

# GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering

visualreasoning.net

Drew A. Hudson  
Stanford University  
353 Serra Mall, Stanford, CA 94305  
dorarad@cs.stanford.edu

Christopher D. Manning  
Stanford University  
353 Serra Mall, Stanford, CA 94305  
manning@cs.stanford.edu

## Abstract

We introduce *GQA*, a new dataset for real-world visual reasoning and compositional question answering, seeking to address key shortcomings of previous *VQA* datasets. We have developed a strong and robust question engine that leverages Visual Genome scene graph structures to create 22M diverse reasoning questions, which all come with functional programs that represent their semantics. We use the programs to gain tight control over the answer distribution and present a new tunable smoothing technique to mitigate question biases. Accompanying the dataset is a suite of new metrics that evaluate essential qualities such as consistency, grounding and plausibility. A careful analysis is performed for baselines as well as state-of-the-art models, providing fine-grained results for different question types and topologies. Whereas a blind LSTM obtains a mere 42.1%, and strong *VQA* models achieve 54.1%, human performance tops at 89.3%, offering ample opportunity for new research to explore. We hope *GQA* will provide an enabling resource for the next generation of models with enhanced robustness, improved consistency, and deeper semantic understanding of vision and language.

## 1. Introduction

It takes more than a smart guess to answer a good question. The ability to assimilate knowledge and use it to draw inferences is among the holy grails of artificial intelligence. A tangible form of this goal is embodied in the task of Visual Question Answering (VQA), where a system has to answer free-form questions by reasoning about presented images. The task demands a rich set of abilities as varied as object recognition, commonsense understanding and relation extraction, spanning both the visual and linguistic domains. In recent years, it has sparked substantial interest throughout the research community, becoming extremely

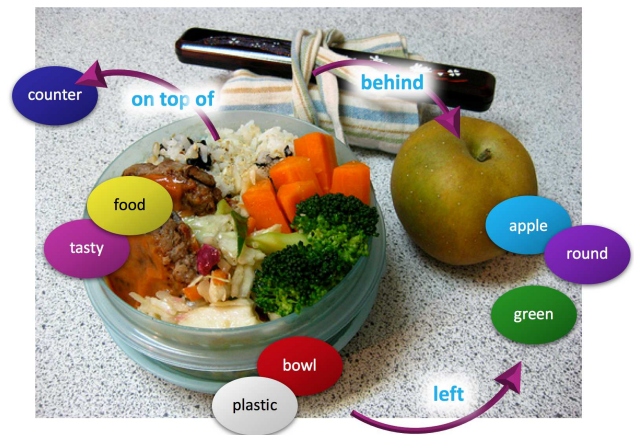


Figure 1: Examples from the new GQA dataset for visual reasoning and compositional question answering:

Is the **bow**l to the right of the **green** **apple**?

What type of **fruit** in the image is **round**?

What color is the **fruit** on the right side, **red** or **green**?

Is there any **milk** in the **bow**l to the left of the **apple**?

popular across the board, with a host of datasets being constructed [3, 9, 13, 36, 17] and numerous models being proposed [4, 33, 5, 8, 10].

The multi-modal nature of the task and the diversity of skills required to address different questions make *VQA* particularly challenging. Yet, designing a good test that will reflect its full qualities and complications may not be that trivial. Despite the great strides that the field recently made, it has been established through a series of studies that existing benchmarks suffer from critical vulnerabilities that render them highly unreliable in measuring the actual degree of visual understanding capacities [34, 9, 1, 7, 2, 11, 16].

Most notable among the flaws of current benchmarks is the strong and prevalent real-world priors displayed throughout the data [34, 9, 2] – most tomatoes are red and most tables are wooden. These in turn are exploited by

VQA models, which become heavily reliant upon such statistical biases and tendencies within the answer distribution to largely circumvent the need for true visual scene understanding [1, 9, 13, 7]. This situation is exacerbated by the simplicity of many of the questions, from both linguistic and semantic perspectives, which in practice rarely require much beyond object recognition [28]. Consequently, early benchmarks led to an inflated sense of the state of scene understanding, severely diminishing their credibility [32]. Aside from that, the lack of annotations regarding question structure and content leaves it difficult to understand the factors affecting models’ behavior and performance and to identify the root causes behind their mistakes.

To address these shortcomings, while retaining the visual and semantic richness of real-world images, we introduce GQA, a new dataset for visual reasoning and compositional question answering. We have developed and carefully refined a robust *question engine*, leveraging **content**: information about objects, attributes and relations provided through Visual Genome *Scene Graphs* [17], along with **structure**: a newly-created extensive linguistic grammar which couples hundreds of structural patterns and detailed lexical semantic resources. Together, they are combined in our engine to generate over 22 million novel and diverse questions, which all come with structured representations in the form of functional programs that specify their contents and semantics, and are visually grounded in the image scene graphs.

GQA questions involve varied reasoning skills, and multi-step inference in particular. We further use the associated semantic representations to greatly reduce biases within the dataset and control for its question type composition, downsampling it to create a 1.7M balanced dataset. Contrary to VQA 2.0, here we balance not only binary questions, but also open ones, by applying a tunable smoothing technique that makes the answer distribution for each question group more uniform. Just like a well-designed exam, our benchmark makes the educated guesses strategy far less rewarding, and demands instead more refined comprehension of both the visual and linguistic contents.

Along with the dataset, we have designed a suite of new metrics, which include consistency, validity, plausibility, grounding and distribution scores, to complement the standard accuracy measure commonly used in assessing methods’ performance. Indeed, studies have shown that the accuracy metric alone does not account for a range of anomalous behaviors that models demonstrate, such as ignoring key question words or attending to irrelevant image regions [1, 7]. Other works have argued for the need to devise new evaluation measures and techniques to shed more light on systems’ inner workings [16, 29, 30, 15]. In fact, beyond providing new metrics, GQA can even directly support the development of more interpretable models, as it provides

a sentence-long explanation that corroborates each answer, and further associates each word from both the questions and the responses with a visual pointer to the relevant region in the image, similar in nature to datasets by Zhu *et al.* [36], Park *et al.* [26], and Li *et al.* [19]. These in turn can serve as a strong supervision signal to train models with enhanced transparency and accessibility.

GQA combines the best of both worlds, having clearly defined and crisp semantic representations on the one hand but enjoying the semantic and visual richness of real-world images on the other. Our three main contributions are (1) the GQA dataset as a resource for studying visual reasoning; (2) development of an effective method for generating a large number of semantically varied questions, which marries scene graph representations with computational linguistic methods; (3) new metrics for GQA, that allow for better assessment of system success and failure modes, as demonstrated through a comprehensive performance analysis of existing models on this task. We hope that the GQA dataset will provide fertile ground for the development of novel methods that push the boundaries of question answering and visual reasoning.

## 2. Related Work

Recent years have witnessed tremendous progress in visual understanding. Multiple attempts have been made to mitigate the systematic biases of VQA datasets as discussed in section 1 [9, 34, 2, 13], but they fall short in providing an adequate solution: Some approaches operate over constrained and synthetic images [34, 13], neglecting the realism and diversity natural photos provide. Meanwhile, Goyal *et al.* [9] associate most of the questions in VQA1.0 with a pair of similar pictures that result in different answers. While offering partial relief, this technique fails to address open questions, leaving their answer distribution largely unbalanced. In fact, since the method does not cover 29% of the questions due to limitations of the annotation process, even within the binary ones biases still remain.<sup>1</sup>

At the other extreme, Agrawal *et al.* [2] partition the questions into training and validation sets such that their respective answer distributions become intentionally dissimilar. While undoubtedly challenging, these adversarial settings penalize models, maybe unjustly, for learning salient properties of the training data. In the absence of other information, making an educated guess is a legitimate choice – a valid and beneficial strategy pursued by machines and people alike [24, 6, 23]. What we essentially need is a balanced test that is more resilient to such gaming strategies, as we strive to achieve with GQA.

<sup>1</sup>For VQA1.0, blind models achieve 50% in accuracy without even considering the images whatsoever [3]. Similarly, for VQA2.0, 67% and 27% of the binary and open questions respectively are answered correctly by such models [9].

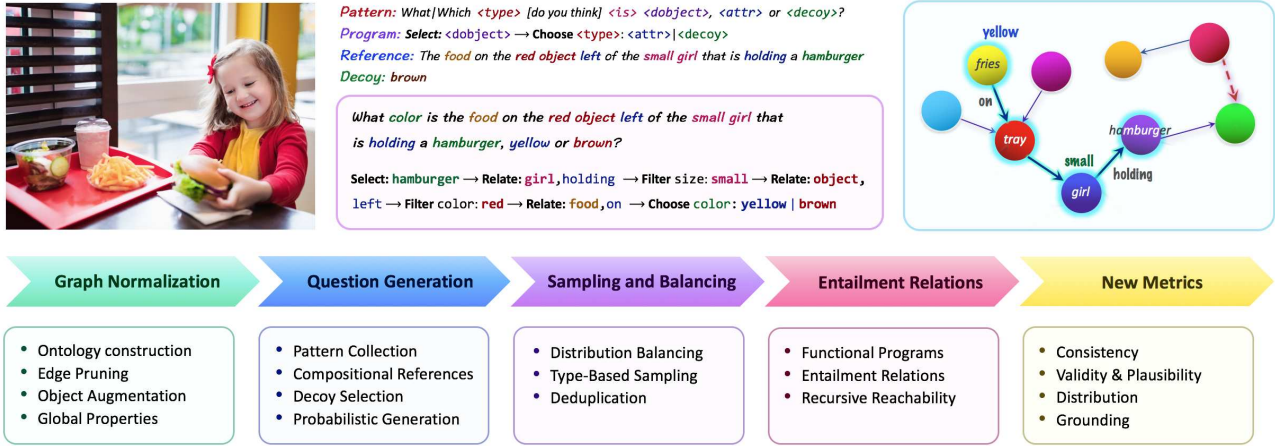


Figure 2: Overview of the GQA construction process. Given an image annotated with a scene graph of its objects, attributes and relations, we produce compositional questions by traversing the graph. Each question has both a standard natural-language form and a functional program representing its semantics. Please refer to section 3 for further detail.

In creating GQA, we drew inspiration from the CLEVR task [13], which consists of compositional questions over synthetic images. However, its artificial nature and low diversity, with only a handful of object classes and properties, makes it particularly vulnerable to memorization of all combinations, thereby reducing its effective degree of compositionality. Conversely, GQA operates over real images and a large semantic space, making it much more challenging. Even though our questions are not natural as in other VQA datasets [9, 36], they display a rich vocabulary and diverse linguistic and grammatical structures. They may serve in fact as a cleaner benchmark to assess models in a more controlled and comprehensive fashion, as discussed below.

The task of question generation has been explored in earlier work, mostly for the purpose of data augmentation. Contrary to GQA, those datasets are either small in scale [22] or use only a restricted set of objects and a handful of non-compositional templates [15, 21]. Neural alternatives to visual question generation have been recently proposed [25, 12, 35], but they aim at a quite different goal of creating engaging but potentially inaccurate questions about the wider context of the image such as subjective evoked feelings or speculative events that may lead to or result from the depicted scenes [25].

### 3. The GQA Dataset

The GQA dataset centers around real-world reasoning, scene understanding and compositional question answering. It consists of 113K images and 22M questions of assorted types and varying compositionality degrees, measuring performance on an array of reasoning skills such as object and attribute recognition, transitive relation tracking, spatial rea-

soning, logical inference and comparisons. Figure 2 provides a brief overview of the GQA components and generation process, and figure 3 presents multiple instances from the dataset. The dataset along with further information are available at [visualreasoning.net](http://visualreasoning.net).

The images, questions and corresponding answers are all accompanied by matching semantic representations: Each *image* is annotated with a dense **Scene Graph** [14, 17], representing the objects, attributes and relations it contains. Each *question* is associated with a **functional program** which lists the series of reasoning steps needed to be performed to arrive at the answer. Each *answer* is augmented with both textual and visual justifications, pointing to the relevant region within the image.

The structured representations and detailed annotations for images and questions offer multiple advantages. They enable tight control over the answer distribution, which allows us to create a balanced set of challenging questions, and support the formulation of a suite of new metrics that aim to provide deeper insight into models’ behavior. They facilitate performance assessment along various axes of question type and topology, and may open the door for the development of novel methods with more grounded and transparent knowledge representation and reasoning.

We proceed by describing the GQA question engine and the four-step dataset construction pipeline: First, we thoroughly clean, normalize, consolidate and augment the Visual Genome scene graphs [17] linked to each image. Then, we traverse the objects and relations within the graphs, and marry them with grammatical patterns gleaned from VQA 2.0 [9] and sundry probabilistic grammar rules to produce a semantically-rich and diverse set of questions. In the third stage, we use the underlying semantic forms to reduce bi-

ases in the conditional answer distribution, resulting in a balanced dataset that is more robust against shortcuts and guesses. Finally, we discuss the question functional representation, and explain how we use it to compute entailment between questions, supporting new evaluation metrics.

### 3.1. Scene Graph Normalization

Our starting point in creating the GQA dataset is the Visual Genome Scene Graph annotations [17] that cover 113k images from COCO [20] and Flickr [31].<sup>2</sup> The scene graph serves as a formalized representation of the image: each node denotes an **object**, a visual entity within the image, like a person, an apple, grass or clouds. It is linked to a bounding box specifying its position and size, and is marked up with about 1–3 **attributes**, properties of the object: e.g., its color, shape, material or activity. The objects are connected by **relation** edges, representing actions (verbs), spatial relations (prepositions), and comparatives.

The scene graphs are annotated with free-form natural language. In order to use them for question generation we first have to normalize the graphs and their vocabulary. We provide here a brief overview of the normalization process, and present a more detailed description in the supplementary. First, we create a clean, consolidated and unambiguous ontology over the graph with 2690 classes including various objects, attributes and relations. We further augment it with semantic and linguistic information which will aid us in creating grammatical questions. Then, we prune inaccurate or unnatural edges, using combination of object detection confidences, n-gram frequencies, co-occurrence statistics, word embedding distances, category-based rules, and manual curation. Finally, we enrich the graph with positional information (absolute and relative) as well as semantic properties (location, weather). By the end of this stage, the resulting scene graphs have clean, unified, rich and unambiguous semantics for both the nodes and the edges.

### 3.2. The Question Engine

At the heart of our pipeline is the question engine, responsible for producing diverse, relevant and grammatical questions with varying degrees of compositionality. The generation process harnesses two resources: one is the scene graphs which fuel the engine with rich content – information about objects, attributes and relationships; the other is the structural patterns, a mold that shapes the content, casting it into a question.

Our engine operates over 524 patterns, spanning 117 question groups, and 1878 answers which are based on the scene graphs. Each group is associated with three components: (1) a functional program that represents its semantics; (2) A set of textual rephrases which express it in nat-

<sup>2</sup>We extend Visual Genome dataset with 5k hidden scene graphs collected through crowdsourcing, used for the test set.



- A1. Is the **tray** on top of the **table** black or light brown? light brown
- A2. Are the **napkin** and the **cup** the same color? yes
- A3. Is the small **table** both oval and wooden? yes
- A4. Is there any **fruit** to the left of the **tray** the **cup** is on top of? yes
- A5. Are there any **cups** to the left of the **tray** on top of the **table**? no
- B1. What is the brown **animal** sitting inside of? **box**
- B2. What is the large **container** made of? cardboard
- B3. What **animal** is in the **box**? **bear**
- B4. Is there a **bag** to the right of the green **door**? no
- B5. Is there a **box** inside the plastic **bag**? no

Figure 3: Examples of questions from the GQA dataset.

ural language, e.g., “What|Which <type> [do you think] <is> <theObject>?”; (3) A pair of short and long answers: e.g., <attribute> and “The <object> <is> <attribute>.” respectively.<sup>3</sup>

We begin from a seed set of 250 manually constructed patterns, and extend it with 274 natural patterns derived from VQA1.0 [3] through templatization of words from our ontology.<sup>4</sup> To increase the question diversity, apart from using synonyms for objects and attributes, we incorporate probabilistic sections into the patterns, such as optional phrases  $[x]$  and alternate expressions  $(x|y)$ , which get instantiated at random.

*It is important to note that the patterns do not strictly limit the structure or depth of each question, but only outline their high-level form, as many of the template fields can be populated with nested compositional references. For instance, in the pattern above, we may replace <theObject> with “the apple to the left of the white refrigerator”.*

To achieve that compositionality, we compute for each object a set of candidate references, which can either be **direct**, e.g. *the bear*, *this animal*, or **indirect**, using modifiers, e.g. *the white bear*, *the bear on the left*, *the animal behind the tree*, *the bear that is wearing a coat*. Direct references are used when the uniqueness of the object can be confidently confirmed by object detectors, making the corresponding references unambiguous. Alternatively, we use indirect references, leading to multi-step questions as varied

<sup>3</sup>Note that the long answers can serve as textual justifications, especially for questions that require increased reasoning such as logical inference, where a question like “Is there a red apple in the picture?” may have the answer: “No, there is an apple, but it is green”

<sup>4</sup>For instance, a question-answer pair in VQA1.0 such as “What color is the apple? red” turns after templatization into “What <type> <is> <theObject>? <attribute>”.

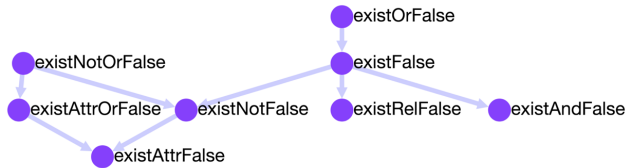


Figure 4: Examples of entailment relations between different question types. Refer to section 3.3 for further detail.

as *Who is looking at the animal that is wearing the red coat in front of the window?*, and thus greatly increasing the patterns’ effective flexibility. This is the key ingredient behind the automatic generation of compositional questions.

Finally, we compute a set of decoys for the scene graph elements. Indeed, some questions, such as negative ones or those that involve logical inference, pertain to the absence of an object or to an incorrect attribute. Examples include *Is the apple green?* for a red apple, or *Is the girl eating ice cream?* when she is in fact eating a cake. Given a triplet  $(s, r, o)$ , (e.g. *(girl, eating, cake)*) we select a distractor  $\hat{o}$  considering its likelihood to be in relation with  $s$  and its plausibility to co-occur in the context of the other objects in the depicted scene. A similar technique is applied in selecting attribute decoys (e.g. a green apple). While choosing distractors, we exclude from consideration candidates that we deem too similar (e.g. *pink* and *orange*), based on a manually defined list for each concept in the ontology.

Having all resources prepared: (1) the clean scene graphs, (2) the structural patterns, (3) the object references and (4) the decoys, we can proceed to generating the questions! We traverse the graph, and for each object, object-attribute pair or subject-relation-object triplet, we produce relevant questions by instantiating a randomly selected question pattern, e.g. “*What <type> is <theObject>, <attribute> or <CAttribute>?*”, populating all the fields with the matching information, yielding, for example, the question: “*What (color) (is) the (apple on the table), (red) or (green)?*”. When choosing object references, we avoid selecting those that disclose the answer or repeat information, e.g. “*What color is the red apple?*” or “*Which dessert sits besides the apple to the left of the cake?*”. We also avoid asking about relations that tend to have multiple instances for the same object, e.g. asking what object is on the table, as there may be multiple valid answers.

By the end of this stage, we obtain a diverse set of 22M interesting, challenging and grammatical questions, pertaining to each and every aspect of the image.

### 3.3. Functional Representation and Entailment

Each question pattern is associated with a structured representation in the form of a functional program. For instance, the question *What color is the apple on the white*

*table?* is semantically equivalent to the following program: *select: table*  $\rightarrow$  *filter: white*  $\rightarrow$  *relate(subject,on): apple*  $\rightarrow$  *query: color*. As we can see, these programs are composed of atomic operations such as object selection, traversal along a relation edge, or an attribute verification, which are then chained together to create challenging reasoning questions.

The semantically unambiguous representations offer multiple advantages over free-form unrestricted questions. For one thing, they enable comprehensive assessment of methods by dissecting their performance along different axes of question textual and semantic lengths, type and topology, thus facilitating the diagnosis of their success and failure modes (section 4.2). Second, they aid us in balancing the dataset distribution, mitigating its question-conditional priors and guarding against educated guesses (section 3.4). Finally, they allow us to identify entailment and equivalence relations between different questions: knowing the answer to the question *What color is the apple?* allows a coherent learner to infer the answer to the questions *Is the apple red?* *Is it green?* etc. The same goes especially for questions that involve logical inference like *or* and *and* operations or spatial reasoning, e.g. *left* and *right*.

As further discussed in section 4.4, this entailment property can be used to measure the coherence and consistency of the models, shedding new light on their inner workings, compared to the widespread but potentially misleading accuracy metric. We define direct entailment relations between the various functional programs and use these to recursively compute all the questions that can be entailed from a given source. A complete catalog of the functions, their associated question types, and the entailment relations between them is provided in the supplementary.

### 3.4. Sampling and Balancing

One of the main issues of existing VQA datasets is the prevalent question-conditional biases that allow learners to make educated guesses without truly understanding the presented images, as explained in section 1. However, precise representation of the question semantics can allow tighter control over these biases, having the potential to greatly alleviate the problem. We leverage this observation and use the functional programs attached to each question to smooth out the answer distribution.

Given a question’s functional program, we derive two labels, global and local: The global label assigns the question to its answer type, e.g. *color* for *What color is the apple?*. The local label further considers the main subject/s of the question, e.g. *apple-color* or *table-material*. We use these labels to partition the questions into groups, and smooth the answer distribution of each group within the two levels of granularity, first globally, and then locally.

For each group, we first compute its answer distribution

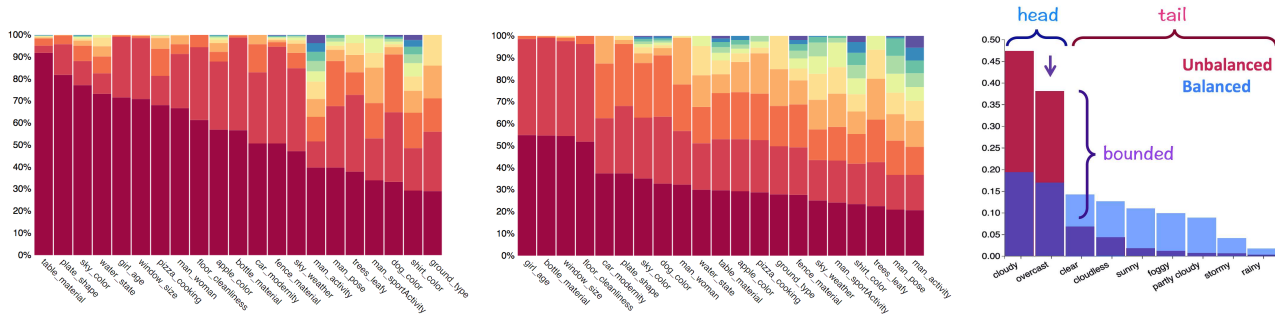


Figure 5: Visualization of the balancing process. The conditional answer distribution before (left) and after (middle) the balancing for a selection of question groups. We show the top 10 answers, where the column height corresponds to the relative frequency of each answer. We can see that on the left the distributions are heavily biased, while on the middle it is more uniform and with heavier tails, while intentionally retaining the original real-world tendencies up to a tunable degree. **Right**: An illustration of the balancing process.

$P$  for each group, which we then downsample (formally, using rejection-sampling) to fit a smoother answer distribution  $Q$  derived through the following procedure: We iterate over the answers of that group in decreasing frequency order, and reweight  $P$ 's head up to the current iteration to make it more comparable to the tail size. While repeating this operation as we go through the answers, iteratively “moving” probability from the head into the tail [27], we also maintain minimum and maximum ratios between each pair of subsequent answers (sorted by frequency). This ensures that the relative frequency-based answer ranking stays the same.

The advantage of this scheme is that it retains the general real-world tendencies, smoothing them out up to a tunable degree to make the benchmark more challenging and less biased. Refer to figure 5 for a visualization and to the supplementary for a precise depiction of the procedure. Since balancing is performed in two granularity levels, the obtained answer distributions are made more uniform both locally and globally. Quantitatively, the entropy of the answer distribution is increased by 72%, confirming the success of this stage.

Finally, we downsample the questions based on their type to control the dataset type composition, and filter out redundant questions that are too semantically similar to existing ones. We split the dataset into 70% train, 10% validation, 10% test and 10% challenge, making sure that all the questions about a given image appear in the same split.

## 4. Analysis and Baseline Experiments

In the following, we provide an analysis of the GQA dataset and evaluate the performance of baselines, state-of-the-art models and human subjects, revealing a large gap from the latter. To establish the diversity and realism of GQA questions, we test transfer performance between the GQA and VQA datasets. We then introduce the new metrics that complement our dataset, present quantitative results

and discuss their implications and merits. In the supplementary, we perform a head-to-head comparison between GQA and the popular VQA 2.0 dataset [9], and proceed with further diagnosis of the current top-performing model, MAC [10], evaluating it along multiple axes such as training-set size, question length and compositionality degree.

### 4.1. Dataset Analysis and Comparison

The GQA dataset consists of 22,669,678 questions over 113,018 images, which cover wide range of reasoning skills and vary in length and number of required inference-steps (figure 6). The dataset has a vocabulary size of 3097 words and 1878 possible answers. While smaller than natural language datasets, further investigation reveals that it covers 88.8% and 70.6% of VQA questions and answers respectively, corroborating its wide diversity. A wide selection of dataset visualizations is provided in the supplementary.

We associate each question with two types: structural and semantic. The **structural type** is derived from the final operation in the question’s functional program. It can be (1) **verify** for yes/no questions, (2) **query** for all open questions, (3) **choose** for questions that present two alternatives to choose from, e.g. “Is it red or blue?”; (4) **logical** which involve logical inference, and (5) **compare** for comparison questions between two or more objects. The **semantic type** refers to the main subject of the question: (1) *object*: for existence questions, (2) *attribute*: consider the properties or position of an object, (3) *category*: related to object identification within some class, (4) *relation*: for questions asking about the subject or object of a described relation (e.g. “what is the girl wearing?”), and (5) *global*: about overall properties of the scene such as weather or place. As shown in figure 6, the questions’ types vary at both the semantic and structural levels.

## 4.2. Baseline Experiments

We analyze an assortment of models as well as human subjects on GQA. The evaluation results are shown in table 1. Baselines include a “blind” LSTM model with access to the questions only, a “deaf” CNN model with access to the images only, an LSTM+CNN model, and two prior models based on the question group, local or global, which return the most common answer for each group, as defined in section 3.3. We can see that they all achieve low results of 17.82%–41.07%. For the LSTM model, inspection of specific question types reveals that it achieves only 22.7% for open *query* questions, and not far above chance for binary question types. We also evaluate the performance of the bottom-up attention model [4] – the winner of the 2017 VQA challenge, and the MAC model [10] – a state-of-the-art compositional attention model for CLEVR [13]. While surpassing the baselines, they still perform well below human scores<sup>5</sup>, offering ample opportunity for further research in the visual reasoning domain.

## 4.3. Transfer Performance

We tested the transfer performance between the GQA and VQA datasets, training on one and testing on the other: A MAC model trained on GQA achieves 52.1% on VQA before fine-tuning and 60.5% afterwards. Compare these with 51.6% for LSTM+CNN and 68.3% for MAC, when both are trained and tested on VQA. These quite good results demonstrate the realism and diversity of GQA questions, showing that the dataset can serve as a good proxy for human-like questions. In contrast, MAC trained on VQA gets 39.8% on GQA before fine-tuning and 46.5% afterwards, illustrating the further challenge GQA poses.

## 4.4. New Evaluation Metrics

Apart from the standard accuracy metric and the more detailed type-based diagnosis our dataset supports, we introduce five new metrics to get further insight into visual reasoning methods and point to missing capabilities we believe coherent reasoning models should possess.

**Consistency.** This metric measures responses consistency across different questions. Recall that in section 3.3, we used the questions’ semantic representation to derive equivalence and entailment relations between them. When being presented with a new question, any learner striving to be trustworthy should not contradict its previous answers. It should not answer *green* to a new question about an apple it has just identified as *red*.

For each question-answer pair  $(q, a)$ , we define a set  $E_q = q_1, q_2, \dots, q_n$  of entailed questions, the answers to

<sup>5</sup>To evaluate human performance, we used Amazon Mechanical Turk to collect human responses for 4000 random questions, taking the majority over 5 answers per question.

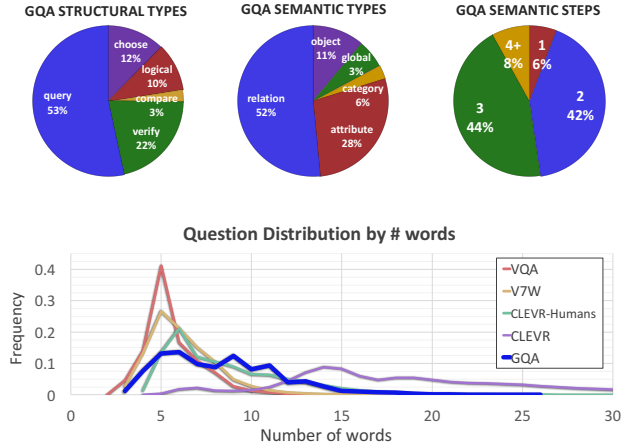


Figure 6: **Top:** Dataset statistics, partitioned into structural types, semantic types, and the number of reasoning steps. **Bottom:** VQA datasets question length distribution.

which can be unambiguously inferred given  $(q, a)$ . For instance, given the question-answer pair *Is there a red apple to the left of the white plate? yes*, we can infer the answers to questions such as *Is the plate to the right of the apple?*, *Is there a red fruit to the left of the plate?*, *What is the white thing to the right of the apple?*, etc. For each question  $q$  in  $Q$  – the set of questions *the model answered correctly*, we measure the model’s accuracy over the entailed questions  $E_q$  and then average these scores across all questions in  $Q$ .

We see that while people have exceptional consistency of 98.4%, even the best models are inconsistent in about 1 out of 5 questions, and models such as LSTM contradict themselves almost half the time. Achieving high consistency may require deeper understanding of the question semantics in the context of the image, and, in contrast with accuracy, is more robust against educated guesses as it inspects connections between related questions, and thus may serve as a better measure of models’ true visual understanding skills.

**Validity and Plausibility.** The validity metric checks whether a given answer is in the question scope, *e.g.* responding some color to a color question. The plausibility score goes a step further, measuring whether the answer is reasonable, or makes sense, given the question (*e.g.* elephant usually do not eat, say, pizza). Specifically, we check whether the answer occurs at least once in relation with the question’s subject, across the whole dataset. Thus, we consider *e.g.*, *red* and *green* as plausible apple colors and, conversely, *purple* as implausible.<sup>6</sup> The experiments show that models fail to respond with plausible or even valid answers at least 5–15% of the times, indicating limited com-

<sup>6</sup>While the plausibility metric may not be fully precise especially for infrequent objects due to potential data scarcity issues, it may provide a good sense of the general level of world knowledge the model has acquired.

Metric	Global Prior	Local Prior	CNN	LSTM	CNN+LSTM	BottomUp	MAC	Humans
Open	16.52	16.99	1.74	22.69	31.80	34.83	38.91	87.4
Binary	42.99	47.53	36.05	61.90	63.26	66.64	71.23	91.2
Query	16.52	16.99	1.55	22.69	31.80	34.83	38.91	87.4
Compare	35.59	41.91	36.34	57.79	56.62	56.32	60.04	93.1
Choose	17.45	26.58	0.85	57.15	61.40	66.56	70.59	94.3
Logical	50.32	50.11	47.18	61.73	62.05	64.03	69.99	88.5
Verify	53.40	58.80	47.02	65.78	67.00	71.45	75.45	90.1
Global	24.70	20.19	8.64	27.22	56.57	60.29	60.82	92.3
Object	49.96	54.00	47.33	74.33	75.90	78.45	81.49	88.1
Attribute	34.89	42.67	22.66	48.28	50.91	53.88	59.82	90.7
Relation	22.88	20.16	11.60	33.24	39.45	42.84	46.16	89.2
Category	15.26	17.31	3.56	22.33	37.49	41.18	44.38	90.3
Distribution	130.86	21.56	19.99	17.93	7.46	5.98	5.34	-
Grounding	-	-	-	-	-	78.47	82.24	-
Validity	89.02	84.44	35.78	96.39	96.02	96.18	96.16	98.9
Plausibility	75.34	84.42	34.84	87.30	84.25	84.57	84.48	97.2
Consistency	51.78	54.34	62.40	68.68	74.57	78.71	81.59	98.4
<b>Accuracy</b>	<b>28.93</b>	<b>31.31</b>	<b>17.82</b>	<b>41.07</b>	<b>46.55</b>	<b>49.74</b>	<b>54.06</b>	<b>89.3</b>

Table 1: Results for baselines and state-of-the-art models on the GQA dataset. All results refer to the test set. Models are evaluated for overall accuracy as well as accuracy per type. In addition, they are evaluated by validity, plausibility, distribution, consistency, and when possible, grounding metrics. Please refer to the text for further detail.

prehension of some questions. Given that these properties are noticeable statistics of the dataset’s conditional answer distribution, not even depending on the specific images, we would expect a sound method to achieve higher scores.

**Distribution.** To get further insight into the extent to which methods manage to model the conditional answer distribution, we define the distribution metric, which measures the overall match between the true answer distribution and the model predicted distribution, using Chi-Square statistic [18]. It allows us to see if the model predicts not only the most common answers but also the less frequent ones. Indeed, the experiments demonstrate that the leading SOTA models score lower than the baselines (for this metric, lower is better), indicating increased capacity in fitting more subtle trends of the dataset’s distribution.

**Grounding.** For attention-based models, the grounding score checks whether the model attends to regions within the image that are relevant to the question. For each dataset instance, we define a pointer  $r$  to the visual region which the question or answer refer to, and measure the model’s visual attention (probability) over that region. This metric allows us to evaluate the degree to which the model grounds its reasoning in the image, rather than just making educated guesses based on question priors or world tendencies.

Indeed, the models mostly attend to the relevant regions in the image, with grounding scores of about 80%. To verify the reliability of the metric, we further perform experiments with spatial features instead of the object-informed ones used by BottomUp [4] and MAC [10], which lead to a much lower 43% score, demonstrating that object-based features provide models with better granularity for the task, allowing them to focus on more pertinent regions than with the coarser spatial features.

## 5. Conclusion

In this paper, we introduced the GQA dataset for real-world visual reasoning and compositional question answering. We described the dataset generation process, provided baseline experiments and defined new measures to get more insight into models’ behavior and performance. We believe this benchmark can help drive VQA research in the right directions of deeper semantic understanding, sound reasoning, enhanced robustness and improved consistency. A potential avenue towards such goals may involve more intimate integration between visual knowledge extraction and question answering, two flourishing fields that oftentimes have been pursued independently. We strongly hope that GQA will motivate and support the development of more compositional, interpretable and cogent reasoning models, to advance research in scene understanding and visual question answering.

## 6. Acknowledgments

We wish to thank Justin Johnson for discussions about the early versions of this work, and Ross Girshick for his inspirational talk at the VQA workshop 2018. We thank Ranjay Krishna, Eric Cosatto, Alexandru Niculescu-Mizil and the anonymous reviewers for helpful suggestions and comments. Stanford University gratefully acknowledges Facebook Inc., Samsung Electronics Co., Ltd., and the Defense Advanced Research Projects Agency (DARPA) Communicating with Computers (CwC) program under ARO prime contract no. W911NF15-1-0462 for generously supporting this work.



## References

- [1] A. Agrawal, D. Batra, and D. Parikh. Analyzing the behavior of visual question answering models. In *EMNLP*, pages 1955–1960, 2016. 1, 2
- [2] A. Agrawal, D. Batra, D. Parikh, and A. Kembhavi. Don’t just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4971–4980, 2018. 1, 2
- [3] A. Agrawal, J. Lu, S. Antol, M. Mitchell, C. L. Zitnick, D. Parikh, and D. Batra. VQA: Visual question answering. *International Journal of Computer Vision*, 123(1):4–31, 2017. 1, 2, 4
- [4] P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, and L. Zhang. Bottom-up and top-down attention for image captioning and VQA. *arXiv preprint arXiv:1707.07998*, 2017. 1, 7, 8
- [5] J. Andreas, M. Rohrbach, T. Darrell, and D. Klein. Neural module networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 39–48, 2016. 1
- [6] Y. Attali and M. Bar-Hillel. Guess where: The position of correct answers in multiple-choice test items as a psychometric variable. *Journal of Educational Measurement*, 40(2):109–128, 2003. 2
- [7] A. Das, H. Agrawal, L. Zitnick, D. Parikh, and D. Batra. Human attention in visual question answering: Do humans and deep networks look at the same regions? *Computer Vision and Image Understanding*, 163:90–100, 2017. 1, 2
- [8] A. Fukui, D. H. Park, D. Yang, A. Rohrbach, T. Darrell, and M. Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016. 1
- [9] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh. 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 (CVPR)*, pages 6325–6334, 2017. 1, 2, 3, 6
- [10] D. A. Hudson and C. D. Manning. Compositional attention networks for machine reasoning. *International Conference for Representation Learning (ICLR)*, 2018. 1, 6, 7, 8
- [11] A. Jabri, A. Joulin, and L. van der Maaten. Revisiting visual question answering baselines. In *European conference on computer vision*, pages 727–739. Springer, 2016. 1
- [12] U. Jain, Z. Zhang, and A. G. Schwing. Creativity: Generating diverse questions using variational autoencoders. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5415–5424, 2017. 3
- [13] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, and R. Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1988–1997, 2017. 1, 2, 3, 7
- [14] J. Johnson, R. Krishna, M. Stark, L.-J. Li, D. Shamma, M. Bernstein, and L. Fei-Fei. Image retrieval using scene graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3668–3678, 2015. 3
- [15] K. Kafle and C. Kanan. An analysis of visual question answering algorithms. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 1983–1991. IEEE, 2017. 2, 3
- [16] K. Kafle and C. Kanan. Visual question answering: Datasets, algorithms, and future challenges. *Computer Vision and Image Understanding*, 163:3–20, 2017. 1, 2
- [17] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, et al. Visual genome: Connecting language and vision using crowd-sourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2017. 1, 2, 3, 4
- [18] H. O. Lancaster and E. Seneta. Chi-square distribution. *Encyclopedia of biostatistics*, 2, 2005. 8
- [19] Q. Li, Q. Tao, S. Joty, J. Cai, and J. Luo. VQA-E: Explaining, elaborating, and enhancing your answers for visual questions. *arXiv preprint arXiv:1803.07464*, 2018. 2
- [20] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 4
- [21] A. Mahendru, V. Prabhu, A. Mohapatra, D. Batra, and S. Lee. The promise of premise: Harnessing question premises in visual question answering. *arXiv preprint arXiv:1705.00601*, 2017. 3
- [22] M. Malinowski and M. Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *Advances in neural information processing systems*, pages 1682–1690, 2014. 3
- [23] J. Millman, C. H. Bishop, and R. Ebel. An analysis of test-wiseness. *Educational and Psychological Measurement*, 25(3):707–726, 1965. 2
- [24] J. J. Mondak and B. C. Davis. Asked and answered: Knowledge levels when we won’t take ‘don’t know’ for an answer. *Political Behavior*, 23(3):199–224, 2001. 2
- [25] N. Mostafazadeh, I. Misra, J. Devlin, M. Mitchell, X. He, and L. Vanderwende. Generating natural questions about an image. *arXiv preprint arXiv:1603.06059*, 2016. 3
- [26] D. H. Park, L. A. Hendricks, Z. Akata, A. Rohrbach, B. Schiele, T. Darrell, and M. Rohrbach. Multimodal explanations: Justifying decisions and pointing to the evidence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [27] Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121, 2000. 6
- [28] A. Suhr, S. Zhou, I. Zhang, H. Bai, and Y. Artzi. A corpus for reasoning about natural language grounded in photographs. *arXiv preprint arXiv:1811.00491*, 2018. 2
- [29] D. Teney, P. Anderson, X. He, and A. van den Hengel. Tips and tricks for visual question answering: Learnings from the 2017 challenge. *arXiv preprint arXiv:1708.02711*, 2017. 2

- [30] D. Teney, L. Liu, and A. van den Hengel. Graph-structured representations for visual question answering. *arXiv preprint*, 2017. 2
- [31] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li. Yfcc100m: The new data in multimedia research. *arXiv preprint arXiv:1503.01817*, 2015. 4
- [32] A. Torralba and A. A. Efros. Unbiased look at dataset bias. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1521–1528. IEEE, 2011. 2
- [33] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21–29, 2016. 1
- [34] P. Zhang, Y. Goyal, D. Summers-Stay, D. Batra, and D. Parikh. Yin and yang: Balancing and answering binary visual questions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5014–5022, 2016. 1, 2
- [35] S. Zhang, L. Qu, S. You, Z. Yang, and J. Zhang. Automatic generation of grounded visual questions. *arXiv preprint arXiv:1612.06530*, 2016. 3
- [36] Y. Zhu, O. Groth, M. Bernstein, and L. Fei-Fei. Visual7W: Grounded question answering in images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4995–5004, 2016. 1, 2, 3