

Counterfactual VQA: A Cause-Effect Look at Language Bias

Yulei Niu¹ Kaihua Tang¹ Hanwang Zhang¹ Zhiwu Lu^{2,3} Xian-Sheng Hua⁴ Ji-Rong Wen^{2,3}

¹Nanyang Technological University, Singapore ²Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China

³Beijing Key Laboratory of Big Data Management and Analysis Methods ⁴Damo Academy, Alibaba Group

yn.yuleiniu@gmail.com, {kaihua001@e., hanwangzhang}@ntu.edu.sg, {luzhiwu, jrwen}@ruc.edu.cn

Abstract

VQA models may tend to rely on language bias as a shortcut and thus fail to sufficiently learn the multi-modal knowledge from both vision and language. Recent debiasing methods proposed to exclude the language prior during inference. However, they fail to disentangle the “good” language context and “bad” language bias from the whole. In this paper, we investigate how to mitigate language bias in VQA. Motivated by causal effects, we proposed a novel counterfactual inference framework, which enables us to capture the language bias as the direct causal effect of questions on answers and reduce the language bias by subtracting the direct language effect from the total causal effect. Experiments demonstrate that our proposed counterfactual inference framework 1) is general to various VQA backbones and fusion strategies, 2) achieves competitive performance on the language-bias sensitive VQA-CP dataset while performs robustly on the balanced VQA v2 dataset without any augmented data. The code is available at <https://github.com/yuleiniu/cfvqa>.

1. Introduction

Visual Question Answering (VQA) [8, 4] has become the fundamental building block that underpins many frontier interactive AI systems, such as visual dialog [15], vision-and-language navigation [6], and visual commonsense reasoning [51]. VQA systems are required to perform visual analysis, language understanding, and multi-modal reasoning. Recent studies [19, 3, 8, 19, 25] found that VQA models may rely on spurious linguistic correlations rather than multi-modal reasoning. For instance, simply answering “tennis” to the sport-related questions and “yes” to the questions “Do you see a ...” can achieve approximately 40% and 90% accuracy on the VQA v1.0 dataset. As a result, VQA models will fail to generalize well if they simply memorize the strong language priors in the training data [2, 19], especially on the recently proposed VQA-CP [3] dataset where the priors are quite different in the training and test sets.

One straightforward solution to mitigate language bias is to enhance the training data by using extra annotations or data augmentation. In particular, visual [14] and textual [23] explanations are used to improve the visual grounding ability [37, 44]. Besides, counterfactual training samples generation [11, 1, 55, 18, 28] helps to balance the training data, and outperform other debiasing methods by large margins on VQA-CP. These methods demonstrate the effect of *debaised training* to improve the generalizability of VQA models. However, it is worth noting that VQA-CP was proposed to validate whether VQA models can disentangle the learned visual knowledge and memorized language priors [3]. Therefore, how to make unbiased inference under *biased training* still remains a major challenge in VQA. Another popular solution [10, 13] is using a separate question-only branch to learn the language prior in the training set. During the test stage, the prior is mitigated by excluding the extra branch. However, we argue that the language prior consists of both “bad” language bias (e.g., binding the color of bananas with the major color “yellow”) and “good” language context (e.g., narrowing the answer space based on the question type “what color”). Simply excluding the extra branch cannot make use of the good context. Indeed, it is still challenging for recent debiasing VQA methods to disentangle the good and bad from the whole.

Motivated by counterfactual reasoning and causal effects [29, 30, 31], we propose a novel counterfactual inference framework called CF-VQA to reduce language bias in VQA. Overall, we formulate language bias as the direct causal effect of questions on answers, and mitigate the bias by subtracting the direct language effect from the total causal effect. As illustrated in Figure 1, we introduce two scenarios, conventional VQA and counterfactual VQA, to estimate the total causal effect and direct language effect, respectively. These two scenarios are defined as follows:

Conventional VQA: *What will answer A be, if machine hears question Q, sees image V, and extracts the multi-modal knowledge K?*

Counterfactual VQA: *What would A be, if machine hears Q, but had not extracted K or seen V?*

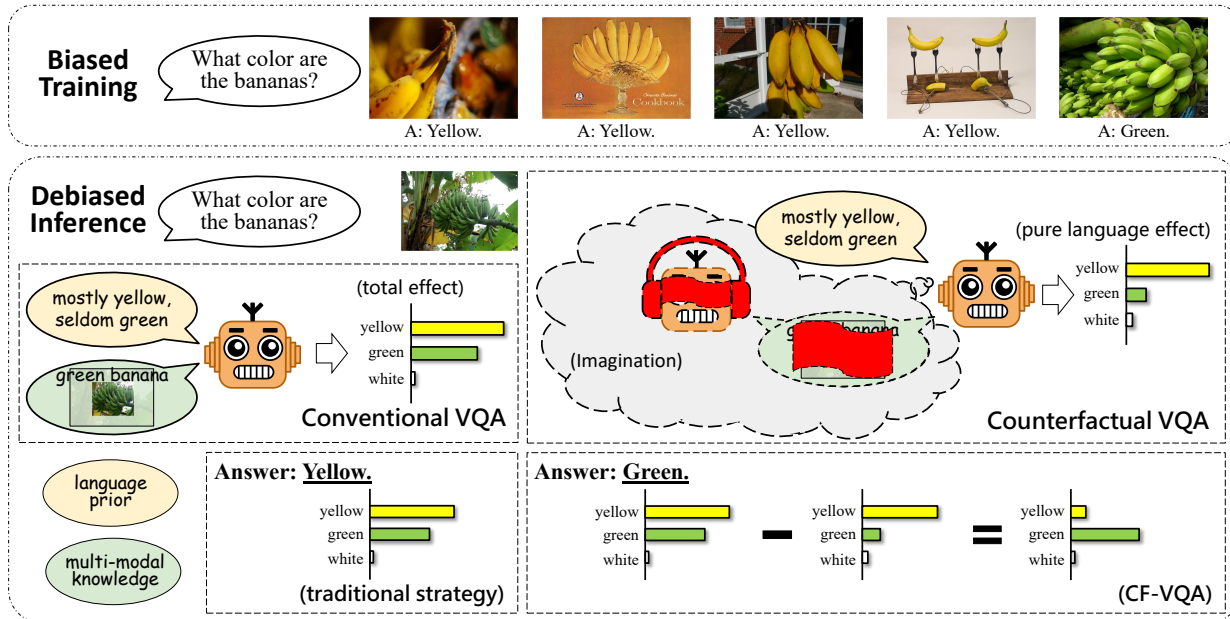


Figure 1: Our cause-effect look at language bias in VQA. Conventional VQA depicts the fact where machine hears the question and extracts the multi-modal knowledge. Counterfactual VQA depicts the scenario where machine hears the question but the knowledge is blocked. We subtract the pure language effect from the total effect for debaised inference.

Intuitively, conventional VQA depicts the scenario where both Q and V are available. In this case, we can estimate the total causal effect of V and Q on A . However, conventional VQA cannot disentangle the single-modal linguistic correlation and multi-modal reasoning, *i.e.*, direct and indirect effects. Therefore, we consider the following counterfactual question: “What would have happened if the machine had not performed multi-modal reasoning?” The answer to this question can be obtained by imagining a scenario where the machine hears Q , but the multi-modal knowledge K is intervened under the no-treatment condition, *i.e.*, V and Q had not been accessible. Since the response of K to Q is blocked, VQA models can only rely on the single-modal impact. Therefore, language bias can be identified by estimating the direct causal effect of Q on A , *i.e.*, pure language effect. The training stages follows language-prior based methods [10, 13] that train an ensemble model with a prevailing VQA model and single-modal branches. During the test stage, CF-VQA uses the debaised causal effect for inference, which is obtained by subtracting the pure language effect from the total effect. Perhaps surprisingly, recent language-prior based methods [10, 13] can be further unified into our proposed counterfactual inference framework as special cases. In particular, CV-VQA can easily improve RUBi [10] by 7.5% with only one more learnable parameter. Experimental results show that CF-VQA outperforms the methods without data argumentation by large margins on the VQA-CP dataset [3] while remaining stable on the balanced VQA v2 dataset [19].

The main contribution of this paper is threefold. First, our counterfactual inference framework is the first to formulate language bias in VQA as causal effects. Second, we provide a novel causality-based interpretation for recent debiasing VQA works [10, 13]. Third, our cause-effect look is general and suitable for different baseline VQA architectures and fusion strategies.

2. Related Work

Language Bias in VQA can be interpreted in two ways. First, there exists strong correlations between questions and answers, which reflects the “*language prior*” [19, 3]. Simply answering “tennis” to the sport-related questions can achieve approximately 40% accuracy on the VQA v1.0 dataset. Second, the questioner tends to ask about the objects seen in the image, which leads to the “*visual priming bias*” [8, 19, 25]. Simply answering “yes” to all the questions “Do you see a ...” achieves nearly 90% accuracy on the VQA v1.0 dataset. In both ways, machines may merely focus on the question rather than the visual content. This serious shortcut limits the generalization of VQA models [2, 53, 19], especially when the test scenario is quite different from the training scenario.

Debiasing Strategies in VQA. Recently, a new VQA split, Visual Question Answering under Changing Priors (VQA-CP) [3], was proposed to evaluate the generalizability of VQA models. In VQA-CP, the distributions of answers for every question type are different during training and

test stages. Most of recent solutions to reduce the language bias in VQA can be grouped into three categories, strengthening visual grounding [37, 44], weakening language prior [33, 10, 13], and implicit/explicit data argumentation [11, 1, 40, 55]. First, human visual [14] and textual [23] explanations are exploited to strengthen the visual grounding in VQA [37, 44]. Second, ensemble-based methods proposed to use a separated QA branch to capture the language prior under adversarial learning [33] or multi-task learning [10, 13]. Third, recent works [11, 1] automatically generate additional question-image pairs to balance the distribution of training data. In this paper, the language-prior based methods [10, 13] can be unified into our proposed counterfactual inference framework as special cases.

Causality-inspired Computer Vision. Counterfactual thinking and causal inference have inspired several studies in computer vision, including visual explanations [20, 7, 42, 48], scene graph generation [12, 38], image recognition [38], video analysis [16, 26], zero-shot and few-shot learning [50, 49] incremental learning [22], representation learning [43, 54], semantic segmentation [52], and vision-language tasks [11, 39, 32, 45, 17, 46]. Especially, counterfactual learning has been exploited in recent VQA studies [11, 39, 1]. Different from these works that generate counterfactual samples for *debiased* training, our cause-effect look focuses on counterfactual inference with even *biased* training data.

3. Preliminaries

In this section, we introduce the used concepts of causal inference [29, 30, 35, 31]. In the following, we represent a random variable as a capital letter (*e.g.*, X), and denote its observed value as a lowercase letter (*e.g.*, x).

Causal graph reflects the causal relationships between variables, which is represented as a directed acyclic graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where \mathcal{V} denotes the set of variables and \mathcal{E} represents the cause-and-effect relationships. Figure 2(a) shows an example of causal graph that consists of three variables. If the variable X has a *direct* effect on the variable Y , we say that Y is the child of X , *i.e.*, $X \rightarrow Y$. If X has an *indirect* effect on Y via the variable M , we say that M acts as a *mediator* between X and Y , *i.e.*, $X \rightarrow M \rightarrow Y$.

Counterfactual notations are used to translate causal assumptions from graphs to formulas. The value that Y would obtain if X is set to x and M is set to m is denoted as:

$$Y_{x,m} = Y(X = x, M = m).^1 \quad (1)$$

In the factual scenario, we have $m = M_x = M(X = x)$. In the counterfactual scenario, X is set as different values for M and Y . For example, $Y_{x,M_{x^*}}$ describes the situation

¹If there is no confounder of X , then we have that $do(X = x)$ is equivalent to $X = x$ and can omit the *do* operator.

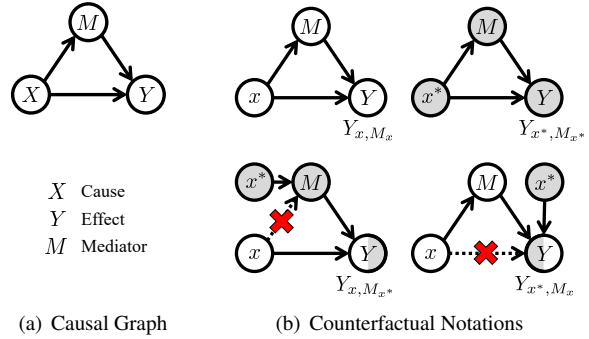


Figure 2: (a) Example of causal graph. (b) Examples of counterfactual notations. White nodes are at the value $X = x$ while gray nodes are at the value $X = x^*$.

where X is set to x and M is set to the value when X had been x^* , *i.e.*, $Y_{x,M_{x^*}} = Y(X = x, M = M(X = x^*))$. Note that X can be simultaneously set to different values x and x^* only in the counterfactual world. Figure 2(b) illustrates examples of counterfactual notations.

Causal effects reflect the comparisons between two potential outcomes of the same individual given two different treatments [36, 34]. Supposed that $X = x$ represents “under treatment condition” and $X = x^*$ represents “under no-treatment condition”². The total effect (TE) of treatment $X = x$ on Y compares two hypothetical situations $X = x$ and $X = x^*$, which is denoted as:

$$TE = Y_{x,M_x} - Y_{x^*,M_{x^*}}. \quad (2)$$

Total effect can be decomposed into natural direct effect (NDE) and total indirect effect (TIE). NDE denotes the effect of X on Y with the mediator M blocked, and expresses the increase in Y with X changing from x^* to x , while M is set to the value it *would have obtained* at $X = x^*$, *i.e.*, the response of M to the treatment $X = x$ is disabled:

$$NDE = Y_{x,M_{x^*}} - Y_{x^*,M_{x^*}}. \quad (3)$$

TIE is the difference between TE and NDE, denoted as:

$$TIE = TE - NDE = Y_{x,M_x} - Y_{x,M_{x^*}}. \quad (4)$$

TE can be also decomposed into natural indirect effect (NIE) and total direct effect (TDE). Similarly, NIE reflects the effect of X on Y through the mediator M , *i.e.*, $X \rightarrow M \rightarrow Y$, while the direct effect on $X \rightarrow Y$ is blocked by setting X as x^* . NIE is denoted as:

$$NIE = Y_{x^*,M_x} - Y_{x^*,M_{x^*}}. \quad (5)$$

In Section 4, we will further discuss the meanings and differences of these effects in VQA.

²For example, to estimate the effect of a drug on a disease, $X = x$ represents taking the drug, while $X = x^*$ represents not taking the drug.

4. Cause-Effect Look at VQA

Following the common formulation, we define the VQA task as a multi-class classification problem. VQA models are required to select an answer from the candidate set $\mathcal{A} = \{a\}$ given an image $V = v$ and a question $Q = q$.

4.1. Cause-Effect Look

The causal graph of VQA is illustrated in Figure 3(a). The effect of V and Q on A can be divided into the single-modal impact and multi-modal impact. The single-modal impact captures the *direct* effect of V or Q on A via $V \rightarrow A$ or $Q \rightarrow A$. The multi-modal impact captures the *indirect* effect of V and Q on A via the multi-modal knowledge K , i.e., $V, Q \rightarrow K \rightarrow A$. We propose to exclude pure language effect on $Q \rightarrow A$ to reduce language bias in VQA.

Following the counterfactual notations in Eq. (1), we denote the score that the answer a (e.g., “green”) would obtain if V is set to v (e.g., an image which includes green bananas) and Q is set to q (e.g., a question “What color are the bananas?”) as

$$Y_{v,q}(a) = Y(a; V=v, Q=q).$$

Without loss of generality, we omit a for simplicity, i.e., $Y_{v,q} = Y(V=v, Q=q)$. Similarly, the counterfactual notation of K is denoted as $K_{v,q} = K(V=v, Q=q)$.

As shown in Figure 3(a), there exist three paths directly connected to A , i.e., $Q \rightarrow A$, $V \rightarrow A$, and $K \rightarrow A$. Therefore, we rewrite $Y_{v,q}$ as the function of Q , V and K :

$$Y_{v,q} = Z_{q,v,k} = Z(Q=q, V=v, K=k), \quad (6)$$

where $k = K_{v,q}$. Following the definition of causal effects in Section 3, the *total effect* (TE) of $V = v$ and $Q = q$ on $A = a$ can be written as:

$$TE = Y_{v,q} - Y_{v^*,q^*} = Z_{q,v,k} - Z_{q^*,v^*,k^*}, \quad (7)$$

where $k^* = K_{v^*,q^*}$. Here v^* and q^* represent the no-treatment condition where v and q are not given.

As we have discussed in Section 1, VQA models may suffer from the spurious correlation between questions and answers, and thus fail to conduct effective multi-modal reasoning. Therefore, we expect VQA models to exclude the direct impact of questions. To achieve this goal, we proposed *counterfactual VQA* to estimate the causal effect of $Q = q$ on $A = a$ by blocking the effect of K and V . Counterfactual VQA describes the scenario where Q is set to q and K would attain the value k^* when Q had been q^* and V had been v^* . Since the response of mediator K to inputs is blocked, the model can only rely on the given question for decision making. Figure 3(b) shows the comparison between conventional VQA and counterfactual VQA. We obtain the *natural direct effect* (NDE) of Q on A by comparing

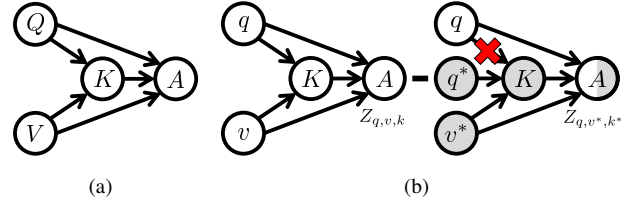


Figure 3: (a) Causal graph for VQA. Q : question. V : image. K : multi-modal knowledge. A : answer. (b) Comparison between conventional VQA (left) and counterfactual VQA (right). White nodes are at the value $V = v$ and $Q = q$ while gray nodes are at the value $V = v^*$ and $Q = q^*$.

counterfactual VQA to the no-treatment conditions:

$$NDE = Z_{q,v^*,k^*} - Z_{q^*,v^*,k^*}. \quad (8)$$

Since the effect of Q on the intermediate K is blocked (i.e., $K = k^*$), NDE explicitly captures the language bias. Furthermore, the reduction of language bias can be realized by subtracting NDE from TE, which is represented as:

$$TIE = TE - NDE = Z_{q,v,k} - Z_{q,v^*,k^*}. \quad (9)$$

We select the answer with the maximum TIE for inference, which is totally different from traditional strategies that is based on the posterior probability i.e., $P(a|v, q)$.

4.2. Implementation

Parameterization. The calculation of the score $Z_{q,v,k}$ in Eq. (6) is parameterized by three neural models \mathcal{F}_Q , \mathcal{F}_V , \mathcal{F}_{VQ} and one fusion function h as:

$$\begin{aligned} Z_q &= \mathcal{F}_Q(q), & Z_v &= \mathcal{F}_V(v), & Z_k &= \mathcal{F}_{VQ}(v, q), \\ Z_{q,v,k} &= h(Z_q, Z_v, Z_k), \end{aligned} \quad (10)$$

where \mathcal{F}_Q is the language-only branch (i.e., $Q \rightarrow A$), \mathcal{F}_V is the vision-only branch (i.e., $V \rightarrow A$), and \mathcal{F}_{VQ} is the vision-language branch (i.e., $V, Q \rightarrow K \rightarrow A$). The output scores are fused by the function h to obtain the final score $Z_{q,v,k}$.

As introduced in Section 4.1, the no-treatment condition is defined as blocking the signal from vision or language, i.e., v or q is not given. We represent the no-treatment conditions as $V = v^* = \emptyset$ and $Q = q^* = \emptyset$. Note that neural models cannot deal with no-treatment conditions where the inputs are void. Therefore, we assume that the model will randomly guess with equal probability under the no-treatment conditions. In this case, Z_q , Z_v and Z_k in Eq. (10) can be represented as:

$$Z_q = \begin{cases} z_q = \mathcal{F}_Q(q) & \text{if } Q = q \\ z_q^* = c & \text{if } Q = \emptyset \end{cases}, \quad (11)$$

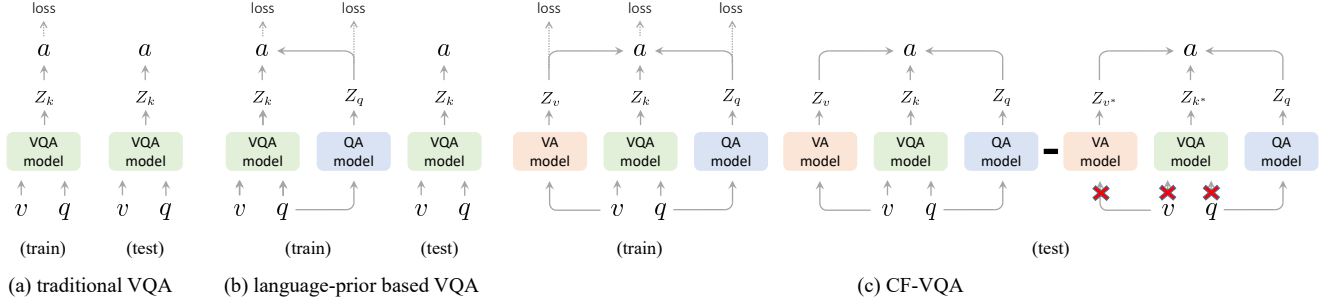


Figure 4: Comparison between our CF-VQA and other VQA methods. (a) Traditional VQA methods use a single VQA model. (b) Language-prior based methods [10, 13] use an additional QA model to capture the language prior during the training stage. The QA model is not used during testing. (c) CF-VQA maintains the QA model during the test stage, and makes inference based on the debiased causal effect. The VA model is optional depending on the causal graph assumption.

$$Z_v = \begin{cases} z_v = \mathcal{F}_V(v) & \text{if } V = v \\ z_v^* = c & \text{if } V = \emptyset \end{cases}, \quad (12)$$

$$Z_k = \begin{cases} z_k = \mathcal{F}_{VQ}(v, q) & \text{if } V = v \text{ and } Q = q \\ z_k^* = c & \text{if } V = \emptyset \text{ or } Q = \emptyset \end{cases}, \quad (13)$$

where c denotes a learnable parameter. We use the uniform distribution assumption for two reasons. Firstly, as for us human, we would like to make a wild guess if we have absolutely no idea about the specific treatments, including question types or topics. Secondly, as z_v^* and z_k^* are used to estimated NDE of Q , the uniform distribution can guarantee a safe estimation. We further empirically validate the assumption in the ablation study.

Fusion Strategies. We expect that the fused score $Z_{q,v,k}$ is a combination of Z_q , Z_v and Z_k . We proposed two non-linear fusion variants, Harmonic (HM) and SUM:

$$\text{(HM)} \quad h(Z_q, Z_v, Z_k) = \log \frac{Z_{\text{HM}}}{1 + Z_{\text{HM}}}, \quad (14)$$

where $Z_{\text{HM}} = \sigma(Z_q) \cdot \sigma(Z_v) \cdot \sigma(Z_k)$.

$$\text{(SUM)} \quad h(Z_q, Z_v, Z_k) = \log \sigma(Z_{\text{SUM}}), \quad (15)$$

where $Z_{\text{SUM}} = Z_q + Z_v + Z_k$.

Training. The training strategy follows [10]. As illustrated in Figure 4(c), given a triplet (v, q, a) where a is the ground-truth answer of image-question pair (v, q) , the branches are optimized by minimizing the cross-entropy losses over the scores $Z_{q,v,k}$, Z_q and Z_v :

$$\mathcal{L}_{cls} = \mathcal{L}_{VQA}(v, q, a) + \mathcal{L}_{QA}(q, a) + \mathcal{L}_{VA}(v, a), \quad (16)$$

where \mathcal{L}_{VQA} , \mathcal{L}_{QA} and \mathcal{L}_{VA} are over $Z_{q,v,k}$, Z_q and Z_v . Note that we introduce a learnable parameter c in Eq. (11)-(13), which controls the sharpness of the distribution of Z_{q,v^*,k^*} like the softmax temperature [21]. We hypothesize that the sharpness of NDE should be similar to that of

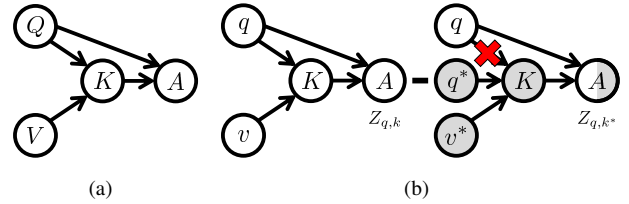


Figure 5: (a) Simplified VQA causal graph. (b) Comparison between conventional VQA and counterfactual VQA.

TE. Otherwise, an improper c would lead to the result that TIE in Eq. (4) is dominated by TE or NDE. Thus, we use Kullback-Leibler divergence to estimate c :

$$\mathcal{L}_{kl} = \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} -p(a|q, v, k) \log p(a|q, v^*, k^*), \quad (17)$$

where $p(a|q, v, k) = \text{softmax}(Z_{q,v,k})$ and $p(a|q, v^*, k^*) = \text{softmax}(Z_{q,v^*,k^*})$. Only c is updated when minimizing \mathcal{L}_{kl} . The final loss is the combination of \mathcal{L}_{cls} and \mathcal{L}_{kl} :

$$\mathcal{L} = \sum_{(v,q,a) \in \mathcal{D}} \mathcal{L}_{cls} + \mathcal{L}_{kl} \quad (18)$$

Inference. As discussed in Section 4.1, we use the debiased causal effect for inference, which is implemented as:

$$\begin{aligned} TIE &= TE - NDE = Z_{q,v,k} - Z_{q,v^*,k^*} \\ &= h(z_q, z_v, z_k) - h(z_q, z_v^*, z_k^*). \end{aligned} \quad (19)$$

4.3. Revisiting RUBi and Learned-Mixin

Note that recent language-prior based methods RUBi [10] and Learned-Mixin [13] use an ensemble model that consists a vision-language branch \mathcal{F}_{VQ} and a question-only branch \mathcal{F}_Q . Besides, they simply exclude

Table 1: **Comparison on VQA-CP v2 test set and VQA v2 val set.** **Best** and **second best** results obtained without extra generated training samples are highlighted in each column. “Base.” indicates the VQA base model. We report the average accuracy of CF-VQA over 5 experiments with different random seeds.

Test set	VQA-CP v2 test					VQA v2 val				
	Methods	Base.	All	Y/N	Num.	Other	All	Y/N	Num.	Other
GVQA [3]	–	31.30	57.99	13.68	22.14	48.24	72.03	31.17	34.65	
SAN [47]	–	24.96	38.35	11.14	21.74	52.41	70.06	39.28	47.84	
UpDn [5]	–	39.74	42.27	11.93	46.05	63.48	81.18	42.14	55.66	
S-MRL [10]	–	38.46	42.85	12.81	43.20	63.10	–	–	–	
<i>methods based on modifying language module follow:</i>										
DLR [24]	UpDn	48.87	70.99	18.72	45.57	57.96	76.82	39.33	48.54	
VGQE [27]	UpDn	48.75	–	–	–	64.04	–	–	–	
VGQE [27]	S-MRL	50.11	66.35	27.08	46.77	63.18	–	–	–	
<i>methods based on strengthening visual attention follow:</i>										
AttAlign [37]	UpDn	39.37	43.02	11.89	45.00	63.24	80.99	42.55	55.22	
HINT [37]	UpDn	46.73	67.27	10.61	45.88	63.38	81.18	42.99	55.56	
SCR [44]	UpDn	49.45	72.36	10.93	48.02	62.2	78.8	41.6	54.5	
<i>methods based on weakening language prior follow:</i>										
AdvReg. [33]	UpDn	41.17	65.49	15.48	35.48	62.75	79.84	42.35	55.16	
RUBi [10]	UpDn	44.23	67.05	17.48	39.61	–	–	–	–	
RUBi [10]	S-MRL	47.11	68.65	20.28	43.18	61.16	–	–	–	
LM [13]	UpDn	48.78	72.78	14.61	45.58	63.26	81.16	42.22	55.22	
LM+H [13]	UpDn	52.01	72.58	31.12	46.97	56.35	65.06	37.63	54.69	
CF-VQA (HM)	UpDn	49.74 \pm 0.26	74.81 \pm 0.51	18.46 \pm 0.39	45.19 \pm 0.27	63.73 \pm 0.07	82.15 \pm 0.08	44.29 \pm 0.54	54.86 \pm 0.07	
CF-VQA (HM)	S-MRL	51.27 \pm 0.25	77.80 \pm 0.83	20.64 \pm 1.01	45.76 \pm 0.21	62.49 \pm 0.06	81.19 \pm 0.09	44.64 \pm 0.20	52.98 \pm 0.14	
CF-VQA (SUM)	UpDn	53.55 \pm 0.10	91.15 \pm 0.06	13.03 \pm 0.21	44.97 \pm 0.20	63.54 \pm 0.09	82.51 \pm 0.12	43.96 \pm 0.17	54.30 \pm 0.09	
CF-VQA (SUM)	S-MRL	55.05 \pm 0.12	90.61 \pm 0.28	21.50 \pm 0.86	45.61 \pm 0.16	60.94 \pm 0.15	81.13 \pm 0.15	43.86 \pm 0.40	50.11 \pm 0.16	
<i>methods based on balancing training data follow:</i>										
CVL [1]	UpDn	42.12	45.72	12.45	48.34	–	–	–	–	
Unshuffling [40]	UpDn	42.39	47.72	14.43	47.24	61.08	78.32	42.16	52.81	
RandImg [41]	UpDn	55.37	83.89	41.60	44.20	57.24	76.53	33.87	48.57	
SSL [55]	UpDn	57.59	86.53	29.87	50.03	63.73	–	–	–	
CSS [11]	UpDn	58.95	84.37	49.42	48.21	59.91	73.25	39.77	55.11	
CSS+CL [28]	UpDn	59.18	86.99	49.89	47.16	57.29	67.27	38.40	54.71	
Mutant [18]	UpDn	61.72	88.90	49.68	50.78	62.56	82.07	42.52	53.28	

\mathcal{F}_Q and use $z_k = \mathcal{F}_{VQ}(v, q)$ during the inference stage. The conceptual comparison between our CF-VQA and language-prior based strategies are illustrated in Figure 4. These methods can be unified into our counterfactual inference framework, which (1) follow a simplified causal graph (Fig. 5(a)) without the direct path $V \rightarrow A$, and (2) use natural indirect effect (NIE) in Eq. (5) for inference. The detailed analysis is provided in the appendix.

5. Experiments

We mainly conduct the experiments on the VQA-CP [3] dataset. VQA-CP is proposed to evaluate the robustness of VQA models when the answer distributions of training and test splits are significantly different. In addition, we also report the results on the balanced VQA v2 dataset to see whether the approach over-corrects language bias. The models are evaluated via accuracy. We conduct experiments with three baseline VQA architectures: Stacked Attention Network (SAN) [47], Bottom-up and Top-down Attention (UpDn) [5], and a simplified MUREL [9] (S-MRL) [10].

5.1. Quantitative Results

We first compare CF-VQA with state-of-the-art methods. Recent approaches can be grouped as follows. (1) Methods that *modify language modules* proposed to decouple the linguistic concepts (DLR) [24] or generate visually-grounded question representations (VGQE) [27]. (2) Methods that *strengthen visual attention* exploit human visual [14] or textual [23] explanations, including AttAlign [37], HINT [37] and SCR [44]. (3) Methods that *weaken language prior* proposed to directly formulate the language prior by a separate question-only branch, including AdvReg. [33], RUBi [10] and Learned-Mixin (LM) [13]. (4) Methods that *balance training data* proposed to change the training distribution for unbiased training, including CVL [1], Unshuffling [40], RandImg [41], SSL [55], CSS [11], CL [28], and Mutant [18]. Unshuffling [40] partitions the training set into multiple invariant subsets. Other methods generate counterfactual training samples by masking or transforming critical words and objects [11, 18] or replacing the image [1, 41, 55].

Table 2: **Ablation of CF-VQA** on VQA-CP v2 test set. “SAN/UpDn/S-MRL” denotes the baseline VQA model. “HM/SUM” represents the strategies that train the ensemble model and test with only the vision-language branch following ensemble-based method [10, 13]. * represents the reproduced results.

	All	Y/N	Num.	Other		All	Y/N	Num.	Other		All	Y/N	Num.	Other
SAN*	33.18	38.57	12.25	36.10	UpDn*	37.69	43.17	12.53	41.72	S-MRL*	37.09	41.39	12.46	41.60
HM	45.89	70.37	23.99	39.07	HM	47.97	69.19	18.80	44.86	HM	49.37	73.20	20.10	44.92
+ CF-VQA	48.10	77.68	22.19	39.71	+ CF-VQA	49.74	74.81	18.46	45.19	+ CF-VQA	51.27	77.80	20.64	45.76
SUM	43.98	68.98	17.32	38.19	SUM	47.29	72.26	12.54	43.74	SUM	48.27	74.60	20.96	41.96
+ CF-VQA	50.15	87.95	16.46	39.59	+ CF-VQA	53.55	91.15	13.03	44.97	+ CF-VQA	55.05	90.61	21.50	45.61

Table 3: **Ablation of CF-VQA with the simplified causal graph** on VQA-CP v2 test set. “SAN/UpDn/S-MRL” denotes the baseline VQA model. “HM/SUM” represents the strategies that train the ensemble model and test with only the vision-language branch following ensemble-based method [10, 13]. * represents the reproduced results.

	All	Y/N	Num.	Other		All	Y/N	Num.	Other		All	Y/N	Num.	Other
SAN*	33.18	38.57	12.25	36.10	UpDn*	37.69	43.17	12.53	41.72	S-MRL*	37.09	41.39	12.46	41.60
HM	45.48	71.32	17.19	39.71	HM	46.50	67.54	12.83	44.72	HM	49.57	72.31	20.28	45.68
+ CF-VQA	49.43	83.82	17.52	40.16	+ CF-VQA	49.53	77.02	12.86	45.18	+ CF-VQA	52.68	82.05	20.76	46.04
SUM	42.35	62.33	16.64	38.94	SUM	47.10	70.00	12.80	44.51	SUM	49.42	74.43	20.52	44.24
+ CF-VQA	49.85	87.75	16.15	39.24	+ CF-VQA	53.55	91.15	12.81	45.02	+ CF-VQA	54.52	90.69	21.84	44.53

Table 4: **Ablation of assumptions for counterfactual outputs** on VQA-CP v2 test set.

		All	Y/N	Num.	Other
S-MRL [10]		38.46	42.85	12.81	43.20
	random	31.27	29.69	42.87	28.91
HM	prior	46.29	61.88	20.03	45.33
	uniform	51.27	77.80	20.64	45.76
	random	27.52	28.00	37.88	24.42
SUM	prior	38.06	41.43	14.90	42.64
	uniform	55.05	90.61	21.50	45.61

The results on VQA-CP v2 and VQA v2 are reported in Table 1. Most of the methods that explicitly generate training samples [41, 55, 11, 28, 18] outperform others by large margins. However, these methods explicitly change the training priors, which violates the original intention of VQA-CP, *i.e.*, to evaluate whether VQA models are driven by memorizing priors in training data [3]. Therefore, we do not directly compare CF-VQA with these methods for fairness. Overall, compared to non-augmentation approaches, our proposed CF-VQA achieves a new state-of-the-art performance on VQA-CP v2. With a deep look at the question type, we find that the improvement on “Yes/No” questions is extremely large (from $\sim 70\%$ to $\sim 90\%$), which indicates that language bias have different effects on different types of questions. Besides, methods with extra annotations or generated training samples effectively improves the accuracy on “Other” questions, while CF-VQA achieves comparable performance to other methods. It is worth noting that LM [13] achieves a competitive performance on VQA-CP v2 with an additional language entropy penalty (LM+H). However, the accuracy drops significantly by $\sim 7\%$ on VQA v2, which indicates that the entropy penalty forces the model to over-correct language

Table 5: The comparison between CF-VQA and RUBi.

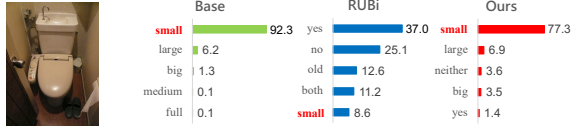
	VQA-CP v2 test				VQA v2 val
	All	Y/N	Num.	Other	All
S-MRL [10]	38.46	42.85	12.81	43.20	63.10
RUBi [10]	47.11 ± 0.51	68.65	20.28	43.18	61.16
+ CF-VQA	54.69± 0.98	89.90	32.39	42.01	60.53

bias, especially on “Yes/No” questions. As a comparison, CF-VQA is more robust on VQA v2.

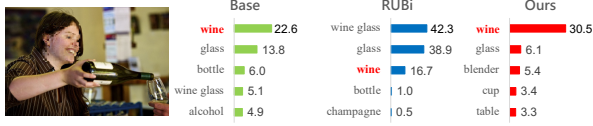
We further conduct ablation studies to validate (1) the generalizability of CF-VQA to both baseline VQA architectures, fusion strategies and causal graphs, and (2) the distribution assumption under the no-treatment conditions. As shown in Table 2 and 3, CF-VQA outperforms ensemble-based strategies by over 2% for HM and over 5% for SUM in all cases. To empirically validate the distribution assumption, we proposed two candidate assumptions. “Random” denotes that c_i for answer a_i are learned without any constraint. “Prior” denotes that $\{c_i\}$ obey the prior distribution of the training set. As shown in Table 4, “random” and “prior” even perform worse than the baselines. As we discussed in Section 4.2, the possible reason is that the uniform distribution assumption guarantees a safe estimation of NDE, *i.e.*, the biased language effect.

Based on our cause-effect look, CF-VQA can easily improve RUBi by replacing NIE with TIE and introducing only one learnable parameter. The details are provided in the appendix. As shown in Table 5, CF-VQA improves RUBi by over 7.5% (*i.e.*, from 47.11 to 54.69) on VQA-CP v2, while the accuracy slightly drops on VQA v2. Compared to our proposed symmetric fusion strategies HM and SUM, the standard deviation of RUBi’s accuracy is larger, which indicates that our proposed symmetric fusion strategies are more stable and robust.

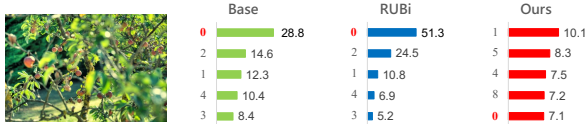
Q: Is this room large or small?



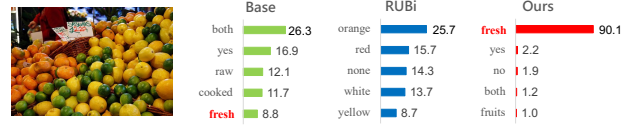
Q: What is being poured?



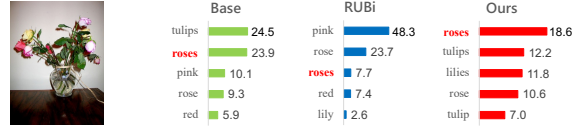
Q: How many birds are depicted?



Q: Is this fruit fresh or frozen?



Q: What type of flowers are these?



Q: How many dolls are on the bed?

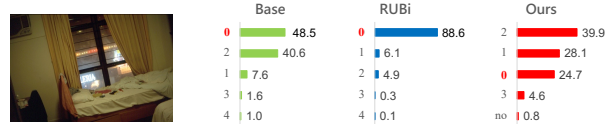


Figure 6: Qualitative comparison on VQA-CP v2 test split. Red bold answer denotes the ground-truth one.

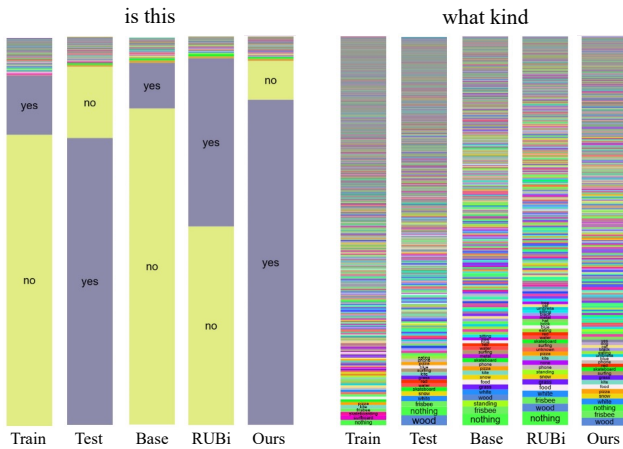


Figure 7: Answer distributions on VQA-CP v2.

5.2. Qualitative Results

The qualitative results are provided to validate whether CF-VQA can effectively reduce language bias and retain language context. As illustrated in Figure 7, CF-VQA can successfully overcome language bias on yes/no questions compared to RUBi, while the baseline model suffers from the memorized language prior on the training set. Besides, for “what kind” questions, RUBi prefers the meaningless answer “none” rather than specific ones. Although CF-VQA cannot recover the answer distribution very well, it attempts to respond with more meaningful answers (e.g., wood, frisbee). Examples in Figure 6 further illustrate how CF-VQA preserves language context for inference. For the left top example, CF-VQA recognizes the correct context “large or small”, while RUBi tends to answer yes/no based

on the wrong context “is this”. For the right example at the second row, although RUBi successfully locates the flowers, it wrongly focuses on visual attributes (i.e., “pink”) rather than categories (i.e., “what type”). These results highlight the importance of language context, which is not considered by language prior based approaches. The examples at the third row show the failure cases on number-related questions. It remains an open problem how to improve the accuracy of number-related questions.

6. Conclusion

In this paper, we proposed a novel counterfactual inference framework CF-VQA to reduce language bias in VQA. The bias is formulated as the direct causal effect of questions on answers and estimated by counterfactual reasoning. The reduction of language bias is realized by subtract the direct linguistic effect from the total causal effect. Experimental results demonstrate the effectiveness and generalizability of CF-VQA. Furthermore, recent debiasing studies [10, 13] can be unified into our proposed counterfactual inference framework. Surprisingly, we can further improve RUBi [10] by simply changing a few lines of code and including only one more learnable parameter based on our cause-effect look. In the future, we will consider how to make balances between robustness and debiasing ability.

Acknowledgements We would like to thank anonymous ACs and reviewers for their valuable discussion and insightful suggestions. This work was supported in part by NTU-Alibaba JRI, MOE AcRF Tier 2 grant, National Natural Science Foundation of China (61976220 and 61832017), and Beijing Outstanding Young Scientist Program (BJJWZYJH012019100020098).

References

- [1] Ehsan Abbasnejad, Damien Teney, Amin Parvaneh, Javen Shi, and Anton van den Hengel. Counterfactual vision and language learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10044–10054, 2020. 1, 3, 6
- [2] Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. Analyzing the behavior of visual question answering models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1955–1960, 2016. 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4971–4980, 2018. 1, 2, 6, 7
- [4] Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, C Lawrence Zitnick, Devi Parikh, and Dhruv Batra. Vqa: Visual question answering. *International Journal of Computer Vision*, 123(1):4–31, 2017. 1
- [5] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6077–6086, 2018. 6
- [6] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3674–3683, 2018. 1
- [7] Lisa Anne Hendricks, Ronghang Hu, Trevor Darrell, and Zeynep Akata. Grounding visual explanations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 264–279, 2018. 3
- [8] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015. 1, 2
- [9] Remi Cadene, Hedi Ben-Younes, Matthieu Cord, and Nicolas Thome. Murel: Multimodal relational reasoning for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1989–1998, 2019. 6
- [10] Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. Rubi: Reducing unimodal biases for visual question answering. *Advances in Neural Information Processing Systems*, 32:841–852, 2019. 1, 2, 3, 5, 6, 7, 8
- [11] Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. Counterfactual samples synthesizing for robust visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10800–10809, 2020. 1, 3, 6, 7
- [12] Long Chen, Hanwang Zhang, Jun Xiao, Xiangnan He, Shiliang Pu, and Shih-Fu Chang. Counterfactual critic multi-agent training for scene graph generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4613–4623, 2019. 3
- [13] Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. Don’t take the easy way out: Ensemble based methods for avoiding known dataset biases. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4060–4073, 2019. 1, 2, 3, 5, 6, 7, 8
- [14] Abhishek Das, Harsh Agrawal, Larry Zitnick, Devi Parikh, and Dhruv Batra. Human attention in visual question answering: Do humans and deep networks look at the same regions? *Computer Vision and Image Understanding*, 163:90–100, 2017. 1, 3, 6
- [15] Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. Visual dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 326–335, 2017. 1
- [16] Zhiyuan Fang, Shu Kong, Charless Fowlkes, and Yezhou Yang. Modularized textual grounding for counterfactual resilience. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6378–6388, 2019. 3
- [17] Tsu-Jui Fu, Xin Wang, Scott Grafton, Miguel Eckstein, and William Yang Wang. Iterative language-based image editing via self-supervised counterfactual reasoning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4413–4422, 2020. 3
- [18] Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. Mutant: A training paradigm for out-of-distribution generalization in visual question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 878–892, 2020. 1, 6, 7
- [19] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6904–6913, 2017. 1, 2
- [20] Yash Goyal, Ziyang Wu, Jan Ernst, Dhruv Batra, Devi Parikh, and Stefan Lee. Counterfactual visual explanations. In *International Conference on Machine Learning*, pages 2376–2384. PMLR, 2019. 3
- [21] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 5
- [22] Xinting Hu, Kaihua Tang, Chunyan Miao, Xian-Sheng Hua, and Hanwang Zhang. Distilling causal effect of data in class-incremental learning. *arXiv preprint arXiv:2103.01737*, 2021. 3
- [23] Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Anna Rohrbach, Bernt Schiele, Trevor Darrell, and Marcus

- Rohrbach. Multimodal explanations: Justifying decisions and pointing to the evidence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8779–8788, 2018. 1, 3, 6
- [24] Chenchen Jing, Yuwei Wu, Xiaoxun Zhang, Yunde Jia, and Qi Wu. Overcoming language priors in vqa via decomposed linguistic representations. In *AAAI*, pages 11181–11188, 2020. 6
- [25] Kushal Kafle and Christopher Kanan. An analysis of visual question answering algorithms. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1965–1973, 2017. 1, 2
- [26] Atsushi Kanehira, Kentaro Takemoto, Sho Inayoshi, and Tatsuya Harada. Multimodal explanations by predicting counterfactuality in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8594–8602, 2019. 3
- [27] Gouthaman KV and Anurag Mittal. Reducing language biases in visual question answering with visually-grounded question encoder. *arXiv preprint arXiv:2007.06198*, 2020. 6
- [28] Zujie Liang, Weitao Jiang, Haifeng Hu, and Jiaying Zhu. Learning to contrast the counterfactual samples for robust visual question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3285–3292, 2020. 1, 6, 7
- [29] Judea Pearl. *Causality: models, reasoning and inference*, volume 29. Springer, 2000. 1, 3
- [30] Judea Pearl. Direct and indirect effects. In *Proceedings of the seventeenth conference on uncertainty in artificial intelligence*, pages 411–420. Morgan Kaufmann Publishers Inc., 2001. 1, 3
- [31] Judea Pearl and Dana Mackenzie. *The book of why: the new science of cause and effect*. Basic Books, 2018. 1, 3
- [32] Jiaxin Qi, Yulei Niu, Jianqiang Huang, and Hanwang Zhang. Two causal principles for improving visual dialog. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10860–10869, 2020. 3
- [33] Sainandan Ramakrishnan, Aishwarya Agrawal, and Stefan Lee. Overcoming language priors in visual question answering with adversarial regularization. In *Advances in Neural Information Processing Systems*, pages 1541–1551, 2018. 3, 6
- [34] James Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling*, 7(9-12):1393–1512, 1986. 3
- [35] James M Robins. Semantics of causal dag models and the identification of direct and indirect effects. *Oxford Statistical Science Series*, pages 70–82, 2003. 3
- [36] Donald B Rubin. Bayesian inference for causal effects: The role of randomization. *The Annals of statistics*, pages 34–58, 1978. 3
- [37] Ramprasaath R Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Shalini Ghosh, Larry Heck, Dhruv Batra, and Devi Parikh. Taking a hint: Leveraging explanations to make vision and language models more grounded. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2591–2600, 2019. 1, 3, 6
- [38] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3716–3725, 2020. 3
- [39] Damien Teney, Ehsan Abbasnejad, and Anton van den Hengel. Learning what makes a difference from counterfactual examples and gradient supervision. *arXiv preprint arXiv:2004.09034*, 2020. 3
- [40] Damien Teney, Ehsan Abbasnejad, and Anton van den Hengel. Unshuffling data for improved generalization. *arXiv preprint arXiv:2002.11894*, 2020. 3, 6
- [41] 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. 6, 7
- [42] Pei Wang and Nuno Vasconcelos. Scout: Self-aware discriminant counterfactual explanations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8981–8990, 2020. 3
- [43] Tan Wang, Jianqiang Huang, Hanwang Zhang, and Qianru Sun. Visual commonsense r-cnn. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10760–10770, 2020. 3
- [44] Jialin Wu and Raymond J Mooney. Self-critical reasoning for robust visual question answering. *arXiv preprint arXiv:1905.09998*, 2019. 1, 3, 6
- [45] Xu Yang, Hanwang Zhang, and Jianfei Cai. Deconfounded image captioning: A causal retrospect. *arXiv preprint arXiv:2003.03923*, 2020. 3
- [46] Xu Yang, Hanwang Zhang, Guojun Qi, and Jianfei Cai. Causal attention for vision-language tasks. *arXiv preprint arXiv:2103.03493*, 2021. 3
- [47] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 21–29, 2016. 6
- [48] Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B Tenenbaum. Clevrer: Collision events for video representation and reasoning. *arXiv preprint arXiv:1910.01442*, 2019. 3
- [49] Zhongqi Yue, Tan Wang, Hanwang Zhang, Qianru Sun, and Xian-Sheng Hua. Counterfactual zero-shot and open-set visual recognition. *arXiv preprint arXiv:2103.00887*, 2021. 3
- [50] Zhongqi Yue, Hanwang Zhang, Qianru Sun, and Xian-Sheng Hua. Interventional few-shot learning. *arXiv preprint arXiv:2009.13000*, 2020. 3
- [51] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6720–6731, 2019. 1
- [52] Dong Zhang, Hanwang Zhang, Jinhui Tang, Xiansheng Hua, and Qianru Sun. Causal intervention for weakly-supervised semantic segmentation. *arXiv preprint arXiv:2009.12547*, 2020. 3

- [53] 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. [2](#)
- [54] Shengyu Zhang, Tan Jiang, Tan Wang, Kun Kuang, Zhou Zhao, Jianke Zhu, Jin Yu, Hongxia Yang, and Fei Wu. De-
vlbert: Learning deconfounded visio-linguistic representations. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4373–4382, 2020. [3](#)
- [55] Xi Zhu, Zhendong Mao, Chunxiao Liu, Peng Zhang, Bin Wang, and Yongdong Zhang. Overcoming language priors with self-supervised learning for visual question answering. *arXiv preprint arXiv:2012.11528*, 2020. [1](#), [3](#), [6](#), [7](#)