

Roses are Red, Violets are Blue... But Should VQA expect Them To?

Corentin Kervadec^{1,2} Grigory Antipov¹ Moez Baccouche¹ Christian Wolf²
¹Orange, Cesson-Sévigné, France. ²LIRIS, INSA-Lyon, UMR CNRS 5205, France

corentinkervadec.github.io {grigory.antipov, moez.baccouche}@orange.com christian.wolf@insa-lyon.fr

Abstract

Models for Visual Question Answering (VQA) are notorious for their tendency to rely on dataset biases, as the large and unbalanced diversity of questions and concepts involved tends to prevent models from learning to “reason”, leading them to perform “educated guesses” instead. In this paper, we claim that the standard evaluation metric, which consists in measuring the overall in-domain accuracy, is misleading. Since questions and concepts are unbalanced, this tends to favor models which exploit subtle training set statistics. Alternatively, naively introducing artificial distribution shifts between train and test splits is also not completely satisfying. First, the shifts do not reflect real-world tendencies, resulting in unsuitable models; second, since the shifts are handcrafted, trained models are specifically designed for this particular setting, and do not generalize to other configurations. We propose the GQA-OOD benchmark designed to overcome these concerns: we measure and compare accuracy over both rare and frequent question-answer pairs, and argue that the former is better suited to the evaluation of reasoning abilities, which we experimentally validate with models trained to more or less exploit biases. In a large-scale study involving 7 VQA models and 3 bias reduction techniques, we also experimentally demonstrate that these models fail to address questions involving infrequent concepts and provide recommendations for future directions of research.

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 problem is notorious for its 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

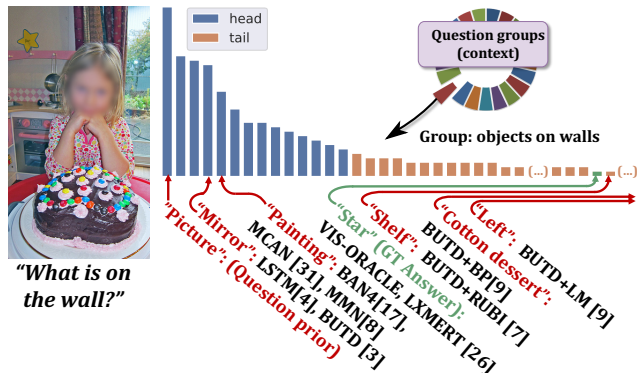


Figure 1. We address bias exploitation in VQA and propose a new benchmark for Out-Of-Distribution evaluation containing distribution shifts tailored to different question groups with highly imbalanced distributions. A new evaluation metric based on rareness inside each question group, here shown for “objects on walls”, is experimentally demonstrated to be less prone to bias exploitation. We show that SOTA methods (7 VQA models and 3 bias reduction methods) reproduce biases in training data.

to learning in the form of “Clever Hans” effects. Unfortunately, and in spite of recent efforts [11, 15], 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 [7, 9]. 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 of 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 al-

lows a better diagnostic of current VQA performance. In particular, our benchmark can be viewed as an alternative to the VQA-CP [2] dataset, which has lead to mixed results (see [25] for a recent comprehensive study). Our benchmark comprises (i) a new fine-grained reorganization of the GQA dataset [15] 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 generalisation 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.

Our contributions — (i) 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. (ii) Compared to competing benchmarks, our dataset features distribution shifts for both, validation and test, allowing to validate models under OOD conditions. (iii) We experimentally evaluate the usefulness of the proposed metric showing its behavior on models trained to, more or less, exploit biases. (iv) 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 datasets was VQA1 [4] with ~76K questions over 25K realistic images. It started a new task, but was soon found to suffer from biases. [11] pointed to strong imbalances among the presented answers and proposed the second (improved) version: VQA2. [16] introduced the fully synthetic CLEVR dataset, designed to evaluate reasoning capabilities. Its strong point is its detailed and structured annotation. In [15], CLEVR is adapted to real-world images resulting in the automatically created GQA dataset (1.7M questions), offering a better control on dataset statistics.

Attempts to reduce the bias-dependency — Despite efforts to design complex architectures, VQA models suffer from significant generalization incapacities [1]. They tend to answer questions without using the image, and even when they do, they do not always exploit relevant visual regions [10]. They tend to overly rely on dataset biases [12],




Figure 2. We re-organize the GQA dataset [15] in a fine-grained way: the benchmark contains a distribution shift in validation and test, allowing to validate *and* evaluate in OOD settings.

and are not able to generalize to unseen distributions [2]. Methods to solve this include *e.g.* setting up an adversarial game against a question-only adversary to regularize training [21], and a question-only branch in addition to a base model during training to prevent it from learning textual biases [7] (RUBi). At test time, this branch is omitted. [9] regularize model predictions using question type statistics from the training set. Other attempts force VQA models to attend to the most significant visual regions from humans’ perspective [28, 22]. While they show promising results when evaluated on unseen distributions [2], they generally slightly degrade performances on standard benchmarks, which tend to favor models relying on dataset biases.

Reinventing VQA evaluation — Propositions of new evaluation methods for VQA went hand in hand with the design of new models. Early work [19] proposed a soft evaluation score based on a lexical database, then replaced by a hard classification score, more prone to favor biased predictions, but easier to use in practice. The authors of GQA [15] proposed several additional metrics, namely: consistency, plausibility, validity, distribution, *etc.* They provide insights on VQA performance, 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 generalization capability of Neural State Machines, [14] proposed interesting splits of GQA, reorganizing the train and validation splits to distinctly evaluate how well the models generalize on visual content and on the linguistic structure of the question. The authors of [5] analyzed generalization 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 generalize in these conditions. This confirms the need of carefully designed benchmarks to evaluate the true capabilities of VQA models. In this context, [2] introduced a reorganisation of the VQA2 [11] dataset splits, namely VQA-CP2, where the training distribution is made explicitly different from the one in the test. We discuss the differences between VQA-CP2 and our benchmark in detail in Section 3.

¹<https://github.com/gqa-ood/GQA-OOD>

Shortcomings of Out-of-distribution testing Recent works have raised some criticisms about OOD evaluation protocols. In particular, [25] points out several pitfalls observed when evaluating VQA in OOD setting using VQA-CP [2]: (1) this protocol does not allow to test a model in both out- and in-distribution settings without having to retrain it on a different set of data; (2) several works rely on known construction procedures of the OOD test split, e.g. the test answers’ distribution is the inverse of the train one, coming up with methods specialized to a very particular setup; and (3) as no validation split is provided, the hyper-parameters are usually selected on the test split. Although these wrong practices are in part to the responsibility of the model designers, they show us the need to rethink the evaluation protocol. In parallel, [23] come up with an interesting negative result while analysing visual grounding bias-reduction methods designed on top of the VQA-CP dataset. These methods [28, 22], consisting in supervising a VQA model to attend to visual regions which are relevant to a human considering the question, are very efficient on VQA-CP. Surprisingly, [23] found that simply enforcing the model to attend to random visual regions was at least as much efficient on out- and in-distribution settings. Why was this negative result not observed before? We think that a more profound empirical evaluation of models’ behavior would help to better judge and compare the efficiency of new VQA methods.

These works unveil the need of a better evaluation protocol, allowing to test in both out- and in-distribution settings without falling into ood testing pitfalls. In this paper, we try to provide this evaluation protocol. We complement it with a large experimental study of the generalisation abilities of VQA architectures and de-bias methods in order to better understand their pros and cons.

3. GQA-OOD: a benchmark for OOD settings

We introduce a new VQA benchmark named *GQA-OOD* designed to evaluate models and algorithms in 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. These rare events might involve concepts which are also present in the training set. Let’s for instance 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, infrequent (correct) test answers would be, for instance, *blue*, requiring models to reason to provide the correct answer. We design a benchmark where this shift is not global but 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 exist-

Dataset	Split	#Quest.	#Groups	#Imgs
GQA-OOD	val	51,045	3,849	9,406
	testdev	2,796	471	388
GQA	val	132,062	36,832	10,234
	testdev	12,578	7,803	398

(a)

Split	Subset	#Quest.	#Groups	#Imgs
GQA-OOD val	head	33,882	3,849	8,664
	tail	17,163	3,849	6,632
GQA-OOD testdev	head	1,733	471	365
	tail	1,063	471	330

(b)

Table 1. Data statistics: (a) GQA-OOD vs. GQA; (b) head vs. tail

ing GQA [15] dataset², which provides more fine-grained annotations compared to competing VQA2 [11] (the questions in GQA have been automatically generated from scene graphs, which allows better control of the context). Figure 2 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 process further described below. 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 distribution shifts, 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 concepts related to the question, e.g. ‘zebra’, ‘violet’, etc. There is a total of $\approx 37K$ local groups related to $\approx 132K$ 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 measure 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 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

²We use version 1.2 of the GQA dataset [15].

close the distribution $p(x)$ is to a uniform distribution of the same dimension. Finally, we keep groups with a normalised entropy smaller than a threshold empirically set to $T=0.9$. This selects all benchmark’s questions, but further work is 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. Figure 1 illustrates how the subsets are selected using the concrete example answer histogram of question group *objects on walls*.

- **Acc-tail:** the accuracy on OOD samples, which are the samples of the tail of the answer class distribution, *i.e.* the rarest answers given the context. We define the tail classes as classes i with $|a_i| \leq \alpha\mu(a)$, 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 $\alpha=1.2$, and in Section 4.2 we analyze and illustrate the impact of the choice of α on *Acc-tail*. Figure 1 provides an example of such a tail question — we can see that the answer *Star* is rare in this group, therefore it belongs to the tail set like the other answers shown in orange.
- **Acc-head:** the accuracy on the distribution head for each local group, given as difference between the whole group and its tail (blue answers in Figure 1).
- **Acc-all:** the overall (classical) accuracy over all *GQA-ODD* samples, *i.e.* the in-domain accuracy. In Figure 1, this corresponds to the blue and orange answers.

Table 1 provides statistics of the proposed benchmark. We also analyzed the nature, distribution and diversity of the questions w.r.t to GQA [15], *c.f.* Appendix.

Difference with VQA-CP2 — The VQA-CP2 [2] 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, questions are grouped according to their first words and the ground-truth answer. The shift is created by splitting according to types. In contrast, our proposed *GQA-ODD* dataset allows fine-grained analysis of the generalization 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*) in out- and in-distribution settings, and (iii) even through the possibility of continuous evaluation along different operating points (see Figure 3).

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 [7, 9, 21, 28, 22], which should be frowned upon, or, alternatively, validate on a subset of train which does not include a shift [26]. Our *GQA-ODD* 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, our proposed dataset requires models to be trained on the existing GQA train split. This forces 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.

Limitations — The proposed benchmark is built on the GQA dataset, whose questions have been automatically generated, resulting in a limited vocabulary (GQA covers 70% of VQA2 answers [15]) and a synthetic syntax. While the images are natural and real, one might argue, that the questions are not “*in the wild*”. However, the benefits of the synthetic nature of the questions largely out-weigh its limitations. In particular, this offers a better control on the data and excludes unmodelled external knowledge, which leads to a better evaluation of reasoning abilities. We made the source code publicly available, and we encourage the field to use it to study robustness in OOD settings.

4. Experiments

In our experiments we used several SOTA VQA models and we compared the proposed *GQA-ODD* benchmark to the standard benchmarks VQA2 [11], GQA [15] and VQA-CP2 [2]. The line-up includes strong models with object-level attention and two Transformer [27]-based model, as well as two blind baseline models and a visual oracle.

Question Prior — this blind baseline returns the most probable answer estimated from training set statistics. Following [2, 11], we use the question types priors when evaluating on VQA-CP and VQA2. For GQA, we use the training global group priors. **LSTM** [4] — this blind baseline takes GloVe embeddings [20] and encodes them using an LSTM [13] followed by a feed-forward network. Neither of the two blind baselines use the input image. **BUTD** [3] — 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. **BAN4** [18] — another object-level attention based model using a bilinear fusion mechanism to models relations between words and objects with multi-hop reasoning. We use the 4-layers version of BAN. **MCAN** [29] — a Transformer [27] architecture 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. **LXMERT** [24] — a SOTA approach combining a Transformer architecture (close to the




Figure 3. Performance (higher is better) (a) and head/tail confusion (lower is better) (b) for different definitions of the tail distribution (α parameter values) on the GQA-OOD benchmark. We compare several VQA models. The x-axis is in log-scale.

one used in MCAN) with a large scale pre-training procedure. LXMERT is pretrained on roughly 9 millions question-image pairs with semi-supervised tasks such as reconstructing masked words and visual objects. It is then finetuned on the GQA dataset. **MMN** [8] — a Meta Module Network for compositional visual reasoning. It is based on a hybrid approach combining neural module networks (NMN) and monolithic architectures (such as MCAN). The former, NMN, is based on hand-crafted neural network program blocs and is supposed to allow a better compositionality and interpretability. The latter, monolithic architecture, performs its operations in a latent space and has been shown to be experimentally more efficient. NMN tries to combine the best of both worlds. **VIS-ORACLE** — a model with a perfect sight, *i.e.* taking as input the question and a set of ground truth objects directly taken from the annotation of GQA³. It allows to evaluate the performance of a model without the imperfection of the visual extractor. It is based on a Transformer architecture similar to the one used for LXMERT with smaller capacity.

For training details of all models we refer to the Appendix. Models evaluated on *GQA-OOD* were trained on the training set of GQA (balanced) and validated on the validation split of *GQA-OOD*, unless otherwise stated. When available, we provide the standard deviation computed over at least four different seeds.

4.1. Evaluation of the proposed metric

We believe that a good evaluation metric satisfies at least two properties: it is easy to interpret, and it provides an estimate for the quality targeted by the evaluation. We argued above on the merits of our proposed tail accuracy (Acc-tail) as a way of estimating VQA performance less influenced by bias. In what follows, we complete this by an experimental

³As GT annotations (scene-graphs) are only available for the train and validation split, we do not evaluate VIS-ORACLE on the testdev split.

Model	Baseline benchm.	Proposed benchmark (Acc-tail)		
	Tot. Acc.	$\alpha=1.2$	$\alpha=0.5$	$\alpha=0.3$
BUTD [3] + <i>bal</i>	60.7 \pm 0.4	45.4 \pm 0.3	33.8 \pm 0.5	24.6 \pm 0.5
BUTD [3] + <i>all</i>	59.8 \pm 0.1	41.9 \pm 0.1	29.5 \pm 0.3	18.3 \pm 0.6
Δ (relative):	-1.4%	-7.7%	-12.9%	-25.7%

Table 2. We compare two different VQA models based on BUTD [3], one of which has been trained on a split known to be biased (BUTD [3]+*all*), and evaluate the proposed metric’s capacity to detect this bias. All scores in % on the GQA-OOD val split.

validation of the metric. To this end, we compared two different VQA models, one of which has been trained in a way known to be biased. In particular, we trained BUTD [3], known to capture training set biases [2], on the GQA and GQA-OOD validation splits. The first version, BUTD+*bal*, is trained on the widely used balanced training set of GQA, which we had also used for all other experiments in this paper. This training set had been created by smoothing the question distribution in order to mitigate dataset biases [15]. The second one variant, BUTD+*all*, is trained on the raw and unbalanced GQA training set, which leads to more spurious biases than the balanced version. As the unbalanced set is ten times bigger than the balanced one, we split it in ten subsets and provide the average score.

Results are given in Table 2, comparing two different metrics, namely the classical total accuracy and our GQA-OOD *acc-tail* metric, with three different values for the α hyper-parameter. First, we observe that the two versions of BUTD obtain similar scores on GQA overall — the relative difference is only 1.4%. This is not a surprise, the classical metric is influenced by biases. As expected, the two VQA models behave differently on our proposed *acc-tail* metric: the model trained on the unbalanced training set is outperformed by the balanced one by a large margin. Moreover, the score difference increases with decreasing α , (*i.e.* when




Figure 4. Acc-tail performance and head/tail confusion, as in Fig. 3), but for different bias-reduction methods on top of BUTD [3].

the metric focuses on the rarer and rarer question-answer pairs, providing valuable evidence that *acc-tail* is indeed well suited for measuring VQA performance undisturbed by bias dependencies.

4.2. Analysis of VQA model error distributions

The *GQA-OOD* benchmark allows us to perform an analysis of the error prediction distributions for various VQA models as shown in Table 3. 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 — We can see that that VQA models (dramatically) fail to generalize on infrequent association of concepts. The two blind models (Question Prior and LSTM in Table 3-a) obtain the highest gap between *acc-tail* and *acc-head*, explained by the fact that they uniquely rely on question biases. The Δ score indicates that BUTD, MMN, MCAN, BAN4 and LXMERT also struggle (in a lesser extent) to generalize on the less frequent question-answer pairs. Nevertheless, we observe that the Transformer-based architecture combined with large-scale BERT training, LXMERT, outperforms all models on the *acc-tail* metric, confirming its superiority. This is corroborated by [12], who show that pretrained Transformers improve OOD robustness in NLP.

In contrast to our proposed *acc-tail* metric, the metric *acc-all*, close to the standard VQA metric, does not reflect the true model’s performances, since it is mechanically increased by the high scores obtained on the most frequent question-answers. This confirms the need of a two-in-one evaluation: measuring the out- and in-distribution performance scores, as we propose.

Visualising the generalisation behavior — The definition of what constitutes a “rare” answer, i.e. the size of the tail, depends on the parameter α . In Figure 3-a, we analyze how VQA model prediction errors (*acc-tail*) depend on this definition, i.e. how models behave w.r.t. to questions whose answers are more and more rare. Increasing α increases the tail — in the extreme case it is equal to the whole distribu-

tion (right side of the plot). With small α , only the most infrequent question-answer pairs are evaluated (left side of the plot). All models follow the same dynamic: starting from a tail size which represents roughly half of the question-answer pairs, tail accuracy starts to linearly decrease until reaching a dramatically low score (about 30 pts lower than the overall accuracy). An exception is VIS-ORACLE: its dynamics is nearly flat, prediction error is almost decorrelated from answer rareness. This provides evidence that a model using perfect visual input is able to learn reasoning with significantly decreased dependency on dataset biases.

We complement this analysis by measuring the confusion between *head* and *tail* as a function of α , shown in Figure 3-b, which provides insights on the causes of the generalisation failure observed in Figure 3-a. This confusion corresponds to the proportion of questions where the model predicts a *head* answer with a *tail* GT answer⁴. For $\alpha=1.2$, LXMERT confuses answers for 25% of questions, which increases up to 42% for $\alpha=0.3$. Similar behavior is observed for the other models, but interestingly *not* for VIS-ORACLE, where curve is nearly flat, again providing evidence for a low dependency on statistical biases in the training set. As a side note, we found in [17] that initialising LXMERT weights with VIS-ORACLE allows to boost the accuracy on *acc-tail*.

Exploiting biases vs. “reasoning” — It is difficult to assess, whether a model reasons or not, in particular since the term “reasoning” has various different definitions. Referring to [6], “*a plausible definition of ‘reasoning’ could be ‘algebraically manipulating previously acquired knowledge in order to answer a new question’*”. In the context of VQA, we could interpret this as “*algebraically manipulating words and visual objects to answer a new question*”. With this definition, using statistics biases cannot be considered reasoning, but should rather be denoted as “educated guesses” [15] or *biased* answers. Using the proposed *GQA-OOD* benchmark, we explore the estimation of three reasoning labels qualifying the mode of operation a model uses for a given input: *bias*, *reason* and *other/unknown*. In absence of GT information, we propose to estimate these labels from proxy rules: a VQA model is estimated to *reason*, when it correctly predicts an answer, which is rare in GT and rare in prediction; it is considered *biased*, when it wrongly predicts an answer, which is *frequent* in its prediction and *rare* in GT.

Figure 5-a shows the calculation of these labels based on the distribution of the *head* and *tail* labels of each answer in the predictions (rows) and GT (columns) for LXMERT [24] on the validation split of *GQA-OOD*. We add a *borderline* label representing the fuzzy frontier between reasoning and bias exploitation⁵. In Figure 5-b, we show the distribution

⁴More information is available in the appendix.

⁵*head*: $\alpha > 1.2$, *borderline*: $0.7 < \alpha < 1.2$, *tail*: $\alpha < 0.7$.

Model	Uses image	acc-all	acc-tail	acc-head	Δ
Quest. Prior	\times	21.6	17.8	24.1	35.4
LSTM [4]	\times	30.7	24.0	34.8	45.0
BUTD [3]	\checkmark	46.4 \pm 1.1	42.1 \pm 0.9	49.1 \pm 1.1	16.6
MCAN [29]	\checkmark	50.8 \pm 0.4	46.5 \pm 0.5	53.4 \pm 0.6	14.8
BAN4 [18]	\checkmark	50.2 \pm 0.7	47.2 \pm 0.5	51.9 \pm 1.0	9.9
MMN [8]	\checkmark	52.7	48.0	55.5	15.6
LXMERT [24]	\checkmark	54.6	49.8	57.7	15.9

(a)

Technique	acc-all	acc-tail	acc-head	Δ
BUTD [3]	46.4 \pm 1.1	42.1\pm0.9	49.1 \pm 1.1	16.6
+RUBi+QB	46.7\pm1.3	42.1\pm1.0	49.4\pm1.5	17.3
+RUBi [7]	38.8 \pm 2.4	35.7 \pm 2.3	40.8 \pm 2.7	14.3
+LM [9]	34.5 \pm 0.7	32.2 \pm 1.2	35.9 \pm 1.2	11.5
+BP [9]	33.1 \pm 0.4	30.8 \pm 1.0	34.5 \pm 0.5	12.0

(b)

Table 3. Comparison of several VQA methods on the GQA-OOD testdev split. *Acc-tail*: OOD settings, *Acc-head*: accuracy on most probable answers (given context), scores in %. (a) different models; (b) BUTD [3] model combined with different bias reduction techniques.

	Head		Borderline		Tail	
	C	W	C	W	C	W
Head	30.0%	9.6%	0.0%	3.1%	0.0%	5.3%
Borderline	0.0%	5.5%	6.3%	2.1%	0.0%	3.2%
Tail	0.0%	5.0%	0.0%	0.8%	11.6%	1.3%

C=Correct, W=Wrong


Rows=predicted labels, columns=GT labels

Blue=Model is estimated to reason

Orange=Model is estimated to exploit bias

Green=Unknown label

(a)



(b)

Figure 5. We estimate “reasoning labels”: the model is estimated to *reason*, when it correctly predicts an answer rare in GT and rare in prediction; it is considered *biased*, when it wrongly predicts an answer, which is *frequent* in its prediction and *rare* in GT. All values are computed over the GQA-OOD validation split. The matrix (a) shows the joint distribution of predicted and GT classes. (b): Distribution the estimated reasoning labels over the GQA [15] question types for the LXMERT [24] model. The model often predicts a biased answer on the *query* and *compare* questions while there is evidence that it may reason on *verify*, *choose* and *logical* questions.

of these reasoning labels over the different GQA structural question types [15]: *verify*, *choose*, *compare* and *query*. We observe that LXMERT seems to “reason” on the *verify*, *choose* and *logical* questions, which are binary questions, while *compare*⁶ and *query* questions are the most prone to bias exploitation. From this, we conclude that future efforts on improvements of model capacities to answer open questions (*e.g* typed as *query*) should be particular fruitful.

4.3. Re-evaluating bias-reduction methods

We use the proposed benchmark to re-evaluate several bias-reduction methods, which have been initially designed on the VQA-CP [2] dataset. As these methods were designed to be model-agnostic [7, 9], we use them together with the BUTD [3] architecture:

RUBi [7] — 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** [9] — is similar to RUBi but differs by directly

⁶only 1% of the tail questions are typed as *compare*.

taking training set statistics to infer question type biases during training⁷. The question type biases are fused with the base model predictions using a product of experts [9], and removed during testing. **LM** [9] — is an improved version of BP [9]. 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 appendix for details).

Surprisingly, none of the three bias-reduction methods succeed to improve *acc-tail* (*cf.* Table 3-b), they even deteriorate *acc-head*. This is unexpected as they have been designed to overcome the dependency on question type biases. For further analysis, we evaluate RUBi while keeping the question-only branch during testing (RUBi+QB). As expected, it outperforms RUBi on *acc-head*, indicating it has better captured frequent patterns. However, it also outperforms RUBi on the OOD settings, demonstrating that preventing from learning frequent patterns does not necessarily increase performances on rare samples.

We provide a visualization of the generalisation behav-

⁷VQA2: biases are over question types; GQA: local groups.

Model	VQA2	GQA		VQA-CP2	GQA-OOD
	overall	overall	dist.	overall	acc-tail
Q. Prior	32.1	27.0	55.6	8.8	17.8
LSTM [4]	43.0	39.1	3.6	22.1	24.0
BUTD [3]	63.5	51.6 \pm 0.3	1.8	40.1	42.1 \pm 0.9
MCAN [29]	66.1	56.3 \pm 0.2	1.6	42.5	46.5 \pm 0.5
BAN4 [18]	65.9	54.7 \pm 0.4	1.6	40.7	47.2 \pm 0.5
MMN [8]	-	59.6	1.8	-	48.0
LXMERT [24]	69.9	59.6	1.5	-	49.8
BUTD [3]	63.5	51.6 \pm 0.3	1.8	40.1	42.1 \pm 0.9
+RUBi+QB	-	51.9 \pm 1.1	1.7	47.6 \pm 3.7	42.1 \pm 1.0
+RUBi [7]	61.2	43.6 \pm 2.0	1.9	44.2	35.7 \pm 2.3
+LM [9]	56.4	39.7 \pm 0.7	2.1	52.0	32.2 \pm 1.2
+BP [9]	63.2	39.6 \pm 0.3	2.2	39.9	30.8 \pm 1.0

Table 4. We compare the proposed *acc-tail* metric with other benchmarks. Results computed on the testdev split of *GQA-OOD* and GQA [15], the test split of VQA-CP2 [2] and the VQA2 [11] validation split. Values in italic: trained and tested by ourselves.

ior on bias-reduction methods in Figure 4-a. For BP [9], LM [9] and, to a lesser extent, RUBi [7], we observe that the right side of the curve has flattened, indicating that overall accuracy, dominated by frequent question-answer pairs, has been reduced by bias-reduction. The left side of the curve, however, corresponding to rare samples, remains almost unchanged, revealing that these methods have somewhat succeeded in preventing the base model from learning dataset biases. As a comparison, the LSTM model in Figure 3-a performs worse than BUTD but conserves the same frequent/rare imbalance. We observe that RUBi+QB responds the same way as BUTD, confirming the effect of bias-reduction; looking at *head/tail* confusion in Figure 4-b, the result is even more pronounced. In short, we demonstrate the effectiveness of bias reduction methods in preventing the base model from learning salient properties of the training set, and occasionally reducing the dependency toward dataset biases. However, this does not necessarily help the model to learn the subtle distributions, required for generalization and for learning to reason.

4.4. Comparison with other benchmarks

We compare the proposed *GQA-OOD* benchmark with the following three standard VQA datasets:

GQA [15] (balanced version) is a dataset with $\sim 1.7M$ question-answer pairs automatically generated from real images. Among all VQA datasets, GQA has the richest annotations allowing to evaluate models with complementary metrics: consistency, validity, plausibility, [15] *etc.* Here, we only discuss overall accuracy and the distribution score on the GQA testdev split as the other metrics are unrelated to the topic of our paper. The distribution score is obtained by measuring the match between the true GT answer distribution and the predicted distribution. **VQA2 [11]** (265K images, ≥ 3 questions each) is composed of questions created by humans, each question is annotated with 10 GT an-

swers. We compare with overall accuracy on the VQA2 validation split. **VQA-CP2 [2]** has been constructed by reorganising the training and validation splits of VQA2 [11] aiming at differences in answer distributions between training and test splits. It has been designed to measure sensitivity to language bias.

Comparison with GQA and VQA2 — In Table 4.3, we compare our *acc-tail* score with the other benchmarks. We can see that overall accuracy on GQA and VQA2 is not sufficient to fully characterize the VQA performances. Our evaluation in OOD settings is the only one to reveal that even SOTA models struggle on infrequent question-answer pairs. The best-performing model LXMERT loses about 10 points in the OOD setting. Our metric also unveils that, despite performing on-par with LXMERT on GQA overall, MMN struggles more on infrequent question-answer pairs. Finally, we argue that *acc-tail* is easier to interpret than the error distribution measure defined in GQA.

Comparison with VQA-CP2 [2] — Comparing *acc-tail* to VQA-CP2 overall accuracy, we observe similar scores on standard VQA models, but a completely different behavior for bias-reduction methods. While they do not improve the scores in the OOD setting (*cf.* Section 4.2), they achieve strong performances on VQA-CP2. The score of LM stands out, achieving the highest overall accuracy on VQA-CP2 (52.0%) but one of the lowest *acc-tail* on *GQA-OOD* (33%), with similar behavior for RUBi and BP. In short, while VQA-CP2 measures to what extent a VQA model struggles to generalize to a specific unseen distributions, the VQA-CP2 evaluation does not reflect the model behaviour on rare question-answer pairs.

Discussion and conclusions

Going beyond previous attempts to reduce the influence of dataset biases in VQA evaluation, our proposed *GQA-OOD* benchmark allows to *both* evaluate (1) whether models have absorbed tendencies in the training data, and (2) how well they generalize to rare/unseen question-answer pairs. This was made possible by (i) a thorough choice of imbalanced question groups, (ii) a new set of metrics and finally, (iii) by allowing to control the amount of distribution shift via the hyper-parameter α . We have provided evidence that the benchmark and metric measure performance and dependency on dataset bias. Our experiments have also shown that neither conventional SOTA VQA models nor dedicated bias reduction methods succeed in all aspects of the proposed evaluation benchmark. We hope that this sheds light on the current shortcomings in vision and language reasoning, and we hope that *GQA-OOD* will contribute to the emergence of new models, less prone to learning spurious biases and more reliable in real-world scenarios.

Acknowledgements — C. Wolf acknowledges support from ANR through grant “Remember” (ANR-20-CHIA-0018).

References

- [1] 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. [2](#)
- [2] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Anirudha 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. [2](#), [3](#), [4](#), [5](#), [7](#), [8](#)
- [3] 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. [4](#), [5](#), [6](#), [7](#), [8](#)
- [4] 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. [2](#), [4](#), [7](#), [8](#)
- [5] Dzmitry Bahdanau, AI Element, Harm de Vries, Shikhar Murty, Philippe Beaudoin, Yoshua Bengio, and Aaron Courville. Closure: Assessing systematic generalization of clevr models. [2](#)
- [6] Léon Bottou. From machine learning to machine reasoning. *Machine learning*, 94(2):133–149, 2014. [6](#)
- [7] 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. [1](#), [2](#), [4](#), [7](#), [8](#)
- [8] Wenhu Chen, Zhe Gan, Linjie Li, Yu Cheng, William Wang, and Jingjing Liu. Meta module network for compositional visual reasoning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 655–664, 2021. [5](#), [7](#), [8](#)
- [9] 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. [1](#), [2](#), [4](#), [7](#), [8](#)
- [10] 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. [2](#)
- [11] 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. [1](#), [2](#), [3](#), [4](#), [8](#)
- [12] 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. [2](#), [6](#)
- [13] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. [4](#)
- [14] Drew Hudson and Christopher D Manning. Learning by abstraction: The neural state machine. In *Advances in Neural Information Processing Systems*, pages 5901–5914, 2019. [2](#)
- [15] 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. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [16] 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. [2](#)
- [17] Corentin Kervadec, Theo Jaunet, Grigory Antipov, Moez Baccouche, Romain Vuillemot, and Christian Wolf. How transferable are reasoning patterns in vqa? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. [6](#)
- [18] Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. Bilinear attention networks. In *Advances in Neural Information Processing Systems*, pages 1564–1574, 2018. [4](#), [7](#), [8](#)
- [19] 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. [2](#)
- [20] 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. [4](#)
- [21] 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. [2](#), [4](#)
- [22] 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. [2](#), [3](#), [4](#)
- [23] Robik Shrestha, Kushal Kafle, and Christopher Kanan. A negative case analysis of visual grounding methods for vqa. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8172–8181, 2020. [3](#)
- [24] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5103–5114, 2019. [4](#), [6](#), [7](#), [8](#)
- [25] Damien Teney, Ehsan Abbasnejad, Kushal Kafle, Robik Shrestha, Christopher Kanan, and Anton van den Hengel.

- On the value of out-of-distribution testing: An example of goodharts law. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 407–417. Curran Associates, Inc., 2020. [2](#), [3](#)
- [26] Damien Teney, Ehsan Abbasnejad, and Anton van den Hengel. Unshuffling data for improved generalization. 2020. [4](#)
- [27] 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. [4](#)
- [28] Jialin Wu and Raymond Mooney. Self-critical reasoning for robust visual question answering. In *Advances in Neural Information Processing Systems*, pages 8601–8611, 2019. [2](#), [3](#), [4](#)
- [29] 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. [4](#), [7](#), [8](#)