

# Reference-free Hallucination Detection for Large Vision-Language Models

Qing Li<sup>1\*</sup>, Jiahui Geng<sup>1\*</sup>, Chenyang Lyu<sup>1</sup>, Derui Zhu<sup>2</sup>, Maxim Panov<sup>1</sup>, Fakhri Karray<sup>1†</sup>

<sup>1</sup> Mohamed bin Zayed University of Artificial Intelligence

{qing.li, jiahui.geng, chenyang.lyu, maxim.panov, fakhri.karray}@mbzuai.ac.ae

<sup>2</sup> Technical University of Munich

derui.zhu@tum.de

## Abstract

Large vision-language models (LVLMs) have made significant progress in recent years. While LVLMs exhibit excellent ability in language understanding, question answering, and conversations of visual inputs, they are prone to producing hallucinations. While several methods are proposed to evaluate the hallucinations in LVLMs, most are reference-based and depend on external tools, which complicates their practical application. To assess the viability of alternative methods, it is critical to understand whether the reference-free approaches, which do not rely on any external tools, can efficiently detect hallucinations. Therefore, we initiate an exploratory study to demonstrate the effectiveness of different reference-free solutions in detecting hallucinations in LVLMs. In particular, we conduct an extensive study on three kinds of techniques: uncertainty-based, consistency-based, and supervised uncertainty quantification methods on four representative LVLMs across two different tasks. The empirical results show that the reference-free approaches are capable of effectively detecting non-factual responses in LVLMs, with the supervised uncertainty quantification method outperforming the others, achieving the best performance across different settings.

## 1 Introduction

Large vision-language models (LVLMs) like LLaVA (Liu et al., 2024), MiniGPT-4 (Zhu et al., 2024b) have demonstrated remarkable capabilities in understanding and generating complex visual and textual content. However, an emerging concern with these models is their tendency towards hallucination (Zhou et al., 2024; Zhu et al., 2024a; Geng et al., 2024b). For instance, a model might describe an object or event in the generated text that

\*Equal contribution.

†Corresponding author.

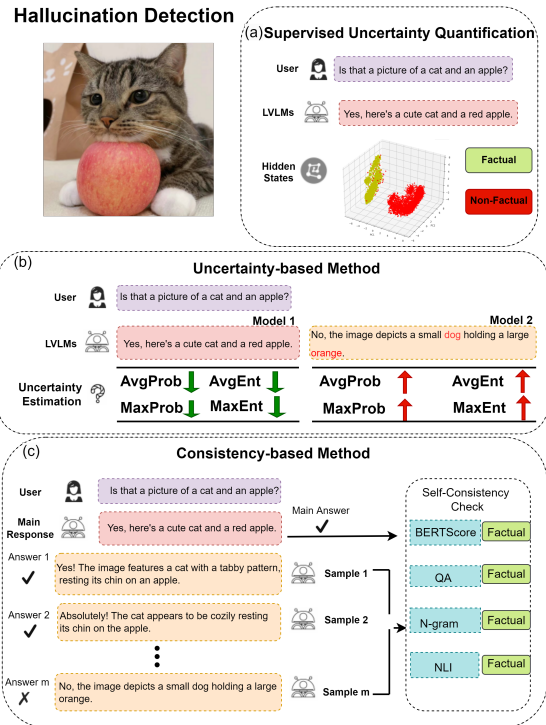


Figure 1: Reference-free Hallucination Detection Methods used in this work.

is not present in the input image. This phenomenon poses challenges for the reliability of LVLMs, highlighting intrinsic limitations of current models in maintaining coherence between visual inputs and textual descriptions (Huang et al., 2023a).

Most existing methods of assessing hallucinations rely on external models. Li et al. (2023) propose POPE, which requires automated segmentation tools like SEEM (Zou et al., 2024) to identify present and absent objects in images. Faithscore (Jing et al., 2023) utilizes a visual entailment model (Xie et al., 2019) to validate the factuality, focusing on object attributes and their relationships. Nonetheless, these approaches increase the computational cost and are restricted by the capabilities of these models.




Figure 2: Two data sources: self-generated and hand-crafted data. The test model assigns a probability of 0.9 to “cat” in self-generated input, while in hand-crafted input, “dog” receives a probability of 0.1.

Reference-free methods are a promising approach to address this issue, primarily utilizing the model’s own knowledge. They have attracted great research interest in Large Language Models (LLMs), with typical methods including the uncertainty-based methods (Hu et al., 2023; Geng et al., 2024a; Huang et al., 2023b; Vazhentsev et al., 2023; Fadeeva et al., 2023), consistency-based methods (Manakul et al., 2023) and supervised uncertainty quantification (SUQ) approach (Chen et al., 2023; Azaria and Mitchell, 2023). However, their potential in the context of LVLMs is unknown. They have not yet been systematically studied, possibly due to the complexity of multimodal data. We are committed to bridging this gap.

In this work, we collect and implement a battery of state-of-the-art methods, including four metrics of uncertainty-based methods, four variants of consistency-based methods, and the supervised uncertainty quantification method. The experiments are conducted on five LVLMs with different types, versions, and sizes on Yes-and-No and Open-ended tasks. The main contributions of this paper are summarized as follows:

1. We comprehensively measure the performance of different reference-free approaches in detecting hallucination.
2. We demonstrate that the supervised uncertainty quantification (SUQ) method performs best across various settings.
3. We contribute the Image-Hallucination Annotation Dataset (IHAD), a manually annotated, sentence-level dataset created by prompting LLaVA-v1.5-7b.

## 2 Reference-free Hallucination Detection Methods

**Uncertainty-based methods** have the hypothesis that when LVLMs are uncertain about generated information, generated tokens often have higher

Method	Mini-GPT-4v	LLaVA-v1.5-7b	LLaVA-v1.6-vicuna-7b	LLaVA-v1.6-mistral-7b
Random	48.73	48.73	48.73	48.73
<b>Adv</b>				
Uncertainty-based Methods				
AvgProb	77.43	86.82	91.80	92.04
AvgEnt	48.89	50.05	48.90	48.92
Supervised Uncertainty Quantification Method				
SUQ	87.23	92.35	93.32	93.50
<b>Ran</b>				
Baseline	48.73	48.73	48.73	48.73
Uncertainty-based Methods				
AvgProb	86.62	95.89	96.23	96.99
AvgEnt	48.76	50.25	49.91	50.41
Supervised Uncertainty Quantification Method				
SUQ	94.44	97.86	98.68	98.33
<b>Pop</b>				
Baseline	48.73	48.73	48.73	48.73
Uncertainty-based Methods				
AvgProb	82.11	92.73	93.57	95.06
AvgEnt	48.50	49.76	50.63	51.39
Supervised Uncertainty Quantification Method				
SUQ	91.02	97.85	98.11	98.50
<b>GQA</b>				
Baseline	48.47	48.47	48.47	48.47
Uncertainty-based Methods				
AvgProb	66.87	78.45	80.31	79.15
AvgEnt	48.75	49.88	48.95	49.33
Supervised Uncertainty Quantification Method				
SUQ	69.14	82.84	84.57	83.25

Table 1: AUC-PR for Yes-or-No tasks.

uncertainty, shown in Figure 1(b). Uncertainty-based methods are extensively studied for hallucination detection in LLMs (Hu et al., 2023; Geng et al., 2024a; Huang et al., 2023b; Fadeeva et al., 2024), and are also used as benchmark techniques (Zhu et al., 2024a; Manakul et al., 2023). We adopt the four proposed metrics: AvgProb, AvgEnt, MaxProb, MaxEnt to aggregate token-level uncertainty to measure sentence-level and passage-level uncertainty. Appendix A.1 provides detailed descriptions of these metrics.

**Consistency-based methods** mainly include self-consistency, cross-question consistency, and cross-model consistency (Zhang et al., 2023). In this work, we focus on four variants – BERTScore, Question Answering (QA), Unigram, Natural Language Inference (NLI) – of the self-consistency method proposed in (Manakul et al., 2023) to detect non-factual information for LVLMs, as demonstrated in Figure 1(c). The hallucination score for the  $i$ -th sentence, denoted by  $\mathcal{S}(i)$ , ranges from 0.0 to 1.0, where a score close to 0.0 indicates the sentence is accurate, and a score near 1.0 suggests the sentence is hallucinated. More detailed information can be found in Appendix A.2.

Method	Mini-GPT-4v	LLaVA-v1.5-7b	LLaVA-v1.6-vicuna-7b	LLaVA-v1.6-mistral-7b	LLaVA-v1.6-vicuna-13b
Baseline	29.29	29.29	29.29	29.29	29.29
Uncertainty-based Methods					
AvgProb	30.52	37.22	38.61	37.54	43.30
AvgEnt	29.92	32.18	39.13	37.22	40.11
MaxProb	33.45	32.05	34.86	34.73	40.03
MaxEnt	29.96	29.71	39.71	39.62	33.14
Consistency-based Methods					
QA	35.08	40.27	41.44	39.48	46.38
BERTScore	37.21	39.13	51.16	51.64	54.12
Unigram (max)	37.63	38.68	39.62	40.08	49.23
NLI	42.14	44.09	52.39	51.91	57.36
Supervised Uncertainty Quantification Method					
SUQ	<b>62.53</b>	<b>61.23</b>	<b>65.26</b>	<b>64.78</b>	<b>70.12</b>
M-Hal					
Baseline	35.71	35.71	35.71	35.71	35.71
Uncertainty-based Methods					
AvgProb	36.16	42.35	50.88	50.75	53.42
AvgEnt	35.81	37.18	48.38	46.86	50.04
MaxProb	35.97	41.89	46.24	45.96	45.21
MaxEnt	35.53	35.88	50.85	46.65	50.33
Consistency-based Methods					
QA	41.34	54.72	48.15	49.19	58.12
BERTScore	40.29	42.37	49.03	47.87	50.33
Unigram (max)	39.26	39.06	43.13	44.10	49.79
NLI	49.59	50.80	50.94	58.41	62.58
Supervised Uncertainty Quantification Method					
SUQ	<b>62.73</b>	<b>63.85</b>	<b>65.74</b>	<b>64.17</b>	<b>69.19</b>
IHAD					
Baseline	83.00	83.00	83.00	83.00	83.00
Uncertainty-based Methods					
AvgProb	84.29	87.50	87.38	87.63	90.23
AvgEnt	83.06	85.64	88.41	87.45	89.41
MaxProb	84.17	89.40	89.89	89.39	91.20
MaxEnt	83.03	87.18	88.21	88.44	92.10
Supervised Uncertainty Quantification Method					
SUQ	<b>85.34</b>	<b>94.08</b>	<b>93.09</b>	<b>94.07</b>	<b>96.03</b>

Table 2: AUC-PR for sentence-level Open-ended tasks.

**Supervised uncertainty quantification (SUQ).** To analyze and understand these internal states, SUQ method (Chen et al., 2023; Azaria and Mitchell, 2023) shown in Figure 1(a) train a classifier, referred to as a *probe*, on a dataset containing labeled examples. This classifier outputs the likelihood of a statement being truthful by analyzing the hidden layer activations of the LVLm during the reading or generation of states.

### 3 Study Design

**Dataset.** Our experiments are conducted on public datasets: POPE (Li et al., 2023), GQA (Hudson and Manning, 2019), and M-HalDetect (Gunjal et al., 2024). POPE contains three subsets: random, popular, and adversarial. For simplicity, we refer to them as Ran, Pop, and Adv, and we also abbreviate M-HalDetect as M-Hal. In addition, we create a manually annotated, sentence-level dataset, IHAD, by using LLaVA-v1.5-7b to generate concise image descriptions for 500 images from MSCOCO (Lin et al., 2014) using prompt “Provide a concise de-

Method	Mini-GPT-4v	LLaVA-v1.5-7b	LLaVA-v1.6-vicuna-7b	LLaVA-v1.6-mistral-7b	LLaVA-v1.6-vicuna-13b
Baseline	64.05	64.05	64.05	64.05	64.05
Uncertainty-based Methods					
AvgProb	70.28	72.78	72.89	73.32	81.23
AvgEnt	64.89	64.44	70.83	70.05	72.97
MaxProb	66.68	68.97	72.67	70.37	73.56
MaxEnt	68.75	64.99	77.17	78.62	80.23
Supervised Uncertainty Quantification Method					
SUQ	<b>84.92</b>	<b>89.41</b>	<b>85.91</b>	<b>86.54</b>	<b>91.22</b>
M-Hal					
Baseline	83.00	83.00	83.00	83.00	83.00
Uncertainty-based Methods					
AvgProb	84.29	87.50	87.38	87.63	90.23
AvgEnt	83.06	85.64	88.41	87.45	89.41
MaxProb	84.17	89.40	89.89	89.39	91.20
MaxEnt	83.03	87.18	88.21	88.44	92.10
Supervised Uncertainty Quantification Method					
SUQ	<b>85.34</b>	<b>94.08</b>	<b>93.09</b>	<b>94.07</b>	<b>96.03</b>
IHAD					
Baseline	83.00	83.00	83.00	83.00	83.00
Uncertainty-based Methods					
AvgProb	84.29	87.50	87.38	87.63	90.23
AvgEnt	83.06	85.64	88.41	87.45	89.41
MaxProb	84.17	89.40	89.89	89.39	91.20
MaxEnt	83.03	87.18	88.21	88.44	92.10
Supervised Uncertainty Quantification Method					
SUQ	<b>85.34</b>	<b>94.08</b>	<b>93.09</b>	<b>94.07</b>	<b>96.03</b>

Table 3: AUC-PR for passage-level Open-ended tasks.

scription of the given image”. Ran, Pop, Adv, and a subset of GQA used in this work evaluate object hallucination as a binary classification task, prompting LVLms to output “Yes” or “No”, categorizing them as Yes-or-No tasks. In contrast, M-Hal and IHAD involve long text answers based on open-ended questions, making them Open-ended tasks. More detailed dataset information is shown in Appendix B.1.

**Models.** We evaluate five representative categories of LVLms: MiniGPT-4v (Zhu et al., 2024b), LLaVA-v1.5-7b (Liu et al., 2023), LLaVA-v1.6-vicuna-7b, LLaVA-v1.6-mistral-7b (Liu et al., 2024) and LLaVA-v1.6-vicuna-13b. These models, featuring different architectures and strong performance, are used as benchmarks in various works (Huang et al., 2023a; Zhou et al., 2024). We only test models of a 7-billion and 13-billion size due to the limited computational resources.

**Experimental Setup.** We combine questions and responses to create inputs for test models. If the response is from this test model, it is considered self-generated input; otherwise, it is regarded as hand-crafted input for this model. Inputs are then processed by test models, from which we extract internal information, including probability, entropy, and embedding states. In this work, all datasets are hand-crafted except for IHAD, which is self-generated data only for LLaVA-v1.5-7b. Figure 2 shows test model assigns different internal information to two different inputs. Detailed experimental setups are presented in Appendix B.2.

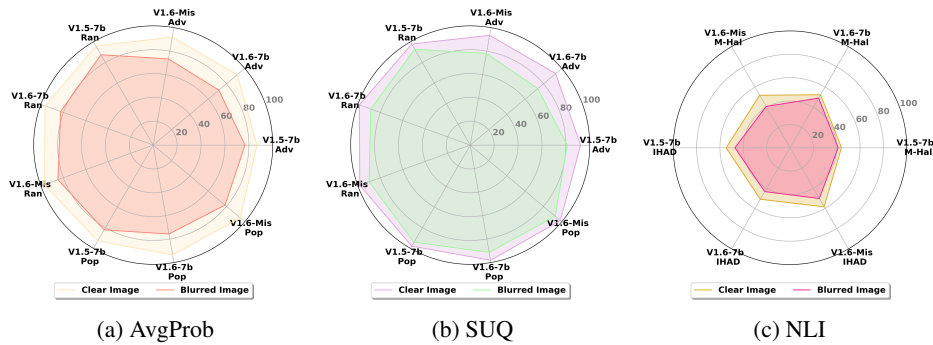


Figure 3: Comparison using images with different clarity.

## 4 Experimental Results and Analysis

The comprehensive results on Open-ended tasks are presented in Tables 1, 2 and 3. The numbers in the pink and yellow background, respectively, represent the best results for the uncertainty-based methods and consistency-based methods, while the numbers in bold indicate the overall best results. We utilize AUC-PR (Davis and Goadrich, 2006), which stands for the area under the precision-recall curve, to objectively evaluate the effectiveness of each model. The higher the value of AUC-PR, the stronger the ability of this method for hallucination detection.

### 4.1 Effectiveness of Various Methods

**Hallucination Detection on Yes-or-No Tasks.** We only compare the uncertainty-based and SUQ methods on Yes-or-No tasks because the model’s answers exclusively contain “Yes” or “No”. The random baseline is calculated as the ratio of non-factual examples to the total number of examples. Our results in Table 1 demonstrate that *SUQ method outperforms the uncertainty-based methods, exceeding AvgProb by about 4.5%. AvgProb achieves AUC-PR values nearly double those of the random baseline across all datasets and models, indicating it is an efficient method since it is an unsupervised method.* In addition, the probability Prob overall performs better than the entropy Ent in Yes-or-No tasks.

**Sentence-level Hallucination Detection.** Based on the data in Table 2, we find that *in most instances, the consistency-based approaches surpasses the uncertainty-based metrics, yet they remain significantly inferior to the SUQ method by approximately 15%. Among consistency-based methods, NLI outperforms best while Unigram shows the least effectiveness in most cases. Specifically, on the M-Hal dataset, the NLI performance*

of LLaVA-v1.6-7b exceeds the highest performing uncertainty metric, MaxEnt, by approximately 12%. BERTScore and QA outperform the uncertain estimation in most setups.

**Passage-level Hallucination Detection.** We calculate the metrics of uncertainty-based methods across all tokens in a passage and take the embedding information of the passage’s last token for SUQ. Consistency-based methods are not used to detect non-factual information in passages as the heavy computation. Our results in Table 3 show that *the SUQ method is still better than the uncertainty-based methods, including probability and entropy methods. Besides, uncertainty-based methods are weakly better than the random baseline at passage-level.*

### 4.2 Robustness and Limitations of SUQ

To investigate the robustness of SUQ methods, we employ two classifiers to assess the impact on each dataset. Each dataset is split into a training set and a test set in a 3:1 ratio. One classifier is trained on the Pop training set, while the other is trained on the training set of each dataset individually. Figure 4 illustrates that *the classifier trained on Pop exhibits comparable performance to classifiers trained on separated datasets, indicating the robustness of the SUQ method.*

Although our experiments confirmed the effectiveness of the SUQ method, it was found to be ineffective in detecting subtle, manually crafted hallucinations. For example, the difference between a correct statement, “There is a red apple and a cute kitten.” and an incorrect one, “There is a red orange and a cute kitten.” has a negligible effect on the model’s hidden state at the final position, posing a challenge for SUQ methods to differentiate between them.




Figure 4: Comparison of using different classifiers.

### 4.3 Impact of Image Clarity

To explore this question, we employ the Gaussian blur technique, which smooths the image by assigning a weighted average to the surrounding pixels of each pixel, resulting in an obscure picture. We set radius = 10 in our experiments. Figure 3 shows the effectiveness of AvgProb, SUQ, and NLI in blurred images are significantly reduced compared to clear images. The phenomenon suggests that *enhancing image clarity is a crucial strategy for boosting the effectiveness of hallucination detection methods.*

### 4.4 Impact of Data Source

Reference-free methods, particularly uncertainty-based and SUQ methods, rely on the internal information of models. A natural question is whether these methods are affected by data sources. M-Hal and IHAD are hand-crafted data and self-generated data, respectively, for LLaVA-v1.5-7b. Therefore, we compare various detection methods on M-Hal and IHAD based on LLaVA-v1.5-7b. Specifically, we calculate  $\Delta\text{AUC-PR}$

$$\Delta\text{AUC-PR}_{i,j} = \text{AUC-PR}_{i,j} - \text{AUC-PR}_{i,\text{baseline}},$$

where  $i$  refers to M-Hal and IHAD,  $j$  denotes different approaches.

Figure 5 shows that in uncertainty-based approaches, some metrics outperform in M-Hal while others excel in IHAD, a trend also observed in consistency-based techniques. This phenomenon illustrates that *the reference-free hallucination detection methods are insensitive to the data source.*

**Others.** We have also explored the impact of unimportant tokens for uncertainty-based methods and the optimal hidden layers for SUQ. They are discussed respectively in Appendices C.1 and C.2.

## 5 Conclusion

This paper systematically studies reference-free approaches for detecting the hallucination of LLMs.




Figure 5: Comparison of using different data sources.

Our results indicate that the SUQ method performs best as a supervised method, but it relies on sufficient training data, which could limit its applicability in certain scenarios. Consistency-based methods outperform uncertainty-based methods; however, they require more computation and are more suitable for black-box models. Uncertainty-based methods are particularly well-suited for tasks similar to Yes-or-No questions.

### Limitations

Three major limitations are identified in this work. First, we did not compare with reference-based methods. The performance of reference-based methods depends on the performance of external models and is specific to certain tasks. In this article, we primarily aim to explore the potential of reference-free methods, which have a broader range of applications. Second, we did not classify types of hallucinations; limited by resources, even with fine-grained annotations, we only marked whether sentences contained hallucinations without distinguishing types such as object existence errors, attribute errors, text recognition errors, and fact-conflicting errors. This might be where SUQ methods have an advantage. Third, we only studied white-box models. Recent research indicates that white-box models can be proxies for black-box models to calculate their confidence regarding specific issues. Unigram method is one such proxy model approach. In the future, we could use more complex models, closer in architecture to the target models, as proxies to explore their potential in hallucination detection.

### Ethics and Broader Impact

We sampled a portion of the data from existing datasets for our experiments, which may affect the accuracy of some of our conclusions.

## References

- Amos Azaria and Tom Mitchell. 2023. [The internal state of an LLM knows when it’s lying](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 967–976, Singapore. Association for Computational Linguistics.
- Nuo Chen, Ning Wu, Shining Liang, Ming Gong, Linjun Shou, Dongmei Zhang, and Jia Li. 2023. [Beyond surface: Probing llama across scales and layers](#). *ArXiv preprint*, abs/2312.04333.
- Jesse Davis and Mark Goadrich. 2006. The relationship between precision-recall and roc curves. In *Proceedings of the 23rd international conference on Machine learning*, pages 233–240.
- Hanyu Duan, Yi Yang, and Kar Yan Tam. 2024. [Do llms know about hallucination? an empirical investigation of llm’s hidden states](#). *ArXiv preprint*, abs/2402.09733.
- Ekaterina Fadeeva, Aleksandr Rubashevskii, Artem Shelmanov, Sergey Petrakov, Haonan Li, Hamdy Mubarak, Evgenii Tsybalov, Gleb Kuzmin, Alexander Panchenko, Timothy Baldwin, Preslav Nakov, and Maxim Panov. 2024. [Fact-checking the output of large language models via token-level uncertainty quantification](#). In *Findings of the Association for Computational Linguistics: ACL 2024*. Association for Computational Linguistics.
- Ekaterina Fadeeva, Roman Vashurin, Akim Tsvigun, Artem Vazhentsev, Sergey Petrakov, Kirill Fedyanin, Daniil Vasilev, Elizaveta Goncharova, Alexander Panchenko, Maxim Panov, Timothy Baldwin, and Artem Shelmanov. 2023. [Lm-polygraph: Uncertainty estimation for language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023 - System Demonstrations, Singapore, December 6-10, 2023*, pages 446–461. Association for Computational Linguistics.
- Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepl, Preslav Nakov, and Iryna Gurevych. 2024a. [A survey of confidence estimation and calibration in large language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6577–6595, Mexico City, Mexico. Association for Computational Linguistics.
- Jiahui Geng, Yova Kementchedjheva, Preslav Nakov, and Iryna Gurevych. 2024b. [Multimodal large language models to support real-world fact-checking](#). *arXiv preprint arXiv:2403.03627*.
- Anisha Gunjal, Jihan Yin, and Erhan Bas. 2024. [Detecting and preventing hallucinations in large vision language models](#). In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, pages 18135–18143. AAAI Press.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. [Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Mengting Hu, Zhen Zhang, Shiwan Zhao, Minlie Huang, and Bingzhe Wu. 2023. [Uncertainty in natural language processing: Sources, quantification, and applications](#). *ArXiv preprint*, abs/2306.04459.
- Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming Zhang, and Nenghai Yu. 2023a. [Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation](#). *ArXiv preprint*, abs/2311.17911.
- Yuheng Huang, Jiayang Song, Zhijie Wang, Huaming Chen, and Lei Ma. 2023b. [Look before you leap: An exploratory study of uncertainty measurement for large language models](#). *ArXiv preprint*, abs/2307.10236.
- Drew A. Hudson and Christopher D. Manning. 2019. [GQA: A new dataset for real-world visual reasoning and compositional question answering](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6700–6709. Computer Vision Foundation / IEEE.
- Liqiang Jing, Ruosen Li, Yunmo Chen, Mengzhao Jia, and Xinya Du. 2023. [Faithscore: Evaluating hallucinations in large vision-language models](#). *ArXiv preprint*, abs/2311.01477.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Xin Zhao, and Ji-Rong Wen. 2023. [Evaluating object hallucination in large vision-language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 292–305, Singapore. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. [Microsoft coco: Common objects in context](#). In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. [Improved baselines with visual instruction tuning](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). In *Advances in Neural Information Processing Systems 36: Annual*

*Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023.*

Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. [SelfCheckGPT: Zero-resource black-box hallucination detection for generative large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9004–9017, Singapore. Association for Computational Linguistics.

Artem Vazhentsev, Gleb Kuzmin, Akim Tsvigun, Alexander Panchenko, Maxim Panov, Mikhail Burtsev, and Artem Shelmanov. 2023. [Hybrid uncertainty quantification for selective text classification in ambiguous tasks](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11659–11681, Toronto, Canada. Association for Computational Linguistics.

Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. 2019. [Visual entailment: A novel task for fine-grained image understanding](#). *ArXiv preprint*, abs/1901.06706.

Jiaxin Zhang, Zhuohang Li, Kamalika Das, Bradley Malin, and Sricharan Kumar. 2023. [SAC<sup>3</sup>: Reliable hallucination detection in black-box language models via semantic-aware cross-check consistency](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15445–15458, Singapore. Association for Computational Linguistics.

Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. 2024. [Analyzing and mitigating object hallucination in large vision-language models](#). In *The Twelfth International Conference on Learning Representations*.

Derui Zhu, Dingfan Chen, Qing Li, Zongxiong Chen, Lei Ma, Jens Grossklags, and Mario Fritz. 2024a. [PoLLMgraph: Unraveling hallucinations in large language models via state transition dynamics](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 4737–4751, Mexico City, Mexico. Association for Computational Linguistics.

Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2024b. [MiniGPT-4: Enhancing vision-language understanding with advanced large language models](#). In *The Twelfth International Conference on Learning Representations*.

Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Wang, Lijuan Wang, Jianfeng Gao, and Yong Jae Lee. 2024. [Segment everything everywhere all at once](#). *Advances in Neural Information Processing Systems*, 36.

## A Reference-free Hallucination Detection Methods

We collect uncertainty-based and consistency-based methods proposed in (Manakul et al., 2023) to test the effectiveness on LVLMs.

### A.1 Uncertainty-based methods

To aggregate the uncertainty information obtained at the token level, we employ four metrics to aggregate token-level uncertainty into sentence level. In particular, a sentence-level uncertainty score can be obtained by taking either the maximum or average of the negative loglikelihood  $-\log p_{ij}$  in a sentence:

$$\text{MaxProb}(i) = \max_j (-\log p_{ij}), \quad (1)$$

$$\text{AvgProb}(i) = -\frac{1}{J_i} \sum_{j=1}^{J_i} \log p_{ij}, \quad (2)$$

where  $p_{ij}$  is the probability of a token at a position  $j$  in the sentence  $i$  and  $J_i$  is the total number of tokens in the considered sentence. Additionally, one can also replace the negative loglikelihood  $-\log p_{ij}$  with the entropy  $\mathcal{H}_{ij}$ :

$$\text{MaxEnt}(i) = \max_j \mathcal{H}_{ij}, \quad (3)$$

$$\text{AvgEnt}(i) = \frac{1}{J_i} \sum_{j=1}^{J_i} \mathcal{H}_{ij}, \quad (4)$$

where  $\mathcal{H}_{ij}$  is the entropy of the token distribution for the  $j$ -th token in the sentence  $i$ .

### A.2 Consistency-based Methods

**BERTScore** denoted as  $\mathcal{S}_{BERT}$ , finds the average BERTScore of the considered sentence with the most similar sentence from each drawn sample. Let  $\mathcal{B}(\cdot, \cdot)$  denote the BERTScore between two sentences.  $\mathcal{S}_{BERT}$  is computed by  $\mathcal{B}(\cdot, \cdot)$  as:

$$\mathcal{S}_{BERT}(i) = 1 - \frac{1}{N} \sum_{n=1}^N \max_k \mathcal{B}(r_i, s_k^n), \quad (5)$$

where  $r_i$  represents the  $i$ -th sentence in main response  $R$ ,  $s_k^n$  represents the  $k$ -th sentence in the  $n$ -th sample  $S^n$  and  $N$  is the total number of sampled sentences.

**Question Answering (QA).** The multiple-choice question answering generation (MQAG) evaluates consistency by creating multiple-choice questions

from the primary generated response and answering them using other sampled responses as reference. MQAG consists of two stages: question generation  $G$  and question answering  $A$ . For the sentence  $r_i$  in the response  $R$ , we draw questions  $q$  and options  $\mathbf{o}$ :

$$q, \mathbf{o} \sim P_G(q, \mathbf{o} \mid r_i, R). \quad (6)$$

The answering stage  $A$  selects the answers:

$$a_R = \arg \max_k [P_A(o_k \mid q, R, \mathbf{o})], \quad (7)$$

$$a_{S^n} = \arg \max_k [P_A(o_k \mid q, S^n, \mathbf{o})]. \quad (8)$$

We compare whether  $a_R$  is equal to  $a_{S^n}$  for each sample in  $\{S^1, \dots, S^N\}$ , yielding  $N_m$  matches and  $N_n$  not-matches. A simple inconsistency score for the  $i$ -th sentence and question  $q$  based on the match/not-match counts is defined:  $S_{QA}(i, q) = \frac{N_n}{N_n + N_m}$ . Manakul et al. (2023) modify the inconsistency score  $S_{QA}(i, q)$  by incorporating the answerability of generated questions. Finally,  $S_{QA}(i)$  is the average of inconsistency scores across  $q$ ,

$$S_{QA}(i) = E_q[S_{QA}(i, q)]. \quad (9)$$

**Unigram.** The concept behind Unigram is to develop a new model that approximates the LVLMs by samples  $\{S^1, \dots, S^N\}$  and get the LVLM’s token probabilities using this model. As  $N$  increases, the new model gets closer to LVLMs. Due to time and cost constraints, we just train a simple  $n$ -gram model using the samples  $\{S^1, \dots, S^N\}$  as well as the main response  $R$ . We then compare the average and maximum of the negative probabilities of the sentence in response  $R$  using the following equations:

$$S_{n\text{-gram}}^{\text{Avg}}(i) = -\frac{1}{J_i} \sum_{j=1}^{J_i} \log \hat{p}_{ij}, \quad (10)$$

$$S_{n\text{-gram}}^{\text{Max}}(i) = \max_j (-\log \hat{p}_{ij}), \quad (11)$$

where  $\hat{p}_{ij}$  is the probability of a token at position  $j$  of a sentence  $i$ .

**Natural Language Inference (NLI)** determines whether a hypothesis follows a premise, classified into either entailment/neutral/contradiction. In this work, we use DeBERTa-v3-large (He et al., 2023) fine-tuned to MNLI as the NLI model. The input for NLI classifiers is typically the premise concatenated to the hypothesis, which for NLI is the

sampled passage  $S^n$  concatenated to the sentence to be assessed  $r_i$  in the response  $R$ . Only the logits associated with the ‘entailment’ and ‘contradiction’ classes are considered,

$$P(\text{contradict} \mid r_i, S^n) = \frac{\exp(z_e^{i,n})}{\exp(z_e^{i,n}) + \exp(z_c^{i,n})},$$

where  $z_e^{i,n} = z_e(r_i, S^n)$  and  $z_c^{i,n} = z_c(r_i, S^n)$  are the logits of the ‘entailment’ and ‘contradiction’ classes. NLI score for sentence  $r_i$  on samples  $\{S^1, \dots, S^N\}$  is then defined as,

$$S_{\text{NLI}}(i) = \frac{1}{N} \sum_{n=1}^N P(\text{contradict} \mid r_i, S^n). \quad (12)$$

## B Dataset and Implementation

### B.1 Dataset

**POPE** (Li et al., 2023) formulates the evaluation of object hallucination as a binary classification task that prompts LVLMs to output “Yes” or “No”, e.g., “Is there a chair in the figure?” POPE adopts three sampling settings to construct negative samples: random, popular, and adversarial. The random setting selects non-present objects at random. In contrast, the popular setting chooses from a list of frequently occurring but absent objects, and the adversarial method selects based on common co-occurrence in contexts despite the absence in the target image.

**GQA** (Hudson and Manning, 2019) is proposed to detect hallucinations, including existence, attribute, and relation. It contains not only Yes-or-No problems but also Open-ended problems. We randomly 9050 some binary problems for 397 images used in our experiments.

**M-HalDetect** (Gunjal et al., 2024) consists of 16k fine-grained annotations on visual question-answering examples, including object hallucination, descriptions, and relationships. We utilize about 3000 questions on 764 figures for our Open-ended tasks, averaging 4 sentences per response.

**IHAD** serves as the evaluation of description for Open-ended tasks. It is a sentence-level dataset created by prompting LLaVA-V1.5-7b to generate concise image descriptions for 500 images from MSCOCO (Lin et al., 2014). On average, each response contains 5 sentences. The creation of IHAD aims to investigate if the data source impacts the internal information of LVLMs and further impacts hallucination detection methods, which rely on this internal information.



Method	Mini GPT-v4	LLaVA v1.5-7b	LLaVA v1.6-vicuna-7b	LLaVA v1.6-mistral-7b	
Baseline	48.73	48.73	48.73	48.73	
Adv	Uncertainty-based Methods				
	AvgProb	55.87	58.53	54.27	81.70
	AvgEnt	51.74	50.87	48.82	49.41
	MaxProb	32.28	36.48	32.11	41.78
	MaxEnt	49.03	50.60	48.35	49.25
Ran	Baseline	48.73	48.73	48.73	48.73
	Uncertainty-based Methods				
	AvgProb	54.57	51.95	47.35	85.10
	AvgEnt	50.49	51.13	53.51	50.80
	MaxProb	39.80	33.77	30.68	41.30
Pop	Baseline	48.73	48.73	48.73	48.73
	Uncertainty-based Methods				
	AvgProb	55.87	57.65	51.52	81.74
	AvgEnt	51.74	50.59	51.83	51.92
	MaxProb	32.28	35.90	31.41	40.97
GQA	Baseline	48.47	48.47	48.47	48.47
	Uncertainty-based Methods				
	AvgProb	55.72	56.56	57.93	65.06
	AvgEnt	48.96	50.12	49.55	48.38
	MaxProb	33.20	39.64	37.27	46.25
MaxEnt	48.16	49.98	49.15	48.39	

Table 4: Yes-or-No tasks with period. The terms **top** denote the highest AUC-PR values for uncertainty-based methods.

## B.2 Experimental Setup

In our experiments, we merge questions and responses from humans or other models to create hand-crafted input. This input is then processed by the evaluation model, from which we extract internal information, including probability, entropy, and embedding states. The evaluation model, capable of recognizing hallucinations (Duan et al., 2024), can assign reasonable internal information to responses not generated by itself. As shown in Figure 2, when generating the “cat” token, the evaluation model assigns probabilities to all tokens and selects the higher probability token for output. Specifically, the probability assigned to “cat” in self-generated input is 0.9, while in hand-crafted input, “dog” receives a probability of 0.1. POPE, GQA, and M-Hal are hand-crafted data for all four evaluation models used in this work. IHAD’s response, however, is generated by LLaVA-v1.5-7b, making IHAD self-generated data specifically for LLaVA-v1.5-7b. For the other three models, IHAD is considered hand-crafted data.

In the Yes-or-No tasks, we simplify the origi-

nal responses to “Yes” for positive answers (e.g., “Yes, there is a cute cat.”) and “No” for all others. To balance positive and negative samples, we create two entries for each original datum, one with a “Yes” response and the other with a “No” response. Thus, the original datasets Adv, Ran, and Pop, each containing 3000 questions for 500 images, are expanded to 6000 questions per set for the same images. Similarly, GQA is increased to 9050 questions for 397 figures. Regarding Open-ended tasks, each response contains an average of 4 sentences in M-HalDetect, while 5 sentences in IHAD.

In the consistency-based method, we set the temperature to 0.0 to obtain the main resource and use standard beam search decoding. In addition, we set the temperature to 1.0 for the stochastically generated samples and generated  $N = 10$  samples.

The classifier used in the SUQ method employs a feedforward neural network featuring three hidden layers with decreasing numbers of hidden units (256, 128, 64), all utilizing ReLU activations. The final layer is a sigmoid output. We use the Adam optimizer. We do not fine-tune any of these hyperparameters for this task. The classifier is trained for 20 epochs. We use about two-thirds of a dataset to train a classifier based on a specific model and then test its accuracy on the other third on the same model. The trained well classifier is required to determine which sentences the LVLm “believes” are true and which it “believes” are false.

## C Experimental Results

### C.1 Impact of Unimportant Tokens for Uncertainty-based Methods

It is observed that the performance of uncertainty methods on the Open-ended tasks is not as good as on the Yes-or-No tasks. To explore the reasons, we append a period, resulting in “Yes.” or “No.” in the Yes-or-No task. In this case, probability  $p$  and entropy  $\mathcal{H}$  of two tokens need to be considered. Table 4 shows the detailed performance for Yes-or-No tasks with period using different methods. Specifically, for LLaVA-v1.6-mistral, the AvgProb scores decline to 81.70 from 92.04 for Adv, to 85.10 from 96.99 for Ran, to 81.74 from 95.06 for Pop, and to 65.06 from 79.15 for GQA datasets, indicating the performance of AvgProb is severely affected by unimportant tokens. The other three metrics, AvgEnt, MaxProb, and MaxEnt, nearly equal to the baseline and play no role in hallucination detec-

tion.

## C.2 Optimal Hidden Layers for Effective SUQ Method

To explore which hidden layer information should be used, we perform 3D TSNE visualizations of the embedding states of LLaVA at the fifth and last layers for popular tasks, as shown in Figure 6. These visualizations show that at the last layer, there are distinct separations between the representations of different labels. Conversely, at the fifth layer, representations of different labels are relatively close and almost blend together, indicating that the shallow layers primarily preserve language-agnostic features. Therefore, we use the last hidden layer to explore the SUQ method in our experiments.




Figure 6: 3D TSNE plot of language embeddings computed from the fifth and last layers of LLaVA-v1.6-7b on popular task.

## C.3 The comparison of clear and blurred images.

The comparison of clear and blurred images used in 4.3 is illustrated in Figure 7.




Figure 7: The comparison of clear and blurred images.