence of each object in the scene based on the variation along the vision and text axes. In addition, we include a discussion on open-generation answers and CLIPScore results. Our key contributions are as follows:

- We propose a benchmark, called BEfore-AFter (BEAF), to evaluate the hallucination of large vision-language models. Our benchmark comprises a dataset that considers both text- and vision-axes, and change-aware metrics, which allow us to perform granular evaluation.
- We demonstrate that previously reported outcomes deemed non-hallucinatory in the prior text-axis-only evaluation benchmarks may be hallucinations.
- We also visualize the impact of relationships between objects within hallucinatory images based on the BEAF results.

2 Related Work

Vision Language Models (VLMs). VLMs [1, 4, 5, 8, 10, 23–25, 40, 43] have demonstrated their impressive image describing and reasoning ability based on powerful large language models (LLMs) [6, 31–33]. Generally, VLMs consist of three parts: vision encoder, LLM, and vision-language alignment module. The vision information extracted from the vision encoder is translated to the LLM-understandable token by passing through the alignment module. LLM takes a prompt text and visual tokens as input and outputs response text tokens according to a given image and prompt. Their training process is divided into multiple steps depending on which parts of VLMs are trained.

Hallucination in VLMs. Hallucination [12, 19, 22, 34, 35, 39] in VLMs refers to the case where the generated text answer does not reflect the true contents of the provided images but rather relies on the internal knowledge of the models. This behavior might cause harmful consequences in high-risk applications such as the medical domain or auto-driving scenarios. This undesirable behavior could be caused by various factors, *e.g.*, data imbalance, image resolution, and model capacity. The assessment of hallucination is also related to the true understanding of VLMs. If VLMs respond to the given input correctly rather than relying on innate knowledge, VLMs are considered to have a perfect understanding of the perceived world. Numerous methodologies [12, 21, 42] have been developed to tackle the issue of hallucination in VLMs by addressing data biases [14, 21], optimization techniques [12, 18, 29], and post-processing methodologies [9, 37, 42].

Evaluating Hallucination in VLMs. Evaluating hallucination performance is essential to guide the research of VLM. Previous studies have introduced hallucination benchmarks to understand and assess these undesirable phenomena of VLMs. Motivated by traditional visual and question benchmarks, existing hallucination benchmarks [14, 19, 34] follow this VQA style that provides a simple QnA evaluation pipeline and intuitive metrics. POPE [19] and CIEM [14] curate images and QnA pairs with metrics such as accuracy, precision, and recall, where they use a specific question type inducing multi-choice or yes/no answers, called discriminative type. AMBER [34] includes an additional question type that induces open-ended caption generation as answers, called generative type. For such open-generation assessment, they extract the objects from the output text answers and compare them with the ground-truth objects in the image. These prior studies focus on manipulating texts of questions and answers only,

Table 1: Data statistics. Our BEAF dataset contains 26K image-question pairs, consisting of the original and manipulated ones. On average, an original image is associated with 3.45 manipulated images and 11.72 questions.

	Image	Img-Q Pair	Yes	No
Original	500	5,653	2,026	3,627
Manipulated	1,727	20,465	6,267	14,198
Total	2,227	26,118	8,293	17,825




Fig. 2: Ratio of removed objects from image and object in question. For convenience, we report the ratio of upper-class categories instead of the object class itself. The total number of removed objects is the same as the number of manipulated images, and one object in question is the same as the number of image-question pairs.

which we call text-axis evaluation. However, if some objects always co-occur with another object, we cannot judge whether VLMs truly understand scene information accurately. Also, VLMs take multi-modal inputs, both images and text prompts; thus, evaluation considering only the text-axis would be insufficient as a hallucination assessment of VLMs. In this work, we manipulate the scene information and questions together, *i.e.*, vision-/text-axes; it enables the fine-grained analysis of hallucination. We call it vision-/text-axes evaluation.

3 BEAF Evaluation Benchmark for Hallucination

We explore another axis, *i.e.*, vision-axis, to make up for the deficiencies in the text-axis evaluation. Our key idea is to manipulate a visual scene by removing and inpainting objects within an image. This allows us to focus on changes and to check whether the VLMs are correctly aware of it. We describe how to curate our dataset in Sec. 3.1, and then elaborate our new metrics in Sec. 3.2

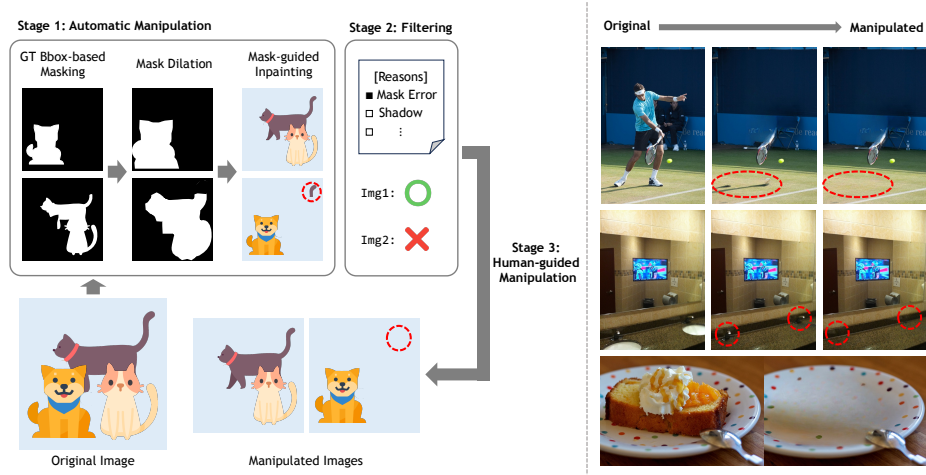


Fig. 3: Image manipulation pipeline. We illustrate the image manipulation pipeline comprised of three stages. In **Stage 1**, we automatically remove target objects sharing the same semantic class from the given images. **Stage 2** is to filter the automatically manipulated results based on the predefined rules, such as mask errors, remained shadows, and low-quality outcomes. Undesirable manipulations are either corrected or discarded during this stage. Finally, in **Stage 3**, human annotators engage in human-guided manipulation of the filtered images to achieve high-quality results.

3.1 Evaluation Dataset

Dataset Overview. Table 1 and Fig. 2 summarize our data statistics. The dataset consists of images and their associated QnA. The image type is divided into original and manipulated. Five hundred original images are sampled from the validation set of the MS-COCO dataset [20], which shares the same split with POPE [19], and their associated 1,727 manipulated images are generated by removing objects from the original images. On average, 3.45 manipulated images are created from the original images. The main question format is “Is there {object} in this image?” and {object} can be a positive object within the image or a negative object not in the image; the positive object is one of the objects in the image, and the negative object is the one not in the image. We adopt the positive and negative objects from the POPE for the original image and adapt the question for the manipulated image. If we remove the positive object from the original image to make the manipulated one, the positive object becomes the negative object in the manipulated one. The total dataset comprised 2,227 images and 26,118 image-question pairs. The “Yes” and “No” ground truth answers are 8,293 (31.75%) and 17,825 (68.25%), respectively.

Image Dataset Construction Pipeline. We manipulate the original images by removing objects within the image. To collect the manipulated images, we pass through a three-stage pipeline (see Fig. 3): the first stage involves auto-

matic manipulation, followed by a filtering stage, and finally, a human-guided manipulation stage. We offer a detailed explanation of each stage as follows.

1) *First Stage: Automatic Manipulation.* For automatic object removal, we use an off-the-shelf inpainting model, LaMa [30], with mask guidance. While the initial option is to use the ground truth masks provided within the original dataset, we opt not to use them because of the low quality ground truth masks in MS-COCO. For example, polygon masks often cannot cover all areas of the corresponding object and can be propagated to the inpainting quality. Therefore, we employ a pre-trained mask extractor called SAM [17], utilizing bounding box annotations. Subsequently, we dilate the extracted mask to have a margin for more effective object removal, and remove the object with the inpainting model given the dilated mask guidance. We empirically choose to use LaMa over diffusion models, *e.g.*, [28, 44], as we find that this specialized model performs better in the context of object removal.

2) *Second Stage: Filtering.* Although we remove objects from the original images, achieving consistent quality in the outcome can be influenced by a myriad of factors. These factors include, but are not limited to, the size of the objects being removed and the complexity of the backgrounds they inhabit. Additionally, the presence of shadows remains unaddressed, as they are not effectively masked using SAM; the ground truth masks also do not capture the shadow. As shown on the right side of Fig. 3, it displays undesirable images due to the presence of shadow or a faucet. The human annotators classified the results obtained from the previous stage into usable or unusable, where only the usable results are further manipulated. The detailed filtering rules can be found in the supplementary material. We directly use the usable results and hand over the results that need further manipulation to the next stage.

3) *Third Stage: Human-guided Manipulation.* The third stage of our workflow relies on human annotators to intervene and refine the images obtained from the second stage. As mentioned above, these filtered images often exhibit undesirable artifacts such as ghost shadows, afterimages, and fragmented objects. Even such a small artifact may provide strong cues indicating the past presence of objects that both humans and machines can deduce, *e.g.*, a cat’s tail suggests the past presence of the cat before the object removal. Thus, human annotators are asked to meticulously manipulate the images, either correcting or enhancing them. Annotators effectively refine visual data, aligning it more closely with real natural images in human perception.

As shown in Fig. 3, our annotation process results in images that are free of the aforementioned artifacts. This manual intervention is crucial, as it ensures the fidelity and reliability of our benchmark dataset.

3.2 Evaluation Metrics

The prior studies [14, 19, 34] suggest using traditional metrics such as accuracy, precision, recall, and F1 score, which we can also use the metrics in our dataset. In addition, we propose four new metrics: True Understanding (TU), IGnorance

(IG), StuBbornness (SB), and InDecision (ID), for more detailed evaluation by exploiting the distinctive configuration of before-/after-changes in our dataset.

Let $\mathcal{S} = \{(I_o, I_m, Q, A_o, A_m, R)_i\}_{i=1}^{|\mathcal{S}|}$ denote our dataset, where each tuple consists of an original image I_o , a version of its manipulated image I_m , their associated question Q , their respective corresponding answers A_o and A_m , and a True/False flag R that indicates whether the question is associated with the object being removed in the manipulated image I_m . For compactness of our metric definitions, we define $\text{Filter}(\cdot)$ to extract the subset of \mathcal{S} satisfying input conditions as:

$$\text{Filter}(b_o, b_m, b_r) = \{t \mid \text{IsCorrect}(A_o) = b_o, \text{IsCorrect}(A_m) = b_m, R = b_r, t \in \mathcal{S}\}$$

where the tuple $t = (I_o, I_m, Q, A_o, A_m, R)$. For example, $\text{Filter}(\text{True}, \text{False}, \text{True})$ is the subset of \mathcal{S} of which the element satisfy the conditions that are 1) the answers corresponding to respective original and manipulated images are correct and wrong, and 2) the question is the one related to the object removed in the manipulated images.

True Understanding. True Understanding (TU) measures whether the models truly understand the scene. It evaluates if the model can accurately answer to the question about the removed object. As shown in Fig. 1, if the apple is removed, a model that truly understands the change should switch its answer from “Yes” to “No” to the question “Is there an apple?”. We define the TU metric as follows:

$$\text{TU} = \frac{|\text{Filter}(\text{True}, \text{True}, \text{True})|}{|\text{Filter}(\text{True} \vee \text{False}, \text{True} \vee \text{False}, \text{True})|} \cdot 100. \quad (1)$$

Compared to the traditional accuracy, our metric specifically considers the correctness of answers before and after manipulation. Therefore, it provides a conditional measure of accuracy, taking into account the model’s response to changes within the scene. A higher value of this metric indicates a more accurate understanding of the scene by the model, as it demonstrates the model’s ability to adapt its answers appropriately to the scene changes.

Ignorance. The second metric is IGnorance (IG), which measures the extent to which models lack knowledge about specific scene information. This metric quantifies ignorance by identifying instances where models provide incorrect answers to the question about the removed objects. As shown in Fig. 1, ignorance is evident if models incorrectly answer “No” to the presence of an apple before its removal and incorrectly answer “Yes” after the apple is removed, indicating a failure to recognize it. IG is defined as follows:

$$\text{IG} = \frac{|\text{Filter}(\text{False}, \text{False}, \text{True})|}{|\text{Filter}(\text{True} \vee \text{False}, \text{True} \vee \text{False}, \text{True})|} \cdot 100. \quad (2)$$

The lower the metric, the less ignorant the model is.

Stubbornness. What we observe is that VLMs tend to output the same answer repeatedly. It means that VLMs are biased toward the specific answer rather than arbitrary one when they do not know. This repetitive behavior is undesirable; thus, we introduce the metric of StuBbornness (SB), which measures the extent

to which models adhere to their initial answers. Furthermore, we categorize SB into SB_p and SB_n where the subscripts of p and n correspond to the consistent positive (“Yes”) and negative (“No”) answers, respectively. We define them as follows:

$$SB_p = \frac{100 \cdot |\text{Filter}(\text{True}, \text{False}, \text{True})|}{|\text{Filter}(\text{True} \vee \text{False}, \text{True} \vee \text{False}, \text{True})|}, \quad SB_n = \frac{100 \cdot |\text{Filter}(\text{False}, \text{True}, \text{True})|}{|\text{Filter}(\text{True} \vee \text{False}, \text{True} \vee \text{False}, \text{True})|}. \quad (3)$$

SB is the summation of SB_p and SB_n and can be computed from $100 - \text{TU} - \text{IG}$. The lower the metric, the less ignorant the model is.

Indecision. Unlike the previous three metrics that focus on the removed objects in the manipulated images, InDecision (ID) focuses on the answers to the questions that are not relevant to the removed objects. These answers should not be changed even after the manipulation. Nevertheless, models, often change their answers about unmanipulated objects, which would imply that the model predicts a random answer without grounding it on the image. We track these behaviors with the ID metric defined as follows:

$$\text{ID} = \frac{|\text{Filter}(\text{True}, \text{False}, \text{False})| + |\text{Filter}(\text{False}, \text{True}, \text{False})|}{|\text{Filter}(\text{True} \vee \text{False}, \text{True} \vee \text{False}, \text{False})|} \cdot 100. \quad (4)$$

The lower the metric, the less ignorant the model is. Additionally, we combine TU and ID using a harmonic mean to make it easier to compare on our benchmark. We refer to it as F1-Score, which is computed as follows: $F1 = \frac{2}{\text{TU}^{-1} + (100 - \text{ID})^{-1}}$.

4 Experiment

In Sec. 4.1, we evaluate the current VLMs using our BEAF dataset. We employ both proposed and traditional metrics and compare their performance with other benchmarks. Furthermore, we conduct an analysis of image-wise dynamic correctness by visualizing image-wise object relationship containing vision and text axes. We provide a report on the object-wise error rate. In Sec. 4.2, we assess the performance of VLMs on the open generation task using our dataset, while also analyzing both the dataset and VLMs based on the CLIP score.

All the reported performance is the zero-shot performance of VLMs, including LLaVA [23], InstructBLIP [8], Shikra [5], and mPLUG-Owl [36]. We use prompt templates provided in their released codebases or papers.

4.1 Main Results

Evaluation on BEAF benchmark. We evaluate several VLMs on our BEAF dataset using the proposed change-aware metrics, TU, IG, SB, ID, and F1. The evaluation results are presented in Table 2. Among the 13B-size models, LLaVA outperforms InstructBLIP in terms of TU, IG, and SB because InstructBLIP consistently answers “Yes” regardless of scene manipulation, and this result is related to its high value of SB_p . InstructBLIP achieves better ID compared to LLaVA. Among 7B-size models, Shikra exhibits the best in TU and SB_p . Shikra

Table 2: Evaluation results on BEAF benchmark. We evaluate various VLMs on our benchmark. F1 is a harmonic mean of TU and 1-ID. Yes ratio is the percentage of total answers that are answered “Yes” for the removed object-related question. The former is for the original image, and the latter stands for the manipulated image. Underline stands for the best.

Size	Model	TU(↑)	IG(↓)	SB _p (↓)	SB _n (↓)	ID(↓)	F1(↑)	Yes (%)
13B	LLaVA-v1.5	<u>24.3</u>	<u>0.2</u>	<u>72.0</u>	<u>3.5</u>	6.4	<u>38.6</u>	96.4 → 72.2
	InstructBLIP	11.1	0.3	84.4	4.2	<u>6.1</u>	19.9	95.5 → 84.7
7B	LLaVA-v1.5	32.6	0.1	61.3	6.0	5.6	48.5	93.9 → 61.4
	InstructBLIP	33.5	0.6	51.3	14.6	<u>5.2</u>	49.5	84.8 → 51.9
	Shikra	<u>52.7</u>	0.4	<u>31.2</u>	15.7	5.7	<u>67.6</u>	83.9 → 31.6
	mPLUG-Owl2	24.6	<u>0.0</u>	72.1	<u>3.3</u>	7.0	38.8	96.7 → 72.1

Table 3: Traditional evaluation results on BEAF dataset. We evaluate various VLMs on our dataset in terms of accuracy, precision, recall, and F1 score. Note that the BEAF dataset includes both original and manipulated data but does not consider the changes for this traditional evaluation. Underline stands for the best.

Size	Model	Acc.(↑)	Precision(↑)	Recall(↑)	F1(↑)	Yes (%)
13B	LLaVA-v1.5	<u>74.9</u>	56.3	93.4	<u>70.3</u>	52.6
	InstructBLIP	71.5	52.8	<u>94.3</u>	67.7	56.7
7B	LLaVA-v1.5	79.7	62.6	89.5	73.7	45.4
	InstructBLIP	82.0	68.3	80.7	74.0	37.5
	Shikra	<u>84.5</u>	<u>74.9</u>	76.9	<u>75.9</u>	32.6
	mPLUG-Owl2	69.1	50.7	<u>93.8</u>	65.8	58.7

is trained using a location-aware strategy, which could be beneficial for hallucination evaluation because knowledge of location aids in judging the existence of an object. The overall results indicate that, although it is rare to get all the answers wrong, it is common for the answers to remain the same even though the image has changed. This indicates that recent VLMs are not adequately reflective of such changes.

Traditional evaluation on BEAF dataset. In addition to the proposed change-aware metrics, we also evaluate using traditional metrics such as accuracy, precision, recall, and F1 score, as shown in Table 3. While the trends in accuracy or F1 score are similar to those shown in TU results of Table 2, the absolute values of TU are much lower than those of the traditional metrics. This discrepancy suggests that, although a model achieves high performance in the existing benchmarks, our benchmark tells us that we cannot conclude low hallucination and the models may still suffer from hallucination. For example, although LLaVA-v1.5-13B achieves high performance on accuracy, the Yes ratio in Table 2 remains high even after the manipulation. It may include instances where the correct answer was guessed without actual reference to the given image. The experimental findings underscore the importance of observing changes in answers along the vision-axis to accurately evaluate hallucination in VLMs.

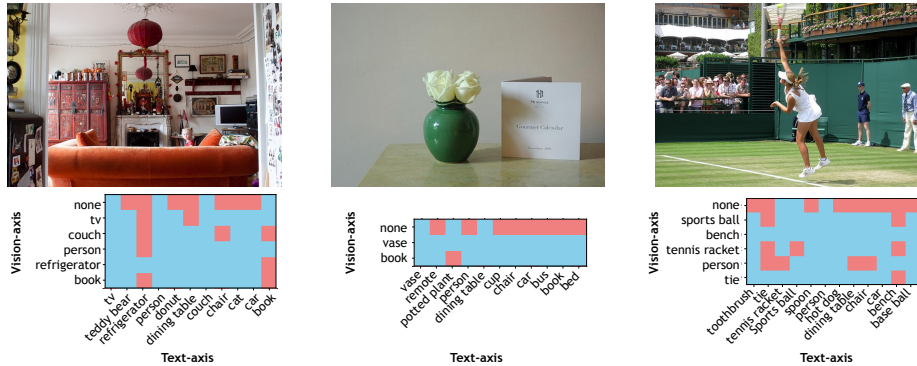


Fig. 4: Visualization of image-wise object relationship. We visualize image-wise object relationship along the text and vision axes from Shikra. [Top] Original image samples. [Bottom] The object relation table given the manipulated images and questions. We color the correct answer blue and the wrong one red. The text axis stands for the target object queried in the question, and the vision axis for the removed object in the image. The “none” in the vision axis means the original image (not manipulated). Thereby, we can analyze influence between objects within a scene at once.

Comparison with other benchmarks.

We compare our metric TU with the performance for other benchmarks of VQAv2 [11] and GQA [15], as shown in Table 4. Each VQA performance value is reported from the proposed papers. The results indicate that, while mPLUG-Owl2 outperforms others on VQAv2 and LLaVA excels on GQA, Shikra leads in performance on our metric. We observe that there is a negative correlation between performance and the incidence of hallucination, which aligns with the previous observation [35]. Although improving performance on various benchmarks is challenging, it is essential to check multiple aspects of the performance, including hallucination. We believe that our benchmark provides valuable insights in conjunction with existing benchmarks.

Image-wise object relationship. We visualize changes in dynamic object relationship along the vision and text axes, as shown in Fig. 4. Compared to our metrics providing comprehensive statistics as summarized scores, this visualization enables to track the changes in VLMs’ responses at object-wise and scene-wise levels. The x-axis is the text axis where each object is associated with the question, “Is there {object}?”; the y-axis is the vision axis where the object is removed from the original image. The blue color is correct, the red one is wrong.

In the first column in Fig. 4, altering the scene affects the textual correctness. Specifically, removing either the couch or the book shifts the book’s textual cor-

Table 4: Comparison of performance on VQA and TU. We compare the accuracy on VQA datasets, VQAv2 and GQA, as well as TU on BEAF dataset.

Model (7B)	VQAv2	GQA	TU
LLaVA-v1.5	78.5	62.0	32.6
InstructBLIP	-	49.2	33.5
Shikra	77.4	-	52.7
mPLUG-Owl2	79.4	56.1	24.6




Fig. 5: Object-wise error rate on the manipulated image-question pairs. We plot the error rate of each object to investigate which object is inaccurately inferred when it is removed. We exclude the objects manipulated less than 20 times. The solid line represents the average error rate of VLMs (7B), while the purple area indicates the 95% confidence interval.

rectness from correct to incorrect. This suggests a correlation between the couch and the book, indicating that complex scenes can lead to model hallucinations even in the absence of the book. In the second scene, when the object vase or book is removed, most of the incorrectness becomes correct. This could imply that the presence of a vase or a book might lead the model to misinterpret the scene. Through the simultaneous analysis of both vision and text axes, we observe immediate changes in the answers and offer interpretations for the dynamic interactions observed in the second and third scenes illustrated in the figure.

Object-wise error rate. We compute the mean and 95% confidence interval of the object-wise error after the manipulation in Fig. 5. It represents whether the VLMs can accurately detect the removal. The results show that VLMs frequently fail to recognize certain objects like traffic lights, vases, and remotes. We speculate that the difficulty in recognizing these objects may stem from their complexity or their infrequent appearance in the training data.

4.2 Additional Study

Evaluation of open generation. In Table 5, we compare the results of the CHAIR [27] and TU metrics to evaluate the hallucination in the open-ended generation answers on our dataset. Briefly, CHAIR_I and CHAIR_S count the hallucinated instances and sentences containing these objects in the generated output, respectively. The lower the score, the better. We observe that if the open generation contains the hallucinated object and sentences, TU is also lower, as seen with InstructBLIP-13B.

Table 5: Comparison of CHAIR metric and TU. We measure the CHAIR score based on the answers derived from the images in BEAF using the prompt “Describe this image.”

Size	Model	CHAIR_I	CHAIR_S	TU
13B	LLaVA-v1.5	21.6	72.3	24.3
	InstructBLIP	28.1	81.1	11.1
7B	LLaVA-v1.5	22.3	73.9	32.6
	InstructBLIP	21.0	67.5	33.5
	Shikra	23.6	77.0	52.7
	mPLUG-Owl2	25.1	78.4	24.6

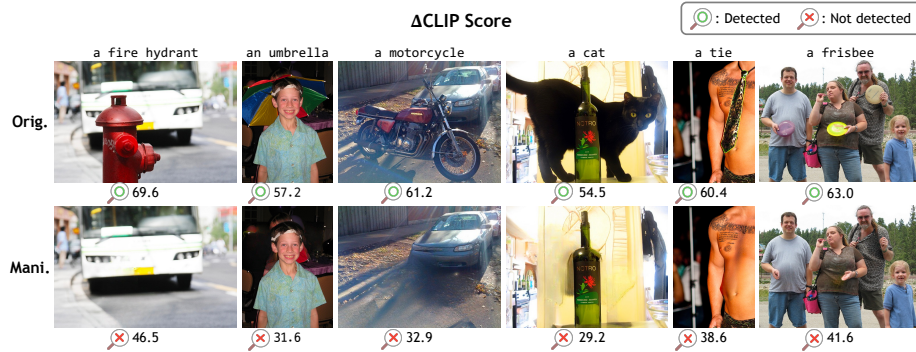


Fig. 6: Samples with CLIP Score. We visualize samples and display the CLIP Score below. Additionally, we report the results of Faster R-CNN [26] using the magnifying glass icon, the removed object is not detected, which demonstrates the manipulation quality.

CLIP Score. We compare the CLIP Score [13] between original and manipulated images to ensure the target object has been well-removed. A lot of VLMs utilize the visual encoder from CLIP, prompting us to analyze CLIP’s performance with our datasets. We measure the CLIP score between an object and a prompt “a photo of {object}.” We use a cropped object image based on the GT bounding box. If the image has multiple objects of the same semantic class, we use the average CLIP score of all the same semantic objects. For the manipulated image, {object} is the removed object. As shown in Table 6, the manipulated images have lower CLIP scores compared to their original counterparts.

We also visualize the images with the CLIP Score and the multi-class classification results in Fig. 6. The multi-class classification correctly identifies the depicted objects. In addition, Fig. 7 visualizes an open generation sample result, where red-colored text indicates hallucinated words. Before manipulation, all the VLMs mention “people” in the background; however, they still mention “people” even though all the persons are removed from the image. This implies that the commonly used visual encoder might be a non-negligible source of hallucination.

Table 6: CLIP Score. We measure the CLIP Score of original and manipulated images with the text prompt.

Type	CLIP Score
Original	53.7
Manipulated	47.9

5 Discussion and Conclusion

We present a comprehensive analysis of hallucination using our BEAF benchmark. Unlike conventional hallucination benchmarks that measure accuracy by changing only the text-axis question, we also manipulate images by removing objects and propose four sophisticated metrics to track the changes in answers

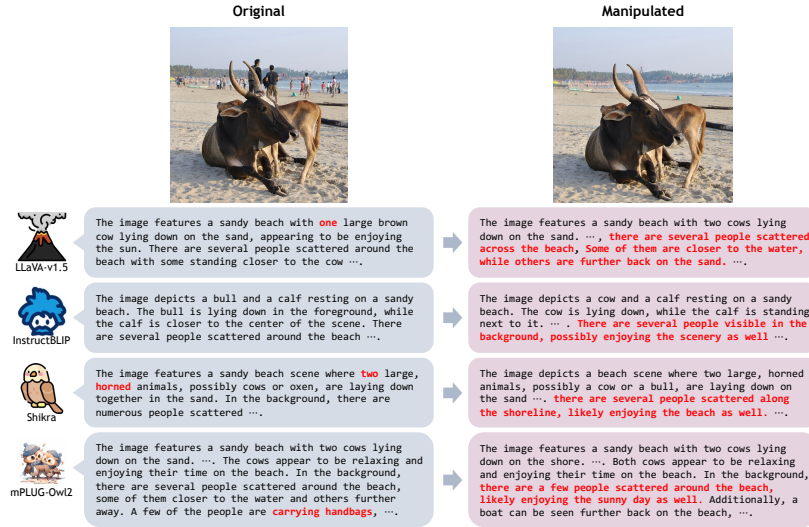


Fig. 7: Qualitative example on open-ended generation. We report the changes in answers for the open-ended generation task with the prompt “Describe this image.” as the image undergoes manipulation. The red color highlights the hallucination parts.

according to image changes. This approach enables us to a more discerning analysis of the hallucination in two-axis views, *e.g.*, vision and text axes. We can find with our metric that traditional accuracy may include instances where the correct answer was guessed without actual reference to the given image. We found that our TU is a more *calibrated* metric, which allows us to observe the intriguing phenomenon that has not been revealed in the prior work: VLMs tend to be confused if the complexity of an image is high. (refer to Sec. A in the supplementary material). Additionally, we introduce the visualization method to track the correctness as the objects are removed. We believe that our benchmark presents a comprehensive analysis.

Limitation. While our benchmark offers a more fine-grained evaluation, it still faces the constraints of the MS-COCO dataset. These include a limited variety of object categories and an uneven distribution of object instances. While we utilized current state-of-the-art in-painting models and mask extractors, we manually adjusted the images to ensure image quality and dataset integrity. This step prevented us from achieving a fully automated creation process. Although such limitations exist, these aspects are expected to naturally diminish or disappear as in-painting models advance, making our dataset construction pipeline fully automatic and our method applicable to all images. Furthermore, we foresee future evaluations expanding in scale through manipulations like object removal, insertion, or resizing. While our current method may appear as a modest evaluation benchmark, we hope that our work becomes the first step toward controllable hallucination evaluation.

Acknowledgment

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2021-0-02068, Artificial Intelligence Innovation Hub; No.2022-0-00124, Development of Artificial Intelligence Technology for Self-Improving Competency-Aware Learning Capabilities; No.RS-2019-II191906, Artificial Intelligence Graduate School Program(POSTECH))

References

1. 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 (2023) [1](#), [4](#)
2. Agrawal, H., Desai, K., Wang, Y., Chen, X., Jain, R., Johnson, M., Batra, D., Parikh, D., Lee, S., Anderson, P.: Nocaps: Novel object captioning at scale. In: IEEE International Conference on Computer Vision (ICCV) (2019) [2](#)
3. Bisk, Y., Zellers, R., Gao, J., Choi, Y., et al.: Piqa: Reasoning about physical commonsense in natural language. In: AAAI Conference on Artificial Intelligence (AAAI) (2020) [1](#)
4. Chen, J., Zhu, D., Shen, X., Li, X., Liu, Z., Zhang, P., Krishnamoorthi, R., Chandra, V., Xiong, Y., Elhoseiny, M.: Minigpt-v2: large language model as a unified interface for vision-language multi-task learning. arXiv (2023) [4](#)
5. Chen, K., Zhang, Z., Zeng, W., Zhang, R., Zhu, F., Zhao, R.: Shikra: Unleashing multimodal llm’s referential dialogue magic. arXiv (2023) [1](#), [4](#), [9](#)
6. Chiang, W.L., Li, Z., Lin, Z., Sheng, Y., Wu, Z., Zhang, H., Zheng, L., Zhuang, S., Zhuang, Y., Gonzalez, J.E., Stoica, I., Xing, E.P.: Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality (March 2023), <https://lmsys.org/blog/2023-03-30-vicuna/> [1](#), [4](#)
7. Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., Hesse, C., Schulman, J.: Training verifiers to solve math word problems. arXiv (2021) [1](#)
8. Dai, W., Li, J., Li, D., Tiong, A.M.H., Zhao, J., Wang, W., Li, B., Fung, P., Hoi, S.: Instructblip: Towards general-purpose vision-language models with instruction tuning. In: Advances in Neural Information Processing Systems (NeurIPS) (2023) [1](#), [4](#), [9](#)
9. Deng, A., Chen, Z., Hooi, B.: Seeing is believing: Mitigating hallucination in large vision-language models via clip-guided decoding. arXiv (2024) [4](#)
10. Gao, P., Han, J., Zhang, R., Lin, Z., Geng, S., Zhou, A., Zhang, W., Lu, P., He, C., Yue, X., Li, H., Qiao, Y.: Llama-adapter v2: Parameter-efficient visual instruction model. arXiv (2023) [4](#)
11. Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., Parikh, D.: Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) [2](#), [11](#)
12. Gunjal, A., Yin, J., Bas, E.: Detecting and preventing hallucinations in large vision language models. In: AAAI Conference on Artificial Intelligence (AAAI) (2024) [2](#), [4](#)

13. Hessel, J., Holtzman, A., Forbes, M., Bras, R.L., Choi, Y.: Clipscore: A reference-free evaluation metric for image captioning. *arXiv* (2021) [13](#)
14. Hu, H., Zhang, J., Zhao, M., Sun, Z.: Ciem: Contrastive instruction evaluation method for better instruction tuning. *arXiv* (2023) [2](#), [4](#), [7](#)
15. Hudson, D.A., Manning, C.D.: Gqa: A new dataset for real-world visual reasoning and compositional question answering. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019) [2](#), [11](#)
16. Hyun, L., Sung-Bin, K., Han, S., Yu, Y., Oh, T.H.: Smile: Multimodal dataset for understanding laughter in video with language models. *arXiv* (2023) [1](#)
17. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. *arXiv* (2023) [7](#)
18. Leng, S., Zhang, H., Chen, G., Li, X., Lu, S., Miao, C., Bing, L.: Mitigating object hallucinations in large vision-language models through visual contrastive decoding. *arXiv* (2023) [4](#)
19. Li, Y., Du, Y., Zhou, K., Wang, J., Zhao, W.X., Wen, J.R.: Evaluating object hallucination in large vision-language models. In: *Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2023) [2](#), [4](#), [6](#), [7](#)
20. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *European Conference on Computer Vision (ECCV)* (2014) [6](#)
21. Liu, F., Lin, K., Li, L., Wang, J., Yacoob, Y., Wang, L.: Mitigating hallucination in large multi-modal models via robust instruction tuning. In: *International Conference on Learning Representations (ICLR)* (2024) [4](#)
22. Liu, H., Xue, W., Chen, Y., Chen, D., Zhao, X., Wang, K., Hou, L., Li, R., Peng, W.: A survey on hallucination in large vision-language models. *arXiv* (2024) [4](#)
23. Liu, H., Li, C., Li, Y., Lee, Y.J.: Improved baselines with visual instruction tuning. *arXiv* (2023) [1](#), [4](#), [9](#)
24. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2023) [1](#), [4](#)
25. Peng, Z., Wang, W., Dong, L., Hao, Y., Huang, S., Ma, S., Wei, F.: Kosmos-2: Grounding multimodal large language models to the world. In: *International Conference on Learning Representations (ICLR)* (2024) [4](#)
26. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2015) [13](#)
27. Rohrbach, A., Hendricks, L.A., Burns, K., Darrell, T., Saenko, K.: Object hallucination in image captioning. In: *Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2018) [12](#)
28. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2022) [7](#)
29. Sun, Z., Shen, S., Cao, S., Liu, H., Li, C., Shen, Y., Gan, C., Gui, L.Y., Wang, Y.X., Yang, Y., et al.: Aligning large multimodal models with factually augmented rlhf. *arXiv* (2023) [4](#)
30. Suvorov, R., Logacheva, E., Mashikhin, A., Remizova, A., Ashukha, A., Silvestrov, A., Kong, N., Goka, H., Park, K., Lempitsky, V.: Resolution-robust large mask inpainting with fourier convolutions. In: *IEEE Winter Conf. on Applications of Computer Vision (WACV)* (2022) [7](#)
31. Taori, R., Gulrajani, I., Zhang, T., Dubois, Y., Li, X., Guestrin, C., Liang, P., Hashimoto, T.B.: Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models*. **3**(6), 7 (2023) [4](#)

32. 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 (2023) [1](#), [4](#)
33. 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 (2023) [1](#), [4](#)
34. Wang, J., Wang, Y., Xu, G., Zhang, J., Gu, Y., Jia, H., Yan, M., Zhang, J., Sang, J.: An llm-free multi-dimensional benchmark for mllms hallucination evaluation. arXiv (2023) [2](#), [4](#), [7](#)
35. Wang, J., Zhou, Y., Xu, G., Shi, P., Zhao, C., Xu, H., Ye, Q., Yan, M., Zhang, J., Zhu, J., et al.: Evaluation and analysis of hallucination in large vision-language models. arXiv preprint arXiv:2308.15126 (2023) [4](#), [11](#)
36. 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 (2023) [9](#)
37. Yin, S., Fu, C., Zhao, S., Xu, T., Wang, H., Sui, D., Shen, Y., Li, K., Sun, X., Chen, E.: Woodpecker: Hallucination correction for multimodal large language models. arXiv (2023) [4](#)
38. Young, P., Lai, A., Hodosh, M., Hockenmaier, J.: From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. Transactions of the Association for Computational Linguistics (TACL) (2014) [2](#)
39. Zhai, B., Yang, S., Zhao, X., Xu, C., Shen, S., Zhao, D., Keutzer, K., Li, M., Yan, T., Fan, X.: Halle-switch: Rethinking and controlling object existence hallucinations in large vision language models for detailed caption. arXiv preprint arXiv:2310.01779 (2023) [2](#), [4](#)
40. Zhang, R., Han, J., Zhou, A., Hu, X., Yan, S., Lu, P., Li, H., Gao, P., Qiao, Y.: Llama-adapter: Efficient fine-tuning of language models with zero-init attention. In: International Conference on Learning Representations (ICLR) (2024) [4](#)
41. Zhao, Y., Li, Z., Jin, Z., Zhang, F., Zhao, H., Dou, C., Tao, Z., Xu, X., Liu, D.: Enhancing the spatial awareness capability of multi-modal large language model. arXiv (2023) [2](#)
42. Zhou, Y., Cui, C., Yoon, J., Zhang, L., Deng, Z., Finn, C., Bansal, M., Yao, H.: Analyzing and mitigating object hallucination in large vision-language models. In: International Conference on Learning Representations (ICLR) (2024) [2](#), [4](#)
43. Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M.: Minigpt-4: Enhancing vision-language understanding with advanced large language models. arXiv (2023) [4](#)
44. Zhuang, J., Zeng, Y., Liu, W., Yuan, C., Chen, K.: A task is worth one word: Learning with task prompts for high-quality versatile image inpainting. arXiv (2023) [7](#)