

# Mind the Uncertainty in Human Disagreement: Evaluating Discrepancies Between Model Predictions and Human Responses in VQA

Jian Lan<sup>1,2</sup>, Diego Frassinelli<sup>1</sup>, Barbara Plank<sup>1,2</sup>

<sup>1</sup> Ludwig Maximilian University of Munich, Germany

<sup>2</sup> Munich Center for Machine Learning (MCML), Munich, Germany

{lan}@dbs.ifi.lmu.de

{frassinelli, bplank}@cis.lmu.de

## Abstract

Large vision-language models struggle to accurately predict responses provided by multiple human annotators, particularly when those responses exhibit high uncertainty. In this study, we focus on a Visual Question Answering (VQA) task and comprehensively evaluate how well the output of the state-of-the-art vision-language model correlates with the distribution of human responses. To do so, we categorize our samples based on their levels (low, medium, high) of human uncertainty in disagreement (HUD) and employ, not only accuracy, but also three new human-correlated metrics for the first time in VQA, to investigate the impact of HUD. We also verify the effect of common calibration and human calibration (Baan et al. 2022) on the alignment of models and humans. Our results show that even BEiT3, currently the best model for this task, struggles to capture the multi-label distribution inherent in diverse human responses. Additionally, we observe that the commonly used accuracy-oriented calibration technique adversely affects BEiT3’s ability to capture HUD, further widening the gap between model predictions and human distributions. In contrast, we show the benefits of calibrating models towards human distributions for VQA, to better align model confidence with human uncertainty. Our findings highlight that for VQA, the alignment between human responses and model predictions is understudied and is an important target for future studies.

**Code, project page** — <https://github.com/mainlp/vqa-hud>

**Extended version** — <https://arxiv.org/pdf/2410.02773>

## Introduction

To strengthen human trust in the task of visual question answering (VQA), large vision-language models should exhibit reliable confidence in their predictions. Regardless of high model accuracy, poor confidence would indicate that “a model does not know when it does not know” (Baan et al. 2022; Pavlick and Kwiatkowski 2019; Yang, Yoo, and Lee 2024). Addressing the importance of model confidence has drawn increasing attention in recent years (Naeini, Cooper, and Hauskrecht 2015; Guo et al. 2017), aiming to make models more reliable and capable of mirroring humans’ uncertainty (Collins et al. 2023; Ilia and Aziz 2024;

Peterson et al. 2019), especially for cases with high human disagreement. We here refer to *human uncertainty in human disagreement* (HUD) as a two-part observation: a) for each sample, different individuals may have different knowledge, and thus it may lead to different (valid) responses (also known as *human label variation* (Plank 2022)); b) the same humans may also have their own degree of uncertainty in responding. HUD is a huge challenge that makes it difficult for models to align with humans (Peterson et al. 2019), as many samples may not have a single ground truth with human agreement and high confidence. This issue is particularly pronounced in visual question answering (VQA), where humans frequently have different opinions, and show uncertainty in their responses. To illustrate the problem, Figure 1 illustrates the task and the responses generated by four human annotators (R). It also provides the confidence (C) of each annotator in the given answer. In the example, the two responses in gray are the same (“blue and gray”) but with different human confidence labels (1.0 vs. 0.5), while those in green and blue have different response labels with the same human confidence labels.

Traditionally, a model would predict the response with the highest probability among the prediction distribution (e.g., “blue and gray” in Figure 1) but it would miss the effect that human confidence plays on the prediction distribution. It will also ignore the existence of multiple valid responses to the same question. For example, “blue” (0.22) is less probable in the model prediction than “blue and gray” (0.72), but, in reality, it has a higher human confidence (0.36 vs. 0.27) and is indeed a highly possible answer. Therefore, training a model to predict the most probable correct answer, while ignoring HUD, does not allow the model to accurately reflect human behavior and distributions in real-world scenarios. It also makes models struggle to assign reliable probabilities to other equally valid answers. Most of the previous studies on VQA have not explicitly addressed HUD, and have predominantly evaluated models based on the accuracy of a single prediction of each sample alone.

This study examines how well state-of-the-art VQA models align with human confidence distributions when HUD exists. Unlike previous work on VQA, we propose to explicitly utilize HUD information when evaluating the discrepancies between model predictions and human responses. Specifically, we use VQA 2.0 (Goyal et al. 2017), a dataset

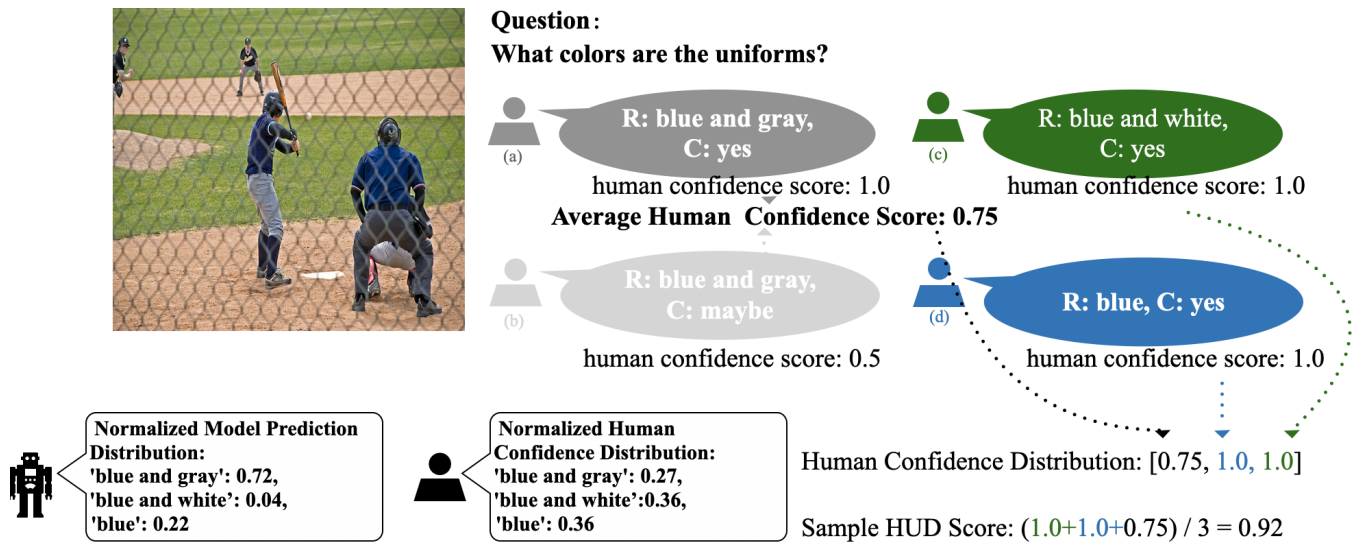


Figure 1: An example of VQA 2.0. We show different people’s annotation in different colors, and compare the model predictions with human distributions. Please note there are 10 annotators for each sample in VQA 2.0, we only use four in the figure as a clear and easy display, where R is human response, and C is the confidence label.

where ten annotators answered each question and also indicated their confidence level in their answer. In this study we aim to answer the following two research questions: **RQ1:** To what extent do different levels of HUD impact the accuracy of the model and its alignment with human confidence? **RQ2:** Does calibration improve the human-model alignment?

To answer these questions, we start by adopting a splitting strategy, categorizing samples into three sets (low, medium, high) based on their HUD levels. We then employ not only traditional accuracy but also propose to use three human-correlated metrics for VQA: Total Variation Distance (*TVD*), Kullback-Leibler divergence (*KL*), and Human Entropy Calibration Error (*EntCE*). These new measures allow us to comprehensively analyze the discrepancies between models and human in the context of HUD. Furthermore, we apply Temperature Scaling (Guo et al. 2017) to verify the effectiveness of the traditional calibration method.

Our results show that even state-of-the-art models like BEiT3 (Wang et al. 2023) struggle to capture the multi-label distribution caused by diverse human responses. Additionally, we observe that commonly used accuracy-oriented calibration techniques adversely affect BEiT3’s confidence, further widening the gap between model prediction distributions and human confidence distributions. Instead, we propose to calibrate models towards human distributions, and thus better align models with humans under HUD. Our key contributions are: **1)** To the best of our knowledge, we are the first to study human uncertainty in disagreement in VQA by investigating the gap between the SOTA model prediction distributions and human response distributions; **2)** We implement a novel comprehensive evaluation of VQA models, evaluating on various human-centered metrics rather than solely relying on VQA accuracy. We demonstrate that, when intrinsic HUD exists, the VQA accuracy cannot adequately

indicate model’s capabilities; finally, **3)** We demonstrate that the variations in HUD levels affect model performances differently. We therefore advocate for an increased future attention on improving the alignment between models and humans under HUD.

## Related Work

**Human Disagreement:** Human disagreement, also known as human label variation (Plank 2022), has been studied in various tasks including natural language inference (Baan et al. 2022), machine translation (Popović 2021), question answering (Kwiatkowski et al. 2019), image classification (Peterson et al. 2019), and social computing (Gordon et al. 2021). All these studies are based on datasets where multiple human annotators give *different plausible answers* for each sample (Nie, Zhou, and Bansal 2020). They then either aggregate human labels, using majority vote for training, or use human soft labels (Peterson et al. 2019), optimizing models towards a human distribution based on label frequency. Human uncertainty here means uncertainty reflected by annotators choosing different labels, and is different from annotators providing indicators of uncertainty in their responses. So far, human uncertainty in responses is an understudied phenomenon in relation to human label variation.

**Visual Question Answering:** The VQA task was first proposed by Antol et al. (2015), where they provide a benchmark including the VQA 1.0 dataset and the tradition evaluation metric VQA-Accuracy (VQA-Acc). VQA 2.0 (Goyal et al. 2017) is based on VQA 1.0, where commonsense questions are removed with the aim to force the model to uniquely rely on the image content for answering.

For this task, large pre-trained vision-language models have become mainstream (Tan and Bansal 2019; Li et al.

2022; Bao et al. 2022; Wang et al. 2023) and have been further improved by optimizing for VQA-Accuracy. Among them, LXMERT (Tan and Bansal 2019) is a comparatively light model which provides the first baseline in the era of pre-trained models. It yields a VQA-Accuracy of 72.5 on the VQA 2.0 test set. BEiT3 (Wang et al. 2023) is the latest and strongest SOTA model with a VQA-Accuracy of 84.03. In VQA, human disagreement is studied implicitly by Jolly, Pezzelle, and Nabi (2021), where they propose a diagnose tool to identify a sample’s difficulty level by calculating the number of different response labels among human annotators and clustering the semantics of responses based on Word2Vec (Mikolov et al. 2018). However, to the best of our knowledge, human uncertainty in the responses has not been yet studied in VQA, where most work only focuses on VQA-Accuracy as the sole evaluation metric.

**Uncertainty in VQA:** VQA is a complex task involving both vision and language and, for this reason, it is normal that a human annotator might be uncertain in providing a label. The level of uncertainty in the responses indicates that the annotators are not completely sure about a given label. Unfortunately, only very few datasets explicitly collect and release annotation on human uncertainty. Two positive examples are VQA 2.0 (Goyal et al. 2017) and VizWiz (Gurari et al. 2018), two widely used datasets for VQA. They contain both labels and uncertainty values generated by ten different human annotators (“How certain are you: yes, maybe, no?”).

Despite the availability of these datasets, previous studies on VQA (Antol et al. 2015; Goyal et al. 2017; Gurari et al. 2018; Tan and Bansal 2019; Wang et al. 2023) have not directly addressed human uncertainty presented in the responses. They simply trained models towards the most frequent label, optimizing for the highest VQA-Accuracy. Only recently, model uncertainty towards multiple answers have become an emerging direction (Yang, Yoo, and Lee 2024), where the community is focusing on measuring the level of uncertainty of large language models (Xiong et al. 2024). However, the main challenge is that, still, human uncertainty information is not often available, making it difficult to evaluate how models correlate with humans.

## Human Uncertainty in Disagreement

We first introduce the VQA 2.0 samples used in our study, and then discuss how we calculate the HUD scores and group the samples in three HUD levels based on their scores.

### Data

In this study we use VQA 2.0 (Goyal et al. 2017), one of the most commonly used datasets for VQA. Given that both BEiT3 (Wang et al. 2023) and LXMERT (Tan and Bansal 2019) are pre-trained and fine-tuned only on VQA 2.0, this is the perfect dataset for a fair comparison between these two models.

As shown in Figure 1, given an image-question-answer triplet  $t = (i, q, \mathcal{A})$ , where  $i$  is an image,  $q$  is a question, and  $\mathcal{A}$  is an answer set,  $\mathcal{A}$  consists of 10 independent humans’ annotations. In each annotation  $h_n$ ,  $n = 1, \dots, 10$ , every

annotator gives their response  $r_n$  and also a confidence level  $c_n$ . The level  $c_n$  corresponds to one of three pre-defined categories  $\langle 'yes', 'no', 'maybe' \rangle$ , indicating whether an annotator is confident in their answer.

### The HUD Score

Since the confidence labels ‘yes’, ‘no’, and ‘maybe’ are expressed in natural language, they can not be directly utilized for comparisons with model prediction distributions. To quantify human uncertainty, we assign different human confidence scores to each response based on its confidence label, as shown in Figure 1. We quantify every ‘yes’ as 1.0, ‘no’ as 0.01, and ‘maybe’ as 0.5 respectively. We want the values of the three categories to be restricted between 0 and 1, as this range aligns with the distribution of the model’s predicted probabilities. Additionally, we assume that the value of ‘maybe’ is the average score of ‘no’ and ‘yes’, ensuring it does not bias towards either end. We assign a very small (non-zero) value to ‘no’ to avoid the situation of a distribution of all zeros for which entropy cannot be calculated.

When multiple annotators provide the same response (e.g., ‘blue and gray’ in Figure 1), we compute the average confidence score (e.g., 0.75) to quantify the degree of human uncertainty for that response. The HUD score is then determined by averaging these human confidence scores across responses to the same question.

### Grouping samples based on HUD Levels

Our hypothesis is that a model would perform differently according to different HUD levels of the samples. For this reason, we divide the samples into three equal portions based on their HUD scores, namely low human uncertainty set (we use ‘low set’ in short for the rest of the paper), medium human uncertainty set (‘medium set’), and high human uncertainty set (‘high set’). By this design, humans are more certain about the samples in low set, while more uncertain in the high set.

### Model Confidence Scores and Calibration

Here we discuss how we obtain the model prediction distributions and how to better align them with humans.

### Model Prediction Distribution

Given a sample  $x$ , the prediction distribution of a model  $\mathcal{M}$ , denoted as  $\mathcal{M}(x)$ , corresponds to the probabilities  $P_Y(Y = y|X = x)$  the model assigns to each class  $y$ . The most standard and most commonly used method to gain  $P_Y$  is to extract the last hidden state of the neural network for each label, and use the *softmax* function, which converts a vector of raw scores (logits), into a probability distribution over multiple classes to get the final normalized probability distribution:  $P_Y = \text{softmax}([l_1, l_2, \dots, l_K])$ , where  $K$  is the number of all labels.

### Calibration Method

Calibration aims to adjust the predicted probabilities  $P_Y$  towards a more reliable distribution (Guo et al. 2017; Desai

and Durrett 2020; Jiang et al. 2021). A well calibrated multi-class model accurately captures the true likelihood of predictions for all possible classes (Baan et al. 2022; Vaicenavicius et al. 2019; Kull et al. 2019). For instance, in VQA, if there are four possibly correct answers for a question and the humans’ uncertainty scores for them are: [0.60, 0.30, 0.05, 0.05] respectively, a well calibrated model  $\mathcal{M}$  should also have a prediction distribution approximately close to [0.60, 0.30, 0.05, 0.05] for same four labels.

Among different calibration techniques, Temperature Scaling (Guo et al. 2017) is one of the most widely used and most traditional methods. Temperature Scaling is a post-processing technique applied to the logits of a network by dividing them by a temperature parameter  $T > 0$ . Then we can use the modified logits and the *softmax* function to get a calibrated distribution:  $P_{Y=y} = \frac{e^{(t_y/T)}}{\sum_{j=1}^K e^{t_j/T}}$ . In this work, we employ Temperature Scaling to test if this traditional calibration method is effective in our experiments, whether it is more effective for either of the two models, and indicate what a good calibration strategy is.

## Experiments

**Models:** Our study focuses on BEiT3 (Wang et al. 2023), the latest SOTA model, and on LXMERT (Tan and Bansal 2019), the previous SOTA and also one of the most commonly used models. We compare their performance on different HUD sets.

**Fine-Tuning Set and Implementation Details:** The open-sourced training set for VQA 2.0 includes 443,757 samples, and the validation set has 213,954 samples. The test set is not open-source and the human uncertainty labels are not available. Therefore, we exclude the test set from our experiments. We follow exactly the same implementation details provided by BEiT3<sup>1</sup> and LXMERT<sup>2</sup>, using their checkpoints (for BEiT3, we use the BEiT3-base model) of the pre-trained models and fine-tune them following the original instructions. More implementation details are provided in the Appendix. It is essential to point out that BEiT3 and LXMERT use slightly different sets for fine-tuning and validation. They partitioned the original validation set in a customized manner, where BEiT3 keeps 5,303 samples from the original validation set and add the remaining set together with the training set for fine-tuning, while LXMERT keeps 25,994 samples from it and do the same. Despite this difference, previous studies have successfully compared the performance of these two models. For this reason, we also follow the settings described in the original papers.

**Validation Set Details:** We are left with a set of 3,248 samples for BEiT3, where 1,083 samples are in the low, 1,083 in the medium, and 1,082 in the high set respectively.

Similarly, we have a total of 15,408 samples for LXMERT, with 5,136 samples for each HUD set.

For BEiT3, the average number of distinct human responses per sample in the low set is 3.06, with an average

HUD score of 0.98. In the medium set, these values are 3.97 and 0.86, respectively. While in the high set, they are 4.82 and 0.64, respectively. Similarly, for LXMERT, the values are: 3.08 and 0.98 for the low set, 4.04 and 0.86 for the medium set, and 4.67 and 0.64 for the high set. Figure 2 shows two very similar distributions of HUD scores for the two validation sets, where they have the same mean value and standard variation value (std). The two black lines in both figures show the split boundaries of the low, medium, and high sets. We point out that even though our HUD set split strategy causes seemingly high confidence scores on each set, we target on revealing the differences in model performances caused by three sets, where humans have comparatively high, medium, and low uncertainty. Moreover, since the data distributions are similar, the performance of the two models remain comparable, even though our main focus is not to directly compare their performances.

**Evaluation metrics:** To conduct our evaluation, we include the following metrics (the corresponding equations except VQA-Accuracy can be found in the Appendix.

**VQA-Accuracy** (Antol et al. 2015): Given a model prediction, VQA-Accuracy is defined as:

$$\text{Acc}(\text{ans}) = \min \left\{ \frac{\#\text{humans that said ans}}{3}, 1 \right\}. \quad (1)$$

VQA-Accuracy takes into account human label frequency, assigning higher scores to answers annotated by a greater number of human annotators. It is maximized (1.0), if at least 3 raters gave the exact answer. The number 3 is a manually set parameter proposed in the original work (Antol et al. 2015).

However, it does not consider human uncertainty scores. In this evaluation scenario, a model response overlapping with an annotator response with a very low confidence would still be considered correct.

**KL-divergence (KL;** Kullback and Leibler (1951)) measures how the model probability distribution diverges from human’s, denoted as  $\text{KL}(\text{Human}||\text{Model})$ . Given a human distribution, the model distribution is the approximate distribution, and the KL score indicates how much the approximation distribution deviates from the true distribution.

**Total Variation Distance (TVD;** Devroye and Lugosi (2001)) measures the absolute difference between the model’s and human’s probability distributions, to indicate how much overlap there is between two distributions. In Baan et al. (2022), this measure was proposed as Human Distribution Calibration Error.

**Human Entropy Calibration Error (EntCE;** Baan et al. (2022)) measures the absolute difference between human’s and a model’s entropy of the their distribution, and indicate the average confusion difference in two distributions.

## Results and Discussion

**Before Calibration.** Table 1 reports the performances of the two models (without calibration) across three different levels of human uncertainty, as well as for the overall evaluation set. Performance metrics include VQA-Accuracy, TVD,

<sup>1</sup><https://github.com/microsoft/unilm/tree/master/beit3>

<sup>2</sup><https://github.com/airsplay/LXMERT>




Figure 2: The data distributions of the two validation sets based on HUD scores. We show the two split boundaries using the black lines, the standard deviations in yellow lines, and the mean values in red lines.

Model	HUD set split	VQA-Acc $\uparrow$	TVD $\downarrow$	KL(H  M) $\downarrow$	EntCE $\downarrow$
LXMERT	All	0.671	0.448	1.886	0.555
	Low	<b>0.761</b>	0.454	2.237	0.580
	Med	0.675	0.450	1.786	0.579
	High	0.576	<b>0.439</b>	<b>1.634</b>	<b>0.506</b>
BEiT3	All	0.842	0.436	1.701	0.541
	Low	<b>0.921</b>	0.455	2.150	0.575
	Med	0.845	0.437	1.544	0.546
	High	0.759	<b>0.416</b>	<b>1.408</b>	<b>0.503</b>

Table 1: Model performances on VQA-Accuracy and three human-correlated metrics before calibration. For each model, the **best results** are highlighted in bold. HUD: High means human express high uncertainty for their responses.

KL, and EntCE. Overall, we observe that BEiT3 consistently outperforms LXMERT in accuracy (0.67 vs. 0.84), thereby confirming that it is the better model in terms of accuracy. When zooming into the three HUD levels, both models show a similar trend: as human uncertainty level increases from low to high, their accuracy decreases. The models obtain the highest accuracy level for the stimuli in the low HUD set, as humans provide answers with high confidence to the stimuli in that category. The continuous decline in both models’ VQA-Accuracy performance supports our hypothesis that the higher the HUD level is (more uncertainty), the more challenging it is for a model to predict the best answer.

Besides accuracy, we look at model confidence and compare it against human label variation (which is now across raters). We provide models’ performances on TVD, KL, and EntCE to reveal the extent to which the model prediction distributions approximate humans’. When we compare the two models, both on the overall set and on the three subsets, the results on all three metrics show that BEiT3 correlates better with humans than LXMERT. For both models, we surpris-

ingly find they achieve their best performance on the high set (high human uncertainty). We provide an in-depth analysis as follows.

As introduced in the validation set details, both models have an average of approximately 3.0, 4.0, and 4.8 distinct human response labels per sample in the low, medium, and high sets, with corresponding HUD scores of 0.98, 0.86, and 0.64. This indicates that higher human uncertainty corresponds to higher levels of disagreement among responses. We believe the reason both models have the highest scores on the low set is that the samples have comparatively fewer human response labels, and thus it is easier for a model to learn a ‘seemingly most correct answer’ but ignore the other possibilities. This is also why models are having higher accuracy on the low set; they learn to predict the most correct answer which is supported by the human majority vote, but they do not learn to align well with the overall human answer distribution, i.e., on all the other possible answers. On the contrary, for samples from the high human uncertainty set, there are comparatively more response labels with a smoother confidence distribution (less abrupt variations between the confidence scores on each label), and this helps the models to pay attention to possible labels rather than focusing on one single label. We further showcase and analyze this in the Case Study.

**Calibration towards VQA-Accuracy** Table 2 shows the effectiveness of Temperature Scaling. The effect of TS is traditionally evaluated by **Expected Calibration Error (ECE)** (Guo et al. 2017), which measures the absolute difference between the accuracy of the predictions of a model and its confidence towards these predictions. This evaluation quantifies how well the model’s predicted probabilities are calibrated towards its accuracy. In Table 2 we report the ECE results before and after TS. We empirically set the Temperature  $t$  by comparing the  $t$  from 0.1 to 2.0 with an in-

Model	HUD set split	ECE Before ↓	Calibration towards VQA-Acc				Calibration towards Human Distributions		
			ECE After ↓	TVD ↓	KL ↓	EntCE ↓	TVD ↓	KL ↓	EntCE ↓
LXMERT	All	0.066	0.065	0.442	1.771	0.542	0.354	0.731	0.353
	Low	0.072	0.070	0.449	2.103	0.569	0.369	0.877	0.397
	Med	0.069	0.069	0.444	1.674	0.564	0.347	0.667	0.351
	High	0.071	0.066	0.433	1.534	0.492	0.345	0.649	0.310
BEiT3	All	0.094	0.039	0.488 (-)	3.232 (-)	0.670 (-)	0.340	0.647	0.333
	Low	0.084	0.039	0.496 (-)	3.999 (-)	0.674 (-)	0.371	0.837	0.389
	Med	0.099	0.042	0.494 (-)	2.982 (-)	0.683 (-)	0.331	0.563	0.318
	High	0.098	0.046	0.473 (-)	2.712 (-)	0.655 (-)	0.319	0.542	0.292

Table 2: Model performances on ECE, TVD, KL(Human || Model), and EntCE after calibration. We compare the calibration direction towards accuracy and towards human. We use (-) to indicate when calibration negatively affects the model performance.

terval of 0.05. We set  $t$  to 1.05 for LXMERT and 0.6 for BEiT3, where they each have the lowest ECE on the overall set after TS. We then fix  $t$  and report the results for TVD, KL, and EntCE. As expected, the ECE scores of all sets drop after calibration. However, using TS towards VQA-Accuracy undermines BEiT3’s confidence (indicated by the minus symbol ‘-’), and the scores for TVD, KL, and EntCE increase e.g., from 1.701 to 3.232 on KL for BEiT3. On the contrary, calibrating towards VQA-Accuracy only helps LXMERT improve the correlation with the three human distributions. We conclude that the traditional calibration technique still works for slightly weaker models like LXMERT, but it does not help the latest strong model BEiT3. Besides, here we do not further compare the ECE results between different sets or between models, but only use ECE to test if using TS to calibrate towards VQA-Accuracy works and how strongly this technique influences the models’ correlation with human. ECE does not consider human disagreement and distributions, and thus we do not further analyze ECE results.

**Calibration towards human distributions.** A crucial aspect of this work involves using TS as a calibration method to align with human distributions by comparing scores using TVD, KL, and EntCE. Similar to what we introduced in the previous paragraph, we empirically set the temperature  $t$  to be 2.0 for both models, as it helps model reach better results. Normally,  $t$  is not supposed to be a large value since a too high value makes the predicted probabilities become closer to a uniform distribution, reducing the differences between the classes, which is not desired. As shown in Table 2 on the right-hand side, on the overall set and also three subsets, calibrating towards human distributions rather than VQA-Accuracy helps both models reach lower scores (better results) on TVD (e.g. from 0.448 to 0.354 on LXMERT, from 0.436 to 0.340 on BEiT3), KL (e.g. from 1.886 to 0.731 on LXMERT, from 1.701 to 0.647 on BEiT3), and EntCE (e.g. from 0.555 to 0.353 on EntCE, from 0.541 to 0.333 on BEiT3) compared with the original results in Table 1. Notably, TVD and EntCE measure the average discrepancies

between two distributions (on each element in the vector), but are not sensitive to element class rankings. KL is sensitive to class rankings, where small changes in probabilities can lead to large increases in the divergence. Therefore, we observe a drastic decrease in KL.

On each set split, BEiT3 still outperforms LXMERT on all three metrics, while both models reach the best performances on the high set, and the worst performances on low set. We conclude that, even though BEiT3 has a seemingly strong performance on all metrics, it is not yet correlated well with humans, especially on the low HUD set. Moreover, by calibrating towards human distributions, it even yields very similar human correlation performances when compared with LXMERT. This indicates that BEiT3 is simply better optimized towards human majority vote based on label frequency, but it is not much better at learning true human distributions. In other words, while BEiT3 is the better model in terms of accuracy compared to LXMERT, it is also less well-aligned to human preferences. This opens up for interesting future research directions on how to strike a good balance between training a high-accuracy model and one that is well aligned with humans.

### Case Study

Figure 3 and figure /reffig4showcase two samples from the low (A(1)-A(3)) and high (B(1)-(B3)) sets and compare models’ specific predictions against humans. In line with what we discussed in the last section, on the low set (see A(1)), the model approximates only towards the most correct answer ‘green’, which has the highest VQA-Accuracy. However, besides ‘green’, in the human distribution, ‘yellow’ and ‘blue’ also have similar weights. This means the model should predict all three answers with similar distributions to mirror humans’ confidence in the real world. However, both BEiT3 and LXMERT show an extremely high probability for ‘green’, and very low probabilities for ‘yellow’ and ‘blue’.

When calibrating towards VQA-Accuracy, in A(2) we see a slight improvement in LXMERT. Once again, this traditional calibration method negatively affects the performance

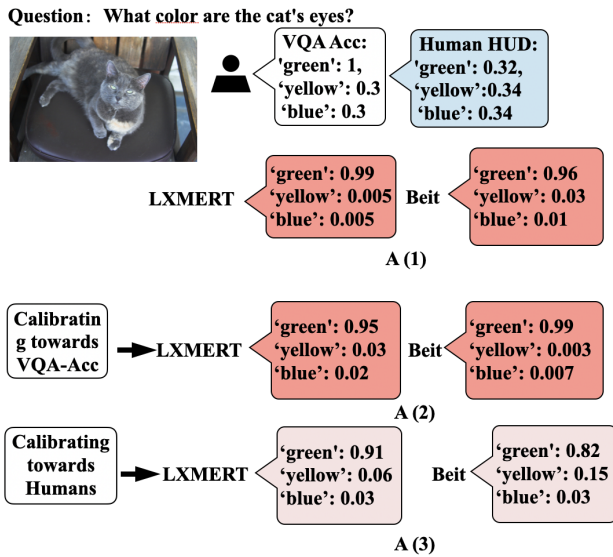


Figure 3: Case Study. The A(1)-A(3) present a sample from the **low set**, displaying each response label’s VQA Accuracy score, Human HUD Distribution, and the prediction distributions of LXMERT and BEiT3. We report all three situations, where models are: not calibrated, calibrated towards VQA-Accuracy, and calibrated towards human distributions.

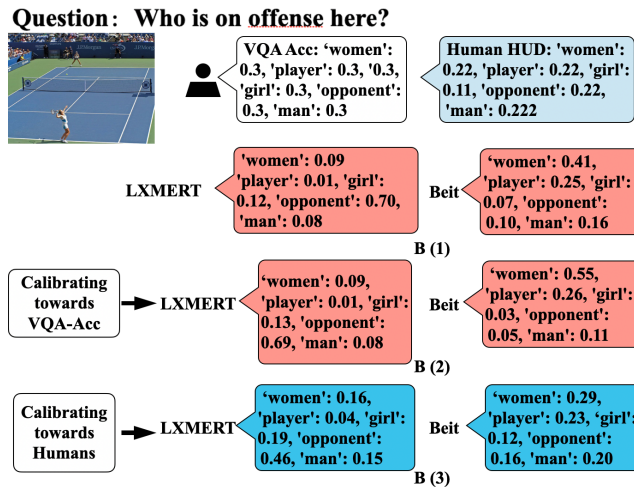


Figure 4: Similarly, the (B(1)-B(3)) present a sample from the **high set**. Human HUD distributions, which are used for measuring human-model correlation, are colored in light blue. In red we highlight unsatisfactory model performances, in light red models that slightly improved but are still not good, while dark blue indicates much better correlation with humans.

of BEiT3. On the contrary, when calibrating towards humans (see A(3)), there is a continuous improvement in both models, even though they are still far from human performance. In B(1), the sample from the high set contains more human response labels, each with lower but similar VQA-

Accuracy scores. In this case, BEiT3 distributes the probability mass across most of the labels provided without the huge gap seen in A(1) (0.96 vs 0.01). Although LXMERT also assigns probability scores to multiple labels, it strongly favors one label ('opponent') with a probability much higher than the others. In B(2), calibrating towards VQA-Accuracy causes tiny changes on the labels 'girl' and 'opponent', while it again does harm BEiT3, causing much higher values for 'women', while decreases the probabilities for 'opponent' and 'man'. On the contrary, calibrating towards humans work well for both models, making them both correlate better with humans.

Finally, by comparing the results in A(3) and in B(3) we see how even when calibrating towards humans, the models still predict one specific label (e.g., 'green') while mostly ignoring the other candidates (e.g., 'yellow' and 'blue'). This clearly shows why model responses correlate better with humans on the high set.

## Conclusion

In this work, we study human uncertainty in Visual Question Answering (VQA). We find that when humans exhibit high levels of uncertainty in disagreement, the models have difficulty in predicting the most correct answer, but are actually better at correlating with human confidence distributions. We further find that the traditional calibration strategy towards accuracy does not work for BEiT3, the latest SOTA model. Instead, we then demonstrate that calibrating it towards human distribution is more effective. We conclude that evaluating VQA models uniquely on the VQA-Accuracy metric is not sufficient, and future studies should focus more on aligning models with human uncertainty and disagreement.

## Limitations

Our study has three limitations. Firstly, while being highly reproducible, our study is restricted to two slightly different validation sets, even though we believe they are very comparable. Future studies could extend our work and fine-tune and validate models on more datasets. Secondly, it is not feasible to explore all the possible strategies to quantify human uncertainty labels. Although our study provides a very standard solution (mapping 'yes', 'maybe', 'no' to 1, 0.5, and 0.01 respectively), this represents a good start point; we advocate for future research to explore more advanced quantification strategies. Thirdly, in our work we use Temperature Scaling; future studies should consider other calibration techniques to directly target human uncertainty in disagreement.

## Ethics Statement

We anticipate no ethical concerns with this work. We utilized open-sourced data and models, which have been cited.

## Acknowledgments

BP is supported by the European Research Council (ERC) Consolidator Grant DIALECT 101043235.

## References

- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, 2425–2433.
- Baan, J.; Aziz, W.; Plank, B.; and Fernandez, R. 2022. Stop Measuring Calibration When Humans Disagree. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 1892–1915. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Bao, H.; Wang, W.; Dong, L.; Liu, Q.; Mohammed, O. K.; Aggarwal, K.; Som, S.; Piao, S.; and Wei, F. 2022. Vlmo: Unified vision-language pre-training with mixture-of-modality-experts. *Advances in Neural Information Processing Systems*, 35: 32897–32912.
- Collins, K. M.; Barker, M.; Espinosa Zarlenga, M.; Raman, N.; Bhatt, U.; Jamnik, M.; Sucholutsky, I.; Weller, A.; and Dvijotham, K. 2023. Human Uncertainty in Concept-Based AI Systems. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 869–889. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702310.
- Desai, S.; and Durrett, G. 2020. Calibration of Pre-trained Transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 295–302. Online: Association for Computational Linguistics.
- Devroye, L.; and Lugosi, G. 2001. *Combinatorial methods in density estimation*. Springer Science & Business Media.
- Gordon, M. L.; Zhou, K.; Patel, K.; Hashimoto, T.; and Bernstein, M. S. 2021. The disagreement deconvolution: Bringing machine learning performance metrics in line with reality. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–14.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6904–6913.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On calibration of modern neural networks. In *International conference on machine learning*, 1321–1330. PMLR.
- Gurari, D.; Li, Q.; Stangl, A. J.; Guo, A.; Lin, C.; Grauman, K.; Luo, J.; and Bigham, J. P. 2018. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3608–3617.
- Ilija, E.; and Aziz, W. 2024. Predict the Next Word: <Humans exhibit uncertainty in this task and language models ----->. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, 234–255. St. Julian's, Malta: Association for Computational Linguistics.
- Jiang, Z.; Araki, J.; Ding, H.; and Neubig, G. 2021. How Can We Know When Language Models Know? On the Calibration of Language Models for Question Answering. *Transactions of the Association for Computational Linguistics*, 9: 962–977.
- Jolly, S.; Pezzelle, S.; and Nabi, M. 2021. EaSe: A Diagnostic Tool for VQA based on Answer Diversity. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2407–2414. Online: Association for Computational Linguistics.
- Kull, M.; Perello Nieto, M.; Kängsepp, M.; Silva Filho, T.; Song, H.; and Flach, P. 2019. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with dirichlet calibration. *Advances in neural information processing systems*, 32.
- Kullback, S.; and Leibler, R. A. 1951. On information and sufficiency. *The annals of mathematical statistics*, 22(1): 79–86.
- Kwiatkowski, T.; Palomaki, J.; Redfield, O.; Collins, M.; Parikh, A.; Alberti, C.; Epstein, D.; Polosukhin, I.; Devlin, J.; Lee, K.; Toutanova, K.; Jones, L.; Kelcey, M.; Chang, M.-W.; Dai, A. M.; Uszkoreit, J.; Le, Q.; and Petrov, S. 2019. Natural Questions: A Benchmark for Question Answering Research. *Transactions of the Association for Computational Linguistics*, 7: 452–466.
- Li, C.; Xu, H.; Tian, J.; Wang, W.; Yan, M.; Bi, B.; Ye, J.; Chen, H.; Xu, G.; Cao, Z.; Zhang, J.; Huang, S.; Huang, F.; Zhou, J.; and Si, L. 2022. mPLUG: Effective and Efficient Vision-Language Learning by Cross-modal Skip-connections. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 7241–7259. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Mikolov, T.; Grave, É.; Bojanowski, P.; Puhersch, C.; and Joulin, A. 2018. Advances in Pre-Training Distributed Word Representations. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Naeini, M. P.; Cooper, G.; and Hauskrecht, M. 2015. Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the AAAI conference on artificial intelligence*.
- Nie, Y.; Zhou, X.; and Bansal, M. 2020. What Can We Learn from Collective Human Opinions on Natural Language Inference Data? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 9131–9143. Online: Association for Computational Linguistics.
- Pavlick, E.; and Kwiatkowski, T. 2019. Inherent disagreements in human textual inferences. *Transactions of the Association for Computational Linguistics*, 7: 677–694.
- Peterson, J. C.; Battleday, R. M.; Griffiths, T. L.; and Rusakovsky, O. 2019. Human uncertainty makes classification more robust. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9617–9626.
- Plank, B. 2022. The “Problem” of Human Label Variation: On Ground Truth in Data, Modeling and Evaluation. In *Proceedings of the 2022 Conference on Empirical Methods in*

*Natural Language Processing*, 10671–10682. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.

Popović, M. 2021. Agree to Disagree: Analysis of Inter-Annotator Disagreements in Human Evaluation of Machine Translation Output. In *Proceedings of the 25th Conference on Computational Natural Language Learning*, 234–243. Online: Association for Computational Linguistics.

Tan, H.; and Bansal, M. 2019. LXMERT: Learning Cross-Modality Encoder Representations from Transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 5100–5111. Hong Kong, China: Association for Computational Linguistics.

Vaicenavicius, J.; Widmann, D.; Andersson, C.; Lindsten, F.; Roll, J.; and Schön, T. 2019. Evaluating model calibration in classification. In *The 22nd international conference on artificial intelligence and statistics*, 3459–3467. PMLR.

Wang, W.; Bao, H.; Dong, L.; Bjorck, J.; Peng, Z.; Liu, Q.; Aggarwal, K.; Mohammed, O. K.; Singhal, S.; Som, S.; et al. 2023. Image as a foreign language: Beit pretraining for vision and vision-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19175–19186.

Xiong, M.; Hu, Z.; Lu, X.; LI, Y.; Fu, J.; He, J.; and Hooi, B. 2024. Can LLMs Express Their Uncertainty? An Empirical Evaluation of Confidence Elicitation in LLMs. In *The Twelfth International Conference on Learning Representations*.

Yang, Y.; Yoo, H.; and Lee, H. 2024. MAQA: Evaluating Uncertainty Quantification in LLMs Regarding Data Uncertainty. arXiv:2408.06816.