

Loss Re-Scaling VQA: Revisiting the Language Prior Problem From a Class-Imbalance View

Yangyang Guo¹, Liqiang Nie¹, *Senior Member, IEEE*, Zhiyong Cheng², Qi Tian³, *Fellow, IEEE*,
and Min Zhang⁴, *Member, IEEE*

Abstract—Recent studies have pointed out that many well-developed Visual Question Answering (VQA) models are heavily affected by the language prior problem. It refers to making predictions based on the co-occurrence pattern between textual questions and answers instead of reasoning upon visual contents. To tackle this problem, most existing methods focus on strengthening the visual feature learning capability to reduce this text shortcut influence on model decisions. However, few efforts have been devoted to analyzing its inherent cause and providing an explicit interpretation. It thus lacks a good guidance for the research community to move forward in a purposeful way, resulting in model construction perplexity towards overcoming this non-trivial problem. In this paper, we propose to interpret the language prior problem in VQA from a class-imbalance view. Concretely, we design a novel interpretation scheme whereby the loss of mis-predicted frequent and sparse answers from the same question type is distinctly exhibited during the late training phase. It explicitly reveals why the VQA model tends to produce a frequent yet obviously wrong answer, to a given question whose right answer is sparse in the training set. Based upon this observation, we further propose a novel loss re-scaling approach to assign different weights to each answer according to the training data statistics for estimating the final loss. We apply our approach into six strong baselines and the experimental results on two VQA-CP benchmark datasets evidently demonstrate its effectiveness. In addition, we also justify the validity of the class imbalance interpretation scheme on other computer vision tasks, such as face recognition and image classification.

Index Terms—Visual question answering, language prior problem, class imbalance, loss re-scaling.

I. INTRODUCTION

VISION and language are two ubiquitous elements in human cognition. One key artificial intelligence effort in bridging these two is to answer natural language questions about a visual scene, the *de facto* Visual Question Answering (VQA) task. VQA has gained great progress over the past few years owing to the considerable development of computer

vision and natural language processing. Conventional methods generally cast it as a **classification problem**, where the image and question are respectively handled via Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), and the fused multi-modal features are classified into candidate answers [1], [2].

Though a lot of VQA approaches have achieved great success over various benchmarks [1], [3], one notorious issue is that they are often hindered by the widely investigated language prior problem. That is, these models answer questions without comprehensively exploiting the visual contents. In fact, existing VQA methods are often brittle and easy to be fooled by some statistical shortcuts owing to the linguistically superficial correlation between questions and answers. For example, the *how many*¹ questions are frequently be answered with answer 2 in the training set [1]. This simple correlation misleads models to overwhelmingly reply to *how many* questions with 2 yet not truly reasoning the given image (see the training answer distribution and prediction in Figure 1 for an example). As a result, it is problematic to assert that the previous benchmark progress is driven by the faithful visual understanding rather than merely learning the dataset bias. To tackle this, a diagnostic benchmark, i.e., VQA-CP (VQA under Changing Priors) [4] has recently been constructed. It re-splits VQA v1 [1] and VQA v2 [3] datasets in a way such that the distribution of answers per question type is significantly distinct between training and testing sets. As can be expected, the performance of many prior state-of-the-art VQA models [2], [5], [6] largely degrades on this dataset.

It is worth mentioning that great success has been achieved by prior efforts pertaining to overcoming the language prior problem in VQA. These methods can be roughly grouped into two categories: *single-branch* and *two-branch*. Specifically, approaches in the first category directly strengthen the visual feature learning for weakening the textual shortcut influence on model decision [7]–[10]. For instance, HINT [7] and SCR [10] leverage additional visual annotations to align the image region-level importance.² Though yielding competitive improvement over their rivals, nonetheless, collecting such human annotations is expensive and burdensome, limiting their practicality in adaptation across datasets. By contrast, the two-branch methods deliver dominated performance due to their explicit counterwork over the language priors [12]–[15]. A typical design is to append an additional question-only

Manuscript received October 29, 2020; revised July 10, 2021; accepted November 9, 2021. Date of publication November 30, 2021; date of current version December 8, 2021. This work was supported by the National Natural Science Foundation of China under Grant U1936203. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Heng Tao Shen. (*Corresponding author: Liqiang Nie.*)

Yangyang Guo and Liqiang Nie are with the School of Computer Science and Technology, Shandong University, Qingdao 266237, China (e-mail: guoyang.eric@gmail.com; nieliqiang@gmail.com).

Zhiyong Cheng is with the Shandong Artificial Intelligence Institute, Qilu University of Technology (Shandong Academy of Sciences), Jinan 250014, China (e-mail: jason.zy.cheng@gmail.com).

Qi Tian is with Huawei Cloud & AI, Shenzhen 518100, China (e-mail: tian.qi1@huawei.com).

Min Zhang is with the School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen 518055, China (e-mail: zhangminmt@hotmail.com).

Digital Object Identifier 10.1109/TIP.2021.3128322

1941-0042 © 2021 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See <https://www.ieee.org/publications/rights/index.html> for more information.

¹We refer *question type* to the first few words in the given question.

²The image region importance is expressed via the model gradient magnitude of the ground-truth answers, i.e., Grad-CAM [11], with human attention map.

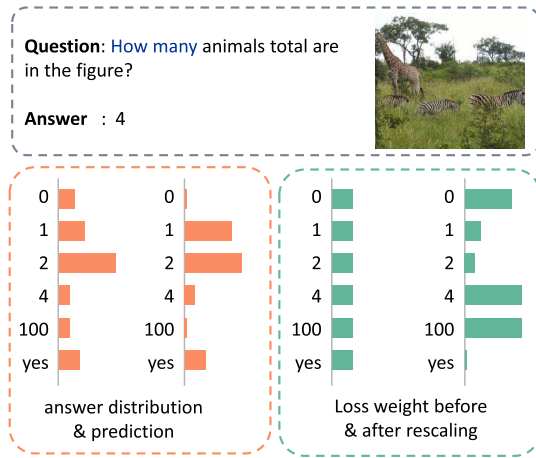


Fig. 1. Illustration of why the language prior problem arises and our solution. On the left orange bars, 1) the answer distribution in the training set for the current question type *how many* is biased, which leads the answer prediction to drift towards frequent answers including an out-of-type one - *yes*. On the right green bars, 2) the original loss weight is evenly distributed over each candidate answer, while our solution is to re-scale it for the model decision balancing.

training branch to the orthodox question-image one, which intentionally captures the language priors and is further suppressed by the latter branch. The computational cost from this category of methods is thereby increased since another new question-only branch is introduced to the backbone model.

In general, the pipeline of the existing studies can be summarized into two consecutive steps: 1) describing the background and consequence of the language prior problem in VQA, and 2) developing effective methods to mitigate it and demonstrating the superiority over their contenders. However, little attention has been paid to interpreting why it drives models to blindly answer questions even though the corresponding image shows very intuitive cues for the right answer. Without an interpretation, the SOTA results on the VQA-CP benchmark mostly manifest to be skeptical. In other words, it is dubious that the performance improvement is from the alleviation of the language priors, or the over-fitting on the new curated datasets. Moreover, another imperative issue of the existing approaches is that they often involve many parameters to the backbone model, resulting in heavier computations. Taking the more advanced two-branch approaches for example, the model at least increases the parameters for the answer prediction of the question-only branch (e.g., proportional to the product of the candidate answer size and the number of question features), which is burdensome and demands further optimization.

To deal with the above two issues, in this paper, we propose to bridge the connection between the language prior in VQA and the class-imbalance problem, and make contributions from the following two aspects. Firstly, to analyze the inherent reasons that cause the language prior problem in VQA models, we offer to revisit the problem from a class-imbalance view and contribute a novel viewpoint to interpret it. Recall that each candidate answer is taken as a single class and the answer distribution under each question type is highly biased. This observation is consistent with that in the class-imbalance classification scenario when restricting the same question type for each instance. In the light of this, we argue that the language

prior problem is essentially aroused by the imbalanced answer class distribution per question type. Accordingly, the VQA models are severely affected by the direct correlation between the question type and its frequent answers, leaving the visual features overlooked for answer prediction. To validate this point, we make two assumptions: 1) the loss values are diverse from the trained models pertaining to answers with different numbers in the training set; and 2) the gradient norm of the parameters from the question encoding layer is much larger than that of the classification layer. The former assumption interprets why the model cannot fit the tail classes (sparse answers per question type) even with more training iterations. And the latter one points out that the parameters from the question encoding layer should be the blame for the language prior learning. This seamlessly matches the application of previous two-branch approaches [12], [14] that overcomes this problem through tuning the question encoding parameters, and supports these methods. After validating both assumptions, we then propose a loss re-scaling approach, which borrows the idea from class-imbalance solutions. The key of our approach is to assign a distinct weight to each answer for estimating the instance loss based on its associated question type. Specifically, smaller weights are assigned to frequent answers while larger weights are attached to sparse ones. Notably, the proposed method is model-agnostic, which means it can be plugged into most existing methods to overcome the language prior problem. Besides, the loss re-scaling introduces **no** training parameters, i.e., **zero** incremental inference time to the backbone model. Compared to previous methods, the increased computational cost is thus negligible.

To testify the effectiveness of the proposed loss re-scaling approach, we conduct extensive experiments with six backbone methods on the VQA-CP v2 and VQA-CP v1 datasets. The experimental results demonstrate the enhanced performance when applying our method to the backbone models under various settings. Moreover, we also evaluate the viability of the class-imbalance interpretation scheme on other Computer Vision (CV) tasks which are also hindered by the class-imbalance problem.

In summary, the contributions of this paper are three-fold:

- We revisit the language prior problem in VQA from a class-imbalance view. The proposed interpretation scheme is practical and can be extended to other relevant CV fields. To the best of our knowledge, this is the first work to provide interpretations for the cause of the language prior problem in VQA.
- Based on this novel viewpoint, we further design a simple yet effective loss re-scaling method to alleviate the language prior effect in VQA.
- Extensive experiments are performed on several benchmark datasets and promising results from the proposed method can be observed. In addition, the code has been released to facilitate other researchers.³

II. RELATED WORK

In this section, we discuss three directions of related literature which are highly relevant to this work: Visual Question

³<https://github.com/guoyang9/class-imbalance-VQA>

Answering, Language Prior Problem in VQA and Data Imbalance Issue.

A. Visual Question Answering

Visual Question Answering constitutes a typical vision-and-language task [16]–[20], involving the inputs of an image and a natural language question. Almost all the existing VQA models follow such a training paradigm: the image and the question are encoded with CNNs and RNNs, respectively, followed by a multi-modal fusion module [21] and a classifier to predict the right answer(s). Based on this paradigm, current approaches can be roughly classified into four categories [22]: *Joint Embedding*, *Attention Mechanism-based*, *Compositional Models*, and *Knowledge-enhanced*. The *Joint Embedding* methods [1], [23] present to encode the image and question into a common latent space, in which the visual reasoning is performed. Later on, the *attention mechanism* is introduced to focus on salient image regions [6], [24] or with question words [5], [25]. The modular structure of questions is also exploited for more explicit reasoning of *Compositional VQA models* [26]. Lastly, it is natural that some questions cannot be directly answered from the image information merely, requiring accessing external knowledge for accurate prediction. *Knowledge Enhanced* approaches [27], [28] thus integrate knowledge bases such as DBpedia [29] or ConceptNet [30] into VQA.

Beyond the simple reasoning on images, researchers have also explored other auxiliary information and built more challenging benchmarks. For instance, Visual Commonsense Reasoning (VCR) requires a VQA model to provide an accurate answer as well a rationale (expressed via text) explaining why the answer is chosen. TextVQA asks models to read and reason the scene text in images. And video question answering [31]–[33] extends the question answering procedure from images to videos.

B. Language Prior Problem in VQA

A large body of studies pointed out that many VQA models are greatly affected by the detrimental language prior problem [4], [12], [34]–[36]. To address it, research have devoted efforts into the following two sides.

1) *From the Data Side*: It is extremely challenging to circumvent the dataset bias when crowd-sourcing with human annotators. In VQA v1 [1], Antol *et al.* collected the first large-scale VQA dataset, which manifests a lot of ‘visual priming biases’. For example, when people ask the question ‘Is there a clock tower in the picture?’, there actually contain clock towers in the given image. This poses negative effects on the VQA model learning since a balanced dataset should at least involves images without clock towers for this kind of questions. To amend this problem, [3], [37] later presented the VQA v2 dataset which introduces complementary samples in such a way that each question is associated with a pair of images yet different answers. However, [4], [34] indicate that this dataset is still limited by the bias problem, which can potentially drive VQA models to learn some language priors. In view of this, VQA-CP [4] was created by re-splitting

datasets following the rule that the answer distribution of training and testing sets is distinct with respect to question types. The performance of many VQA models drops dramatically on the VQA-CP datasets. Different from [4] re-organizing the existing biased datasets, Johnson *et al.* [38] curated a new diagnostic 3D shape dataset, to control the question-conditional bias via rejection sampling within families of related questions. The answer distribution is therefore balanced.

2) *From the Model Side*: Studies in this side make efforts to directly designing mechanisms for reducing the language prior effect, which can be summarized into two groups: single-branch and two-branch models. Specifically, single-branch methods are proposed to directly enhance the visual feature learning in VQA. For example, HINT [7] and SCR [10] align the image region-level importance with the additional human attention maps. VGQE [9] considers the visual and textual modalities equally when encoding the question, where the question features include not only the linguistic information from the question but also the visual information from the image. Zhu *et al.* [8] proposed a self-supervised framework to generate balanced question-image pairs. Differently, two-branch methods intend to explicitly counter the language prior learned by the model. A common processing is to introduce a question-only branch for deliberately capturing the language priors, followed by the question-image branch training to restrain it. For instance, Q-Adv [12] trains the above two models in an adversarial way, which minimizes the loss of the question-image model while maximizing the question-only one. More recent fusion-based methods [13]–[15] employ the late fusion strategy to combine the two answer predictions and guide the model to focus more on the answers which cannot be correctly addressed by the question-only branch.

Though the language prior problem in VQA has been extensively studied in recent years, nevertheless, few work contributes to analyzing its inherent cause, i.e., what arouses this problem is largely untapped in literature. Hence, in this work, we tentatively proposed to interpret the problem from a class-imbalance view, and developed a simple yet effective approach to tackle it.

C. Data Imbalance Issue

The data imbalance issue has been long recognized as a troublesome problem in a variety of research fields. Recent efforts mainly pursue solutions along three directions [39]: 1) *Data re-weighting* methods attach different weights to classes for estimating the final loss. A common solution is to assign large weights for tail classes (i.e., less frequent ones) and small weights for head classes (i.e., more frequent ones). For example, Focal loss [40] is proposed to automatically put more focus on hard and mis-classified samples. 2) *Data re-sampling* approaches target at re-sampling the highly biased data to attain a more balanced distribution. In particular, over-sampling methods [41] simply repeat the minority classes several times to alleviate the training bias, while under-sampling ones [42] randomly choose a subset of the majority classes. Nevertheless, the two sampling methods all have some drawbacks: over-sampling is easy to over-fit and under-sampling undermines the model generalization capability. And

TABLE I
TOY EXAMPLE OF WHY THE LANGUAGE PRIOR PROBLEM OCCURS UNDER
THE QUESTION TYPE *How Many*

ground truth	proportion in training set	prediction	loss value	endowed name
2	80%	4	large	hard mistake
4	4%	2	small	easy mistake

3) *Transfer learning*. Different from the above two lines of efforts, methods along this line aim to transfer the knowledge learned from abundant head classes to the tail ones [43]. The knowledge can be intra-class variance [44] or deep semantic features [45].

III. REVISITING THE LANGUAGE PRIOR PROBLEM

The language prior problem in VQA has been extensively studied in recent years, however, few work tries to interpret it from the angle of model learning, which makes the advancement on benchmarks unreliable to some extent. In light of this, we intend to fill this gap by revisiting it from a class-imbalance view. The reason is that current VQA models often treat VQA as a classification problem and the answer distribution per question type is highly imbalanced. In the following, we start with two assumptions and then empirically justify them. We further extend this interpretation scheme to two other CV tasks.

A. New Viewpoint

We firstly present a novel viewpoint for the class-imbalance issue according to the loss computation.

Assumption 1: The loss value of hard mistake samples is larger than that of easy mistake ones during the late training phase (when the learning converges).

We use a toy example in Table I to illustrate this assumption.⁴ As shown in Table I, 2 takes 80% of the answers to questions under the question type *how many*, while 4 accounts for only 4% in the training set. After the model is trained and the learning goes convergence, we assume that the loss value of these instances which are incorrectly predicted to 4 from 2 (hard mistake) is much larger than that of the ones predicted to 2 from 4 (easy mistake) - the name is given according to the relative number of true and predicted answers. Specifically, during the early training phase, the loss of both easy and hard mistakes is numerically similar. However, when it comes to the late training phase or the learning converges, the study on how the class imbalance issue affects the model prediction is particularly sparse in literature to our knowledge. In fact, there demands a deeper analysis on whether the loss computation is consistent between the early and late training phases. It is because that if the model learning is special during late training, it can provide us with insights on how the class-imbalance problem affects the model prediction when testing, and enlighten us to design directed methods.

Towards this end, we would like to empirically study it in this work. Particularly, we leverage a trained model, e.g., UpDn [2], to observe the loss value with respect to different

⁴Note that we only consider those instances which are mis-predicted, since the loss from the correctly predicted ones would be zero and back-propagate no gradients.

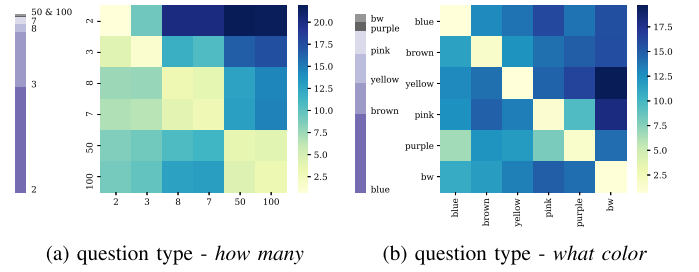


Fig. 2. Answer distribution and loss values from two biased question types *how many* on the left and *what color* on the right. The bar figure corresponds to the proportion of each answer. Regarding the confusion matrix, the x-axis denotes the ground-truth answers from frequent to sparse, y-axis represents the predicted answers with the same order, and value signifies the loss. Note that the numbers of each case are kept identical for fair analysis.

mistakes. To begin with, we fix the model parameters when the learning goes convergence. Thereafter, we reuse the instances in the training set and save the resultant prediction. The statistics and results are demonstrated in Figure 2. It can be observed that the loss values from the confusion matrix in the upper triangle are much larger than that of the lower one. Especially for *how many* questions, the values predicted from sparse answers (i.e., 50, 100) to the most frequent one 2 are the highest among others, denoting that the model penalizes these hard mistakes to a great extent. In contrast, the loss values of easy mistakes (corresponding to the lower triangle) tend to be much smaller, signifying that the model seems to overlook them during the late training. This finding is quite beneficial to understand the class-imbalance problem and we will show its validity in other domains in Section III-B.

Based on the first assumption, we further analyze the overall loss values and gradient norms with respect to training iterations. The second assumption is given as:

Assumption 2: The gradient norm from the parameters of the question encoding layer is larger than that of the prediction layer.

It is widely accepted that parameters with larger gradient norm often go with more updating when performing back-propagation. This assumption indicates that the question encoding layer affects more to the language prior learning than the prediction layer. To validate this, we plot the gradient norm from two groups of parameters with respect to the training iterations in Figure 3. It can be seen that the gradient norm of the parameters from the question encoding layer continues to increase, while the parameters from the prediction layer decreases a lot with more iterations. Note that the VQA models are affected more by the language prior problem in the late training phase. As a result, the parameters fluctuate sharper should contribute more to the learning of language priors. Namely, the question encoding layer should be the blame for the language prior problem, rather than the prediction one. This observation is consistent with most two-branch studies [12], [14], which aim to tune the question encoder for reducing the language priors. And this is the first time to evidently interpret why we should follow such a paradigm to alleviate this problem.

B. Extension to Other CV Tasks

To explore whether the first class-imbalance interpretation scheme works in other highly imbalanced CV scenarios,

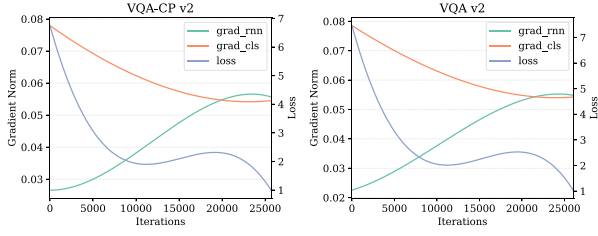


Fig. 3. Gradient norm of two groups of parameters with respect to training iterations. *grad_rnn* denotes the gradient norm from the question encoding layer parameters, while *grad_cls* corresponds to the prediction layer.

we additionally tested the face recognition and image classification tasks and show the results in Figure 4. The method and dataset employed for face recognition are respectively the Center Loss [46] the LFW dataset (Labeled Face in the Wild) [47], where we sampled seven persons with distinctive numbers in the training set. While for the image classification, we leveraged the imbalanced CIFAR-10 dataset [48] and the Resnet [49] Network for evaluation. From Figure 4, we can observe that for both tasks, the results show a similar manifestation with the ones in Figure 2. For example, when incorrectly predicting from the sparse *Chavez* class to other more frequent ones in face recognition, i.e., hard mistakes, the loss values become much larger than others. This can draw the parameters to avoid these mistakes and predict to other classes albeit yielding incorrect predictions. Similar observations can also be found on the image classification task. Another interesting point from the imbalanced image classification one is that the model often ignores the mistakes from more frequent classes while penalizes the ones from sparser ones.

IV. PROPOSED METHOD

In this section, we present our method to overcome the language prior problem in VQA. Our method borrows the idea from the approaches attacking the class-imbalance issue in other domains. Specifically, in this work, we adopt the loss re-weighting approach to deal with this problem. In particular, we design a novel mechanism to assign different weights over different answers (classes) under each question type to estimate the final loss. Before delving into the details of the loss re-scaling method, we firstly recap the basic learning functions in VQA.

A. Preliminary

The goal of VQA is to provide an accurate answer \hat{a} to a given question Q upon an image I , which can be achieved by,

$$\hat{a} = \arg \max_{a \in A} p(a|Q, I; \Theta), \quad (1)$$

where Θ denotes the model parameters and A represents all the candidate answers. Existing methods mostly tackle it as a classification task. There are actually two available solutions to predict the correct answer(s): Sigmoid activation function followed by Binary Cross Entropy loss and Softmax activation function followed by the Cross Entropy loss. We use Sigm-BCE and Soft-CE to denote the two solutions, respectively. It is worth noting the ground-truth for each image-question pair is not necessarily unique, that is, one question is answered by ten annotators and then the label is softened by

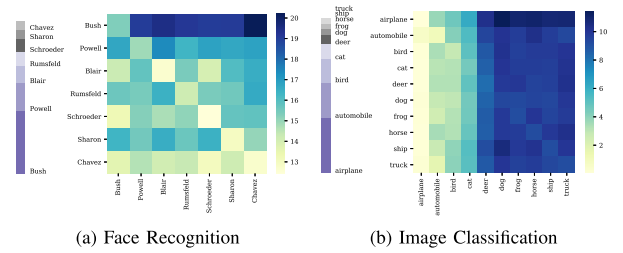


Fig. 4. Class distribution and loss values from two CV tasks. The left one corresponds to seven persons from the LFW face recognition dataset. And the right one denotes the ten classes in the imbalanced CIFAR-10 dataset.

ten. In this way, for Sigm-CE, taking each candidate answer as a single label for the binary classification problem,⁵ is given as,

$$\begin{aligned} p &= W_p h_f + b_p, \\ \hat{p} &= \sigma(p), \\ \mathcal{L}_{bce} &= - \sum_{i=1}^{|A|} a_i \log \hat{p}_i + (1 - a_i) \log(1 - \hat{p}_i), \end{aligned} \quad (2)$$

where h_f is the fused multi-modal feature, $\sigma(\cdot)$ represents the Sigmoid activation function, a_i denotes the ground-truth answer label, W_p and b_p are the learnable matrix and bias, respectively.

In contrast, Soft-CE tackles the multi-label classification in VQA via,

$$\begin{aligned} p &= W_p h_f + b_p, \\ \hat{p} &= \text{Softmax}(p), \\ \mathcal{L}_{ce} &= - \sum_{i=1}^{|A|} a_i \log \hat{p}_i, \end{aligned} \quad (3)$$

where $\text{Softmax}(\cdot)$ implies the Softmax activation function.

Based on the aforementioned two solutions, we propose to assign distinctive weights for each answer when computing the final loss. To this end, in the following, we firstly design an *Answer Mask* module to constrain the predicted answers fall into the correct set pertaining to each question type. We then employ a novel strategy to compute the answer weights for the loss estimation.

B. Answer Mask

Let's revisit the example in Figure 1, for the question type *how many*, there even exists an out-of-type answer *yes*. This perplexes a lot as very rare *how many* questions should be answered with binary answers *yes* or *no*. One possible reason is that the *yes* answer takes large parts of the answers, which makes the model unexpectedly drift to this frequent label. To overcome this problem and ensure the predicted answers lie into the specific answer set of the given question type, we introduce an answer mask in this section. The motivation is supported by the findings in Figure 5, where only a very small number of instances in the validation set are distributed to the out-of-type answers from the corresponding answer set. Concretely, the answer mask module takes the question feature

⁵Notice that we omit the mini-batch summation in all loss functions for simplicity.

as inputs, and then employs a fully-connected layer equipped with a non-linear activation function,

$$\mathbf{m}_p = f(\mathbf{W}_q \mathbf{q}), \quad (4)$$

where the size of $\mathbf{m}_p \in \mathbb{R}^{|A|}$ is the same with the candidate answer size, $f(x)$ denotes the activation function, and \mathbf{W}_q are the learnable weights. To enable the answer mask modeling, we collect the true mask label $\mathbf{m}_a \in \mathbb{R}^{|A|}$ where the labels belonging to the given question type are set to 1 and others are 0. In this way, the BCE loss is used to train this module,

$$\mathcal{L}_{mask} = \sum_{i=1}^{|A|} m_a^i \log m_p^i + (1 - m_a^i) \log(1 - m_p^i), \quad (5)$$

where m_a^i and m_p^i represent the i -th element in \mathbf{m}_a and \mathbf{m}_p , respectively.

In fact, a straightforward design of $f(x)$ is to adopt the sigmoid activation function, which can map the features to a mask value within $(0, 1)$. However, we favor the output to be larger than 0.5 in most times so that the final predicted answer would be grounded more by this mask. Note that deploying all-ones mask will lead to over-fitting and is not flexible. We therefore devise a generalized softplus function in Equation 6,

$$f(x) = \max(1, \frac{1}{\alpha} \log(1 + \exp \alpha x)), \quad (6)$$

where α is a hyper-parameter which is fixed to 1 in our experiment.

C. Loss Re-Scaling

The key to our loss re-scaling approach is to attach distinctive weight to each answer. To achieve this, we intuitively introduce a fixed weight, where it can prevent the hard mistakes to update the parameters harshly while guide the easy mistakes for more influential model learning. In particular, the weight is obtained via,

$$\mu_i = \frac{\sum_k^{|A|} n_k^j - n_i^j}{n_i^j}, \quad (7)$$

where n_i^j represents the number of answer a_i under the question type qt_j (for other answers not in the answer set of the current question type, μ_i is equal to 1). Under this situation, recollect the previous example in Table I, the loss weights for answer ‘2’ and ‘4’ are $(1 - 0.8)/0.8 = 0.25$, and $(1 - 0.04)/0.04 = 24$, respectively. In this way, the easy mistakes become relatively hard to propagate information, while hard mistakes gain more chance to update parameters, especially for earlier layers. In implementation, we found that large weights made the training more unstable, we therefore employ a smoothing function and truncate too large values to be 100,

$$\mu_i = \max(100, \log(1 + \exp \mu_i)). \quad (8)$$

This approach is specially designed for avoiding VQA models to learn language priors when training. During testing, the loss re-scaling is deactivated and only the backbone

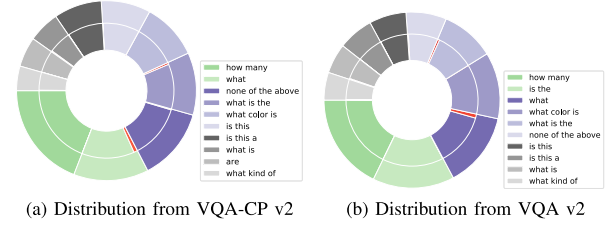


Fig. 5. Answer distribution in the training set (outer circle) and the validation set (inner circle) of the most frequent question types. The red samples represent the answers which are out of the corresponding answer set in the training set.

model remains. This assures that there is **no incremental computation** at all when evaluation compared with the backbone model.⁶ In addition, we provide the losses and partial derivatives in Appendix.

D. Method Application

Most mainstream VQA architectures are compatible with our method. To show how the above two modules work, we leverage a general backbone and illustrate its learning pipeline and ours in Figure 6. Specifically, for each input image, it utilizes the pre-trained object detection networks (e.g., Faster R-CNN [50]) to obtain the most salient objects; meanwhile, the question words are sequentially processed via a textual encoder (e.g., Gated Recurrent Network [51]) to capture the semantic information. The visual reasoning is then performed, followed by the classifier for predicting answers. As can be seen in Figure 6a, the BCE losses for answers 4, 2 and yes are 1.0, 18.0 and 5.0, respectively. With our answer mask module in Figure 6b, the yes answer no longer takes effects since it is an out-of-type outlier for the question type *how many*. Our loss re-scaling module then assigns distinct loss weights to each candidate answer according to the statistics in Table I. One can see that the loss for the sparse answer 4 is augmented while for the frequent answer 2 is diminished.

V. EXPERIMENTAL SETUP

We conducted extensive experiments on three benchmark datasets to validate the effectiveness of the proposed method. In particular, the experiments are mainly performed to answer the following research questions:

- **RQ1:** Can the proposed loss re-scaling method improve the state-of-the-art VQA models?
- **RQ2:** How do the answer mask and loss re-scaling modules perform on both solutions in Section IV-A?
- **RQ3:** Is the proposed loss re-scaling method superior than other loss re-weighting approaches?
- **RQ4:** Why does the proposed method outperform the backbone method?

In the following, we first provide the basic information of the evaluated benchmark datasets. We then detail the standard evaluation metric, followed by the key baselines used in this work, and we end this section with some implementation details.

⁶By zero incremental inference time, we explicitly refer to the loss re-scaling module excluding the answer mask one, since the loss re-scaling module produces a primary effect as shown in Table V.

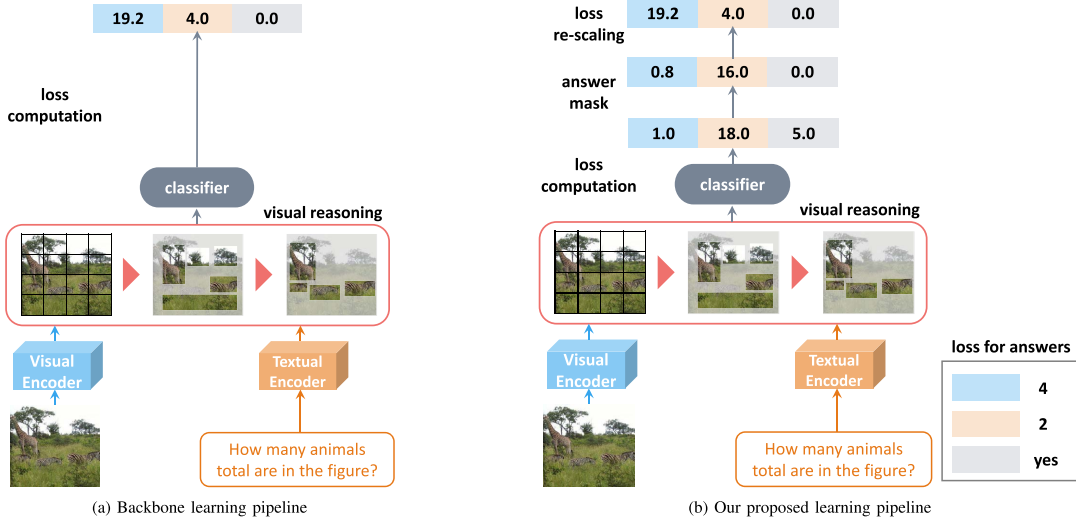


Fig. 6. Visual comparison of loss estimation between the backbone learning pipeline and ours. We use three answers, i.e., 4, 2 and yes, to intuitively illustrate how the proposed two modules work.

TABLE II

ACCURACY COMPARISON BETWEEN THE PROPOSED METHOD AND BASELINES OVER THE VQA-CP v2 TEST AND VQA v2 VAL SETS. REGARDING THE METHOD CATEGORY, ‘PLAIN’ REPRESENTS VQA METHODS WITHOUT SPECIALIZATION FOR OVERCOMING THE LANGUAGE PRIOR PROBLEM. ‘TWO-BRANCH’ AND ‘SINGLE-BRANCH’ DISTINGUISH WHETHER THERE IS ANOTHER QUESTION-ONLY TRAINING BRANCH. ‘—’ AND ‘†’ DENOTE THE NUMBER IS NOT AVAILABLE AND OUR IMPLEMENTATION, RESPECTIVELY. SYMBOLS ‘†’ DENOTES THE STATISTICAL SIGNIFICANCE WITH TWO-SIDED T-TEST OF $p < 0.01$, COMPARED TO THE CORRESPONDING BACKBONE. THE BEST PERFORMANCE IN CURRENT SPLITS IS HIGHLIGHTED IN BOLD

Method Category	Method	VQA-CP v2 test				VQA v2 val			
		Y/N	Num.	Other	All	Y/N	Num.	Other	All
Plain	Question-only [4]	35.09	11.63	07.11	15.95	67.95	30.97	27.20	43.01
	SAN [6]	38.35	11.14	21.74	24.96	70.06	39.28	47.84	52.41
	NMN [26]	38.94	11.92	25.72	27.47	73.38	33.23	39.93	51.62
	MCB [52]	41.01	11.96	40.57	36.33	77.91	37.47	51.76	59.71
	HAN [53]	52.25	13.79	20.33	28.65	-	-	-	-
	UpDn [2]	42.27	11.93	46.05	39.74	81.18	42.14	55.66	63.48
	UpDn [2]†	42.31	11.84	44.35	38.80	80.59	40.51	53.69	62.07
Two-branch	LXMERT [54]‡	46.70	27.14	61.20	51.78	88.34	56.64	65.78	73.06
	AdvReg [12]	65.49	15.48	35.48	41.17	79.84	42.35	55.16	62.75
	Rubi [14]	68.65	20.28	43.18	47.11	-	-	-	61.16
	LM [15]‡	73.65	41.44	42.82	51.68	80.37	40.62	53.53	61.92
	LMH [15]‡	70.29	44.10	44.86	52.15	74.18	38.43	53.05	59.07
	CSS [13]	43.96	12.78	47.48	41.16	-	-	-	-
	CSS+LMH [13]	84.37	49.42	48.21	58.95	73.25	39.77	55.11	59.91
Single-branch	CSS [13]‡	40.88	12.32	45.23	38.95	80.28	41.90	53.81	62.20
	CSS+LMH [13]‡	77.73	50.92	45.30	55.61	42.16	35.05	52.01	46.08
	GVQA [4]	57.99	13.68	22.14	31.30	72.03	-	-	48.24
	HINT [7]	67.27	10.61	45.88	46.73	81.18	42.99	55.56	63.38
Ours	SCR [10]	72.36	10.93	48.02	49.45	78.8	41.6	54.5	62.2
	VGQE [7]	66.35	27.08	46.77	50.11	-	-	-	63.18
	UpDn+Ours	68.42†	21.71†	42.88†	47.09†	64.22†	39.61†	53.09†	55.50†
	LM+Ours	73.74†	45.14†	44.59†	53.17†	76.28†	36.61†	52.71	59.45†
	LMH+Ours	72.82	48.00†	44.46†	53.26†	68.21†	36.37†	52.29†	56.81†
	CSS+Ours	68.45†	40.79†	44.18	50.73†	77.73†	41.39†	53.76†	61.14
	CSS+LMH+Ours	83.95	47.81†	44.59†	56.55†	70.52†	33.56†	50.83†	55.96†
	LXMERT+Ours	79.77†	59.06†	61.41	66.40†	85.32†	52.07†	62.60†	69.76†

A. Datasets

We evaluated our proposed method mainly on the two VQA-CP datasets: VQA-CP v2 and VQA-CP v1 [4], which are well-known benchmarks for evaluating the models’ capability to overcome the language prior problem. The VQA-CP v2 and VQA-CP v1 datasets consist of $\sim 122\text{K}$ images, $\sim 658\text{K}$ questions and $\sim 6.6\text{M}$ answers, and $\sim 122\text{K}$ images, $\sim 370\text{K}$ questions and $\sim 3.7\text{M}$ answers, respectively. Moreover, the answer distribution per question type is significantly different

between training and testing sets. Other than VQA-CP, we also reported results on the VQA v2, a more biased benchmark, for completeness. For all the datasets, the answers are divided into three categories: *Y/N*, *Num.* and *Other*.

B. Evaluation Metric

We adopted the standard metric in VQA for evaluation [1]. For each predicted answer a , the accuracy is computed as,

$$Acc = \min(1, \frac{\#\text{humans that provide answer } a}{3}). \quad (9)$$

Note that each question is answered by ten annotators, and this metric takes the human disagreement into consideration [1], [3].

C. Tested Backbones

We mainly tested our method over the following six backbones. Thereinto, UpDn is the most popular VQA baseline, LXMERT is a recently proposed cross-modality transformer-based framework. LM, LMH, CSS and CSS+LMH are developed to intentionally alleviate the language prior problem in VQA.

UpDn [2] firstly leverages the pre-trained object detection frameworks to obtain salient object features for high-level reasoning. It then employs a simple attention network to focus on the most important objects which are highly related with the given question.

LXMERT [54] is built upon the Transformer encoders to learn the connections between vision and language. It is pre-trained with diverse pre-training tasks on a large-scale dataset of image and sentence pairs. It is worth noting that we did not apply the answer mask module to LXMERT to avoiding undermining its pre-trained linguistic features.

LM [15] is a typical two-branch method which takes UpDn as baseline and combine the answer prediction from the two branches (i.e., question-image and question-only) in an ensemble mode.

LMH [15] extends LM by introducing another entropy loss to encourage the question-only branch prediction to be non-uniform, posing a greater impact on the ensemble.

CSS [13] generates counterfactual training samples by masking critical image objects or question words. These samples are then trained with negative labels.

CSS+LMH [13] combines CSS and LMH together to achieve better performance.

D. Implementation Details

We implemented our method based on the publicly available codes for all the backbones. For the answer mask module, we appended a dropout layer with the dropout rate tuned from 0.0 to 1.0 with a step size of 0.1. And the answer mask loss weight is fixed to 1.0. For the loss re-scaling approach, we first trained the answer mask module until converging, and then fine-tuned the model with the re-scaling weights for another ten epochs. Note that we kept all the other settings being unchanged, such as the learning rate, optimizer, mini-batch size.

VI. EXPERIMENTAL RESULTS

A. Overall Performance Comparison (RQ1)

Table II and III present the performance results on VQA-CP v2 testing, VQA v2 validation and VQA-CP v1 testing sets. The key observations from these two tables are as follows:

- Our methods achieve the best over most answer categories on the VQA-CP v2 and VQA-CP v1 datasets. Since these two datasets are deemed as an effective protocol for evaluating the language prior overcoming capability,

TABLE III
ACCURACY COMPARISON BETWEEN THE PROPOSED METHOD AND BASELINES OVER THE VQA-CP v1 DATASET. ‘—’ AND ‘†’ DENOTE THE NUMBER IS NOT AVAILABLE AND OUR IMPLEMENTATION, RESPECTIVELY. SYMBOLS ‘†’ DENOTES THE STATISTICAL SIGNIFICANCE WITH TWO-SIDED T-TEST OF $p < 0.01$, COMPARED TO THE CORRESPONDING BACKBONE. THE BEST PERFORMANCE IN CURRENT SPLITS IS HIGHLIGHTED IN BOLD

Method	VQA-CP v1 test			
	Y/N	Num.	Other	All
Question-only [4]	35.72	11.07	08.34	20.16
SAN [6]	35.34	11.34	24.70	26.88
NMN [26]	38.85	11.23	27.88	29.64
MCB [52]	37.96	11.80	39.90	34.39
UpDn [2]†	43.76	12.49	42.57	38.02
GVQA [4]	64.72	11.87	24.86	39.23
AdvReg [12]	74.16	12.44	25.32	43.43
LM [15]†	75.01	28.86	42.22	53.59
LMH [15]†	76.61	29.05	43.38	54.76
CSS+LMH [13]	85.60	40.57	44.62	60.95
CSS [13]†	43.15	12.79	43.88	38.36
CSS+LMH [13]†	83.55	29.79	43.59	57.86
LXMERT [54]†	54.08	25.05	62.72	52.82
UpDn+Ours	44.25†	18.04†	43.03†	39.34†
LM+Ours	75.51†	26.93†	43.67†	54.07†
LMH+Ours	78.26†	29.80†	42.76	55.32
CSS+Ours	74.39†	37.02†	36.22†	52.20†
CSS+LMH+Ours	84.56†	32.89†	43.86	58.91†
LXMERT+Ours	86.82†	45.03†	62.03†	69.47†

we can conclude that the proposed method is capable of dealing with this issue and superior than the existing methods.

- For all the six backbones, with our answer mask and loss re-scaling modules, they can consistently obtain a significant performance improvement on the VQA-CP datasets. For example, in the VQA-CP v2 dataset, the overall accuracy of UpDn boosts from 38.80% to 47.09%, with a 8.29% absolute gain.
- The performance gap between VQA v2 val and VQA-CP v2 test in Table II quantifies the bias learning level since the VQA v2 is a more biased dataset. It can be seen that the methods specially designed for tackling the language prior problem outperform the traditional methods, and our approaches perform even better, indicating less language priors are learned.
- For the comparison of two-branch and single-branch methods, we can observe that two-branch ones often outperform their contemporary rivals, such as AdvReg v.s. GVQA, Rubi v.s. HINT, and LMH even surpasses the succeeding method VGQE. One noticeable reason for this is that the question-only model in the two-branch methods explicitly captures the language priors and is then suppressed by the question-image model.

In addition, we further computed the Cohen’s kappa to measure the agreement between labels and predictions in Table IV. From the results, it can be observed that our method outperforms the backbones in most cases. However, the kappa values mostly manifest to be small. One possible reason to this is that the number of classes is too large and all the three datasets are negatively affected by the bias distribution to some extent.

TABLE IV
COHEN’S KAPPA COMPARISON BETWEEN BACKBONE MODELS AND OUR METHODS ON THREE BENCHMARK DATASETS

Method	VQA-CP v2 test				VQA-CP v1 test				VQA v2 val			
	Y/N	Num.	Other	All	Y/N	Num.	Other	All	Y/N	Num.	Other	All
UpDn	47.86	6.81	0.80	4.96	43.01	4.40	0.53	8.57	49.96	14.13	0.72	8.14
UpDn+Ours	55.61	11.89	0.78	5.66	63.90	38.68	0.75	14.24	51.46	15.34	0.81	9.25
LM	50.51	24.79	0.92	6.06	53.52	17.48	0.59	11.34	49.49	13.28	0.73	8.19
LM+Ours	50.61	24.87	0.92	6.98	56.35	13.60	0.58	11.34	48.85	10.56	0.68	5.91
LMH	50.97	24.99	0.85	5.42	55.55	16.70	0.60	11.03	49.88	11.59	0.67	7.56
LMH+Ours	53.27	26.03	0.76	5.74	54.70	19.10	0.59	11.21	49.60	9.89	0.65	6.15
CSS	47.67	7.02	0.81	4.90	41.66	4.54	0.57	8.57	50.19	14.27	0.74	8.26
CSS+Ours	51.79	20.98	0.78	5.19	66.42	29.16	0.67	13.44	51.21	14.37	0.76	8.97
CSS+LMH	50.58	28.62	0.81	5.19	59.89	16.18	0.61	11.00	45.91	10.52	0.65	4.76
CSS+LMH+Ours	51.97	29.21	0.82	6.23	61.33	19.60	0.59	11.93	49.75	10.03	0.64	7.64
LXMERT	46.32	8.83	0.83	4.78	43.22	6.28	0.59	8.51	50.03	13.17	0.69	7.99
LXMERT+Ours	49.95	22.51	0.85	5.43	56.14	17.94	0.63	11.25	50.02	11.47	0.62	7.88

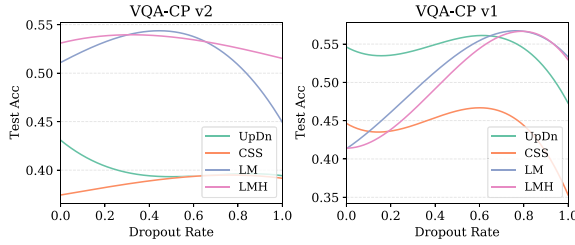


Fig. 7. Testing Acc of four evaluated methods over two datasets with respect to the dropout rate.

B. Ablation Study (RQ2)

We conducted detailed experiments to validate the effectiveness of each module and reported the results in Table V and Figure 7. Note that the LM and LMH are particularly developed based on BCE loss function which cannot be adapted with the Soft-CE solution. And we did not apply the answer mask to LXMERT for avoiding undermining its pre-trained linguistic features. We can find that:

- The loss re-scaling scheme is effective under a series of settings. When directly applying it to the backbones, the accuracy is enhanced with a large margin. For instance, CSS and LXMERT achieve 8.6% and 14.62% absolute gains under the BCE loss, respectively. With the pre-training of the answer mask module, our method can further obtain a higher accuracy.
- For the answer mask module, the softplus-G activation function does surpass the sigmoid one. The reason is that the mask tends to be larger than 0.5 with a high probability, making less answers out of the corresponding answer set in the training set.
- We further evaluated the dropout rate effect in the answer mask module and illustrated the results in Figure 7. It can be observed that UpDn prefers a smaller dropout rate, while LM, CSS and LMH fit larger ones.
- Surprisingly, our methods can attain a substantially better results over the backbones when using the CE loss function. For example, the final absolute performance improvement of our method over the UpDn backbone is over 16.53%! It is even better than the recent strong baselines LMH and CSS+LMH. This highlights the effectiveness of the loss re-scaling scheme and the potential

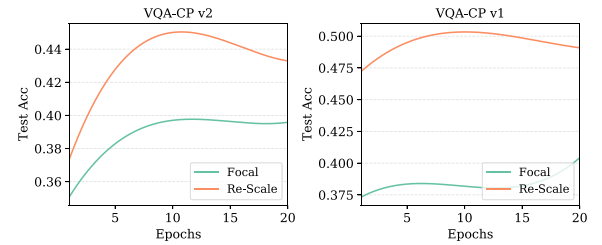


Fig. 8. Testing Acc of loss re-weighting methods on the UpDn baseline over two datasets with respect to training epochs.

of the class-imbalance view on alleviating the language prior problem in VQA.

C. Other Loss Re-Weighting Method (RQ3)

To examine whether the loss re-scaling scheme is superior than other commonly used loss re-weighting approaches, we adopted the Focal Loss [40], [55] and tested its performance in VQA. Focal Loss is designed to put more focus on hard, mis-classified examples, where the loss weights are auto-learned from the training data,

$$FL(p_i) = -\alpha(1 - p_i)^\gamma \log(p_i), \quad (10)$$

where α and γ are hyper-parameters and p_i denotes the prediction for binary classification. We extended this loss function to the VQA multi-class case and plotted the results in Figure 8. We can observe that there exists a quite large performance margin between our method and the Focal Loss. One possible reason is that our loss weights is particularly computed under each question type, which provides more explicit guidance than the auto-learning strategy in Focal Loss.

D. Case Study (RQ4)

We visualized the attention map and predicted answers of UpDn and our method in Figure 9. Specifically, for the first, second and the last instances, the backbone all provides the incorrect answers due to the language priors. The reason is that answers *black*, *baseball* and *2* are the dominated answers for question types *what color*, *what* and *how many*, respectively. For the last instance, the attention map from the baseline focuses on both the *man* and the *child*, which misleads the model to predict that there are 2 children. In contrast, our

TABLE V

EFFECTIVENESS VALIDATION OF THE ANSWER MASK AND LOSS RE-SCALING MODULES. ‘-’ REPRESENTS THE NUMBER IS NOT AVAILABLE, SINCE LM AND LMH DIRECTLY PERFORMS ON THE BCE LOSS. AND SOFTPLUS-G IMPLIES THE ACTIVATION FUNCTION IN EQUATION 6

Solution	Module	UpDn	LM	LMH	CSS	CSS+LMH	LXMERT
Sigm-BCE	baseline	38.80	51.68	52.15	38.95	55.61	51.78
	+answer mask (sigmoid)	40.94	51.71	51.56	38.76	56.12	-
	+answer mask (softplus-G)	41.69	52.76	52.70	38.72	56.55	-
	+loss re-scaling	40.84	51.96	52.68	47.55	50.76	66.40
	+both	47.09	53.17	53.26	50.73	56.55	-
Soft-CE	baseline	40.49	-	-	35.23	-	58.07
	+answer mask (sigmoid)	39.32	-	-	35.21	-	-
	+answer mask (softplus-G)	40.30	-	-	34.80	-	-
	+loss re-scaling	55.93	-	-	52.05	-	69.37
	+both	57.02	-	-	53.50	-	-



Fig. 9. Visualization results from UpDn and our method on the VQA-CP v2 dataset. Regions with lighter colors denote what the model focuses on when making prediction.

method attends solely on the *child* object and yields the correct answer. For the third instance, the baseline only looks an asleep *cat* and concludes that both cats are asleep, while our method notices the awake *cat*. And for the fourth instance, though the baseline also predicts the right answer, nevertheless, its attention is focused on both the *women* and the *dog*. Instead, our method puts more attention on the animal beneath the bench - *dog*.

VII. CONCLUSION AND FUTURE WORK

There are few studies on analyzing the inherent cause of the language prior problem in VQA. In this work, we propose to fill this gap via interpreting it from a class-imbalance view. To begin with, we present two assumptions and empirically validate their viability. We then develop a simple yet effective loss re-scaling approach to attach distinct weights to each answer according to the given question type. Extensive experiments on three publicly available datasets validate the effectiveness of the proposed method.

In the future, we would like to exploit other typical methods overcoming the class-balance issue, i.e., data re-balancing or transfer learning, for alleviating the language prior problem.

Moreover, the class-imbalance view can be generalized to other vision-and-language tasks which are also hindered by the language bias problem, e.g., image captioning.

APPENDIX

LOSSES AND PARTIAL DERIVATIVES

Based on the re-scaling weight, we further derive the losses and its partial derivatives over the prediction for both the Soft-CE and Sigm-BCE below. With this, the gradient of other learnable parameters can be easily computed.

1) For Sigm-BCE, the loss function becomes,

$$\mathcal{L}_{bce}^{\mu} = - \sum_{i=1}^{|A|} \mu_i (a_i \log \sigma(p_i) + (1 - a_i) \log(1 - \sigma(p_i))), \quad (11)$$

where the partial derivative can be obtained by,

$$\begin{aligned} \frac{\partial \mathcal{L}_{bce}^{\mu}}{\partial p_i} &= -\mu_i (a_i \frac{1}{\sigma(p_i)} (\sigma(p_i)(1 - \sigma(p_i))) \\ &\quad + (1 - a_i) \frac{1}{1 - \sigma(p_i)} (-\sigma(p_i)(1 - \sigma(p_i)))) \\ &= \mu_i ((1 - a_i) \sigma(p_i) - a_i (1 - \sigma(p_i))). \end{aligned} \quad (12)$$

2) For the Soft-CE, the loss function is given as,

$$\mathcal{L}_{ce}^{\mu} = - \sum_{i=1}^{|A|} \mu_i a_i \log \frac{\exp p_i}{\sum_{j=1}^A \exp p_j}. \quad (13)$$

To compute the partial derivative p_i of from \mathcal{L}_{ce} , we firstly provide the partial derivative of the Softmax function,

$$\frac{\partial \hat{p}_i}{\partial p_i} = \begin{cases} \hat{p}_i(1 - \hat{p}_j), & \text{if } i == j \\ -\hat{p}_i \hat{p}_j, & \text{otherwise.} \end{cases} \quad (14)$$

And then the partial derivative is given by,

$$\begin{aligned} \frac{\partial \mathcal{L}_{ce}^{\mu}}{\partial p_i} &= - \sum_{i=1}^{|A|} \mu_i a_i \frac{\partial \log \hat{p}_i}{\partial p_i} \\ &= - \sum_{i=1}^{|A|} \mu_i a_i \frac{1}{\hat{p}_i} \frac{\partial \hat{p}_i}{\partial p_i} \\ &= \mu_i a_i (\hat{p}_i - 1) + \sum_{k \neq i} \mu_k a_k \hat{p}_i \\ &= (\mu_i a_i + \sum_{k \neq i} \mu_k a_k) \hat{p}_i - \mu_i a_i \\ &= \sum_{k=1}^{|A|} \mu_k a_k \hat{p}_i - \mu_i a_i, \end{aligned} \quad (15)$$

where line 3 is derived from Equation 14. In this way, the final loss becomes,

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{mask}, \quad (16)$$

where the \mathcal{L}_{cls} can be either \mathcal{L}_{ce}^{μ} or \mathcal{L}_{bce}^{μ} .

REFERENCES

- [1] S. Antol *et al.*, "VQA: Visual question answering," in *Proc. ICCV*, Dec. 2015, pp. 2425–2433.
- [2] P. Anderson *et al.*, "Bottom-up and top-down attention for image captioning and visual question answering," in *Proc. CVPR*, Jun. 2018, pp. 6077–6086.
- [3] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh, "Making the V in VQA matter: Elevating the role of image understanding in visual question answering," in *Proc. CVPR*, Jul. 2017, pp. 6904–6913.
- [4] A. Agrawal, D. Batra, D. Parikh, and A. Kembhavi, "Don't just assume; look and answer: Overcoming priors for visual question answering," in *Proc. CVPR*, Jun. 2018, pp. 4971–4980.
- [5] J. Lu, J. Yang, D. Batra, and D. Parikh, "Hierarchical question-image co-attention for visual question answering," in *Proc. NeurIPS*, 2016, pp. 289–297.
- [6] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola, "Stacked attention networks for image question answering," in *Proc. CVPR*, 2016, pp. 21–29.
- [7] R. R. Selvaraju *et al.*, "Taking a HINT: Leveraging explanations to make vision and language models more grounded," in *Proc. ICCV*, Oct. 2019, pp. 2591–2600.
- [8] X. Zhu, Z. Mao, C. Liu, P. Zhang, B. Wang, and Y. Zhang, "Overcoming language priors with self-supervised learning for visual question answering," in *Proc. IJCAI*, Jul. 2020, pp. 1–7.
- [9] K. V. Gouthaman and A. Mittal, "Reducing language biases in visual question answering with visually-grounded question encoder," 2020, *arXiv:2007.06198*.
- [10] J. Wu and R. Mooney, "Self-critical reasoning for robust visual question answering," in *Proc. NeurIPS*, 2019, pp. 1–11.
- [11] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. ICCV*, Oct. 2017, pp. 618–626.
- [12] S. Ramakrishnan, A. Agrawal, and S. Lee, "Overcoming language priors in visual question answering with adversarial regularization," in *Proc. NeurIPS*, 2018, pp. 1–11.
- [13] L. Chen, X. Yan, J. Xiao, H. Zhang, S. Pu, and Y. Zhuang, "Counterfactual samples synthesizing for robust visual question answering," in *Proc. CVPR*, Jun. 2020, pp. 10800–10809.
- [14] R. Cadene, C. Dancette, H. Ben-Younes, M. Cord, and D. Parikh, "RUBI: Reducing unimodal biases for visual question answering," in *Proc. NeurIPS*, 2019, pp. 1–12.
- [15] C. Clark, M. Yatskar, and L. Zettlemoyer, "Don't take the easy way out: Ensemble based methods for avoiding known dataset biases," in *Proc. EMNLP*, 2019, pp. 1–14.
- [16] Z. Zhao, Z. Zhang, X. Jiang, and D. Cai, "Multi-turn video question answering via hierarchical attention context reinforced networks," *IEEE Trans. Image Process.*, vol. 28, no. 8, pp. 3860–3872, Aug. 2019.
- [17] J. Ji, C. Xu, X. Zhang, B. Wang, and X. Song, "Spatio-temporal memory attention for image captioning," *IEEE Trans. Image Process.*, vol. 29, pp. 7615–7628, 2020.
- [18] M. Zhang, Y. Yang, H. Zhang, Y. Ji, H. T. Shen, and T.-S. Chua, "More is better: Precise and detailed image captioning using online positive recall and missing concepts mining," *IEEE Trans. Image Process.*, vol. 28, no. 1, pp. 32–44, Jan. 2019.
- [19] L. Peng, Y. Yang, Z. Wang, Z. Huang, and H. T. Shen, "MRA-Net: Improving VQA via multi-modal relation attention network," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Jun. 25, 2020, doi: [10.1109/TPAMI.2020.3004830](https://doi.org/10.1109/TPAMI.2020.3004830).
- [20] L. Peng, Y. Yang, X. Zhang, Y. Ji, H. Lu, and H. T. Shen, "Answer again: Improving VQA with cascaded-answering model," *IEEE Trans. Knowl. Data Eng.*, early access, Jun. 1, 2020, doi: [10.1109/TKDE.2020.2998805](https://doi.org/10.1109/TKDE.2020.2998805).
- [21] Y. Guo, Z. Cheng, L. Nie, X.-S. Xu, and M. Kankanhalli, "Multi-modal preference modeling for product search," in *Proc. MM*, 2018, pp. 1865–1873.
- [22] Q. Wu, D. Teney, P. Wang, C. Shen, A. Dick, and A. van den Hengel, "Visual question answering: A survey of methods and datasets," *Comput. Vis. Image Understand.*, vol. 163, pp. 21–40, Oct. 2017.
- [23] M. Malinowski, M. Rohrbach, and M. Fritz, "Ask your neurons: A neural-based approach to answering questions about images," in *Proc. ICCV*, Dec. 2015, pp. 1–9.
- [24] V. Kazemi and A. Elqursh, "Show, ask, attend, and answer: A strong baseline for visual question answering," 2017, *arXiv:1704.03162*.
- [25] Z. Yu, J. Yu, J. Fan, and D. Tao, "Multi-modal factorized bilinear pooling with co-attention learning for visual question answering," in *Proc. ICCV*, Oct. 2017, pp. 1821–1830.
- [26] J. Andreas, M. Rohrbach, T. Darrell, and D. Klein, "Neural module networks," in *Proc. CVPR*, 2016, pp. 39–40.
- [27] Q. Wu, P. Wang, C. Shen, A. Dick, and A. van den Hengel, "Ask me anything: Free-form visual question answering based on knowledge from external sources," in *Proc. CVPR*, Jun. 2016, pp. 4622–4630.
- [28] P. Wang, Q. Wu, C. Shen, A. Dick, and A. van den Hengel, "FVQA: Fact-based visual question answering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 10, pp. 2413–2427, Oct. 2018.
- [29] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. Ives, "DBpedia: A nucleus for a web of open data," in *Proc. ISWC*. Springer, 2007.
- [30] H. Liu and P. Singh, "ConceptNet—A practical commonsense reasoning tool-kit," *BT Technol. J.*, vol. 22, no. 4, pp. 211–226, Oct. 2004.
- [31] L. Zhu, Z. Xu, Y. Yang, and A. G. Hauptmann, "Uncovering the temporal context for video question answering," *Int. J. Comput. Vis.*, vol. 124, no. 3, pp. 409–421, Sep. 2017.
- [32] J. Lei, L. Yu, M. Bansal, and T. Berg, "TVQA: Localized, compositional video question answering," in *Proc. EMNLP*, 2018, pp. 1–13.
- [33] Z. Zhao, Q. Yang, D. Cai, X. He, and Y. Zhuang, "Video question answering via hierarchical spatio-temporal attention networks," in *Proc. IJCAI*, Aug. 2017, pp. 3518–3524.
- [34] Y. Guo, Z. Cheng, L. Nie, Y. Liu, Y. Wang, and M. Kankanhalli, "Quantifying and alleviating the language prior problem in visual question answering," in *Proc. SIGIR*, 2019, pp. 75–84.
- [35] C. Jing, Y. Wu, X. Zhang, Y. Jia, and Q. Wu, "Overcoming language priors in VQA via decomposed linguistic representations," in *Proc. AAAI*, 2020, pp. 11181–11188.
- [36] D. Teney, E. Abbasnejad, and A. van den Hengel, "Unshuffling data for improved generalization," 2020, *arXiv:2002.11894*.
- [37] P. Zhang, Y. Goyal, D. Summers-Stay, D. Batra, and D. Parikh, "Yin and Yang: Balancing and answering binary visual questions," in *Proc. CVPR*, Jun. 2016, pp. 5014–5022.

- [38] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, and R. Girshick, "CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning," in *Proc. CVPR*, Jul. 2017, pp. 2901–2910.
- [39] B. Kang *et al.*, "Decoupling representation and classifier for long-tailed recognition," in *Proc. ICLR*, 2020, pp. 1–16.
- [40] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, "Focal loss for dense object detection," in *Proc. ICCV*, Oct. 2017, pp. 2980–2988.
- [41] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," *J. Artif. Intell. Res.*, vol. 16, pp. 321–357, Jun. 2002.
- [42] C. Drummond and R. C. Holte, "C4.5, class imbalance, and cost sensitivity: Why under-sampling beats over-sampling," in *Proc. Workshop Learn. Imbalanced Datasets II*, 2003, pp. 1–8.
- [43] X.-S. Wei, P. Wang, L. Liu, C. Shen, and J. Wu, "Piecewise classifier mappings: Learning fine-grained learners for novel categories with few examples," *IEEE Trans. Image Process.*, vol. 28, no. 12, pp. 6116–6125, Dec. 2019.
- [44] X. Yin, X. Yu, K. Sohn, X. Liu, and M. Chandraker, "Feature transfer learning for face recognition with under-represented data," in *Proc. CVPR*, Jun. 2019, pp. 5704–5713.
- [45] Z. Liu, Z. Miao, X. Zhan, J. Wang, B. Gong, and S. X. Yu, "Large-scale long-tailed recognition in an open world," in *Proc. CVPR*, 2019, pp. 2537–2546.
- [46] Y. Wen, K. Zhang, Z. Li, and Y. Qiao, "A discriminative feature learning approach for deep face recognition," in *Proc. ECCV*, 2016, pp. 499–515.
- [47] G. B. Huang, M. Mattar, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," in *Proc. Workshop Faces 'Real-Life' Images, Detection, Alignment, Recognit.*, 2008.
- [48] H. Kang, T. Vu, and C. D. Yoo, "Learning imbalanced datasets with maximum margin loss," in *Proc. NeurIPS*, Sep. 2021, pp. 1269–1273.
- [49] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, 2016, pp. 770–778.
- [50] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. NeurIPS*, 2015, pp. 1–9.
- [51] K. Cho *et al.*, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in *Proc. EMNLP*, 2014, pp. 1–15.
- [52] A. Fukui, D. H. Park, D. Yang, A. Rohrbach, T. Darrell, and M. Rohrbach, "Multimodal compact bilinear pooling for visual question answering and visual grounding," in *Proc. EMNLP*, 2016, pp. 1–12.
- [53] M. Malinowski, C. Doersch, A. Santoro, and P. Battaglia, "Learning visual question answering by bootstrapping hard attention," in *Proc. ECCV*, 2018, pp. 3–20.
- [54] H. Tan and M. Bansal, "LXMERT: Learning cross-modality encoder representations from transformers," in *Proc. EMNLP*, 2019, pp. 1–14.
- [55] Q. Zhou *et al.*, "Fine-grained spatial alignment model for person re-identification with focal triplet loss," *IEEE Trans. Image Process.*, vol. 29, pp. 7578–7589, 2020.



AND DATA ENGINEERING (TKDE), and *ACM Transactions on Multimedia Computing, Communications, and Applications* (TOMM).

Yangyang Guo is currently pursuing the Ph.D. degree with the School of Computer Science and Technology, Shandong University. He is currently working as a Research Fellow with the National University of Singapore. He has published several papers in top conferences and journals, such as ACM MM and IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING (TKDE). He has served as a Regular Reviewer for journals, including IEEE TRANSACTIONS ON MULTIMEDIA (TMM), IEEE TRANSACTIONS ON KNOWLEDGE



as the ACM MM and SIGIR Best Paper Honorable Mention Award in 2019, the SIGMM Rising Star Award in 2020, the TR35 China 2020, the DAMO Academy Young Fellow in 2020, and the SIGIR Best Student Paper Award in 2021. He is an Associate Editor of IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING (TKDE), IEEE TRANSACTIONS ON MULTIMEDIA (TMM), *ACM Transactions on Multimedia Computing, Communications, and Applications* (TOMM), and *Information Science*. He is also the Area Chair of ACM MM from 2018 to 2021.

Liqiang Nie (Senior Member, IEEE) received the B.Eng. degree from Xi'an Jiaotong University and the Ph.D. degree from the National University of Singapore (NUS). After Ph.D. degree, he continued his research at NUS as a Research Fellow for three years. He is currently a Professor with Shandong University. He has coauthored more than 200 papers and four books. He received more than 12,000 Google Scholar citations. His research interests include multimedia computing and information retrieval. He has received many awards, such



TRANSACTIONS ON IMAGE PROCESSING (TIP), IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING (TKDE), and IEEE TRANSACTIONS ON MULTIMEDIA (TMM).

Zhiyong Cheng received the M.S. degree from Xi'an Jiaotong University in 2010 and the Ph.D. degree in computer science from Singapore Management University in 2016. He was a Research Fellow with the National University of Singapore. He is currently a Professor with the Shandong Artificial Intelligence Institute, Qilu University of Technology (Shandong Academy of Sciences). His research interests mainly focus on large-scale multimedia content analysis and retrieval. He has served as a Regular Reviewer for journals, including IEEE



a Full Professor with the Department of Computer Science, The University of Texas at San Antonio (UTSA), a Changjiang Chaired Professor of the Ministry of Education, an Overseas Outstanding Youth, and an Overseas Expert by the Chinese Academy of Sciences.

Qi Tian (Fellow, IEEE) received the B.E. degree in electronic engineering from Tsinghua University and the Ph.D. degree in electrical and computer engineering (ECE) from the University of Illinois at Urbana-Champaign (UIUC), Champaign, IL, USA. He was the Visiting Chaired Professor with the Center for Neural and Cognitive Computation, Tsinghua University, and a Lead Researcher with the Media Computing Group, Microsoft Research Asia (MSRA). He is currently the Chief Scientist of artificial intelligence with Huawei Cloud & AI,



received the bachelor's and Ph.D. degrees in computer science from the Harbin Institute of Technology, Harbin, in 1991 and 1997, respectively. He is currently a Professor with the Harbin Institute of Technology (Shenzhen), Shenzhen, China. His current research interests include natural language processing, multi-model, artificial intelligence, and cryptology.