

Check It Again: Progressive Visual Question Answering via Visual Entailment

Qingyi Si^{1,2}, Zheng Lin^{1*}, Mingyu Zheng^{1,2}, Peng Fu¹, Weiping Wang¹

¹Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China

²School of Cyber Security, University of Chinese Academy of Sciences, Beijing, China
{siqingyi, linzheng, zhengmingyu, fupeng, wangweiping}@iie.ac.cn

Abstract

While sophisticated Visual Question Answering models have achieved remarkable success, they tend to answer questions only according to superficial correlations between question and answer. Several recent approaches have been developed to address this language priors problem. However, most of them predict the correct answer according to one best output without checking the authenticity of answers. Besides, they only explore the interaction between image and question, ignoring the semantics of candidate answers. In this paper, we propose a select-and-rerank (SAR) progressive framework based on Visual Entailment. Specifically, we first select the candidate answers relevant to the question or the image, then we rerank the candidate answers by a visual entailment task, which verifies whether the image semantically entails the synthetic statement of the question and each candidate answer. Experimental results show the effectiveness of our proposed framework, which establishes a new state-of-the-art accuracy on VQA-CP v2 with a 7.55% improvement.¹

1 Introduction

Visual Question Answering (VQA) task is a multi-modal problem which requires the comprehensive understanding of both visual and textual information. Presented with an input image and a question, the VQA system tries to determine the correct answer in the large prediction space. Recently, some studies (Jabri et al., 2016; Agrawal et al., 2016; Zhang et al., 2016; Goyal et al., 2017) demonstrate that VQA systems suffer from the superficial correlation bias (i.e. language priors) caused by accidental correlations between answers and questions. As a result, traditional VQA models always output the

*Corresponding author: Zheng Lin.

¹The code is available at <https://github.com/PhoebusSi/SAR>

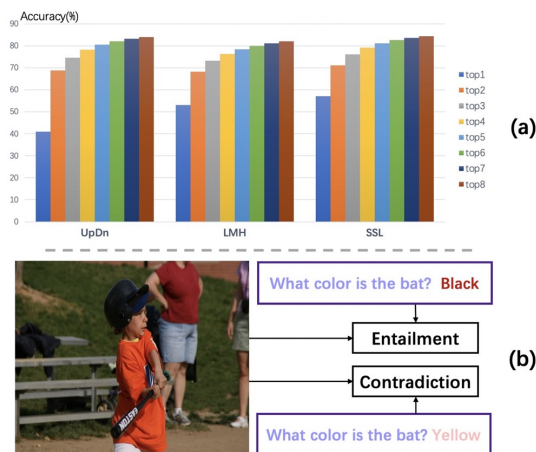


Figure 1: (a) We evaluate the performance of UpDn, LMH, SSL on the VQA-CP v2 test. $topN$ represents the topN accuracy. (b) Visual verification utilizing answer semantics.

most common answer (Selvaraju et al., 2019) of the input sample’s question category, no matter what image is given. To address this language priors problem, various approaches have been developed.

However, through exploring the characteristics of the existing methods, we find that whether the general VQA models such as UpDn (Anderson et al., 2018) and LXMERT (Tan and Bansal, 2019) or models carefully designed for language priors, as LMH (Clark et al., 2019) and SSL (Zhu et al., 2020), yield a non-negligible problem. Both models predict the correct answer according to one best output without checking the authenticity of answers. Besides, these models have not made good use of the semantics information of answers that could be helpful for alleviating the language-priors.

As presented in Figure 1(a), quite a few correct answers usually occur at top N candidates rather than top one. Meanwhile, if the top N candidate answers are given, the image can further verify the visual presence/absence of concepts based on the combination of the question and the candidate

answer. As shown in Figure 1(b), the question is about the color of the bat and two candidate answers are “yellow” and “black”. After checking the correctness of candidate answers, the wrong answer “yellow” which is contradicted with the image can be excluded and the correct answer “black” which is consistent with the image is confirmed. Nevertheless, this visual verification, which utilizes answer semantics to alleviate language priors, has not been fully investigated.

In this paper, we propose a select-and-rerank (SAR) progressive framework based on Visual Entailment. The intuition behind the proposed framework comes from two observations. First, after excluding the answers unrelated to the question and image, the prediction space is shrunken and we can obtain a small number of candidate answers. Second, on the condition that a question and one of its candidate answer is bridged into a complete statement, the authenticity of this statement can be inferred by the content of the image. Therefore, after selecting several possible answers as candidates, we can utilize the visual entailment, consisting of image-text pairs, to verify whether the image semantically entails the synthetic statement. Based on the entailment degree, we can further rerank candidate answers and give the model another chance to find the right answer. To summarize, our contributions are as follows:

1. We propose a select-and-rerank progressive framework to tackle the language priors problem, and empirically investigate a range of design choices for each module of this framework. In addition, it is a generic framework, which can be easily combined with the existing VQA models and further boost their abilities.

2. We highlight the verification process between text and image, and formulate the VQA task as a visual entailment problem. This process makes full use of the interactive information of image, question and candidate answers.

3. Experimental results demonstrate that our framework establishes a new state-of-the-art accuracy of 66.73%, outperforming the existing methods by a large margin.

2 Related Work

Language-Priors Methods To address the language prior problem of VQA models, a lot of approaches have been proposed, which can be roughly categorized into two lines: (1) *Design-*

ing Specific Debiasing Models to Reduce Biases. Most works of this line are ensemble-based methods (Ramakrishnan et al., 2018; Grand and Belinkov, 2019; Belinkov et al., 2019; Cadene et al., 2019; Clark et al., 2019; Mahabadi and Henderson, 2019), among these, LMH(Clark et al., 2019) reduces all biases between question-answer pairs by penalizing the samples that can be answered without utilizing image content. (2) *Data Augmentation to Reduce Biases.* The main idea of such works (Zhang et al., 2016; Goyal et al., 2017; Agrawal et al., 2018) is to carefully construct more balanced datasets to overcome priors. For example, the recent method SSL(Zhu et al., 2020) first automatically generates a set of balanced question-image pairs, then introduces an auxiliary self-supervised task to use the balanced data. CSS(Chen et al., 2020a) balances the data by adding more complementary samples which are generated by masking objects in the image or some keywords in the question. Based on CSS, CL(Liang et al., 2020) forces the model to utilize the relationship between complementary samples and original samples. Unlike SSL and CSS which do not use any extra manual annotations, MUTANT(Gokhale et al., 2020) locates critical objects in the image and critical words in the question by utilizing the extra object-name labels, which directly helps the model to ground the textual concepts in the image. However, above methods only explore the interaction between the image and the question, ignoring the semantics of candidate answers. In this paper, we propose a progressive VQA framework SAR which achieves better interaction among the question, the image and the answer.

Answer Re-ranking Although Answer Re-ranking is still in the infancy in VQA task, it has been widely studied for QA tasks like open-domain question answering, in which models need to answer questions based on a broad range of open-domain knowledge sources. Recent works (Wang et al., 2018b,a; Kratzwald et al., 2019) address this task in a two-stage manner: extract candidates from all passages, then focus on these candidate answers and rerank them to get a final answer. RankVQA(Qiao et al., 2020) introduces Answer Re-ranking method to VQA task. They design an auxiliary task which reranks candidate answers according to their matching degrees with the input image and off-line generated image captions. However, RankVQA still predicts the final answer from

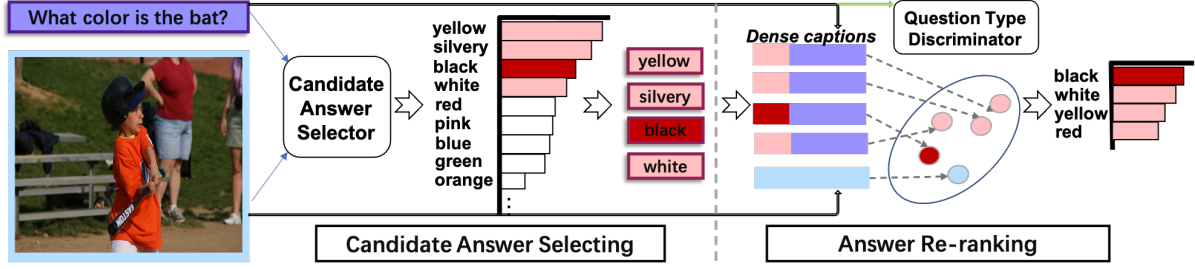


Figure 2: Overview of the progressive framework SAR.

the huge prediction space rather than selected candidate answers.

3 Method

Figure 2 shows an overview of the proposed select-and-rerank (SAR) framework, which consists of a Candidate Answer Selecting module and an Answer Re-ranking module. In the Candidate Answer Selecting module, given an image and a question, we first use a current VQA model to get a candidate answer set consisting of top N answers. In this module, the answers irrelevant to the question can be filtered out. Next, we formulate the VQA as a VE task in the Answer Re-ranking module, where the image is premise and the synthetic dense caption(Johnson et al., 2016) (combination of the answer and the question) is hypothesis. We use the cross-domain pre-trained model LXMERT(Tan and Bansal, 2019) as VE scorer to compute the entailment score of each image-caption pair, and thus the answer corresponding to the dense caption with the highest score is our final prediction.

3.1 Candidate Answer Selecting

The Candidate Answer Selector (CAS) selects several answers from all possible answers as candidates and thus shrinks the huge prediction space. Given a VQA dataset $D = \{I_i, Q_i\}_{i=1}^M$ with M samples, where $I_i \in I$, $Q_i \in Q$ are the image and question of the i_{th} sample and A is the whole prediction space consisting of thousands of answer categories. Essentially, the VQA model applied as CAS is a $|A|$ -class classifier, and is a free choice in our framework. Given an image I_i and a question Q_i , CAS first gives the regression scores over all optional answers: $P(A|Q_i, I_i)$. Then CAS chooses N answers A_i^* with top N scores as candidates, which is concluded as follows:

$$A_i^* = \text{top}N(\text{argsort}(P(A|Q_i, I_i))) \quad (1)$$

N (hyper-parameter) candidate answers $A_i^* = [A_i^1, A_i^2, \dots, A_i^N]$ are selected for each (I_i, Q_i) pair by CAS, forming a dataset $D' = \{I_i, Q_i, A_i^n\}_{i=1, n=1}^{M, N}$ with $M * N$ instances, where $A_i^n \in A_i^*$, for the next Answer Re-ranking module. In this paper, we mainly use SSL as our CAS. We also conduct experiments to analyze the impact of different CAS and different N .

3.2 Answer Re-ranking

3.2.1 Visual Entailment

Visual Entailment (VE) task is proposed by Xie et al. (2019), where the premise is a real-world image, denoted by P_{image} , and the hypothesis is a text, denoted by H_{text} . Given a sample of (P_{image}, H_{text}) , the goal of VE task is to determine whether the H_{text} can be concluded based on the information of P_{image} . According to following protocols, the label of the sample is assigned to (1) *Entailment*, if there is enough evidence in P_{image} to conclude H_{text} is true. (2) *Contradiction*, if there is enough evidence in P_{image} to conclude H_{text} is false. (3) *Neutral*, if there is no sufficient evidence in P_{image} to give a conclusion about H_{text} .

3.2.2 VQA As Visual Entailment

A question Q_i and each of its candidate answers A_i^* can be bridged into a complete statement, and then the image could verify the authenticity of each statement. More specifically, the visual presence of concepts (e.g. “black bat”/“yellow bat”) based on the combination of the question and the correct/wrong candidate answer can be entailed/contradicted by the content of the image. In this way, we achieve better interaction among question, image and answer.

Therefore, we formulate VQA as a VE problem, in which the image I_i is premise, and the synthetic statement of an answer A_i^n in A_i^* and question Q_i , represented as (Q_i, A_i^n) , is hypothesis. For an image, synthetic statements of different

questions describe different regions of the same image. Following Johnson et al. (2016), we also refer to the synthetic statement as “dense caption”. We use A_i^+ to represent the A_i^n if A_i^n is the correct answer of Q_i , use A_i^- otherwise. There is enough evidence in I_i to prove (Q_i, A_i^+) is true, i.e. the visual linguistic semantically entails (Q_i, A_i^+) . And there is enough evidence in I_i to prove (Q_i, A_i^-) is false, i.e. the visual linguistic semantically contradicts (Q_i, A_i^-) . Note that, there is no *Neutral* label in our VE task and we only have two labels: *Entailment* and *Contradiction*.

3.2.3 Re-Ranking based on VE

We re-rank dense captions by contrastive learning, that is, (Q_i, A_i^+) should be more semantically similar to I_i than (Q_i, A_i^-) . The right part of Figure 2 illustrates this idea. The more semantically similar I_i to (Q_i, A_i^n) , the deeper the visual entailment degree is. We score the visual entailment degree of I_i to each $(Q_i, A_i^n) \in (Q_i, A_i^*)$ and rerank the candidate answers A_i^* by this score. The ranking-first answer is our final output.

Question-Answer Combination Strategy The answer information makes sense only when combine it with the question. We encode the combination of question and answer text to obtain the joint concept.

We design three question-answer combination strategies: **R**, **C** and **R→C** to combine question and answer into synthetic dense caption C_i :

R: Replace question category prefix with answer. The prefix of each question is the question category such as “are there”, “what color”, etc. For instance, given a question “How many flowers in the vase?”, its answer “8” and its question category “how many”, the resulting dense caption is “8 flowers in the vase”. Similarly, “No a crosswalk” is the result of question “Is this a crosswalk?” and answer “No”. We build a dictionary of all question categories of the train set, then we adopt a Forward Maximum Matching algorithm to determine the question category for every test sample.

C: Concatenate question and answer directly. For two cases above, the resulting dense captions are “8 How many flowers in the vase?” and “No Is this a crosswalk?”. The resulting dense captions after concatenation are actually rhetorical questions. We deliberately add answer text to the front of question text in order to avoid the answer being deleted when trimming dense captions to the same length.

R→C: We first use strategy **R** at training, which is aimed at preventing the model from excessively focusing on the co-occurrence relation between question category and answer, and then use strategy **C** at testing to introduce more information for inference.

Adopting any strategy above, we combine Q_i and each answer in A_i^* to derive the dense captions C_i^* . And thus we have a dataset $D'' = \{I_i, C_i^n\}_{i=1, n=1}^{M, N}$ with $M * N$ instances for VE task.

VE Scorer We use the pre-trained model LXMERT to score the visual entailment degree of (I_i, C_i^n) . LXMERT separately encodes image and caption text in two streams. Next, the separate streams interact through co-attentional transformer layers. In the textual stream, the dense caption is encoded into a high-level concept. Then the visual representations from visual stream can verify the visual presence/absence of the high-level concept.

We represent the VE score for the i_{th} image and its n_{th} candidate caption as: $sigmoid(Trm(I_i, C_i^n))$, where $Trm()$ is the 1-dimensional output from the dense layers following LXMERT, $\delta()$ denotes the sigmoid function. The larger score represents higher entailment degree. We optimize parameters by minimizing the multi-label soft loss:

$$L_{VE} = \frac{-1}{M * N} \sum_{i=1}^M \sum_{n=1}^N [t_i^n \log(\delta(Trm(I_i, C_i^n))) + (1 - t_i^n) \log(1 - \delta(Trm(I_i, C_i^n)))] \quad (2)$$

where t_i^n is the soft target score of the n_{th} answer.

Combination with Language-Priors Method

After Candidate Answer Selecting, the amount of candidate answers decreases from all possible answers to top N . Although some unrelated answers are filtered out, the dataset D'' for VE system is still biased. Therefore, we can optionally apply existing language-priors methods to our framework for further reducing language priors. Take the SSL as an example, we apply the loss function of its self-supervised task to our framework by adjusting the loss function to:

$$L_{ssl} = \frac{\alpha}{M * N} \sum_{i=1}^M \sum_{n=1}^N P(I_i', C_i^n) \quad (3)$$

where (I_i', C_i^n) denotes the irrelevant image-caption pairs, α is a down-weighting coefficients.

The probability $P(I'_i, C_i^n)$ could be considered as the confidence of (I'_i, C_i^n) being a relevant pair. We can reformulate the overall loss function:

$$L = L_{VE} + L_{ssl} \quad (4)$$

3.3 Inference Process

Question Type Discriminator Intuitively, most “Yes/No” questions can be answered by the answer “Yes” or “No”. There is no need to provide too many candidate answers for “Yes/No” questions at the test stage. Therefore, we propose a Question Type Discriminator (QTD) to determine the question type and then correspondingly set different numbers of candidate answers, denoted as N' . Specifically, we roughly divided question types (including “Yes/No”, “Num” and “Other”) into yes/no and non-yes/no. A GRU binary classifier is trained with cross-entropy loss and evaluated with 5-fold cross-validation on the train split of each dataset. Then, the trained QTD model with an accuracy about 97% is implemented as an off-line module during the test stage. We will further investigate the effect of N' on each question type in the next section.

Final Prediction In the inference phase, we search for the best dense caption \hat{C}_i among all candidates C_i^* for the i_{th} image.

$$\hat{C}_i = \underset{n \in N'}{\operatorname{argmax}} \delta(\operatorname{Trm}(I_i, C_i^n)) \quad (5)$$

The answer \hat{A}_i corresponding to \hat{C}_i is the final prediction.

4 Experiments

4.1 Setting

Datasets Our models are trained and evaluated on the VQA-CP v2 (Agrawal et al., 2018) dataset, which is well-crafted by re-organizing VQA v2 (Goyal et al., 2017) training and validation sets such that answers for each question category (65 categories according to the question prefix) have different distributions in the train and test sets. Therefore, VQA-CP v2 is a natural choice for evaluating VQA model’s generalizability. The questions of VQA-CP v2 include 3 types: “Yes/No”, “Num” and “Other”. Note that the question type and question category (e.g. “what color”) are different. Besides, we also evaluate our models on the VQA v2 validation set for completeness, and compare the accuracy difference between two datasets

with the standard VQA evaluation metric (Antol et al., 2015).

Baselines We compare our method with the following baseline methods: UpDn (Anderson et al., 2018), AReg (Ramakrishnan et al., 2018), RUBi (Cadene et al., 2019), LMH (Clark et al., 2019), RankVQA (Qiao et al., 2020), SSL (Zhu et al., 2020), CSS (Chen et al., 2020a), CL (Liang et al., 2020) and LXMERT (Tan and Bansal, 2019). Most of them are designed for the language priors problem, while LXMERT represents the recent trend towards utilizing BERT-like pre-trained models (Li et al., 2019; Chen et al., 2020b; Li et al., 2020) which have top performances on various downstream vision and language tasks (including VQA-v2). Note that MUTANT (Gokhale et al., 2020) uses the extra object-name label to ground the textual concepts in the image. For fair comparison, we do not compare with MUTANT.

4.2 Implementation Details

In this paper, we mainly choose SSL as our CAS and set $N=12$ and $N=20$ for training. To extract image features, we follow previous work and use the pre-trained Faster R-CNN to encode each image as a set of fixed 36 objects with 2048-dimensional feature vectors. We use the tokenizer of LXMERT to segment each dense caption into words. All the questions are trimmed to the same length of 15 or 18, respectively for R or C question-answer combination strategy. In the Answer Re-ranking Module, we respectively incorporate two language-priors methods, SSL and LMH, into our proposed framework SAR, which is dubbed as SAR+SSL and SAR+LMH. Our models are trained on two TITAN RTX 24GB GPUs. We train SAR+SSL for 20 epochs with batch size of 32, SAR and SAR+LMH for 10 epochs with batch size of 64. For SAR+SSL, we follow the same setting as the original paper (Zhu et al., 2020), except that we don’t need to pre-train the model with the VQA loss before fine-tuning it with the self-supervised loss. The Adam optimizer is adopted with the learning rate $1e-5$.

For Question Type Discriminator, we use 300-dimensional Glove (Pennington et al., 2014) vectors to initialize word embeddings and feed them into a unidirectional GRU with 128 hidden units. When testing on the VQA-CP v2, N' ranges from 1-2 for yes/no questions and 5-15 for non-yes/no questions. As for VQA v2, N' ranges from 1-2 for yes/no

Model	VQA-CP v2 test(%) \uparrow				VQA-v2 val(%) \uparrow				GAP (%) \downarrow
	ALL	Yes/No	Num	Other	All	Yes/No	Num	Other	
UpDN(Anderson et al., 2018)	39.74	42.27	11.93	46.05	63.48	81.18	42.14	55.66	23.74
Areg(Ramakrishnan et al., 2018)	41.17	65.49	15.48	35.48	62.75	79.84	42.35	55.16	21.58
RUBI(Cadene et al., 2019)	47.11	68.65	20.28	43.18	61.16	-	-	-	14.05
LMH(Clark et al., 2019)	52.45	69.81	44.46	45.54	61.64	77.85	40.03	55.04	9.19
RankVQA(Qiao et al., 2020)	43.05	42.53	13.91	51.32	65.42	82.51	57.75	45.35	22.37
LXMERT(Tan and Bansal, 2019)	46.23	42.84	18.91	55.51	74.16	89.31	56.85	65.14	27.93
SSL(Zhu et al., 2020)	57.59	86.53	29.87	50.03	63.73	-	-	-	6.14
CSS(Chen et al., 2020a)	58.95	84.37	49.42	48.21	59.91	73.25	39.77	55.11	0.96
CL(Liang et al., 2020)	59.18	86.99	49.89	47.16	-	-	-	-	-
Top12-SAR(R \rightarrow C) (Ours)	64.55	83.03	50.05	58.8	70.41	87.87	54.34	61.38	5.86
Top20-SAR(R \rightarrow C) (Ours)	65.44	83.13	54.52	59.16	<u>70.63</u>	<u>87.91</u>	<u>54.93</u>	<u>61.64</u>	5.19
Top12-SAR+SSL(R \rightarrow C) (Ours)	64.29	82.86	51.98	57.94	69.84	87.22	54.41	60.70	5.55
Top20-SAR+SSL(R \rightarrow C) (Ours)	65.32	83.41	54.32	58.85	70.03	87.47	54.59	60.85	4.71
Top12-SAR+LMH(R) (Ours)	65.93	85.38	62.30	56.73	69.13	87.61	50.43	60.03	3.20
Top20-SAR+LMH(R) (Ours)	66.73	<u>86.00</u>	62.34	57.84	69.22	87.46	51.20	60.12	<u>2.49</u>

Table 1: Results on VQA-CP v2 test and VQA-v2 validation set. Overall best scores are bold, our best are underlined. The gap represents the accuracy difference between VQA v2 and VQA-CP v2.

questions and 2-5 for non-yes/no questions.

4.3 Results and Analysis

4.3.1 Main Results

Performance on two benchmarks VQA-CP-v2 and VQA-v2 is shown in Table 1. We report the best results of SAR, SAR+SSL and SAR+LMH among 3 question-answer combination strategies respectively. “TopN-” represents that N candidate answers (selected by CAS) feed into the Answer Re-ranking Module for training. Our approach is evaluated with two settings of N (12 and 20).

From the results on VQA-CP v2 shown in Table 1, we can observe that: (1) Top20-SAR+LMH establishes a new state-of-the-art accuracy of 66.73% on VQA-CP v2, beating the previous best-performing method CL by 7.55%. Even without combining language-priors methods in Answer Re-ranking module, our model Top20-SAR outperforms CL by 6.26%. These show the outstanding effectiveness of our proposed SAR framework. (2) SAR+SSL and SAR+LMH achieve much better performance than SSL and LMH, which demonstrates that SAR is compatible with current language-priors methods and could realize their full potential. (3) Compared with another reranking-based model RankVQA, our method elevates the performance by a large margin of 23.68%. This shows the superiority of our proposed progressive select-and-rerank framework over RankVQA which only uses the answer reranking as an auxiliary task. (4) Previous models did not generalize well on all question types. CL is the previ-

ous best on the “Yes/No”, “Num” questions and LXMERT on the “Other” questions. In comparison, our model not only rivals the previous best model on the “Yes/No” questions but also improves the best performance on the “Num” and “Other” questions by 12.45% and 3.65%. The remarkable performance on all question types demonstrates that our model makes a significant progress toward a truly comprehensive VQA model.

We also evaluate our method on the VQA v2 which is deemed to have strong language biases. As shown in Table 1, our method achieves the best accuracy of 70.63% amongst baselines specially designed for overcoming language priors, and is the closest to the SOTA established by LXMERT which is trained explicitly for the biased data setting. For completeness, the performance gap between two datasets is also compared in Table 1 with the protocol from Chen et al. (2020a). Compared with most previous models which suffer severe performance drops between VQA v2 and VQA-CP v2 (e.g., 27.93% in LXMERT), the Top20-SAR+LMH significantly decreases the performance drop to 2.49%, which demonstrates the effectiveness of our framework to further overcome the language biases. Though CSS achieves a better performance gap, it sacrifices the performance on the VQA v2. Meanwhile, as N rises from 12 to 20, our models achieve better accuracy on both datasets along with a smaller performance gap. This demonstrates that, unlike previous methods, our method can alleviate language priors while maintaining an excellent capability of answering questions. Nonetheless, we

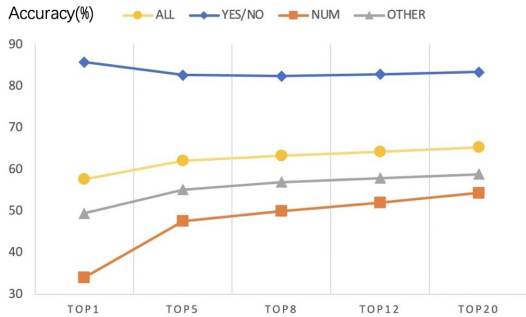


Figure 3: Results from model SAR+SSL(R→C) in VQA-CP v2 with different N during training.

Model/CAS	UpDn	LMH	SSL
w/o SAR*	41.04	53.03	57.66
SAR	61.71	61.65	64.55
SAR+SSL	63.52	61.78	64.29
SAR+LMH	64.98	62.72	65.14

Table 2: Results based on different CAS in VQA-CP v2. We set $N=12$. * indicates the results come from our reimplementing using official released codes.

believe that, how to improve the model’s generality and further transform the trade-off between eliminating language priors and answering questions into win-win outcomes, is a promising research direction in the future.

4.3.2 The Effect of N

From Figure 3, we can observe that the overall performance is getting better as N increases. The performance improvement on the “Num” and “Other” questions is especially obvious, and there is a very slight drop on the “Yes/No” questions. We believe that SAR can further get better performance by properly increasing N . Due to the resource limitation, the largest N we use is 20 in this paper.

4.3.3 The Effect of Different CAS

To find out the potential performance limitation of CAS models, we show the accuracy of 3 CAS models on the VQA-CP v2 test set. As shown in Figure 1 (a), the Top3 accuracy (acc) of 3 models is about 70% and Top6 acc is 80%, which guarantees that sufficient correct answers are recalled by CAS. And thus, the performance limitation of CAS is negligible.

We also conduct experiments to investigate the effect of different CAS on SAR. From the results shown in Table 2, we can observe that: (1) Choosing a better VQA model as CAS does not guarantee a better performance, e.g. performance based on

Top N	Model	R	C	R→C
Top12	SAR	59.51	60.24	64.55
	SAR+SSL	62.12	62.87	64.29
	SAR+LMH	65.93	65.23	65.14
Top20	SAR	60.43	61.86	65.44
	SAR+SSL	62.29	63.94	65.32
	SAR+LMH	66.73	65.19	66.71

Table 3: Results on the VQA-CP v2 test set based on different question-answer combination strategies: R, C and R→C. The major difference between R and C is whether keeping question prefix which includes 65 categories.

UpDn outperforms that based on LMH, but LMH is a better VQA model in overcoming language priors compared with UpDn. This is because a good Candidate Answer Selector has two requirements: (a) It should be able to recall more correct answers. (b) Under the scenario of language biases, wrong answers recalled by CAS at training time should have superficial correlations with the question as strong as possible. However, the ensemble methods, such as LMH, are trained to pay more attention to the samples which are not correctly answered by the question-only model. This seriously reduces the recall rate of those language-priors wrong answers, which leads to the training data for VE is too simple and thus hurts the model’s capability of reducing language priors. (2) If CAS is the general VQA model UpDn rather than LMH and SSL, the improvement brought from the combination with language-priors method in Answer Re-ranking module is more obvious. (3) Even we choose the UpDn, a backbone model of most current works, as our CAS and **do not involve any language-priors methods**, SAR still achieves a much better accuracy than the previous SOTA model CL by 2.53%, which shows that our basic framework already possesses outstanding capability of reducing language priors.

4.3.4 The Effect of Question-Answer Combination Strategies

From the results shown in Table 3, we can observe that: (1) From overall results, R→C achieves or rivals the best performance on three models. On average, R→C outperforms C by 2.02% which demonstrates avoiding the co-occurrence of question category and answer during training time could effectively alleviate language priors; R→C outperforms R by 2.41% which indicates that the informa-

Model	All	Yes/No	Num	Other
LXM	46.23	42.84	18.91	55.51
LXM+SSL	53.09	55.07	29.60	58.50
CAS+LXM(R)	55.58	70.91	29.14	54.81
CAS+LXM+SSL(R)	59.41	76.60	40.81	55.51
CAS+LXM+QTD(R)	59.51	83.20	29.17	55.42
CAS+LXM+SSL+QTD(R)	62.12	85.14	41.63	55.68

Table 4: Ablation study to investigate the effect of each component of Top12-SAR+SSL: Candidate Answer Selector (CAS), LXMERT (LXM), Question Type Discriminator (QTD) and SSL.

tion of question category is useful in inference. (2) On the SAR and SAR+SSL, C consistently outperforms R, but on the SAR+LMH, we see opposite results. This is probably because our method and the balancing-data method SSL could learn the positive bias resulted from the superficial correlations between question category and answer, which is useful for generalization, but the ensemble-based method LMH will attenuate positive bias during de-biasing process. (3) Even without language priors method, SAR with R→C rivals or outperforms the SAR+SSL and SAR+LMH with R or C, which shows that R→C strategy could help the model to alleviate language priors. As a result, compared with R or C, our framework with R→C only gains a slight performance improvement after using the same language-priors methods.

4.3.5 Ablation Study

“CAS+” represents we use the select-and-rerank framework. From Table 4, we can find that: (1) LXM+SSL represents directly applying SSL to LXMERT. Its poor performance shows that the major contribution of our framework does not come from the combination of the language-priors method SSL and pre-trained model LXMERT. (2) Compared with LXM and LXM+SSL, CAS+LXM and CAS+LXM+SSL respectively gain prominent performance boost of 9.35% and 6.32%, which demonstrates the importance and effectiveness of our proposed select-and-rerank procedure. (3) CAS+LXM+QTD(R) and CAS+LXM+SSL+QTD(R) respectively outperform CAS+LXM(R) and CAS+LXM+SSL(R) by 3.93% and 2.71%, which shows the contribution of QTD module. This further demonstrates that choosing appropriate N' for different question types is a useful step for model performance. (4) CAS+LXM+SSL+QTD improves the performance of CAS+LXM+QTD by 2.61%, which shows that

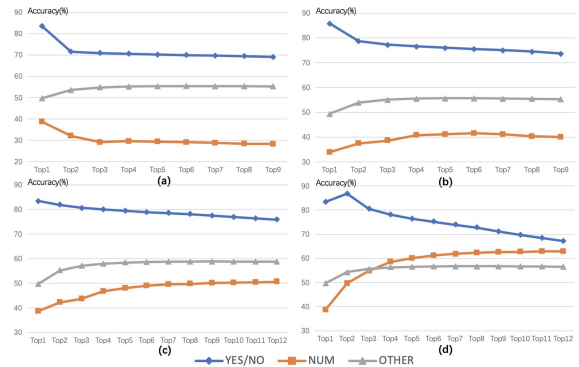


Figure 4: Results from SAR(R), SAR+SSL(R), SAR(R→C) and SAR+LMH(R) with different N' during test. To better investigate the impact of N' on each question type, we report the results without Question Type Discriminator.

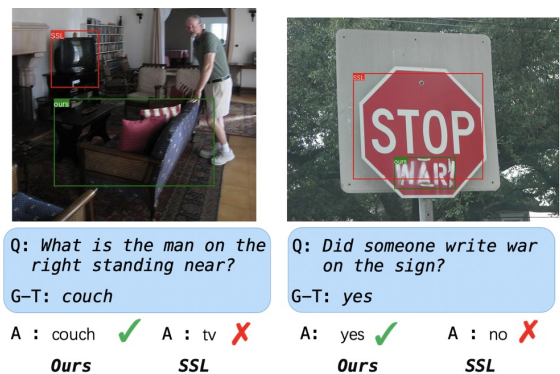


Figure 5: Qualitative comparison between our Top20-SAR(R→C) and the baseline SSL. The green/red bounding boxes indicate the most important regions resulting from ours/SSL. G-T is ground-truth.

current language-priors methods fit our framework well and could further improve performance.

4.3.6 The Effect of N'

From Figure 4, we can find that: (1) The best N' for yes/no questions is smaller than that for non-yes/no questions due to the nature of yes/no question. (2) As N' increases, the accuracy of “Num” and “Other” questions rises first and then decreases. There is a trade-off behind this phenomenon: when N' is too small, the correct answer may not be recalled by CAS; when N' is too large, the distraction from wrong answers makes it more difficult for model to choose the correct answer.

4.3.7 Qualitative Examples

We qualitatively evaluate the effectiveness of our framework. As shown in Figure 5, compared with SSL, SAR performs better not only in question answering but also in visual grounding. With the

help of answer semantics, SAR can focus on the region relevant to the candidate answer and further use the region to verify its correctness.

5 Conclusion

In this paper, we propose a select-and-rerank (SAR) progressive framework based on Visual Entailment. Specifically, we first select candidate answers to shrink the prediction space, then we rerank candidate answers by a visual entailment task which verifies whether the image semantically entails the synthetic statement of the question and each candidate answer. Our framework can make full use of the interactive information of image, question and candidate answers. In addition, it is a generic framework, which can be easily combined with the existing VQA models and further boost their abilities. We demonstrate advantages of our framework on the VQA-CP v2 dataset with extensive experiments and analyses. Our method establishes a new state-of-the-art accuracy of 66.73% with an improvement of 7.55% on the previous best.

Acknowledgments

This work was supported by National Natural Science Foundation of China (No. 61976207, No. 61906187)

References

- Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. 2016. Analyzing the behavior of visual question answering models. In *EMNLP*.
- Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. 2018. 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.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. 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.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433.
- Yonatan Belinkov, Adam Poliak, Stuart M Shieber, Benjamin Van Durme, and Alexander M Rush. 2019. Don't take the premise for granted: Mitigating artifacts in natural language inference. In *ACL (1)*.
- Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. 2019. Rubi: Reducing unimodal biases for visual question answering. *Advances in Neural Information Processing Systems*, 32:841–852.
- Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. 2020a. Counterfactual samples synthesizing for robust visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10800–10809.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020b. Uniter: Universal image-text representation learning. In *European Conference on Computer Vision*, pages 104–120. Springer.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2019. 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.
- Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. 2020. 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.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. 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.
- Gabriel Grand and Yonatan Belinkov. 2019. Adversarial regularization for visual question answering: Strengths, shortcomings, and side effects. *NAACL HLT 2019*, page 1.
- Allan Jabri, Armand Joulin, and Laurens Van Der Maaten. 2016. Revisiting visual question answering baselines. In *European conference on computer vision*, pages 727–739. Springer.
- Justin Johnson, Andrej Karpathy, and Li Fei-Fei. 2016. Denscap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Bernhard Kratzwald, Anna Eigenmann, and Stefan Feuerriegel. 2019. Rankqa: Neural question answering with answer re-ranking. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6076–6085.

- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. 2020. Oscar: Object-semantic aligned pre-training for vision-language tasks. In *European Conference on Computer Vision*, pages 121–137. Springer.
- Zujie Liang, Weitao Jiang, Haifeng Hu, and Jiaying Zhu. 2020. 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.
- Rabeeh Karimi Mahabadi and James Henderson. 2019. Simple but effective techniques to reduce biases. *arXiv preprint arXiv:1909.06321*, 2(3):5.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Yanyuan Qiao, Zheng Yu, and Jing Liu. 2020. Rankvqa: Answer re-ranking for visual question answering. In *2020 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE.
- Sainandan Ramakrishnan, Aishwarya Agrawal, and Stefan Lee. 2018. Overcoming language priors in visual question answering with adversarial regularization. In *NeurIPS*.
- Ramprasaath R Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Shalini Ghosh, Larry Heck, Dhruv Batra, and Devi Parikh. 2019. 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.
- Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5103–5114.
- Shuohang Wang, Mo Yu, Jing Jiang, Wei Zhang, Xiaoxiao Guo, Shiyu Chang, Zhiguo Wang, Tim Klinger, Gerald Tesauro, and Murray Campbell. 2018a. Evidence aggregation for answer re-ranking in open-domain question answering. In *International Conference on Learning Representations*.
- Zhen Wang, Jiachen Liu, Xinyan Xiao, Yajuan Lyu, and Tian Wu. 2018b. Joint training of candidate extraction and answer selection for reading comprehension. In *ACL (1)*.
- Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. 2019. Visual entailment: A novel task for fine-grained image understanding. *arXiv preprint arXiv:1901.06706*.
- Peng Zhang, Yash Goyal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2016. Yin and yang: Balancing and answering binary visual questions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5014–5022.
- Xi Zhu, Zhendong Mao, Chunxiao Liu, Peng Zhang, Bin Wang, and Yongdong Zhang. 2020. Overcoming language priors with self-supervised learning for visual question answering.