Reducing Language Biases in Visual Question Answering with Visually-Grounded Question Encoder

Gouthaman KV and Anurag Mittal

Indian Institute of Technology Madras, India
 {gkv,amittal}@cse.iitm.ac.in

Abstract. Recent studies have shown that current VQA models are heavily biased on the language priors in the train set to answer the question, irrespective of the image. E.g., overwhelmingly answer "what sport is" as "tennis" or "what color banana" as "yellow." This behavior restricts them from real-world application scenarios. In this work, we propose a novel model-agnostic question encoder, Visually-Grounded Question Encoder (VGQE), for VQA that reduces this effect. VGQE utilizes both visual and language modalities equally while encoding the question. Hence the question representation itself gets sufficient visual-grounding, and thus reduces the dependency of the model on the language priors. We demonstrate the effect of VGQE on three recent VQA models and achieve state-of-the-art results on the bias-sensitive split of the VQAv2 dataset; VQA-CPv2. Further, unlike the existing bias-reduction techniques, on the standard VQAv2 benchmark, our approach does not drop the accuracy; instead, it improves the performance.

Keywords: Deep-learning, Visual Question Answering, Language bias

1 Introduction

Visual Question Answering (VQA) is a good benchmark for context-specific reasoning and scene understanding that requires the combined skill of Computer Vision (CV) and Natural Language Processing (NLP). Given an image and a question in natural language, the task is to answer the question by understanding cues from both the question and the image. Tackling the VQA problem requires a variety of scene understanding capabilities such as object and activity recognition, enumerating objects, knowledge-based reasoning, fine-grained recognition, and common sense reasoning, etc. Thus, such a multi-domain problem yields a good measure of whether computers are reaching capabilities similar to humans.

With the success of deep-learning in CV and NLP, several datasets [3, 14, 21, 15, 19, 42], and models [2, 4, 6, 5, 22] have been proposed to solve VQA. Most of these models perform well in the existing benchmark datasets where the train and test sets have similar answer distribution. However, recent studies show that these models often rely on the statistical correlations between the question

patterns & the most frequent answer in the train set, and shows poor generalized performance (i.e., poor performance on a test set with different answer distributions than the train set [1]). In other words, they are heavily biased on the language modality. E.g., most of the existing models overwhelmingly answer "tennis" by seeing the question pattern "what sport is.." or answer the question "what color of the banana?" as "yellow" even though it is a "green banana," By definition and design, the VQA models need to merge the visual and textual modalities to answer the question. However, in practice, they often answer the question without considering the image modality, i.e., they tend to have less image grounding and rely more on the question. This undesirable behavior restricts the existing VQA models being applicable in practical scenarios.

One reason for this behavior is the strong language biases that exist in the available VQA datasets [3, 14]. Hence, the models trained on these datasets will often rely on such biases [1, 32, 7, 33, 39, 9]. E.g., in the most popular and large VQAv2 dataset [14], a majority of the "what sport is" question is linked to "tennis." Hence, upon training, the model will blindly learn the superficial correlations between the "what sport is" language pattern from the question, and the most frequent answer "tennis." Unfortunately, it is hard to avoid these biases from the train set when collecting real-world samples due to the annotation cost. Instead, we require methods that force the model to look at the image and question equally to predict the answer. Since the main source of bias is the over-dependency of the models on the language side, some approaches such as [32, 7, 9], tried to overcome this by reducing the contribution from the language side. On the other hand, some approaches [33, 1, 39] tried to improve the visual-grounding of the model in order to reduce the language-bias. All of these existing bias reduction methods cause a performance reduction in the standard benchmark VQA datasets.

A common pipeline in existing VQA models is to encode the image and question separately and then fuse them to predict the answer. Typically the image is encoded using a pre-trained CNN (e.g., ResNet [16], VGGNet [35], etc.), and the question is encoded by sending the word-level features to an RNN (GRU [10] or LSTM [18]). In this approach, while encoding the question, it only considers the language modality, and such a question representation has not contained any distinguishable power based on the content in the image. In other words, the question representation is not grounded in the image. E.g., with this scheme, the question "what sport is this?" has the same encoded representation irrespective of whether the image contains "tennis," "baseball," or "skateboarding," etc. Since the majority of such questions are linked to the answer "tennis" in the train set, the model will learn a strong correlation between the question pattern "what sport is" to the answer "tennis," irrespective of the image. This leads to overfitting the question representation to the most frequent answer, and the model will ignore the image altogether.

In this paper, we propose a generic question encoder for VQA, the Visually-Grounded Question Encoder (VGQE), that reduces this problem. VGQE encodes the question by considering not only the linguistic information from the



(a) A generic VQA model with a traditional question encoder.

(b) A generic VQA model with VGQE.

Fig. 1. A VQA model with a traditional question encoder vs. with VGQE: A traditional question encoder uses only the language modality while encoding the question, whereas VGQE uses both language and visual modalities to encode the question.

question but also the visual information from the image. A visual comparison of a generic VQA model with a traditional question encoder and with VGQE is shown in Fig 1. The VGQE explicitly forces the model to look at the image while encoding the question. For each question word, VGQE finds the important visual feature from the image and generates a visually-grounded word-embedding vector. This word embedding vector contains the language information from the question word and the corresponding visual information from the image. These visually-grounded word-embedding vectors are then passed to a sequence of RNN cells (inside VGQE) to encode the question. Since VGQE considers both the modalities equally, the encoded question representation will get sufficient distinguishing power based on the visual counterpart. For e.g., with VGQE, in the case of the "what sport is this" question, the model can easily distinguish the question for "baseball," "tennis," or "skateboarding," etc. As a result, the question representation itself is heavily influenced by the image, and the learning of the correlation between specific language patterns and the most frequent answers in the train set can be reduced.

The VGQE is generic and easily adaptable to any existing VQA models, i.e., replace the existing traditional language-based question encoder in the model with VGQE. In this paper, we demonstrate the ability of VGQE to reduce the language bias on three recent best performing baseline VQA models, i.e., MUREL [6], UpDn [2] and BAN [22]. We did extensive experiments on the VQA-CPv2 dataset [1] and demonstrate the ability of VGQE to push the baseline models to achieve state-of-the-art results. The VQA-CPv2 dataset is specifically designed to test the VQA model's capacity to overcome the language biases from the train set. Further, we show the effect of VGQE in the standard VQAv2 benchmark as well. Unlike existing bias reduction techniques [1, 32, 7, 9, 39], VGQE does not show any drop in the accuracy on the standard VQAv2 dataset; instead, it improves the performance.

2 Related Works

A major role in the success of deep-learning models lies in the availability of large datasets. Most of the existing real-world datasets for various problems will

have some form of biases, and it is hard to avoid this at the time of dataset collection, due to the annotation cost. Consequently, the models trained on these datasets often over-fit to the biases between the inputs and ground truth annotations from the train set. Researchers tried different procedures to mitigate this problem in various domains. For instance, some methods focused on the biases in captioning models [17], gender biases in multi-label object classification [43], biases in ConvNets and ImageNet [36], etc. Since our work focuses on reducing the language biases in VQA models, in the following, we discuss the related works lying on the same line.

Approaches in the dataset side: In VQA, various datasets have been proposed so far such as [3, 14, 21, 24]. The main source of bias in these datasets are from the language side, i.e., the question. There exist strong correlations between some of the word patterns in the question and the frequent answer words in the dataset. For e.g., in the large and popular VQAv1 dataset [3], there exist strong language priors such as "what color.." to "white," "is there.." to "yes," "how many.." to "2", etc. To reduce the language biases in VQAv1, in [14], the authors introduced a more balanced VQAv2 dataset, by adding at least two similar images with different answers for each question. However, with this additional refinement also, there still exists a considerable amount of language biases that can be leveraged by the model during training. In [1, 27, 1]8], the authors empirically show that a question-only model (without using the image) trained on VQAv2 shows reasonable performance, indicating the strong language priors in the dataset. In [1], the authors show that since the train and test answer distributions of VQAv2 are still similar, and the models that solely memorize language biases in the train set show acceptable performance while testing. In this regard, they introduced a diagnostic dataset, the VQA-CP (VQA under Changing Priors), to measure the language bias learned by the VQA models. This dataset is constructed with vastly different answer distributions between the train and test splits. Hence the models that rely heavily on the language priors in the train set will show poor performance while testing. [1] empirically shows that most of the existing best performing VQA models on the VQAv2 dataset show a significant drop in the accuracy when tested on the new test split from VQA-CPv2. E.g., a model with $\approx 66\%$ accuracy in the standard VQAv2 shows only $\approx 40\%$ accuracy in the bias sensitive VQA-CPv2.

Approaches in the architecture side: Several approaches have been considered in the literature to remove such biases. In [1], the authors propose a specific VQA model built upon [40], the GVQA model, consists of specially designed architectural restrictions to prevent the model from relying on the question-answer correlations. This model is complicated in design and requires multi-stage training, which makes it challenging to adapt to existing VQA models. In [32, 7], the authors propose different regularization techniques that can be applied via the question encoder to reduce the bias. In [32], the authors added an adversary question-only branch to the model, and it is supposed to find the amount of bias that exists in the question. Then, a gradient negation of the question-only branch loss is applied to remove answer discriminative features

from the question representation of the original model. They also propose a difference of entropy (DoE) loss captured between the output distributions of the VQA model and the additional question-only branch. In [7], the authors propose a similar approach as in [32], named RUBi, where instead of using the DoE and question-only loss, they mask (using the *sigmoid* operation) the output from the question-only branch and element-wise multiplying it with the output of the original model, to dynamically adapt the value of the classification loss to reduce the bias. However, both of these regularization approaches show a decrease in the performance on the standard VQAv2 benchmark, while reducing the bias. In [33], the authors propose a tuning method, called HINT, that is used to improve the visual-grounding of the existing model. Specifically, they tuned the model with manually annotated attention maps from the VQA-HAT dataset [11]. Similar to this, in [39], the authors propose a tuning approach called SCR, which uses additional manually annotated attention maps (from the VQA-HAT [11]) or textual explanations (from the VQA-X dataset [20]) to improve the visual-grounding of the model. However, both of these approaches [33, 39] require additional manually annotated data, and collecting the same is expensive. Also, these approaches reduce the performance on the standard VQAv2. Recently, in [9], the authors proposed an approach based on language-attention and is built upon the GVQA model [1]. The language attention module splits the question into three language phrases: question type, referring objects, and specific features of the referring object. These phrases further utilized to infer the answers from the image. They claim that splitting the question into different language phrases reduces the learning of the bias. However, this approach also shows a reduction in the performance on VQAv2.

The prior works are trying to overcome the bias either by reducing the contribution from the language side as in [32, 7, 9], or by improving the visualgrounding by tuning with additional manually annotated data as in [33, 39]. One common drawback of all of these methods is that they show a significant drop in the performance on the standard VQAv2 benchmark while reducing the bias. On the contrary, our approach improves the representation power of the question encoder to make visually-grounded question encodings to reduce the bias. The advantage of such an approach is that it is not only model-agnostic but also does not sacrifice the performance on the standard VQAv2 benchmark. Also, it does not require any additional manually annotated data or tuning.

Traditional question encoders and pitfall: The widely adopted question encoding scheme in VQA is passing the word-level features to a recurrent sequence model (LSTM [18] or GRU [10]) and taking the final state vector as the encoded question. Some early works use one-hot representations of question words and pass through an LSTM such as in [3, 13, 41], or GRU as in [4]. The drawback of using the one-hot vector representations of question words is the manual creation of the word vocabulary. Later, the usage of pre-trained word embedding vectors such as GloVe [30], BERT [12], etc., becomes popular, since they do not require any word vocabulary and also rich in linguistic understanding. This question encoding scheme, i.e., pass the pre-trained word embeddings of the

question words to some RNNs, is the currently popular approach in VQA [2, 22, 6, 5, 28, 29]. However, recently, with the success of Transformer based models in various NLP tasks [12], usage of Multi-modal Transformers (MMT) are becoming popular, such as ViLBERT [26], LXMERT [37], etc. These are co-attention models working on top of the Transformers [38], where while encoding the question, the words are prioritized with the guidance of the visual information.

The above-mentioned question-encoding schemes only use the language modality, i.e., the word-embeddings of the question words. As a result, the encoded question contains only the linguistic information from the question, and it cannot distinguish the questions based on their visual counterpart. This situation forces the model to learn the unwanted correlations between the language patterns in the questions, and the most frequent answers in the train set without considering the image which leads to a language side biased model. On the contrary, our question encoder, the VGQE, considers both visual and language modalities equally and generate visually-grounded question representations. A VQA model with such a question representation can reduce the over-dependency on the language priors in the train set.

The VGQE is also related to approaches such as FiLM [31], in terms of the early usage of the complementary information. The FiLM uses the question context to influence the image encoder, whereas VGQE uses the visual information inside the question encoder.

3 Visually-Grounded Question Encoder (VGQE)

We follow the widely-adopted RNN based question encoding scheme in VQA as the base. In VGQE, instead of the traditional RNN cell, a specially designed cell called the VGQE cell is used. A VGQE cell takes the word embedding of the current question word and finds its relevant visual counterpart feature from the image, then creates a visually-grounded question word embedding. This new word embedding is passed to an RNN cell to encode the sequence information.

Before going into the VGQE cell details, we explain the image and question feature representations used by the model. The image is represented by two sets of features $V = \{v_i \in \mathbb{R}^{d_v}\}_{i \in [1,k]}$ and $L = \{l_i \in \mathbb{R}^{d_w}\}_{i \in [1,k]}$ corresponding to the CNN features of objects and embeddings of the class labels respectively, where k is the total number of objects. The question is represented as a sequence of words with corresponding word embedding vectors, the word embedding of the t^{th} question word is being denoted by $q_t \in \mathbb{R}^{d_w}$.

3.1 VGQE cell

The basic building block of the proposed question encoder is the VGQE cell. This is analogous to the RNN cell (LSTM [18] or GRU [10]) in the traditional approach. An illustration of the VGQE cell at time t is shown in Fig. 2. Each VGQE cell takes the object-level features V, L, the current question word embedding q_t and the previous state context vector h_{t-1} , and then outputs the



Fig. 2. VGQE cell at time t: VGW module finds the visually-grounded word embedding g_t for the current question word q_t and the RNN cell (GRU or LSTM) encodes the sequential information in the question. V and L are the sets of object-level features.

current state context vector h_t . The context vector h_t is then pass to the next VGQE cell in sequence. Mathematically we represent a VGQE cell as follows:

$$h_t = \text{VGQE}(V, L, q_t, h_{t-1}) \tag{1}$$

A VGQE cell consists of two modules: the Visually-Grounded Word embedding (VGW) module and a traditional RNN cell (GRU [10] or LSTM [18]). The VGW module is responsible for finding the visually-grounded word embedding vector g_t for the current question word embedding q_t . Then, g_t is passed to the RNN cell to encode the question sequence information.

VGW module: This module is the crux of the VGQE cell, where the visuallygrounded word embedding g_t for the current question word q_t is extracted. It works in two stages: 1) Attention and 2) Fusion. The first stage calculates the visual counterpart feature f_t of the question word embedding q_t , while in the fusion stage, f_t and q_t are fused to generate the visually-grounded word embedding vector g_t .

Attention: This module takes the set of object-label features L and the current question word q_t as inputs and outputs a relevance score vector $\alpha_t \in \mathbb{R}^k$. Each of the k values in α_t tells the relevance of the corresponding object from the image in the context of the current question word q_t . Then, one single visual feature vector f_t is extracted by taking the weighted sum over the set of object CNN features V, where the weights are defined by α_t . Mathematically, the above steps are formulated as follows:

7

$$f_t = \sum_{i=1}^k \alpha_t[i] * v_i$$

$$\alpha_t = softmax(w_a^T(W(\mathbf{L} * (\mathbb{1}q_t))^T))$$
(2)

where * is element-wise multiplication, $v_i \in V$ is the i^{th} object CNN feature vector, $\alpha_t \in \mathbb{R}^k$ and $f_t \in \mathbb{R}^{d_v}$ are the object-relevance score vector and the visual counter part feature for the current question word q_t respectively, $\alpha_t[i]$ is the relevance of the i^{th} object at time $t, \mathbf{L} \in \mathbb{R}^{k \times d_w}$ is the matrix representation of the set of object-label features, $w_a \in \mathbb{R}^{d_w}$ and $W \in \mathbb{R}^{d_w \times d_w}$ are learnable parameters and $\mathbb{1}$ is a column vector of length k consists all ones.

Fusion: In this stage, the word-level visual feature f_t is grounded to the wordlevel language feature q_t . This results in the visually-grounded word embedding vector g_t . For fusion, any multi-modal fusion function F_m such as element-wise multiplication, bi-linear fusion [13, 4, 23, 41, 5] etc. can be used. In this paper, we use the BLOCK fusion as F_m , which is a recently proposed best performing bilinear multi-modal fusion method [5]. The formal representation of the visuallygrounded word embedding vector g_t is as follows:

$$g_t = F_m(f_t, \hat{q}_t; \Theta) \tag{3}$$

where F_m is the multi-modal fusion function (in our case BLOCK) with Θ as the learnable parameters, the vector $\hat{q}_t = W_q(q_t)$ is the fine-tuned word embedding vector with W_q as the learnable parameters of a two-layer network that projects from d_w -space to d-space, and f_t is the visual feature as defined in Eq. (2). The fused vector g_t contains relevant language information from the word embedding q_t and visual information from f_t ; in other words, g_t is a visually-grounded word embedding vector. For the same question word (e.g., banana), the g_t vector will vary according to their visual counterpart (as a result, the "green banana" and "yellow banana" will get distinguishable representations in the question.).

The output from the VGW module, the visually-grounded word embedding vector g_t along with the previous state context vector h_{t-1} , is then passed to a traditional RNN cell to generate the current state question context vector h_t .

3.2 Using VGQE cell to encode the question

In the proposed question encoder, we use the VGQE cell instead of the traditional RNN cell. An illustration of a generic VQA model with VGQE is shown in Fig. 3. The visual encoder encodes the image as in the original model. The question is encoded using VGQE instead of the existing question encoder. The question word embeddings are passed through a sequence of VGQE cells along with the object-level features (V and L) and take the final state representation as the encoded question. In this paper, we use GRU as the RNN cell inside the VGQE cell. Other question-encoder adaptations using variants of RNN, such as LSTM, are also possible, with appropriate changes in the RNN cell of VGQE. Since, in



Fig. 3. An illustration of a generic VQA model with VGQE

VGQE, the question words are grounded on its visual counterpart, the encoded questions itself contain sufficient visual-grounding. Hence, they become robust enough to reduce the correlation between specific language patterns and the most frequent answers in the train set.

3.3 Baseline VQA architecture

In this paper, we adopt the same baseline as used in prior work [7], to perform various experiments, also to make a fair comparison. This baseline is a simplified version of the MUREL VQA model [6]. The word embeddings are passed to a GRU [10] to encode the question. Then, each of the object CNN features $v_i \in V$ are bi-linearly fused (using BLOCK fusion [5]) with the encoded question to get the question-aware object features. A max-pool operation over these features gives a single vector that later used by the answer prediction network.

4 Experiments and Results

4.1 Experimental setup

We train and evaluate the models on the VQA-CPv2 dataset [1]. This dataset is designed to test the robustness of VQA models in dealing with language biases. The train and test sets of VQA-CPv2 have totally different answer distributions; hence a model that strongly depends on the language priors in the train set will perform poorly on the test set. We also use the VQAv2 dataset to train and evaluate the models on the standard VQA benchmark. We use the VQA accuracy [3] as the evaluation metric.

We use the object-level CNN features provided in [2] as V ($d_v = 2048$). We use the pre-trained Glove [30] for the object-label features L and question word embeddings q_t ($d_w = 300$). We train the model using Adamw [25] with weightdecay = $2 * 10^{-5}$ and cross-entropy loss. Further implementation details are provided in the supplementary material.

Model	Overall	Yes/No	Number	Other
GVQA [1]	31.30	57.99	13.68	22.14
RAMEN [34]	39.21	-	-	-
BAN [22]	39.31	-	-	-
MUREL [6]	39.54	42.85	13.17	45.04
UpDn [2]	39.74	42.27	11.93	46.05
UpDn+Q-Adv+DoE [32]	41.17	65.49	15.48	35.48
UpDn+HINT [33]	46.73	67.27	10.61	45.88
UpDn+LangAtt [9]	48.87	70.99	18.72	45.57
UpDn+SCR (VQA-X) [39]	49.45	72.36	10.93	48.02
Baseline [7]	38.46	42.85	12.81	43.20
Baseline+RUBi [7]	47.11	68.65	20.28	43.18
Baseline+VGQE	50.11	66.35	27.08	46.77

Table 1. Comparison with existing models on VQA-CPv2.

4.2 Results

Comparison with state-of-the-art models: In Table 4.1, we compare VGQE against existing state-of-the-art bias reduction techniques in VQA-CPv2. We can see that our approach achieves a new state-of-the-art. With VGQE, we improved the baseline model (that uses the traditional GRU based question encoder) accuracy from 38.46 to 50.11. This performance also corresponds to a notable improvement from the RUBi approach [7], where they also use the same baseline. Our approach outperforms the UpDn based approaches [32, 33, 9, 39] as well. Comparing with GVQA, a specific model designed for VQA-CP, our generic approach outperforms with a +18.81 gain in the accuracy.

Question-type-wise results: In Table 2, we show the comparison of some of the question-type-wise results of the baseline and with VGQE (+VGQE)models. The baseline model shows comparatively poor performance than the model with VGQE. Since most of the time, the baseline model with a traditional question encoder tries to memorize the language biases in the train set. We can see that, in all the question-types, the incorporation of VGQE reduces such biases and improves performance. To further clarify this, in Fig. 4, we visualize the answer distributions of some of the question-types in the VQA-CPv2 train & test sets and outputs of the baseline & baseline+VGQE models (less frequent answers are ignored for better visualization). It is clear from the visualizations that the baseline model learns the language biases in the train set and predicts the answer without considering the image. The incorporation of VGQE helps to reduce such biases and to answer by grounding the question onto the image. E.g., from Fig. 4, consider the case of the "Do you" type questions. Most of such questions are linked to the answer "no" in the train set; in other words, there is a strong correlation between the "Do you" language pattern and the answer "no." We can see that most of the time, the baseline model predicts the answer "no" upon seeing the pattern "Do you," without considering the image.

Table 2. Performance comparison between the baseline and VGQE corresponding to some of the question-types in the VQA-CPv2 test set. VGQE reduces the learning of unwanted question pattern-answer correlations and improves the performance. All numbers are VQA accuracy values.

Question Type	Baseline	+VGQE
Do you	30.7	93.19
Can you	29.04	52.54
What are the	39.8	48.52
What color are the	28.32	64.42
What sport is	88.39	93.19
What is the person	57.40	63.65
What time	39.28	57.88
What room is	80.34	92.87
What is in the	29.34	35.84
Where are the	19.87	31.88
What brand	28.54	47.19
How many	12.91	19.65
How many people	36.62	50.03
Does the	31.91	77.39
Which	26.73	39.0
Why	12.28	14.76



Fig. 4. Answer distributions from train & test sets of VQA-CPv2, and outputs of baseline & baseline+VGQE models, for some question types from Table 2. We can see that VGQE helps the model to reduce the learning of bias from the train set.

This indicates that the baseline model relies on the strong correlation between the language pattern "Do you" and the answer "no," rather than looking at the image. In contrast, the incorporation of VGQE helps to reduce such learning of biases and to predict the answer by considering the image as well.

On comparing the performance gain among the question types, we can see that the questions that start with "why" show relatively less improvement. Since these questions usually require common sense reasoning (e.g., Why people wear the hat?), and in such cases, the VGQE cannot improve much from the baseline.

Qualitative results: In Fig. 5, we show some qualitative comparisons between the baseline (Base) and baseline+VGQE (Ours) models. Fig 5 (a), shows the case of the "What color are the" question pattern, where the train set has heavy biases towards the colors "black", "red", "blue", "brown", and "yellow" (see Fig.4). In Fig 5 (b) and (c), we show the cases of the same question asked for different images. The question "What sport is being played?" has a bias on "tennis" ($\approx 63.2\%$) and "soccer" ($\approx 17.5\%$) whereas "What color is the fire hydrant?" is biased towards the colors "yellow" ($\approx 46\%$), "white" ($\approx 12\%$) and "black" ($\approx 10\%$) in the train set [1]. These visualizations further show that incorporating VGQE helps the baseline model to reduce the learning of biases.



Q: What color are the
bananas. GT: green,Q: What color are the
umbrellas. GT: white,Q: What color are the
cabinets. GT: white,Ours:green,Base: yellow.Ours:white,Base: blue.Ours:white, Base: brown.



(a) Question Pattern: What color are the



GT: red, GT: red, GT: red, Ours: red, Base: yellow Ours: red, Base: white. Ours: red, Base: black (c) Common Question: What color is the fire hydrant?

Fig. 5. Some qualitative comparisons of our model (Baseline+VGQE) with Baseline. Ours, Base and GT represents the baseline+VGQE, baseline and the ground truth, respectively. Best viewed in color.

In Fig. 6, we show a visualization of the grounded image regions at each time step of VGQE with the same question ("What sport is being played?") but different images. We can see that the question words are grounded onto the relevant visual counterparts. Hence, the same words will get different representations based on the image. For e.g., the word "sport" is grounded on the "baseball" player in the first case, whereas it is grounded on the "frisbee" player in the second case. Hence, even though the question is the same, the VGQE generates different visually-grounded question representations.

4.3 Performance of VGQE on other baselines

The proposed question encoder, the VGQE, can be easily incorporated into any existing VQA model, i.e., replace the existing question encoder with VGQE (See



Fig. 6. A visualization of the grounded visual regions at each time step of VGQE with the same question ("What sport is being played?") but different images. The same words will get different visually-grounded embeddings based on the visual counterpart.

Fig. 3 for an illustration). We have already shown the effectiveness of VGQE on the MUREL baseline (see Table 4.1). In this section, we show (see Table 3) the impact of VGQE on two more recent best performing models as well:

- UpDn [2]: This VQA model is based on question guided visual attention. The word embeddings are pass to a GRU to encode the question. The image is encoded using visual attention guided by the question representation, over the set of object CNN features V. Then, both encoded image and question are combined using element-wise multiplication and pass to the answer predictor.
- **BAN** [22]: This model is based on bi-linear co-attention maps between the objects and question words. The question word embeddings are passed through a GRU. Then all the GRU cell outputs and the object-level CNN features are given to a co-attention module, the BAN. This module will give a combined vector which later used by the answer predictor.

In both models, we replaced the traditional GRU-based question encoder with VGQE. Note that, unlike other model-agnostic bias reduction methods, VGQE does not require a new question-only branch as in [32, 7] or tuning with additional annotated data as in [33, 39], while training. The results are shown in Table 3. We can see that, in both models, VGQE consistently improves the performance (39.74 to 48.75 in UpDn [2] and 39.31 to 50.0 in BAN [22]) on VQA-CPv2, showing that here also VGQE reduces the language-bias.

4.4 Performance of VGQE on the standard VQAv2 benchmark

Existing bias reduction techniques such as [32, 7, 9, 33, 39] show a reduction in the performance on the standard VQAv2 benchmark. For instance, in Table 3, the best performing UpDn based models on VQA-CPv2, such as UpDn+LangAtt [9] reduce the accuracy in the VQAv2 val set from 63.48 to 57.96 & UpDn+SCR [39]

			Table 4. Co	mparison (
Model	VQA-CPv2	VQAv2 val	performance	of the bas
UpDn [2] UpDn+Q-Adv [32] UpDn+DoE [32] UpDn+Q-Adv+DoE [32] UpDn+RUBi [7]	$39.74 \\ 40.08 \\ 40.43 \\ 41.17 \\ 44.23 \\ 42.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.23 \\ 44.2$	63.48 60.53 63.43 62.75	line model and VGQE val set. RUI drop in th whereas VGC show any dro	with RUI on VQAy Bi shows e accurad QE does no p.
UpDn+HINT [33] UpDn+LangAtt [9] UpDn+SCR [39] UpDn+SCR (HAT) [39] UpDn+SCR (VQA-X) [39]	46.73 48.87 48.47 49.17 49.45	$\begin{array}{c} 63.38 \\ 57.96 \\ 62.3 \\ 62.2 \\ 62.2 \end{array}$	Model Baseline [7] +RUBi [7]	VQAv2 va 63.10 61.16
UpDn+VGQE	48.75	64.04	+VGQE	63.18
BAN [22,7] BAN+VGQE	39.31 50.00	65.36 65.73		

Table 3. Effect of VGQE on UpDn [2] and BAN [22] models, on VQA-CPv2 and VQAv2 val.

reduces it from 63.48 to 62.2. Note that, UpDn+SCR reduces the performance even though it uses additional manually annotated data (HAT or VQA-X). Also, in Table 4, adding RUBi to the baseline model decreases the accuracy from 63.10 to 61.16. On the contrary, VGQE does not sacrifice the performance on the standard VQAv2 benchmark. Interestingly, VGQE slightly improves the performance of the respective models; from 63.48 to 64.04, 63.10 to 63.18 and 65.36 to 65.73 in UpDn [2], baseline [7] and BAN [22] respectively. We ascribe this to the following; VGQE adds a robust and visually-grounded question representation to the model without damaging its existing reasoning power.

5 Conclusion

Current VQA models rely heavily on the language priors that exists in the train set, without considering the image. Such models that fail to utilize both modalities equally would likely perform poorly in real-world scenarios. We propose VGQE, a novel question-encoder that utilizes both the modalities equally and generates visually-grounded question representations. Such question representations have sufficient distinguishing power based on the visual counterpart and help the models to reduce the learning of language biases from the train set. We did extensive experiments on the bias sensitive VQA-CPv2 dataset and achieved a new state-of-the-art. The VGQE is model agnostic and can easily incorporate into existing VQA models without the need for additional manually annotated data and training. We experimented with three best performing VQA models, and deliver consistent performance improvement in all of them. Further, unlike existing bias reduction techniques, VGQE does not sacrifice the original model's performance on the standard VQAv2 benchmark.

References

- Agrawal, A., Batra, D., Parikh, D., Kembhavi, A.: 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. pp. 4971–4980 (2018)
- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L.: Bottom-up and top-down attention for image captioning and visual question answering. In: IEEE conference on computer vision and pattern recognition. pp. 6077–6086 (2018)
- Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Lawrence Zitnick, C., Parikh, D.: Vqa: Visual question answering. In: Proceedings of the IEEE international conference on computer vision. pp. 2425–2433 (2015)
- Ben-Younes, H., Cadene, R., Cord, M., Thome, N.: Mutan: Multimodal tucker fusion for visual question answering. In: Proc. IEEE Int. Conf. Comp. Vis. vol. 3 (2017)
- Ben-Younes, H., Cadene, R., Thome, N., Cord, M.: Block: Bilinear superdiagonal fusion for visual question answering and visual relationship detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 8102–8109 (2019)
- Cadene, R., Ben-Younes, H., Cord, M., Thome, N.: Murel: Multimodal relational reasoning for visual question answering. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 1989–1998 (2019)
- Cadene, R., Dancette, C., Cord, M., Parikh, D., et al.: Rubi: Reducing unimodal biases for visual question answering. In: Advances in Neural Information Processing Systems. pp. 839–850 (2019)
- Chao, W.L., Hu, H., Sha, F.: Being negative but constructively: Lessons learnt from creating better visual question answering datasets. arXiv preprint arXiv:1704.07121 (2017)
- Chenchen Jing, Yuwei Wu, X.Z.Y.J., Wu., Q.: Overcoming language priors in vqa via decomposed linguistic representations. In: The Thirty-Fourth AAAI Conference on Artificial Intelligence, 2020. (2020)
- Das, A., Agrawal, H., Zitnick, L., Parikh, D., Batra, D.: 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)
- Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (2019)
- Fukui, A., Park, D.H., Yang, D., Rohrbach, A., Darrell, T., Rohrbach, M.: Multimodal compact bilinear pooling for visual question answering and visual grounding. In: Conference on Empirical Methods in Natural Language Processing. pp. 457–468 (2016)
- Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., Parikh, D.: Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 6904–6913 (2017)

- 16 Gouthaman KV and Anurag Mittal
- Gurari, D., Li, Q., Stangl, A.J., Guo, A., Lin, C., Grauman, K., Luo, J., Bigham, J.P.: Vizwiz grand challenge: Answering visual questions from blind people. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3608–3617 (2018)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Hendricks, L.A., Burns, K., Saekno, K., Darrell, T., Rohrbach, A.: Women also snowboard: Overcoming bias in captioning models. European Conference on Computer Vision (ECCV) (2018)
- Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation 9(8), 1735–1780 (1997)
- Hudson, D.A., Manning, C.D.: Gqa: A new dataset for real-world visual reasoning and compositional question answering. Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Huk Park, D., Anne Hendricks, L., Akata, Z., Rohrbach, A., Schiele, B., Darrell, T., Rohrbach, M.: Multimodal explanations: Justifying decisions and pointing to the evidence. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8779–8788 (2018)
- Kafle, K., Kanan, C.: An analysis of visual question answering algorithms. In: ICCV (2017)
- Kim, J.H., Jun, J., Zhang, B.T.: Bilinear attention networks. In: Advances in Neural Information Processing Systems. pp. 1564–1574 (2018)
- Kim, J.H., On, K.W., Lim, W., Kim, J., Ha, J.W., Zhang, B.T.: Hadamard product for low-rank bilinear pooling. ICLR (2017)
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. International Journal of Computer Vision 123(1), 32–73 (2017)
- 25. Loshchilov, Ι., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (2019),https://openreview.net/forum?id=Bkg6RiCqY7
- Lu, J., Batra, D., Parikh, D., Lee, S.: Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In: Advances in Neural Information Processing Systems. pp. 13–23 (2019)
- Manjunatha, V., Saini, N., Davis, L.S.: Explicit bias discovery in visual question answering models. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9562–9571 (2019)
- Nguyen, D.K., Okatani, T.: Improved fusion of visual and language representations by dense symmetric co-attention for visual question answering. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 6087–6096 (2018)
- Peng Gao, Haoxuan You, Z.Z.: Multi-modality latent interaction network for visual question answering. In: Proceedings of the IEEE International Conference on Computer Vision (2019)
- Pennington, J., Socher, R., Manning, C.: Glove: Global vectors for word representation. In: Conference on Empirical Methods in Natural Language Processing. pp. 1532–1543 (2014)
- Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)

- Ramakrishnan, S., Agrawal, A., Lee, S.: Overcoming language priors in visual question answering with adversarial regularization. In: Advances in Neural Information Processing Systems. pp. 1541–1551 (2018)
- 33. Selvaraju, R.R., Lee, S., Shen, Y., Jin, H., Ghosh, S., Heck, L., Batra, D., Parikh, D.: Taking a hint: Leveraging explanations to make vision and language models more grounded. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2591–2600 (2019)
- Shrestha, R., Kafle, K., Kanan, C.: Answer them all! toward universal visual question answering models. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10472–10481 (2019)
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
- Stock, P., Cisse, M.: Convnets and imagenet beyond accuracy: Understanding mistakes and uncovering biases. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 498–512 (2018)
- 37. Tan, H., Bansal, M.: 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). pp. 5103–5114 (2019)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017)
- Wu, J., Mooney, R.: Self-critical reasoning for robust visual question answering. In: Advances in Neural Information Processing Systems. pp. 8601–8611 (2019)
- 40. Yang, Z., He, X., Gao, J., Deng, L., Smola, A.: Stacked attention networks for image question answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 21–29 (2016)
- Yu, Z., Yu, J., Fan, J., Tao, D.: Multi-modal factorized bilinear pooling with coattention learning for visual question answering. In: Proc. IEEE Int. Conf. Comp. Vis. vol. 3 (2017)
- Zellers, R., Bisk, Y., Farhadi, A., Choi, Y.: From recognition to cognition: Visual commonsense reasoning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6720–6731 (2019)
- Zhao, J., Wang, T., Yatskar, M., Ordonez, V., Chang, K.W.: Men also like shopping: Reducing gender bias amplification using corpus-level constraints. arXiv preprint arXiv:1707.09457 (2017)