

Differential Attention for Visual Question Answering

Badri Patro and Vinay P. Namboodiri
IIT Kanpur
{ badri,vinaypn }@iitk.ac.in

Abstract

In this paper we aim to answer questions based on images when provided with a dataset of question-answer pairs for a number of images during training. A number of methods have focused on solving this problem by using image based attention. This is done by focusing on a specific part of the image while answering the question. Humans also do so when solving this problem. However, the regions that the previous systems focus on are not correlated with the regions that humans focus on. The accuracy is limited due to this drawback. In this paper, we propose to solve this problem by using an exemplar based method. We obtain one or more supporting and opposing exemplars to obtain a differential attention region. This differential attention is closer to human attention than other image based attention methods. It also helps in obtaining improved accuracy when answering questions. The method is evaluated on challenging benchmark datasets. We perform better than other image based attention methods and are competitive with other state of the art methods that focus on both image and questions.

1. Introduction

Answering questions regarding images requires us to obtain an understanding about the image. We can gain insights into a method by observing the region of an image the method focuses on while answering a question. It has been observed in a recent work that humans also attend to specific regions of an image while answering questions [3]. We therefore expect strong correlation between focusing on the “right regions while answering questions and obtaining better semantic understanding to solve the problem. This correlation exists as far as humans are concerned [3]. We therefore aim in this paper to obtain image based attention regions that correlate better with human attention. We do that by obtaining a differential attention. The differential attention relies on an exemplar model of cognition.

In cognition studies, the exemplar theory suggests that humans are able to obtain generalisation for solving cogni-

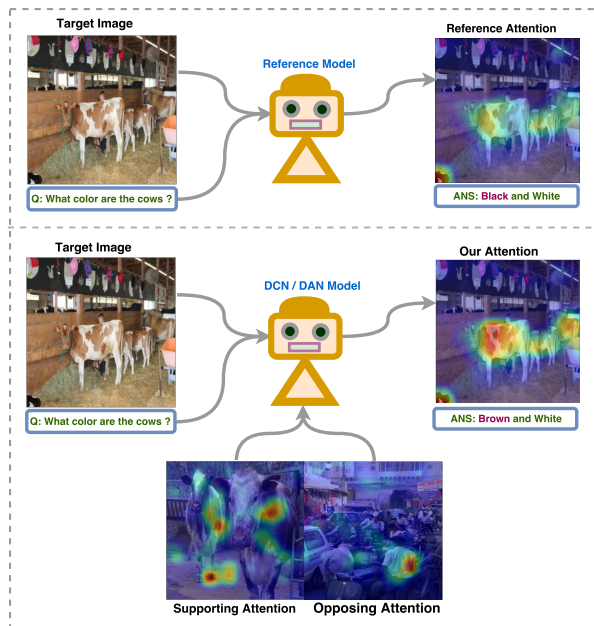


Figure 1. Illustration of improved attention obtained using Differential Context Network. Using the baseline reference we got answer as: “Black and White” But, using our methods DAN or DCN we get answer as “Brown and White”, that is actually the color of the cow. We provide the attention map that indicates the actual improvement in attention.

tive tasks by relying on an exemplar model. In this model, individuals compare new stimuli with the instances already stored in memory [9][20] and obtain answers based on these exemplars. We would like an exemplar model to provide attention. We want to focus on the specific parts in a nearest exemplar that distinguishes it from a far example. We do that by obtaining the differential attention region that distinguishes a supporting exemplar from an opposing exemplar. Our premise is that the difference between a nearest semantic exemplar and a far semantic exemplar can guide attention on a specific image region. We show that by using this differential attention mechanism we are able to obtain significant improvement in solving the visual question answering task. Further, we show that the obtained attention

regions are more correlated with human attention regions both quantitatively and qualitatively. We evaluate this on the challenging VQA-1 [2], VQA-2 [7] and HAT [3]

The main flow of the method followed is outlined in figure 1. Given an image and an associated question, we use an attention network to combine the image and question to obtain a reference attention embedding. This is used to order the examples in the database. We obtain a near example as a supporting exemplar and a far example as an opposing exemplar. These are used to obtain a differential attention vector. We evaluate two variants of our approach, one we term as ‘differential attention network’ (DAN) where the supporting and opposing exemplars are used only to provide a better attention on the image. The other we term as ‘differential context network’ (DCN) that obtains a differential context feature. This is obtained from the difference between supporting and opposing exemplars with respect to the original image to provide a differential feature. The additional context is used in answering questions. Both variants improve results over the baseline with the differential context variant being better.

Through this paper we provide the following contributions

- We adopt an exemplar based approach to improve visual question answering (VQA) methods by providing a differential attention.
- We evaluate two variants for obtaining differential attention - one where we only obtain attention and the other where we obtain differential context in addition to attention.
- We show that this method correlates better with human attention and results in an improved visual question answering that improves the state-of-the-art for image based attention methods. It is also competitive with respect to other proposed methods for this problem.

2. Related Work

The problem of Visual Question Answering (VQA) is a recent problem that was initiated as a new kind of visual Turing test. The aim was to show progress of systems in solving even more challenging tasks as compared to the traditional visual recognition tasks such as object detection and segmentation. An initial work in this area was by German *et al.* [6] that proposed this visual Turing test. Around the same time Malinowski *et al.* [16] proposed a multi-world based approach to obtain questions and answer them from images. These works aimed at answering questions of a limited type. In this work we aim at answering free-form open-domain[2] questions as was attempted by later works.

An initial approach towards solving this problem in the open-domain form was by [17]. This was inspired by the

work on neural machine translation that proposed translation as a sequence to sequence encoder-decoder framework [22]. However, subsequent works [19][2] approached the problem as a classification problem using encoded embeddings. They used soft-max classification over an image embedding (obtained by a CNN) and a question embedding (obtained using an LSTM). Further work by Ma *et al.* [15] varied the way to obtain an embedding by using CNNs to obtain both image and question embeddings. Another interesting approach [18] used dynamic parameter prediction where weights of the CNN model for the image embedding are modified based on the question embedding using hashing. These methods however, are not attention based. Use of attention enables us to focus on specific parts of an image or question that are pertinent for an instance and also offer valuable insight into the performance of the system.

There has been significant interest in including attention to solve the VQA problem. Attention based models comprises of image based attention models, question based attention and some that are both image and question based attention. In image based attention approach the aim is to use the question in order to focus attention over specific regions in an image [21]. An interesting recent work [25] has shown that it is possible to repeatedly obtain attention by using stacked attention over an image based on the question. Our work is closely related to this work. There has been further work [12] that considers a region based attention model over images. The image based attention has allowed systematic comparison of various methods as well as enabled analysis of the correlation with human attention models as shown by [3]. In our approach we focus on image based attention using differential attention and show that it correlates better with image based attention. There has been a number of interesting works on question based attention as well [26][24]. An interesting work obtains varied set of modules for answering questions of different types [1]. Recent work also explores joint image and question based hierarchical co-attention [14]. The idea of differential attention can also be explored through these approaches. However, we restrict ourselves to image based attention as our aim is to obtain a method that correlates well with human attention [3]. There has been an interesting work by [5] that advocates multimodal pooling and obtains state of the art in VQA. Interestingly, we show that by combining it with the proposed method further improves our results.

3. Method

In this paper we adopt a classification framework that uses the image embedding combined with the question embedding to solve for the answer using a softmax function in a multiple choice setting. A similar setting is adopted in the Stacked Attention Network (SAN) [25], that also aims at obtaining better attention and several other state-of-the-art

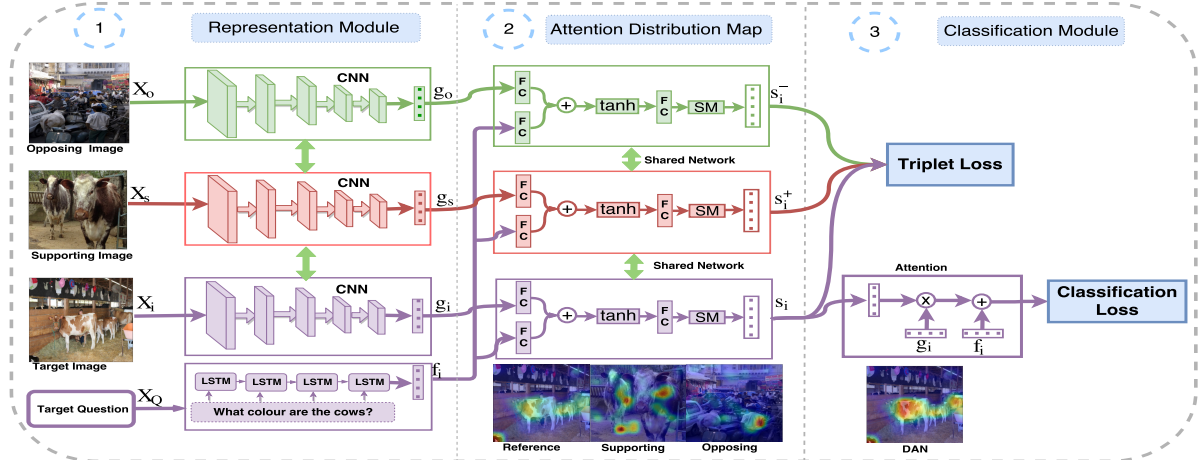


Figure 2. Differential Attention Network

methods. We provide two different variants for obtaining differential attention in the VQA system. We term the first variant a ‘Differential Attention Network’ (DAN) and the other a ‘Differential Context Network’ (DCN). We explain both the methods in the following sub-sections. A common requirement for both these tasks is to obtain nearest semantic exemplars.

3.1. Finding Exemplars

In our method, we use semantic nearest neighbors. Image level similarity does not suffice as the nearest neighbor may be visually similar but may not have the same context implied in the question (for instance, ‘Are the children playing?’ produces similar results for images with children based on visual similarity, whether the children are playing or not). In order to obtain semantic features we use a VQA system [13] to provide us with a joint image-question level embedding that relates meaningful exemplars. We compared image level features against the semantic nearest neighbors and observed that the semantic nearest neighbors were better. We used the semantic nearest neighbors in a k-nearest neighbor approach using a K-D tree data structure to represent the features. The ordering of the data-set features is based on the Euclidean distance. In section 4.1 we provide the evaluation with several values of nearest neighbors that were used as supporting exemplars. For obtaining opposing exemplar we used a far neighbor that was an order of magnitude further than the nearest neighbor. This we obtained through a coarse quantization of training data into bins. We specified the opposing exemplar as one that was around 20 clusters away in a 50 cluster ordering. This parameter is not stringent and it only matters that the opposing exemplar be far from the supporting exemplar. We show that using these supporting and opposing exemplars aids the method and any random ordering adversely affects the method.

3.2. Differential Attention Network (DAN)

In the DAN method, we use a multi-task setting. As one of the tasks we use a triplet loss [8] to learn a distance metric. This metric ensures that the distance between the attention weighted regions of near examples is less and the distance between attention weighted far examples is more. The other task is the main task of VQA. More formally, given an image x_i we obtain an embedding g_i using a CNN that we parameterize through a function $G(x_i, W_c)$ where W_c are the weights of the CNN. Similarly the question q_i results in a question embedding f_i after passing through an LSTM parameterised using the function $F(q_i, W_l)$ where W_l are the weights of the LSTM. This is illustrated in part 1 of figure 2. The output image embedding g_i and question embedding f_i are used in an attention network that combines the image and question embeddings with a weighted softmax function and produces an output attention weighted vector s_i . The attention mechanism is illustrated in figure 2. The weights of this network are learnt end-to-end learning using the two losses, a triplet loss and a soft-max classification loss for the answer (shown in part 3 of figure 2). The aim is to obtain attention weight vectors that bring the supporting exemplar attention close to the image attention and far from the opposing exemplar attention. The joint loss function used for training is given by:

$$L(\mathbf{s}, \mathbf{y}, \theta) = \frac{1}{N} \sum_{i=1}^N (L_{cross}(\mathbf{s}, \mathbf{y}) + \nu T(s_i, s_i^+, s_i^-)) \quad (1)$$

$$L_{cross}(\mathbf{s}, \mathbf{y}) = -\frac{1}{C} \sum_{j=1}^C y_j \log \mathbb{P}(c_j | \mathbf{s})$$

where θ is the set of model parameters for the two loss functions, y is the output class label and s is the input sample. C is the total number of classes in VQA (consists of the set

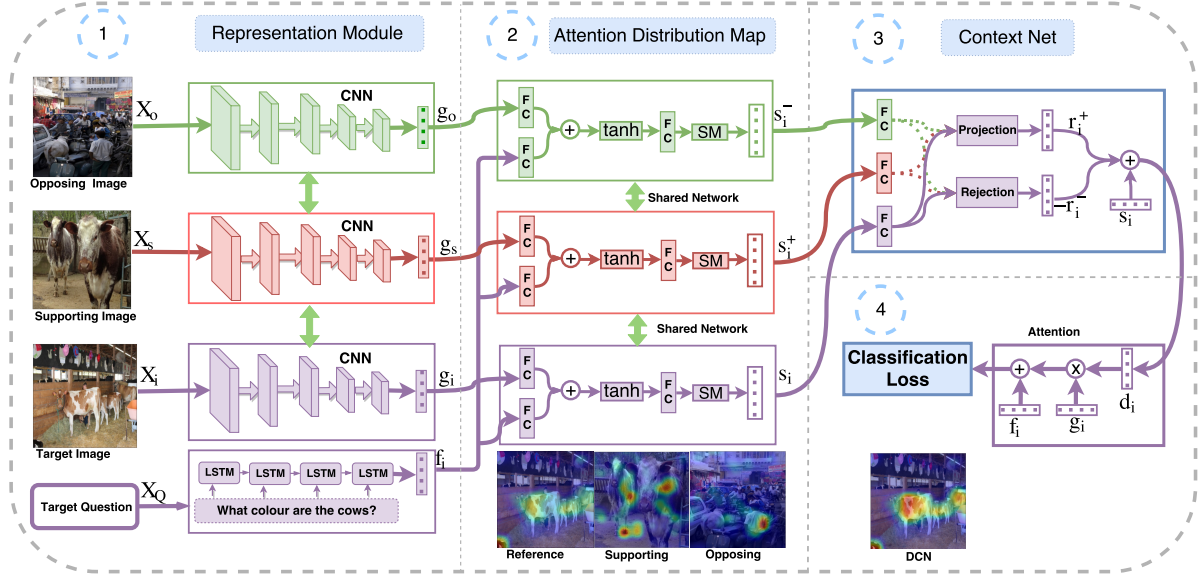


Figure 3. Differential Context Network

of total number of output classes including color, count etc.) and N is the total number of samples. The first term is the classification loss and the second term is the triplet loss. ν is a constant that controls the ratio between classification loss and triplet loss. $T(s_i, s_i^+, s_i^-)$ is the triplet loss function that is used. This is decomposed into two terms, one that brings the positive sample closer and one that pushes the negative sample farther. This is given by

$$T(s_i, s_i^+, s_i^-) = \max(0, \|t(s_i) - t(s_i^+)\|_2^2 + \alpha - \|t(s_i) - t(s_i^-)\|_2^2) \quad (2)$$

The constant α controls the separation margin between supporting and opposing exemplars. The constants ν and α are obtained through validation data.

The method is illustrated in figure 2. We further extend the model to a quintuplet setting where we bring two supporting attention weights closer and two opposing attention weights further in a metric learning setting. We observe in section 4.1 that this further improves the performance of the DAN method.

3.3. Differential Context Network (DCN)

We next consider the other variant that we propose where the differential context feature is added instead of only using it for obtaining attention. The first two parts are same as that for the DAN network. In part 1, we use the image, the supporting and the opposing exemplar and obtain the corresponding image and question embedding. This is followed by obtaining attention vectors s_i, s_i^+, s_i^- for the image, the supporting and the opposing exemplar. While in DAN, these were trained using a triplet loss function, in

DCN, we obtain two context features, the supporting context r_i^+ and the opposing context r_i^- . This is shown in part 3 in figure 3. The supporting context is obtained using the following equation

$$r_i^+ = (s_i \bullet s_i^+) \frac{s_i}{\|s_i\|_{L_2}^2} + (s_i \bullet s_i^-) \frac{s_i}{\|s_i\|_{L_2}^2} \quad (3)$$

where \bullet is the dot product. This results in obtaining correlations between the attention vectors.

The first term of the supporting context r_i^+ is the vector projection of s_i^+ on s_i and second term is the vector projection of s_i^- on s_i . Similarly, for opposing context we compute vector projection of s_i^+ on s_i and s_i^- on s_i . The idea is that the projection measures similarity between the vectors that are related. We subtract the vectors that are not related from the resultant. While doing so, we ensure that we enhance similarity and only remove the feature vector that is not similar to the original semantic embedding. This equation provides the additional feature that is supporting and is relevant for answering the current question q_i for the image x_i .

Similarly, the opposing context is obtained by the following equation

$$r_i^- = (s_i^+ - (s_i \bullet s_i^+) \frac{s_i}{\|s_i\|_{L_2}^2}) + (s_i^- - (s_i \bullet s_i^-) \frac{s_i}{\|s_i\|_{L_2}^2}) \quad (4)$$

We next compute the difference between the supporting and opposing context features i.e. $r_i^+ - r_i^-$ that provides us with the differential context feature \hat{d}_i . This is then either added with the original attention vector (DCN-Add) or multiplied with the original attention vector s_i (DCN-Mul)

Table 1. Analysis network parameter for DAN

Models	VQA1.0 Open-Ended (test-dev)				HAT val dataset
	All	Yes/No	Number	others	Rank-correlation
LSTM Q+I+ Attention(LQIA)	56.1	80.3	37.4	40.46	0.2142
DAN(K=1)+LQIA	59.2	80.1	36.1	46.6	0.2959
DAN(K=2)+LQIA	59.5	80.9	36.6	47.1	0.3090
DAN(K=3)+LQIA	59.9	80.6	37.2	47.5	0.3100
DAN(K=4)+LQIA	60.2	80.9	37.4	47.2	0.3206
DAN(K=1)+MCB	64.8	82.4	38.1	54.2	0.3284
DAN(K=2)+MCB	64.8	82.9	38.0	54.3	0.3298
DAN(K=3)+MCB	64.9	82.6	38.2	54.6	0.3316
DAN(K=4)+MCB	65.0	83.1	38.4	54.9	0.3326
DAN(K=5)+LQIA	58.1	79.4	36.9	45.7	0.2157
DAN(K=1,Random)+LQIA	56.4	79.3	37.1	44.6	0.2545

Table 2. Analysis network parameter for DCN

Models	VQA1.0 Open-Ended (test-dev)				HAT val dataset
	All	Yes/No	Number	others	Rank-correlation
LSTM Q+I+ Attention(LQIA)	56.1	80.3	37.4	40.46	0.2142
DCN Add.v1(K=4)+(LQIA)	60.4	81.0	37.5	47.1	0.3202
DCN Add.v2(K=4)+(LQIA)	60.4	81.2	37.2	47.3	0.3215
DCN Mul.v1(K=4)+(LQIA)	60.6	80.9	37.8	47.9	0.3229
DCN Mul.v2(K=4)+(LQIA)	60.9	81.3	37.5	48.2	0.3242
DCN Add.v1(K=4)+MCB	65.1	83.1	38.5	54.5	0.3359
DCN Add.v2(K=4)+MCB	65.2	83.4	39.0	54.6	0.3376
DCN Mul.v1(K=4)+MCB	65.2	83.9	38.7	54.9	0.3365
DCN Mul.v2(K=4)+MCB	65.4	83.8	39.1	55.2	0.3389

providing us with the final differential context attention vector d_i . This is then the final attention weight vector multiplied to the image embedding g_i to obtain the vector v_i that is then used with the classification loss function. This is shown in part 4 in the figure 3. The resultant attention is observed to be better than the earlier differential attention feature obtained through DAN as the features are also used as context.

The network is trained end-to-end using the following soft-max classification loss function

$$L(\mathbf{v}, \mathbf{y}, \theta) = - \sum_{j=1}^C y_j \log p(c_j | \mathbf{v}) \quad (5)$$

4. Experiments

The experiments have been conducted using the two variants of differential attention that are proposed and compared against baselines on standard datasets. We first analyze the different parameters for the two variants DAN and DCN that are proposed. We further evaluate the two networks by comparing the networks with comparable baselines and evaluate the performance against the state of the art methods. The main evaluation is conducted to evaluate the performance in terms of correlation of attention with human correlation where we obtain state-of-the-art in terms of correlation with human attention. Further, we observe

that its performance in terms of accuracy for solving the VQA task is substantially improved and is competitive with the current state of the art results on standard benchmark datasets. We also analyse the performance of the network on the recently proposed VQA2 dataset.

4.1. Analysis of Network Parameters

In the proposed DAN network, we have a dependency on the number of k-nearest neighbors that should be considered. We observe in table 1, that using 4 nearest neighbors in the triplet network we obtain the highest correlation with human attention as well as accuracy using VQA-1 dataset. We therefore use 4 nearest neighbors in our experiments. We observe that increasing nearest neighbors beyond 4 nearest neighbors results in reduction in accuracy. Further, even using a single nearest neighbor results in substantial improvement that is marginally improved as we move to 4 nearest neighbors.

We also evaluate the effect of using the nearest neighbors as obtained through a baseline model [2] versus using a random assignment of supporting and opposing exemplar. We observe that using DAN with a random set of nearest neighbors decreases the performance of the network. While comparing the network parameters, the comparable baseline we use is the basic model for VQA using LSTM and CNN [2]. This however does not use attention and we evaluate this

Table 3. Open-Ended VQA1.0 accuracy on test-dev

Models	All	Yes/No	Number	others
LSTM Q+I [2]	53.7	78.9	35.2	36.4
LSTM Q+I+ Attention(LQIA)	56.1	80.3	37.4	40.4
DPPnet [18]	57.2	80.7	37.2	41.7
SMem [24]	58.0	80.9	37.3	43.1
SAN [25]	58.7	79.3	36.6	46.1
QRU(1)[12]	59.3	81.0	35.9	46.0
DAN(K=4)+ LQIA	60.2	80.9	37.4	47.2
DMN+[23]	60.3	80.5	36.8	48.3
QRU(2)[12]	60.7	82.3	37.0	47.7
DCN Mul.v2(K=4)+LQIA	60.9	81.3	37.5	48.2
HieCoAtt [14]	61.8	79.7	38.9	51.7
MCB + att [5]	64.2	82.2	37.7	54.8
MLB [11]	65.0	84.0	37.9	54.7
DAN(K=4)+ MCB	65.0	83.1	38.4	54.9
DCN Mul.v2(K=4)+MCB	65.4	83.8	39.1	55.2

Table 4. VQA2.0 accuracy on Validation set for DCN and DAN

Models	All	Yes/No	Number	others
SAN-2	52.82	-	-	-
DAN(K=1) +LQIA	52.96	70.08	34.06	44.20
DCN Add.v1(K=1)+LQIA	53.01	70.13	33.98	44.27
DCN Add.v2(K=1) +LQIA	53.07	70.46	34.30	44.10
DCN Mul.v1(K=1) +LQIA	53.18	70.24	34.53	44.24
DCN Mul.v2(K=1)+LQIA	53.26	70.57	34.61	44.39
DCN Add.v1(K=4)+MCB	65.30	81.89	42.93	55.56
DCN Add.v2(K=4)+MCB	65.41	81.90	42.88	55.99
DCN Mul.v1(K=4)+MCB	65.52	82.07	42.91	55.97
DCN Mul.v2(K=4)+MCB	65.90	82.40	43.18	56.81

method with attention. With the best set of parameters the performance improves the correlation with human attention by 10.64%. We also observe that correspondingly the VQA performance improves by 4.1% over the comparable baseline. We further then incorporate this model with the model from MCB [5] which is a state of the art VQA model. This further improves the result by 4.8% more on VQA and a further increase in correlation with human attention by 1.2%.

In the proposed DCN network we have two different configurations, one where we use the add module (DCN-add) for adding the differential context feature and one where we use the (DCN-mul) multiplication module for adding the differential context feature. We further have a dependency on the number of k-nearest neighbors for the DCN network as well. This is also considered. We next evaluate the effect of using a fixed scaling weight (DCN_v1) for adding the differential context feature against learning a linear scaling weight (DCN_v2) for adding the differential context feature. All these parameter results are compared in table 2.

As can be observed from table 2 the configuration that

obtains maximum accuracy on VQA dataset [2] and in correlation with human attention is the version that uses multiplication with learned weights and with 4 nearest neighbors being considered. This results in an improvement of 11% in terms of correlation with human attention and 4.8% improvement in accuracy on the VQA-1 dataset [2]. We also observe that incorporating DCN with MCB [5] further improves the results by 4.5% further on VQA dataset and results in an improvement of 1.47% improvement in correlation with attention. These configurations are used in comparison with the baselines.

4.2. Comparison with baseline and state of the art

We obtain the initial comparison with the baselines on the rank correlation on human attention (HAT) dataset [3] that provides human attention by using a region deblurring task while solving for VQA. Between humans the rank correlation is 62.3%. The comparison of various state-of-the-art methods and baselines are provided in table 5. The baseline we use [2] is the method used by us for obtaining exem-

Table 5. Rank Correlation on HAT Validation Dataset for DAN and DCN

Models	Rank-correlation
LSTM Q+I+ Attention(LQIA)	0.214 ± 0.001
SAN[3]	0.249 ± 0.004
HieCoAtt-W[14]	0.246 ± 0.004
HieCoAtt-P [14]	0.256 ± 0.004
HieCoAtt-Q[14]	0.264 ± 0.004
MCB + Att.	0.279 ± 0.004
DAN (K=4) +LQIA	0.321± 0.001
DCN Mul_v2(K=4) +LQIA	0.324± 0.001
DAN (K=4) +MCB	0.332± 0.001
DCN Mul_v2(K=4) +MCB	0.338± 0.001
Human [3]	0.623 ± 0.003

plars. This uses a question embedding using an LSTM and an image embedding using a CNN. We additionally consider a variant of the same that uses attention. We have also obtained results for the stacked attention network [25]. The results for the Hierarchical Co-Attention work [14] are obtained from the results reported in Das *et al.* [3]. We observe that in terms of rank correlation with human attention we obtain an improvement of around 10.7% using DAN network (with 4 nearest neighbors) and using DCN network (4 neighbors with multiplication module and learned scaling weights) we obtain an improvement of around 11% over the comparable baseline. We also obtain an improvement of around 6% over the Hierarchical Co-Attention work [14] that uses co-attention on both image and questions. Further incorporating MCB improves the results for both DAN and DCN resulting in an improvement of 7.4% over Hierarchical co-attention work and 5.9% improvement over MCB method. However, as noted by [3], using a saliency based method [10] that is trained on eye tracking data to obtain a measure of where people look in a task independent manner results in more correlation with human attention (0.49). However, this is explicitly trained using human attention and is not task dependent. In our approach, we aim to obtain a method that can simulate human cognitive abilities for solving tasks.

We next evaluate the different baseline and state of the art methods on the VQA dataset [2] in table 3. There have been a number of methods proposed for this benchmark dataset for evaluating the VQA task. Among the notable different methods, the Hierarchical Co-Attention work [14] obtains 61.8% accuracy on VQA task, the dynamic parameter prediction [18] method obtains 57.2% and the stacked attention network [25] obtains 58.7% accuracy. We observe that the differential context network performs well outperforming all the image based attention methods and results in an accuracy of 60.9%. This is a strong result and we ob-

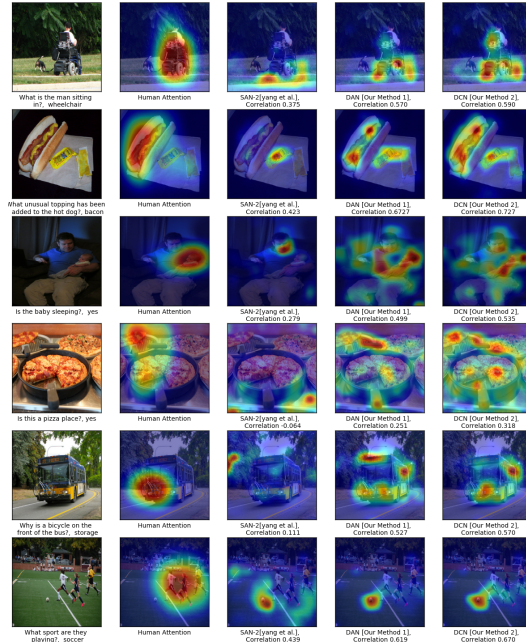


Figure 4. In this figure, the first column indicates target question and corresponding image, second column indicates reference human attention map in HAT dataset, third column refer to generated attention map for SAN, fourth column refers to rank correlation of our DAN model and final column refers to rank correlation for our DCN model.

serve that the performance improves across different kinds of questions. Further, on combining the method with MCB, we obtain improved results of 65% and 65.4% using DAN and DCN respectively improving over the results of MCB by 1.2%. This is consistent with the improved correlation with human attention that we observe in table 5.

We next evaluate the proposed method on a recently proposed VQA-2 dataset [7]. The aim in this new dataset is to remove the bias in different questions. It is a more challenging dataset as compared to the previous VQA-1 dataset [2]. We provide a comparison of the proposed DAN and DCN methods against the stacked attention network (SAN) [25] method. As can be observed in table 4, the proposed methods obtain improved performance over a strong stacked attention baseline. We observe that our proposed methods are also able to improve the result over the SAN method. DCN with 4 nearest neighbors when combined with MCB obtains an accuracy of 65.90%

4.3. Attention Visualization

The main aim of the proposed method is to obtain an improved attention that correlates better with human attention. Hence we visualize the attention regions and compare them. In attention visualization we overlay the attention probability distribution matrix, that is the most prominent part of a given image based on the query question. The procedure



Figure 5. In this figure, the first row indicates the given target image, supporting image and opposing image. second row indicates the attention map for human[3], reference attention map, supporting attention map, opposing attention map, DAN and DCN attention map respectively. Third row generates result by applying attention map on corresponding images.

followed is same as that followed by Das *et al.* [3]. We provide the results of the attention visualization in figure 4. We obtain significant improvement in attention by using DCN as compared to the SAN method [25]. Figure 5 provides how the supporting and opposing attention map helps to improve the reference attention using DAN and DCN. We have provided more results for attention map visualization on the project website ¹.

5. Discussion

In this section we further discuss different aspects of our method that are useful for understanding the method in more detail

We first consider how exemplars improve attention. In differential attention network, we use the exemplars and train them using a triplet network. It is known that using a triplet ([8] and earlier by [4]), that we can learn a representation that accentuates how the image is closer to the supporting exemplar as against the opposing exemplar. The attention is obtained between the image and language representations. Therefore the improved image representation helps in obtaining an improved attention vector. In DCN the same approach is used with the change that the differential exemplar feature is also included in the image representation using projections. More analysis in terms of understanding how the methods qualitatively improves attention is included in the project website.

We next consider whether improved attention implies improved performance. In our empirical analysis we ob-

served that we obtain improved attention *and* improved accuracies in VQA task. While there could be other ways of improving performance on VQA (as suggested by MCB[5]) these can be additionally incorporated with the proposed method and these do yield improved performance in VQA

Lastly we consider whether image (I) and question embedding (Q) are both relevant. We had considered this issue and had conducted experiments by considering I only, by considering Q only, and by considering nearest neighbor using the semantic feature of both Q & I. We had observed that the Q & I embedding from the baseline VQA model performed better than other two. Therefore we believe that both contribute to the embedding.

6. Conclusion

In this paper we propose two different variants for obtaining differential attention for solving the problem of visual question answering. These are differential attention network (DAN) and differential context network (DCN). Both the variants provide significant improvement over the baselines. The method provides an initial view of improving VQA using an exemplar based approach. In future, we would like to further explore this model and extend it to joint image and question based attention models.

7. Acknowledgment

We acknowledge the help provided by our Delta Lab members and our family who have supported us in our research activity.

¹project website: <https://badripatro.github.io/DVQA/>

References

- [1] J. Andreas, M. Rohrbach, T. Darrell, and D. Klein. Learning to compose neural networks for question answering. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, 2016. 2
- [2] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. VQA: Visual Question Answering. In *International Conference on Computer Vision (ICCV)*, 2015. 2, 5, 6, 7
- [3] A. Das, H. Agrawal, C. L. Zitnick, D. Parikh, and D. 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. 1, 2, 6, 7, 8
- [4] A. Frome, Y. Singer, F. Sha, and J. Malik. Learning globally-consistent local distance functions for shape-based image retrieval and classification. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8. IEEE, 2007. 8
- [5] 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 *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Austin, Texas, USA, 2016. 2, 6, 8
- [6] D. Geman, S. Geman, N. Hallonquist, and L. Younes. Visual turing test for computer vision systems. *Proceedings of the National Academy of Sciences of the United States of America*, 112(12):3618–3623, 03 2015. 2
- [7] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2017. 2, 7
- [8] E. Hoffer and N. Ailon. Deep metric learning using triplet network. In *International Workshop on Similarity-Based Pattern Recognition*, pages 84–92. Springer, 2015. 3, 8
- [9] F. Jäkel, B. Schölkopf, and F. Wichmann. Generalization and similarity in exemplar models of categorization: Insights from machine learning. *Psychonomic Bulletin and Review*, 15(2):256–271, Apr. 2008. 1
- [10] T. Judd, K. Ehinger, F. Durand, and A. Torralba. Learning to predict where humans look. In *Computer Vision, 2009 IEEE 12th international conference on*, pages 2106–2113. IEEE, 2009. 7
- [11] J.-H. Kim, K. W. On, W. Lim, J. Kim, J.-W. Ha, and B.-T. Zhang. Hadamard Product for Low-rank Bilinear Pooling. In *The 5th International Conference on Learning Representations*, 2017. 6
- [12] R. Li and J. Jia. Visual question answering with question representation update (qru). In *Advances in Neural Information Processing Systems*, pages 4655–4663, 2016. 2, 6
- [13] J. Lu, X. Lin, D. Batra, and D. Parikh. Deeper lstm and normalized cnn visual question answering model. https://github.com/VT-vision-lab/VQA_LSTM_CNN, 2015. 3
- [14] J. Lu, J. Yang, D. Batra, and D. Parikh. Hierarchical question-image co-attention for visual question answering. In *Advances In Neural Information Processing Systems*, pages 289–297, 2016. 2, 6, 7
- [15] L. Ma, Z. Lu, and H. Li. Learning to answer questions from image using convolutional neural network. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016. 2
- [16] M. Malinowski and M. Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *Advances in Neural Information Processing Systems*, 2014. 2
- [17] M. Malinowski, M. Rohrbach, and M. Fritz. Ask your neurons: A neural-based approach to answering questions about images bibtex. In *International Conference on Computer Vision (ICCV)*, 2015. 2
- [18] H. Noh, P. Hongsuck Seo, and B. Han. Image question answering using convolutional neural network with dynamic parameter prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 30–38, 2016. 2, 6, 7
- [19] M. Ren, R. Kiros, and R. Zemel. Exploring models and data for image question answering. In *Advances in Neural Information Processing Systems*, pages 2953–2961, 2015. 2
- [20] R. Shepard. Toward a universal law of generalization for psychological science. *Science*, 237(4820):1317–1323, 1987. 1
- [21] K. J. Shih, S. Singh, and D. Hoiem. Where to look: Focus regions for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4613–4621, 2016. 2
- [22] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *Proceedings of the 27th International Conference on Neural Information Processing Systems, NIPS’14*, pages 3104–3112, 2014. 2
- [23] C. Xiong, S. Merity, and R. Socher. Dynamic memory networks for visual and textual question answering. In *Proceedings of International Conference on Machine Learning (ICML)*, 2016. 6
- [24] H. Xu and K. Saenko. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In *European Conference on Computer Vision*, pages 451–466. Springer, 2016. 2, 6
- [25] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 21–29, 2016. 2, 6, 7, 8
- [26] Y. Zhu, O. Groth, M. Bernstein, and L. Fei-Fei. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4995–5004, 2016. 2