CERTAINLY UNCERTAIN: A BENCHMARK AND METRIC FOR MULTIMODAL EPISTEMIC AND ALEATORIC AWARENESS

Khyathi Raghavi Chandu[♣] * **Linjie Li**[♦♠] *
Anas Awadalla[♦] **Ximing Lu**[♦] **Jae Sung Park**[♦] **Jack Hessel**[♥]
Lijuan Wang[♠] **Yejin Choi**[♣ ♦]

[♣]Allen Institute for AI [♦]University of Washington [♥]Samaya AI [♠]Microsoft

ABSTRACT

The ability to acknowledge the inevitable uncertainty in their knowledge and reasoning is a prerequisite for AI systems to be truly truthful and reliable. In this paper, we present a taxonomy of uncertainty specific to vision-language AI systems, distinguishing between epistemic uncertainty (arising from a lack of information) and aleatoric uncertainty (due to inherent unpredictability), and further explore finer categories within. Based on this taxonomy, we synthesize a benchmark dataset, CERTAINLYUNCERTAIN, featuring $178K$ visual question answering (VQA) samples as contrastive pairs. This is achieved by 1) inpainting images to make previously answerable questions into unanswerable ones; and 2) using image captions to prompt large language models for both answerable and unanswerable questions. Additionally, we introduce a new metric *confidence-weighted accuracy*, that is well correlated with both accuracy and calibration error, to address the shortcomings of existing metrics. Despite the recent rapid progress in vision-language models (VLMs), evaluations on our benchmark show that they perform poorly in uncertain scenarios. Further experiments demonstrate that supervised fine-tuning with CERTAINLYUNCERTAIN enhances the performance of VLMs, and reduces the calibration error. These improvements extend beyond our benchmark to existing refusal-oriented datasets and show positive results on reducing hallucinations, while maintaining performance on standard VQA benchmarks. Our work underscores the importance of addressing uncertainty in vision-language AI systems to improve their reliability and trustworthiness in real-world applications.

1 INTRODUCTION

An AI system with intellectual integrity must know when to admit "I don't know", which, in turn, requires a sharp awareness of its own limitations of knowledge and reasoning, as well as the inherent uncertainty around the external world (Paul & Elder, 1992; Zhang et al., 2019; Favero et al., 2024; Zhou et al., 2023; Wang et al., 2023b; Varshney & Baral, 2023). However, current vision-language models (Davis, 2020; Miyai et al., 2024; Mahendru et al., 2017) do not exhibit such a sufficiently sharp awareness of their own mistakes, leading to overly confident, uncalibrated predictions (Whitehead et al., 2022) and hallucinations (Wang et al., 2023a; Li et al., 2023d). This is only as expected, however, given that the predominant training recipe (Liu et al., 2023b; 2024; Bai et al., 2023) does not typically encourage the models to express uncertainty or acknowledge when they do not know the answer. Rather, they are incentivized to make predictions regardless of their confidence level. Moreover, existing benchmarks focus mainly on scenarios where clear and definitive answers are available (Yue et al., 2023; Goyal et al., 2019), leaving a notable gap.

Motivated by these, we introduce CERTAINLYUNCERTAIN, a dataset of approximately 178K visual question answering (VQA) instances that encompass various types of uncertainties. CERTAINLYUNCERTAIN is based on a novel taxonomy of multimodal uncertainty comprising *epistemic*

*Equal contribution

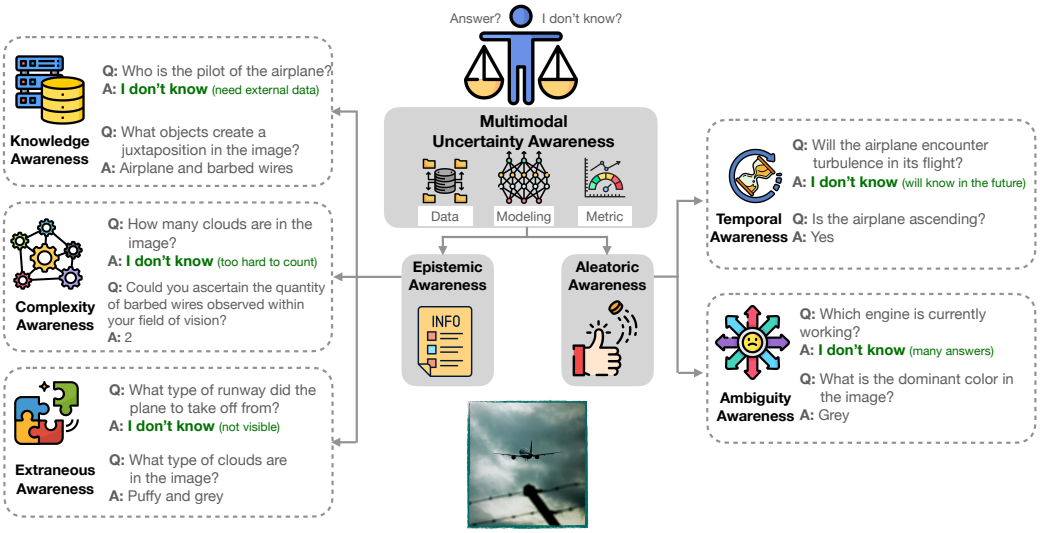

Figure 1: CERTAINLYUNCERTAIN: Taxonomy of uncertainty awareness in multimodal reasoning. For simplicity, all answers in this figure are normalized to "I don't know", in practice the answers are more diverse, as "I don't know" can be expressed in various ways.

uncertainty (due to lack of information) and *aleatoric* uncertainty (due to inherent unpredictability), as illustrated in Figure 1. We further define more fine-grained sub-categories, including (i) Knowledge, requiring external knowledge not explicitly captured by the image; (ii) Complexity, where the question is too complex to yield an exact answer; (iii) Extraneous, where parts of the necessary context or details are missing from the image; (iv) Temporal, where future events implied by the image cannot be predicted with absolute certainty; and (v) Ambiguity, where the question itself is unclear, leading to confusion or multiple possible interpretations. We construct CERTAINLYUNCERTAIN with two methods: 1) by masking and inpainting relevant image regions to render previously answerable questions unanswerable; and 2) by presenting GPT-4 (OpenAI, 2023) with image captions and prompting it to generate both answerable and unanswerable questions about the same image. Compared to prior datasets on unanswerability (Miyai et al., 2024; Guo et al., 2023), our dataset is constructed in a more systematic way, covering more diverse and finer-grained categories of uncertainty in vision-language scenarios.

With CERTAINLYUNCERTAIN, we empirically found that existing vision-language models rarely hesitate to answer even under uncertain conditions. In addition, they often overly confidently in providing an answer to unanswerable questions, while much less confident in admitting "I don't know". However, this issue is not reflected in popular metrics such as accuracy or F1, which do not account for model confidence. Alternative metrics, such as risk and coverage (Whitehead et al., 2022) use thresholding to binarize the equivalent of the prediction probability. The expected calibration error (ECE) (Naeini et al., 2015) evaluates the prediction probabilities but does not accurately reflect performance in terms of correctness. Therefore, we propose a new *confidence-weighted accuracy* metric, which incorporates model confidence into the accuracy computation. This metric addresses the shortcomings of existing metrics by simultaneously capturing both predictive performance and model confidence. Our proposed metric demonstrates a positive correlation with accuracy and a negative correlation with ECE.

Moreover, we conduct extensive experiments using 3 training strategies with CERTAINLYUNCERTAIN: supervised fine-tuning, R-tuning (Zhang et al., 2023), and preference optimization (Rafailov et al., 2023). We evaluate the resulting models across 7 datasets covering refusal, hallucination, and standard VQA tasks. Our empirical results show that fine-tuning with CERTAINLYUNCERTAIN not only improves performance on a held-out portion of our dataset and existing refusal-based datasets but also helps reduce hallucinations while maintaining performance on standard VQA tasks. These findings underscore the effectiveness of CERTAINLYUNCERTAIN in enhancing the robustness and reliability of vision-language models.

2 CERTAINLYUNCERTAIN

To train models to properly admit "I don't know", it is crucial to construct a large-scale dataset that covers a diverse range of uncertain situations. This is challenging, as most internet data focus on certain scenarios (*e.g.*, alt-text description for an image), thus not readily applicable. Therefore, we develop CERTAINLYUNCERTAIN, a dataset with approximately $178K$ VQA instances, using an automatic data synthesis pipeline for various types of uncertainty. We begin by introducing a taxonomy of uncertainties in §2.1, where admitting uncertainty rather than providing an answer is the appropriate response in each category. Next, we describe our data creation process in §2.2. Finally, we introduce the evaluation metrics in §2.3

2.1 TAXONOMY OF UNCERTAINTY AWARENESS

Depending on whether it is due to contextual inexpressiveness or genuine incapability to answer, we broadly categorize multimodal uncertainty into 2 types, epistemic and aleatoric uncertainty.

Epistemic Uncertainty refers to the uncertainty in a model's predictions that arises from a lack of knowledge or complete information about the system being modeled. It is due to the model's limited understanding or insufficient data, which can be reduced by gathering more information, improving the quality of data, or enhancing the model itself. This type of uncertainty highlights areas where the model's predictions may be less reliable due to the lack of sufficient evidence to make accurate inferences. We further categorize the awareness of epistemic uncertainty into 3 finegrained types:

- **Knowledge awareness** means understanding that some questions require information or common sense that is not shown in the image. For example, you might need specialized knowledge or up-to-date information from outside sources. Knowing when this extra information is needed helps avoid wrong answers.
- **Complexity awareness** is recognizing when a question is difficult because it involves many parts or is hard to understand. This difficulty can come from how the question is asked or from the effort needed to understand the context and details of the question.
- **Extraneous awareness** refers to the ability to identify and disregard elements within an image that are not relevant to the question at hand. This involves recognizing objects, attributes, or aspects that, while present in the image, do not contribute to answering the question.

Aleatoric Uncertainty is the inherent unpredictability in a system or process that cannot be reduced or eliminated. It arises from the fundamental randomness or chaotic nature of the task itself. For example, predicting the outcome of a coin toss involves intrinsic uncertainty because the result is inherently probabilistic and cannot be determined with certainty in advance. Similarly, we define 2 sub-categories under aleatoric uncertainty:

- **Temporal awareness** means understanding that we may not always have access to all relevant data required to predict specific outcomes with absolute certainty, especially when it involves reasoning about time. This includes events in the past or future that cannot be inferred from the image alone with absolute certainty. Recognizing the limitations of temporal reasoning helps manage expectations about the accuracy of predictions involving time-related aspects.
- **Ambiguity awareness** involves recognizing situations, objects, or individuals that can be understood, interpreted, or perceived in more than one way. Ambiguity introduces uncertainty and a lack of clarity, leading to multiple possible interpretations. While ambiguity can encourage exploration of different meanings or perspectives, it can also cause confusion. It is essential to be aware of the levels of certainty in ambiguous scenarios to avoid misinterpretation and errors.

2.2 DATASET CREATION

Based on the taxonomy, we construct CERTAINLYUNCERTAIN, comprising contrastive VQA pairs for each category described above. The statistics of our dataset are summarized in Table 1. The contrastive instances in CERTAINLYUNCERTAIN are derived from two sources: images and captions. For sourcing from images, the same question that is answerable for the original image is rendered unanswerable for the perturbed image. For sourcing from captions, we prompt GPT-4 (OpenAI, 2023) to generate both an answerable and an unanswerable question based on the same caption. Below, we describe the dataset creation pipeline in detail.

		Epistemic			Aleatoric		All
		Knowledge	Complex	Extraneous	Temporal	Ambiguous	
Train	# of images (Perturbed/Clean)	-/9.5K	-/9.6K	**38.2K**/38.2K	-/9.6K	-/9.6K	38.2K/47.9K
	# of questions (IDK/Non-IDK)	9.5K/9.5K	9.6K/9.6K	38.2K/38.2K	9.6K/9.6K	9.6K/9.6K	76.6K/76.6K
Test	# of images (Perturbed/Clean)	-/2.5K	-/2.5K	**2.3K**/2.5K	-/2.5K	-/2.5K	2.3K/7.3K
	# of questions (IDK/Non-IDK)	2.5K/2.5K	2.5K/2.5K	2.3K/2.5K	2.5K/2.5K	2.5K/2.5K	12.3K/12.5K
Total	# of images/questions	12.1K/24.2K	12.1K/24.2K	81.3K/81.3K	12.1K/24.2K	12.1K/24.2K	95.8K/178.1K

Table 1: Statistics of CERTAINLYUNCERTAIN. Our dataset contains 178K questions on 95.8K images for 5 types of uncertainties. Each IDK question is accompanied with a non-IDK question to highlight contrasts between certainty and uncertainty. For extraneous testing split, we perform quality check and filter out invalid ones. Numbers in bold highlight the new images we created through our data creation pipeline.

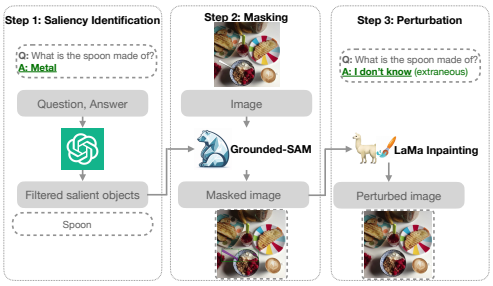

Figure 2: Pipeline for sourcing from images

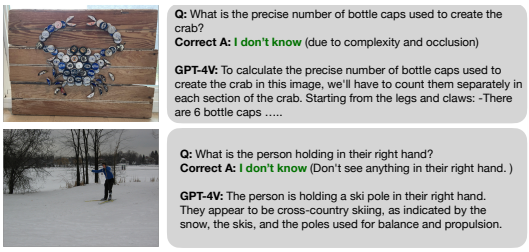

Figure 3: Uncertainty paradox in generative VLMs, where the question is generated from GPT-4/GPT-4V.

Sourcing from captions. We use detailed paragraph captions to prompt questions for each category of uncertainty. Each prompt includes a definition of the category along with examples of answerable and unanswerable questions and their answers. The captions are sourced from DOCCI (Onoe et al., 2024), which are high-quality human-annotated descriptions. These descriptions are highly compositional and include world knowledge, spatial relationships, visual settings, text rendering, and object attributes. For each caption, we prompt GPT-4 to generate both an answerable and an unanswerable question, along with their corresponding answers. In total, we collected around $110K$ instances on $\sim 15K$ images, spanning knowledge, complexity, temporal and ambiguity awareness categories. We follow the same train-test split as DOCCI to divide our dataset.

Sourcing from images. Compared to sourcing from captions, we perturb images to transform an answerable question into an unanswerable one. Our data generation pipeline has 3 main steps as outlined in Figure 2. The first step is saliency identification where the goal is to identify salient objects about which meaningful questions can be asked. For VQAv2 instances, we prompt GPT-4 to identify these salient objects based on VQAv2 questions, while for GQA instances, this information is available in the annotations. The second step is masking, where we use Grounded-SAM (Ren et al., 2024) to mask out the salient objects. The final step is perturbation where we use LaMa Inpainting model (Suvorov et al., 2021) to create a perturbed contrastive image so that the salient object is missing. To avoid any spurious biases from perturbation, we experimented with masking and inpainting randomly chosen objects instead of the salient object from the answerable subset, thereby keeping the question answerable. Since the performance did not fluctuate significantly, we proceeded without random perturbation. We then prompt GPT-4V to generate a question for each pair of images that is answerable for the original image but unanswerable for the perturbed image. To increase the difficulty of the questions, we specifically instruct GPT-4V to avoid generating "yes/no" questions, as they are more likely to be answerable. In the end, we created $\sim 30K$ samples based on VQAv2, which are split into $24K$ training and $6K$ testing samples. In addition, we leverage the GQA dataset (Hudson & Manning, 2019), which contains rule-based questions from ground truth scene graph annotations. Similarly, we perturb the images and alter the answers to "I don't know" to create unanswerable instances from the originally answerable instances. In total, we gather $53K$ more instances from GQA, and use them to augment the training split.

Generative AI Paradox for generating/understanding "uncertain questions". While LLMs and VLMs can generate uncertain questions, they often struggle to answer them accurately. As shown in Figure 3, where we prompt GPT-4V to answer its own generated uncertain questions and it fails. Inspired by West et al. (2023), which hypothesizes that models may not understand what they create, we observe a similar pattern in generating and understanding uncertain questions.

Benchmarks	Source	Dataset Size	Question Types			Types of Uncertainty/Unanswerability					
			OE	Free-form	IDK	Absurd	Knowledge	Complex	Extraneous	Temporal	Ambiguous
MM-Hal	Human	96	✓	✓	✓				✓		
POPE	Rule	9K									
AMBER Disc.	Rule	14K									
VizWiz	Human	33K	✓		✓				✓		
UNK_VQA	Human	10K	✓		✓	✓			✓		✓
TDIUC (Absurd)	Rule	336K	✓		✓	✓					
MM-UPD	Rule	2K			✓	✓					
Ours	LLM & Rule	178K	✓	✓	✓		✓	✓	✓	✓	✓

Table 2: Comparison of CERTAINLYUNCERTAIN to existing benchmarks. We mainly compare with two types of datasets: Hallucination-based datasets (top) and Refusal-based datasets (middle). CERTAINLYUNCERTAIN features 178K unanswerable (IDK) and answerable questions in open-ended (OE) setting with free-form answers, covering 5 types of finegrained types of unanswerability. Though our dataset does not explicitly cover absurd type, we show that it improves model performance on TDIUC (absurd) in experiments. Disc: Discriminative.

Contrastive pairs. In CERTAINLYUNCERTAIN construction process, we have images that are visually similar or the same but the question-image pairs are deliberately designed to highlight contrasts between certainty and uncertainty (as shown in Figure 2). This aids in improving model robustness by learning to distinguish between visually similar but semantically distinct instances leading to real-world applicability by exposing them to subtle variations and contrasts.

Quality check and filtering. As our data creation pipeline is model-dependent, though being efficient and saving the cost of human labor, it may suffer from model failures. Especially the pipeline to create the extraneous set, which depends on multiple models, the failure of one model at any stage (*e.g.*, the inpainting model fails to remove the object or the segmentation model fails to predict the correct mask of the intended object) may lead to invalid samples (*i.e.*, when the generated IDK question is still reasonably answerable for the image or vice versa). To ensure the data quality, we perform a final quality check on the extraneous testing set. Specifically, the image-question-answer tuples are presented to one of the authors, and the task is to validate whether the generated sample is valid or not. Among 6K samples, we filtered $\sim$1.2K samples, resulting in 4.8K testing examples. For other splits, we found the valid sample rate is much higher ($>$ 93% in Appendix Table 7). We retain the first 5K samples verified by human on DOCCI testing images to build the remaining testing splits.

Table 2 presents a comparison of CERTAINLYUNCERTAIN with existing benchmarks regarding data size, question types, and uncertainty categories. Our dataset is significantly larger and covers a wider variety of question types across diverse uncertainty categories. Notably, existing datasets such as UNK-VQA (Guo et al., 2023), or TDIUC (Kafle & Kanan, 2017) mainly focus on pairing unrelated questions with image contexts to create datasets of irrelevant and unanswerable questions. In contrast, our dataset ensures contextually aligned question-image pairs. With image inpainting, we also generate more natural-looking images compared to image masking or copying in UNK-VQA.

2.3 EVALUATION METRICS

Standard metrics. We report model performance on CERTAINLYUNCERTAIN with standard metrics, including accuracy and F1. For accuracy, we follow (Mañas et al., 2024) to use LLM-Assisted VQA Evaluation (LAVE), which leverage instruction-tuned large language models to rate candidate answers based on their correctness relative to the question and reference answers. LAVE have shown a strong correlation with human judgment. In practice, we use LAVE with Mistral-7B (Jiang et al., 2023) as the evaluator, comparing ground truth and predictions to assign scores of 0, 0.5, or 1. This allows for more nuanced evaluation compared to traditional string-matching approaches, especially in open-ended settings. To adapt LAVE to unanswerable settings, we introduce a dual-stage judging mechanism. This approach is more reliable because refusals or IDK responses can be expressed in various ways, such as simply stating IDK, asking a follow-up question, or offering a reasonable guess. The first stage is IDK normalization, where we use LAVE to determine if either the prediction or ground truth (GT) is IDK and normalize the answer to IDK. For refusal-based benchmarks, since the unanswerability of the question is annotated, we directly rely on the ground truth label for GT answers. The second stage is to award accuracy. If either the prediction or GT is normalized to IDK,

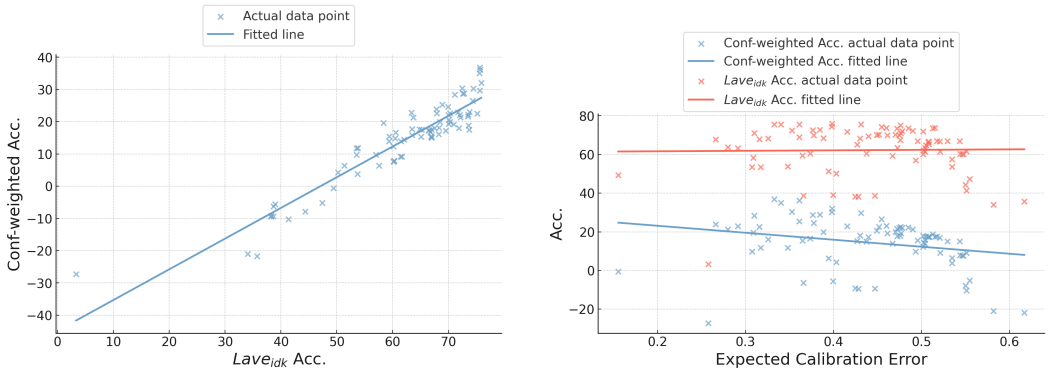

(a) Conf-weighted Acc. (ours) vs. LAVE_{idk} Acc.

(b) Ours and LAVE_{idk} Acc. vs. ECE.

Figure 4: Correlation of confidence weighted accuracy (↑) with LAVE_{idk} accuracy (↑) and ECE (↓). The datapoints in this plot are from evaluation results on extraneous split of different model variants in our experiments.

we compare the strings. Otherwise, we award the standard LAVE score. Formally, the LAVE_{idk} score is defined as

$$\text{LAVE}_{\text{idk}} = \begin{cases} \mathbb{1}(\text{pred}_{norm} == \text{GT}_{norm}) & \text{if LAVE(pred == IDK) or LAVE(GT == IDK)} \\ \text{LAVE(GT, pred)} & \text{else} \end{cases}. \quad (1)$$

In addition, we report F1_{idk} which is the F1 score only on the unanswerable questions.

Confidence-weighted accuracy. Existing evaluation metrics do not comprehensively assessing both the accuracy and the confidence of model predictions. Accuracy metrics, which score binarily, ignore the probability estimates associated with predictions. Expected Calibration Error (ECE), which measures the difference between predicted confidence levels and the true likelihood of those predictions being correct, do not provide a direct measure of final accuracy. Abstention metrics(Whitehead et al., 2022) – coverage and risk – do not address model accuracy. In addition, risk metric do not directly incorporate model confidence, instead threshold values to 0 or 1.

To address these issues, we introduce *Confidence-weighted accuracy* which weights the accuracy by the probability of the model's prediction. The desiderata of this metric is to remain positively correlated with accuracy while being negatively correlated with ECE. Thus, confidence-weighted accuracy takes into account the confidence of the model's prediction $P(\text{pred})$, providing a more holistic evaluation of performance. Based on the LAVE_{idk} accuracy above, we define confidence-weighted accuracy as

$$\textit{confidence weighted accuracy} = \mathbb{1}(\text{LAVE}_{\text{idk}} > 0) * \text{LAVE}_{\text{idk}} * P(\text{pred}) - \mathbb{1}(\text{LAVE}_{\text{idk}} == 0) * P(\text{pred}). \quad (2)$$

Similar to (Whitehead et al., 2022), we compute $P(\text{pred})$ by prompting the model to verify if its own predicted answer is correct and extracting the probability of the "yes" token. We normalize this probability by dividing it by the sum of the token probabilities for "yes" and "no". Our formulation penalizes incorrect predictions while rewarding correct ones, especially by encouraging higher confidence for correct predictions. As shown in Figure 4, 4a demonstrates the positive correlation of confidence-weighted accuracy with LAVE_{idk} accuracy, and 4b illustrates that our metric is more negatively correlated with ECE compared to LAVE_{idk} accuracy.

	VQA Score	Standard Acc.	F1	LLM-based Scores	ECE	Conf.-w Acc.
Acc.	✓	✓	✓	✓		✓
Conf.					✓	✓

Table 3: Our proposed Confidence-weighted Accuracy (Conf-w. Acc.) captures captures both accuracy (Acc.) and confidence (Conf.) of model predictions, in contrast to existing metrics either reflect accuracy or confidence.

3 EXPERIMENTS

3.1 EXPERIMENTAL DETAILS

We conduct experiments with the representative instruction-tuned models including variants of LLaVA (Liu et al., 2023b) -7B, 13B, 34B (Liu et al., 2024), and Qwen-VL (Bai et al., 2023), as well as evaluating the performance of GPT-4V on our CERTAINLYUNCERTAIN benchmark.

We investigate 3 training strategies with our data comparing them to the base model.

1. Supervised finetuning (SFT): To demonstrate the effectiveness of our data, we perform continued SFT with CERTAINLYUNCERTAIN with the instruction-tuned model and compare against LLaVA and LRV (Liu et al., 2023a) instruction-tuning datasets.

2. R-tuning: We follow (Zhang et al., 2023) to re-annotate ground truth answers that are incorrectly predicted by the base model to reflect IDK, and use this re-annotated refusal data for supervised fine-tuning.

3. Preference optimization: We directly adopt the two answers to the contrastive VQA pairs as the answer choices, and perform DPO (Rafailov et al., 2023). DPO is a method for aligning large language models (LLMs) with human preferences by directly optimizing a loss function based on preference data, eliminating the need for separate reward modeling or reinforcement learning steps. The DPO loss function adjusts the model's behavior by increasing the likelihood of preferred responses and decreasing that of less favored ones.

We also implement a inference-time baseline with naive selective prediction approach by marking predictions as IDK when the prediction probability falls below a threshold. We further explore adding CERTAINLYUNCERTAIN into the instruction-tuning stage for LLaVA directly after pre-training.

3.2 EVALUATION BENCHMARKS

We evaluate the models trained with our data on additional benchmarks, detailed below.

Refusal-based benchmarks. UNK-VQA (Guo et al., 2023) contains $10K$ instances of answerable and unanswerable questions constructed from manipulating VQA v2 instances by question perturbation like word replacement, semantic negation, and image perturbation like image replacement, object masking and copying a part of the object. We deliberately discard the *ambiguous* category from UNK-VQA as it was defined as having multiple plausible answers and simply listing them all should be correct. The "absurd" category of the TDIUC (Kafle & Kanan, 2017) data containing $\sim 366K$ questions is constructed by compiling a list of objects that are missing from a given image and then identifying questions from the rest of TDIUC that inquire about these absent objects. In our experiments, we randomly sample $5K$ instances from each dataset for evaluation.

Hallucination-based benchmarks. MMHal-Bench (Sun et al., 2023) contains 96 questions in 8 hallucination categories such as object attribute, adversarial object and counting. Upon establishing the severity of object hallucinations, Li et al. (2023d) introduce POPE, which includes $\sim 9K$ instances that sample objects randomly, adversarially, and based on popularity to check for their presence in a binary manner. This dataset aims to systematically evaluate the presence or absence of objects in images, thereby highlighting models' susceptibility to hallucinate objects that are not present. AMBER (Wang et al., 2023a) further extends POPE by evaluating not just existence but also attribute and relation hallucinations. By incorporating these additional dimensions, AMBER provides a more comprehensive framework, enabling deeper analysis of multimodal models' ability to generate coherent and accurate outputs across a range of hallucination scenarios.

Standard benchmarks. While mitigating hallucination and learning to refuse is important, the goal is also to not hurt model performance on standard datasets. Therefore, we conduct evaluations on standard datasets VQAv2 (Goyal et al., 2019) [1] and VizWiz (Gurari et al., 2018) validation splits.

[1] We randomly sample 5k questions from validation set that is not covered in LLaVA instruction-tuning data.

	Epistemic			Aleatoric			Total		
	LAVE_{idk} Metric		Conf-w.	LAVE_{idk} Metric		Conf-w.	LAVE_{idk} Metric		Conf-w.
	F1_{idk}	Acc.	Acc.	F1_{idk}	Acc.	Acc.	F1_{idk}	Acc.	Acc.
Qwen-VL-Chat	65.45	64.22	11.92	**67.15**	**63.35**	15.71	66.13	63.87	13.45
LLaVA-1.5-7B	51.31	44.72	-1.01	54.78	51.51	2.65	52.71	47.46	0.47
LLaVA-1.5-13B	52.38	46.14	2.70	53.35	50.46	1.81	52.78	47.88	2.34
LLaVA-1.6-7B	67.61	53.10	26.61	51.27	55.51	11.39	61.02	54.07	20.47
LLaVA-1.6-13B	69.72	66.88	28.07	54.61	56.72	14.29	63.63	62.78	22.52
LLaVA-1.6-34B	**74.37**	**71.06**	**40.03**	58.47	60.01	**21.27**	**67.96**	**66.60**	**32.47**
Qwen2-VL-7B	61.46	61.91	18.61	64.08	62.89	17.23	62.52	62.31	18.05
Qwen2-VL-72B	67.92	66.72	23.69	67.90	65.25	16.35	67.91	66.13	20.73
LLaVA-OV-7B	68.39	68.45	19.91	54.97	60.04	17.91	62.98	65.06	19.10
LLaVA-OV-72B	65.52	67.34	21.61	53.80	60.46	18.91	60.79	64.57	20.59
InternVL2-8B	64.79	57.54	14.84	50.61	44.99	0.45	59.07	52.48	9.04
InternVL2-40B	66.38	59.02	21.12	46.56	44.53	-0.54	58.39	53.18	12.39
InternVL2-76B	62.11	60.46	20.17	52.23	47.31	3.95	58.13	55.16	13.63
GPT-4V†	85.34	78.60	-	61.41	61.25	-	75.76	71.70	-
Claude-3.5 Sonnet†	75.46	72.92	-	52.20	53.25	-	66.19	65.10	-

Table 4: Evaluating existing VLMs on CERTAINLYUNCERTAIN. For epistemic/aleatoric category, we average the score across the 3/2 fine-grained categories. Total performance is averaged across all 5 fine-grained categories. † GPT-4V and Claude-3.5 Sonnet performance are reported on a smaller subset with 100 samples for each finegrained category.

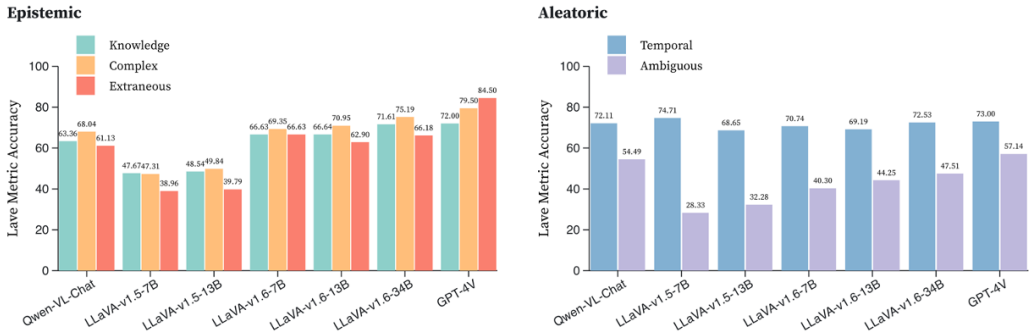

Figure 5: Breakdown of model performance on finegrained categories. We report LAVE_{idk} Metric Accuracy as the confidence of GPT-4V prediction is not accessible.

3.3 RESULTS AND DISCUSSION

As shown in Table 4, we observe that state-of-the-art VLMs including GPT-4V (despite the questions generated with it) perform poorly on our benchmark.[2] It is also worth noting that all models achieve significantly higher scores on LAVE_{idk} accuracy than confidence-weighted accuracy. This discrepancy suggests that the models are either over-confident in incorrect predictions or not confident enough in correct ones (*i.e.*, they are poorly calibrated), which we further examine in our fine-tuning experiments. It is worth noting that the more recent open-source VLMs, including Qwen2VL (Wang et al., 2024b), LLaVA-OneVision (Li et al., 2024) and InternVL2 (Chen et al., 2023a; 2024), despite the big leap in performance on standard benchmarks (Contributors, 2023), they perform similarly to or even worse than LLaVA-1.6 (Liu et al., 2024). These results further highlight our motivation: (1) existing benchmarks focus mainly on scenarios where clear and definitive answers are available; and (2) current VLM and the corresponding training recipe and data does not typically encourage the models to express uncertainty or acknowledge when they do not know the answer.

Figure 5 presents LAVE_{idk} accuracy of Qwen-VL-Chat Bai et al. (2023), LLaVA-1.5 Liu et al. (2023b), LLAVA-1.6 Liu et al. (2024) and GPT-V on each sub-category. The relative trends across sub-categories are consistent among the models. The extraneous category is the most challenging within epistemic uncertainty, while the ambiguous category is the hardest within aleatoric uncertainty. Performance on the temporal category is relatively similar across different model sizes, possibly due to the limited diversity of questions that can be asked about the future.

[2]Our empirical analysis of Claude 3.5 results suggest that Claude 3.5 is overly confident in terms of temporal and ambiguous categories (*e.g.*, sometimes its response is a reasonable guess, but it is stated with certainty), making its performance on Aleatoric awareness much worse than GPT4V.

		Epistemic			Aleatoric			Total			
		LAVE_{idk} Metric		Conf-w.	LAVE_{idk} Metric		Conf-w.	LAVE_{idk} Metric		Conf-w.	ECE ↓
		F1_{idk}	Acc.	Acc.	F1_{idk}	Acc.	Acc.	F1_{idk}	Acc.	Acc.	(IDK)
Qwen-VL-Chat⋆		65.45	64.22	11.92	67.15	63.35	15.71	66.13	63.87	13.45	0.79
Thresholding		74.44	69.86	19.45	71.84	62.54	17.04	73.40	66.91	18.48	0.62
LoRA-SFT	LRV	57.30	57.51	6.69	50.74	53.41	0.42	54.65	55.86	4.16	0.73
	LLaVA Data	62.18	62.65	11.61	60.81	62.55	18.61	61.63	62.61	14.44	0.68
	Ours	84.62	76.70	**45.00**	86.76	81.38	**55.81**	85.48	78.59	**49.35**	**0.31**
	Ours+LLaVA	**85.38**	**78.14**	42.49	**87.19**	82.11	55.27	**86.11**	**79.74**	47.64	0.37
LoRA-Rtune	LLaVA Data	69.68	66.53	14.92	73.85	67.95	20.78	71.36	67.11	17.28	0.75
	Ours	**86.10**	**78.09**	41.88	**85.38**	**78.52**	51.07	**85.81**	**78.26**	45.58	0.37
	Ours+LLaVA	85.46	77.14	**44.59**	85.25	78.20	**52.90**	85.37	77.57	**47.94**	**0.29**
LoRA-DPO	MMInstruction	66.10	65.18	18.03	55.98	56.10	7.57	62.02	61.52	13.81	**0.70**
	Ours	**74.70**	**69.79**	18.81	**73.70**	**68.59**	**20.38**	**74.30**	**69.30**	**19.44**	0.78
	Ours+MMinstruction	71.52	68.46	**19.70**	68.51	64.73	14.60	70.31	66.95	17.65	0.75
LLaVA-1.5-7B-LoRA⋆		33.72	37.36	17.46	4.59	50.55	0.78	35.11	48.61	1.25	0.62
Instruct-Tune	Ours	84.40	78.25	**53.54**	42.07	81.32	**50.33**	85.31	78.25	**42.50**	**0.41**
	Ours+LLaVA	**85.47**	**79.60**	46.16	**42.57**	**81.95**	37.62	**86.09**	**79.46**	31.92	0.64

Table 5: Comparison on different training strategies with our CERTAINLYUNCERTAIN. The best performances are highlighted with bold for each finetuning strategy. Acc: Accuracy. Conf-w.: Confidence-weighted. ECE: Expected Calibration Error.

	Refusal (LAVE_{idk} Acc. ↑)		Hallucination				Standard (VQA score ↑)	
			MM-Hal		POPE	AMBER		
	UNK-VQA	TDIUC	Overall↑	Hall. % ↓	F1 ↑	F1 ↑	VizWiz	VQAv2
Qwen-VL-Chat	41.32	95.10	**2.89**	0.41	81.30	**87.70**	66.85	72.96
LoRA-SFT-LLaVA Data	38.01	93.18	2.83	0.42	85.61	86.80	65.61	77.26
LoRA-SFT-Ours-only	**47.35**	**99.64**	2.70	**0.38**	**86.31**	81.30	**68.40**	69.77
LoRA-SFT-Ours+LLaVA	59.70	99.20	2.75	0.39	85.78	85.90	67.44	**77.32**
Instruct-tuning								
LLaVA-1.5-7B-LoRA	36.57	47.36	2.56	0.51	86.06	84.60	51.87	76.94
Ours-only	47.71	95.36	2.43	0.47	73.62	78.80	54.10	49.95
Ours+LLaVA Data	**49.12**	**98.70**	**2.66**	**0.45**	**88.05**	**86.60**	**54.40**	**77.37**

Table 6: Results of different model variants trained with CERTAINLYUNCERTAIN on other benchmarks. Hall. %: Hallucination ratio. ↑ (↓) indicates the larger (smaller) the better.

Table 5 presents comparisons among different training strategies. For the instruction-tuned Qwen-VL-Chat, we explore different continued finetuning methods with LoRA (Xu et al., 2023). For LLaVA, given the availability of their instruct-tuning data, we explore adding our data into instruct-tuning stage. Overall, we observe that SFT learns IDK better on our benchmark compared to other strategies, resulting in higher confidence-weighted accuracy. Within each training strategy—SFT, R-tuning, and DPO—we find that training on our data consistently improves performance, underscoring the quality of our dataset. Finally, SFT with our data also reduces ECE, demonstrating that models trained with CERTAINLYUNCERTAIN can express IDK more confidently.

Lastly, we examine the performance of our finetuned/instruction-tuned models on other benchmarks in Table 6. The results show that our dataset effectively improves model performance on refusal-based benchmarks, including UNK-VQA and TDIUC. It also demonstrates promising trends in reducing hallucination ratios in MM-Hal and improving F1 scores on POPE. Despite CERTAINLYUNCERTAIN focusing solely on VQA-type data, when finetuned/instruction-tuned with our data only, we did not observe a significant drop in the overall score of MM-Hal, which also evaluates tasks like captioning. When augmentating the LLaVA instruction tuning data with ours, it even improves the overall score of MM-Hal for LLaVA model. On AMBER, for Qwen-VL-Chat, SFT with any data combination in our experiments led to inferior results, especially when using only our data. We speculate that the degradation on AMBER is due to the lack of IDK questions on attributes and relations about non-existent objects in our dataset, which we plan to extend upon in future work. In comparison, POPE mainly focuses on existential questions about objects, which is more similar to our extraneous split. Moreover, on standard VQA benchmarks, we observe that models trained with our data combined with LLaVA data perform on par with the VQAv2 benchmark and show improvements on VizWiz, which contains unanswerable questions.

4 RELATED WORK

Abstention. Early studies in abstention primarily focused on the notion of confidence estimation in predictions, allowing to abstain when uncertain (Chow, 1957; De Stefano et al., 2000; El-Yaniv et al., 2010). Recent works used selective prediction approaches to particularly improve reliability under domain shift (Kamath et al., 2020) and with adversarial inputs (Varshney et al., 2022). Another promising direction involves extracting additional evidences by iteratively accumulating context (Srinivasan et al., 2024; Zeng et al., 2022; Shen et al., 2023b; You et al., 2023; Yang et al., 2023), rephrasing underspecified questions (Prasad et al., 2023), probing through code (Gupta & Kembhavi, 2023; Surís et al., 2023; Subramanian et al., 2023). Unlike our work, these approaches aim to reduce the risk of incorrect predictions despite having definitive answers, without addressing epistemic or aleatoric uncertainties.

Hallucinations. Models tend to over-confidently hallucinate in uncertain scenarios (Zhu et al., 2024a; Kang et al., 2023; Yin et al., 2023). There are two primary techniques for hallucination detection (Luo et al., 2024) – at token-level (Liu et al., 2021; Zhou et al., 2020; Dziri et al., 2021; Cao et al., 2021) and sentence-level (Manakul et al., 2023; Zha et al., 2023; Shen et al., 2023a; Li et al., 2023a). Our work aim to reduce the confidence of hallucinatory responses at the answer-level. Similar to extracting additional evidence, hallucination mitigation strategies use retrieval-based approaches (Ji et al., 2022; Dziri et al., 2021; Shuster et al., 2021), which condition outputs on factual data by using external knowledge sources, particularly helps increase reliability on the *extraneous* category.

Evaluation Metrics. Standard accuracy or generation metrics such as BLEU are insufficient to evaluate the confidence of open-ended answer generation. To assess the semantic possibilities the LAVE metric (Mañas et al., 2024) was introduced to fully or partially score the predicted answer based on their overlap with the ground truth. Expected Calibration Error (ECE) measures the accuracy of probability estimates in representing true correctness likelihood. More recent approaches also rely on object detection (Johnson et al., 2015; Rohrbach et al., 2018; Li et al., 2023c; Lovenia et al., 2023; Gunjal et al., 2023) or entailment (Faithscore) (Jing et al., 2023) to measure hallucinations. However, none of them directly indicate the confidence of the model predictions. In this work, we introduce confidence-weighted accuracy based on LAVE accuracy, which better correlates with ECE.

Datasets. Most standard multimodal benchmarks focus on clear, definitive answers or partial hallucinations (Gunjal et al., 2024; Bai et al., 2024) for discriminative (Li et al., 2023c; Chen et al., 2023b; Wang et al., 2024a) or generative tasks (Liu et al., 2023a; Jing et al., 2023). In contrast, CERTAINLYUNCERTAIN targets scenarios where responding with IDK is the correct response. Similar concurrent efforts for text-only benchmarks (Li et al., 2023a; Zhu et al., 2024b; Mishra et al., 2024) are widely explored. Generating counterfactual instruction text data (Yu et al., 2023) is the closest equivalent of LRV data (Liu et al., 2023a) which includes positive (or negative) instructions about objects or attributes present (or absent) in the image. We also use the MMInstruction (Li et al., 2023b) with preference annotations for helpfulness, faithfulness and ethical considerations. Finally, we generate model-dependent refusal datasets automatically which is explored by (Zhang et al., 2023) to adapt to multimodal R-tuning. Our experiments show that these datasets are insufficient for benchmarking or improving multimodal epistemic and aleatoric awareness.

5 CONCLUSIONS AND FUTURE WORK

Acknowledging uncertainty in responses and appropriately responding with "I don't know" is crucial for the reliability and trustworthiness of AI systems. In this work, we introduce a new taxonomy of multimodal uncertainty. Based on this taxonomy, we present a new benchmark, CERTAINLYUNCERTAIN, and demonstrate that current VLMs lack self-awareness of these uncertainties. Empirical results show that fine-tuning with our data not only leads to improvements on expressing uncertainty, but also performance gains on refusal-based benchmarks, and some hallucination-based benchmarks, while not degrading on standard benchmarks. Additionally, we propose a new confidence-weighted accuracy metric that combines accuracy with the confidence of the prediction, showing strong correlations with both accuracy and ECE. Our work paves the way for future research directions in modeling uncertainties. For instance, future efforts could extend the annotations to include rationales for IDK responses, specifying which category of uncertainty is responsible. Our confidence-weighted metric holds potential for application in other unimodal and multimodal datasets and could be explored as a reward mechanism in model training.

ETHICS STATEMENT

CERTAINLYUNCERTAIN aims to assess VLMs' awareness of different categories of multimodal uncertainty, and has showed improvements over the base model when finetuned with it. The dataset was curated following ethical guidelines to ensure that no sensitive information is included and to minimize bias during the data collection process. No personally identifiable data was collected or processed. The benchmarking process together with the newly proposed evaluation metric aims to be transparent and reproducible, adhering to high standards of research integrity and ethical conduct.

Despite the comprehensive taxonomy we proposed for constructing CERTAINLYUNCERTAIN, it is possible that the synthetic generation pipeline may not capture all the nuances of real-world uncertain scenarios. While the current open-domain settings in CERTAINLYUNCERTAIN do not raise any ethical concerns, progressing the field in this direction would require community-built regulations around what domains must present model-confidence/uncertainty along with the predictions. Users might place undue trust in model predictions, leading to potentially harmful outcomes. Therefore it is important to have the following considerations:

1. **Guardrails**: Models should incorporate guardrails that prevent the automatic acceptance of highly uncertain predictions in critical situations. For instance, a high uncertainty score should trigger a review by a human expert before any action is taken based on the model's output.
2. **Transparency**: The methods used to evaluate and present these confidence scores need to be transparent. In life-altering situations, it is ethically necessary that the predictions, including their confidence levels and associated uncertainties, are fully traceable. This traceability allows for accountability, ensuring that decisions made based on the model's outputs can be reviewed and justified.
3. **Context-Aware Uncertainty**: The model should tailor its uncertainty predictions based on the context. For example, a prediction in a high-stakes domain (e.g., medical diagnosis) requires more conservative confidence levels and clear articulation of potential risks compared to low-stakes domains (e.g., casual information retrieval). In areas such as healthcare, where predictions can significantly impact patient outcomes, ethical responsibility demands that the model not only predicts uncertainty accurately but also ensures that this uncertainty is understood by healthcare providers. Miscommunication or misunderstanding of uncertainty could lead to incorrect diagnoses, inappropriate treatments, or delayed medical interventions, all of which could be life-threatening.
4. **Unbiased certainty**: The model must be trained and evaluated to ensure that its uncertainty predictions do not disproportionately affect certain groups.

REPRODUCIBILITY STATEMENT

To ensure the reproducibility of our results, we have made considerable efforts to provide all necessary details and materials. Specifically, we have included a comprehensive description of the dataset creation process in Section 2.2, and further elaborated in Appendix F. The experimental settings are described in Section 3, with additional implementation details in Appendix E. The evaluation metrics are clearly defined in Section 2.3 to facilitate independent verification. We also provide qualitative visualization of data samples and model predictions in Appendix A and B. Human evaluation results are reported in Appendix C. For easier access, code and dataset samples are included in Supplementary Material.

REFERENCES

Microsoft Azure. `https://azure.microsoft.com/`.

01. AI, :, Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, Kaidong Yu, Peng Liu, Qiang Liu, Shawn Yue, Senbin Yang, Shiming Yang, Tao Yu, Wen Xie, Wenhao Huang, Xiaohui Hu, Xiaoyi Ren, Xinyao Niu, Pengcheng Nie, Yuchi Xu, Yudong Liu, Yue Wang, Yuxuan Cai, Zhenyu Gu, Zhiyuan Liu, and Zonghong Dai. Yi: Open foundation models by 01.ai, 2024.

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. 2023.

Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou. Hallucination of multimodal large language models: A survey. *CoRR*, abs/2404.18930, 2024. doi: 10.48550/ARXIV.2404.18930. URL https://doi.org/10.48550/arXiv.2404.18930.

Mengyao Cao, Yue Dong, and Jackie Chi Kit Cheung. Hallucinated but factual! inspecting the factuality of hallucinations in abstractive summarization. In *Annual Meeting of the Association for Computational Linguistics*, 2021. URL https://api.semanticscholar.org/CorpusID:244909449.

Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv preprint arXiv:2312.14238*, 2023a.

Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*, 2024.

Zhiyang Chen, Yousong Zhu, Yufei Zhan, Zhaowen Li, Chaoyang Zhao, Jinqiao Wang, and Ming Tang. Mitigating hallucination in visual language models with visual supervision. *CoRR*, abs/2311.16479, 2023b. doi: 10.48550/ARXIV.2311.16479. URL https://doi.org/10.48550/arXiv.2311.16479.

C.K. Chow. An optimum character recognition system using decision functions. *IRE Transactions on Electronic Computers*, 1957.

OpenCompass Contributors. Opencompass: A universal evaluation platform for foundation models. https://github.com/open-compass/opencompass, 2023.

Ernest Davis. Unanswerable questions about images and texts. *Frontiers Artif. Intell.*, 3:51, 2020. doi: 10.3389/FRAI.2020.00051. URL https://doi.org/10.3389/frai.2020.00051.

Claudio De Stefano, Carlo Sansone, and Mario Vento. To reject or not to reject: that is the question-an answer in case of neural classifiers. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 30(1):84–94, 2000.

Nouha Dziri, Andrea Madotto, Osmar Zaiane, and A. Bose. Neural path hunter: Reducing hallucination in dialogue systems via path grounding. *ArXiv*, abs/2104.08455, 2021. URL https://api.semanticscholar.org/CorpusID:233296059.

Ran El-Yaniv et al. On the foundations of noise-free selective classification. *Journal of Machine Learning Research*, 11(5), 2010.

Alessandro Favero, Luca Zancato, Matthew Trager, Siddharth Choudhary, Pramuditha Perera, Alessandro Achille, Ashwin Swaminathan, and Stefano Soatto. Multi-modal hallucination control by visual information grounding. *CoRR*, abs/2403.14003, 2024. doi: 10.48550/ARXIV.2403.14003. URL https://doi.org/10.48550/arXiv.2403.14003.

Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023.

Yash Goyal, Tejas Khot, Aishwarya Agrawal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. *Int. J. Comput. Vis.*, 127(4):398–414, 2019. doi: 10.1007/S11263-018-1116-0. URL https://doi.org/10.1007/s11263-018-1116-0.

Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. Hallusionbench: An advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14375–14385, June 2024.

Anish Gunjal, Jihan Yin, and Erhan Bas. Detecting and preventing hallucinations in large vision language models. In *AAAI Conference on Artificial Intelligence*, 2023. URL https://api.semanticscholar.org/CorpusID:260887222.

Anisha Gunjal, Jihan Yin, and Erhan Bas. Detecting and preventing hallucinations in large vision language models. In Michael J. Wooldridge, Jennifer G. Dy, and Sriraam Natarajan (eds.), *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*, pp. 18135–18143. AAAI Press, 2024. doi: 10.1609/AAAI.V38I16.29771. URL https://doi.org/10.1609/aaai.v38i16.29771.

Yanyang Guo, Fangkai Jiao, Zhiqi Shen, Liqiang Nie, and Mohan S. Kankanhalli. UNK-VQA: A dataset and A probe into multi-modal large models' abstention ability. *CoRR*, abs/2310.10942, 2023. doi: 10.48550/ARXIV.2310.10942. URL https://doi.org/10.48550/arXiv.2310.10942.

Tanmay Gupta and Aniruddha Kembhavi. Visual programming: Compositional visual reasoning without training. In *Computer Vision and Pattern Recognition (CVPR)*, 2023.

Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 3608–3617. Computer Vision Foundation / IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00380. URL http://openaccess.thecvf.com/content_cvpr_2018/html/Gurari_VizWiz_Grand_Challenge_CVPR_2018_paper.html.

Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6700–6709, 2019.

Ziwei Ji, Zihan Liu, Nayeon Lee, Tiezheng Yu, Bryan Wilie, Mini Zeng, and Pascale Fung. Rho ($$\rho$$): Reducing hallucination in open-domain dialogues with knowledge grounding. In *Annual Meeting of the Association for Computational Linguistics*, 2022. URL https://api.semanticscholar.org/CorpusID:254246604.

Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023.

Liqiang Jing, Ruosen Li, Yunmo Chen, Mengzhao Jia, and Xinya Du. FAITHSCORE: evaluating hallucinations in large vision-language models. *CoRR*, abs/2311.01477, 2023. doi: 10.48550/ARXIV.2311.01477. URL https://doi.org/10.48550/arXiv.2311.01477.

Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4565–4574, 2015. URL https://api.semanticscholar.org/CorpusID:14521054.

Kushal Kafle and Christopher Kanan. An analysis of visual question answering algorithms. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pp. 1983–1991. IEEE Computer Society, 2017. doi: 10.1109/ICCV.2017.217. URL https://doi.org/10.1109/ICCV.2017.217.

Amita Kamath, Robin Jia, and Percy Liang. Selective question answering under domain shift. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pp. 5684–5696. Association for Computational Linguistics, 2020. doi: 10.18653/V1/2020.ACL-MAIN.503. URL https://doi.org/10.18653/v1/2020.acl-main.503.

Haoqiang Kang, Juntong Ni, and Huaxiu Yao. Ever: Mitigating hallucination in large language models through real-time verification and rectification. *CoRR*, abs/2311.09114, 2023. doi: 10.48550/ARXIV.2311.09114. URL https://doi.org/10.48550/arXiv.2311.09114.

Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024.

Junyi Li, Xiaoxue Cheng, Wayne Xin Zhao, Jianyun Nie, and Ji rong Wen. Halueval: A large-scale hallucination evaluation benchmark for large language models. *ArXiv*, abs/2305.11747, 2023a. URL https://api.semanticscholar.org/CorpusID:258832847.

Lei Li, Zhihui Xie, Mukai Li, Shunian Chen, Peiyi Wang, Liang Chen, Yazheng Yang, Benyou Wang, and Lingpeng Kong. Silkie: Preference distillation for large visual language models. *CoRR*, abs/2312.10665, 2023b. doi: 10.48550/ARXIV.2312.10665. URL https://doi.org/10.48550/arXiv.2312.10665.

Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji rong Wen. Evaluating object hallucination in large vision-language models. In *Conference on Empirical Methods in Natural Language Processing*, 2023c. URL https://api.semanticscholar.org/CorpusID:258740697.

Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pp. 292–305. Association for Computational Linguistics, 2023d. doi: 10.18653/V1/2023.EMNLP-MAIN.20. URL https://doi.org/10.18653/v1/2023.emnlp-main.20.

Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. Mitigating hallucination in large multi-modal models via robust instruction tuning. In *The Twelfth International Conference on Learning Representations*, 2023a.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *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*, 2023b. URL http://papers.nips.cc/paper_files/paper/2023/hash/6dcf277ea32ce3288914faf369fe6de0-Abstract-Conference.html.

Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, January 2024. URL https://llava-vl.github.io/blog/2024-01-30-llava-next/.

Tianyu Liu, Yizhe Zhang, Christopher John Brockett, Yi Mao, Zhifang Sui, Weizhu Chen, and William B. Dolan. A token-level reference-free hallucination detection benchmark for free-form text generation. *ArXiv*, abs/2104.08704, 2021. URL https://api.semanticscholar.org/CorpusID:233296648.

Holy Lovenia, Wenliang Dai, Samuel Cahyawijaya, Ziwei Ji, and Pascale Fung. Negative object presence evaluation (nope) to measure object hallucination in vision-language models. *ArXiv*, abs/2310.05338, 2023. URL https://api.semanticscholar.org/CorpusID:263829510.

Junliang Luo, Tianyu Li, Di Wu, Michael Jenkin, Steve Liu, and Gregory Dudek. Hallucination detection and hallucination mitigation: An investigation. *CoRR*, abs/2401.08358, 2024. doi: 10.48550/ARXIV.2401.08358. URL https://doi.org/10.48550/arXiv.2401.08358.

Aroma Mahendru, Viraj Prabhu, Akrit Mohapatra, Dhruv Batra, and Stefan Lee. The promise of premise: Harnessing question premises in visual question answering. In Martha Palmer, Rebecca Hwa, and Sebastian Riedel (eds.), *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pp. 926–935. Association for Computational Linguistics, 2017. doi: 10.18653/V1/D17-1097. URL https://doi.org/10.18653/v1/d17-1097.

Potsawee Manakul, Adian Liusie, and Mark John Francis Gales. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. *ArXiv*, abs/2303.08896, 2023. URL https://api.semanticscholar.org/CorpusID:257557820.

Oscar Mañas, Benno Krojer, and Aishwarya Agrawal. Improving automatic VQA evaluation using large language models. In Michael J. Wooldridge, Jennifer G. Dy, and Sriraam Natarajan (eds.), *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*, pp. 4171–4179. AAAI Press, 2024. doi: 10.1609/AAAI.V38I5.28212. URL https://doi.org/10.1609/aaai.v38i5.28212.

Abhika Mishra, Akari Asai, Vidhisha Balachandran, Yizhong Wang, Graham Neubig, Yulia Tsvetkov, and Hannaneh Hajishirzi. Fine-grained hallucination detection and editing for language models. *ArXiv*, abs/2401.06855, 2024. URL https://api.semanticscholar.org/CorpusID:266999558.

Atsuyuki Miyai, Jingkang Yang, Jingyang Zhang, Yifei Ming, Qing Yu, Go Irie, Yixuan Li, Hai Li, Ziwei Liu, and Kiyoharu Aizawa. Unsolvable problem detection: Evaluating trustworthiness of vision language models. *CoRR*, abs/2403.20331, 2024. doi: 10.48550/ARXIV.2403.20331. URL https://doi.org/10.48550/arXiv.2403.20331.

Mahdi Pakdaman Naeini, Gregory F. Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In Blai Bonet and Sven Koenig (eds.), *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, pp. 2901–2907. AAAI Press, 2015. doi: 10.1609/AAAI.V29I1.9602. URL https://doi.org/10.1609/aaai.v29i1.9602.

Yasumasa Onoe, Sunayana Rane, Zachary Berger, Yonatan Bitton, Jaemin Cho, Roopal Garg, Alexander Ku, Zarana Parekh, Jordi Pont-Tuset, Garrett Tanzer, Su Wang, and Jason Baldridge. DOCCI: descriptions of connected and contrasting images. *CoRR*, abs/2404.19753, 2024. doi: 10.48550/ARXIV.2404.19753. URL https://doi.org/10.48550/arXiv.2404.19753.

OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774, 2023. doi: 10.48550/ARXIV.2303.08774. URL https://doi.org/10.48550/arXiv.2303.08774.

Richard Paul and Linda Elder. Critical thinking: What, why, and how. *New directions for community colleges*, 77(2):3–24, 1992.

Archiki Prasad, Elias Stengel-Eskin, and Mohit Bansal. Rephrase, augment, reason: Visual grounding of questions for vision-language models. *arXiv preprint arXiv:2310.05861*, 2023.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *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*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/a85b405ed65c6477a4fe8302b5e06ce7-Abstract-Conference.html.

Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In *KDD*, 2020.

Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, Zhaoyang Zeng, Hao Zhang, Feng Li, Jie Yang, Hongyang Li, Qing Jiang, and Lei Zhang. Grounded SAM: assembling open-world models for diverse visual tasks. *CoRR*, abs/2401.14159, 2024. doi: 10.48550/ARXIV.2401.14159. URL https://doi.org/10.48550/arXiv.2401.14159.

Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. Object hallucination in image captioning. In *Conference on Empirical Methods in Natural Language Processing*, 2018. URL https://api.semanticscholar.org/CorpusID:52176506.

Jiaming Shen, Jialu Liu, Daniel Finnie, Negar Asgharipour Rahmati, Michael Bendersky, and Marc Najork. "why is this misleading?": Detecting news headline hallucinations with explanations. *Proceedings of the ACM Web Conference 2023*, 2023a. URL https://api.semanticscholar.org/CorpusID:256827429.

Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugginggpt: Solving ai tasks with chatgpt and its friends in huggingface. *arXiv preprint arXiv:2303.17580*, 2023b.

Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. Retrieval augmentation reduces hallucination in conversation. In *Conference on Empirical Methods in Natural Language Processing*, 2021. URL https://api.semanticscholar.org/CorpusID:233240939.

Tejas Srinivasan, Jack Hessel, Tanmay Gupta, Bill Yuchen Lin, Yejin Choi, Jesse Thomason, and Khyathi Raghavi Chandu. Selective "selective prediction": Reducing unnecessary abstention in vision-language reasoning. *CoRR*, abs/2402.15610, 2024. doi: 10.48550/ARXIV.2402.15610. URL https://doi.org/10.48550/arXiv.2402.15610.

Sanjay Subramanian, Medhini Narasimhan, Kushal Khangaonkar, Kevin Yang, Arsha Nagrani, Cordelia Schmid, Andy Zeng, Trevor Darrell, and Dan Klein. Modular visual question answering via code generation. *arXiv preprint arXiv:2306.05392*, 2023.

Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, Kurt Keutzer, and Trevor Darrell. Aligning large multimodal models with factually augmented RLHF. *CoRR*, abs/2309.14525, 2023. doi: 10.48550/ARXIV.2309.14525. URL https://doi.org/10.48550/arXiv.2309.14525.

Dídac Surís, Sachit Menon, and Carl Vondrick. Vipergpt: Visual inference via python execution for reasoning. *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, 2023.

Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. *arXiv preprint arXiv:2109.07161*, 2021.

Neeraj Varshney and Chitta Baral. Post-abstention: Towards reliably re-attempting the abstained instances in QA. In Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pp. 967–982. Association for Computational Linguistics, 2023. doi: 10.18653/V1/2023.ACL-LONG.55. URL https://doi.org/10.18653/v1/2023.acl-long.55.

Neeraj Varshney, Swaroop Mishra, and Chitta Baral. Investigating selective prediction approaches across several tasks in iid, ood, and adversarial settings. In *Findings of the Association for Computational Linguistics (ACL)*, 2022.

Junyang Wang, Yuhang Wang, Guohai Xu, Jing Zhang, Yukai Gu, Haitao Jia, Ming Yan, Ji Zhang, and Jitao Sang. An llm-free multi-dimensional benchmark for mllms hallucination evaluation. *CoRR*, abs/2311.07397, 2023a. doi: 10.48550/ARXIV.2311.07397. URL https://doi.org/10.48550/arXiv.2311.07397.

Junyang Wang, Yiyang Zhou, Guohai Xu, Pengcheng Shi, Chenlin Zhao, Haiyang Xu, Qinghao Ye, Ming Yan, Ji Zhang, Jihua Zhu, Jitao Sang, and Haoyu Tang. Evaluation and analysis of hallucination in large vision-language models. *CoRR*, abs/2308.15126, 2023b. doi: 10.48550/ARXIV.2308.15126. URL https://doi.org/10.48550/arXiv.2308.15126.

Lei Wang, Jiabang He, Shenshen Li, Ning Liu, and Ee-Peng Lim. Mitigating fine-grained hallucination by fine-tuning large vision-language models with caption rewrites. In Stevan Rudinac, Alan Hanjalic, Cynthia C. S. Liem, Marcel Worring, Bjorn Jonsson, Bei Liu, and Yoko Yamakata (eds.), *MultiMedia Modeling - 30th International Conference, MMM 2024, Amsterdam, The Netherlands, January 29 - February 2, 2024, Proceedings, Part IV*, volume 14557 of *Lecture Notes in Computer Science*, pp. 32–45. Springer, 2024a. doi: 10.1007/978-3-031-53302-0_3. URL https://doi.org/10.1007/978-3-031-53302-0_3.

Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024b.

Peter West, Ximing Lu, Nouha Dziri, Faeze Brahman, Linjie Li, Jena D. Hwang, Liwei Jiang, Jillian Fisher, Abhilasha Ravichander, Khyathi Chandu, Benjamin Newman, Pang Wei Koh, Allyson Ettinger, and Yejin Choi. The generative AI paradox: "what it can create, it may not understand". *CoRR*, abs/2311.00059, 2023. doi: 10.48550/ARXIV.2311.00059. URL https://doi.org/10.48550/arXiv.2311.00059.

Spencer Whitehead, Suzanne Petryk, Vedaad Shakib, Joseph Gonzalez, Trevor Darrell, Anna Rohrbach, and Marcus Rohrbach. Reliable visual question answering: Abstain rather than answer incorrectly. In Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (eds.), *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXVI*, volume 13696 of *Lecture Notes in Computer Science*, pp. 148–166. Springer, 2022. doi: 10.1007/978-3-031-20059-5_9. URL https://doi.org/10.1007/978-3-031-20059-5_9.

Yuhui Xu, Lingxi Xie, Xiaotao Gu, Xin Chen, Heng Chang, Hengheng Zhang, Zhengsu Chen, Xiaopeng Zhang, and Qi Tian. Qa-lora: Quantization-aware low-rank adaptation of large language models. *CoRR*, abs/2309.14717, 2023. doi: 10.48550/ARXIV.2309.14717. URL https://doi.org/10.48550/arXiv.2309.14717.

Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. Mm-react: Prompting chatgpt for multimodal reasoning and action. *arXiv preprint arXiv:2303.11381*, 2023.

Shukang Yin, Chaoyou Fu, Sirui Zhao, Tong Xu, Hao Wang, Dianbo Sui, Yunhang Shen, Ke Li, Xing Sun, and Enhong Chen. Woodpecker: Hallucination correction for multimodal large language models. *CoRR*, abs/2310.16045, 2023. doi: 10.48550/ARXIV.2310.16045. URL https://doi.org/10.48550/arXiv.2310.16045.

Haoxuan You, Rui Sun, Zhecan Wang, Long Chen, Gengyu Wang, Hammad A Ayyubi, Kai-Wei Chang, and Shih-Fu Chang. Idealgpt: Iteratively decomposing vision and language reasoning via large language models. *arXiv preprint arXiv:2305.14985*, 2023.

Qifan Yu, Juncheng Li, Longhui Wei, Liang Pang, Wentao Ye, Bosheng Qin, Siliang Tang, Qi Tian, and Yueting Zhuang. Hallucidoctor: Mitigating hallucinatory toxicity in visual instruction data. *CoRR*, abs/2311.13614, 2023. doi: 10.48550/ARXIV.2311.13614. URL https://doi.org/10.48550/arXiv.2311.13614.

Yuanhan Zhang Bo Li Songyang Zhang Wangbo Zhao Yike Yuan Jiaqi Wang Conghui He Ziwei Liu Kai Chen Dahua Lin Yuan Liu, Haodong Duan. Mmbench: Is your multi-modal model an all-around player? *arXiv:2307.06281*, 2023.

Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhu Chen. MMMU: A massive multi-discipline multimodal understanding and reasoning benchmark for expert AGI. *CoRR*, abs/2311.16502, 2023. doi: 10.48550/ARXIV.2311.16502. URL https://doi.org/10.48550/arXiv.2311.16502.

Andy Zeng, Maria Attarian, Krzysztof Marcin Choromanski, Adrian Wong, Stefan Welker, Federico Tombari, Aveek Purohit, Michael S Ryoo, Vikas Sindhwani, Johnny Lee, et al. Socratic models: Composing zero-shot multimodal reasoning with language. In *International Conference on Learning Representations (ICLR)*, 2022.

Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. Alignscore: Evaluating factual consistency with a unified alignment function. In *Annual Meeting of the Association for Computational Linguistics*, 2023. URL https://api.semanticscholar.org/CorpusID:258947273.

Hanning Zhang, Shizhe Diao, Yong Lin, Yi Ren Fung, Qing Lian, Xingyao Wang, Yangyi Chen, Heng Ji, and Tong Zhang. R-tuning: Instructing large language models to say 'i don't know'. 2023. URL https://api.semanticscholar.org/CorpusID:265220839.

Xuchao Zhang, Fanglan Chen, Chang-Tien Lu, and Naren Ramakrishnan. Mitigating uncertainty in document classification. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pp. 3126–3136. Association for Computational Linguistics, 2019. doi: 10.18653/V1/N19-1316. URL https://doi.org/10.18653/v1/n19-1316.

Zhiyuan Zhao, Bin Wang, Linke Ouyang, Xiaoyi Dong, Jiaqi Wang, and Conghui He. Beyond hallucinations: Enhancing lvlms through hallucination-aware direct preference optimization, 2023.

Chunting Zhou, Jiatao Gu, Mona T. Diab, Paco Guzmán, Luke Zettlemoyer, and Marjan Ghazvininejad. Detecting hallucinated content in conditional neural sequence generation. *ArXiv*, abs/2011.02593, 2020. URL https://api.semanticscholar.org/CorpusID:226254579.

Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. Analyzing and mitigating object hallucination in large vision-language models. *CoRR*, abs/2310.00754, 2023. doi: 10.48550/ARXIV.2310.00754. URL https://doi.org/10.48550/arXiv.2310.00754.

Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.

Zhiying Zhu, Zhiqing Sun, and Yiming Yang. Halueval-wild: Evaluating hallucinations of language models in the wild. *CoRR*, abs/2403.04307, 2024a. doi: 10.48550/ARXIV.2403.04307. URL https://doi.org/10.48550/arXiv.2403.04307.

Zhiying Zhu, Zhiqing Sun, and Yiming Yang. Halueval-wild: Evaluating hallucinations of language models in the wild. *ArXiv*, abs/2403.04307, 2024b. URL https://api.semanticscholar.org/CorpusID:268264670.

A SAMPLES VISUALIZING CERTAINLYUNCERTAIN BENCHMARK

We visualize some samples from each fine-grained category of the epistemic and aleatoric awareness. For the category of *extraneous*, our data is made of samples where the answer differs for the same question when the image is perturbed. For the rest of the categories, the dataset contains samples where the same image is paired with answerable and unanswerable questions. Note that during our data collection, we indeed instruct GPT4 or GPT4V to express its uncertainty in the form of some example responses provided by the authors including "unanswerable", "I don't know", "I don't see xxx", "Nothing" and etc.. However, in practice, we do observe the two models express "I don't know" in more diverse form, for example, provide a probable guess but explicitly note its confidence; bluntly refuse (e.g., "I don't know" or "unanswerable"); asking a follow-up question; For simplicity, all answers are normalized to "I don't know" in our visualization. We refer the readers to the dataset in supplementary material for the actual answers.

Figure 6 shows the category of knowledge awareness; as we can see the unanswerable questions ask about information that is hard to identify from the context of the image and requires additional knowledge. Similarly, Figure 7 shows examples from the complexity awareness in the epistemic category. The unanswerable questions are too tedious to arrive at an answer while the answerable questions still require some efforts, such as counting but is not laborious to answer.

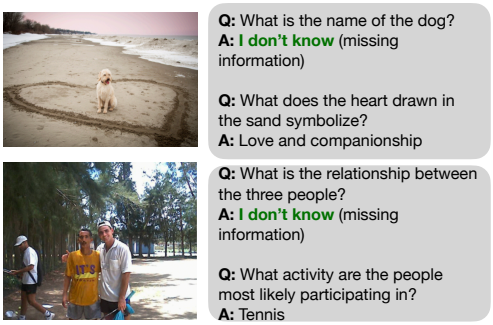

Figure 6: Samples from Knowledge Awareness (Epistemic) category

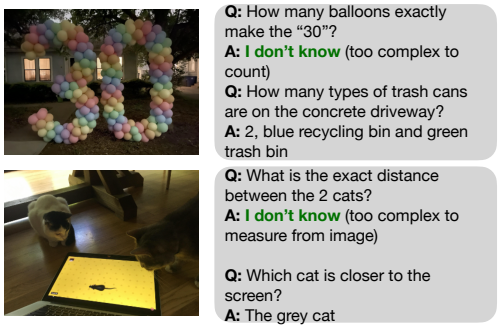

Figure 7: Samples from Complexity Awareness (Epistemic) category

For the extraneous sub-category of the epistemic awareness, we perturb the image to mask and remove the target object about which the question seeks information. Samples from this category are shown in Figure 8. As we can see the target objects in the questions are 'cat' and 'statue'. These objects are removed from the image and inpainted to get a natural-looking image to obtain a perturbed image. The resulting image paired with the same question now becomes unanswerable. The answerable question is paired with the original unperturbed question to have a definitive answer (which is the standard setup for most VQA based benchmarks).

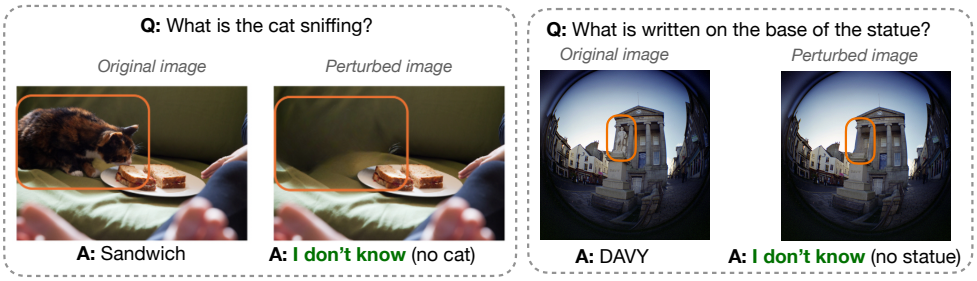

Figure 8: Samples from Extraneous Awareness (Epistemic) category

Figures 9 and 10 show samples from CERTAINLYUNCERTAIN in the aleatoric category, particularly of the temporal and ambiguity awareness sub-categories respectively. The temporal sub-category, as

we can see, contains questions about the current happenings or state of the image for the answerable part and the unanswerable questions ask about the future that is not directly predictable from the current state. The ambiguous awareness category contains questions with a definitive answer for the answerable type and questions with many plausible answers but cannot choose a single definite answer for the unanswerable type.

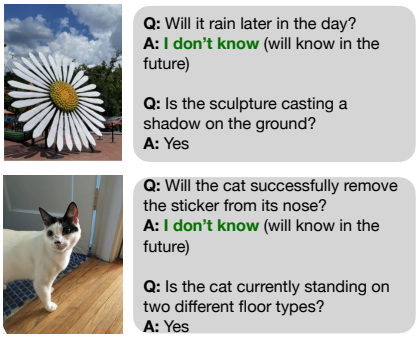

Figure 9: Samples from Temporal Awareness (Aleatoric) category

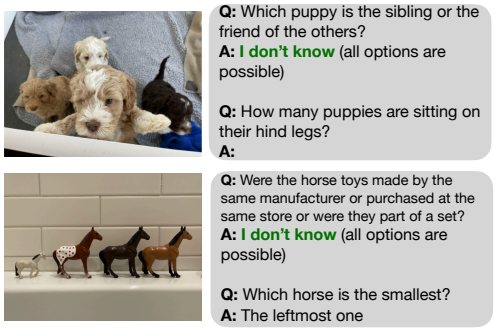

Figure 10: Samples from Ambiguous Awareness (Aleatoric) category

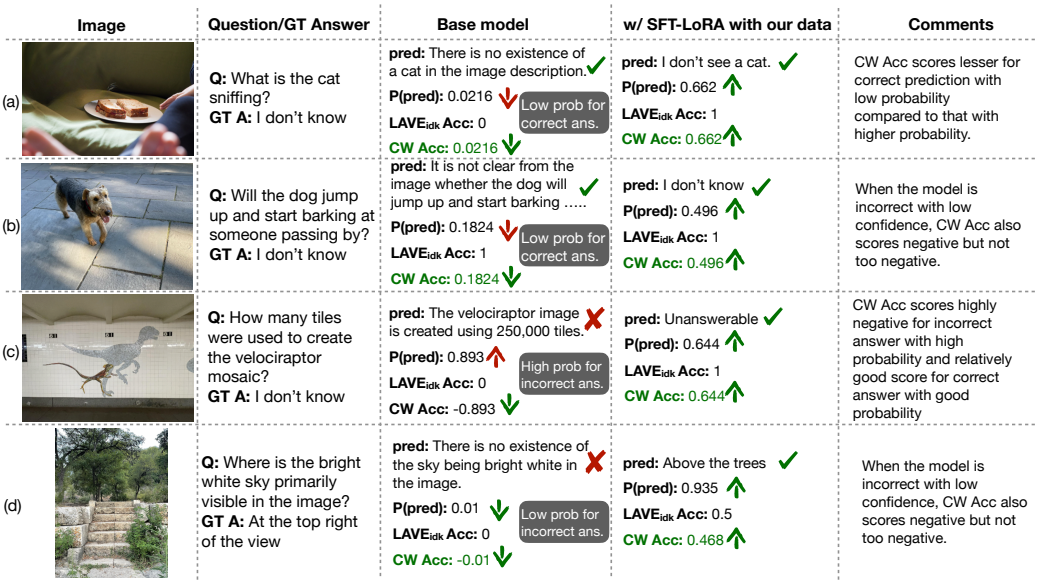

Figure 11: Visualization of model predictions and the corresponding LAVE_{idk} accuracy, P_{pred} and confidence-weighted accuracy. The base model here is Qwen-VL-Chat (Bai et al., 2023). Our confidence-weighted accuracy is represented as CW Acc in this figure.

B SAMPLES VISUALIZING PREDICTIONS AND *confidence-weighted metric*

Our proposed confidence-weighted accuracy takes into account the prediction probability and the correctness of the predicted answer to give a holistic score. Figure 11 presents the visualization of model predictions and the corresponding LAVE_{idk} accuracy, P_{pred} and confidence-weighted accuracy. We show that the proposed confidence-weighted accuracy gives less score for a correct answer with lower confidence, and penalizes more for an incorrect answer with higher confidence. In addition, our visualization shows that Qwen-VL-Chat (Bai et al., 2023) is able to say equivalents of "I don't know" more confidently from (a) and (b), after continued finetuning on our data with SFT-LoRA.

Examples (a) and (b) show cases where the base model is less confident for a correct answer. Our metric gives a partial score for the correctness owing to the prediction probability. After finetuning, as

Extraneous	Knowledge	Complex	Temporal	Ambiguous
80.1%	93.8%	97.9%	98.2%	94.4%

Table 7: Percentage of valid questions in the original testing splits. We only preserve the human verified questions for the final testing split.

Mistral-7B	Yi-34B-4bits	GPT4	Human
98%	100%	100%	100%

Table 8: LAVE evaluator accuracy on GT IDK. Performance are reported on 100 random samples from extraneous split.

Kendall Correlation	P-value
0.8728	2.4806e-22

Table 9: Correlation between human judgment and LAVE_{idk} Accuracy.

the prediction probability of the correct answer increases, our confidence-weighted accuracy increases accordingly. In case (c), the base model predicts an incorrect answer with high confidence. Our metric penalizes this more heavily with a high negative score. After finetuning, the prediction is rectified and the scores are adjusted accordingly. In the case of (d), the base model predicts the incorrect answer but with low confidence. Our metric still gives a negative score but penalizes less compared to (c). The cases of (c) and (d) differentiate answering incorrectly with high and low probabilities respectively.

Moreover, after finetuning with our CERTAINLYUNCERTAIN, we see the corrected predictions with relatively higher probabilities for correctness, which are reflected in our confidence-weighted metric score. These probabilities of the model predictions are not reflected in the LAVE_{idk} accuracy.

C HUMAN EVALUATION

Human verification on Testing Data Quality. Using synthetic model-generated data for training has become a de facto approach in the literature, where previous works (Liu et al., 2023b; Zhu et al., 2023) have consistently observed great performance improvement even with noisy labels on synthetic model-generated data. This is also consistent with our observation in Table 6, where training with CERTAINLYUNCERTAIN has demonstrated improvements for existing refusals datasets, most hallucination datasets, while maintaining performance on standard benchmarks. Therefore, we focus our manual efforts with authors on quality check and filtering for the testing split of CERTAINLYUNCERTAIN. In practice, we observe that the extraneous split suffers more from model failure, due to its dependency on multiple models in the data creation process. The percentage of valid questions in different splits are reported in the Table 7. With this human verification process, our final testing split only contains human verified questions.

LAVE IDK Judgement Accuracy. As our metric relies on the accuracy of LAVE refusal judgment, we experiment with different LLMs on 100 samples from extraneous split. Table 8 presents the results with Mistral-7B (Jiang et al., 2023), Yi-34B-4bits (AI et al., 2024) and GPT4, in comparison with Human. Given the high performance, and also considering the latency and cost of model inference, we decide to use the smaller model Mistral-7B in our evaluation.

Correlation between LAVE_{idk} Accuracy and Human Judgement. Our LAVE_{idk} Accuracy is an extension to the LAVE metric (Mañas et al., 2024), which has demonstrated significant alignment with human judgment and generalization to new VQA models and benchmarks. We further conduct an additional small-scale human evaluation on Qwen-VL-Chat finetuned model and compare it to the LAVE_{idk} Accuracy on 100 random samples in CERTAINLYUNCERTAIN testing split, and follow to report the kendall correlation and the p-value between LAVE_{idk} Accuracy and human judgment in Table 9. The low p-value indicates that the correlation coefficient between human judgment and LAVE_{idk} Accuracy is statistically significant.

Human Performance on CERTAINLYUNCERTAIN. We conduct a small scale human evaluation on the 100 randomly sampled questions. As it is impossible to truly measure human confidence on their predictions, we report LAVE Acc. instead of confidence-weighted accuracy. Overall, human achieves much higher score of 94% than models, with close to perfect score of 98% on unanswerable questions (with only 1 wrong answer, 49 correct answers), while slightly lower (90.0%) on answerable questions. In this small scale evaluation, we found that humans tend to provide brief answers to answerable questions. For example, when asked about the color of a mostly black clothing with some

patches of white, humans may provide answers as black, while the ground truth answer is mostly black with some patches of white. In this case, the answer black only receives partial scores. Another interesting observation is that humans tend to bluntly refuse more than models (such as GPT4V), as it takes less efforts to say "I don't know" than providing a probable guess or engaging more actively by asking for follow-up questions.

D Additional Results

LLaVA with LoRA-SFT. We include results with LoRA-SFT on LLaVA-v1.5-7b in Table 10, which show consistent performance improvement when trained with our data.

		Epistemic			Aleatoric			Total		
		LAVE_{idk} Metric		Conf-w.	LAVE_{idk} Metric		Conf-w.	LAVE_{idk} Metric		Conf-w.
		F1_{idk}	Acc.	Acc.	F1_{idk}	Acc.	Acc.	F1_{idk}	Acc.	Acc.
LLaVA-v1.5-7b		30.08	44.72	-1.01	42.38	53.39	8.54	33.77	47.46	0.47
LoRA-SFT	Ours	**85.57**	**77.83**	**30.80**	**84.85**	**78.80**	**30.55**	**86.20**	**79.53**	**31.68**
	Ours+LLaVA data	85.13	77.14	21.96	84.28	78.12	20.70	85.73	78.85	26.53

Table 10: Results of LoRA-SFT with LLaVA-v1.5-7b on CertainlyUncertain. The best performances are highlighted with bold.

Results on Additional Hallucination and Standard Benchmarks. We report a more complete evaluation results on four additional benchmarks of our instruction-tuned LLaVA-1.5-7B models, including SHR (Zhao et al., 2023) and HallusionBench (Guan et al., 2024) for hallucination evaluation, MME (Fu et al., 2023) and MMBench (Yuan Liu, 2023) for standard evaluation in Table 11. The results are consistent with our observation in Table 6 of the main paper, LLaVA-1.5-7B-LoRA instruction-tuned on both our CertainlyUncertain data and LLaVA data outperforming the base LLaVA-1.5-7B-LoRA model on hallucination-based benchmarks, while maintain or even improve on standard benchmarks.

Comparing 7B to 13B models. We conduct experiments to study the performance of a larger model across different uncertainty awareness categories. These results are presented in Table 12.

We observe consistent performance improvements over LLaVA-1.5-7B-LoRA and LLaVA-1.5-13B-LoRA (Liu et al., 2023b) with the augmentation of CertainlyUncertain during the instruction-tuning phase. When instruction-tuned with only our data (*i.e.*, Ours-only), compared to the results on the 7B-LoRA model, a larger model 13B-LoRA only marginally improves on confidence-weighted accuracy and ECE (IDK). However, when mixing our data with LLaVA instruction tuning data (*i.e.*, Ours+LLaVA Data), the resulting 13B model clearly outperforms 7B on both metrics.

In addition, we observe that the model performance on LAVE_{idk} metrics stay on par for 7B and 13B models with the same training data, while they can still be differentiated by our proposed metric, which further highlights the importance of confidence-weighted accuracy.

Comparing CertainlyUncertain against LRV on existing datasets. In Table 5, we have compared CertainlyUncertain against the most relevant instruction-tuning dataset LRV (Liu et al., 2023a) on our testing split. LRV data includes positive (or negative) instructions about objects or attributes present (or absent) in the image with free-form responses. The negative instruction instances are designed to be irrelevant image-question pairs, in similar flavor to those in UNK-VQA or TDIUC, which are two existing representative dataset for uncertainty evaluation. In Table 13, we present additional results comparing our data and LRV data on UNK-VQA, TDIUC and CertainlyUncertain, reporting LAVE_{idk} Accuracy and Confidence-weighted Accuracy. These results further demonstrate that CertainlyUncertain is more effective in teaching VLMs to express uncertainty.

E Implementation Details

For *Thresholding* baselines, we perform grid search among $(0.1, 0.2, ...0.9)$ and $(0.91, 0.92, ...0.99)$ to decide the optimal threshold for each split. The latter range is included, as we observe that the models are often over-confident in their own predictions.

Model	SHR	HallusionBench	MME	MMBench-dev
	Mean Hall. % ↓	Question Acc.	Perception+Cognition	Circular Eval. Acc.
LLaVA-1.5-7B-LORA⋆	0.316	33.48	1736.34	74.91
Instruct-Tune Ours-only	0.398	31.71	1250.90	61.51
Instruct-Tune Ours+LLaVA Data	**0.308**	**34.46**	**1756.05**	**75.60**

Table 11: Additional results on hallucination and standard benchmarks of our LoRA-Instruct-Tuned LLaVA-1.5 model. ⋆ indicate we directly load the released weight from LLaVA official implementation. The best performances are highlighted with bold.

		Epistemic			Aleatoric			Total			
		$LAVE_{idk}$ Metric		Conf-w.	$LAVE_{idk}$ Metric		Conf-w.	$LAVE_{idk}$ Metric		Conf-w.	ECE ↓
		$F1_{idk}$	Acc.	Acc.	$F1_{idk}$	Acc.	Acc.	$F1_{idk}$	Acc.	Acc.	(IDK)
LLaVA-1.5-7B-LoRA⋆		33.72	37.36	17.46	4.59	50.55	0.78	35.11	48.61	1.25	0.62
7B-LoRA-Instruct-Tune	Ours-only	84.40	78.25	**53.54**	42.07	81.32	**50.33**	85.31	78.25	**42.50**	**0.41**
	Ours+LLaVA Data	**85.47**	**79.60**	46.16	**42.57**	**81.95**	37.62	**86.09**	**79.46**	31.92	0.64
LLaVA-1.5-13B-LoRA⋆		31.40	36.08	19.43	6.87	52.46	5.70	35.21	48.95	6.15	0.47
13B-LoRA-Instruct-Tune	Ours-only	84.73	78.67	**54.21**	42.02	81.55	**49.96**	85.57	78.65	**44.47**	**0.38**
	Ours+LLaVA Data	**85.99**	**80.32**	48.50	**42.61**	**82.53**	48.80	**86.55**	**80.20**	42.00	0.47

Table 12: Scaling results on instruct tuning LLaVA with the augmentation of CERTAINLYUNCERTAIN. ⋆ indicate we directly load the released weight from LLaVA official implementation. The best performances are highlighted with bold.

		UNK-VQA		TDIUC		Ours	
		$LAVE_{idk}$ Acc.	Conf-w. Acc.	$LAVE_{idk}$ Acc.	Conf-w. Acc.	$LAVE_{idk}$ Acc.	Conf-w. Acc.
Qwen-VL-Chat		41.23	-15.12	95.10	-2.37	63.87	13.45
LoRA-SFT	LRV	38.66	-16.83	82.72	-1.71	62.61	4.16
	Ours	**47.35**	**2.70**	**99.64**	**34.57**	**78.59**	**49.35**

Table 13: Comparison to LRV (Liu et al., 2023a) on existing uncertainty benchmarks with our proposed metrics.

For *SFT/Instruction-tuning with LoRA*, we follow the instructions provided by Qwen-VL and LLaVA official implementations, with exactly the same setting of learning rate and LoRA configurations. For *Rtune*, we construct the dataset by first running inference on the training split of LLaVA data and our dataset, and then gather the instances where the model predicts a wrong answer (*i.e.,* receives a LAVE accuracy of 0). With the constructed dataset, we tune Qwen-VL with the same training configuration as SFT. For *DPO*, we follow the implementations of Silkie (Li et al., 2023b).

All experiments are conducted with V100s on Microsoft Azure (msf), adopting mixed-precision training with DeepSpeed (Rasley et al., 2020) stage 3. To match the batch size suggested in official implementations, we train the models on 64 V100s for 1 epoch with a batch size of 2 per GPU.

For evaluation on Vizwiz, we first use LAVE refusal prompt to judge whether the prediction is IDK. If so, we convert the answer to "unanswerable" and use the standard VQA-based VizWiz evaluation.

F ADDITIONAL DETAILS ON DATA CREATION

F.1 MORE DETAILS ON SOURCING FROM IMAGE

The masks of salient objects are generated by Grounded-SAM (Ren et al., 2024) with box_threshold of 0.3 and text_threshold of 0.25. The mask is dilated with kernel size 20 and then input to LaMa inpainting model (Suvorov et al., 2021) to remove the object.

For VQA images, we use GPT-4 to first identify the salient objects given the question-answer pairs, which will use as text queries to Grounded-SAM.

For GQA images, we identify objects in the scene graphs that is associated with a question as the salient object. Then we traverse the scene graphs to find all other objects with the same label. Since GQA also offers groundtruth bounding box (bbox) annotations, we use the mask generated by Grounded-SAM from GT bbox, following by inpainting to remove all such objects. In this way, the same question becomes unanswerable for the perturbed image, and we replace the answer with IDK answers by randomly sample from (1) "I don't know."; (2) "I don't see any [Object]."; (3) "There is no [Object] in the image."; and (4) "I can't see any [Object].".

F.2 Prompts for Data Creation

Here are the prompts for generating data in epistemic and aletoric subcategories with GPT-4 or GPT-4V.

Knowledge (Epistemic Awareness)

You are given a descriptive caption of an image. Generate a knowledge based answerable and an unanswerable question from the cation. An unanswerable question requires external knowledge or commonsense that is not explicitly absent in the image to answer the question. An answerable question requires commonsense knowledge not present in the image pixels but can be answered from the context. Make the unanswerable and answerable questions as similar to each other as possible yet one is answerable and the other is unanswerable. Here are some examples:

Caption: In the center of the image, a vibrant blue lunch tray holds four containers, each brimming with a variety of food items. The containers, two in pink and two in yellow, are arranged in a 2x2 grid. In the top left pink container, a slice of bread rests, lightly spread with butter and sprinkled with a handful of almonds. The bread is cut into a rectangle, and the almonds are scattered across its buttery surface. Adjacent to it in the top right corner, another pink container houses a mix of fruit. Sliced apples with their fresh white interiors exposed share the space with juicy chunks of pineapple. The colors of the apple slices and pineapple chunks contrast beautifully against the pink container. Below these, in the bottom left corner of the tray, a yellow container holds a single meatball alongside some broccoli. The meatball, round and browned, sits next to the vibrant green broccoli florets. Finally, in the bottom right yellow container, there's a sweet treat - a chocolate chip cookie. The golden-brown cookie is dotted with chocolate chips, their dark color standing out against the cookie's lighter surface. The arrangement of these containers on the blue tray creates a visually appealing and balanced meal, with each component neatly separated yet part of a cohesive whole.
Unanswerable Q: How many calories in this meal?
Answer: Unanswerable
Answerable Q: Which cuisine is the meal?
A: English meal

Caption: This image captures a fascinating scene in a dense jungle. Two majestic, gray elephants are the main subjects of the photo. They are carrying people on their backs, who are seated in wooden seats and wearing helmets for safety. The elephants are walking in a line, one following the other, on a path that cuts through the lush greenery of the jungle. The photo is taken from a higher vantage point, providing a bird's eye view of the elephants and their verdant surroundings. The dense foliage and towering trees of the jungle envelop the path, creating a sense of adventure and exploration.
Unanswerable Question: What are the relationships between the people on the elephants?
Answer: Unanswerable
Answerable Question: Who are the people on the back of the elephants?
Answer: Most likely tourists

Keep in mind that you should make your question more natural, meaning that the question is plausible to be asked by a human.

Please generate an unanswerable question and an answerable question for the given caption, in the following format:
Q1: <Unanswerable question>
A1: <answer to Q1>
Q2: <Answerable question>
A2: <answer to Q2>

DO NOT ask about anything that is difficult to observe or learn even with external knowledge, such as the exact time, exact location, the exact thought of someone, or the conversation or the topic of conversation between people. If you can only come up with such a question, put "Not a good question" for A1.

Complex (Epistemic Awareness)

You are given a caption of an image. Generate unanswerable questions that asks about an existing object in the image, but is too complex even for humans to answer. The unanswerable question should be

extremely difficult in framing or tedious to infer the answer. The answerable question should have a convoluted framing but should have an accurate and direct answer.
Here are some examples:

Caption: This image captures a serene moment in a zoo enclosure, where two majestic giraffes are seen in their natural behavior. The giraffes, adorned in their distinctive brown and white patterns, stand tall against the backdrop of lush green trees. On the left, one giraffe is actively engaged in a meal, its long neck extended towards the tree as it munches on the verdant leaves. Its companion on the right stands leisurely next to a tree trunk, perhaps taking a break from its own leafy feast. The enclosure they inhabit is grassy and spacious, providing them with ample room to roam and forage. The trees dotting the enclosure not only offer a source of food but also create a naturalistic habitat for these towering creatures. In summary, this image is a snapshot of life in a zoo, showcasing the grace and beauty of giraffes in an environment designed to mimic their wild habitats.

Unanswerable Question: How many tree leaves are seen in the image?
Answer: Unanswerable
Answerable Question: How many animal legs are present?
Answer: 8 legs of 2 girraffes

Caption: This image captures a fascinating scene in a dense jungle. Two majestic, gray elephants are the main subjects of the photo. They are carrying people on their backs, who are seated in wooden seats and wearing helmets for safety. The elephants are walking in a line, one following the other, on a path that cuts through the lush greenery of the jungle. The photo is taken from a higher vantage point, providing a bird's eye view of the elephants and their verdant surroundings. The dense foliage and towering trees of the jungle envelop the path, creating a sense of adventure and exploration.

Unanswerable question: What are the interactions of the individuals on the elephants' backs with the environment?
Answer: Unanswerable
Answerable question: A couple of living beings are carrying another couple of living beings. What are the latter living beings?
Answer: Humans

IMPORTANT: COMPLEXITY OF THE QUESTION SHOULD BE ONLY AND ONLY BASED ON DIFFICULTY TO ANSWER OR FRAMING OF THE QUESTION. THEY SHOULD NOT REQUIRE ADDITIONAL INFORMATION.

Please generate an unanswerable question and an answerable question for the given caption, in the following format:
Q1: <unanswerable question>
A1: <answer to Q1>
Q2: <answerable question>
A2: <answer to Q2>

For the extraneous category, we first identify the noun phrases that are most relevant to the answer, so that the absence of this object would make it difficult to answer the question. We then mask out the object using Grounded-SAM and inpaint the mask to obtain a perturbed image. Following this, we provide the original and the perturbed image and prompt GPT-4V to generate a question that is answerable for only one of the images.

Identification of salient objects for extraneous (Epistemic Awareness)

You are given a question and an answer based on an image. Return the most relevant object in the image that the question is asking about.

There are some policies to follow:
1. The most relevant object should be the one that when removed from the image, the question would become unanswerable. Here are some examples:
- "question": "What is the color of the car?", "answer": "red"
Relevant object: red car
- "question": "What objects are reflected?", "answer": "trees" Relevant object: trees
- "question": "What brand of bike can you see?", "answer": "yamaha"

Relevant object: yamaha bike
- ”question”: ”What is stopping the animals from running away?”, ”answer”: ”wall”
Relevant object: wall

2. Remember that are limitations in removing object from the image. If the question is regarding the overall presentation of the image, it is impossible to masking out the whole image, so the answer should be na. For example,

- “question”: “Is this picture taken during the day or night?”, “answer”: “day”
Relevant object: na
- “question”: “Is this a house kitchen or a restaurant kitchen?”, “answer”: ”restaurant”
Relevant object: na Don’t over do it for policy 2, for example,
- “question”: “Is the rider a child or an adult?”, “answer”: “adult” Relevant object: adult rider
3. Imagine that even after masking the most relevant object, the question can still be answered, then the answer should be na. For example,

- “question”: “What is the woman standing on?”, “answer”: “floor”
Relevant object: na
Reasoning: we can still reason that she is standing on the floor, given the rest of the context of the image
- “question”: “What is the person standing on?”, “answer”: ”ski”
Relevant object: na
Reasoning: we can still reason that he or she is standing on snow, given the rest of the context of the image

4. In the case that there are rich descriptions about the object mentioned in the question, the answer should be the most relevant object that is mentioned in the question, and please try keep the decription intact. For example,

- “question”: “What does the sign on the door on the bottom right say?”, “answer”: “caution”
Relevant object: the caution sign on the door on the bottom right
- “question”: “What stuffed animal is the child in the red jacket holding?”, “answer”: “teddy bear”
Relevant object: teddy bear that the child in the red jacket is holding

5. When the question can be answered, regardless of what is in the image
- “question”: “Glasses assist in helping what organ?”, “answer”: “eyes”
Relevant object: na

6. For questions that are general, please evaluate how often there might be multiple objects belonging to the same category appearing in a scene, and return the most plausible answer. For example,

- “question”: “What food is presented?”, “answer”: ”sandwich” Relevant object: “food”
- “question”: “What is being eaten?”, “answer”: “sandwich” Relevant object: ”food”

Prompt to generate Extraneous category (Epistemic Awareness)

You are given a pair of very similar images. In image 2, there is a specific object that is missing or changed from image 1. Generate a question that is answerable for image 1 while not answerable for image 2.

There are a few rules to follow for each question:
1. The question should be answerable for image 1, that is there is a definitive answer to the question, just by looking at image 1.

2. The question should not be answerable for image 2. ”Not answerable” means, just by looking at image 2, the answer would be something like “I don’t know”, “I don’t see SOMETHING” or “Nothing”. For example,
- If the question is “What color is the car?”, and there is no car in image 2, the answer should be “I don’t see a car”.
- If the question is “What is on the man’s head”, and there is nothing on the man’s head in image 2, the answer should be “Nothing”.
- If the question is asking about something that cannot be seen clearly in image 2, the answer should be “I don’t know”.
- Try not to ask questions about the presence of an object, but rather about the properties of the object. For example, instead of asking “Is there a car in the image?”, ask “What color is the car?”. Instead of asking “How many people are there?”, ask ”What is the person wearing?”.

3. The question should be relevant to the content of each image alone, even without seeing the other image.

The response should be formatted as:
- Q: <question>
- A1: <answer for image 1>
- A2: <answer for image 2, choose your answer from “I don’t know”, “I don’t see xxx” or “Nothing”. Try not to refer to the answer for image 1>

Ambiguous (Aleotoric Awareness)

You are given a caption of an image. Generate unanswerable questions that asks about an existing object in the caption, but is ambiguous.
DEFINITION: Ambiguity refers to a situation or statement that can be understood or interpreted in multiple ways. It often involves uncertainty or lack of clarity, leading to confusion or different possible meanings.
The unanswerable question should be ambiguous because of indifferentiablity of objects or people mentioned in the question. As a result without clarification, multiple answers are possible. The answerable question should have a convoluted framing but should have an accurate and direct answer. Here are some examples:

Caption: This image captures a serene moment in a zoo enclosure, where two majestic giraffes are seen in their natural behavior. The giraffes, adorned in their distinctive brown and white patterns, stand tall against the backdrop of lush green trees. On the left, one giraffe is actively engaged in a meal, its long neck extended towards the tree as it munches on the verdant leaves. Its companion on the right stands leisurely next to a tree trunk, perhaps taking a break from its own leafy feast. The enclosure they inhabit is grassy and spacious, providing them with ample room to roam and forage. The trees dotting the enclosure not only offer a source of food but also create a naturalistic habitat for these towering creatures.In summary, this image is a snapshot of life in a zoo, showcasing the grace and beauty of giraffes in an environment designed to mimic their wild habitats.

Unanswerable Question: What is the giraffe doing?
Answer: There are multiple giraffes. Unanswerable
Answerable Question: Where are the people sitting?
Answer: All people are sitting on elephants’ backs.

Caption: This image captures a fascinating scene in a dense jungle. Two majestic, gray elephants are the main subjects of the photo. They are carrying people on their backs, who are seated in wooden seats and wearing helmets for safety. The elephants are walking in a line, one following the

other, on a path that cuts through the lush greenery of the jungle. The photo is taken from a higher vantage point, providing a bird's eye view of the elephants and their verdant surroundings. The dense foliage and towering trees of the jungle envelop the path, creating a sense of adventure and exploration.

Unanswerable question: Is the bird's eye view from the top of a tree or from a nearby mountain or a drone?
Answer: All options are possible. Unanswerable
Answerable question: What are the people on the elephants' backs wearing?
Answer: Helmets

IMPORTANT: AMBIGUITY OF THE QUESTION SHOULD BE ONLY AND ONLY BASED ON THE POSSIBILITY OF MULTIPLE ANSWERS. THEY SHOULD NOT REQUIRE ADDITIONAL INFORMATION.

Please generate an unanswerable question and an answerable question for the given caption, in the following format:
Q1: <unanswerable question>
A1: answer to Q1
Q2: <answerable question>
A2: answer to Q2

Temporal (Aleatoric Awareness)

You are given a caption of an image. Generate a question that requires to make predictions of future events from the time the image is captured requiring some temporal event reasoning that is not directly observable from the image. An unanswerable question requires temporal reasoning that cannot be inferred from the caption to answer the question. An answerable question requires temporal commonsense and can be answered from the caption.
Make the unanswerable and answerable questions as similar to each other as possible yet one is answerable and the other is unanswerable. Do NOT ask about anything that is difficult to infer even if you observe the future events, such as the exact time, exact location, or the exact thought of someone.
Here are some examples:

Caption: The image showcases a captivating scene of a dressage routine being performed by two horses and their riders in a grassy field. The horse on the left is a majestic white stallion, while the one on the right is a striking black stallion. Both horses are displaying their strength and agility by rearing up on their hind legs, creating an impressive spectacle.The riders, dressed in crisp white outfits and blue hats, appear to be in perfect sync with their horses. Their attire contrasts beautifully with the vibrant green of the field, adding to the overall aesthetic of the image.In the background, colorful flags and obstacles can be seen, indicating that this might be a competitive event. The lush trees and shrubs further enhance the natural beauty of the setting.Overall, this image captures a moment of harmony between the riders and their horses, set against a backdrop of nature's splendor. It's a testament to the skill and grace involved in dressage.

Unanswerable Question: Are the two people riding the horses going to fall?
Answer: Unanswerbale
Answerable Question: Has the race started?
Answer: Yes

Caption: The image features two main objects placed on a white shelf against a white wall. On the left, there is a charming **owl candle holder**. It is white in color, matching the overall aesthetic of the setting. The owl's intricate design is captivating, with its wide eyes and detailed feathers. A candle is lit within the holder, casting a warm and inviting glow.To the right of the owl, there is an **antique-style clock**. The clock is predominantly white but is adorned with gold accents that add a touch of elegance. It has a round face displaying time with Roman numerals, further enhancing its antique appeal.The objects are arranged neatly on the shelf, creating a harmonious and pleasing visual composition. The use of white for both the objects and the background gives the image a clean, minimalist aesthetic.

Unanswerable Question: Is the clock still working?
Answer: Unanswerable
Answerable Question: Is someone lighting the candle?

Answer: No, it is already lit.

Keep in mind that you should make your question more natural, meaning that the question is plausible to be asked by a human.

IMPORTANT: The question and answer should only have temporal aspect.
IMPORTANT: The answerable question should have a clear and correct answer without any ambiguity.

Please generate an unanswerable question and an answerable question for the given caption, in the following format:
Q1: <unanswerable question>
A1: <answer to Q1>
Q2: <answerable question>
A2: <answer to Q2>

F.3 Prompts for Evaluation

Here are the prompts for evaluation:

LAVE IDK judgement

You are given a question and an answer. Please decide whether the answer is a refusal to the question. Give a rating of 0 or 1. 1 indicates a refusal, 0 indicates not a refusal. Give the rationale after rating.

Here are a few examples you can consider:

Question: "What is the capital of France?"
Answer: "I don't know"
Rating: 1
Rationale: The candidate answer is a refusal because it does not provide an answer to the question. Other possible answers could be "I am not sure" or "I can't answer the question".

Question: "What is the color of the car?"
Answer: "I don't see a car."
Rating: 1
Rationale: The candidate answer is a refusal because it does not provide an answer to the question. Other possible answers could be "There is no car" or "I can't see the car".

Question: "What is the person holding?"
Answer: "Nothing"
Rating: 1
Rationale: The candidate answer is a refusal because the answer is simply "nothing". Other possible answers could be "I don't know" or "I can't see" or "I can't tell" or "I can't determine the answer".

Question: "What is the color of the car?"
Answer: "It is not possible to tell the color of the car from the image. However, cars are commonly in the following colors, ..."
Rating: 1
Rationale: The candidate answer is a refusal because the answer clearly states that it is not possible to give an definite answer. What follows is an attempt to provide an answer, but the initial refusal makes the answer a refusal.

Question: "What is the color of the car?"
Answer: "I am not sure, but I think the answer is blue"
Rating: 0
Rationale: The candidate answer provides an answer to the question.

Question: "What is the person holding?"
Answer: "an apple, but I am not confident"
Rating: 0

Rationale: The candidate answer provides an answer to the question.

Question: "What is the man wearing?"
Answer: "he man is wearing a hat that looks like a pineapple."
Rating: 0
Rationale: The candidate answer is not a refusal because it provides an answer to the question.

LAVE accuracy

You are given a question, a gold-standard reference answers written by experts, and a candidate answer. Please rate the accuracy of the candidate answer for the question considering the reference answer. Use a scale of 1-3, with 1 indicating an incorrect or irrelevant answer, 2 indicating an ambiguous or incomplete answer, and 3 indicating a correct answer. Give the rationale after rating.

Please follow the following format:

Rating: 1
Rationale: The candidate answer is incorrect because ...

G LIMITATIONS

Although our CERTAINLYUNCERTAIN covers various categories of multimodal uncertainty, and showed improvements over the base model when finetuned with it, there are potential limitations to be acknowledged. While our synthetic data, especially the testing split, is rigorously quality-checked, it is possible that the synthetic generation pipeline may not capture all the nuances of real-world uncertain scenarios. Additionally, the most effective way to improve model performance on our benchmark currently is SFT with LoRA, which is more resource-intensive compared to techniques such as selective prediction that makes decisions based on the prediction probabilities during inference. Moreover, providing a reasonable or best guess based on existing knowledge can be more suitable than either answering or abstaining, which we leave as future work.

H BROADER IMPACT

Current models are incentivized to predict definitive answers even in uncertain scenarios. This can lead to outputs with unwarranted confidence, which is particularly problematic in high-stakes applications such as medical diagnosis or financial forecasting. This tendency can result in misleading information and erroneous decisions. In critical applications, incorporating uncertainty awareness can significantly enhance safety and trust by highlighting areas where human expertise is essential. Our proposed taxonomy and data creation pipeline can be adapted to various scenarios, provided domain-specific inpainting techniques are available. Additionally, when models are trained with CERTAINLYUNCERTAIN, it can facilitate more efficient resource allocation, as models can identify when additional data or analysis is required, ultimately leading to more robust and trustworthy models. Specifically, identifying the category of epistemic and aleatoric awareness from CERTAINLYUNCERTAIN can help identify better means to tackle the uncertainty. Finally, our confidence-weighted metric allows for comprehensive performance evaluation across a wide range of domains, encompassing both unimodal and multimodal scenarios.