

# Generative Bias for Robust Visual Question Answering

Jae Won Cho<sup>1</sup> Dong-Jin Kim<sup>2</sup> Hyeonggon Ryu<sup>1</sup> In So Kweon<sup>1</sup>

<sup>1</sup>KAIST <sup>2</sup>Hanyang University

## Abstract

The task of Visual Question Answering (VQA) is known to be plagued by the issue of VQA models exploiting biases within the dataset to make its final prediction. Various previous ensemble based debiasing methods have been proposed where an additional model is purposefully trained to be biased in order to train a robust target model. However, these methods compute the bias for a model simply from the label statistics of the training data or from single modal branches. In this work, in order to better learn the bias a target VQA model suffers from, we propose a generative method to train the bias model directly from the target model, called GenB. In particular, GenB employs a generative network to learn the bias in the target model through a combination of the adversarial objective and knowledge distillation. We then debias our target model with GenB as a bias model, and show through extensive experiments the effects of our method on various VQA bias datasets including VQA-CP2, VQA-CP1, GQA-OOD, and VQA-CE, and show state-of-the-art results with the LXMERT architecture on VQA-CP2.

## 1. Introduction

Visual Question Answering (VQA) [5] is a challenging multi-modal task that requires a model to correctly understand and predict an answer given an input pair of image and question. Various studies have shown that VQA is prone to biases within the dataset and tend to rely heavily on language biases that exists within the dataset [2, 17, 51], and VQA models tend to predict similar answers only depending on the question regardless of the image. In response to this, recent works have developed various bias reduction techniques, and recent methods have exploited ensemble based debiasing methods [7, 13, 19, 41] extensively.

Among ensemble based methods, additional models are introduced to concurrently learn biases that might exist within each modality or dataset. For example, in works such as [7, 19], the Question-Answer model is utilized to determine the language prior biases that exist when a model is asked to give an answer based solely off of the question.

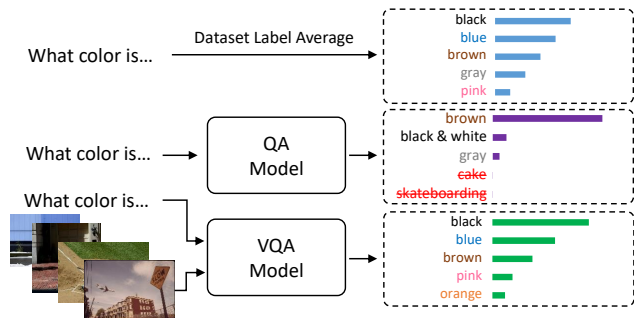


Figure 1. Given a Question Type (“What color is...”), we show all of the averaged answers within the training dataset. The answer computed from the entire training dataset is the known dataset label average or dataset bias as in [13, 19]. We see that the averaged model predictions of the Question-Answer Model and Visual-Question-Answer Model are significantly different.

This Question-Answer model is then utilized to train a robust “target” model, which is used for inference. The key purpose of an ensemble “bias” model is to capture the biases that are formed with its given inputs (*i.e.*, language prior biases from the Question-Answer model). In doing so, if this model is able to represent the bias well, this bias model can be used to teach the target model to avoid such biased answers. In other words, the better the bias model can learn the biases, the better the target model can avoid such biases.

Existing ensemble based methods either use pre-computed label statistics of training data (GGE-D [19] and LMH [13]), or single modal branches that compute the answer from either the question or image [7, 13, 19, 37]. However, we conjecture that there is a limit to the bias representation that can be obtained from such methods, as the model’s representative capacity is limited by its inputs. In addition, pre-computed label statistics represents only part of the bias [19]. As shown in Fig. 1, given a question type, the pre-computed label statistics (or known dataset bias) are noticeably different to the predictions of a model trained with the question or with the image and question. This discrepancy signifies that there is a part of the bias that we cannot fully model simply with the previous methods. Therefore, we propose a novel stochastic bias model that learns

the bias *directly from the target model*.

More specifically, to directly learn the bias distribution of the *target model*, we model the bias model as a Generative Adversarial Network (GAN) [16] to stochastically mimic the target model’s answer distribution given the same question input by introducing a random noise vector. As seen through literature, most biases are held within the question [2], so we use questions as the main bias modality. To further enforce this, we utilize knowledge distillation [20] on top of adversarial training to force the bias model to be as close as possible to the target model, so that the target model learns from harder negative supervision from the bias model. Finally, with our generative bias model, we then use our modified debiasing loss function to train our target model. Our final bias model is able to train the target model that outperforms previous uni-modal and multi-modal ensemble based debiasing methods by a large margin. To the best of our knowledge, we are the first to train the bias model by directly leveraging the behavior of the target model using a generative model for VQA.

To show the efficacy and robustness of our method, we perform extensive experiments on commonly used robustness testing VQA datasets and various different VQA architectures. Our method show the state-of-the-art results on all settings without the use of external human annotations or dataset reshuffling methods.

Our contributions are as follows:

- We propose a novel bias model for ensemble based debiasing for VQA by directly leveraging the target model that we name *GenB*.
- In order to effectively train GenB, we employ a Generative Adversarial Network and knowledge distillation loss to capture both the dataset distribution bias and the bias from the target model.
- We achieve state-of-the-art performance on VQA-CP2, VQA-CP1 as well as the more challenging GQA-OOD dataset and VQA-CE using the simple UpDn baseline without extra annotations or dataset reshuffling and state-of-the-art VQA-CP2 performance on the LXMERT backbone.

## 2. Related Work

VQA [5, 18, 25] has been actively studied in recent years with performance reaching close to human performance [11, 34, 36, 44, 45] in the most recent works. Even still, the VQA dataset has been notorious for its reliance on language biases as shown by [3]. To this date, many VQA datasets and evaluation protocols have been derived from the original VQA dataset and have been released to the public as a means to test and understand the biases and

robustness of VQA models such as VQA-CP2 and CP1 [3], GQA-OOD [26], and VQA-CE [14].

In light of this, many methods have been proposed. One line of work strives to improve visual attention through visual grounding by including additional human annotations or explanations [42, 48], which show limited improvements. Another line of work changes the distribution of the training set by either randomly or systematically replacing the images and questions [10, 46, 47, 49, 52] or by augmenting the dataset through counterfactual samples [9] or through the use of an external inpainting network [15]. Although powerful, according to [37], changing the distribution of the training set does not agree with the philosophy of the creation of this dataset. Therefore, we do not directly compare experimentally with this line of works. Most recently, another line of work has been released where an additional objective of Visual Entailment is added to further boost performance [43], which we do not compare for fairness. Another recent work focuses on distilling knowledge from the some of the aforementioned methods to train a robust student model [38]. As this method is built upon other methods, we compare to this method where possible.

One of the most effective line of work is ensemble based methods [1, 7, 19, 37] and we place our line of work here. Within ensemble based methods, AReg [1], RUBi [7], LMH [13], GGE [19] tackle the language prior directly and only use a question-only model. Unlike these methods, GGE [19] shows the possibility of using the same model as the target model to learn biases but to a limited extent. On the other hand, CF-VQA [37] uses both question and image, but uses the modalities individually without combining them. Our work is distinct from all previous ensemble based methods as we use a generative network with a noise input to aid the bias model in learning the bias directly from the target model. We also believe that our work can be applied to other multi-modal [6, 22, 23, 28, 31–33] and uni-modal research [39] tasks in the future.

## 3. Methodology

In this section, we explain VQA briefly and describe in detail our method GenB.

### 3.1. Visual Question Answering Baseline

With an image and question as a pair of inputs, a VQA model learns to correctly predict an answer from the whole answer set  $\mathcal{A}$ . A typical VQA model  $F(\cdot, \cdot)$  takes both a visual representation  $\mathbf{v} \in \mathbb{R}^{n \times d_v}$  (a set of feature vectors computed from a Convolutional Neural Network given an image where  $n$  is the number of objects in an image and  $d_v$  being the vector dimension) and a question representation  $\mathbf{q} \in \mathbb{R}^{d_q}$  (a single vector computed from a GloVe [40] word embedding followed by a Recurrent Neural Network given a question) as input. Then, an atten-




Figure 2. (a) shows how we train our Bias Model and Discriminator. The Bias Model is trained with 3 different losses including the ground truth BCE (Eq. (1)), knowledge distillation (Eq. (3)), and adversarial Eq. (2)) losses. (b) shows how our Target Model is trained with the Bias model with debiasing loss functions (refer to existing works). Note, steps (a) and (b) happen concurrently and note that we only use the Target Model during inference.

tion module followed by a multi-layer perceptron classifier  $F : \mathbb{R}^{n \times d_v} \times \mathbb{R}^{d_q} \rightarrow \mathbb{R}^{|\mathcal{A}|}$  which generates an answer logit vector  $\mathbf{y} \in \mathbb{R}^{|\mathcal{A}|}$  (i.e.,  $\mathbf{y} = F(V, Q)$ ). Then, after applying a sigmoid function  $\sigma(\cdot)$ , our goal is to make an answer probability prediction  $\sigma(\mathbf{y}) \in [0, 1]^{|\mathcal{A}|}$  close to the ground truth answer probability  $\mathbf{y}_{gt} \in [0, 1]^{|\mathcal{A}|}$ . In this work, we adopt one of the popular state-of-the-art architectures UpDn [4] that is widely used in VQA debiasing research.

### 3.2. Ensembling with Bias Models

In this work, our scope is bias mitigation through ensembling bias models similar to previous works [7, 13, 19]. In ensemble based methods, there exist a “bias” model that generates  $\mathbf{y}_b \in \mathbb{R}^{|\mathcal{A}|}$ , which we define as  $F_b(\cdot, \cdot)$ , and a “target” model, defined as  $F(\cdot, \cdot)$ . Note that, we discard  $F_b(\cdot, \cdot)$  during testing and only use  $F(\cdot, \cdot)$ . As previously mentioned, the goal of the existing bias models is to overfit to the bias as much as possible. Then, given the overfitted bias model, the target model is trained with a debiasing loss function [7, 13, 19] to improve the robustness of the target model. Ultimately, the target model learns to predict an unbiased answer by avoiding the biased answer from the bias model. The bias model  $F_b(\cdot, \cdot)$  can either be the same or different from the original  $F(\cdot, \cdot)$  and there could be multiple models as well [37]. Although previous works try to leverage the bias from the individual modalities [7, 13, 19, 37], we propose that this limits the ability of the model to represent biases. Hence, in order to represent the biases *similar to the target model*, we set the architecture of  $F_b(\cdot, \cdot)$  to be the same as  $F(\cdot, \cdot)$  and we use the UpDn [4] model.

### 3.3. Generative Bias

As mentioned in the Sec. 1, as our goal is to train a bias model that can generate stochastic bias representations, we use a random noise vector in conjunction with a given modality to learn both the dataset bias and the bias that the target model could exhibit. As the question is known to be prone to bias, we keep the question modality and use it as the input to our bias model  $F_b(\cdot, \cdot)$ . But instead of using the image features, we introduce a random noise

vector  $\mathbf{z} \in \mathbb{R}^{n \times 128}$  in addition to a generator network  $G : \mathbb{R}^{n \times 128} \rightarrow \mathbb{R}^{n \times d_v}$  to generate the corresponding input to the bias model  $F_b(\cdot, \cdot)$ . Formally, given a random Gaussian noise vector  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ , a generator network  $G(\cdot)$  synthesizes a vector that has the same dimension as the image feature representation, i.e.,  $\hat{\mathbf{v}} = G(\mathbf{z}) \in \mathbb{R}^{n \times d_v}$ . Ultimately, our model takes in the question  $\mathbf{q}$  and  $G(\mathbf{z})$  as its input and generates the bias logit  $\mathbf{y}_b$  in the form  $F_b(G(\mathbf{z}), \mathbf{q}) = \mathbf{y}_b$ . Note, this can be done on another modality, (i.e.,  $F_b(G(\mathbf{z}), \mathbf{v}) = \mathbf{y}_b$ ), but we found this is unhelpful. For simplicity, we consider generator and bias model as one network and rewrite  $F_b(G(\mathbf{z}), \mathbf{q})$  in the form  $F_{b,G}(\mathbf{z}, \mathbf{q})$  and call our “Generative Bias” method **GenB**.

### 3.4. Training the Bias Model

In order for our bias model GenB to learn the biases given the question, we use the traditional VQA loss, the Binary Cross Entropy Loss:

$$\mathcal{L}_{GT}(F_{b,G}) = \mathcal{L}_{BCE}(\sigma(F_{b,G}(\mathbf{z}, \mathbf{q})), \mathbf{y}_{gt}). \quad (1)$$

However, unlike existing works, we want the bias model to also capture *the biases in the target model*. Hence, in order to mimic the bias of the target model as a random distribution of the answers, we propose adversarial training [16] similar to [29] to train our bias model. In particular, we introduce a discriminator that distinguishes the answers from the target model and the bias model as “real” and “fake” answers respectively. The discriminator is formulated as  $D(F(\mathbf{v}, \mathbf{q}))$  and  $D(F_{b,G}(\mathbf{z}, \mathbf{q}))$  or rewritten as  $D(\mathbf{y})$  and  $D(\mathbf{y}_b)$ . The objective of our adversarial network with generator  $F_{b,G}(\cdot, \cdot)$  and  $D(\cdot)$  can be expressed as follows:

$$\begin{aligned} & \min_{F_{b,G}} \max_D \mathcal{L}_{GAN}(F_{b,G}, D), \text{ where} \\ & \mathcal{L}_{GAN}(F_{b,G}, D) \\ &= \mathbb{E}_{\mathbf{v}, \mathbf{q}} \left[ \log \left( D(F(\mathbf{v}, \mathbf{q})) \right) \right] \\ & \quad + \mathbb{E}_{\mathbf{q}, \mathbf{z}} \left[ \log \left( 1 - D(F_{b,G}(\mathbf{z}, \mathbf{q})) \right) \right] \\ &= \mathbb{E}_{\mathbf{y}} \left[ \log \left( D(\mathbf{y}) \right) \right] + \mathbb{E}_{\mathbf{y}_b} \left[ \log \left( 1 - D(\mathbf{y}_b) \right) \right]. \end{aligned} \quad (2)$$

The generator ( $F_{b,G}$ ) tries to minimize the objective ( $\mathcal{L}_{GAN}$ ) against an adversarial discriminator ( $D$ ) that tries to maximize it. Through alternative training of  $D$  and  $F_{b,G}$ , the distribution of the answer vector from the bias model ( $\mathbf{y}_b$ ) should be close to that from the target model ( $\mathbf{y}$ ).

In addition, to further enforce bias model to capture the intricate biases present in the target model, we add an additional knowledge distillation objective [20] similar to [12, 27, 30, 32, 33] to ensure that the model bias model is able to follow the behavior of the target model with only the  $\mathbf{q}$  given to it. We empirically find that it is beneficial to include a sample-wise distance based metric such as KL divergence. This method is similar to the approaches in the image to image translation task [21]. Then, the goal of the generator is not only to fool the discriminator but also to try to imitate the answer output of the target model in order to give the target model more challenging supervision in the form of *hard negative* sample synthesis. We add another objective to our adversarial training for  $F_{b,G}(\cdot, \cdot)$ :

$$\mathcal{L}_{distill}(F_{b,G}) = \mathbb{E}_{\mathbf{v}, \mathbf{q}, \mathbf{z}} \left[ D_{KL}(F(\mathbf{v}, \mathbf{q}) \| F_{b,G}(\mathbf{z}, \mathbf{q})) \right]. \quad (3)$$

Ultimately, the final training loss for the bias model, or GenB, is as follows:

$$\begin{aligned} & \min_{F_{b,G}} \max_D \mathcal{L}_{GenB}(F_{b,G}, D), \text{ where} \\ & \mathcal{L}_{GenB}(F_{b,G}, D) = \\ & \mathcal{L}_{GAN}(F_{b,G}, D) + \lambda_1 \mathcal{L}_{distill}(F_{b,G}) + \lambda_2 \mathcal{L}_{GT}(F_{b,G}), \end{aligned} \quad (4)$$

where  $\lambda_1$  and  $\lambda_2$  are the loss weight hyper-parameters.

### 3.5. Debiasing the Target Model

Given a generated biased answer  $\mathbf{y}_b$ , there are several debiasing loss functions that can be used such as [7, 13, 19], and we show the effects of each one in Table 5. The GGE [19] loss is one of the best performing losses without the use of label distribution. The GGE loss takes the bias predictions/distributions and generates a gradient in the opposite direction to train the target model. With this starting point, we modify this equation with the ensemble of the biased model GenB in this work as follows:

$$\mathcal{L}_{target}(F) = \mathcal{L}_{BCE}(\mathbf{y}, \mathbf{y}_{DL}), \quad (5)$$

where the  $i$ -th element of the pseudo-label  $\mathbf{y}_{DL}$  is defined as follows:

$$\mathbf{y}_{DL}^i = \min(1, 2 \cdot \mathbf{y}_{gt}^i \cdot \sigma(-2 \cdot \mathbf{y}_{gt}^i \cdot \mathbf{y}_b^i)), \quad (6)$$

where  $\mathbf{y}_{gt}^i$  and  $\mathbf{y}_b^i$  are the  $i$ -th element of the ground truth and the output of the biased model respectively. The key point of difference is that unlike [19] that suppresses the output of the biased model with the sigmoid function, we

use  $\mathbf{y}_b$  without using the sigmoid function. In this case, as the value of  $\mathbf{y}_{DL}$  can exceed 1, we additionally clip the value so that the value of  $\mathbf{y}_{DL}$  is bounded in  $[0, 1]$ . We empirically find these simple modifications on the loss function significantly improves the performance. We conjecture the unsuppressed biased output  $\mathbf{y}_b$  allows our target model to better consider the *intensity* of the bias, leading to a more robust target model. In addition, when we train the target model, we empirically find a trend where instead of using the noise inputs as in  $F_{b,G}(\mathbf{z}, \mathbf{q})$ , using the real images as such  $F_b(\mathbf{v}, \mathbf{q})$ , is more effective, hence we use the output of  $F_b(\mathbf{v}, \mathbf{q})$  to train our target model. When the bias model is trained, it is trained with a noise vector to hallucinate the possible biases when only given the question, however, when we give it the real image, we find from Fig. 3 that the output changes more drastically.

## 4. Experiments

In this section, we explain our experimental setting and show our quantitative and qualitative experimental results.

### 4.1. Setup

**Dataset and evaluation metric.** We conduct our experiments within the VQA datasets that are commonly used for diagnosing bias in VQA models. In particular, we test on the VQA-CP2 and VQA-CP1 datasets [3] and the GQA-OOD dataset [26]. For evaluation on all datasets, we take the standard VQA evaluation metric [5]. In addition to this, we also evaluate our method on the VQA-CE [14], which is based on the VQA v2 dataset, and is a newly introduced VQA evaluation protocol for diagnosing how reliant VQA models are on shortcuts.

**Baseline architecture.** Unless otherwise stated, we adopt a popular VQA baseline architecture UpDn [4] as both our ensemble bias model  $F_b$  and our target model  $F$ . We list the details of the generator and discriminator in the supplementary material. During training, we train both the bias model and target model together, then we use the target model only for inference.

### 4.2. Results on VQA-CP2 and VQA-CP1

We compare how our method GenB performs in relation to the recent state-of-the-art methods that focus on bias reduction as shown in Table 1. For VQA-CP2, we first list the *baseline architectures* in the first section. Then, we compare our method to the methods that *modify language modules* (DLR [24], VGQE [35]), *strengthen visual attention* (HINT [42], SCR [48]), *ensemble based methods*, (AReg [41], RUBi [7], LMH [13], CF-VQA [37], GGE [19]) and *balance training data* by changing the training distribution (CVL [1], RandImg [46], SSL [52], CSS [9], Mutant [15], D-VQA [47], KDDAug [10]) in the respective sections. Among the balancing training data

Method	Base	VQA-CP2 test				VQA-CP1 test			
		All	Yes/No	Num	Other	All	Yes/No	Num	Other
SAN [50]	-	24.96	38.35	11.14	21.74	32.50	36.86	12.47	36.22
GVQA [3]	-	31.30	57.99	13.68	22.14	39.23	64.72	11.87	24.86
S-MRL [7]	-	38.46	42.85	12.81	43.20	36.38	42.72	12.59	40.35
UpDn [4]	-	39.94	42.46	11.93	45.09	36.38	42.72	42.14	40.35
<i>Methods based on modifying language modules</i>									
DLR [24]	UpDn	48.87	70.99	18.72	45.57	-	-	-	-
VGQE [35]	UpDn	48.75	-	-	-	-	-	-	-
VGQE [35]	S-MRL	50.11	66.35	27.08	46.77	-	-	-	-
<i>Methods based on strengthening visual attention</i>									
HINT [42]	UpDn	46.73	67.27	10.61	45.88	-	-	-	-
SCR [48]	UpDn	49.45	72.36	10.93	48.02	-	-	-	-
<i>Methods based on ensemble models</i>									
AReg [41]	UpDn	41.17	65.49	15.48	35.48	43.43	74.16	12.44	25.32
RUBi [7]	UpDn	44.23	67.05	17.48	39.61	50.90	80.83	13.84	36.02
LMH [13]	UpDn	52.45	69.81	<b>44.46</b>	45.54	55.27	76.47	<b>26.66</b>	<b>45.68</b>
CF-VQA(SUM) [37]	UpDn	53.55	<b>91.15</b>	13.03	44.97	57.03	<b>89.02</b>	17.08	41.27
CF-VQA(SUM) [37]	S-MRL	55.05	90.61	21.50	45.61	<b>57.39</b>	<b>88.46</b>	14.80	43.61
CF-VQA(SUM) [37] + IntroD [38]	S-MRL	55.17	<b>90.79</b>	17.92	46.73	-	-	-	-
GGE [19]	UpDn	<b>57.32</b>	87.04	27.75	<b>49.59</b>	-	-	-	-
<b>GenB (Ours)</b>	UpDn	<b>59.15</b>	88.03	<b>40.05</b>	<b>49.25</b>	<b>62.74</b>	86.18	<b>43.85</b>	<b>47.03</b>
<i>Methods based on balancing training data</i>									
CVL [1]	UpDn	42.12	45.72	12.45	48.34	-	-	-	-
RandImg [46]	UpDn	55.37	83.89	41.60	44.20	-	-	-	-
SSL [52]	UpDn	57.59	86.53	29.87	50.03	-	-	-	-
CSS [9]	UpDn	58.95	84.37	49.42	48.21	60.95	85.60	40.57	47.03
CSS [9] + IntroD [38]	UpDn	60.17	89.17	46.91	48.62	-	-	-	-
MUTANT [15]	UpDn	61.72	88.90	49.68	50.78	-	-	-	-
D-VQA [47]	UpDn	61.91	88.93	52.32	50.39	-	-	-	-
KDDAug [10]	UpDn	60.24	86.13	55.08	48.08	-	-	-	-

Table 1. Experimental results on VQA-CP2 and VQA-CP1 test set. **Best** and **second best** results are styled in this manner within the column. We do not directly compare to methods that change the distribution of the training data as it does not go with the philosophy of VQA-CP [37] (which is to test whether VQA models rely on prior training data [3]). Among the compared baselines, our method GenB shows the best performance by a noticeable margin.

methods, while some methods swap the image and questions from the supposed pairs [46, 47, 52], counterfactuals based methods generate counterfactual samples masking critical words or objects [9] or by using an external inpainting network to create a new subset of data [15]. In addition, as IntroD [38] is seen as a technique built on top of certain previous methods, we include the scores in their respective categories. Our model (GenB) is in line with the ensemble models listed. Following previous works [37], we do not compare to methods that change training distribution as these methods conflict with the original intent of VQA-CP (which is to test whether VQA models rely on prior training data [3]) and are listed in the bottom of Table 1.

In Table 1, our method achieves state-of-the-art per-

formance on VQA-CP2, surpassing the second best (GGE [19]) by 1.83%. The performance of our model on all three categories (“Yes/No,” “Num,” “Other”) are within top-3 consistently for the same backbone architecture. Our method also performs highly favorably on “Other”.

We also show how our method performs on the VQA-CP1 dataset, which is a subset of the VQA-CP2 dataset. Note that not all of the baselines are listed as we only list the scores that are available in the respective papers. Our method also shows the state-of-the-art results on this dataset with a significant performance improvement over the second best among the methods compared, CF-VQA(SUM) [37]. Even when compared to the available method CSS [8], that we do not compare as it is considered

Method	GQA-OOD Test			
	All	Tail	Head	Avg
UpDn [4]	46.87	42.13	49.16	45.65
RUBi [7]	45.85	43.37	47.37	45.37
LMH [13]	43.96	40.73	45.93	43.33
CSS [9]	44.24	41.20	46.11	43.66
<b>GenB (Ours)</b>	<b>49.43</b>	<b>45.63</b>	<b>51.76</b>	<b>48.70</b>

Table 2. Experimental results on the GQA-OOD dataset. Our method shows the best performance in all the metrics compared to the state-of-the-art methods by a significant margin. The results show that our method robust on GQA-OOD as well.

Method	VQA-CE		
	Overall	Counter	Easy
UpDn [4]	63.52	33.91	76.69
RUBi [7]	61.88	32.25	75.03
LMH [13]	61.15	34.26	73.13
RandImg [46]	63.34	34.41	76.21
CSS [9]	53.55	34.36	62.08
<b>GenB (Ours)</b>	<b>57.87</b>	<b>34.80</b>	<b>68.15</b>

Table 3. Evaluation on the VQA-CE protocol. Ours shows the best performance in counterexamples (listed as Counter) which is the main scope of the VQA-CE.

unfair according to [37], our method shows a significant gap in performance. Compared to the second best method, CF-VQA(Sum) [37] on UpDn, our method improves the overall performance by 5.60% while also having the best performance in both “Num” and “Other” category, by 3.28% and 2.41% performance improvements respectively.

### 4.3. Results on GQA-OOD

Recently, a new dataset for VQA debiasing, the GQA-OOD [26] dataset, has been released and to test the debiasing ability of our method, we show our results in Table 2. We compare our method to available recent state-of-the-art ensemble based methods RUBi [7], LMH [13], and CSS [9]. Our method shows the best performance in all metrics compared to the state-of-the-art methods by a significant margin. Even when compared to methods that show similar performance to GenB in VQA-CP2 like CSS, GenB significantly outperforms it in GQA-OOD by 5.19% in Overall. Interestingly, although all of the listed previous methods outperform the base UpDn model in other datasets, they show a performance degradation on GQA-OOD. Unlike these methods, our method GenB is able to increase performance on GQA-OOD, showing the robustness of GenB.

### 4.4. Results on VQA-CE

We further evaluate our method on the newly introduced evaluation protocol that measures how much a VQA

Training Loss	Bias Model	VQA-CP2 test			
		All	Yes/No	Num	Other
BCE	UpDn	39.94	42.46	11.93	45.09
BCE	GenB	56.98	88.82	19.39	<b>49.86</b>
BCE + DSC	GenB	56.54	<b>89.06</b>	21.29	49.79
BCE + Distill	GenB	57.06	88.91	23.24	49.65
BCE + DSC + Distill	GenB	<b>59.15</b>	88.03	<b>40.05</b>	49.25

Table 4. We ablate the different losses we use to train the GenB bias model. All inferences scores are based on the target model except the first row. BCE loss is Eq. (1), which is the ground truth VQA loss. DSC refers to the discriminator loss Eq. (4) and Distill refers to the KL Divergence loss Eq. (3). Although the DSC and Distill losses independently do not show large improvement, our final model with all losses show a large margin of improvement.

model depends on shortcuts called VQA-CounterExamples (VQA-CE) [14] in Table 3. This dataset is based on the VQA v2 dataset and lists three categories: Overall, Easy, and Counter; which are the total score on VQA v2 validation set, the subset of samples where the shortcuts of image/question give the correct answer, and the subset of samples where the shortcuts of image/question give incorrect answers respectively. Although the overall score for VQA-CE is dependent on the VQA v2 performance, we show that our model shows the best performance on Counter, which is the main point of interest in this dataset.

### 4.5. Ablation Studies

Our method (GenB) includes several different components as shown from Sec. 3.3 to Sec. 3.5. To understand the effects of each component, we run an ablation study on the VQA-CP2 dataset. For all experiments, the results are of the target model and as the purpose of the bias model is to capture bias instead of correctly predicting the answers, we do not consider the predictions of the bias model. For all our ablation tables, we also add the UpDn baseline in the first row, the model in which our target model and bias model is based on for comparison. To further understand whether GenB can be applied to other networks, we further include an ablation study of GenB on other VQA architectures.

**Bias model training analysis.** We ablate on the different losses we use to train the bias model from the target model and list our findings in Table 4 with BCE denoting the ground truth VQA loss Eq. (1), DSC denoting the discriminator loss Eq. (4), and Distill denoting the KL Divergence distillation loss Eq. (3). GenB trained on the BCE loss is already significantly helpful in debiasing the target model. Adding the DSC and Distill losses individually show similar performance to the GenB trained with BCE loss. However, by doubling down on the two losses, DSC and Distill, the model is better able to capture the biases held within the model, hence giving the significant performance boost. In addition, as the bias model and target model are trained

Ensemble Debias Loss	Bias Model	VQA-CP2 test			
		All	Yes/No	Num	Other
-	UpDn	39.94	42.46	11.93	45.09
GGE [19]	UpDn	47.40	64.45	13.96	47.64
Our Loss	UpDn	52.47	88.20	30.09	40.38
RUBi [7]	GenB	30.77	72.78	12.15	13.87
LMH [13]	GenB	53.99	75.89	<b>44.62</b>	45.08
GGE [19]	GenB	49.51	70.63	14.08	48.16
Ours Loss	GenB	<b>59.15</b>	<b>88.03</b>	40.05	<b>49.25</b>

Table 5. Ablation of ensemble debiasing loss functions and our debiasing loss function Eq. (5). Note that our loss improves the score of using the vanilla UpDn model when compared to GGE by a large margin. Note that GenB works best with our loss and shows a large performance improvement from GGE + GenB.

Bias Model	VQA-CP2 test			
	All	Yes/No	Num	Other
UpDn	39.94	42.46	11.93	45.09
UpDn	52.47	88.20	30.09	40.38
Visual-Answer	41.03	42.69	12.66	47.93
Question-Answer	56.68	<b>89.30</b>	20.78	<b>49.43</b>
GenB Visual	49.54	72.05	12.58	47.89
<b>GenB Question (Ours)</b>	<b>59.15</b>	88.03	<b>40.05</b>	49.25

Table 6. Ablation of different bias models with our modified loss. We fix the target model architecture to be the UpDn model and show the performance of the target model by changing the bias model architecture between UpDn, Question-Answer model, Visual-Answer model, our GenB model but with image and generated noise, and our GenB model, which has question and generated noise. We show that our loss works with other networks, and although other networks can be helpful in debiasing, our model, GenB with question and noise is the best performing model by far.

concurrently, without a ground truth anchor (BCE loss), the model performs extremely poorly, and we conjecture that the bias model struggles to learn any meaningful bias. Note that BCE + DSC shows the best Yes/No score while our final model shows the best score in Num score. We conjecture that the bias model trained with our adversarial training framework is good at modeling the bias for simple answers like Yes/No while the bias with more complex answers like counting is hard to learn only with the adversarial training. We conclude that combining our knowledge distillation loss to the adversarial training is essential to learn the bias with complex answers including the Num category.

**Debiasing loss function analysis.** We list in Sec. 3.5 the different losses we can use for debiasing and we show our results in Table 5. We also list our new debiasing loss function Eq. (5) and how it performs when a vanilla UpDn model is used instead of the GenB.

We analyze how GGE’s loss function fairs when the vanilla UpDn model is used as a bias model. In the second and third row, although the GGE loss improves the score

Architecture	VQA-CP2 test				$\Delta$ Gap
	All	Yes/No	Num	Other	
UpDn [4]	39.94	42.46	11.93	45.09	
UpDn [4] + GenB	<b>59.15</b>	<b>88.03</b>	<b>40.05</b>	<b>49.25</b>	<b>+19.21</b>
BAN <sup>†</sup> [34]	37.35	41.96	12.08	41.71	
BAN <sup>†</sup> [34] + GenB	<b>57.37</b>	<b>89.11</b>	<b>29.52</b>	<b>48.37</b>	<b>+20.02</b>
SAN <sup>†</sup> [50]	38.65	40.59	12.98	44.67	
SAN <sup>†</sup> [50] + GenB	<b>56.72</b>	<b>88.84</b>	<b>19.04</b>	<b>50.22</b>	<b>+18.07</b>
LXMERT [45]	46.23	42.84	18.91	55.51	
LXMERT [45] + GenB (Ours Best)	<b>71.16</b>	<b>92.24</b>	<b>64.71</b>	<b>61.89</b>	<b>+24.93</b>
Reported LXMERT Performance					
LXMERT [45] + MUTANT [15]	69.52	93.15	67.17	57.78	
LXMERT [45] + D-VQA [47]	69.75	80.43	58.57	67.23	
LXMERT [45] + SAR [43]	62.12	85.14	41.63	55.68	

Table 7. Ablation of GenB on different architectures. <sup>†</sup> signifies our re-implementation. We find that adding GenB to other backbones show consistent improvement, showing that GenB is applicable on other baselines. Note that we include the reported LXMERT based performance in the last 4 rows. LXMERT + GenB is state-of-the-art among all reported VQA-CP2 [15, 43, 47].

from the baseline by introducing a bias model, our loss exceeds the performance by 5.07%. In the last two sections, we show the debiasing losses with the full GenB model with all training losses and other debiasing losses. Compared to GGE with GenB, our loss with GenB shows a 9.64% performance gap. Given the GGE loss, the performance increase from the bias model is only 2.11%, whereas our debiasing loss function is able to boost the score from the vanilla UpDn to GenB by 6.68%. We conjecture that as our loss function accounts for the amount of bias a model captures, and the generative bias model is able to stochastically model the bias, our loss function captures the bias in a much more dynamic and helpful manner.

**Effect of different bias models.** We run an ablation experiment to understand the effects of different bias models and the addition of the Generator  $G(\cdot)$  and its effects on the performance of the target model in Table 6. We fix the target model as the UpDn model and show the performance of the target model by changing the bias model. Note that the results are from the predictions of the target model, not the bias model, and the top row is the UpDn baseline. We include the UpDn model as a bias model, the Visual-Answer model, which is a model that takes an image in as its input and predicts the answer, the Question-Answer model, which takes in a question and predicts the answer, GenB Visual model, which takes in an image together with noise to predict the answer, and our model GenB Question, which is the model we have shown up until this point. These models are used to debias the target model with our modified debiasing loss function. We show that even adding noise to the image can be helpful, however, we find that our final model, the model with question and noise, outperforms other methods significantly.

**Effect on different backbones.** We further experiment




Figure 3. Qualitative results showing the predictions and attention scores of our target model and the bias model. We show the attention of the bias model with different noise to see where it would attend to. We also input the real image into the bias model to visualize the attention. The bias model’s attention changes each time and does not attend to the critical areas of the image.

on the robustness of our method GenB and its applicability to other VQA backbones (SAN [50], BAN [34], and LXMERT [45]) as shown in Table 7. For each architecture, we set both the target model and GenB model to be the same architecture. We show that GenB consistently shows significant performance improvements on all backbones, signifying that GenB is not simply limited to the UpDn setting. Note that LXMERT + GenB shows the highest LXMERT performance as reported in [15, 43, 47], which is the absolute state-of-the-art VQA-CP2 score so far. We believe this opens up a possibility for future works where different backbones can be mixed and matched for further debiasing.

#### 4.6. Qualitative Results

We visualize the attention and predictions of our target and bias models in Fig. 3. We run the bias model three times with different noise to show the varying attention and biased predictions. We also input the corresponding image with the question to see where the model would attend, and find that the attention is once again random as the image is previously unseen. As expected, the bias model’s predictions change each time. Even with the same question type (e.g., what color), we find the biased model’s predictions are noticeably different. We find the target model is able

to correctly attend to the salient region while predicting the correct answer by virtue of the bias model.

## 5. Conclusion

In this paper, we started with the intuition “the better the bias model, the better we can debias the target model. Then how can we best model the bias?” In response, we present a simple, effective, and novel generative bias model that we call GenB. We use this generative model to learn the bias that may be inhibited by both the distribution and target model with the aid of generative networks, adversarial training, and knowledge distillation. In addition, in conjunction with our modified loss function, our novel bias model is able to debias our target model, and our target model achieves state-of-the-art performance on various bias diagnosing datasets and architectures.

## Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2020-0-01373, Artificial Intelligence Graduate School Program(Hanyang University))



## References

- [1] Ehsan Abbasnejad, Damien Teney, Amin Parvaneh, Javen Shi, and Anton van den Hengel. Counterfactual vision and language learning. In *CVPR*, 2020. 2, 4, 5
- [2] Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. Analyzing the behavior of visual question answering models. In *EMNLP*, 2016. 1, 2
- [3] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Anirudha Kembhavi. Don't just assume; look and answer: Overcoming priors for visual question answering. In *CVPR*, 2018. 2, 4, 5
- [4] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018. 3, 4, 5, 6, 7
- [5] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *ICCV*, 2015. 1, 2, 4
- [6] Dawit Mureja Argaw, Fabian Caba Heilbron, Joon-Young Lee, Markus Woodson, and In So Kweon. The anatomy of video editing: A dataset and benchmark suite for ai-assisted video editing. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VIII*, pages 201–218. Springer, 2022. 2
- [7] Remi Cadene, Corentin Dancette, Hedi Ben-younes, Matthieu Cord, and Devi Parikh. Rubi: Reducing unimodal biases in visual question answering. In *NeurIPS*, 2019. 1, 2, 3, 4, 5, 6, 7
- [8] David L Chen and William B Dolan. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies—Volume 1*. ACL, 2011. 5
- [9] Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. Counterfactual samples synthesizing for robust visual question answering. In *CVPR*, 2020. 2, 4, 5, 6
- [10] Long Chen, Yuhang Zheng, and Jun Xiao. Rethinking data augmentation for robust visual question answering. *arXiv preprint arXiv:2207.08739*, 2022. 2, 4, 5
- [11] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, 2020. 2
- [12] Jae Won Cho, Dong-Jin Kim, Jinsoo Choi, Yunjae Jung, and In So Kweon. Dealing with missing modalities in the visual question answer-difference prediction task through knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 4
- [13] Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. Don't take the easy way out: Ensemble based methods for avoiding known dataset biases. In *EMNLP*, 2019. 1, 2, 3, 4, 5, 6, 7
- [14] Corentin Dancette, Remi Cadene, Damien Teney, and Matthieu Cord. Beyond question-based biases: Assessing multimodal shortcut learning in visual question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1574–1583, 2021. 2, 4, 6
- [15] Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. Mutant: A training paradigm for out-of-distribution generalization in visual question answering. In *EMNLP*, 2020. 2, 4, 5, 7, 8
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 2, 3
- [17] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *CVPR*, 2017. 1
- [18] Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *CVPR*, 2018. 2
- [19] Xinzhe Han, Shuhui Wang, Chi Su, Qingming Huang, and Qi Tian. Greedy gradient ensemble for robust visual question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1584–1593, 2021. 1, 2, 3, 4, 5, 7
- [20] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2, 4
- [21] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 4
- [22] Youngjoon Jang, Youngtaek Oh, Jae Won Cho, Dong-Jin Kim, Joon Son Chung, and In So Kweon. Signing outside the studio: Benchmarking background robustness for continuous sign language recognition. In *BMVC*, 2022. 2
- [23] Youngjoon Jang, Youngtaek Oh, Jae Won Cho, Myungchul Kim, Dong-Jin Kim, In So Kweon, and Joon Son Chung. Self-sufficient framework for continuous sign language recognition. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023. 2
- [24] Chenchen Jing, Yuwei Wu, Xiaoxun Zhang, Yunde Jia, and Qi Wu. Overcoming language priors in VQA via decomposed linguistic representations. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7–12, 2020*, pages 11181–11188. AAAI Press, 2020. 4, 5
- [25] 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 *CVPR*, 2017. 2
- [26] Corentin Kervadec, Grigory Antipov, Moez Baccouche, and Christian Wolf. Roses are red, violets are blue... but should vqa expect them to? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2776–2785, 2021. 2, 4, 6

- [27] Dong-Jin Kim, Jae Won Cho, Jinsoo Choi, Yunjae Jung, and In So Kweon. Single-modal entropy based active learning for visual question answering. In *BMVC*, 2021. 4
- [28] Dong-Jin Kim, Jinsoo Choi, Tae-Hyun Oh, and In So Kweon. Dense relational captioning: Triple-stream networks for relationship-based captioning. In *CVPR*, 2019. 2
- [29] Dong-Jin Kim, Jinsoo Choi, Tae-Hyun Oh, and In So Kweon. Image captioning with very scarce supervised data: Adversarial semi-supervised learning approach. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2019. 3
- [30] Dong-Jin Kim, Jinsoo Choi, Tae-Hyun Oh, Youngjin Yoon, and In So Kweon. Disjoint multi-task learning between heterogeneous human-centric tasks. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2018. 4
- [31] Dong-Jin Kim, Tae-Hyun Oh, Jinsoo Choi, and In So Kweon. Dense relational image captioning via multi-task triple-stream networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):7348–7362, 2021. 2
- [32] Dong-Jin Kim, Xiao Sun, Jinsoo Choi, Stephen Lin, and In So Kweon. Detecting human-object interactions with action co-occurrence priors. In *ECCV*, 2020. 2, 4
- [33] Dong-Jin Kim, Xiao Sun, Jinsoo Choi, Stephen Lin, and In So Kweon. Acp++: Action co-occurrence priors for human-object interaction detection. *IEEE TIP*, 30:9150–9163, 2021. 2, 4
- [34] Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. Bilinear attention networks. In *NeurIPS*, 2018. 2, 7, 8
- [35] Gouthaman KV and Anurag Mittal. Reducing language biases in visual question answering with visually-grounded question encoder. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, pages 18–34. Springer, 2020. 4, 5
- [36] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vlb: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *arXiv preprint arXiv:1908.02265*, 2019. 2
- [37] Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. Counterfactual vqa: A cause-effect look at language bias. In *CVPR*, 2021. 1, 2, 3, 4, 5, 6
- [38] Yulei Niu and Hanwang Zhang. Introspective distillation for robust question answering. *Advances in Neural Information Processing Systems*, 34:16292–16304, 2021. 2, 5
- [39] Youngtaek Oh, Dong-Jin Kim, and In So Kweon. Daso: Distribution-aware semantics-oriented pseudo-label for imbalanced semi-supervised learning. In *CVPR*, 2022. 2
- [40] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014. 2
- [41] Sainandan Ramakrishnan, Aishwarya Agrawal, and Stefan Lee. Overcoming language priors in visual question answering with adversarial regularization. In *NeurIPS*, 2018. 1, 4, 5
- [42] 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 *ICCV*, 2019. 2, 4, 5
- [43] Qingyi Si, Zheng Lin, Ming yu Zheng, Peng Fu, and Weiping Wang. Check it again: Progressive visual question answering via visual entailment. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4101–4110, 2021. 2, 7, 8
- [44] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019. 2
- [45] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019. 2, 7, 8
- [46] Damien Teney, Kushal Kafle, Robik Shrestha, Ehsan Abbasnejad, Christopher Kanan, and Anton van den Hengel. On the value of out-of-distribution testing: An example of goodhart’s law. *arXiv preprint arXiv:2005.09241*, 2020. 2, 4, 5, 6
- [47] Zhiqian Wen, Guanghui Xu, Minghui Tan, Qingyao Wu, and Qi Wu. Debaised visual question answering from feature and sample perspectives. *Advances in Neural Information Processing Systems*, 34, 2021. 2, 4, 5, 7, 8
- [48] Jialin Wu and Raymond J Mooney. Self-critical reasoning for robust visual question answering. In *NeurIPS*, 2019. 2, 4, 5
- [49] Yike Wu, Yu Zhao, Shiwan Zhao, Ying Zhang, Xiaojie Yuan, Guoqing Zhao, and Ning Jiang. Overcoming language priors in visual question answering via distinguishing superficially similar instances. *arXiv preprint arXiv:2209.08529*, 2022. 2
- [50] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *CVPR*, 2016. 5, 7, 8
- [51] Peng Zhang, Yash Goyal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Yin and yang: Balancing and answering binary visual questions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5014–5022, 2016. 1
- [52] Xi Zhu, Zhendong Mao, Chunxiao Liu, Peng Zhang, Bin Wang, and Yongdong Zhang. Overcoming language priors with self-supervised learning for visual question answering. *IJCAI*, 2021. 2, 4, 5