Why Does a Visual Question Have Different Answers?

Nilavra Bhattacharya*, Qing Li+, Danna Gurari*

* University of Texas at Austin, + University of California, Los Angeles

Abstract

Visual question answering is the task of returning the answer to a question about an image. A challenge is that different people often provide different answers to the same visual question. To our knowledge, this is the first work that aims to understand why. We propose a taxonomy of nine plausible reasons, and create two labelled datasets consisting of ~45,000 visual questions indicating which reasons led to answer differences. We then propose a novel problem of predicting directly from a visual question which reasons will cause answer differences as well as a novel algorithm for this purpose. Experiments demonstrate the advantage of our approach over several related baselines on two diverse datasets. We publicly share the datasets and code at `https://vizwiz.org`.

1. Introduction

Visual question answering (VQA), the task of returning the answer to a question about an image, is of widespread interest across both industry and academia. For example, many people who are blind find existing VQA solutions are indispensable assistants for providing answers to their daily visual questions; e.g., they use vision-based assistants such as VizWiz [7] and BeSpecular [1] to snap a photo with their mobile phones and then receive answers from remote workers. In the artificial intelligence research community, the VQA problem has emerged as an iconic challenge for emulating a human's vision and language capabilities [5, 25].

Despite the tremendous social impact and progress with VQA solutions, a limitation is that most solutions lack a way to handle when a visual question elicits different answers from different people. The prevailing assumption is that the goal is to *return a single answer*. Yet prior work [18] has shown that it is common for visual questions to elicit different answers—it occurs for over half of nearly 500,000 visual questions in three VQA datasets.

Our goal is to identify why different answers arise. Accordingly, our work is premised on the assumption that there can *exist multiple answers* for a visual question. Our work extends prior work [5, 18, 31, 39] which has suggested

reasons why answers can differ—including that visual questions are difficult, subjective, or ambiguous as well as synonymous answers. First, we propose a taxonomy of nine plausible reasons why answers can differ, which are exemplified in Figure 1. We next ask crowd workers to identify which of these reasons led to answer differences for each of ~45,000 visual questions asked by people who are blind and sighted. Finally, we propose a novel problem of predicting which reasons will lead to answer differences directly from a visual question and propose a novel algorithm for this task. Our findings offer compelling evidence that an algorithmic framework must learn to simultaneously model and synthesize different individuals' (potentially conflicting) perceptions of images and language.

We offer our work as a valuable foundation for improving VQA services, by empowering system designers and users to know how to prevent, interpret, or resolve answer differences. Specifically, a solution that anticipates why a visual question will lead to different answers (summarized in Figure 1) could (1) help users identify how to modify their visual question in order to arrive at a single, unambiguous answer; e.g., retake an image when it is low quality or does not show the answer versus modify the question when it is ambiguous or invalid; (2) increase users' awareness for what reasons, if any, trigger answer differences when they are given a single answer; or (3) reveal how to automatically aggregate different answers [2, 19, 24, 26, 43] when multiple answers are collected.

More generally, to our knowledge, this is the first work in the computer vision community to characterize, quantify, and model reasons why annotations differ. We believe it will motivate and facilitate future work on related problems that similarly face annotator differences, including image captioning, visual storytelling, and visual dialog. We publicly share the datasets and code to encourage community progress in developing algorithmic frameworks that can account for the diversity of perspectives in a crowd.

2. Related Work

Visual Question Answering Datasets. Numerous dataset challenges have been posed to encourage the development of algorithms that automatically answer visual ques-

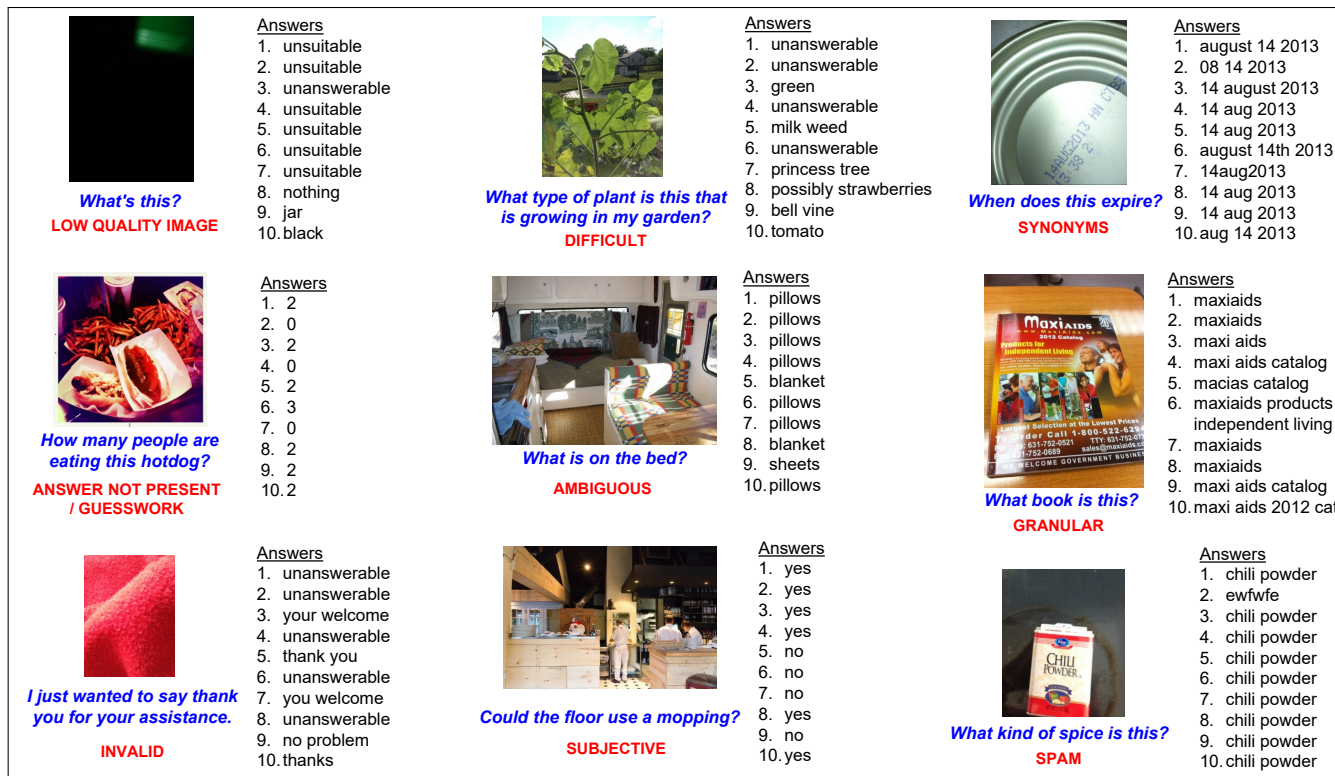|	Answers
What's this? LOW QUALITY IMAGE	1. unsuitable 2. unsuitable 3. unanswerable 4. unsuitable 5. unsuitable 6. unsuitable 7. unsuitable 8. nothing 9. jar 10. black
What type of plant is this that is growing in my garden? DIFFICULT	1. unanswerable 2. unanswerable 3. green 4. unanswerable 5. milk weed 6. unanswerable 7. princess tree 8. possibly strawberries 9. bell vine 10. tomato
When does this expire? SYNONYMS	1. august 14 2013 2. 08 14 2013 3. 14 august 2013 4. 14 aug 2013 5. 14 aug 2013 6. august 14th 2013 7. 14aug2013 8. 14 aug 2013 9. 14 aug 2013 10. aug 14 2013
How many people are eating this hotdog? ANSWER NOT PRESENT / GUESSWORK	1. 2 2. 0 3. 2 4. 0 5. 2 6. 3 7. 0 8. 2 9. 2 10. 2
What is on the bed? AMBIGUOUS	1. pillows 2. pillows 3. pillows 4. pillows 5. blanket 6. pillows 7. pillows 8. blanket 9. sheets 10. pillows
What book is this? GRANULAR	1. maxiaids 2. maxiaids 3. maxi aids 4. maxi aids catalog 5. macias catalog 6. maxiaids products independent living 7. maxiaids 8. maxiaids 9. maxi aids catalog 10. maxi aids 2012 cat
I just wanted to say thank you for your assistance. INVALID	1. unanswerable 2. unanswerable 3. your welcome 4. unanswerable 5. thank you 6. unanswerable 7. you welcome 8. unanswerable 9. no problem 10. thanks
Could the floor use a mopping? SUBJECTIVE	1. yes 2. yes 3. yes 4. yes 5. no 6. no 7. no 8. yes 9. no 10. yes
What kind of spice is this? SPAM	1. chili powder 2. ewfwfe 3. chili powder 4. chili powder 5. chili powder 6. chili powder 7. chili powder 8. chili powder 9. chili powder 10. chili powder

Figure 1: Examples of visual questions (VQs) asked by people who are blind and sighted, and answers from 10 different people. As shown, the answers can differ for a variety of reasons, including because of the VQ (first and second columns) or the answers (third column). We propose a novel problem of predicting why answers will differ for a VQ and a solution.

tions [5, 17, 20, 25]. The shared goal of these challenges is to return a single answer to each visual question. Yet, visual questions often lead to different answers from different people [5, 18, 31, 39]. Prior work [5, 31] tried to mitigate this problem on the performance metric side, by using consensus metrics. We instead introduce the first VQA dataset that fosters research in learning why different answers will arise. Specifically, for two popular VQA datasets, VizWiz [20] and VQA_2.0 [5], we label each visual question with metadata indicating which among nine options are the reasons for the observed answer differences. Experiments demonstrate these datasets are valuable for training algorithms to predict why answers will differ for any visual question.

Challenges/Obstacles for Answering Visual Questions. Our work relates to the body of literature aimed at understanding what can make a visual question challenging, or even impossible, to answer. One work examined the issue of *difficulty* (aka - required skill level), by identifying the minimum age needed to successfully answer a visual question [5]. Another set of works explored the issue of *relevance*, and in particular identifying when questions are unrelated to the contents of images [30]. Another work examined the issue of *answerability*, with an emphasis on when

questions cannot be answered due to extreme image quality issues including blur, saturation, and fingers obstructing the camera view [20]. Our work complements prior work in that we found that each of these issues are commonly associated with visual questions that evoke different answers; e.g., see examples in Figure 1 for "Difficult", "Answer Not Present", and "Low Quality Image" respectively. Our experiments demonstrate a strong advantage of designing algorithms to directly predict the reason why answers will differ over relying on related methods that predict *relevance* [30] or *answerability* [20] alone.

Understanding Why Crowd Responses Differ. More broadly, related work includes literature about why a crowd's annotations can differ and how to resolve those differences (largely in domains outside VQA). Commonly, it is assumed there is a single true answer and that any observed differences stem from insufficient annotator performance, for example because the task is too difficult for some annotators [3, 6, 13, 21, 22, 36, 37, 38, 44] or because malicious workers submit spam [14, 15, 42]. Others have embraced the possibility that the task may be ambiguous [2, 19, 24, 26] or subjective [43] and so multiple annotations can be valid. While each work embeds assumptions

Reason	Label	Description
		Issues with the Question-Image (QI) pair
Low Quality Image	**LQI**	image is too small, out of focus, having poor quality, or nothing is visible
Answer Not Present	**IVE**	good image, but answer to the question is not present in the image (Insufficient Visual Evidence), so some answers reflect guesses
Invalid	**INV**	a proper or semantically correct question is absent
Difficult	**DFF**	questions that require domain expertise (*e.g.*, identifying if a skin-rash is due to bug bite), special skills, or too much effort (*e.g.*, counting the number of sheep in a field full of sheep)
Ambiguous	**AMB**	good image and valid question, but taken together they have more than one valid interpretation, leading to multiple answers
Subjective	**SBJ**	opinion-driven questions, such as assessing beauty, fashion sense, emotions
		Issues with the answers
Synonyms	**SYN**	answers present the same idea, but using different words having similar meaning (*e.g.*, 'round' versus 'circular')
Granular	**GRN**	answers present the same idea, but in different levels of detail or specialization (*e.g.*, 'plane' versus 'Boeing')
Spam	**SPM**	a person inadequately answers a simple, straight-forward visual question

Table 1: Proposed taxonomy of nine reasons that evoke differing answers for visual questions.

regarding why answers differ, a challenge remains of knowing which assumptions are valid when. Thus, we conduct a systematic study to enumerate plausible reasons and propose an algorithm to decide which reasons apply when.

Learning to Anticipate Annotation Differences. Related works have trained algorithms to anticipate when annotation differences will arise for numerous vision problems. Some methods recognize when visual content is *ambiguous* and so will lead to diverse human interpretations, including for captioning images [24], interpreting visual attributes [26], and locating the most prominent foreground object [19]. Other methods recognize to what extent visual content will evoke people's *subjective* perceptions, including for visual humor [10] and memorability [23]. Other methods anticipate whether a crowd will offer different responses to a visual question [18] and to what extent [45]. Complementing prior work, we propose the first solution for deciphering which of multiple reason(s) will lead to annotation differences. A key insight for our approach is to employ a VQA system's predicted answers as predictive cues, motivated by the belief that VQA models trained to optimize for multiple correct answers (e.g., using "soft target" scores) embed some understanding for why there may be uncertainty around a single true answer [39].

3. Labeled Datasets

We now introduce how we created the datasets, which consist of VQAs paired with labels indicating why answers

differ. In particular, in this section, we describe our taxonomy and data labeling approach.

Reasons for Differing Answers. We developed a taxonomy of nine reasons for why answers may differ, which are summarized in Table 1. Six of the nine reasons are inspired by the crowdsourcing literature – INV [32], DFF [44], AMB [24, 26, 43], SBJ [33, 43, 9], SYN [32], and SPM [41, 42, 14, 15]. Two of the reasons are inspired by prior visual question answering work [20] – LQI and IVE. The final category is inspired by our inspection of a random subset of visual questions with their answers – GRN. We checked whether this taxonomy provided full coverage of plausible reasons for answer differences by conducting a pilot crowdsourcing study, and found no additional categories were identified[1]. This taxonomy is divided into reasons that arise from the visual question versus answers.

Approach for Labeling VQAs. We designed our user interface for this annotation task to show the visual question (used interchangeably with Question-Image/QI pair) with the ten answers, and to require a user to select all reasons why the answers differed. The user was required to select at least one reason. We grouped reasons based on whether the issue stemmed from the QI pair versus answers to assist

---

[1]Using the setup from the next subsection, we crowdsourced five labels per VQA for 100 VQAs randomly selected from *VizWiz* [20]. In this user interface, users were provided an "OTHER" category and an optional, open-ended feedback comment box. No additional reasons were identified either by users selecting "OTHER" or providing feedback.

the user in deciding which reason(s) to select. We also included a reason called ***Other*** (**OTH**), linked to a free-entry text-box, so users could suggest what they thought was the relevant reason when they felt no options were suitable.

We chose to label 44,955 VQAs coming from two popular VQA datasets. Each VQA consists of an image and a question paired with 10 answers crowdsourced from Amazon Mechanical Turk (AMT) workers. We include the *VizWiz* [20] dataset to address immediate, real-world needs of people who are blind. It originates from people who are blind who used mobile phones to snap photos and record questions about them [7, 8]; e.g., "what type of beverage is in this bottle?" or "has the milk expired?". We used the entire dataset, excluding VQs where all answers are identical with exact string matching (i.e., no answer differences), resulting in 29,921 VQAs. We also included a sample from the popular *VQA_2.0* [17] dataset for comparison. It differs from *VizWiz* in part because the images and questions were created separately. The images originate from the MS-COCO dataset [28] and the questions came from crowd workers instructed to ask a question about the image that can 'stump' a 'smart robot' [5]. We used a subset of 15,034 randomly selected QI pairs from the training set, for which the ten crowdsourced answers were not identical using exact string matching. Together, VizWiz [20] and VQA_2.0 [17] represent a diversity of users and use case scenarios.

We employed crowd workers from Amazon Mechanical Turk to label the VQAs. For quality control, we restricted workers to those who previously completed over 500 jobs with at least a 95% approval rating, and are from the US to try to ensure English proficiency. We also collected five labels per visual question from five workers.

4. Understanding Why Answers Differ

We analyzed the 224,775 crowdsourced annotations to answer the following: (1) what reasons (un)commonly evoke differing answers?, (2) how many unique reasons typically lead to differing answers for a visual question?, and (3) which reasons commonly occur together versus alone?

4.1. (Un)Common Reasons for Answer Differences

QI Pair Versus Answers. We first quantified the tendency for answer differences to arise because of issues with the QI pair versus answers. To do so, we quantified the proportion of VQAs having issues only with the QI pair (i.e., due to LQI, IVE, INV, DFF, AMB, or SBJ), only with the answers (i.e., due to SYN, GRN, or SPM), and with both. To enrich our analysis, we examined the influence of different levels of trust in the crowd by tallying the valid reasons observed when requiring a minimum of 1, 2, or 3 members of the crowd to offer the same reason for the reason to be valid. Results for each dataset are shown in Figure 2.

We found that answer differences arise most often because of issues with both the QI pair and the answers. For example, 75% and 80% of visual questions from *VizWiz* and *VQA_2.0* datasets, respectively, have issues arising from both sources (Figure 2; 2 person threshold). In contrast, a very small fraction of answer differences arise because of the answers alone. A larger percentage arise because of issues with the QI pair alone, affecting roughly 20% of VQAs for both datasets (Figure 2; 2 person threshold). This latter finding highlights that a considerable portion of answer differences could be avoided by modifying the visual question.

Frequency of Each Reason. We next examined the tendency for each of the nine reasons to lead to different answers. To do so, we calculated the percentage of VQAs assigned with each reason for different levels of trust in the crowd (i.e., requiring a minimum of 1, 2, or 3 members of the crowd to offer the same reason for the reason to be valid). Results are shown in Figure 3.

The most common reasons match for both datasets: ambiguous QIs (AMB), followed by synonymous answers (SYN), and finally varying levels of answer granularity
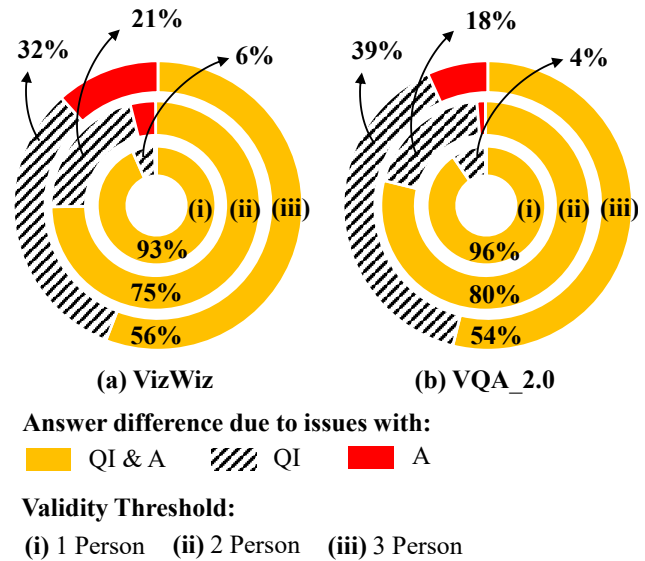

(a) VizWiz   **(b) VQA_2.0**

Answer difference due to issues with:
▮ QI & A   ▨ QI   ▮ A

Validity Threshold:
(i) 1 Person   **(ii)** 2 Person   **(iii)** 3 Person

Figure 2: Percentage of VQAs where answer differences arise due to issues with both the QI pair and the 10 answers (**QI & A**, yellow), issues with the QI pair only (**QI**, striped), or issues with the 10 answers only (**A**, red), for the **(a)** *VizWiz* and **(b)** *VQA_2.0* datasets. Results are shown with respect to different levels of trust in the crowd workers: **(i)** *Trust All*: only one worker has to select the reason (1 person validity threshold); **(ii)** *Trust Any Pair*: at least two workers must agree on the reason (2 person validity threshold); and **(iii)** *Trust Majority*: at least three workers must agree on the reason (3 person validity threshold).
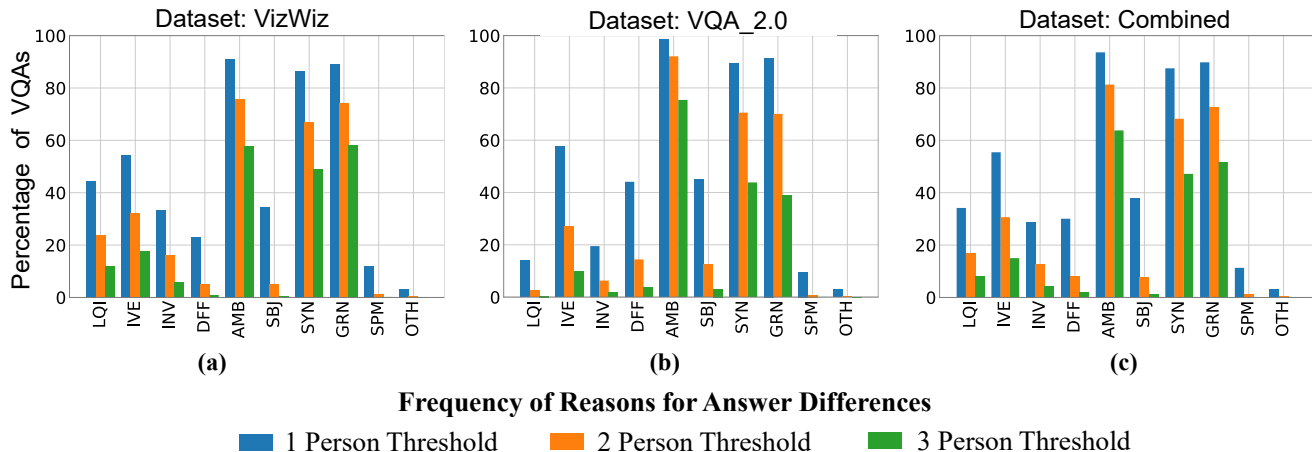
Figure 3: Histograms showing the frequency of each reason leading to answer differences for (a) 29,921 visual questions asked by blind people, (b) 15,034 visual questions asked by sighted people, and (c) combination of the two. The plots are computed based on increasing thresholds of inter-worker agreement required to make a reason valid ranging from requiring at least one worker selecting it up to at least three workers. The most popular reasons are ambiguous visual questions (AMB), synonymous answers (SYN), and varying answer granularity (GRN) whereas the most rare are spam (SPM) and other (OTH).

(GRN). Ambiguity (AMB) accounts for 81.3% of answer disagreements across both datasets (Figure 3c; 2 person threshold). Ambiguous examples in the *VizWiz* dataset often arise because the question "What (object) is this . . . ?" is asked about an image showing multiple objects (e.g., 'store', 'shopping area', 'shopping cart'). In the *VQA_2.0* dataset, we found ambiguity can arise for lengthy questions which may leave individuals confused about how to interpret the question (*e.g.* 'What weather related event can be seen under the clouds in the horizon?') as well as for visual questions seemingly designed to be ambiguous in order to "stump a robot" [5] (*e.g.* 'Q: Where are the baby elephants? Ans 1: right, Ans 2: on the grass, Ans 3: next to mom and dad, etc.). The closely-following second and third most common reasons are answer granularity (GRN) and synonyms (SYN) which account for 72.9% and 68.3% of VQAs across both datasets (Figure 3c; 2 person threshold). These findings highlight that most answer differences can be resolved by disambiguating visual questions or resolving synonyms and differing granularity [16, 27, 43].

The least common reason for answer differences is **spam (SPM)**, with it accounting for approximately 1.1% of VQAs across both datasets (Figure 3c; 2 person threshold). This is interesting because the issue of spam has received a lot of attention in the crowdsourcing literature (e.g., [14, 15, 42] to name a few). Our findings suggest that improving spam detection solutions will lead to considerably less impact than improving approaches addressing the other reasons.

Despite that the approaches for curating the VQAs in the *VizWiz* and *VQA_2.0* datasets are very different—with *VizWiz* arising from daily visual challenges of blind users,

and *VQA_2.0* containing visual questions designed to be hard for machines to answer—we found overall that the ranking and prevalence of reasons for answer difference across the two datasets is very similar. The key differences lie in difficult VQs (DFF) and low quality images (LQI). For example, the percentage of difficult visual questions is four times more in *VQA_2.0* than *VizWiz*; i.e., ~3% versus ~12% (2-person threshold). Additionally, the percentage of low quality images in *VizWiz* accounts for approximately nine times more than that observed for *VQA_2.0*; i.e., 23.8% versus 2.6% of all visual questions respectively (2-person threshold). Despite such dataset differences, we will show in Section 5 that prediction models can still learn to predict which reason(s) will lead to answer differences.

Examples of VQAs that were most confidently voted by the crowd workers as belonging to each of the nine reasons for answer differences are shown in Figure 1.

4.2. Co-Occurring Reasons for Answer Differences

We now examine how many reasons typically lead to answer differences for a given VQA as well as to what extent reasons co-occur. For the following analysis, we assume a reason occurs for a VQA if at least two crowd workers flag that reason as occurring.

Number of Unique Reasons. We first tallied the number of unique reasons leading to answer differences for each VQA. Across both datasets, there are most commonly three unique reasons; i.e., for more than 55% of the *VizWiz* and *VQA_2.0* VQAs. Two and four reasons also are common, accounting for 15% and 16% VQAs respectively across both datasets. The remaining ~20% arise from one unique

Reason	Co-occurs with (%)										Reason Clarity (%)
	LQI	IVE	INV	DFF	AMB	SBJ	SYN	GRN	SPM	OTH	
LQI	0	44	17	3	0	0	0	0	0	0	56
IVE	33	0	44	9	0	0	0	0	0	0	56
INV	25	89	0	9	0	1	0	0	0	0	11
DFF	13	59	28	0	0	0	0	0	0	0	41
AMB	0	0	0	0	0	3	76	83	0	0	17
SBJ	0	0	3	0	49	0	0	0	0	0	51
SYN	0	0	0	0	86	0	0	92	0	0	8
GRN	0	0	0	0	85	0	83	0	0	0	15
SPM	8	0	0	0	0	0	0	0	0	2	92
OTH	0	0	0	0	0	0	0	0	7	0	93

(a)

Reason	Co-occurs with (%)										Reason Clarity (%)
	LQI	IVE	INV	DFF	AMB	SBJ	SYN	GRN	SPM	OTH	
LQI	0	28	8	3	0	0	0	0	1	0	72
IVE	3	0	20	10	0	5	0	0	1	0	80
INV	3	87	0	0	0	5	0	0	0	0	13
DFF	0	19	0	0	0	0	0	0	0	0	81
AMB	0	0	0	0	0	0	65	62	0	0	35
SBJ	0	12	2	0	0	0	0	0	0	0	88
SYN	0	0	0	0	85	0	0	76	0	0	15
GRN	0	0	0	0	82	0	76	0	0	0	18
SPM	4	19	1	8	0	0	0	0	0	1	81
OTH	4	0	0	0	0	0	0	0	3	0	96

(b)

Figure 4: Tendency for each reason to co-occur with every other reason as well as to occur alone (i.e., reason clarity) for the (a) *VizWiz* and (b) *VQA_2.0* datasets (2-person threshold). Across both datasets, 'spam' and 'other' most often occur alone.

reason, followed by five and six unique reasons respectively, across both datasets. These findings motivate representing the problem of predicting which reasons lead to answer differences as a multi-label classification problem.

Reasons Occurring Together. We next examined the extent to which the various reasons co-occur. To do so, we computed an adaptation of causal power [11, 29], measuring the co-occurrence of two reasons d_i and d_j as:

$$co\_occurrence\ (d_i, d_j) = \frac{P(d_j|d_i) - P(d_j|\bar{d}_i)}{1 - P(d_j|\bar{d}_i)} \quad (1)$$

where $P(d)$ is the probability that reason d is present for a VQA and $P(\bar{d})$ is the probability that it is not present. Intuitively, this metric indicates how often d_j arises when d_i occurs for a VQA. Results are shown in Figure 4.

Across both *VizWiz* and *VQA_2.0*, we observe the reasons with the highest co-occurrences ($\geq 80\%$) are answer synonyms (SYN), answer granularity (GRN), invalid questions (INV), and ambiguity (AMB). For example, in *VizWiz*, for all the VQAs where SYN was chosen, GRN co-occurs for 92% of those VQAs, followed by AMB for 86%. Likewise, in all the question where GRN occurs, AMB occurs in 85% of them, and SYN occurs in 83% of them. Our results offer strong evidence that ambiguity in the QI pair can cause people to be uncertain both about what level of detail to provide and what word to use among valid synonyms; e.g., 'money', 'currency', '10 dollar bill'. We hypothesize that one promising way a VQA system could greatly reduce the frequency of answer differences is to instruct individuals asking QIs to clarify the level of detail they are seeking whenever answer differences are expected to arise from QI ambiguity or differing answer granularity.

We also observe that across both datasets, invalid question (INV) is paired with insufficient visual evidence to answer the question (IVE). This is true for 89% of VQAs in *VizWiz* and 87% of VQAs in *VQA_2.0*. This suggests that if the answer to the question is not present in the image, people think that the question is invalid.

Reasons Occurring Alone. We next measured how often a reason occurs on its own. To do so, we compute *clarity* of a reason d as follows: the percentage of all VQAs where d is chosen for which no other reason is chosen. Results are shown in Figure 4. Across both datasets, spam (SPM) commonly occurs alone; i.e., for at least 92% and 81% of the VQAs in *VizWiz* and *VQA_2.0*, in which this reason is valid. For *VizWiz*, as reflected by small percentage values for clarity, most of the reasons commonly co-occur with at least one other label. In contrast, *VQA_2.0* only has four of the nine labels (INV, SYN, GRN, AMB) with low clarity ($\leq 35\%$), while the rest have high clarity ($\geq 72\%$). This reveals that co-occurrence in *VQA_2.0* centers around fewer reasons than for *VizWiz*.

5. Predicting Why Answers Will Differ

We now introduce a novel machine learning task of predicting why a visual question will lead to different answers.

5.1. Prediction Model

We pose the task as a multi-label classification problem.

Ground Truth. We compute binary ground truth for each of the nine reasons described in Table 1, as well as the "Other" category. Consequently, our ground truth consists of 10 labels. For each label, we consider it as present (i.e., '1') for a visual question only if at least two of the five crowd workers selected that label as present.

Proposed Model. The motivation for the design of our model is that it employs as predictive cues the input image (I), input question (Q), *and* answers to the QI pair (A). This is because we know that answer differences can arise from the QI pair alone (i.e., first six reasons in Table 1) as well as from the answers alone (i.e., last three reasons in Table 1).
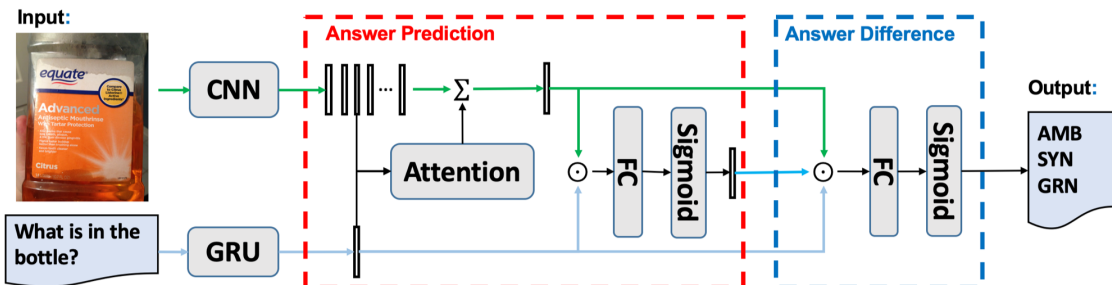
Figure 5: Summary of the proposed model (Q+I+A) for predicting which among 10 reasons will lead to answer differences.

Our key challenge is how to represent the answers, since the ground truth answers are not known in practice.

Our model is summarized in Figure 5. It takes as input the *image*, encoded as the last convolutional layer of the Faster R-CNN model [35], and the *question*, encoded as a 300 dimension word from the pre-trained GloVe vectors [34] that is then passed to a single-layer GRU with 1024 hidden units. Our key design decision was to pass this input to an "Answer Prediction" module, with the goal that it outputs all answers. We employed a VQA algorithm [4, 40] that is trained to optimize for multiple correct answers and output a "soft target" vector representation of its confidence on all answer candidates; e.g., for the second example in second row of Figure 1, the ideal prediction for the 10 answers of 7 "pillows", 2 "blanket", and 1 "sheets" would be 0.7 for "pillows", 0.2 for "blanket", 0.1 for "sheets", and 0 otherwise.[2] By including the "Answer Prediction" module, the subsequent "Answer Difference" module can then make predictions based not only on the image and question features but also using the (predicted) answers. The "Answer Difference" module passes its input (Q, I, & A) through a fully-connected layer of 1024 units to a sigmoid function in order to predict the probability of each of the 10 labels.

For training, we adopt the binary cross entropy loss as:

$$L = \sum_{i=1}^{N} y_i \log(p_i) + (1 - y_i) \log(1 - p_i),$$

where N is the total number of labels, y_i is the ground truth label, and p_i is the predicted probability from the sigmoid function. We initialize the model using the pre-trained weights of [34] for the "GRU", pre-trained weights of [4] for the "CNN" and "Answer Prediction" modules, and random values for the "Answer Difference" module. We fine-tune the whole network using the Adam solver with a fixed learning rate of 0.001 and batch size of 128. Dropout and early stopping (five epochs) are used to reduce overfitting.

5.2. Evaluation

Dataset Train/Validation/Test Split. We used the whole *VizWiz* dataset, including the QI pairs where all answers were identical (i.e., 3% of the VQAs from the original dataset) so that trained algorithms can work well in the presence of QI pairs that do not lead to answer differences. Using similar train/validation/test splits from [20], we have 19,969 training (64%), 3166 validation (10%), and 7983 test (26%) samples. For the 15,034 VQs from the *VQA_2.0* dataset, we introduced a 65/10/25 split which resulted in 9772 training, 1504 validation, and 3758 test examples.

Evaluation Metrics. We report the average precision for each label and the mean average precision across all labels.

Baselines. To our knowledge, no prior work has tried to predict the reason(s) why a visual question will have different answers. Therefore, we evaluate the benefit of three related baseline methods to reveal the value of re-purposing existing approaches for our new problem.[3] We include random guessing (i.e., **Random**), since this is the best a user can achieve today. We also include an algorithm for the related task of predicting whether a question is relevant to the given image, which we call **QI-relevance**. We adapt the pretrained question-image-relevance system [30] to predict among the plausible reasons: if the QI pair is predicted as relevant, the "LQI", "IVE", and "AMB" are predicted as 0, and 1 otherwise. Other labels are predicted randomly since they are unrelated to QI-relevance. We also include an algorithm that predicts whether a VQ is **unanswerable** for the *VizWiz* dataset ("Q+I" in [20]). Similar to the QI-relevance baseline, if the VQ is predicted as answerable, the "LQI", "IVE", and "AMB" are predicted as 0, and 1 otherwise.

We also evaluate five variants of our **Q+I+A** model. We examine the power of different predictive cues by predicting only from the question-image pair (i.e., **Q+I**), question (i.e., **Q**), and image (i.e., **I**). We examine fine-tuning the VQA model [17, 20] on our datasets after replacing the last layer

---

[2]To avoid learning on the same data observed at test time, we ensure the "Answer Prediction" module is pretrained on a different VQA dataset.

[3]To avoid overlap between the training data and our test data, we retrained all baselines with the test set samples excluded (when needed).

Table 2: Average precision for predicting why answers to visual questions will differ for the VQA_2.0 and VizWiz datasets.

	Model	Overall	LQI	IVE	INV	DFF	AMB	SBJ	SYN	GRN	SPM	OTH
VQA_2.0	Random	30.24	3.71	22.43	15.09	14.62	95.19	14.18	64.99	69.42	0.52	**2.25**
	QI-Relevance [30]	32.23	4.01	43.16	15.09	14.62	94.11	14.18	64.99	69.42	0.52	2.25
	I	31.88	4.31	29.46	9.28	17.02	92.91	17.99	74.55	72.56	0.5	0.24
	Q	43.47	7.65	58.89	44.56	28.15	96.42	24.04	88.63	84.67	**1.36**	0.38
	Q+I	43.16	9.05	58.03	41.95	28.22	96.25	24.29	88.26	84.02	1.27	0.26
	Q+I+A	**44.55**	**11.58**	59.95	**46.03**	30.27	96.47	24.88	89.69	85.62	0.8	0.26
	Q+I+A_FT	44.46	8.11	**60.67**	43.36	**31.35**	**96.98**	**25.31**	**90.49**	**86.89**	1.03	0.48
	Q+I+A_GT	44.09	8.94	59.64	45.21	30.04	96.60	23.83	89.75	85.80	0.82	0.30
VizWiz	Random	30.15	23.59	33.69	18.15	5.70	74.70	5.14	66.61	71.94	1.35	0.62
	QI-Relevance [30]	31.71	30.56	40.52	18.15	5.7	76.53	5.14	66.61	71.94	1.35	0.62
	Unanswerable [20]	35.31	44.82	58.63	18.15	5.7	80.14	5.14	66.61	71.94	1.35	0.62
	I	40.54	55.42	50.66	30.12	8.77	83.39	8.64	79.76	86.29	1.71	0.61
	Q	40.5	35.87	54.66	39.24	12.32	84.41	11.00	79.46	85.10	**2.15**	0.76
	Q+I	45.73	57.81	62.47	43.24	13.77	87.81	11.14	86.36	92.01	2.00	0.75
	Q+I+A	**50.02**	**65.58**	**77.42**	56.54	**10.49**	**89.70**	11.26	90.42	95.44	1.98	**1.31**
	Q+I+A_FT	50.01	64.93	77.40	**56.78**	10.10	89.48	**13.16**	**90.52**	**95.50**	1.84	1.28
	Q+I+A_GT	50.68	66.25	77.71	57.20	13.55	90.01	12.46	90.53	95.51	1.96	1.62

of the pre-trained model to a fully-connected layer representing the answer difference reasons (i.e., **Q+I+A_FT**). Finally, we examine what happens when we use the ground truth (GT) instead of the "Answer Prediction" module ; i.e., directly use a "soft target" representation of the GT answers rather than predicted answers (i.e., **Q+I+A_GT**). Following [40], only answers that appear more than eight times in the training dataset are considered in the answer representation.

Results. Results for both datasets are shown in Table 2.

As observed, the proposed model (i.e., **Q+I+A**) outperforms the existing baselines by a large margin overall. For example, across both datasets, the performance gain is more than 12 percentage points compared to the next best baseline of **QI-relevance**. Compared to the status quo of **Random** guessing, the gains are even greater. The results reveal that re-purposing existing algorithms for our new problem is inadequate, which motivates the need for new algorithmic frameworks that directly learn the "answer difference" task.

Our findings also reveal the benefit of different predictive cues. Compared with **Q+I**, the **Q+I+A** improves the performance by 1.4% and 4.3% on *VQA_2.0* and *VizWiz* datasets respectively. This verifies the effectiveness of adding the predicted answers as one more signal for prediction. Interestingly, we observe **I** has a much larger impact for the *VizWiz* dataset than for *VQA_2.0*. We attribute this difference largely to its significant advantage for the LQI and IVE categories for *VizWiz*, providing over a 20% and 30% gain respectively over **Random** guessing.

When comparing the performance of **Q+I+A** with **Q+I+A_FT** and **Q+I+A_GT**, we observe similar perfor-

mance. For **Q+I+A_FT**, this could be because the predicted answers have already captured the knowledge encoded in the pre-trained VQA model. For **Q+I+A_GT**, this might arise because the answer probabilities from VQA models offer richer information than human-annotated answers and because uncommon answers are excluded.

Overall, the models perform worst across both datasets for VQs that are subjective ($< 26\%$) and difficult ($< 32\%$). This highlights a need for models to learn abstract concepts such as common sense. Other challenges for the model, despite considerable improvement of our models over related baselines (i.e., by typically over 20%), are recognizing invalid questions and low quality images.

6. Conclusions

We proposed a taxonomy of nine reasons why answers to visual questions can differ and a novel problem of predicting why answers will differ. Our experiments demonstrate the promise of algorithms that predict directly from a visual question for this novel task. The datasets and code are publicly shared at https://vizwiz.org to facilitate future extensions of this work. Valuable future work includes user studies to identify how to employ such algorithms to guide users in modifying their visual questions so they arrive at a single answer or to retroactively aggregate answers [37].

References

[1] BeSpecular. https://www.bespecular.com. 1

[2] Ehsan Amid and Antti Ukkonen. Multiview triplet embedding: Learning attributes in multiple maps. In *International Conference on Machine Learning*, pages 1472–1480, 2015. 1, 2

[3] Hossein Amirkhani and Mohammad Rahmati. Agreement/disagreement based crowd labeling. *Applied intelligence*, 41(1):212–222, 2014. 2

[4] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018. 7

[5] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. VQA: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2425–2433, 2015. 1, 2, 4, 5

[6] Lora Aroyo and Chris Welty. Crowd Truth: Harnessing disagreement in crowdsourcing a relation extraction gold standard. *WebSci2013. ACM*, 2013, 2013. 2

[7] Jeffrey P Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samual White, and others. VizWiz: Nearly real-time answers to visual questions. In *Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology*, pages 333–342. ACM, 2010. 1, 4

[8] Erin Brady, Meredith Ringel Morris, Yu Zhong, Samuel White, and Jeffrey P. Bigham. Visual challenges in the everyday lives of blind people. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 2117–2126, New York, NY, USA, 2013. ACM. 4

[9] Michele A Burton, Erin Brady, Robin Brewer, Callie Neylan, Jeffrey P Bigham, and Amy Hurst. Crowdsourcing subjective fashion advice using vizwiz: challenges and opportunities. In *Proceedings of the 14th international ACM SIGACCESS conference on Computers and accessibility*, pages 135–142. ACM, 2012. 3

[10] Arjun Chandrasekaran, Ashwin K. Vijayakumar, Stanislaw Antol, Mohit Bansal, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. We are humor beings: Understanding and predicting visual humor. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4603–4612, 2016. 3

[11] Patricia W Cheng. From covariation to causation: a causal power theory. *Psychological review*, 104(2):367, 1997. 6

[12] Jacob Cohen. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46, 1960. 11, 12

[13] Anca Dumitrache, Lora Aroyo, and Chris Welty. Crowdsourcing ground truth for medical relation extraction. *arXiv preprint arXiv:1701.02185*, 2017. 2

[14] Carsten Eickhoff and Arjen P de Vries. Increasing cheat robustness of crowdsourcing tasks. *Information retrieval*, 16(2):121–137, 2013. 2, 3, 5

[15] Ujwal Gadiraju, Ricardo Kawase, Stefan Dietze, and Gianluca Demartini. Understanding malicious behavior in crowdsourcing platforms: The case of online surveys. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 1631–1640. ACM, 2015. 2, 3, 5

[16] Ujwal Gadiraju, Jie Yang, and Alessandro Bozzon. Clarity is a worthwhile quality: On the role of task clarity in microtask crowdsourcing. In *Proceedings of the 28th ACM Conference on Hypertext and Social Media*, pages 5–14. ACM, 2017. 5

[17] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *CVPR*, volume 1, page 9, 2017. 2, 4, 7

[18] Danna Gurari and Kristen Grauman. CrowdVerge: Predicting If People Will Agree on the Answer to a Visual Question. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 3511–3522. ACM, 2017. 1, 2, 3

[19] Danna Gurari, Kun He, Bo Xiong, Jianming Zhang, Mehrnoosh Sameki, Suyog Dutt Jain, Stan Sclaroff, Margrit Betke, and Kristen Grauman. Predicting Foreground Object Ambiguity and Efficiently Crowdsourcing the Segmentation (s). *International Journal of Computer Vision*, 126(7):714–730, 2018. 1, 2, 3

[20] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3608–3617, 2018. 2, 3, 4, 7, 8, 16

[21] Oana Inel, Lora Aroyo, Chris Welty, and Robert-Jan Sips. Domain-independent quality measures for crowd truth disagreement. *Detection, Representation, and Exploitation of Events in the Semantic Web*, page 2, 2013. 2

[22] Oana Inel, Khalid Khamkham, Tatiana Cristea, Anca Dumitrache, Arne Rutjes, Jelle van der Ploeg, Lukasz Romaszko, Lora Aroyo, and Robert-Jan Sips. Crowdtruth: Machine-human computation framework for harnessing disagreement in gathering annotated data. In *International Semantic Web Conference*, pages 486–504. Springer, 2014. 2

[23] Phillip Isola, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. What makes an image memorable? In *CVPR 2011*, pages 145–152. IEEE, 2011. 3

[24] Mainak Jas and Devi Parikh. Image specificity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2727–2736, 2015. 1, 2, 3

[25] Kushal Kafle and Christopher Kanan. An analysis of visual question answering algorithms. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1983–1991. IEEE, 2017. 1, 2

[26] Adriana Kovashka and Kristen Grauman. Discovering attribute shades of meaning with the crowd. *International Journal of Computer Vision*, 114(1):56–73, 2015. 1, 2, 3

[27] Yining Li, Chen Huang, Xiaoou Tang, and Chen Change Loy. Learning to Disambiguate by Asking Discriminative Questions. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 3439–3448, Oct. 2017. ISSN:. 5

[28] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014. 4

[29] Christian C Luhmann and Woo-kyoung Ahn. The meaning and computation of causal power: Comment on cheng (1997) and novick and cheng (2004). 2005. 6

[30] Aroma Mahendru, Viraj Prabhu, Akrit Mohapatra, Dhruv Batra, and Stefan Lee. The promise of premise: Harnessing question premises in visual question answering. *EMNLP*, 2017. 2, 7, 8, 16

[31] Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz. Ask Your Neurons: A Neural-Based Approach to Answering Questions about Images. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1–9, Dec. 2015. ISSN:. 1, 2

[32] Julian Michael, Gabriel Stanovsky, Luheng He, Ido Dagan, and Luke Zettlemoyer. Crowdsourcing question-answer meaning representations. *arXiv preprint arXiv:1711.05885*, 2017. 3

[33] An Thanh Nguyen, Matthew Halpern, Byron C Wallace, and Matthew Lease. Probabilistic modeling for crowdsourcing partially-subjective ratings. In *Fourth AAAI Conference on Human Computation and Crowdsourcing*, 2016. 3

[34] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014. 7

[35] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 7

[36] Viktoriia Sharmanska, Daniel Hernández-Lobato, Jose Miguel Hernandez-Lobato, and Novi Quadrianto. Ambiguity Helps: Classification with Disagreements in Crowdsourced Annotations. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2194–2202, June 2016. ISSN:. 2

[37] Aashish Sheshadri and Matthew Lease. Square: A benchmark for research on computing crowd consensus. In *First AAAI Conference on Human Computation and Crowdsourcing*, 2013. 2, 8

[38] Guillermo Soberón, Lora Aroyo, Chris Welty, Oana Inel, Hui Lin, and Manfred Overmeen. Measuring crowd truth: Disagreement metrics combined with worker behavior filters. In *CrowdSem 2013 Workshop*, 2013. 2

[39] Damien Teney, Peter Anderson, Xiaodong He, and Anton van den Hengel. Tips and tricks for visual question answering: Learnings from the 2017 challenge. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4223–4232, 2018. 1, 2, 3

[40] Damien Teney, Peter Anderson, Xiaodong He, and Anton van den Hengel. Tips and tricks for visual question answering: Learnings from the 2017 challenge. *CVPR*, pages 4223–4232, 2018. 7, 8

[41] Jeroen Vuurens, Arjen P de Vries, and Carsten Eickhoff. How much spam can you take? an analysis of crowdsourcing results to increase accuracy. In *Proc. ACM SIGIR Workshop on Crowdsourcing for Information Retrieval (CIR'11)*, pages 21–26, 2011. 3

[42] Jeroen BP Vuurens and Arjen P De Vries. Obtaining high-quality relevance judgments using crowdsourcing. *IEEE Internet Computing*, 16(5):20–27, 2012. 2, 3, 5

[43] Mengting Wan and Julian McAuley. Modeling Ambiguity, Subjectivity, and Diverging Viewpoints in Opinion Question Answering Systems. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pages 489–498, Dec. 2016. ISSN:. 1, 2, 3, 5

[44] Peter Welinder, Steve Branson, Pietro Perona, and Serge J Belongie. The multidimensional wisdom of crowds. In *Advances in Neural Information Processing Systems*, pages 2424–2432, 2010. 2, 3

[45] Chun-Ju Yang, Kristen Grauman, and Danna Gurari. Visual Question Answer Diversity. In *HCOMP*, pages 184–192, 2018. 3

Appendix

This document supplements our methods and results provided in the main paper.

A. Crowdsourcing Task (supplements section 3 of the main paper)

A.1. User Interface

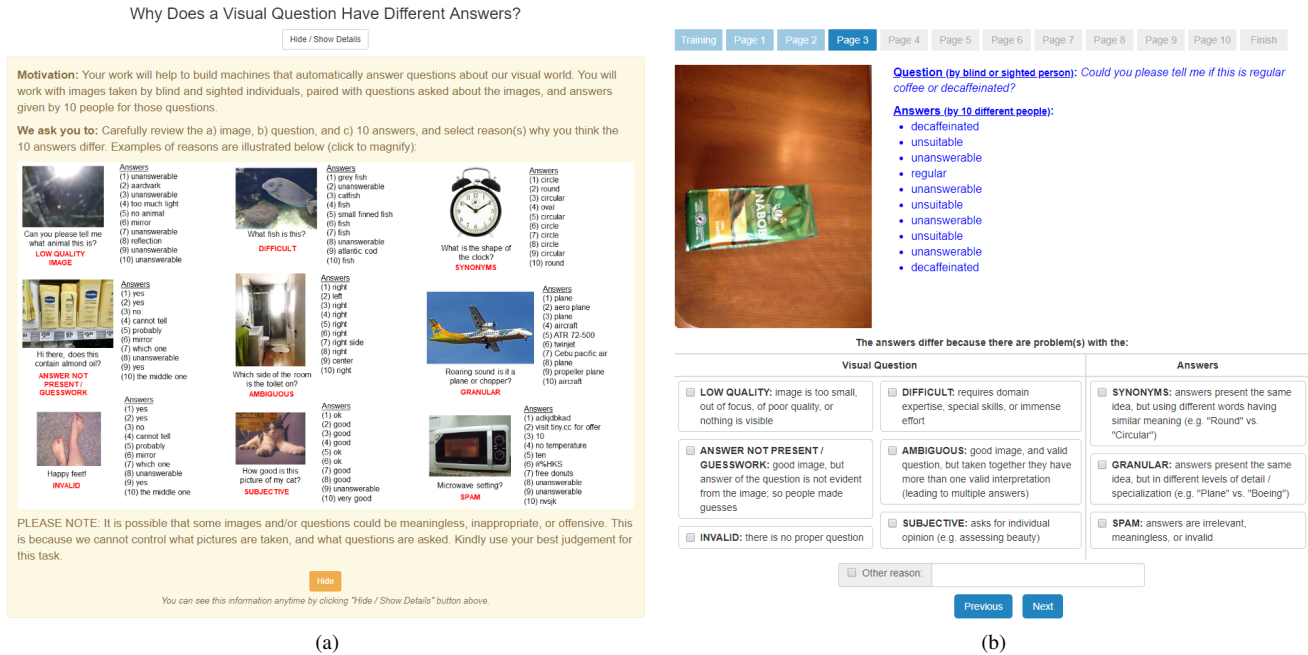The crowdsourcing user interface is shown in Figure 6.

Figure 6: (a) Task instructions to train crowd workers about the reasons that can lead to different answers. (b) The interface crowd workers used to choose why different answers are observed for a given QI pair with its 10 corresponding answers.

A.2. Quality Control

We included a training example that each crowd worker had to complete prior to completing our task. The authors identified the correct labels beforehand for this example. For each HIT posted to Amazon Mechanical Turk, the worker had to select these correct labels in order to proceed to the actual task.

B. Dataset Analysis (supplements section 4 of the main paper)

B.1. Inter-Annotator Agreement for Reasons Labels

We examine inter-annotator agreement among crowd workers. To do so, we measure the Worker-Worker Similarity (WWS) as the pairwise annotation similarity between two workers across all the VQAs they have annotated in common. The WWS measure indicates how close a worker performs to the group of workers who have solved the same task. We calculate WWS between two crowd workers w_i and w_j using three approaches: (a) number of common labels selected, (b) cosine similarity, and (c) Cohen's κ [12].

WWS - Common Labels
This metric is defined as

$$wws(w_i, w_j) = \frac{\sum\limits_{t \in T_{i,j}} numCommonLabels(w_i, w_j, t)}{\sum\limits_{t \in T_{i,j}} numAnnotations(w_i, t)}$$

where $T_{i,j}$ is the subset of all VQA tasks T annotated by both workers; $numCommonLabels(w_i, w_j, t)$ is the number of identical labels selected by both workers w_i and w_j on a VQA task t; and $numAnnotations(w_i, t)$ is the total number of labels selected by a worker w_i for a single VQA task t.
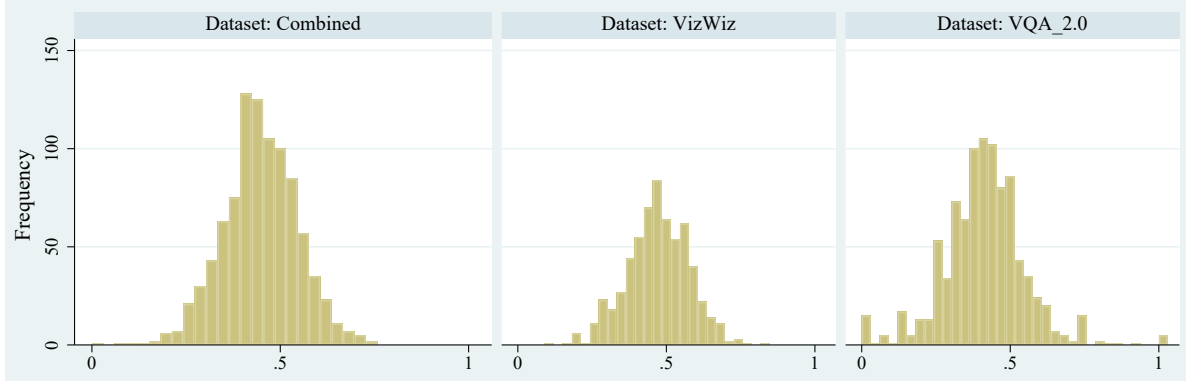

Figure 7: Distribution of 'WWS - Common Labels' for all crowd workers across both datasets alone as well as combined.

WWS - Cosine Similarity
This metric is defined as follows:

$$avg(\cos(V_{t,w_i}, V_{t,w_j})) \ \forall \ \text{worker } j, j \neq i$$

where V_{t,w_i} is the 'Task Vector' of worker w_i annotating VQA task t. A Task Vector for a worker annotating a VQA task is defined as a vector whose length is 10 (i.e. equal to the number of labels available), and whose individual elements are either 0 or 1, depending on whether the worker selected the label or not. E.g. if a worker selects the labels LQI, AMB, and SBJ, and the ordering of the labels in the Task Vector are LQI, IVE, INV, DFF, AMB, SBJ, SYN, GRN, SPM, OTH, then the Task Vector becomes: `[1,0,0,0,1,1,0,0,0,0]`.

WWS - Cohen's κ
This metric is defined as follows:

$$wws_\kappa = avg(\kappa)$$

where κ is the Cohen's kappa coefficient [12] used to measure inter-rater agreement.
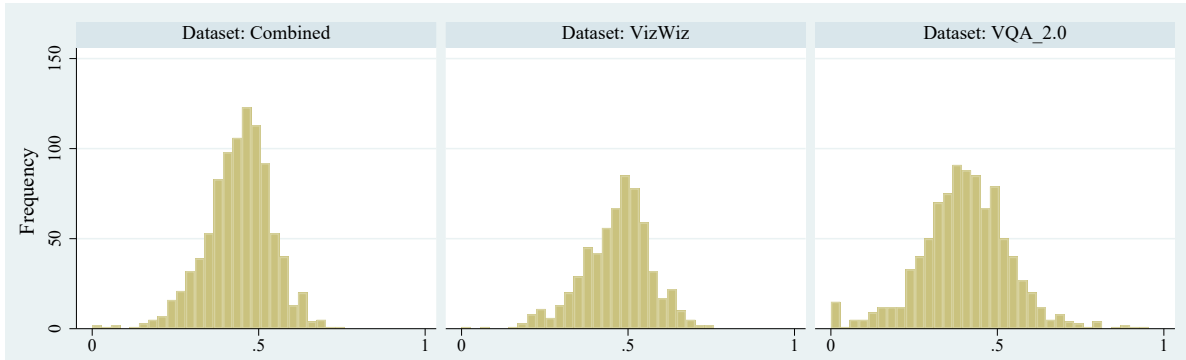

Figure 8: Distribution of 'WWS - Cosine Similarity' for all crowd workers across both datasets alone as well as combined.
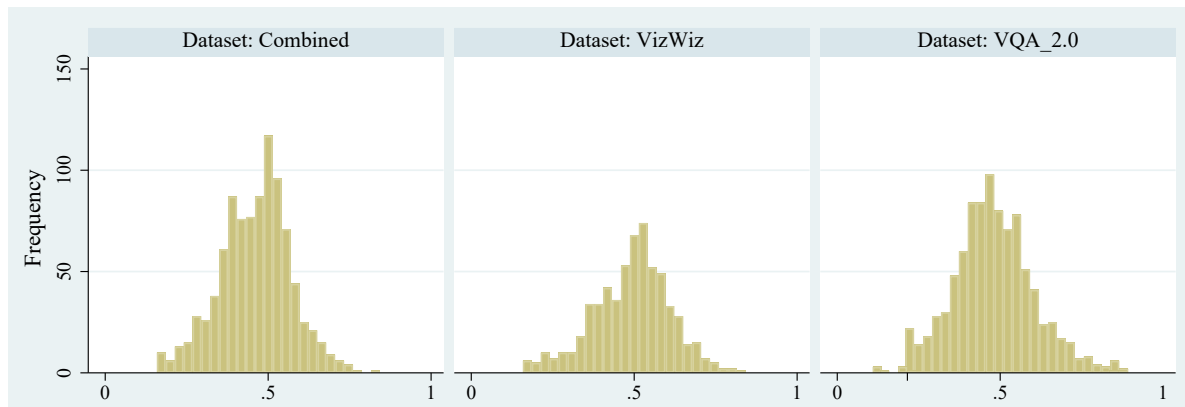
Figure 9: Distribution of 'WWS - Cohen's κ' for all crowd workers across both datasets alone as well as combined.

Figures 7, 8, and 9 show the distribution of the three WWS metrics for the 934 distinct crowdworkers who provided annotations for our dataset, averaged for each worker. Among them, 615 distinct workers annotated VQAs from the *VizWiz* dataset, while 928 distinct workers annotated the *VQA_2.0* dataset.

All the distributions assume an approximately normal form, with peaks at 0.5. This suggests that most workers agreed with 50% of the other workers with whom they shared common annotation tasks.

In the case of *VQA_2.0*, there seems to be a small yet distinct percentage of workers who did not agree with anyone. This is characterized by a small lump near the 0 value in the plots for *VQA_2.0*, of all the three WWS metrics (Figures 7, 8, & 9).

B.2. Analysis Using All Validity Thresholds
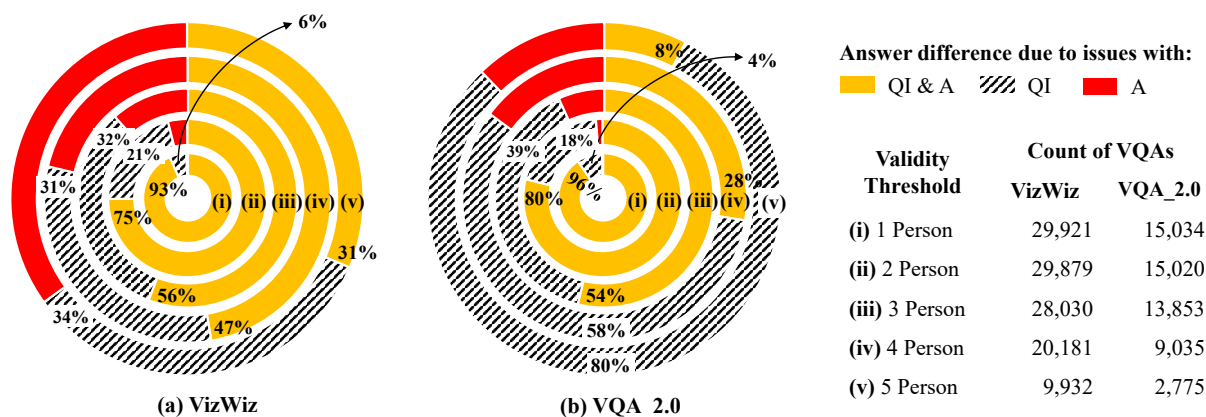

(a) VizWiz   (b) VQA_2.0

Figure 10: Relative proportion of the various sources of answer disagreement (augmented from Figure 2 in the paper).

We tallied the number of reasons leading to answer differences for each VQA, employing various levels of trust in crowd workers: from 1 person threshold to 5 person thresholds.
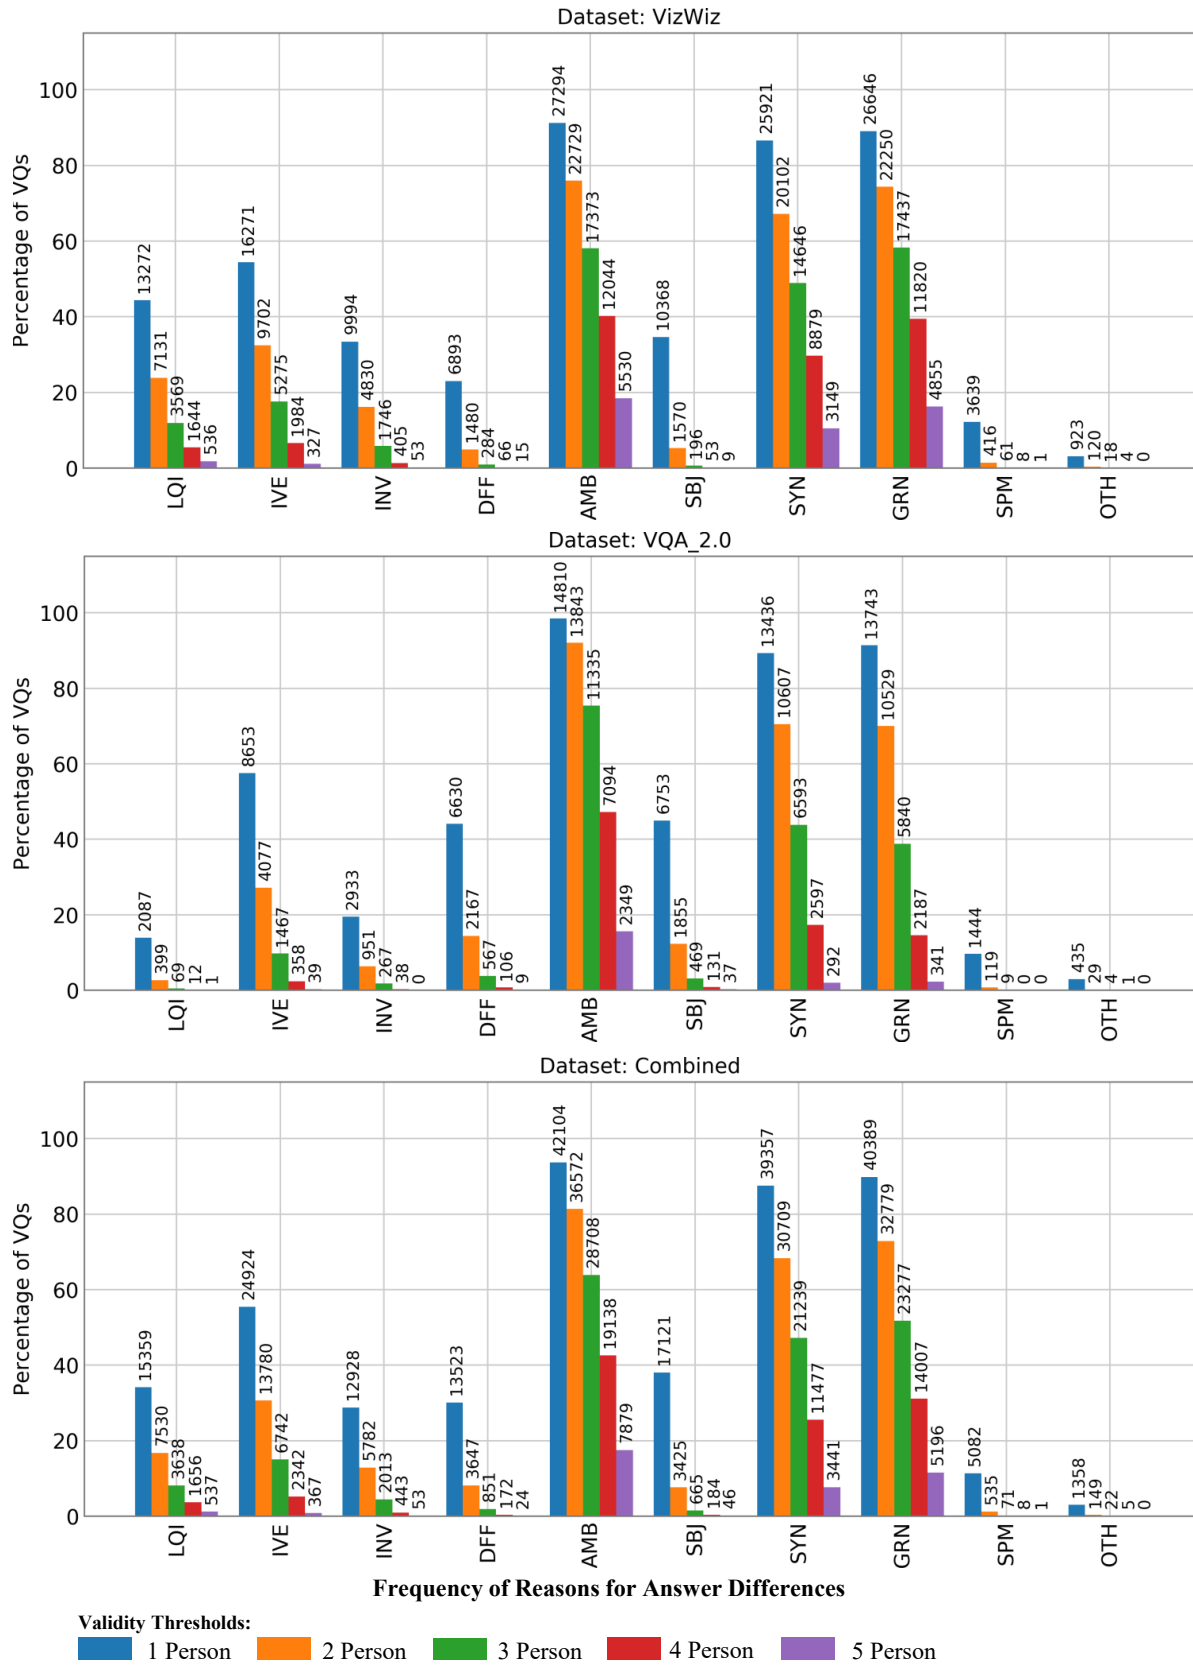
Figure 11: Histograms showing the frequency of each reason leading to answer differences (augmented from Figure 3 in the paper). Data labels show counts of VQAs matching the validity threshold.
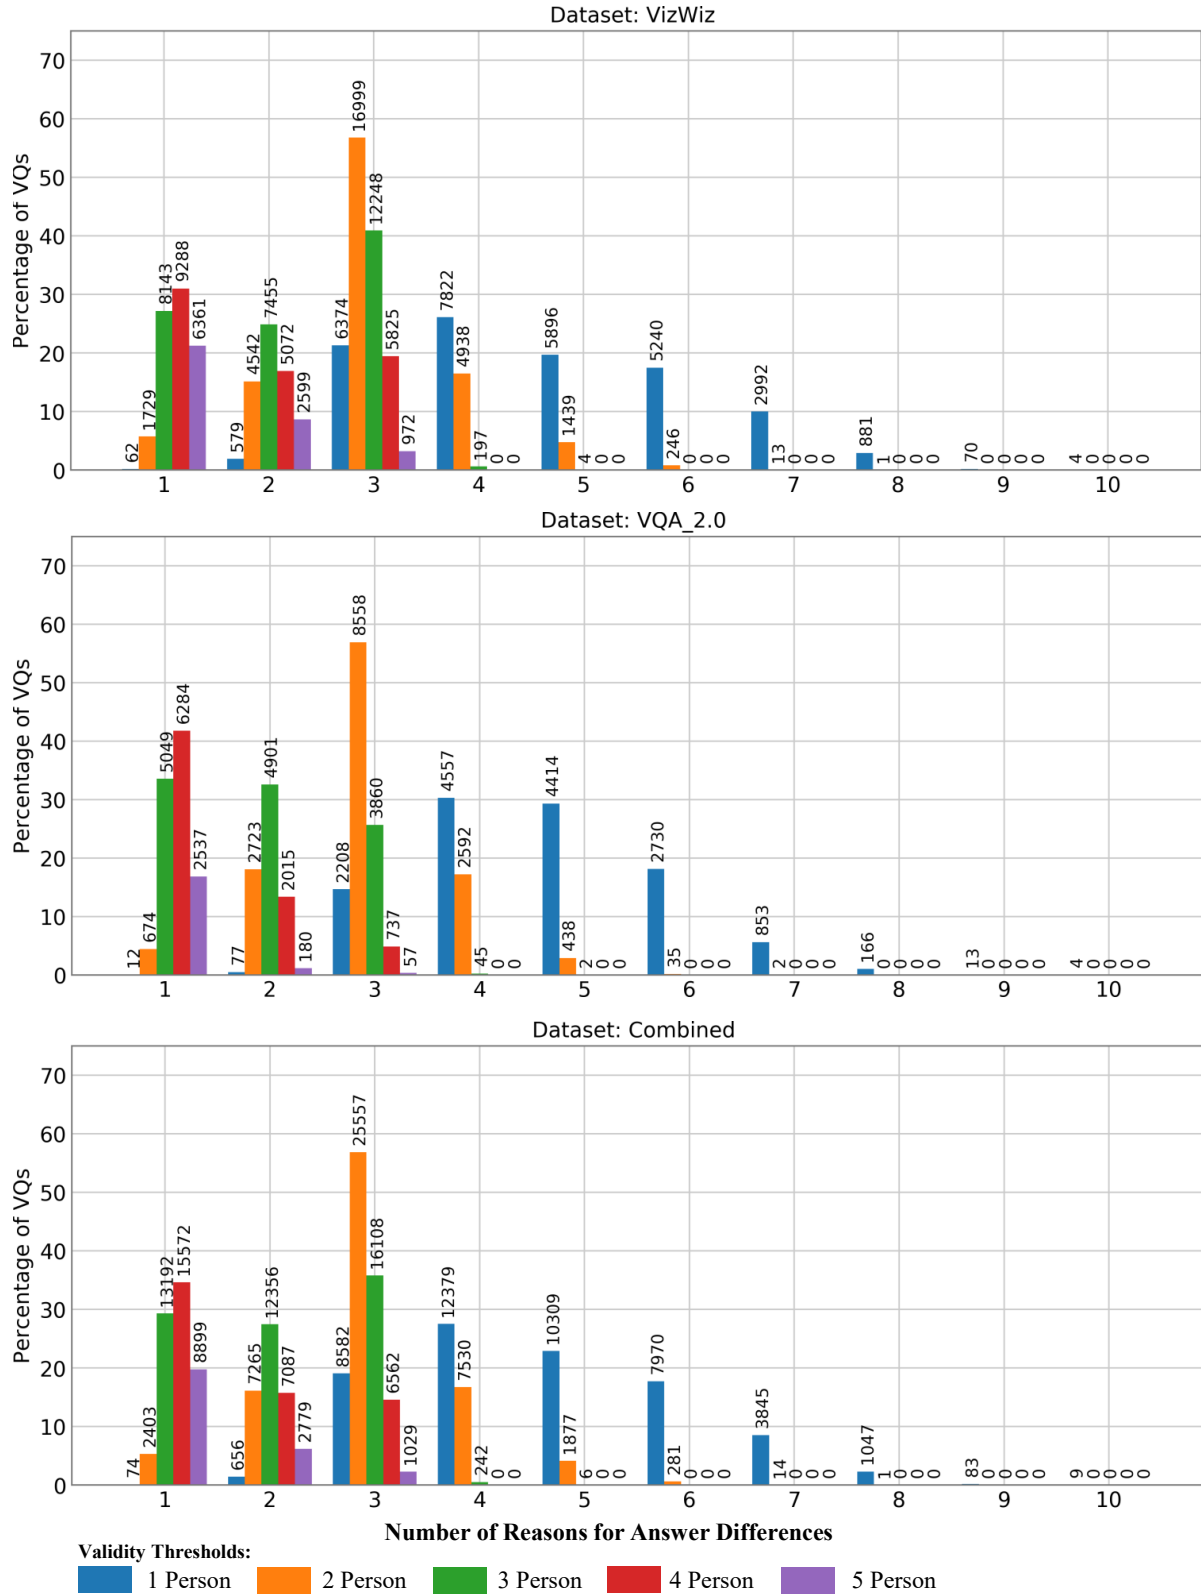
Figure 12: Histograms showing the number of unique reasons of answer differences identified for each visual question. Data labels show counts of VQAs matching the validity threshold.

Figure 10 shows the percentage of visual questions, where answer differences arise due to issues with both the QI pair and the 10 answers (**QI & A**, yellow), issues with the QI pair only (**QI**, striped), or issues with the 10 answers only (**A**, red), for the **(a)** *VizWiz* and **(b)** *VQA_2.0* datasets. Results are shown with respect to different levels of trust in the crowd workers: **(i)** *Trust All*: only one worker has to select the reason (1 person validity threshold); **(ii)** *Trust Any Pair*: at least two workers must agree on the reason (2 person validity threshold); **(iii)** *Trust Majority*: at least three workers must agree on the reason (3 person validity threshold); **(iv)**: at least four workers must have to select the reason (4 person validity threshold); and **(v)** *Trust Consensus*: all five workers must agree on the reason (5 person validity threshold).

Figure 11 shows histograms of the frequency of each reason leading to answer differences for (a) 29,921 visual questions asked by blind people (*VizWiz*), (b) 15,034 visual questions asked by sighted people (*VQA_2.0*), and (c) combination of the previous two. The plots are computed based on increasing thresholds of inter-worker agreement required to make a reason valid, ranging from requiring only one worker selecting it (1 person validity threshold) up to all workers agreeing (5 person threshold). The most popular reasons are ambiguous visual questions (AMB), synonymous answers (SYN), and varying answer granularity (GRN) whereas the most rare are spam (SPM) and other (OTH).

Figure 12 shows the summary of how many unique reasons are identified as the sources of answer differences for 29,921 VQs asked by blind people (*VizWiz*), 15,034 VQs asked by sighted people (*VQA_2.0*), and their combination. Across both datasets, most commonly there are three unique reasons for answer differences. Visual inspections show that these are the three most popular reasons: 'ambiguous', 'synonyms', and 'granularity'.

C. Prediction Model Analysis (supplements section 5 of the main paper)

Table 3: Average precision for predicting why answers to visual questions will differ for the VQA_2.0 and VizWiz datasets when we exclude the "spam" category for training the models.

	Model	Overall	LQI	IVE	INV	DEF	AMB	SBJ	SYN	GRN	OTH
VQA_2.0	Random	33.54	3.71	22.43	15.09	14.62	95.19	14.18	64.99	69.42	**2.25**
	QI-Relevance [30]	35.76	4.01	43.16	15.09	14.62	94.11	14.18	64.99	69.42	2.25
	I	35.38	3.66	29.55	10.06	17.04	93.04	18.29	74.50	72.09	0.18
	Q	48.11	7.91	59.23	43.43	28.02	**96.70**	23.69	88.85	84.78	0.36
	Q+I	47.87	9.18	58.83	40.83	28.18	96.48	24.20	88.39	84.37	0.37
	Q+I+A	48.05	7.09	59.46	**45.18**	27.99	96.60	21.89	88.97	85.05	0.26
	Q+I+A_FT	**48.89**	**9.11**	**59.78**	44.99	**30.11**	96.52	**24.33**	**89.49**	**85.39**	0.25
	Q+I+A_GT	48.96	8.65	60.30	46.00	28.91	96.63	24.13	89.73	86.00	0.26
VizWiz	Random	33.35	23.59	33.69	18.15	5.70	74.70	5.14	66.61	71.94	0.62
	QI-Relevance [30]	38.76	30.56	40.52	18.15	5.7	76.53	5.14	66.61	71.94	0.62
	Unanswerable [20]	43.26	44.82	58.63	18.15	5.7	80.14	5.14	66.61	71.94	0.62
	I	44.79	55.23	50.38	29.85	8.17	83.42	9.19	79.96	86.34	0.62
	Q	44.76	35.38	54.43	38.91	13.59	84.44	10.59	79.68	85.15	0.65
	Q+I	50.59	56.54	61.91	45.25	13.80	87.55	**11.55**	85.97	91.42	**1.36**
	Q+I+A	55.18	**65.51**	**77.36**	**55.76**	10.38	**89.77**	10.83	**90.39**	**95.50**	1.14
	Q+I+A_FT	**55.35**	65.30	77.18	54.19	**14.24**	89.60	11.32	89.99	95.24	1.07
	Q+I+A_GT	55.97	66.03	77.80	56.55	12.94	90.03	12.51	90.41	95.51	1.97