

Beyond Logit Lens: Contextual Embeddings for Robust Hallucination Detection & Grounding in VLMs

Anirudh Phukan³, Divyansh^{1*}, Harshit Kumar Morj^{2*}, Vaishnavi^{1*},
Apoorv Saxena³, Koustava Goswami³,

¹IIT Kanpur, ²IIT Bombay, ³Adobe Research

Correspondence: phukan@adobe.com

Abstract

The rapid development of Large Multimodal Models (LMMs) has significantly advanced multimodal understanding by harnessing the language abilities of Large Language Models (LLMs) and integrating modality-specific encoders. However, LMMs are plagued by hallucinations that limit their reliability and adoption. While traditional methods to detect and mitigate these hallucinations often involve costly training or rely heavily on external models, recent approaches utilizing internal model features present a promising alternative. In this paper, we critically assess the limitations of the state-of-the-art training-free technique, the logit lens, in handling generalized visual hallucinations. We introduce *ContextualLens*, a refined method that leverages contextual token embeddings from middle layers of LMMs. This approach significantly improves hallucination detection and grounding across diverse categories, including actions and OCR, while also excelling in tasks requiring contextual understanding, such as spatial relations and attribute comparison. Our novel grounding technique yields highly precise bounding boxes, facilitating a transition from Zero-Shot Object Segmentation to Grounded Visual Question Answering. Our contributions pave the way for more reliable and interpretable multimodal models.

1 Introduction

Recent advancements in multimodal understanding have been significantly driven by Large Multimodal Models (LMMs), which capitalize on the language capabilities of Large Language Models (LLMs) and integrate modality-specific understanding by training adapters that connect LLMs to pre-trained modality-specific encoders (Jin et al., 2024). However, critical issues inherent to LLMs, such as the tendency to produce highly confident incorrect answers—known as hallucinations—are also

Authors marked with * were interns at Adobe Research when this work was done.

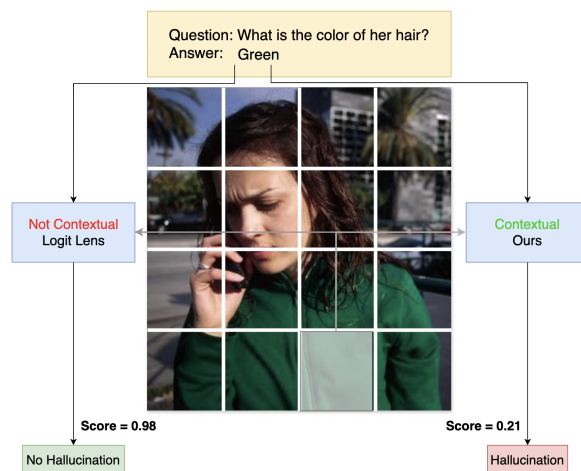


Figure 1: **Visual Hallucination Detection.** Logit lens, which only verifies the presence of the token "Green" in the image, mistakenly considers the answer correct. In contrast, *ContextualLens* assigns a low score to "Green" as it correctly contextualizes the color to the jacket rather than the hair.

transferred to LMMs (Bai et al., 2024). Moreover, LMMs introduce additional hallucinations specific to the integrated modalities (Liu et al., 2024a). Addressing these hallucinations, both by detecting their occurrence and mitigating their effects, as well as by providing evidence to support the generated responses, is crucial for fostering user trust and facilitating the widespread adoption of these technologies (Bohnet et al., 2022).

Detecting and mitigating hallucinations has been extensively explored in both language (Tonmoy et al., 2024) and vision domains (Liu et al., 2024a). In the context of language, attribution and citation have been proposed as approaches for providing evidence to support model output (Gao et al., 2023; Huo et al., 2023). In computer vision, similar methods are typically studied under Grounded Visual Question Answering (Zhang et al., 2024; Khosh-sirat and Kambhamettu, 2023). However, many existing techniques necessitate either training from

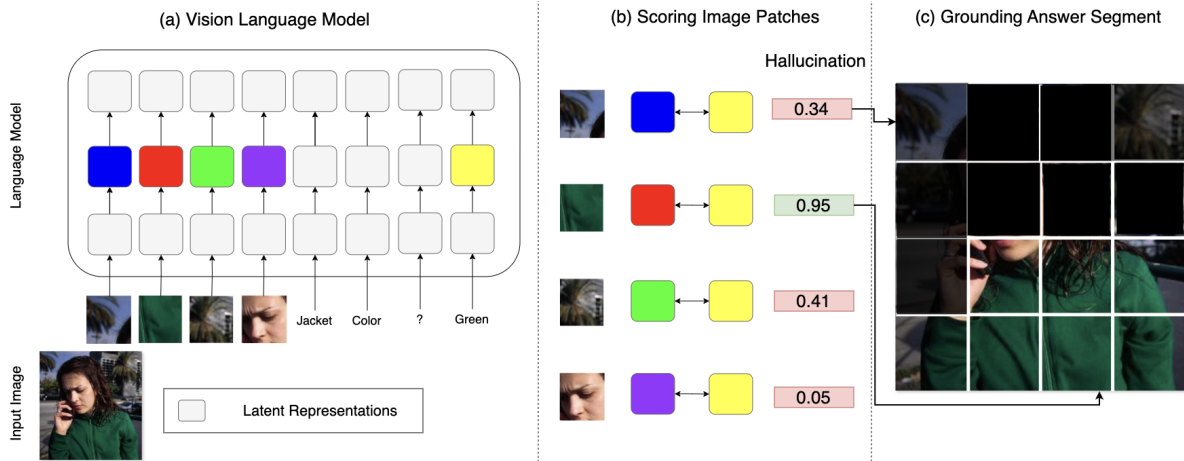


Figure 2: (a) We extract latent representations of image patches and answer tokens from intermediate layers, (b) compute the cosine similarity between each image patch and the average embedding of the answer tokens to score patches for hallucination detection, and (c) ground the answer in specific image patches.

scratch or fine-tuning (Jiang et al., 2024), and frequently rely on external models such as retrievers or object detectors. These requirements impose significant training costs or result in increased latency during inference, posing challenges to their practical deployment in real-world scenarios.

Recent studies have increasingly leveraged internal model features, such as latent representations, logits, and attention weights, to tackle the objectives of detecting hallucinations (Azaria and Mitchell, 2023; Varshney et al., 2023) and ensuring the grounding of generated outputs in LLMs (Phukan et al., 2024; Qi et al., 2024). Notably, one such study demonstrates the efficacy of a training-free interpretability technique, known as the logit lens, in identifying and mitigating object hallucinations within Visual-Language Models (VLMs) (Jiang et al., 2024).

The logit lens technique, originally introduced in the context of language models, involves directly mapping intermediate activations to the vocabulary space using the unembedding layer, allowing for an interpretable view of token predictions at different layers (nostalgebraist, 2020). By applying this technique to VLMs, the authors probe each image patch to determine the presence of objects.

However, we identify several fundamental limitations with the application of the logit lens in this context, which hinder its effectiveness in addressing more general forms of visual hallucinations as categorized by Yan et al. (2024). We contend that the reliance on token embeddings from the unembedding layer—which are neither contextual nor easily combinable to form multi-token con-

cepts—results in its failure to handle more generalized hallucination scenarios. Consequently, we propose leveraging token embeddings from the middle layers, which have been shown to be contextual (Phukan et al., 2024) and effective in representing concepts (Wendler et al., 2024).

By employing contextual answer token embeddings, we are able to detect hallucinations in categories that were previously near random performance and improve detection accuracy in other categories. Additionally, this method enables a transition from Zero-Shot Object Segmentation using the logit lens technique to performing the more general Grounded Visual Question Answering task. This advancement underscores the potential of our approach in enhancing the reliability and interpretability of multimodal models while operating in the training-free paradigm.

Our contributions are threefold:

- We investigate the robustness of a *SOTA* training-free VLM Object Hallucination Detection and Segmentation method on VQA datasets. We find that the method does not generalize well to actions and OCR, also completely failing on tasks requiring contextual understanding such as attributes, spatial relations, and comparisons.
- We extend the method by introducing *ContextualLens*, which replaces the use of the logit lens with contextual embeddings from the middle layers of the VLM, leading to successful hallucination detection in categories that previously performed near random and

improving detection accuracy in other categories.

- We propose a novel grounding technique that returns highly precise bounding boxes, enabling a transition from Zero-Shot Object Segmentation to performing the more general Grounded Visual Question Answering task.

2 Related Work

2.1 Hallucination Detection and Mitigation in Large Foundation Models

Recent advancements in hallucination detection within LLMs have been driven by studying output logits (Varshney et al., 2023), activations (Chen et al., 2024a), and latent representations (Azaria and Mitchell, 2023; Su et al., 2024). Despite the progress in detection, predominant strategies for hallucination mitigation encompass Retrieval Augmented Generation (Gao et al., 2022; Peng et al., 2023a), Iterative Prompting (Ji et al., 2023), Supervised Fine-tuning (Tian et al., 2023), and Alternative Decoding Strategies (Shi et al., 2023; Chuang et al., 2023), requiring specific training or external models.

Additionally, VLMs necessitate distinct approaches to tackle issues inherent to the vision modality. Effective methods span increasing visual resolution (Bai et al., 2023; Li et al., 2024), integrating segmentation and depth maps (Jain et al., 2024), enhancing connection modules (Chen et al., 2024b), and optimizing the decoding process (Huang et al., 2024). Jiang et al. (2024) propose an innovative, training-free method, using the logit lens to probe individual image patch embeddings for the detection and mitigation of object hallucinations.

2.2 Attribution and Grounded Visual Question Answering

In the language domain, citing or attributing generated texts to the sources are commonly studied for building user trust. Methods fall into three major categories: attribution using fine-tuned/trained models (Gao et al., 2023; Sun et al., 2022), attribution using external retrievers or auxiliary models (Huo et al., 2023; Lee et al., 2019; Ramu et al., 2024; Sancheti et al., 2024), and methodologies that harness internal model features directly (Phukan et al., 2024; Qi et al., 2024). Phukan et al. (2024) leverage contextual embeddings from the intermediate layers of LLMs to pair answer

tokens with document tokens by utilizing embedding similarity. Similarly, Qi et al. (2024) employ KL-divergence measures between logits with and without contextual information to pinpoint context-sensitive output tokens.

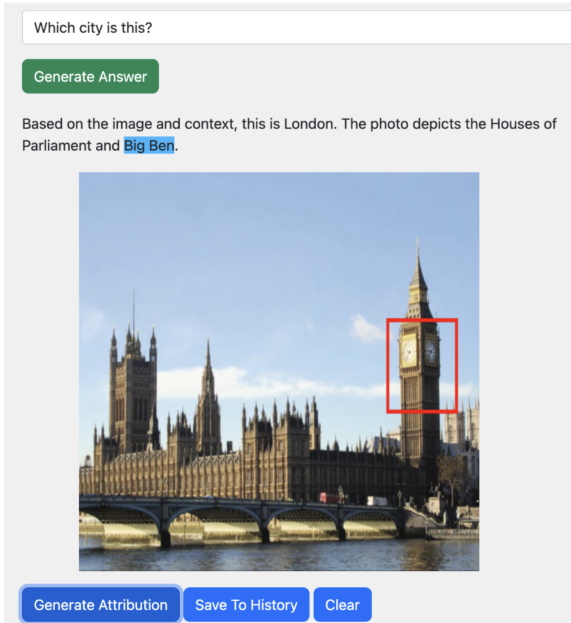
In the vision-language domain, the task of providing evidence for model outputs has been extensively explored under Grounded Visual Question Answering (GVQA). Prior to the development of LMMs, GVQA relied on custom, end-to-end trained architectures (Tan and Bansal, 2019; Zhang et al., 2021) or specialized attention mechanisms (Urooj et al., 2021; Khoshsirat and Kambhamettu, 2023). More recent advancements have given rise to grounding LMMs (Peng et al., 2023b; Rasheed et al., 2024), which generate segmentation masks or bounding boxes for objects and attributes in related tasks such as Grounded Conversation Generation (GCG), Grounded Image Captioning (GIC), and Reference Expression Segmentation (RES). Notably, *GROUNDHOG* (Zhang et al., 2024) is a grounding LMM capable of performing GVQA.

ContextualLens is inspired by the training-free object hallucination detection and mitigation method proposed by Jiang et al. (2024), but with a crucial distinction. We utilize output token embeddings from intermediate layers instead of employing the logit lens probe. This modification enables us to overcome several limitations of the previous work, particularly by enhancing the representation of multi-token concepts and contextual relationships (§4.2).

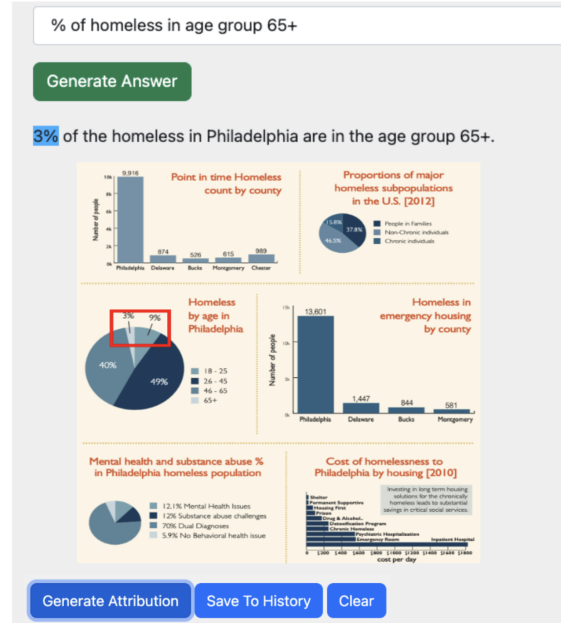
3 Preliminaries

3.1 Vision Language Model Architecture

Vision-Language Models (VLMs) such as Llava (Liu et al., 2024b) use the following general recipe to blend visual and textual inputs. Formally, let \mathcal{M}_{vl} denote a VLM, comprising three primary components: a vision encoder \mathcal{V} , a connection module \mathcal{C} , and a Large Language Model (LLM) \mathcal{M}_l . Given an input image I , the vision encoder \mathcal{V} processes this image to produce n image features, each corresponding to a distinct patch within the image. These features are then projected by the connection module \mathcal{C} into the d -dimensional input embedding space of the LLM \mathcal{M}_l . Subsequently, a textual prompt p is concatenated with the projected image features to generate a sequence of m text features. This combined sequence traverses through the L layers of \mathcal{M}_l , where each layer refines and



(a) Big Ben



(b) Chart

Figure 3: Qualitative examples for Grounded Visual Question Answering. Here is a GUI where we pass an image and ask a question. After the answer is generated, we select the span and click "Generate Attribution". Observe, the bounding boxes correctly ground the selected span to the image.

integrates the multimodal information. For a given input embedding x , we denote the latent representation at layer l as $h_l(x)$. Lastly, at the final layer, these latent representations are mapped back to the vocabulary space using the Unembedding Matrix $W_U \in \mathbb{R}^{|V| \times d}$, where V represents the vocabulary of \mathcal{M}_l .

3.2 Interpreting Latent Representations

The logit lens technique involves applying the Unembedding Matrix W_U to intermediate latent representations $h_l(x)$ in order to obtain logit distributions over the entire vocabulary. Jiang et al. (2024) utilized this technique to estimate the probability of individual image patches aligning with specific tokens in the vocabulary by performing $W_U \cdot h_l(t_i)$, where t_i represents the intermediate representation of an image patch.

Phukan et al. (2024) highlighted how to map answers tokens to document tokens using contextual embeddings from intermediate latent representations. Given a span of tokens, the embedding for the entire span is computed by averaging the embeddings of the individual tokens within that span.

To compare the similarity between spans, cosine similarity is employed. The scoring enables the effective comparison of contextual embeddings derived from different spans, facilitating the iden-

tification of semantically related segments within the text.

3.3 Grounded Visual Question Answering

Grounded Visual Question Answering (GVQA) extends traditional VQA by incorporating the explicit requirement for the answers to be justified with visual evidence from the image. Given an input image I , a textual question Q , and a VLM \mathcal{M}_{vl} , the model generates an answer A and a grounding map G . The grounding map G identifies the regions within the image I that correspond to the elements of the answer A . The critical challenge in GVQA lies in ensuring that the grounding map G faithfully represents the visual basis for each part of the generated answer A .

4 Proposed Work

4.1 Motivation

Hallucination detection and mitigation in both language and vision modalities have mostly involved extensive re-training or fine-tuning, substantially increasing computational costs and latency, which hampers their real-world applicability. The logit lens methodology, which probes individual image patches using token embeddings from the unembedding layer, marks a step forward by offering

a training-free alternative in the vision-language domain.

However, as demonstrated in our analysis (§6.1), the logit lens method exhibits significant performance degradation when applied to more complex hallucination scenarios, such as those involving attributes, comparisons, and relations among objects. This limitation is primarily due to its reliance on non-contextual embeddings, which are incapable of effectively representing multi-token concepts and contextual elements such as spatial relationships and attribute-based discrepancies.

We aim to address the shortcomings of the logit lens by leveraging middle-layer, contextual embeddings known for their ability to encapsulate richer semantic information. By doing so, we hope to enhance hallucination detection across a more comprehensive range of categories and provide precise visual grounding, while maintaining the computational efficiency and training-free nature of the logit lens.

4.2 Methodology

4.2.1 Hallucination Detection

To perform hallucination detection, we utilize the contextual embeddings from intermediate layers of the VLM. Our method proceeds as follows:

Given the answer tokens generated by the model, we compute the average embedding at a specific layer l_T , denoted as h_{Ans,l_T} . Formally, let $\{t_1, t_2, \dots, t_k\}$ be the sequence of answer tokens, and $h_{l_T}(t_i)$ be the embedding of token t_i at layer l_T . The average embedding h_{Ans,l_T} is computed as:

$$h_{\text{Ans},l_T} = \frac{1}{k} \sum_{i=1}^k h_{l_T}(t_i)$$

Next, we evaluate the similarity between the answer embedding h_{Ans,l_T} and each image patch embedding from a specific layer l_I . Let $\{p_1, p_2, \dots, p_n\}$ denote the set of image patches, and $h_{l_I}(p_j)$ be the embedding of patch p_j at layer l_I . We compute the score for each patch p_j as the cosine similarity:

$$\text{Score}(p_j) = \text{CosineSim}(h_{\text{Ans},l_T}, h_{l_I}(p_j))$$

The resulting scores, denoted as Scores, represent the relevance of each image patch with respect

to the answer embedding. To determine the confidence for hallucination, we consider the maximum score in Scores:

$$\text{Confidence}_{\text{max}} = \max(\text{Scores})$$

A high $\text{Confidence}_{\text{max}}$ indicates a low likelihood of hallucination, as it suggests a strong correspondence between the answer and at least one image patch. Conversely, a low $\text{Confidence}_{\text{max}}$ signals potential hallucination, implying that the answer may not be visually supported by any segment of the image.

4.2.2 Grounded Visual Question Answering

To effectively ground the generated answers in corresponding image regions, we introduce two techniques: a refined version of (Jiang et al., 2024) based on contextual embeddings and an alternative approach that directly returns bounding boxes for visual grounding.

Basic Technique: In our first approach, we compute the average embedding of the answer tokens across all layers, denoted as h_{Ans} , with a shape of (L, d) , where L represents the number of layers and d the embedding dimension. Formally, given the sequence of answer tokens $\{t_1, t_2, \dots, t_k\}$, the answer embedding h_{Ans} is computed as follows:

$$h_{\text{Ans},l} = \frac{1}{k} \sum_{i=1}^k h_l(t_i), l \in \{1, 2, \dots, L\}$$

This results in a layer-wise answer embedding $h_{\text{Ans}} = (h_{\text{Ans},1}, h_{\text{Ans},2}, \dots, h_{\text{Ans},L})$. Simultaneously, for each image patch p_j , we extract its embedding across all layers, denoted as $h_{p_j} = (h_1(p_j), h_2(p_j), \dots, h_L(p_j))$.

For each layer, we compute the cosine similarity between the layer-specific answer embedding $h_{\text{Ans},l}$ and the corresponding image patch embedding $h_l(p_j)$. This gives us a set of scores for each patch across layers: $\text{Scores}_l = \{\text{Score}_l(p_1), \text{Score}_l(p_2), \dots, \text{Score}_l(p_n)\}$.

The final score for each patch p_j is determined by taking the maximum score across all layers:

$$\text{FinalScore}(p_j) = \max_{l \in \{1, \dots, L\}} (\text{Score}_l(p_j))$$

This method allows us to evaluate the relevance of each image patch to the answer tokens, providing a grounding map based on the similarity score. This map is resized to the original image dimensions.

| Category | InternlmVL-7B | | | | Qwen2VL-7B | | | |
|-------------|---------------|-------|--------------|--------------|------------|-------|--------------|--------------|
| | Random | LL | Out Probs | CL (Ours) | Random | LL | Out Probs | CL (Ours) |
| Action | 0.776 | 0.795 | 0.787 | <u>0.796</u> | 0.604 | 0.636 | 0.710 | 0.752 |
| Attribute | 0.796 | 0.786 | 0.820 | <u>0.825</u> | 0.812 | 0.830 | 0.839 | 0.911 |
| Comparison | 0.576 | 0.580 | 0.563 | 0.623 | 0.548 | 0.558 | 0.567 | 0.685 |
| Count | 0.856 | 0.898 | 0.946 | 0.885 | 0.804 | 0.860 | 0.956 | 0.889 |
| Environment | 0.748 | 0.771 | <u>0.835</u> | 0.811 | 0.600 | 0.465 | 0.633 | 0.682 |
| Relation | 0.656 | 0.668 | <u>0.755</u> | <u>0.755</u> | 0.592 | 0.572 | 0.647 | <u>0.655</u> |
| OCR | 0.740 | 0.769 | 0.856 | 0.772 | 0.740 | 0.793 | 0.860 | <u>0.871</u> |

Table 1: Comparison of mAP scores across different hallucination categories in the HQH dataset. The best-performing method in each category is **bolded** if it is significantly better than the second-best, and underlined if the difference is marginal.

Bounding Box Technique: In our second approach, we directly identify bounding boxes that correspond to answer tokens, bypassing the need for thresholding each patch. This approach is more robust for practical applications.

Given the average embedding of the answer tokens at a specific layer l_b , denoted as h_{Ans} , we aim to identify the most relevant bounding box from a set of potential bounding boxes \mathcal{S} . Let W and H be the number of image patches across the width and height of the image, respectively. Thus, $|\mathcal{S}| = W^2 \times H^2$, representing all possible bounding boxes within the image.

For each bounding box $s \in \mathcal{S}$, we define the embedding h_s as the average embedding of all the image patches within the bounding box at layer l_b :

$$h_s = \frac{1}{|P_s|} \sum_{p_j \in P_s} h_{l_b}(p_j)$$

where P_s is the set of patches contained in the bounding box s , and $h_{l_b}(p_j)$ represents the embedding of patch p_j at layer l_b .

Next, we compute the cosine similarity between the answer embedding h_{Ans} and the embedding of each bounding box h_s . The bounding box s^* that maximizes this cosine similarity is selected as the grounded region:

$$s^* = \arg \max_{s \in \mathcal{S}} \text{CosineSim}(h_{Ans}, h_s)$$

This method harnesses the average embeddings of the answer tokens and bounding boxes to directly locate the most relevant region within the image that corresponds to the generated answer.

5 Experimental Setup

5.1 Datasets

We conducted experiments on three datasets to evaluate methods for both tasks outlined in §4.2. More

details can be found in Appendix A.

Hallucination Detection: We experimented on High-Quality Hallucination Benchmark (HQH) (Yan et al., 2024) dataset, comprising 4,000 image-question pairs accompanied by ground-truth answers and categorized into eight distinct types of potential hallucination scenarios, allowing us to test for the generalization of methods.

Grounded VQA: Experiments on two datasets TextVQA-X (Rao et al., 2021) and VizWiz-G (Chen et al., 2022) highlight the capability of our method’s grounding performance on GVQA task. The datasets have 3,620 and 1,131 image-question pairs respectively.

5.2 Metrics

We use the Mean Average Precision (mAP) metric for hallucination detection and plot the Precision-Recall Curve to evaluate the grounding aspect of GVQA. More details can be found in Appendix B.

5.3 Baselines

5.3.1 Hallucination Detection

We employ two primary baselines for comparative evaluation. The first baseline leverages the maximum probability across the generated answer tokens, referred to as Output Probabilities (Out Probs). The second baseline is the work by Jiang et al. (2024) (LL). They evaluate the hallucination likelihood by taking the maximum value of the softmax-normalized output token logits over all layers. Logit computation is described in §3.2.

5.3.2 Grounded Visual Question Answering

We adapt the method proposed by (Jiang et al., 2024) for zero-shot segmentation, which to the best of our knowledge is the only training-free grounding technique for VLMs. We take the mean internal confidence for tokens comprising the answer, similar to how they perform hallucination detection.

We resize the set of internal confidence values per image patch back to the size of the image.

6 Results and Analysis

We perform experiments for our approach with the following VLMs: Qwen2-VL-7B (Wang et al., 2024) and internlm-xcomposer2-v1-7b (Dong et al., 2024). All of our experiments were run on a A100 machine with 4 80GB GPUs.

6.1 Hallucination Detection

The performance of *ContextualLens* (§4.2.1) and the baseline methods (§5.3.1) on the HQH dataset is summarized in Table 1. Each category comprises 500 image-question pairs. To select the best layer combination for our method, we created a test and validation split with 250 examples each. The scores in Table 1 are reported for the test split. Further details on how to perform layer selection in real-world settings, including a robustness analysis using adversarial validation, are provided in Appendix C.

Notably, the logit lens technique, although effective for object hallucination detection as shown by (Jiang et al., 2024), performs nearly at random in the Attribute, Comparison, and Relation categories across both models. Further investigation into the types of questions in these categories reveals that the Attribute category often involves questions about the color of objects, the Comparison category includes questions comparing two objects (e.g., which is larger), and the Relation category consists of questions about objects' spatial relations. The logit lens technique falls short in these scenarios because it relies on non-contextual embeddings and thus can only determine the presence of objects or features, not their spatial relationships or comparative attributes.

For instance, in the Attribute category, if the question is about the color of an object (e.g., "What is the color of the ball?") and the model answers "blue," the logit lens can only verify the presence of the color blue in the image but not whether the color pertains to the ball or another object like the sky. This underscores the limitation of using non-contextual embeddings for complex hallucination detection tasks.

Our method, which employs contextual embeddings, fares significantly better in these categories compared to the logit lens approach and often outperforms even output probabilities, sometimes by

a substantial margin. We refer readers to Figures 6 and 7 for some qualitative examples.

For categories such as Action and OCR, where grounding actions (e.g., "Looking at the Television") or text (e.g., "Melrose Ave.") is necessary, the logit lens technique performs better than random but still falls short compared to output probabilities. This is likely due to the logit lens's poor handling of multi-token objects. While objects which are usually single token may align relatively well, the complexity of representing actions or texts across multiple tokens poses challenges. Our approach, leveraging middle-layer embeddings known to better represent multi-token concepts (Wendler et al., 2024), consistently outperforms the logit lens and, in many instances, surpasses output probabilities. See Figure 8 for a qualitative example.

Interestingly, output probabilities excel in the Count category compared to both *ContextualLens* and the logit lens. This category involves questions about the number of instances of a specific object. The inherent challenge for our approach and the logit lens is that visual elements often have low semantic overlap with numerical tokens (e.g., a "car" in the image is unlikely to semantically match "two" in the text). Consequently, output probabilities provide more accurate hallucination detection in such scenarios. A failure case is illustrated in Figure 9.

6.2 Grounded Visual Question Answering

The performance of our methods (§4.2.2) and the baseline method (§5.3.2) on the TextVQA-X dataset and VizWiz-G dataset is summarized in Figure 4 and 5 respectively.

From the plots, it is evident that our adaptation of Jiang et al. (2024)'s method for object segmentation, by replacing the logit lens with contextual middle layer representations, consistently leads to better grounding for VQA across the tested models and datasets. The TextVQA-X dataset is analogous to the OCR category in the HQH dataset, while VizWiz-G encompasses attributes, OCR, and general object segmentation. Our method's improved grounding performance highlights the utility of using contextual and conceptual embeddings, which also elucidates why we achieve superior hallucination detection.

An essential use case of grounded VQA is not only to understand what image features contribute to answer generation but also to guide users to the

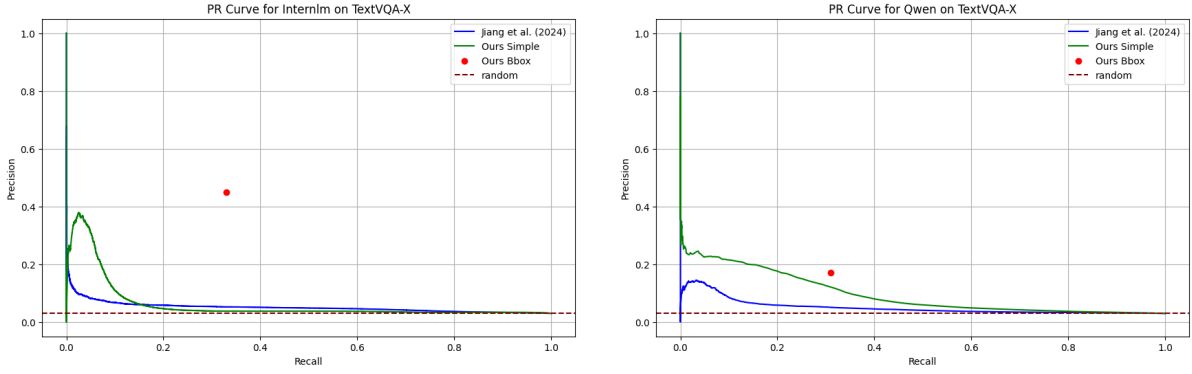


Figure 4: PR Curves on the TextVQA-X dataset.

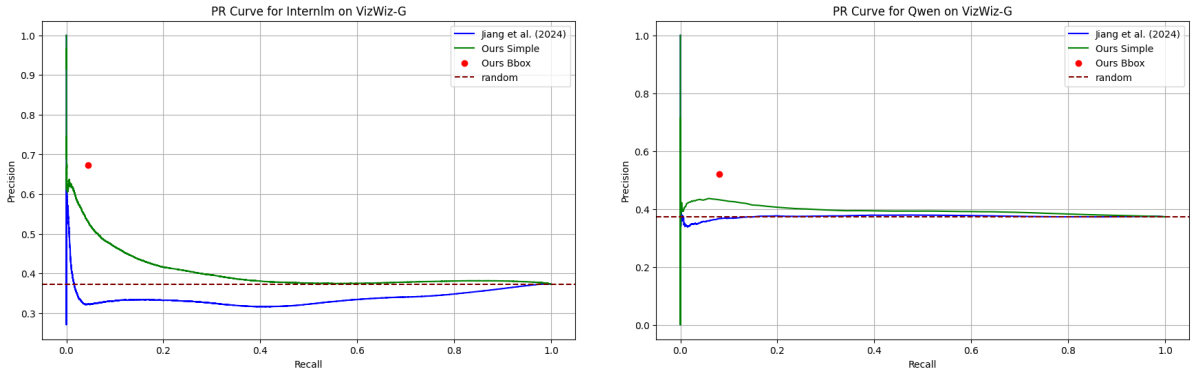


Figure 5: PR Curves on the VizWiz-G dataset.

general location of the evidence within the image. This need parallels the concept of attribution explored in LLMs and textual domains. We term this broader application in VLMs as multimodal attribution, where precision is more critical than recall. Users prioritize accurate identification of the general evidence location over highlighting every relevant pixel. This is precisely where our bounding box method proves valuable.

As illustrated in Figures 4 and 5, the precision of the bounding box method is often significantly higher than what can be achieved by navigating the PR-Curve for the basic version. This precision makes it particularly useful for multimodal attribution, providing users with accurate general locations of evidence, which they can then verify. The higher precision and comparatively lower recall is because we identify the optimal bounding box, which perfectly encloses the relevant evidence. The recall is higher on TextVQA-X since we ground text, and the optimal bounding box often encloses the entire textual element. In VizWiz-G, however, questions about object attributes like color could be satisfied by highlighting a small patch of that color on the object, leading to a lower recall.

6.2.1 Qualitative Examples

Some qualitative results of our bounding box visual grounding technique are showcased in Figure 3. For instance, in Figure 3a, we ground the phrase "Big Ben" to its corresponding location in the image. Our method effectively highlights the relevant portion by leveraging the latent conceptual representation of "Big Ben" intrinsic to the LMM, facilitating the mapping between textual descriptions and visual concepts. Another example is depicted in Figure 3b. Remarkably, our method proficiently grounds answers within charts and infographics, a novel capability within the community to the best of our knowledge, achievable in a training-free manner. Additional qualitative examples are available in Appendix D.

7 Discussion & Conclusion

In this paper, we presented *ContextualLens*, an approach for detecting hallucinations in LMMs by leveraging contextual token embeddings from intermediate layers. We identified limitations in the state-of-the-art training-free technique, the logit lens, notably its poor handling of complex visual hallucinations involving attributes, comparisons,

and spatial relations due to its reliance on non-contextual embeddings. Our proposed method employs contextual embeddings known to capture richer semantic information, improving hallucination detection across diverse categories.

Our experimental results on the HQH benchmark showed that *ContextualLens* enables detection in categories that previously performed near random and improves detection for multi-token concept representation. Furthermore, by introducing a novel grounding technique with highly precise bounding boxes, we advanced Zero-Shot Object Segmentation to a more general GVQA task, validating our approach’s effectiveness on TextVQA-X and VizWiz-G datasets.

Our method’s precision and training-free paradigm circumvent the computational costs of re-training or fine-tuning, promoting its practical applicability in real-world multimodal attribution tasks. Our contributions pave the way for more reliable and user-trustworthy LMMs.

8 Limitations & Future Work

While our work significantly advances the state-of-the-art in training-free hallucination detection and grounding in LMMs, it has some limitations that provide exciting avenues for future research.

Firstly, the validation of both hallucination detection and grounding has been conducted primarily on factual short answer VQA datasets. Extending this to more diverse and complex datasets remains an area for exploration. Secondly, *ContextualLens* is outperformed by output probabilities in the Count category. Future work could explore extensions to the method for abstractive scenarios such as counting. Lastly, although our method excels at highlighting precise evidence regions sufficient for human verification, improving recall remains an open challenge. Subsequent research could investigate leveraging our scoring mechanisms as priors for pre-trained segmentation models or improving the evaluation metrics for grounded visual question answering.

References

Amos Azaria and Tom Mitchell. 2023. The internal state of an llm knows when it’s lying. *arXiv preprint arXiv:2304.13734*.

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A frontier large

vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.

- Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou. 2024. Hallucination of multimodal large language models: A survey. *arXiv preprint arXiv:2404.18930*.
- Bernd Bohnet, Vinh Q Tran, Pat Verga, Roei Aharoni, Daniel Andor, Livio Baldini Soares, Massimiliano Ciaramita, Jacob Eisenstein, Kuzman Ganchev, Jonathan Herzig, et al. 2022. Attributed question answering: Evaluation and modeling for attributed large language models. *arXiv preprint arXiv:2212.08037*.
- Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. 2024a. Inside: Llms’ internal states retain the power of hallucination detection. *arXiv preprint arXiv:2402.03744*.
- Chongyan Chen, Samreen Anjum, and Danna Gurari. 2022. Grounding answers for visual questions asked by visually impaired people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19098–19107.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024b. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198.
- Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. 2023. Dola: Decoding by contrasting layers improves factuality in large language models. *arXiv preprint arXiv:2309.03883*.
- Xiaoyi Dong, Pan Zhang, Yuhang Zang, Yuhang Cao, Bin Wang, Linke Ouyang, Xilin Wei, Songyang Zhang, Haodong Duan, Maosong Cao, et al. 2024. Internlm-xcomposer2: Mastering free-form text-image composition and comprehension in vision-language large model. *arXiv preprint arXiv:2401.16420*.
- Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Y Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, et al. 2022. Rarr: Researching and revising what language models say, using language models. *arXiv preprint arXiv:2210.08726*.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. Enabling large language models to generate text with citations. *arXiv preprint arXiv:2305.14627*.
- 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13418–13427.

- Siqing Huo, Negar Arabzadeh, and Charles Clarke. 2023. Retrieving supporting evidence for generative question answering. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 11–20.
- Jitesh Jain, Jianwei Yang, and Humphrey Shi. 2024. Vcoder: Versatile vision encoders for multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27992–28002.
- Ziwei Ji, Tiezheng Yu, Yan Xu, Nayeon Lee, Etsuko Ishii, and Pascale Fung. 2023. Towards mitigating hallucination in large language models via self-reflection. *arXiv preprint arXiv:2310.06271*.
- Nick Jiang, Anish Kachinthaya, Suzie Petryk, and Yossi Gandelsman. 2024. Interpreting and editing vision-language representations to mitigate hallucinations. *arXiv preprint arXiv:2410.02762*.
- Yizhang Jin, Jian Li, Yexin Liu, Tianjun Gu, Kai Wu, Zhengkai Jiang, Muyang He, Bo Zhao, Xin Tan, Zhenye Gan, et al. 2024. Efficient multimodal large language models: A survey. *arXiv preprint arXiv:2405.10739*.
- Seyedalireza Khoshsirat and Chandra Kambhamettu. 2023. Sentence attention blocks for answer grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6080–6090.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. Latent retrieval for weakly supervised open domain question answering. *arXiv preprint arXiv:1906.00300*.
- Zhang Li, Biao Yang, Qiang Liu, Zhiyin Ma, Shuo Zhang, Jingxu Yang, Yabo Sun, Yuliang Liu, and Xiang Bai. 2024. Monkey: Image resolution and text label are important things for large multi-modal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26763–26773.
- Hanchao Liu, Wenyuan Xue, Yifei Chen, Dapeng Chen, Xiutian Zhao, Ke Wang, Liping Hou, Rongjun Li, and Wei Peng. 2024a. A survey on hallucination in large vision-language models. *arXiv preprint arXiv:2402.00253*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024b. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- nostalgebraist. 2020. Interpreting gpt: The logit lens. <https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens>. Aug 2020.
- Baolin Peng, Michel Galley, Pengcheng He, Hao Cheng, Yujia Xie, Yu Hu, Qiuyuan Huang, Lars Liden, Zhou Yu, Weizhu Chen, et al. 2023a. Check your facts and try again: Improving large language models with external knowledge and automated feedback. *arXiv preprint arXiv:2302.12813*.
- Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. 2023b. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*.
- Anirudh Phukan, Shwetha Somasundaram, Apoorv Saxena, Koustava Goswami, and Balaji Vasani Srinivasan. 2024. Peering into the mind of language models: An approach for attribution in contextual question answering. *arXiv preprint arXiv:2405.17980*.
- Jirui Qi, Gabriele Sarti, Raquel Fernández, and Arianna Bisazza. 2024. Model internals-based answer attribution for trustworthy retrieval-augmented generation. *arXiv preprint arXiv:2406.13663*.
- Pritika Ramu, Koustava Goswami, Apoorv Saxena, and Balaji Vasani Srinivasan. 2024. Enhancing post-hoc attributions in long document comprehension via coarse grained answer decomposition. *arXiv preprint arXiv:2409.17073*.
- Varun Nagaraj Rao, Xingjian Zhen, Karen Hovsepian, and Mingwei Shen. 2021. A first look: Towards explainable textvqa models via visual and textual explanations. *arXiv preprint arXiv:2105.02626*.
- Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M Anwer, Eric Xing, Ming-Hsuan Yang, and Fahad S Khan. 2024. Glamm: Pixel grounding large multimodal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13009–13018.
- Abhilasha Sancheti, Koustava Goswami, and Balaji Vasani Srinivasan. 2024. Post-hoc answer attribution for grounded and trustworthy long document comprehension: Task, insights, and challenges. *arXiv preprint arXiv:2406.06938*.
- Weijia Shi, Xiaochuang Han, Mike Lewis, Yulia Tsvetkov, Luke Zettlemoyer, and Scott Wen-tau Yih. 2023. Trusting your evidence: Hallucinate less with context-aware decoding. *arXiv preprint arXiv:2305.14739*.
- Weihang Su, Changyue Wang, Qingyao Ai, Yiran Hu, Zhijing Wu, Yujia Zhou, and Yiqun Liu. 2024. Unsupervised real-time hallucination detection based on the internal states of large language models. *arXiv preprint arXiv:2403.06448*.
- Zhiqing Sun, Xuezhi Wang, Yi Tay, Yiming Yang, and Denny Zhou. 2022. Recitation-augmented language models. *arXiv preprint arXiv:2210.01296*.
- Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*.

- Katherine Tian, Eric Mitchell, Huaxiu Yao, Christopher D Manning, and Chelsea Finn. 2023. Fine-tuning language models for factuality. *arXiv preprint arXiv:2311.08401*.
- SM Tonmoy, SM Zaman, Vinija Jain, Anku Rani, Vipula Rawte, Aman Chadha, and Amitava Das. 2024. A comprehensive survey of hallucination mitigation techniques in large language models. *arXiv preprint arXiv:2401.01313*.
- Aisha Urooj, Hilde Kuehne, Kevin Duarte, Chuang Gan, Niels Lobo, and Mubarak Shah. 2021. Found a reason for me? weakly-supervised grounded visual question answering using capsules. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8465–8474.
- Neeraj Varshney, Wenlin Yao, Hongming Zhang, Jian-shu Chen, and Dong Yu. 2023. A stitch in time saves nine: Detecting and mitigating hallucinations of llms by validating low-confidence generation. *arXiv preprint arXiv:2307.03987*.
- 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*.
- Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2024. Do llamas work in english? on the latent language of multilingual transformers. *arXiv preprint arXiv:2402.10588*.
- Bei Yan, Jie Zhang, Zheng Yuan, Shiguang Shan, and Xilin Chen. 2024. Evaluating the quality of hallucination benchmarks for large vision-language models. *arXiv preprint arXiv:2406.17115*.
- Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5579–5588.
- Yichi Zhang, Ziqiao Ma, Xiaofeng Gao, Suhaila Shakiah, Qiaozi Gao, and Joyce Chai. 2024. Groundhog: Grounding large language models to holistic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14227–14238.

A Datasets

High-Quality Hallucination Benchmark (HQH) (Yan et al., 2024): This dataset comprises 4,000 image-question pairs accompanied by ground-truth answers and is categorized into eight distinct types of potential hallucination scenarios: Attribute, Action, Counting, Environment, Comparison, Relation, OCR, and Existence. The questions in HQH are open-ended, eliciting concise and direct answers. Evaluation within this benchmark is conducted using GPT-3.5 prompting, which has demonstrated high reliability and validity. This benchmark allows us to assess the generalization capabilities of our proposed strategies across multiple hallucination categories. Notably, we exclude the existence category from our evaluation, as our focus is on determining the presence of objects in the image, rather than assessing exhaustive object enumeration.

TextVQA-X (Rao et al., 2021): This dataset features human-annotated multimodal explanations, including ground truth segmentation maps and multiple references for textual explanations containing text within images. For our experiments, we utilize the validation split, which consists of 3,620 image-question pairs related to scene-text, accompanied by corresponding ground-truth answers. This dataset enables the assessment of grounding performance on GVQA task.

VizWiz-G (Chen et al., 2022): This dataset focuses on visually grounding answers to visual questions posed by individuals with visual impairments. Another GVQA dataset, VizWiz-G’s validation split includes 1,131 image-question pairs, each paired with a ground-truth answer and a corresponding segmentation mask.

B Metrics

B.1 Hallucination Detection

Mean Average Precision (mAP): We measure hallucination detection by framing it as a binary classification task. mAP provides a comprehensive evaluation by considering the precision-recall trade-offs across various threshold values.

B.2 Grounded Visual Question Answering

Precision-Recall Curve: To evaluate the grounding aspect of GVQA, we use Precision-Recall (PR) curves. For methods yielding a confidence score for grounding, we compute precision and recall at various threshold levels and plot the corresponding

PR curve. This provides insights into the trade-offs between precision and recall across different settings, offering a detailed evaluation of grounding performance. For methods that directly output a segmentation mask, precision and recall are computed, and the corresponding point is plotted on the PR curve for comparison.

C Optimal Layer Selection

We use the hallucination detection task to discuss optimal layer selection, as its diverse categories make it ideal for evaluating robustness. The results in Table 1 were obtained using category-specific validation sets (250 examples per category). To test robustness, we conducted experiments with adversarial validation sets, where the category being tested was excluded from the validation data (1,500 examples: 250×6 categories). For InternLM-VL, we observed that Image layer embedding 13 and Text layer embedding 27 consistently ranked among the top two combinations across all categories. These results, summarized in Table 2, demonstrate that for all categories—except OCR—the performance closely matches that achieved with task-specific validation, indicating the robustness of layer selection. In the absence of task-specific data, a robust method for selecting optimal layers would be to identify those that rank highest during adversarial validation for available categories.

| Category | Task Specific | Adversarial |
|-------------|---------------|-------------|
| Action | 0.796 | 0.792 |
| Attribute | 0.825 | 0.833 |
| Comparison | 0.623 | 0.638 |
| Count | 0.885 | 0.896 |
| Environment | 0.811 | 0.813 |
| Relation | 0.755 | 0.752 |
| OCR | 0.772 | 0.744 |

Table 2: InternLM-VL mAP scores on HQH dataset for task-specific and adversarial validation.

D Additional Qualitative examples for Grounded Visual Question Answering

Figure 10 shows some scenarios in which our bounding box based multimodal attribution system could be used.



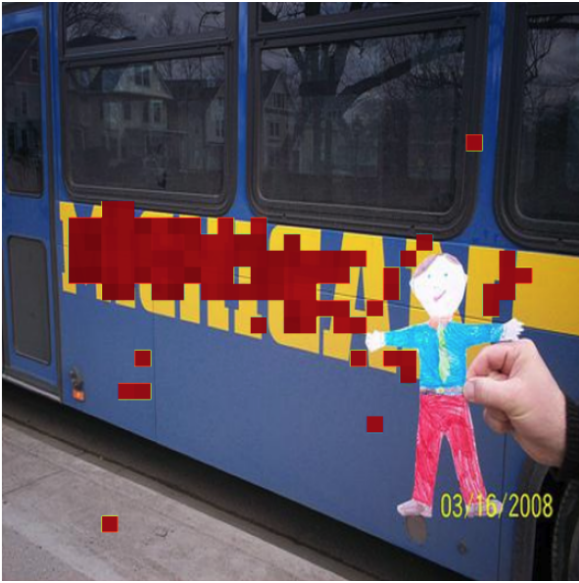
(a) Question: What is the color of the **doll**? Answer: Brown (b) Question: What is the color of the **pillow**? Answer: Brown

Figure 6: Top 20 image patches while detecting hallucination in *attribute* category using *ContextualLens*. We see that text tokens "Brown" are contextualized by the objects they refer to.



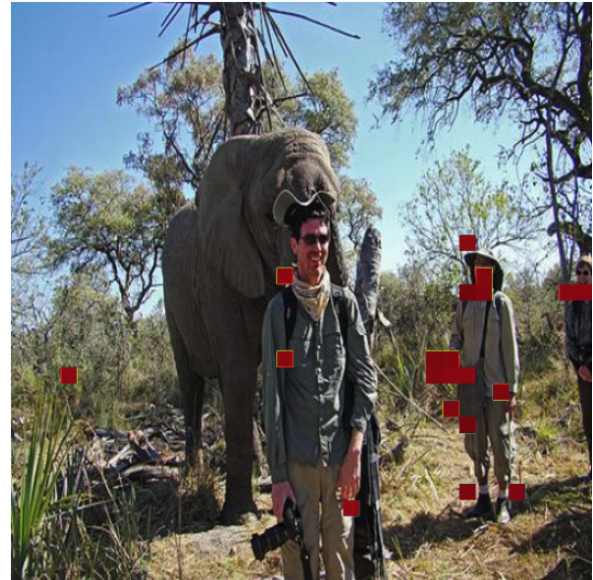
(a) Question: What is the **older** boy holding? Answer: Ball (b) Question: What is the **younger** boy holding? Answer: Ball

Figure 7: Top 5 image patches while detecting hallucination in *comparison* category using *ContextualLens*. We see that image patches corresponding to the younger boy's hand are also highlighted in the second case.



(a) Question: What is written on the bus? Answer: **Michigan** (b) Question: What is written on the bus? Answer: **New York**

Figure 8: Top 100 image patches while detecting hallucination in *OCR* category using *ContextualLens*. We see that image patches corresponding to multi-token "Michigan" are highlighted in the first case.



(a) Question: How many people are in this picture? Answer: **3** (b) Question: How many people are in this picture? Answer: **2**

Figure 9: Top 20 image patches while detecting hallucination in *count* category using *ContextualLens*. We see that image patches highlighted in both cases are similar and the score for the second case is higher ($0.891 < 0.892$).

Describe all the details of this image

Generate Answer

The image shows a highway directional sign in China. The sign indicates: 1. Three different routes: - Route S30 leading to Beijing and the Capital International Airport (首都机场) - Route G2 leading to Beijing and Langfang (廊坊) - Route S40 leading to the Yin Hai New District (滨海新区) 2. An exit number, which is 57 (出口57). 3. The destination is 1 km ahead.




Generate Attribution Save To History Clear

Describe this image

Generate Answer

The image is of a cheque from the State Bank of India. It is **made out to** "Prateek Agrawal" for the amount of fifty lakh twenty-five thousand rupees only (₹50,25,000). The cheque is dated April 5, 2019, and is signed by "Deepak Choudhary." The cheque number is 230270.



Generate Attribution Save To History Clear

(a) The image shows a direction board with Chinese text. The (b) The image shows a hand-filled bank cheque. The red box highlights the Chinese phrase, linking it to the English phrase "made out to," which means "to pay someone."

When is the No-School day

Generate Answer

The No-School days are on Friday, the 7th, and Monday, the 10th.


| Monday | Tuesday | Wednesday | Thursday | Friday |
|--|--|--|--|---|
| 3 Egg with Cheese Taco* Applesauce* Milk | 4 French Toast Bake* Grapefruit Milk | 5 Scrambled Eggs* Wheat Biscuits Banana Milk | 6 Banana Pancakes* Turkey Bacon Oranges Milk Agave Nectar | 7 No School :) |
| 10 No School :) | 11 Cinnamon Toast Bagel Applesauce* Milk | 12 Chicken Patty Wheat Biscuits Banana Milk | 13 Wheat Pancakes* Turkey Bacon Oranges Milk Agave Nectar | 14 Cinnamon Eggs Taco* Assorted Fruit Milk |
| 17 Egg with Cheese Taco* Applesauce* Milk | 18 Peach Baked Crisp Grapefruit Milk | 19 Turkey Sausage Patty Wheat Biscuits Banana Milk | 20 Pumpkin Pancakes* Turkey Bacon Oranges Milk Agave Nectar | 21 Egg and Bacon Taco* Assorted Fruits Milk |
| 24 Egg with Cheese Taco* Applesauce* Milk | 25 Cinnamon Sugar Grapefruit Milk | 26 Scrambled Eggs* Wheat Biscuits Banana Milk | 27 Wheat Pancakes* Turkey Bacon Oranges Milk Agave Nectar | 28 Egg and Sausage Taco* Assorted Fruit Milk |
| 31 Scrambled Gobbler Eggs with milk hair Tacos* Ectoplasma Liquid Chees | Did You Know? Cantaloupe is called a lot of different names all over the world. In Australia they call it a rockmelon, but in South Africa they call it a sweet melon. | | Good Eats At: *Gluten Free *Dairy-Free *Peb *Peb *Apple Menu subject to change based on availability and costs. mailto:info@schoolnutrition.com kathryn@schoolnutrition.com | |

Generate Attribution Save To History Clear

Did she make notes for her trip

Generate Answer

Yes, she made **notes** for her trip.



Generate Attribution Save To History Clear

(c) The image shows a timetable sheet. The red box highlights the square with the schedule for the date corresponding to the school holiday. (d) The image shows a snapshot of luggage being packed. The red box highlights the notes associated with the packed luggage.

Figure 10: Additional qualitative examples for Grounded Visual Question Answering.