

Cross-Modal Generative Augmentation for Visual Question Answering

Zixu Wang¹

zixu.wang@imperial.ac.uk

Yishu Miao¹

y.miao20@imperial.ac.uk

Lucia Specia^{1,2}

l.specia@imperial.ac.uk

¹ Department of Computing

Imperial College London

London, UK

² Department of Computer Science

University of Sheffield

Sheffield, UK

Abstract

Data augmentation has been shown to effectively improve the performance of multimodal machine learning models. This paper introduces a generative model for data augmentation by leveraging the correlations among multiple modalities. Different from conventional data augmentation approaches that apply low-level operations with deterministic heuristics, our method learns a generator that generates samples of the target modality conditioned on observed modalities in the variational auto-encoder framework. Additionally, the proposed model is able to quantify the confidence of augmented data by its generative probability, and can be jointly optimised with a downstream task. Experiments on Visual Question Answering as downstream task demonstrate the effectiveness of the proposed generative model, which is able to improve strong UpDn-based models to achieve state-of-the-art performance.

1 Introduction

Multimodal machine learning is a multidisciplinary field combining language, vision, and speech processing to address a multitude of tasks [5, 6, 24, 36, 49]. However, a major bottleneck in multimodal learning is the need for multi-way parallel data, i.e. data with all modalities for all samples. For example, Visual Question Answering (VQA) [3, 9, 13, 30, 42, 45, 57, 59] requires learning from parallel data across three different sources – image, question and answer, a costly resource if created at large scale. In this paper we propose a generative model that leverages the joint distribution of multiple modalities from existing datasets to carry out data augmentation for VQA.

There are two major challenges for data augmentation. First, the augmented data should contain meaningful **variations** and less repetition, and second, the **reliability** of the augmented data should be guaranteed and effectively evaluated. The repetitive, insignificant variants or unreliable augmented data could have a negative effect on downstream tasks. Conventional data augmentation approaches apply low-level operations, e.g. adding noise or using replacements via heuristics. For example, Chen et al. [8] and Gokhale et al. [14] apply visual and semantic transformations which create additional counterfactual samples to




Figure 1: Illustration of the proposed generative process for data augmentation. The left one describes the conventional VQA pipeline consisting of an image and annotated QA pairs. The right one is our approach to generating QA pairs for data augmentation.

improve the model’s sensitivity to trivial noise in VQA. This provides enough variations but requires specific design and extensive preprocessing to fit the task domain.

Our proposed model explores the generative abilities of conditional distributions for augmenting question-answer (QA) pairs given only images. Figure 1 shows the difference between a conventional VQA model and our proposed generative framework. The annotated QA pairs are directly used in supervised learning to predict answers in a VQA model, while our model aims at generating QA pairs from unlabelled images. We first learn the conditional distributions of different modalities using annotated images and QA pairs. Then, we introduce Q and A as two discrete latent variables, and construct a generative model for unlabelled images to generate QA pairs as augmented data. Finally, the augmented data with reliability scores are used for training to improve on strong VQA base models. Without any specific low-level operations on the data, we utilise the dynamics of the generative distributions and the compositionality of multiple modalities to explore the novel QA pairs given an image. Additionally, the generators are optimised by the REINFORCE algorithm [54] while minimising the variational lower bound.

The generative approach has the following advantages. First, the model has strong generalisation ability to incorporate additional unseen and unlabelled images. Diverse QA pairs can be generated by exploring the dynamics of the generator and the compositionality of multiple modalities. Second, the answers are sampled from images before generating the questions, hence it is less prone to exploit linguistic priors in questions and to generate trivial QA pairs that are irrelevant to the given images. Third, the augmented data can be quantified by the generative distribution, which acts as reliability scores of QA pairs for downstream VQA training. Our approach opens a promising new direction for multimodal learning, with strong potential for cross-modal understanding and generalisation.

Using VQA as downstream task, our approach outperforms base models (UpDn and LXMERT), and substantially improves strong UpDn models, leading to the state-of-the-art performance on the task.

2 Related Work

Data Augmentation in VQA There has been extensive work on improving VQA robustness with data augmentation [12, 19, 59, 46, 47, 60]. The pioneer work is introduced by Kafle et al. [19] where they generate new questions by using semantic annotations on images. Shah et al. [42] proposes a cycle consistency training scheme where it generates questions and trains the model with question and answer consistency. Augmenting VQA training data with counterfactual samples has also been proven effective [11, 20] with complementary samples by masking critical objects in images or words in questions, and assigning different

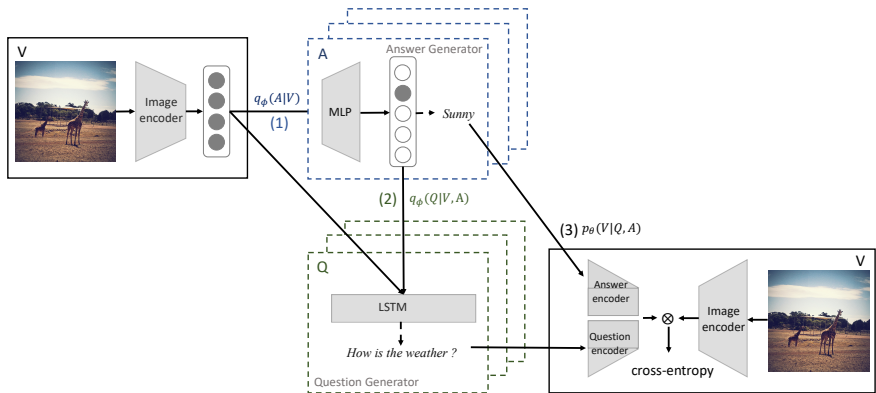


Figure 2: Architecture of the proposed generative model for data augmentation in the VQA classification task. For the training of the generator $q_\phi(Q,A|V)$, (1) given an image, we first build an image classification module $q_\phi(A|V)$ to predict possible answers; (2) with a sampled answer (e.g. “sunny”) from answer distribution, we build a conditional question generator $q_\phi(Q|V,A)$ using an LSTM to generate questions (e.g. “how is the weather ?”) corresponding to the image and the sampled answer; then (3) we apply a neural network $p_\theta(V|Q,A)$ to measure the score of relevance between generated QA pairs and the image.

ground-truth answers (e.g., CSS [8] and MUTANT [14]). However, most data augmentation work focuses on explicitly injecting semantic and visual perturbations on raw inputs or using adversarial samples to provide contrastive learning signals. Our work augments the original dataset with implicitly generated QA pairs especially for unlabelled images. This helps the model learn correlations between modalities.

Visual Question Generation Generating questions from images has received equal interests as VQA [18, 27, 28, 65, 60, 66]. Several recent works have explored the task of visual question generation with variational auto-encoders and maximising mutual information with answer categories [18, 26]. Our work draws inspiration from VQG by incorporating a generation module in the generative framework.

Our generative model incorporates a cross-modal retrieval module [15, 63] to measure the correlations between generated QA pairs and the image. Cross-modal retrieval performance relies on appropriate representations for multi-modal data. Most of the existing studies on cross-modal retrieval mainly focus on learning a high-level common space and exploiting visual-semantic embedding to calculate the similarities between image and sentence features with ranking loss [16, 24, 23, 61, 62]. We build a classifier to output the reliability scores of generated QA pairs and images.

3 Model

Variational auto-encoders (VAE) are generally applied for unsupervised representation learning of language and images [22, 23, 61, 62, 63, 41]. Inspired by the idea of VAE for data augmentation [16, 17, 63], we propose a generative framework in order to explore cross-modal interactions in the multimodal data for the task of VQA.

Here, V , Q , A are used to denote the input image, question, and the answer respectively. Figure 2 shows the structure of variational auto-encoder for VQA. Q , A are the two latent variables introduced to represent questions and answers, respectively. The training objective consists of the variable lower bound and the cross-entropy (CE) loss for VQA classification:

$$E_{q_\phi(Q,A|V)}[\log p_\theta(V, Q, A) - \log q_\phi(Q, A|V)] + p_\psi(A|V, Q) \quad (1)$$

where $q_\phi(Q, A|V)$ is the generator that generates corresponding QA pairs for given images and can be factorised into answer generator $q_\phi(A|V)$ and question generator $q_\phi(Q|V, A)$; $p_\theta(V, Q, A)$ is the cross-modal distribution that regularises the QA samples, and $p_\psi(A|V, Q)$ is the multi-label classification loss for VQA. Note that the lower bound will act as a confidence measure of generate QA pairs and will be used to reweigh the cross-entropy loss while training the downstream VQA model.

The overall architecture of our generative model is illustrated in Figure 2. The training steps are summarised as follows: the joint distribution of three modalities (V , Q , and A) is modelled to learn prior knowledge on cross-modal interactions; optimise the generator to generate QA pairs by first predicting answers from the image and then generating question; sample QA pairs for additional unlabelled images to assist large-scale training for VQA.

3.1 Generative Learning

The first term in the training objective is the variational lower bound that is optimised for learning cross-modal distributions. In order to construct this lower bound, we need to model both the generative distribution $q_\phi(Q, A|V)$ and the cross-modal distribution $p_\theta(V, Q, A)$.

3.1.1 Generative Distribution

The training of the generator $q_\phi(Q, A|V)$ can be decomposed into:

$$q_\phi(Q, A|V) = q_\phi(Q|V, A)q_\phi(A|V) \quad (2)$$

Specifically, given a triplet (V, Q, A) , we build an answer generator $q_\phi(A|V)$ and a conditional question generator $q_\phi(Q|V, A)$. The original image V and the generated Q are then combined to obtain an answer prediction A using the conventional VQA model $p_\psi(A|V, Q)$.

The design of our generative components is based on two hypotheses. First, a model that can predict possible answer candidates directly from a given image has better understanding of the visual content and the dependency between the image and the answers. Second, assuming the predicted answers have high correlation with the image, the conditional language model can generate valid questions considering both the image and the predicted answers.

Answer Generator - $q_\phi(A|V)$ As illustrated in Figure 2 step (1), first we need to build an image classification model to predict possible answers to a given image. Following typical setups in VQA models, we encode the input images as regional visual feature representations. Specifically, these features are extracted from a bottom-up approach [9], where the input image is passed through a ResNet CNN within the Faster R-CNN framework to obtain a vector representation. We take pretrained visual features as a preprocessing step for efficiency, and keep them fixed during the training of the VQA model. We follow the same setup for both labelled and unlabelled images in our experiments.

In our generative framework, the predicted answer distributions given images can be modelled as $q_\phi(A|V)$. We build the classifier by using a non-linear fully connected layer and batch normalisation on the pretrained image features, followed by a projection to the space of all possible answers and a softmax layer. This can be seen as a multi-label image classification task where the labels are possible VQA answers corresponding to the given image instead of image classification categories. Without seeing the questions with strong language priors, the model can directly learn the mapping from the image to answer candidates.

Question Generator - $q_\phi(Q|V,A)$ The second stage of our generative process is to generate questions conditioned on the given image and sampled answers, as shown in Figure 2 step (2). The goal of our question generator is to define a conditional language model $q_\phi(Q|V,A)$ to learn the transformations from one modality (*i.e.* image) to another (*i.e.* question) with the help from possible answers. This is similar to a visual question generation (VQG) task that is aimed at generating not only relevant but diverse questions to each image. However, typical VQG models only take images as input, and thus are not goal-driven and do not guarantee that the generated question corresponds to a specific type of answer. We follow the common setup in [26, 42] to encode the answer along with the image before generating the question. Such an approach allows the model to condition its question on the answer.

We model the question generator by an image encoder, answer encoder, and an LSTM decoder. With predicted answer distributions, the model can generate various plausible questions for the sampled answers. The image encoder transforms the attended image features to lower dimensional feature vectors, and the answer encoder takes the distribution over the answer space as input and outputs the vector representation of the answer. The image and answer representation are then fused together and passed through the LSTM decoder to generate a question, a process that is optimised by minimising the negative log likelihood with teacher-forcing.

3.1.2 Cross-Modal Distribution

We need to model the cross-modal distributions on the existing VQA dataset, where all the modalities are observed. The cross-modal distribution can be decomposed into:

$$\begin{aligned} p_\theta(V,Q,A) &= p_\theta(V|Q,A)p_\theta(Q,A) \\ &= p_\theta(V|Q,A)p_\theta(Q|A)p_\theta(A) \end{aligned} \quad (3)$$

We first learn to model the joint distribution of QA pairs. We build the model to learn the prior knowledge of QA pairs $p_\theta(Q,A)$ through $p_\theta(Q|A)p_\theta(A)$. Firstly, the prior distribution of answer candidates $p_\theta(A)$ is obtained directly from the dataset. With a sampled answer from the prior distribution, we build a question generator conditioned on the answer. Note that instead of obtaining a high quality generative distribution for QA pairs, we learn a joint distribution of the two modalities as the prior knowledge of QA pairs.

The QA pairs are expected to have high dependency and correlation with the given image. However, not all the QA pairs are coherent and consistent with the image, which leads to negative noises during training. Inspired by Shah et al. [42], we overcome this issue by modelling the conditional distribution of images given QA pairs. As shown in Figure 2 part (3), we build a neural network to model $p_\theta(V|Q,A)$ and score the correlation/relevance between each QA pair and the image. We add a layer to fuse question and answer representations together, then multiply the fused QA vector with the image representation through another

fusion layer. Having the representations of both QA pairs and images, we feed them into a binary classifier to output a relevance score of the image and QA pairs after a sigmoid layer.

3.2 VQA Classification

Following the conventional setup in VQA, the image feature V and question representation Q are extracted from the image encoder and the question encoder. The VQA task is constructed as a classification problem to output the most likely answer A from a fixed set of answers based on the content of the image V and question Q . Following Teney et al. [18] and Zhu et al. [6], instead of softmax we use sigmoid outputs in our VQA training to cast it as a multi-label training objective:

$$p_{\psi}(A|V, Q) = \text{CE}(f_{\text{VQA}}(V, Q), A) \quad (4)$$

where CE represents the cross-entropy loss in VQA models.

$$p_{\psi}(A|V, Q) = \text{CE}(f_{\text{VQA}}(V, Q), A) * \mathcal{R} \quad (5)$$

One of the assumptions of our proposed generative training scheme is that the generated QA pairs are always semantically and syntactically correct and have high correlation with the image. However, this is not always the case. In order to overcome this issue, we propose a reweighing mechanism based on a reliability score obtained from our generative objective, as in Equation 5. We define the reliability and confidence of each generated QA pair as \mathcal{R} using the generative objective in Equation 1 and use it to reweigh the cross-entropy loss during VQA training. We add this reliability score to make the model aware of the confidence and quality of the augmented samples.

4 Experimental Setup

4.1 Datasets

We train and evaluate our models on the VQA-CP-v2 [4] and VQA-v2 [9] datasets. These are the most commonly used VQA benchmarks, where VQA-CP-v2 is created from VQA-v2 to evaluate robustness and generalizability of VQA models, by reorganising the train and validation splits. We also report the results of our model on the validation split of VQA-v2, which contains a strong language prior. VQA-CP-v2 contains 121K images, 438K questions, and 4.4M answers for training and 98K images, 220K questions and 2.2M answers for testing. We evaluate our model on the validation split of VQA-v2, which contains 83K images, 444K questions, and 4.4M answers for training and 41K images, 210K questions and 2.1M answers for validation.

The unlabelled images come from the Visual Genome (VG) [25] and MS COCO 2017-unlabelled [24] datasets. We directly use the images from these datasets and generate QA pairs relating to the images using proposed generator. The number of additional images is 101K from VG and 120K from MS COCO 2017-unlabelled data, which are combined to provide 221K unlabelled images in total. Note that we do not filter out VG images that overlap with original VQA dataset as our generator can generate new QA pairs relating to these images.

4.2 Benchmark Models

Most prior works on VQA, such as the ones we compared our model against – GVQA [2], RUBI [7], SCR [55], LMH [11], CSS [8] – are all built based on UpDn [9]. Therefore, we also build our generative model using UpDn [9] as its backbone and investigate the efficacy of the extension under the generative paradigm. UpDn incorporates bottom-up attention in VQA by extracting features associated with image regions proposed by Faster RCNN [40] trained on Visual Genome [25]. We use the evaluation code from official VQA challenge [9].

4.3 Implementation Details

Following previous work, we take the pretrained image features extracted from a ResNet CNN within a Faster R-CNN framework. Each image is transformed into a $K \times 2048$ dimensional vector representation where $K = 36$ in our setup represents a set of objects in the image. Each question is trimmed into a sequence with a maximum length of 14 words and initialised using the 300-dimensional word embeddings from GloVe [58]. A sentence-level representation is obtained by feeding the word embeddings into a GRU [11] with 1280 dimensions. While using sampled answers to guide question generation, we add an answer encoder to transform the samples answer to a representation of dimension 512.

We pre-train the model with the proposed generative objective for 10 epochs (~ 4 hours) and fix it for generating QA pairs to assist downstream VQA training. This number of epochs is chosen because our generative model converges and the quality of generated questions does not improve after these epochs of training. The number of parameters is 90M and the model is trained on 2 RTX 2080Ti. We use batch size of 256 and adapt the Adam optimizer [21] with the initial learning rate of 0.001.

5 Results

In this section we present our experimental results including quantitative and qualitative analysis. We first compare the results with previous approaches including strong baseline models and SOTA benchmarks; then we validate the efficacy of each component proposed in our framework in ablation studies; finally we provide a qualitative analysis of the generated QA pairs.

5.1 Comparison with Alternative Approaches

In Table 1, we compare our proposed framework with previous SOTA approaches on two benchmarks: VQA-CP-v2 and VQA-v2. We compare our generative VQA model against existing models. RUBi[7], SCR[55], LMH[11], and CSS[8] are built on UpDn[9] by adding different de-biasing components to mitigate superficial language biases [37] and improve robustness of VQA models. Besides, CSS[8] and MUTANT[12] make use of data augmentation to provide large-scale data sizes for VQA training.

We show that our generative VQA model outperforms most of the alternative approaches above. For VQA-CP-v2, our method achieves 60.70 accuracy on all question types, competitive with the top results from MUTANT. Our method shows improvements with 3.20 on the Yes-No category, 15.52 on Number-based questions; also 0.83 compared to MUTANT

Model	VQA-CP-v2 test (%) \uparrow				VQA-v2 test (%) \uparrow			
	All	Yes/No	Num	Other	All	Yes/No	Num	Other
GVQA [10]	31.30	57.99	13.68	22.14	48.24	72.03	31.17	34.65
RUBi [11]	47.11	68.65	20.28	43.18	63.10	-	-	-
SCR [12]	48.47	70.41	10.42	47.29	62.30	77.40	40.90	56.50
LMH [13]	52.45	69.81	44.46	45.54	61.64	77.85	40.03	55.04
CSS [8]	58.95	84.37	49.42	48.21	59.91	73.25	39.77	55.11
LXMERT [14]	46.23	42.84	18.91	55.51	74.16	89.31	56.85	65.14
MUTANT [14]	61.72	88.90	49.68	50.78	62.56	82.07	42.52	53.28
Ours	46.93	43.25	20.03	54.56	59.32	80.10	40.23	49.58
UpDn [9]	41.58	43.07	13.58	48.48	63.48	81.18	42.14	55.66
UpDn + SSL [6]	57.59	86.53	29.87	50.03	63.73	-	-	-
Ours	<u>60.70</u>	89.73	<u>45.89</u>	48.36	<u>64.09</u>	<u>81.78</u>	<u>44.57</u>	<u>55.84</u>

Table 1: Accuracies on VQA-CP-v2 test and VQA-v2 validation sets. ‘‘Ours’’ represents the final model build on LXMERT/UpDn with augmentation sampler, cross-modal joint distribution and reweighing augmented loss. Overall best scores are **bold**, and our best ones are underlined.

on the Yes-No category. Our model improves substantially on ‘‘Num’’ questions. When built on UpDn and evaluated on the VQA-v2 val dataset, our framework achieves better performance (%64.09) than UpDn and outperforms MUTANT (%62.56). This indicates that our generative augmentation model has higher generalisation ability to different VQA datasets, compared to MUTANT.

To further prove the effectiveness of our generative model, we conduct experiments based on pre-trained visual-language framework - LXMERT [14], which is also the backbone of the state-of-the-art model on VQA-CP-v2 - MUTANT [14] using data augmentation. Specifically, our generated QA pairs are used to fine-tune the pre-trained LXMERT model. Our model reaches some improvement over LXMERT on VQA-CP-v2, which further shows the efficacy of our generated QA pairs. Although MUTANT is also focused on data augmentation and achieves significant improvements over LXMERT, their augmented data is from manually introduced mutations such as removing object instances and colour inversion for a strong contrastive learning signal. The augmentation relies on low-level operations with deterministic heuristics and applies costly extensive transformations on original samples independently of VQA training. We believe this is intentionally designed for VQA-CP-v2 dataset and limits the generalisation ability, while our goal is to introduce a generative approach that can be generalised to different multimodal tasks for cross-modal data augmentation.

5.2 Evaluation of Model Components

In order to evaluate the efficacy of each component of our generative model, we conduct an ablation study. Table 2 shows the results against two strong base models: UpDn [9] and SSL [6]. SSL is built on UpDn and uses negative sampling to replace images in the triplets of (image, question, answer) to create irrelevant samples. We can observe that in both cases, the components in our proposed framework can be beneficial and improve the performance by a reasonable margin.

Specifically, the incorporation of generative training for augmentation, i.e. answer gen-

	VQA-CP-v2 test (%)			
	All	Yes/No	Num	Other
UpDn	41.58	43.07	13.58	48.48
+ $q_\phi(Q, A V)$	42.84	44.57	24.93	47.05
+ $p_\theta(V, Q, A)$	42.96	45.88	24.82	47.66
+ reweighted	43.12	45.76	25.71	48.51
SSL	57.59	86.53	29.87	50.03
+ $q_\phi(Q, A V)$	59.80	89.54	43.80	48.61
+ $p_\theta(V, Q, A)$	60.37	89.49	44.43	48.72
+ reweighted	60.70	89.73	45.89	48.36

Table 2: Study on the benefits of each component of the proposed approach: augmentation sampler, cross-modal joint distribution, and reweighted loss. The best accuracies are in **bold**.

eration $q(A|V)$ and question generation $q(Q|V, A)$, improves 1.26 and 2.21 accuracy points on the two base models, respectively. When introducing the cross-modal joint modelling to provide prior knowledge of multimodal distribution, the performance is further improved by 0.12 and 0.57 accuracy points, respectively. The overall better results observed with SSL can be attributed to the negative sampling feature in this model. The most significant improvement across all the categories is observed in “Num” questions. With the generative objective as reward to reweigh the VQA loss on the augmented data, the best performance of our model can achieve is 60.70 accuracy, outperforming the strong SSL model by 3.11 points.

5.3 Examples of Generated QA pairs

To qualitatively demonstrate the effectiveness of the proposed approach, Figure 3 shows examples of generated QA pairs and compares them with the original ones from the VQA-CP-v2 dataset.

In the four examples, we present five generated QA pairs from the proposed generator and five QA pairs from the original dataset, for each image. For the generated QA pairs, we have a column $q_\phi(A|V)$ showing the probabilities of sampled answers given only the input image, in descending order. It can be seen from the Answer column that most of the answers can be directly related to the image, but some are not. This is because we use ground-truth answers from the original dataset to train the image-answer pipeline, which contains some answers that have high correlations with the questions but not with the image.

In the column of Questions, we show the generated questions that are conditioned on the sampled answers and the image, which have high diversity. To quantify diversity, we measure type-token ratio on a per-image basis. The average ratio is 0.53, which is – as expected – lower than the average ratio for the original QA pairs (0.79), but Figure 3 shows that generated questions are meaningful and related to the answers and the image.

Additionally, the Reliability \mathcal{R} of each QA pair is illustrated, showing that most QA pairs have high reliability and relevance to the image. Note that high probabilities of answers given an image cannot guarantee that the augmented QA pairs have high correlation with the image. In the column of Novel, we use green check to indicate the novel generated questions when compared to the set of question in the column of Original QA pairs. The questions are novel when they convey different information from the original ones.





Images	Augmented QA pairs					Original QA pairs	
	$q_p(A V)$	Answer	Question	Reliability \mathcal{R}	Novel	Answer	Question
	0.32	many	how many vegetables are shown ?	0.87	✓	white	what color is the plate ?
	0.23	vegetable	what is in the picture ?	0.79	✗	no	are there utensils on the table ?
	0.10	yes	does this look healthy ?	0.82	✗	vegetable	what is on the plate ?
	0.03	maybe	what kind of protein is there ?	0.41	✓	eggplant	what is the purple vegetable called ?
	0.03	left	where is the banana ?	0.20	✗	no	can you see any bananas ?
	0.16	yes	is the man wearing a short ?	0.42	✗	skateboard	what is the man riding on ?
	0.12	no	is the skateboard upside down ?	0.75	✓	no	is the man looking at the camera ?
	0.08	jump	what is the woman doing ?	0.57	✗	yes	is the man wearing a hat ?
	0.02	man	what is this ?	0.32	✗	black	what color is the hat ?
	0.01	hat	what is the man wearing ?	0.88	✗	plain	What design is on girl's socks ?
	0.29	sunny	what is the weather condition ?	0.82	✗	1	how many people are on the scooter ?
	0.23	1	how many people are there ?	0.86	✗	spring	what season was this picture taken ?
	0.12	yes	is the scooter in motion ?	0.60	✓	sunny	overcast or sunny ?
	0.005	station	what is the building ?	0.53	✗	suzuki	what kind of motorcycle is that ?
	0.001	outside	can I open this door ?	0.43	✓	cosmo	what brand of oil is on the gas pumps?
	0.23	3	how many people in the picture ?	0.32	✗	ball	what is the girl holding ?
	0.22	1	how many girls in the picture ?	0.83	✗	1	how many frisbee are in the picture ?
	0.08	grass	what is the woman watching ?	0.80	✓	skirt	what is the little girl wearing ?
	0.01	frisbee	what is the girl playing ?	0.73	✗	0	how many hydrants are there ?
	0.001	no	is she running ?	0.63	✗	yes	has the grass been mowed recently ?

Figure 3: **Qualitative Analysis:** Examples comparing generated QA pairs against those from the original dataset annotations for the same image. Generated QA pairs are informative if they are out of distribution compared with the original QA pairs.

We also inspect the generated data for a better understanding of the generated QA pairs for different types of questions. The improvements over “NUM” questions can be due to the answer generator. Our answer generator can sample diverse possible numbers first given an image and the subsequently generated questions are always “how many” questions, which is not always the case for “Yes/No” and “Other” types. This provides diverse “NUM” QA pairs which are not always “1” as the answers in the original dataset.

6 Conclusions

This paper introduces a generative model for cross-modal data augmentation on VQA. We learn a generator to generate reliable QA pairs given images under a generative framework. The augmented QA pairs are trained and evaluated by the generative distribution pretrained on VQA dataset, which are in turn employed in a downstream VQA task with confidence scores to selectively improve the classification performance. Without low-level operations or specific heuristics, our proposed model is able to augment unlabelled images with large-scale reasonable QA pairs, which boosts a vanilla model to achieve benchmark results. The strong generalisation ability of our model open avenues for extension to other multimodal machine learning tasks.

Acknowledgements

This work received support from the MultiMT project (H2020 ERC Starting Grant No. 678017) and the Air Force Office of Scientific Research (under award number FA8655-20-1-7006).

References

- [1] V. Agarwal, Rakshith Shetty, and M. Fritz. Towards causal vqa: Revealing and reducing spurious correlations by invariant and covariant semantic editing. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9687–9695, 2020.
- [2] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. Don’t just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [3] 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 (CVPR)*, June 2018.
- [4] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. VQA: Visual Question Answering. In *International Conference on Computer Vision (ICCV)*, 2015.
- [5] T. Baltrušaitis, C. Ahuja, and L. Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):423–443, 2019. doi: 10.1109/TPAMI.2018.2798607.
- [6] Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. Experience grounds language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8718–8735, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.703. URL <https://www.aclweb.org/anthology/2020.emnlp-main.703>.
- [7] Remi Cadene, Corentin Dancette, Hedi Ben younes, Matthieu Cord, and Devi Parikh. Rubi: Reducing unimodal biases for visual question answering. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, pages 841–852. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/51d92be1c60d1db1d2e5e7a07da55b26-Paper.pdf>.
- [8] Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. Counterfactual samples synthesizing for robust visual question answering. In *CVPR*, 2020.
- [9] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, 2020.
- [10] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.

- [11] Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. Don't take the easy way out: Ensemble based methods for avoiding known dataset biases. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4069–4082, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1418. URL <https://www.aclweb.org/anthology/D19-1418>.
- [12] Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K. Srivastava, Li Deng, Piotr Dollar, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C. Platt, C. Lawrence Zitnick, and Geoffrey Zweig. From captions to visual concepts and back. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [13] Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. Large-scale adversarial training for vision-and-language representation learning. In *NeurIPS*, 2020.
- [14] Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. MUTANT: A training paradigm for out-of-distribution generalization in visual question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 878–892, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.63. URL <https://www.aclweb.org/anthology/2020.emnlp-main.63>.
- [15] Jiuxiang Gu, Jianfei Cai, Shafiq R. Joty, Li Niu, and Gang Wang. Look, imagine and match: Improving textual-visual cross-modal retrieval with generative models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [16] Yutai Hou, Yijia Liu, W. Che, and Ting Liu. Sequence-to-sequence data augmentation for dialogue language understanding. *ArXiv*, abs/1807.01554, 2018.
- [17] Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. Toward controlled generation of text. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1587–1596, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/hu17e.html>.
- [18] Unnat Jain, Ziyu Zhang, and Alexander G. Schwing. Creativity: Generating diverse questions using variational autoencoders. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [19] Kushal Kafle, Mohammed Yousefhusien, and Christopher Kanan. Data augmentation for visual question answering. In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 198–202, Santiago de Compostela, Spain, September 2017. Association for Computational Linguistics. doi: 10.18653/v1/W17-3529. URL <https://www.aclweb.org/anthology/W17-3529>.
- [20] Yash Kant, A. Moudgil, Dhruv Batra, D. Parikh, and Harsh Agrawal. Contrast and classify: Alternate training for robust vqa. *ArXiv*, abs/2010.06087, 2020.

- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *ICLR*, 2014.
- [23] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27, pages 3581–3589. Curran Associates, Inc., 2014. URL <https://proceedings.neurips.cc/paper/2014/file/d523773c6b194f37b938d340d5d02232-Paper.pdf>.
- [24] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014.
- [25] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2017. URL <https://doi.org/10.1007/s11263-016-0981-7>.
- [26] Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Information maximizing visual question generation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [27] Dong Bok Lee, Seanie Lee, Woo Tae Jeong, Donghwan Kim, and Sung Ju Hwang. Generating diverse and consistent QA pairs from contexts with information-maximizing hierarchical conditional VAEs. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 208–224, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.20. URL <https://www.aclweb.org/anthology/2020.acl-main.20>.
- [28] Yikang Li, Nan Duan, Bolei Zhou, Xiao Chu, Wanli Ouyang, Xiaogang Wang, and Ming Zhou. Visual question generation as dual task of visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [29] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft coco: Common objects in context, 2014. URL <http://arxiv.org/abs/1405.0312>. cite arxiv:1405.0312Comment: 1) updated annotation pipeline description and figures; 2) added new section describing datasets splits; 3) updated author list.
- [30] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29, pages 289–297. Curran Associates, Inc., 2016. URL <https://proceedings.neurips.cc/paper/2016/file/9dcb88e0137649590b755372b040afad-Paper.pdf>.

- [31] Yishu Miao and Phil Blunsom. Language as a latent variable: Discrete generative models for sentence compression. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 319–328, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1031. URL <https://www.aclweb.org/anthology/D16-1031>.
- [32] Yishu Miao, Lei Yu, and Phil Blunsom. Neural variational inference for text processing. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1727–1736, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <http://proceedings.mlr.press/v48/miao16.html>.
- [33] Andriy Mnih and Karol Gregor. Neural variational inference and learning in belief networks. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1791–1799, Beijing, China, 22–24 Jun 2014. PMLR. URL <http://proceedings.mlr.press/v32/mnih14.html>.
- [34] Aditya Mogadala, M. Kalimuthu, and D. Klakow. Trends in integration of vision and language research: A survey of tasks, datasets, and methods. *ArXiv*, abs/1907.09358, 2019.
- [35] Nasrin Mostafazadeh, Ishan Misra, Jacob Devlin, Margaret Mitchell, Xiaodong He, and Lucy Vanderwende. Generating natural questions about an image. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1802–1813, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1170. URL <https://www.aclweb.org/anthology/P16-1170>.
- [36] J. Ngiam, A. Khosla, Mingyu Kim, Juhan Nam, H. Lee, and A. Ng. Multimodal deep learning. In *ICML*, 2011.
- [37] Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. Counterfactual vqa: A cause-effect look at language bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [38] Jeffrey Pennington, Richard Socher, and Christopher Manning. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1162. URL <https://www.aclweb.org/anthology/D14-1162>.
- [39] Arijit Ray, Karan Sikka, Ajay Divakaran, Stefan Lee, and Giedrius Burachas. Sunny and dark outside?! improving answer consistency in VQA through entailed question generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5860–5865, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1596. URL <https://www.aclweb.org/anthology/D19-1596>.

- [40] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28, pages 91–99. Curran Associates, Inc., 2015. URL <https://proceedings.neurips.cc/paper/2015/file/14bfa6bb14875e45bba028a21ed38046-Paper.pdf>.
- [41] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1278–1286, Beijing, China, 22–24 Jun 2014. PMLR. URL <http://proceedings.mlr.press/v32/rezende14.html>.
- [42] Meet Shah, Xinlei Chen, Marcus Rohrbach, and Devi Parikh. Cycle-consistency for robust visual question answering. In *2019 Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2019.
- [43] Richard Socher, Andrej Karpathy, Quoc V. Le, Christopher D. Manning, and Andrew Y. Ng. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics*, 2: 207–218, 2014. doi: 10.1162/tacl_a_00177. URL <https://www.aclweb.org/anthology/Q14-1017>.
- [44] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VI-bert: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SygXPaEYvH>.
- [45] Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1514. URL <https://www.aclweb.org/anthology/D19-1514>.
- [46] Ruixue Tang, Chao Ma, Wei Emma Zhang, Qi Wu, and Xiaokang Yang. Semantic equivalent adversarial data augmentation for visual question answering. In *European Conference on Computer Vision (ECCV)*, 2020.
- [47] Ruixue Tang, Chifeng Ma, W. Zhang, Qi Wu, and Xiaokang Yang. Semantic equivalent adversarial data augmentation for visual question answering. In *ECCV*, 2020.
- [48] Damien Teney, Peter Anderson, Xiaodong He, and Anton Van Den Hengel. Tips and tricks for visual question answering: Learnings from the 2017 challenge. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4223–4232, 2018.
- [49] Shagun Uppal, Sarthak Bhagat, Devamanyu Hazarika, Navonil Majumdar, Soujanya Poria, R. Zimmermann, and Amir Zadeh. Emerging trends of multimodal research in vision and language. *ArXiv*, abs/2010.09522, 2020.

- [50] Nihir Vedd, Zixu Wang, Marek Rei, Yishu Miao, and Lucia Specia. Guiding visual question generation. *arXiv preprint arXiv:2110.08226*, 2021.
- [51] Ivan Vendrov, Ryan Kiros, S. Fidler, and R. Urtasun. Order-embeddings of images and language. *CoRR*, abs/1511.06361, 2016.
- [52] Liwei Wang, Yin Li, and Svetlana Lazebnik. Learning deep structure-preserving image-text embeddings. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [53] Zihao Wang, Xihui Liu, Hongsheng Li, Lu Sheng, Junjie Yan, Xiaogang Wang, and Jing Shao. Camp: Cross-modal adaptive message passing for text-image retrieval. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [54] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.*, 8(3–4):229–256, May 1992. ISSN 0885-6125. doi: 10.1007/BF00992696. URL <https://doi.org/10.1007/BF00992696>.
- [55] Jialin Wu and Raymond Mooney. Self-critical reasoning for robust visual question answering. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, pages 8604–8614. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/33b879e7ab79f56af1e88359f9314a10-Paper.pdf>.
- [56] Peixi Xiong and Ying Wu. Ta-student vqa: Multi-agents training by self-questioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [57] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alexander J. Smola. Stacked attention networks for image question answering. In *CVPR*, pages 21–29. IEEE Computer Society, 2016. URL <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2016.html#YangHGDS16>.
- [58] Kang Min Yoo, Youhyun Shin, and Sang-goo Lee. Data augmentation for spoken language understanding via joint variational generation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 7402–7409, 2019.
- [59] Yu Jiang*, Vivek Natarajan*, Xinlei Chen*, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Pythia v0.1: the winning entry to the vqa challenge 2018. *arXiv preprint arXiv:1807.09956*, 2018.
- [60] Xi Zhu, Zhendong Mao, Chunxiao Liu, Peng Zhang, Bin Wang, and Yongdong Zhang. Overcoming language priors with self-supervised learning for visual question answering. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 1083–1089. International Joint Conferences on Artificial Intelligence Organization, 7 2020. doi: 10.24963/ijcai.2020/151. URL <https://doi.org/10.24963/ijcai.2020/151>. Main track.