



HAL
open science

Roses Are Red, Violets Are Blue... but Should Vqa Expect Them To?

Corentin Kervadec, Grigory Antipov, Moez Baccouche, Christian Wolf

► **To cite this version:**

Corentin Kervadec, Grigory Antipov, Moez Baccouche, Christian Wolf. Roses Are Red, Violets Are Blue... but Should Vqa Expect Them To?. 2020. hal-02860823v1

HAL Id: hal-02860823

<https://hal.science/hal-02860823v1>

Preprint submitted on 8 Jun 2020 (v1), last revised 8 Apr 2021 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ROSES ARE RED, VIOLETS ARE BLUE... BUT SHOULD VQA EXPECT THEM TO?

Corentin Kervadec

Orange Labs, Cesson-Sévigné, France
Université de Lyon, INSA-Lyon, LIRIS
UMR CNRS 5205, Villeurbanne, France
corentin.kervadec@orange.com

Grigory Antipov

Orange Labs,
Cesson-Sévigné, France
grigory.antipov@orange.com

Moez Baccouche

Orange Labs,
Cesson-Sévigné, France
moez.baccouche@orange.com

Christian Wolf

Université de Lyon, INSA-Lyon, LIRIS
UMR CNRS 5205, Villeurbanne, France
christian.wolf@insa-lyon.fr

ABSTRACT

To be reliable on rare events is an important requirement for systems based on machine learning. In this work we focus on Visual Question Answering (VQA), where, in spite of recent efforts, datasets remain imbalanced, causing shortcomings of current models: tendencies to overly exploit dataset biases and struggles to generalise to unseen associations of concepts. We focus on a systemic evaluation of model error distributions and address fundamental questions: How is the prediction error distributed? What is the prediction accuracy on infrequent vs. frequent concepts? In this work, we design a new benchmark based on a fine-grained reorganization of the GQA dataset [1], which allows to precisely answer these questions. It introduces distributions shifts in both validation and test splits, which are defined on question groups and are thus tailored to each question. We performed a large-scale study and we experimentally demonstrate that several state-of-the-art VQA models, even those specifically designed for bias reduction, fail to address questions involving infrequent concepts. Furthermore, we show that the high accuracy obtained on the frequent concepts alone is mechanically increasing overall accuracy, covering up the true behavior of current VQA models.

1 Introduction

Visual Question Answering (VQA), i.e. the task of answering a question posed over an image, is often seen as a testbed for the capability of learning-based systems to perform high-level reasoning. This multimodal (language and image) problem is notorious for its large diversity meaning that VQA models are required to learn various high-level general representations of concepts of the physical world as well as their interactions.

Efforts to learn the necessary high-level reasoning from large-scale datasets depend on the absence of harmful biases in the data, which could provide unwanted shortcuts to learning in the form of “Clever Hans” effects. Unfortunately, and in spite of recent efforts [2, 1], most VQA datasets remain very imbalanced. Common concepts are significantly more frequent, e.g the presence of a “*red rose*”, compared to out of context concepts like the presence of a “*zebra in a city*”. This causes the tendency of models to overly rely on biases, hindering generalisation [3, 4]. Despite a general consensus on this diagnostic, systemic evaluations of error distributions are rare. In particular, overall accuracy is still the major, and often unique, metric used to evaluate models and methods, although it is clearly insufficient. Several questions remain open. How is error distributed? Are true positives due to reasoning or to exploitation of bias? What is the prediction accuracy on infrequent vs. frequent concepts? How can we validate models in Out Of Distribution (OOD)-settings?

In this work we propose a new benchmark and a study on State-Of-The-Art (SOTA) VQA models, which allows to precisely answer these questions. The proposed new evaluation protocol is complementary to existing ones, but allows

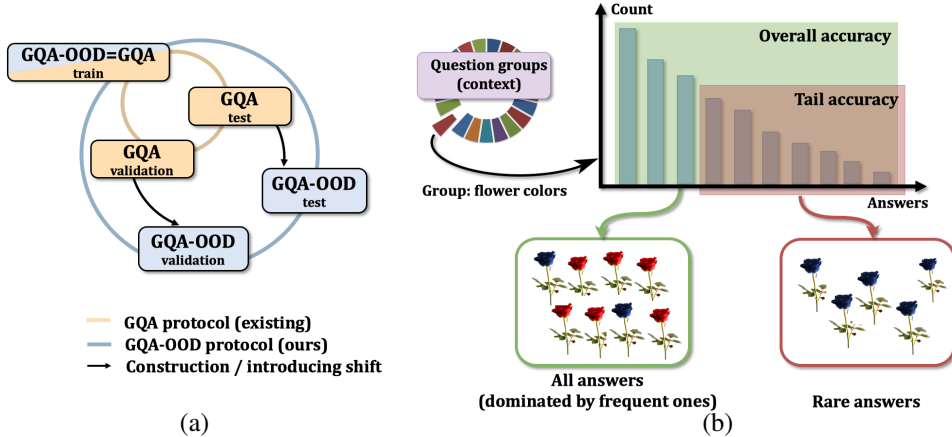


Figure 1: We re-organize the GQA dataset [1] in a fine-grained way: (a) the benchmark contains a distribution shift in validation and test, allowing to validate *and* evaluate in OOD settings; (b) our distribution shifts are tailored to different question groups with highly imbalanced distributions. Illustrated here: the group raising questions on rose colors.

a better diagnostic of current VQA performance. The benchmark comprises (i) a new fine-grained reorganization of the GQA dataset [1] introducing distribution shifts in both validation and test sets (see Figure 1-a); (ii) a set of evaluation metrics; (iii) new evaluation plots illustrating the reasoning behavior of VQA models on different operating points. The choice of the GQA dataset is motivated by its useful structuring into question groups, which allows to capture biases precisely, to select groups with strong biases and to create distribution shifts tailored to the exact nature of each question (see Figure 1-b). It also makes it possible to analyze how errors are distributed over different associations of concepts according to their frequency in the dataset.

We validate the benchmark with a large study and we experimentally demonstrate that several SOTA VQA models fail to address questions involving infrequent concepts. We also test several SOTA methods for bias reduction and come to similar conclusions. Furthermore, we show that the high accuracy obtained on the frequent concepts alone is mechanically increasing overall accuracy, covering up the true behavior of current VQA models.

The contributions of this paper are summarized as follows:

1. We propose and make public¹ a new fine-grained re-organization of the GQA dataset and a set of the respective evaluation metrics allowing to precisely evaluate the reasoning behavior of VQA models and to characterise and visualise their generalisation behavior on different operating points w.r.t distribution shifts.
2. Compared to competing benchmarks, our dataset features distribution shifts for both, validation and test, allowing to validate models under OOD conditions.
3. In a large study, we evaluate several recent VQA models and show that they struggle to generalise in OOD conditions; we also test several SOTA bias reduction methods and show that there is still room for improvement in addressing bias in VQA.

2 Related Work

VQA corpuses —One of the first large-scale VQA datasets was VQA1 [5] gathering about 76K questions posed over 25K realistic images. It started a new task, but was soon found to suffer from biases. The authors of [2] pointed to strong imbalances among the presented answers and proposed the second (improved) version of the dataset: VQA2. In parallel, [6] introduced the fully synthetic CLEVR dataset, conceived to evaluate the reasoning capabilities in a simple environment. While being synthetic its strong point are detailed and structured annotations. The authors of [1] adapted CLEVR to real-world images and constructed the GQA dataset, gathering 1.7M balanced questions. As this corpus is automatically generated, it offers a better control on the dataset statistics.

Attempts to reduce the bias-dependency — Despite efforts to design complex architectures, VQA models suffer from a significant generalisation incapabilities. As shown in [7], VQA models tend to answer the questions without using the image, and even when they do, they do not always exploit the relevant visual regions [8]. Moreover, VQA models tend to overly rely on dataset biases [9], and are not able to generalise to unseen distributions [10]. In this context,

¹<https://github.com/gqaood/GQA-OOD>

several methods have been designed to alleviate this dependency towards dataset biases. For example, one can set up an adversarial game against a question-only adversary to regularise training [11]. In a similar way, RUBi [3] uses a question-only branch in addition to a base model during training to prevent it from learning textual biases. At test time, the question-only branch is omitted. In concurrent work, [4] adopted a similar strategy by regularizing the model predictions using question type statistics from the training set. Other works [12, 13] attempted to force VQA models to attend to the most significant visual regions (from humans’ perspective). While these methods show promising results when evaluated on unseen distributions [10], they generally slightly degrade performances on the standard benchmarks, which tend to favor models relying on dataset biases.

Reinventing VQA evaluation — In parallel to the design of more and more complex but biased VQA architectures, we observe the need of new ways to evaluate VQA models. Early work [14] proposed to evaluate the answer prediction using a soft evaluation score based on a lexical database. However, this metric has been replaced by a hard classification score, more prone to favor biased predictions, but easier to use in practice. The authors of GQA [1] proposed several additional metrics along with their dataset, namely: consistency, plausibility, validity, distribution, *etc.* These metrics give a better insight on the VQA performances, but do not evaluate its ability to predict correct answers in OOD setting, especially when applied to the original balanced GQA dataset. To evaluate the generalisation capability of Neural State Machines, [15] proposed interesting splits of GQA. In particular, they reorganized the train and validation splits to distinctly evaluate how well the models generalize on the content of the visual scene, and on the linguistic structure of the question. The authors of [16] analysed the generalisation to unseen associations of known linguistic structures using their CLOSURE benchmark. They demonstrated that SOTA models (including those which were explicitly designed to overcome this issue) fail to generalise in these conditions. This confirms the need of carefully designed benchmarks to evaluate the true capabilities of VQA models. In this context, [10] introduced a reorganisation of the VQA2 [2] dataset splits, namely VQA-CP2, where the training distribution is made explicitly different from the one in the test. This new setup allows to evaluate the generalisation on an unseen distribution. However, as shown in Section 3, our proposed benchmark is complementary and conceptually different from VQA-CP2, especially in its construction process, since it introduces structured distribution shifts defined on local question groups.

3 GQA-OOD: a benchmark for OOD settings

We introduce a new VQA benchmark named *GQA-OOD* designed to evaluate VQA models and algorithms in Out-Of-Distribution (OOD) configurations. We here define OOD samples as rare events, in particular measured w.r.t. to a base distribution, *e.g.* a training distribution. It is worth noticing that these rare events might involve concepts which are also present in the training set. As an illustration, consider the question: ‘*What color is this rose?*’. If the image represents a rose, then *red* would be a common color, but in an OOD setting, frequent (correct) test answers would be, for instance, *blue*, requiring models to reason to provide the correct answer. We insist that in our benchmark this shift in distributions is not global and depends on the context. If the context changes, and the flower type is a violet then a (correct) OOD answer would now be *red* instead of *blue*.

The *GQA-OOD* benchmark consists of a dataset and new evaluation metrics. The dataset itself is based on the existing GQA [1] dataset, which provides more fine-grained annotations compared to competing VQA2 [2]. GQA questions have been automatically generated from scene graphs, which allows better control of the question context. Figure 1-a shows how the proposed protocol compares to the existing GQA protocol: the two share the same (existing) training set, but we introduce fine-grained shifts into both the validation and the test sets applying the same process as further developed below².

The shifted subsets have been constructed in 3 steps: (i) dividing questions into groups according to their contexts; (ii) extracting the most imbalanced question groups considering their answer distributions; (iii) selecting OOD samples among the remaining questions.

Question groups — To structure the process introducing the distribution shift, we use the notion of *local groups* provided in the GQA annotation. They allow to precisely define the type of a question, *e.g.* ‘*What color ...?*’, ‘*Where is ...?*’, *etc.* They also depend on the specific concept related to the question, *e.g.* ‘*zebra*’, ‘*violet*’, *etc.* There is a total of roughly 37K local groups related to about 132K questions in the GQA validation split. We use the balanced version of GQA, whose question distribution has been smoothed in order to obtain a more uniform answer distribution. However, this does not impact the imbalanced nature of the dataset, which is often due to real world tendencies, *e.g.* that ‘*roses are red*’.

Measuring group imbalance — We extract a subset of the most imbalanced question groups, as we are interested in evaluating the prediction error specifically in the context, where shifts in distribution are meaningful and strong. We

²We use version 1.2 of the GQA dataset [1].

Table 1: Statistics of the proposed dataset: (a) GQA-OOD vs. GQA; (b) head vs. tail split .

Dataset	Split	#Quest.	#Groups	#Imgs	Split	Subset	#Quest.	#Groups	#Imgs
GQA-OOD	val	51,045	3,849	9,406	val	head	33,882	3,849	8,664
	testdev	2,796	471	388		tail	17,163	3,849	6,632
GQA	val	132,062	36,832	10,234	testdev	head	1,733	471	365
	testdev	12,578	7,803	398		tail	1,063	471	330

(a)

(b)

measure the balance through Shannon entropy given as $e(x) = -\sum_{i=0}^d p(x_i) \log p(x_i)$, where $p(x_i)$ is the estimated probability of the class i . As entropy depends on the number of answer classes, which is highly variable between different question groups, we normalize entropy w.r.t. the to number d of possible answers in the group: $\bar{e}(x) = \frac{e(x)}{\log(d)}$, where $\log(d)$ is equal to the entropy of a uniform distribution of size d . Normalized entropy $\bar{e}(x)$ thus measures how close the distribution $p(x)$ is to a uniform distribution of the same dimension. On a side note, roughly 70% of questions groups have a dimension $d=1$, i.e. only one possible answer has been found. As they do not represent a large proportion of the total amount of questions, we automatically discard them. And, finally, we keep groups with a normalised entropy smaller than a threshold T , which we empirically set to $T=0.9$. This selects all the questions of the *GQA-OOD* benchmark, but further work has been done in order to select specific answer classes for each group.

OOD setting and metrics — We introduce a shift in distribution by selecting a subset of answer classes for each question group according to their frequencies, and introduce three different metrics according to which classes are used for evaluation. All these metrics are defined over the aforementioned imbalanced local groups.

- **Acc-tail:** the accuracy on OOD samples, which are the samples of the tail of the answer class distribution, *i.e.* the most rare answers given the context. We define the tail classes as classes i with $|a_i| \leq (1 + \alpha)\mu(a)$, where $|a_i|$ is the number of samples belonging to the class i and $\mu(a)$ is the average sample count for the group. We empirically set the parameter $\alpha = 0.2$, and in Section 4.1 we analyse and illustrate the impact of the choice of α on *Acc-tail*.
- **Acc-head:** the accuracy on the head of distribution for each local group (by the head of the group we understand the set difference between the whole group and its tail).
- **Acc-all:** the overall (classical) accuracy over all *GQA-OOD* samples.

Table 1 provides statistics of the proposed benchmark. We also analysed the nature, distribution and the diversity of the questions w.r.t to GQA [1], *c.f.* supplementary material (Section A).

Difference with VQA-CP2 — The excellent VQA-CP2 [10] dataset was a first of its kind and paved the way for follow-up work on bias reduction methods in VQA. However, its construction is conceptually different from our work, partially due to the restrictions of the base dataset VQA2 w.r.t. to GQA, but also due to key design choices.

Lacking annotations on group structure in the base dataset, in VQA-CP questions are grouped together according to their first words and the same ground-truth answer. The shift is created by avoiding repeated types between splits. In contrast, our proposed *GQA-OOD* dataset allows fine-grained analysis of the generalisation behavior of a VQA model by (i) question group, and via (ii) different metrics corresponding to different amounts of shifts (*acc-tail* vs. *acc-head*), and (iii) even through the possibility of continuous evaluation along different operating points (see Figure 2).

VQA-CP2 comprises two splits only (train+test), lacking the possibility of validating model hyper-parameters. Most techniques therefore seem to optimize their hyper-parameters on the test split [3, 4, 11, 12, 13], which should be frowned upon, as the evaluation overfits on the test set, or, alternatively, validate on a subset of train which does not include a shift [17]. Our *GQA-OOD* dataset contains a validation set with a shift w.r.t. to the train set, which allows to validate hyper-parameters in OOD settings.

Finally, unlike VQA-CP, the proposed *GQA-OOD* dataset requires models to be trained on the existing GQA train split. This requires models to reduce bias in their test results while being exposed to natural tendencies and biases captured in the training corpus, favoring work on bias reduction through methodology instead of through cleaning of training data.

4 Experiments

We evaluated several SOTA models on the proposed *GQA-OOD* benchmark and compared with evaluations on the benchmarks VQA2 [2], GQA [1] and VQA-CP2 [10]. The selected line-up includes strong models with object-level

attention and one Transformer [18]-based model, as well as two blind baseline models (more training details are given in the supplementary material section B):

Question Prior — this blind baseline returns the most probable answer estimated from training set statistics. Following [10, 2], we use the question types priors when evaluating on VQA-CP and VQA2. For GQA, we use the training global group priors.

LSTM [5] — this blind baseline takes GloVe embeddings [19] and encodes them using an LSTM [20] followed by a feed-forward network.

BUTD [21] — a strong VQA baseline based on object-level attention, in particular, bounding boxes with dense visual feature vectors extracted from the image using an object detector.

MCAN [22] — this SOTA approach is based on a Transformer [18] architecture and designed to model both intra-modality interactions and the inter-modality ones. It allows complex multi-hop reasoning through several stacked self-attention blocks. In our experiments, we use the 6-layers version of MCAN.

In addition to these four VQA baselines, we also evaluate three bias-reduction methods widely studied on the VQA-CP [10] dataset. These methods are designed to be model-agnostic [3, 4]. In this work, we use them in conjunction with the BUTD architecture.

RUBi [3] — adds a question-only branch to the base model during training to prevent it from learning question biases. This branch is omitted during evaluation. To better analyze bias dependencies, we also study a modified version of RUBi, which we refer to as **RUBi+QB** below. In this variant, the question-only branch is kept during evaluation.

BP [4] — is similar to RUBi but differs by directly taking training set statistics to infer question type biases during training³. The question type biases are fused with the base model predictions using a product of experts [4], and removed during testing.

LMH [4] — is an improved version of BP [4]. In this version, the question bias is dynamically weighted by the base model in order to control its influence. In the original setup, an entropy penalty is added to the loss to prevent the model to ignore the bias. Nevertheless, when training on GQA, we obtain better results without this penalty (see supp material section B for details).

Training details for the aforementioned models and methods can be found in the supp material section B. Models evaluated on *GQA-OOD* are trained on the training set of GQA (balanced) and validated on the validation split of *GQA-OOD*. When available, we provide the standard deviation computed over at least four different seeds.

4.1 Analysing the prediction error distribution

The *GQA-OOD* benchmark allows us to perform an analysis of the error prediction distributions for various VQA models. Table 2 shows the results obtained when evaluating several VQA models on *GQA-OOD*. We provide the three metrics introduced in Section 3: *acc-tail*, *acc-head* and *acc-all*. We also measure the difference $\Delta(\text{tail}, \text{head}) = \frac{\text{acc-head} - \text{acc-tail}}{\text{acc-tail}}$ to illustrate how much is the error prediction imbalanced between frequent and rare answers.

Models fail on rare question-answer pairs — As one can notice, the two blind models (Question Prior and LSTM in Table 2-a) obtain the highest gap between *acc-tail* and *acc-head*. Indeed, as they do not have access to the image, they rely on question biases. The Δ score indicates that both BUTD and MCAN also struggle (in a lesser extent) to generalise on the less frequent question-answer pairs. However, MCAN outperforms BUTD on both metrics, confirming the superiority of the Transformer-based architecture.

Overall accuracy is dominated by frequent question-answer pairs — It is interesting to observe that *acc-all*, which is the standard metric in VQA, does not reflect the true model’s performances, since it is mechanically increased by the high score obtained on the most frequent question-answers. This confirms the need of a specific evaluation in OOD settings as provided by *acc-tail*.

Bias-reduction methods are not efficient on rare samples — Surprisingly, none of the three bias-reduction methods (namely RUBi [3], BP [4] and LMH [4]) succeed to improve *acc-tail* (cf. Table 2-b). They even seem to deteriorate *acc-head*. This is unexpected as these methods have been conceived to overcome the dependency toward question type biases. In order to better understand these results, we evaluate RUBi while keeping the question-only branch during testing (RUBi+QB). Expectedly, it outperforms RUBi on *acc-head* showing it has better captured frequent patterns.

³When training on VQA2 and VQA-CP2, biases are computed over question types. On GQA, biases are computed over local groups.

Table 2: Comparison of several methods on the GQA-OOD testdev split. *Acc-tail*: OOD settings, *Acc-head*: accuracy on the most probable answers given the context. (a) Different models; (b) different bias reduction techniques on the BUTD [21] model. All scores in %.


Model	acc-all	acc-tail	acc-head	Δ
Quest. Prior	21.6	17.8	24.1	35.4
LSTM [5]	30.7	24.0	34.8	45.0
BUTD [21]	46.4 \pm 1.1	42.1 \pm 0.9	49.1 \pm 1.1	16.6
MCAN [22]	50.8 \pm 0.4	46.5 \pm 0.5	53.4 \pm 0.6	14.8

(a)

Technique	acc-all	acc-tail	acc-head	Δ
BUTD [21]	46.4 \pm 1.1	42.1 \pm 0.9	49.1 \pm 1.1	16.6
+RUBi+QB	46.7 \pm 1.3	42.1 \pm 1.0	49.4 \pm 1.5	17.3
+RUBi [3]	38.8 \pm 2.4	35.7 \pm 2.3	40.8 \pm 2.7	14.3
+LMH [4]	34.5 \pm 0.7	32.2 \pm 1.2	35.9 \pm 1.2	11.5
+BP [4]	33.1 \pm 0.4	30.8 \pm 1.0	34.5 \pm 0.5	12.0

(b)

Figure 2: Performance for different definitions of the tail distribution (α parameter values): (a) different VQA models; (b) different bias reduction techniques. The x-axis is in log-scale.



However, it also outperforms RUBi on the OOD settings, demonstrating that preventing from learning frequent patterns does not necessarily increase performances on rare samples.

Visualising the generalisation behavior — In Figure 2, we analyse how the prediction error is distributed on the validation split. In particular, we plot accuracy *acc-tail* as a function of the tail’s size, which is controlled using the α parameter described in Section 3. As shown in Figure 2-a, LSTM, BUTD [21] and MCAN [22] follow the same dynamic: starting from a tail size which represents roughly half of the question-answer pairs, the accuracy starts to linearly decrease until reaching a dramatically low score (about 30 pts lower than the overall accuracy). This demonstrates the real need to take the prediction error distribution into consideration when benchmarking VQA models. We perform the same analysis on bias-reduction methods in Figure 2-b. For BP [4], LMH [4] and, to a lesser extent, RUBi [3], we observe that the right side of the curve has been flattened. This demonstrates that the overall accuracy, dominated by frequent question-answer pairs, has been reduced because of the bias-reduction method. On the other hand, the left side of the curve, corresponding to the rare samples, remains almost unchanged. This reveals that these methods have somewhat succeed in preventing the base model from learning dataset biases as the error distribution between frequent and rare samples is more uniform. As a comparison, the LSTM model in Figure 2-a performs worse than BUTD [21] but conserves the same frequent/rare imbalance. We observe that RUBi+QB responds the same way as BUTD [21], confirming the effect of the bias-reduction method. Thus, although it is clear that these methods succeed in preventing the base model from learning salient properties of the training set, we observe that they do not help it to learn subtle distributions.

4.2 Comparison with other benchmarks

Below, we compare *GQA-OOD* with the following three standard VQA benchmarks:

GQA [1] (balanced version) is a VQA dataset gathering about 1.7M question-answer pairs automatically generated from real images. Among all other VQA datasets, GQA has the richest annotations allowing to evaluate models with complimentary metrics: consistency, validity, plausibility, [1] *etc.* In this paper, we only exploit the overall accuracy and the distribution score on the GQA [1] testdev split as the other metrics are unrelated w.r.t. the subject of our paper.

Table 3: We compare the proposed *acc-tail* metric with other benchmarks. Results are computed on the testdev split of *GQA-OOD* and GQA [1], the test split of VQA-CP2 [10] and the VQA2 [2] validation split. Values in italic have been trained and tested by ourselves.

Model	VQA2		GQA	VQA-CP2	GQA-OOD
	overall	overall	distribution	overall	acc-tail
Question Prior	32.1	<i>27.0%</i>	<i>55.62</i>	8.8%	<i>17.8%</i>
LSTM [5]	43.0	<i>39.1%</i>	<i>3.6</i>	<i>22.1%</i>	<i>24.0%</i>
BUTD [21]	63.5	<i>51.6%±0.3</i>	<i>1.8</i>	40.1%	<i>42.1%±0.9</i>
MCAN [22]	<i>66.1%</i>	<i>56.3%±0.2</i>	<i>1.59</i>	<i>42.5%</i>	<i>46.5%±0.5</i>
BUTD [21]	<i>63.5%</i>	<i>51.6%±0.3</i>	<i>1.8</i>	40.1%	<i>42.1%±0.9</i>
+RUBi+QB	-	<i>51.9%±1.1</i>	<i>1.68</i>	<i>47.6%±3.7</i>	<i>42.1%±1.0</i>
+RUBi [3]	<i>61.2%</i>	<i>43.6%±2.0</i>	<i>1.88</i>	44.2%	<i>35.7%±2.3</i>
+LMH [4]	<i>56.4%</i>	<i>39.7%±0.7</i>	<i>2.1</i>	52.0%	<i>32.2%±1.2</i>
+BP [4]	<i>63.2%</i>	<i>39.6%±0.3</i>	<i>2.2</i>	39.9%	<i>30.8%±1.0</i>

The distribution score is obtained by measuring the match between the true answer distribution and the model predicted distribution.

VQA2 [2] is composed of questions posed by humans about real-world images. The dataset contains 265K images with at least 3 questions for each. Each question is annotated with 10 ground-truth answers. We compare our benchmark with the overall accuracy on the VQA2 [2] validation split.

VQA-CP2 [10] has been constructed by reorganising the training and validation splits of the VQA2 [2] dataset in order to explicitly make the distributions of answers different between the training and test splits. The VQA-CP2 dataset has been designed to measure the sensitivity of a VQA model to the language bias, in other words, as a test for measuring the ability of a model to generalise on unseen data.

Comparison with GQA and VQA2 — In Table 3, we compare our *acc-tail* score with the other benchmarks. One can observe that overall accuracy on GQA and VQA2 is not sufficient to fully characterise the VQA performances of evaluated models. Indeed, our evaluation in OOD settings is the only one to reveal that even SOTA models struggle on infrequent question-answer pairs. Thus, the best-performing model MCAN (reaching 56.3% on GQA) looses about 10 points in the OOD setting. Finally, the error distribution measure defined in GQA is hardly interpretable in terms of VQA performances comparing to *acc-tail*, which makes it difficult to exploit.

Comparison with VQA-CP2 [10] — When comparing *acc-tail* and VQA-CP2 overall accuracy, we observe that the four VQA models obtain similar scores. However, VQA-CP2 [10] penalises more the MCAN [22] architecture. We note a completely different behavior when comparing scores obtained with bias-reduction methods. Indeed, while we observe that these methods do not improve the scores in the OOD setting (*cf.* Section 4.1), they achieve strong performances on VQA-CP2 [10]. In this context, the score of LMH [4] is the most notable. Whereas it achieves the highest overall accuracy on VQA-CP2 (52.0%), it obtains the lowest *acc-tail* on *GQA-OOD* (33%). In a less extent, we observe the same behaviour for RUBi [3] and BP [4].

Summarizing, VQA-CP2 [10] allows to measure to what extent a VQA model is relying on question types biases. It demonstrates that VQA systems experience difficulties to generalise to unseen distribution which illustrates their sensitivity to biases. However, the VQA-CP2 evaluation does not reflect the model behaviour on rare question-answer pairs. Moreover, evaluating on VQA-CP2 [10] requires to train on its specific training set. Not only this is restrictive, but, more importantly, it results in the emergence of models specifically conceived to address this dataset. Finally, unlike our *GQA-OOD*, VQA-CP2 does not have a dedicated validation split, forcing the researchers to validate their models directly on the test split [3, 4, 11, 12, 13] (which obviously has a negative impact on the objectiveness of such evaluation).

Conclusion

Going beyond previous attempts to reduce the influence of dataset biases in VQA evaluation, the proposed *GQA-OOD* benchmark allows to *both* evaluate (1) whether VQA models have absorbed the overall tendencies of the training data, and (2) how well they generalise to rare (or unseen) question-answer pairs. This was made possible (i) by a thorough choice of the most imbalanced question groups, (ii) by designing a new set of the evaluation metrics and finally, (iii) by leaving the researchers with a possibility of controlling the amount of distribution shift via the hyper-parameter α . Our experiments have shown that neither the conventional SOTA VQA models nor the dedicated bias reduction methods succeed in all aspects of the proposed evaluation benchmark. This leaves a big room of improvement for future work,

and we hope that our *GQA-OOD* benchmark will contribute to emergence of new VQA models which are less prone to learning spurious biases and thus more reliable in real-world scenarios.

References

- [1] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6700–6709, 2019.
- [2] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6904–6913, 2017.
- [3] Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. Rubi: Reducing unimodal biases for visual question answering. In *Advances in Neural Information Processing Systems*, pages 839–850, 2019.
- [4] Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. Don’t take the easy way out: Ensemble based methods for avoiding known dataset biases. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4060–4073, 2019.
- [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.
- [6] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2901–2910, 2017.
- [7] Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. Analyzing the behavior of visual question answering models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1955–1960, 2016.
- [8] Abhishek Das, Harsh Agrawal, C. Lawrence Zitnick, Devi Parikh, and Dhruv Batra. Human Attention in Visual Question Answering: Do Humans and Deep Networks Look at the Same Regions? In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
- [9] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *European Conference on Computer Vision*, pages 793–811. Springer, 2018.
- [10] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. Don’t just assume; look and answer: Overcoming priors for visual question answering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [11] Sainandan Ramakrishnan, Aishwarya Agrawal, and Stefan Lee. Overcoming language priors in visual question answering with adversarial regularization. In *Advances in Neural Information Processing Systems*, pages 1541–1551, 2018.
- [12] Jialin Wu and Raymond Mooney. Self-critical reasoning for robust visual question answering. In *Advances in Neural Information Processing Systems*, pages 8601–8611, 2019.
- [13] Ramprasaath R Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Shalini Ghosh, Larry Heck, Dhruv Batra, and Devi Parikh. Taking a hint: Leveraging explanations to make vision and language models more grounded. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2591–2600, 2019.
- [14] Mateusz Malinowski and Mario 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.
- [15] Drew Hudson and Christopher D Manning. Learning by abstraction: The neural state machine. In *Advances in Neural Information Processing Systems*, pages 5901–5914, 2019.
- [16] Dzmitry Bahdanau, AI Element, Harm de Vries, Shikhar Murty, Philippe Beaudoin, Yoshua Bengio, and Aaron Courville. Closure: Assessing systematic generalization of clevr models.
- [17] Damien Teney, Ehsan Abbasnejad, and Anton van den Hengel. Unshuffling data for improved generalization. 2020.

- [18] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [19] Jeffrey Pennington, Richard Socher, and Christopher D 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.
- [20] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [21] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018.
- [22] Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6281–6290, 2019.
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 2014.

A Dataset statistics

We provide some analysis and statistics to assess the reliability of the proposed benchmark. In particular, we analyse the nature and the distribution of the questions involved in *GQA-OOD* and demonstrate that it preserves the original question diversity of GQA [1].

Question diversity — Figure 6 and Figure 7 show the distribution of question structure type as defined in GQA [1] on the validation split. As one can observe, the process implemented to construct *GQA-OOD* does not alter the question diversity of the original split. However, the proportion of open questions – ‘query’ in Figure 4 and Figure 5 – has increased in *GQA-OOD*. Indeed, open questions – such as color questions – generally accept a wider diversity of answer, therefore it is prone to be more imbalanced. At contrary, other types such as ‘choose’, ‘verify’ or ‘compare’ usually accept only two possible answers and are easier to balance. Figure 4 and Figure 5 details the distribution of the structure types in the validation in *GQA-OOD* compared to GQA.

B Training details

Training hyper-parameters — All models evaluated on GQA and *GQA* have been trained on the balanced training set of GQA. For MCAN and BUTD we use publicly available implementations at <https://github.com/MILVLG/openvqa>. LSTM, BUTD [21], RUBi [3], BP [4] and LMH [4] are trained during 20 epochs with a batch size equals to 512 and Adam [23] optimizer. At the beginning of the training we linearly increase the learning rate from $2e^{-3}$ to $2e^{-1}$ during 3 epochs, followed by a decay by a factor of 0.2 at epochs 10 and 12. MCAN [22] is trained during 11 epoch with a batch size equals to 64 and Adamax [23] optimizer. At the beginning of the training we linearly increase the learning rate from $1e^{-4}$ to $2e^{-1}$ during 3 epochs, followed by a decay by a factor of 0.2 at epochs 10 and 12.

LMH hyper-parameters — Figure 3 details the hyper-parameter search for the entropy penalty weight in LMH [4]. We found that the entropy penalty was degrading the *GQA-OOD* accuracy when training on GQA. In particular, the flattening of the right side of the curve (most frequent samples) is even more present for higher penalty weight.




Figure 3: Influence of the LMH entropy penalty weight on the prediction error distribution.




Figure 4: Distribution of the semantic types in GQA.




Figure 5: Distribution of the semantic types in *GQA-OOD* (tail).




Figure 6: Distribution of the structural types in GQA.




Figure 7: Distribution of the semantic types in *GQA-OOD* (tail).