



Estimating semantic structure for the VQA answer space

Corentin Kervadec, Grigory Antipov, Moez Baccouche, Christian Wolf

► To cite this version:

Corentin Kervadec, Grigory Antipov, Moez Baccouche, Christian Wolf. Estimating semantic structure for the VQA answer space. 2020. hal-02862763

HAL Id: hal-02862763

<https://hal.science/hal-02862763>

Preprint submitted on 9 Jun 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ESTIMATING SEMANTIC STRUCTURE FOR THE VQA ANSWER SPACE

Corentin Kervadec

Orange Labs, Cesson-Sévigné, France
Université de Lyon, INSA-Lyon, LIRIS
UMR CNRS 5205, Villeurbanne, France
corentin.kervadec@orange.com

Grigory Antipov

Orange Labs,
Cesson-Sévigné, France
grigory.antipov@orange.com

Moez Baccouche

Orange Labs,
Cesson-Sévigné, France
moez.baccouche@orange.com

Christian Wolf

Université de Lyon, INSA-Lyon, LIRIS
UMR CNRS 5205, Villeurbanne, France
christian.wolf@insa-lyon.fr

ABSTRACT

Since its appearance, Visual Question Answering (VQA, i.e. answering a question posed over an image), has always been treated as a classification problem over a set of predefined answers. Despite its convenience, this classification approach poorly reflects the semantics of the problem limiting the answering to a choice between independent proposals, without taking into account the similarity between them (e.g. equally penalizing for answering “cat” or “German shepherd” instead of “dog”). We address this issue by proposing (1) two measures of proximity between VQA classes, and (2) a corresponding loss which takes into account the estimated proximity. This significantly improves the generalization of VQA models by reducing their language bias. In particular, we show that our approach is completely model-agnostic since it allows consistent improvements with three different VQA models. Finally, by combining our method with a language bias reduction approach, we report SOTA-level performance on the challenging VQAv2-CP dataset.

Keywords VQA, vision and language, visual reasoning

1 Introduction

Visual Question Answering (VQA) is a task which requires to provide a textual answer given a question and an image as input. When properly formulated, this problem requires a high-level understanding of the content of the image as well as the problem statement (the question), and therefore is often considered to be a proxy task allowing to evaluate the visual reasoning abilities of a system.

While the problem itself requires the prediction of textual output (the answer word or sentence), which is an output space with rich structure, most, if not all known benchmarks and evaluation protocols deal with it as a classification problem, for instance VQAv1 [1], VQAv2 [2], GQA [3]. This is done by creating a dictionary of output answer classes constructed from the most frequent answers of the training set. Models addressing the problem are asked to predict a probability distribution over this dictionary and are usually trained with cross-entropy loss plus eventual self-supervised auxiliary losses.

The merit of this approach is its ease of use, its straightforward definition of loss function and its empirical success. It, however, raises two fundamental issues: (i) the dictionary highly depends on the training set and hurts generalization to unseen data; (ii) the answer classes of the dictionary are considered to be independent, without taking into account their semantic relationships. This results in models highly dependant on question biases, as observed in [4].

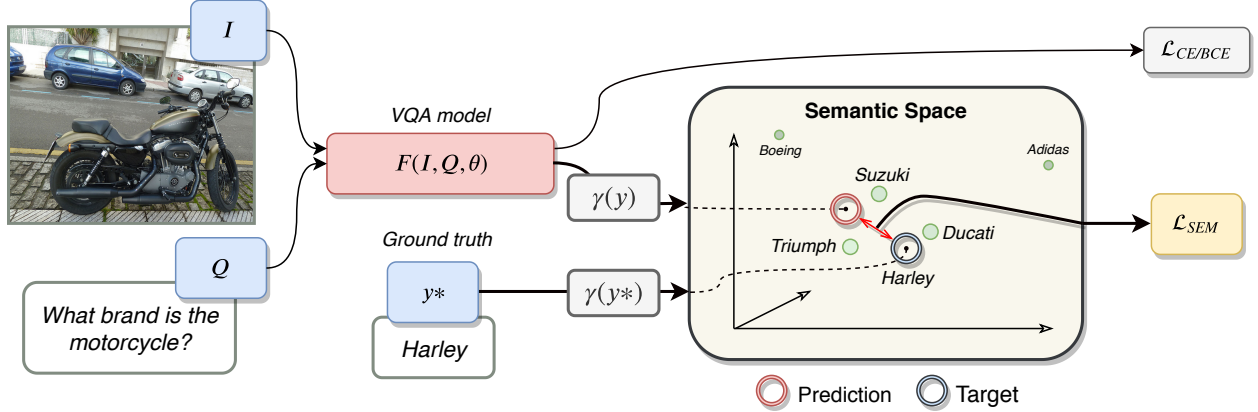


Figure 1: We introduce a new loss structuring the semantic space of VQA output classes. The prediction y and the target label y^* are both projected into a semantic space using $\gamma(\cdot)$. Then, the model is trained to minimize the distance between the prediction and the target in that space using our semantic loss \mathcal{L}_{SEM} .

While the long-term objective of the community is arguably to move to a direct structured prediction of the textual output, in this work we focus on the second issue (ii), arguing that properly structuring the semantic space of output classes can overcome some of the shortcomings of the classification strategy. We address this through a new loss function (that we call *semantic loss*), which measures the semantic distance between the prediction and the ground truth answer. If, for instance, the question is “Who is on the car?” and the expected answer is “woman”, we propose to penalize the wrong answer “girl” less than an also wrong answer “boy”, while the classical cross-entropy loss would penalize both wrong answers equally (see Figure 1).

We show that this new loss provides two different benefits: (i) the more direct benefit of pushing error cases to more favorable answer classes, like the “girl” instead of “boy” example above. This alone should provide a quality improvement in many applications, albeit not measured through the standard metrics *Accuracy*, *Recall*, *Precision*, not adapted to measure the reasoning abilities of an agent; (ii) we also show that structuring the output space improves performance in absolute terms as measured through *Accuracy*, therefore drives the model to make fewer errors whatever their type might be.

While the intuitive notion of semantic proximity is easy to grasp for a human, its formal definition is less clear. We propose two different methods to estimate the proximity of two answer words or sentences from data: the first one exploits the relationship between answer class and answer text and makes use of classical self-supervised pre-training of word embeddings, while the second one extracts statistics from the multiple redundant ground truth annotations of VQA datasets.

During training, we combine the new loss along with the standard cross-entropy discriminative loss. In the experiments we show that this improves accuracy on the VQA task and, in particular, helps reducing the dependency on language biases. More so, we show that the performance improvements obtained by our contribution are complementary to other efforts in removing language bias [5], providing strong combined performance. We believe this further indicates the interest of this new learning signal.

Our contributions are as follows:

1. We design a semantic loss for VQA, which helps the model to better structure its answer space and to learn the semantic relations between the answers.
2. We propose two different methods for estimating semantic proximity between answer classes based on word embeddings and annotation statistics, respectively.
3. We demonstrate the effectiveness of our method on the VQAv2-CP [4] and VQAv2 [2] datasets with three different neural models, showing consistency of the improvement over datasets and models.
4. When combined with other efforts in addressing language bias, our performance improvements add up and achieve performance on-par with the State-Of-The-Art (SOTA) on VQAv2-CP with reasonably complex model architectures.

2 Related work

VQA as a classification task — The broad majority of works are approaching VQA as a classification task. This strategy simplifies the supervision part of the training and makes the VQA approaches easily comparable between each other. Indeed, many VQA approaches, including attention-based networks [6], object-based attention mechanisms [7], bilinear fusion methods [8], and, more recently, Transformer [9]-based models [10], have been introduced and evaluated on the VQA classification benchmarks.

Biases in VQA datasets — The success of these works is in part due to the creation of large annotated VQA corpus. The VQAv1 [1] dataset gathers more than 200k real-world images annotated with 760K questions in natural language. Moreover, each question is annotated with 10 ground truth answers in order to model the ambiguity and the disagreement between annotators. Nevertheless, as the data collection is fastidious and costly, VQAv1 suffers from numerous language biases as pointed out by many works such as [2], and [11]. In order to downplay the influence of the language bias, [2] released an updated version of the same dataset (baptized VQAv2) by carefully balancing the answer distribution per question type. [12] went one step further by designing a fully synthetic dataset called CLEVR in which automatically generated questions are asked on synthetically generated 3D images. This corpus is thought as a diagnostic dataset aimed at evaluating the visual reasoning capability of the models. However, the limited environment of CLEVR [12] prevents from generalizing to complex realistic images such as in VQAv2 [2]. Finally, [3] built the GQA dataset, a semi-synthetic dataset where automatically generated questions are asked on real images, which can be seen as an intermediate step between CLEVR and VQAv2.

The generalization curse of VQA — Nevertheless, despite the efforts made on data collection, the language bias issue persists and VQA models continue to suffer from the generalization curse. Early works raised an alarm by diagnosing many drawbacks on VQA models, such as their tendency to pay little attention to the image and to “only read half of the question” [11]. Similarly, [13] showed that attention-based models do not attend to the same visual regions as humans do. More recently, [14] pointed out the gender biases learned by image captioning models. To better diagnose the generalization gap on VQA, [4] reorganized VQAv1 [1] and VQAv2 [2] into the VQA-CP (VQA under Changing Priors), a new dataset where the per-question answer distribution of the train split is made explicitly different from the one in the test split. In particular, they show that a blind model (which only has access to the question without seeing the image) achieves a surprisingly high accuracy on VQAv2 – 44% – whereas it reaches only 16% on VQAv2-CP [4]. Moreover, many of the successful models on VQA datasets fall short on VQA-CP unveiling their lack in visual understanding and their tendency to rely on question biases.

Reducing language biases on VQA — Therefore, several works have recently tried to tackle this generalisation issue. [15] trained their model using an adversarial game against a question-only adversary in order to discourage the base model to rely on language prior. Similarly, the authors of RUBi [5] added a question-only branch to the base model during the training to adapt its prediction in order to prevent it from learning question biases. Other methods make use of additional annotated supervision to improve the generalization capability. Using the VQA-HAT dataset’s annotations [13], the HINT [16] model is supervised to attend to the same visual regions as humans. [17] built upon HINT proposing even a more sophisticated approach by carefully designing a three-steps learning strategy, named Self-Critical Reasoning (SRC). SRC accentuates the model sensitivity to the important visual regions. It should be noted that SRC requires additional data annotations such as human attention maps [13] or textual explanations [18]. Finally, [19] introduced a Decomposed Linguistic Representation¹ (DLR) approach consisting in learning to decompose the question into the “type representation”, “object representation” and “concept representation”. Although it allows to improve the model’s accuracy on VQAv2-CP [4], it causes a significant drop of performance on VQAv2 [2]. In this paper, we contribute to these efforts as we demonstrate how a semantic loss, helping the model to structure its output space, permits to reduce dependency towards language biases.

Our method has also connections with distributed encoding approaches which have been successfully applied on age estimation from faces [20] and more generally on label distribution learning [21]. Additionally, an early VQA work, the DAQUAR [22] dataset, pioneers the use of soft evaluation for VQA. They use a variant of the Wu Palmer similarity [23] over a lexical database to compute a soft prediction score. This way, a prediction semantically close to the target answer is no longer considered as false. This allows to obtain finer evaluation of performances on VQA. However, it is only used for evaluation and such metric has fundamental drawbacks, such as its inability to discriminate colors.

¹To the best of our knowledge, at the time of writing of our paper, paper [19] has not yet been accepted to a peer-reviewed conference or journal.

3 Structuring the answer space

As our contribution is agnostic w.r.t. particular model architectures, we consider a VQA model as a whitebox function f taking into account an image I , a question q and producing an answer y , which is an output vector over the output alphabet:

$$y = f(I, q; \theta). \quad (1)$$

In the literature, the cross-entropy loss is frequently used to measure the prediction error during training, which casts the task as a classification problem with a single unique correct answer per problem instance:

$$\mathcal{L}_{CE} = CE(\text{softmax}(y), y^*), \quad (2)$$

$$= \sum_i^{N_c} y_i^* \log(\text{softmax}(y)_i) \quad (3)$$

where $CE(\cdot, \cdot)$ is the cross-entropy function, N_c is the size of the answer dictionary (hence the number of classes), and y^* is the one-hot encoded vector of ground truth answers. Alternatively, some datasets (such as VQAv1 [1] and VQAv2 [2]) admit more than one correct answer for a given question. This allows to take into account ambiguities in the question formulations and in the annotation uncertainties. In that case, an appropriate formulation is a model allowing to predict more than one answer class combined with a soft binary cross-entropy loss given as:

$$\mathcal{L}_{BCE} = \frac{1}{N_c} \sum_i^{N_c} BCE(\text{sigmoid}(y_i), y_i^*), \quad (4)$$

$$= \frac{1}{N_c} \sum_i^{N_c} y_i^* \cdot \log(\text{sigmoid}(y_i)) + (1 - y_i^*) \cdot \log(1 - \text{sigmoid}(y_i)) \quad (5)$$

where $BCE(\cdot, \cdot)$ is the binary cross-entropy, y_i and y_i^* are respectively the predicted and the ground truth probability of answer i of the dictionary.

\mathcal{L}_{CE} and \mathcal{L}_{BCE} are, both, widely used learning signals [7, 2, 10, 5, 8], and despite their difference, they share the common shortcoming of not taking into account eventual differences in the semantic proximity of answers.

To address this issue, in this work, we introduce a new semantic loss, which provides additional structure to the output space. Defining the semantic loss requires to set up (i) a semantic space which embeds the answers, and (ii) a distance function measuring the semantic similarity between two answers.

Projection to a semantic space — a semantic space is required to satisfy several proprieties. It needs:

- a - to be structured, i.e. to take into account semantic proximity estimated from some data source,
- b - to be able to cope with the continuous nature of neural networks, which provide continuous estimates for each output class (estimates of posterior probabilities if trained with cross-entropy) as opposed to discrete symbolic predictions.

We address requirement (b) by defining a function γ , which projects the continuous prediction and the discrete ground truth label, respectively, into a joint semantic space:

$$\gamma(y) = \sum_{i \in \text{top}_k(y)}^{N_c} y_i g(i), \quad (6)$$

The sum in (6) weights the different class mappings g by their predicted output, and is defined over the top- k highest predictions y_i , where k is a hyper-parameter. This allows the mapping to be dominated by the highest-probably answer classes and eliminate spurious non-probable influences.

The function γ depends on the mapping g , which needs to address requirement (a). We propose two different semantic spaces, which estimate the output space structure from two different data sources, namely, respectively, pre-trained word embeddings and redundancies in VQA dataset groundtruth annotations. These two learning signals have fundamentally different origins.

Glove — Word embeddings are widely-used projections from a discrete symbol space to a continuous vectorial space [24, 25]. They are trained in a self-supervised way, minimizing the error in predicting words from a context, i.e.

groups of words around the predicted target. In other words, the semantic space emerges as a side product of the estimation of co-occurrences in language space. The Glove representation [24] has been shown to capture fine-grained semantic and syntactic regularities, and as such is thus a natural choice for our semantic space. We use the GloVe embedding directly to project output class i to its vectorial representation:

$$g(i) = \text{Glove}(a_i) \quad (7)$$

where a_i is the textual representation of answer the class i .

Co-oc — While the word embeddings mentioned beforehand exploit statistical regularities in large text corpus to estimate semantic proximity, we propose an alternative which directly taps into human assessments. Annotations on semantic distances are hard to come by, so we provide estimates from alternative human annotations. The two datasets VQAv2 [2] and VQAv2-CP [4] provide 10 ground truth answers per question obtained by 10 different people. We define the semantic proximity of a pair of answers classes i and j as the amount of coherence between these two classes in terms of human annotations. To be more precise, we estimate it from the co-occurrences c_{ij} of the two answer classes in different question instances:

$$c_{ij} = \frac{\log |(i, j)|}{\log |(i)| \log |(j)|} \quad (8)$$

where (i) and (j) , are the sets of question instances where answers i and j , respectively, occurs, and (i, j) is the set of question instances where both answers i and j occurs. The Co-oc embedding vector is then defined as follows:

$$g(i) = [c_{i1}, \dots, c_{ij}, \dots, c_{iN_c}] \quad (9)$$

In Co-oc space, two answers are close if they are likely to be used as answers for a same question.

Differences of the spaces Glove and Co-oc — The two embeddings are of different nature, one being estimated from word occurrences in large language corpuses, the second one directly exploiting human annotation. While Co-oc has been estimated in a goal-driven way and, for this reason, arguably could be more adapted to structuring output spaces, we should note that its coverage is smaller. Glove is estimated from large-scale corpuses and it can therefore be expected that any reasonable combination of words has received a significant statistical support during its estimation. On the other hand, the size of the human annotations in the VQA datasets is limited, and equation (9) estimating the Co-oc vector is dominated by co-occurrences caused by ambiguities in question formulation and human interpretation of questions, content and reasoning. Pairs of answer classes, which are far from each other semantically, therefore receive a small support for statistical estimation making the estimates noisy for these cases. It is difficult to say which one of these two effects — statistical support and goal-driven nature of estimate — will take the upper hand. So, this question is answered in the experimental section.

However, in order to qualitatively confirm that the designed embeddings are correctly encoding the targeted semantic proximity, we show in Figure 2 the 2D t-SNE projection [26] of Glove and Co-oc embeddings corresponding to the answers classes in the VQAv2-CP dataset [4]. In Figure 2, we manually selected four classes' categories which are semantically close, namely *colors* (10 classes), *dog species* (9 classes), *motorcycle brands* (4 classes) and *tree species* (4 classes). As we can see in the Figure, these categories correspond to spatially grouped clusters in the 2D projection space, both for Glove and Co-oc embeddings, thus confirming their relevance.

Distances in the semantic space — To compensate for differences in normalization, we chose the cosine similarity as a measure for proximity in the embedding spaces:

$$s(x, y) = \frac{x \cdot y}{\|x\| \|y\|}, \quad (10)$$

The semantic loss \mathcal{L}_{SEM} penalizing miss-classifications between predictions \mathbf{y} and targets \mathbf{y}^* according to semantic proximity is then given as:

$$\mathcal{L}_{SEM} = 1 - s(\gamma(\mathbf{y}), \gamma(\mathbf{y}^*)), \quad (11)$$

As described in Figure 3, \mathcal{L}_{SEM} allows to take into account the semantic proximity between the prediction and the ground truth, whereas $\mathcal{L}_{CE/BCE}$ does not make the difference between wrong predictions which are close or unrelated to the ground truth (respectively case B and C in the Figure 3). Finally, we combine the semantic loss with the classical cross-entropy (or binary cross-entropy) as follows:

$$L = \mathcal{L}_{CE/BCE} + \lambda \mathcal{L}_{SEM}, \quad (12)$$

where λ is a hyper-parameter.

At test time, we remove the semantic loss component and predict an answer as the highest logit/probability. Our method can thus be applied to any VQA architecture, and only requires the supplementary loss \mathcal{L}_{SEM} during the training.




Figure 2: Visual comparison of the proposed semantic spaces for structuring the VQA class answers: Glove (on the left) and Co-oc (on the right). The embeddings of all classes from the VQAv2-CP dataset in the respective spaces are illustrated via t-SNE [26]. The big circles of various hues represent 4 different categories of the VQAv2-CP answer classes which have been chosen for the sake of illustration. The table lists the different categories, the color in which they are represented in the figure as well as the VQA answer classes of each category.

4 Experiments

We evaluate our contributions on the following two VQA corpuses:

VQAv2 has been proposed in [2] and contains open-ended questions in natural language about real images. The corpus gathers 265K images with at least 3 questions for each. Each question is annotated with 10 ground-truth answers.

VQAv2-CP has been introduced in [4]. It has been constructed by reorganizing the training and validation splits of the VQAv2 [2] dataset in order to explicitly make the distribution of answers different between the training and test. In other words, the VQAv2-CP dataset has been designed to measure the sensitivity of a VQA model to the language bias, and therefore is a test for measuring the ability of a model to generalize to unseen situations.

The evaluation of our contribution on VQAv2-CP is a particularly interesting setup, as obtaining good results there requires an agent to reason beyond exploiting biases. This is important in the context of our contribution. While classical auxiliary losses add additional difficulty to a learning task, as for instance self-supervised contrastive losses [27], the proposed semantic loss is a different case as it is inherently still based on classification of the answer classes, albeit on a restructured output space. Testing this auxiliary loss in a setting, where the training distribution equals the test distribution would make the evaluation unfavorable by introducing a difference between the minimized objective and the evaluation metric. We claim that the new semantic loss increases reasoning and decreases the dependence

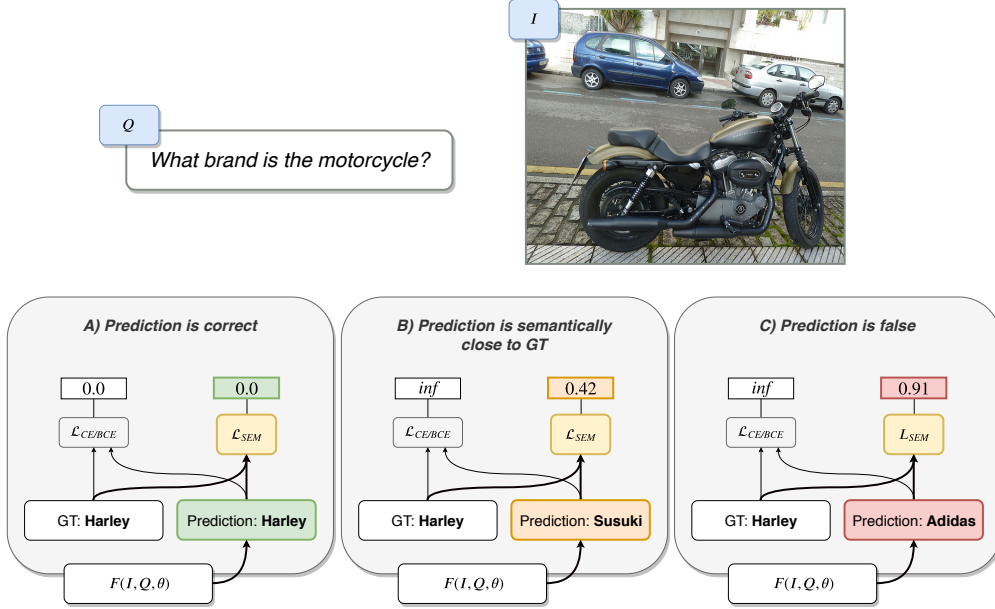


Figure 3: Schematic illustration of the proposed loss \mathcal{L}_{SEM} , which penalizes the semantic distance between the prediction and the ground truth answer (GT). When the prediction is false but semantically related to the ground truth – ‘Suzuki’ vs. ‘Harley’ in case B – we obtain a lower \mathcal{L}_{SEM} value than when the prediction is unrelated to the ground truth – ‘Adidas’ vs. ‘Harley’ in case C. At the same time, the traditional cross-entropy loss $\mathcal{L}_{CE/BCE}$ does not make the difference between case B and C. Numerical values are obtained using the semantic loss with Co-oc semantic space.

on biases, which is better evaluated on datasets with shifts in distributions such as VQAv2-CP. We nevertheless also include comparisons on VQAv2.

To demonstrate that the semantic loss is model-agnostic, we test it on three different standard VQA architectures:

Bottom-Up-Top-Down (UpDn) [7], is a strong baseline architecture for VQA. It introduced the use of a bottom-up (from pixels to visual objects) visual attention to the standard top-down mechanism. In particular, UpDn uses an object detector — R-CNN [28] — to extract bounding boxes along with dense visual features for each object in the image. Thereby, the question attention is computed over a set of objects rather than over standard grid features as has been done before in the literature.

Bilinear Attention Network (BAN) [8] adds a bilinear attention operator on top of the bottom-up top-down mechanism introduced in [7]. Moreover, this model allows multi-hop reasoning by stacking multiple bilinear attention layers with residual connections.

Deep Modular Co-Attention Networks (MCAN) [10] is a Transformer-based [9] multi-modal architecture aiming at modeling both the interactions inside one modality (between words or visual objects) and the interactions between the two modalities (between words and objects) using self-attention mechanisms. Like BAN [10], this architecture contains several stacked self-attention blocks in order to perform complex multi-hop reasoning.

We complement our evaluation by showing the complementarity of our method with RUBi [5], a SOTA approach focusing on the reduction of language bias. It consists of a training procedure adding a question-only branch with a masking mechanism to the base VQA model during the training. The RUBi module adapts the prediction of the base model in order to prevent it from fully exploiting a language-only bias. At test time, the question-only branch is removed.

Training details — We train all our models during 40 epochs. We use the 6-layers version of MCAN [10]. At the beginning of the training we linearly increase the learning rate from $1e^{-4}$ to $2e^{-1}$ during 8 epochs, followed by a decay by a factor of 0.2 at epochs 10 and 20. We set the batch size to 64. For UpDn [7] and BAN [8] we set the batch size to 512 and increase the learning rate from $2e^{-3}$ to $2e^{-1}$ during the first 8 epochs, followed by a decay by a factor of 0.2 at

Table 1: Consistency of the performance gains over multiple neural model architectures: performance using MCAN [10], BAN [8] and UpDn [7] on VQAv2-CP. Baselines marked with † have been trained by ourselves.

Model	Semantic loss	Embedding	Test Acc.
MCAN [10] †	✗		42.5
MCAN	✓	Glove	44.5
MCAN	✓	Co-oc	45.4
BAN [8] †	✗		40.6
BAN	✓	Glove	41.5
BAN	✓	Co-oc	41.4
UpDn [7] †	✗		40.4
UpDn	✓	Glove	40.9
UpDn	✓	Co-oc	40.5

epochs 10 and 20. We use the 4-layers implementation of BAN. We use binary cross entropy along with the Adam optimizer [29] for MCAN [10] and Adamax [29] for BAN [8] and UpDn [7]².

All of our experiments are run on two NVIDIA P100 GPUs with half precision training using the apex library.³ Note that all of our models are trained on the training split only, without the help of any external dataset such as in [8] and [10]. We set the two hyper-parameters $\lambda = 10.0$ and $k = 10$ using grid search.

4.1 Results

Model-agnosticity — Table 1 shows the effectiveness of our semantic loss on VQAv2-CP with three models. The proposed approach improves accuracy by +2.0, +0.9 and +0.5 points on, respectively, MCAN, BAN and UpDn models when using the Glove semantic space. The gain is in general even higher when using Co-oc semantic space, leading to improvements of +2.9 (MCAN), 0.8 (BAN) and +0.1 (UpDn). The on-par performance of Co-oc with respect to Glove (or arguably superior), despite the estimation of Glove from large-scale datasets, illustrates the strength of the goal-directed estimation strategy of Co-oc and the thus learning signal directly derived from human annotations.

We observe that the impact of the semantic loss on the UpDn architecture is less significant than on BAN and MCAN, both for the Co-oc and Glove semantic spaces. We conjecture that this is due to the higher dependency of UpDn on the question bias. To further investigate this, we performed experiments combining the semantic loss with a state-of-the-art bias reduction method.

Complementarity of gains with bias-reduction methods — We combine the semantic loss with RUBi [5], a state-of-the-art method designed to reduce language biases in VQA models. RUBi combines standard VQA models with a second question-only branch, whose objective is the explicit estimation of language biases. During training time, the prediction of the question-only branch is used as a mask combined with the VQA branch by element-wise multiplication, which drives the VQA model to overcome the inherent language bias. The masking is removed during testing.

As shown in Table 2, the semantic loss (in the using Glove variant) improves upon the combination of UpDn model architecture + RUBi training with a margin of +3.3 points and reaches an accuracy of 47.5% on the VQAv2-CP test split⁴. This indicates that the proposed loss is complementary to existing bias-reduction approaches and improves the generalization and reasoning abilities of the model.

Impact on VQAv2 dataset — As discussed before, the VQAv2-CP dataset has been proposed by [15] with the goal to evaluate the performance of VQA models in a condition where they cannot fully rely on question biases. Indeed, as shown in [15], the original VQAv2 [2] dataset contains numerous question biases (*e.g.* the question “what color is the banana in the picture” can be correctly answered as “yellow” without even analyzing the picture in VQAv2). At the same time, it is very important to verify that our semantic loss, which is effective for training VQA models on the *unbiased* VQAv2-CP dataset, does not hinder the model’s performances when the training is done on the *biased* VQAv2 dataset.

²For the baseline models, we use publicly available implementations at <https://github.com/MILVLG/openvqa>

³<https://github.com/NVIDIA/apex>

⁴We observe an instability when using UpDn+RUBi which occasionally prevents the model from converging. As a consequence, we provide the average accuracy over four converged models with random seeds along with the standard deviation.

Table 2: Complementary of gains: the semantic loss can be combined with SOTA methods decreasing language biases like RUBi [5], showing combined gains on VQAv2-CP [4]. The semantic space is Glove [24]. Baselines with † have been trained by us.

Model	Test Acc.
UpDn [7] †	40.4
+ RUBi [5]	44.23
+ RUBi [5] + Semantic Loss (ours)	47.5 \pm 0.3

Table 3: Comparison on VQAv2 [2]. Among the work focusing on reducing language biases, the proposed semantic loss is the method which degrades performance the least on VQAv2. *For a fair comparison, we display the accuracy of the base model used by the authors of the different methods. We only display models which does not rely on additional annotation such as [16] and [17]. Baselines with † have been trained by us.

Model	VQAv2-val	Δ
UpDn*	62.9	
UpDn + DLR [19]	58.0	-4.9
Base*	63.1	
Base + RUBi [5]	61.2	-2.6
UpDn*	63.5	
UpDn + SCR (QA) [17]	62.3	-1.2
UpDn*	63.5	
UpDn + Q-Adv + DoE [15]	62.8	-0.7
MCAN †	66.1	
MCAN + Semantic Loss (ours)	66.0	-0.1

Table 4: Comparison of our method combined with RUBi [5] against state-of-the-art on VQAv2-CP. We divide the table in two groups: methods based on the UpDn [7] architecture – on the bottom part – and the others – on the top part. Column ‘Supp. ann.’ specifies models trained with additional annotations. Models with † have been trained by us. Base* correspond to the baseline model used in [5]

Model	Test Acc.	Supp. ann.
Question-Only [4]	15.95	
BAN [8] †	40.6	
Q-type Balanced Sampling [5]	42.1	
MCAN [10] †	42.5	
NSM [30]	45.8	
Base* + RUBi [5]	47.1	
UpDn [7] †	40.4	
UpDn + Q-Adv + DoE [15]	41.2	
UpDn + RUBi [5]	44.2	
UpDn + HINT [16]	46.7	✓
UpDn + RUBi + Semantic Loss (ours)	47.5	
UpDn + SCR (QA) [17]	48.5	
UpDn + DLR [19]	48.9	
UpDn + SCR (VQA-X) [17]	49.5	✓

Therefore, Table 3 analyzes the impact of our semantic loss on the VQAv2 dataset and compares it with recent approaches designed to remove the language bias in VQA. More precisely, we compare the accuracies on the VQAv2

validation dataset of baseline VQA models (the original baselines from the respective works are taken) with and without one of the SOTA approaches aiming to reduce the question biases. For a fair comparison, we only compare with methods which does not rely on extra annotated supervision such as HINT [16] and SCR(VQA-X) [17]. The impacts of the compared approaches on the respective baselines are highlighted in the “ Δ ” column of Table 3.

When combining the semantic loss with MCAN [10] in the Co-oc variant, we observe a marginal drop of -0.1 in accuracy. On the contrary, SCR (QA) [17] and RUBi [5] cause significant drops of performance of respectively -1.2 and -2.6 points. The drop of the recent DLR [19] method is even more impressive reaching almost -5 accuracy points. All in all, contrary to other SOTA methods, our semantic loss allows to reduce the dependency on question biases (*cf.* results presented in Table 1) without sacrificing accuracy on the biased VQAv2 dataset.

Comparison with the state-of-the-art — We compare our method with SOTA approaches on the VQAv2-CP dataset in Table 4. For a fair comparison, we divide Table 4 into two groups: methods based on the UpDn [7] architecture and the others. When combining the proposed loss along with another bias reduction method – namely RUBi [5] – using the UpDn architecture, we achieve a SOTA-level accuracy of 47.5%. Note that, contrary to HINT [16] or SCR(VQA-X) [17], our approach does not require any additional annotations. DLR [19] achieves a high accuracy (48.5%) on VQAv2-CP. However, unlike our approach, DLR causes a significant drop in accuracy (of -4.9 points) on the biased VQAv2 dataset as highlighted in Table 3.

5 Conclusions

VQA as almost always been treated as a classification task. However, despite its convenience, this strategy does not take into account the semantic relationships between the answers. We showed that suitably structuring the semantic space of output classes can overcome some of the shortcomings of the classification strategy widely used in VQA. We propose a new loss based on proximity in a semantic space and we suggest two different ways to estimate semantic proximity. One is based on word embeddings, the second one directly taps into human assessments and exploits the ambiguities of question formulations and their interpretation. We show that while this proximity space is estimated from less statistical coverage, it is not less effective.

We experimentally demonstrate the effectiveness of the semantic loss in reducing dependency over language biases on VQAv2-CP [4] and its consistency over several standard VQA architectures. Moreover, we show that contrary to other SOTA methods, this gain in accuracy is not made at a cost of a degradation of the performance on the classic VQAv2 [2] dataset. Finally, when combined with another bias reduction method, our semantic loss allows to achieve an accuracy on par with SOTA on VQAv2-CP [4].

In future work, we aim to continue to pave the way to the long-term objective of the community in moving to a direct structured prediction of the textual output, which will bring VQA closer to more traditional models in NLP. Even in this case of direct prediction of text sequences, however, it is far from clear whether classification as an auxiliary loss could not eventually provide an additional useful and complementary learning signal.

References

- [1] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.
- [2] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6904–6913, 2017.
- [3] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6700–6709, 2019.
- [4] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. Don’t just assume; look and answer: Overcoming priors for visual question answering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [5] Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. Rubi: Reducing unimodal biases for visual question answering. In *Advances in Neural Information Processing Systems*, pages 839–850, 2019.
- [6] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 21–29, 2016.
- [7] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018.
- [8] Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. Bilinear attention networks. In *Advances in Neural Information Processing Systems*, pages 1564–1574, 2018.
- [9] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [10] Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6281–6290, 2019.
- [11] Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. Analyzing the behavior of visual question answering models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1955–1960, 2016.
- [12] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2901–2910, 2017.
- [13] Abhishek Das, Harsh Agrawal, C. Lawrence Zitnick, Devi Parikh, and Dhruv Batra. Human Attention in Visual Question Answering: Do Humans and Deep Networks Look at the Same Regions? In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
- [14] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *European Conference on Computer Vision*, pages 793–811. Springer, 2018.
- [15] Sainandan Ramakrishnan, Aishwarya Agrawal, and Stefan Lee. Overcoming language priors in visual question answering with adversarial regularization. In *Advances in Neural Information Processing Systems*, pages 1541–1551, 2018.
- [16] Ramprasaath R Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Shalini Ghosh, Larry Heck, Dhruv Batra, and Devi Parikh. Taking a hint: Leveraging explanations to make vision and language models more grounded. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2591–2600, 2019.
- [17] Jialin Wu and Raymond Mooney. Self-critical reasoning for robust visual question answering. In *Advances in Neural Information Processing Systems*, pages 8601–8611, 2019.
- [18] Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Anna Rohrbach, Bernt Schiele, Trevor Darrell, and Marcus Rohrbach. Multimodal explanations: Justifying decisions and pointing to the evidence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8779–8788, 2018.

- [19] Chencheng Jing, Yuwei Wu, Xiaoxun Zhang, Yunde Jia, and Qi Wu. Overcoming language priors in vqa via decomposed linguistic representations.
- [20] Xin Geng, Chao Yin, and Zhi-Hua Zhou. Facial age estimation by learning from label distributions. In *IEEE transactions on pattern analysis and machine intelligence*, pages 2401–2412, 2013.
- [21] Xin Geng. Label distribution learning. *IEEE Transactions on Knowledge and Data Engineering*, 28(7):1734–1748, 2016.
- [22] Mateusz Malinowski and Mario Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *Advances in neural information processing systems*, pages 1682–1690, 2014.
- [23] Zhibiao Wu and Martha Palmer. Verbs semantics and lexical selection. In *Proceedings of the 32nd annual meeting on Association for Computational Linguistics*, pages 133–138. Association for Computational Linguistics, 1994.
- [24] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- [25] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
- [26] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. pages 2579–2605, 2008.
- [27] Yazhe Li Aaron van den Oord and Oriol Vinyals. Representation learning with contrastive predictive coding. In *Advances in Neural Information Processing Systems*, 2018.
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 2014.
- [30] Drew Hudson and Christopher D Manning. Learning by abstraction: The neural state machine. In *Advances in Neural Information Processing Systems*, pages 5901–5914, 2019.