

Progressive Attention Networks for Visual Attribute Prediction

Paul Hongsuck Seo¹
hsseo@postech.ac.kr

Zhe Lin²
zlin@adobe.com

Scott Cohen²
scohen@adobe.com

Xiaohui Shen²
xshen@adobe.com

Bohyung Han³
bhhan@snu.ac.kr

¹ POSTECH
South Korea

² Adobe Research
USA

³ Seoul National University
South Korea

Abstract

We propose a novel attention model that can accurately attends to target objects of various scales and shapes in images. The model is trained to gradually suppress irrelevant regions in an input image via a progressive attentive process over multiple layers of a convolutional neural network. The attentive process in each layer determines whether to pass or block features at certain spatial locations for use in the subsequent layers. The proposed progressive attention mechanism works well especially when combined with hard attention. We further employ local contexts to incorporate neighborhood features of each location and estimate a better attention probability map. The experiments on synthetic and real datasets show that the proposed attention networks outperform traditional attention methods in visual attribute prediction tasks.

1 Introduction

Attentive mechanisms often play important roles in modern neural networks (NNs) especially in computer vision tasks. Many visual attention models have been introduced in the previous literature, and they have shown that attaching an attention to NNs improves the accuracy in various tasks such as image classification [1, 2, 3], image generation [4], image caption generation [5] and visual question answering [6, 7, 8].

There are several motivations for incorporating attentive mechanisms in NNs. One of them is analogy to the human perceptual process. The human visual system often pays attention to a region of interest instead of processing an entire scene. Likewise, in a neural attention model, we can focus only on attended areas of the input image. This is beneficial in terms of computational cost; the number of hidden units may be reduced since the hidden activations only need to encode the region with attention [9].

Another important motivation is that various high-level computer vision tasks should identify a particular region for accurate attribute prediction. Figure 1 illustrates an example




Figure 1: Intermediate attention maps of our progressive attention method to solve a reference problem (with query 7 and answer red). It shows that attention is gradually refined through the network layers for resolving the reference problem. Distracting patterns in smaller scales are suppressed at earlier layers while those in larger scales (e.g., 9) are suppressed at later layers with larger receptive fields. All attended images are independently rescaled for better visualization.

task to predict the color (answer) of a given input number (query). The query specifies a particular object in the input image ('7' in this example) for answering its attribute (red). To address this type of tasks, the network architecture should incorporate an attentive mechanism either explicitly or implicitly.

One of the most popular attention mechanisms for NNs is soft attention [13], which aggregates responses in a feature map weighted by their attention probabilities. This process results in a single attended feature vector. Since the soft attention method is fully differentiable, the entire network can be trained end-to-end using a standard backpropagation. However, it can only model attention to local regions with a fixed size depending on the receptive field of the layer chosen for attention. This makes the soft attention method inappropriate for complicated cases, where objects involve significant variations in their scale and shape.

To overcome this limitation, we propose a novel attention network, referred to as *progressive attention network* (PAN), which enables precise attention over objects of different scales and shapes by attaching attentive mechanisms to multiple layers within a convolutional neural network (CNN). More specifically, the proposed network predicts attentions on intermediate feature maps and forwards the attended feature maps in each layer to the subsequent layers in CNN. Moreover, we improve the proposed progressive attention by replacing feature aggregation, which may distort the original feature semantics, with hard attention via likelihood marginalization. The contribution of this work is four-fold:

- A novel attention model (progressive attention network) learned to handle accurate scale and shape of attentions for a target object,
- Integration of hard attention with likelihood marginalization into the proposed progressive attention networks,
- Use of local contexts to improve stability of the progressive attention model,
- Significant performance improvement over traditional single-step attention models in query-specific visual attribute prediction tasks.

The rest of this paper is organized as follows. We first review related work in Section 2. Section 3 describes the proposed model with local context information. We then present our experimental results on several datasets in Section 4 and conclude the paper in Section 5.

2 Related Work

Attention on features The most straightforward attention mechanism is a feature based method, which selects a subset of features by explicitly attaching an attention model to NN

architectures. This approach has improved performance in many tasks [11, 12, 13, 14, 15, 16, 17, 18, 19]. For example, it has been used to handle sequences of variable lengths in neural machine translation models [12, 14], speech recognition [11] and handwriting generation [13], and manage memory access mechanisms for memory networks [14] and neural turing machines [15]. When applied to computer vision tasks to resolve reference problems, these models are designed to pay attention to CNN features corresponding to subregions in input images. Image caption generation and visual question answering are often benefited from this attention mechanism [11, 16, 17, 18, 19, 20].

Attention by image transformation Another stream of attention model is image transformation. This approach identifies transformation parameters fitting region of interest. [21] and [22] transform an input image with predicted translation parameters (t_x and t_y) and a fixed scale factor ($\delta < 1$) for image classification or multiple object recognition. Scale factor is also predicted in [22] for image generation. Spatial transformer networks (STNs) predict all six parameters of affine transformation, and even extend it to a projective transformation and a thin plate spline transformation [23]. However, STN is limited to attending a single candidate region defined by a small number of parameters in an image. Our model overcomes this issue by formulating attention as progressive filtering on feature maps instead of assuming that objects are roughly aligned by a constrained spatial transformation.

Multiple attention processes There have been several approaches iteratively performing attentive processes to utilize relations between objects. For example, [24] iteratively attends to images conditioned on the previous attention states for visual question answering. Iterative attention mechanisms to memory cells is incorporated to retrieve different values stored in the memory [9, 19]. In [25], an extension of STN to multiple transformer layers has also been presented but is still limited to rigid shape of attention. Our model is similar to these approaches, but aims to attend to target regions via operating on multiple CNN layers in a progressive manner; attention information is predicted progressively from feature maps through multiple layers of CNN to capture the detailed shapes of target objects.

Training attention models The networks with soft attention are fully differentiable and thus trainable end-to-end by backpropagation. Stochastic hard attention has been introduced in [23, 26], where networks explicitly select a single feature location based on the predicted attention probability map. Because the explicit selection (or sampling) procedure is not differentiable, hard attention methods employ REINFORCE learning rule [27] for training. Transformation-based attention models [9, 26] are typically learned by REINFORCE as well while STN [23] proposes a fully differentiable formulation suitable for end-to-end learning. The proposed network has advantage of end-to-end trainability by a standard backpropagation without any extra technique.

3 Progressive Attention Networks (PANs)

To alleviate the limitations of existing attention models in handling variable object scales and shapes, we propose a progressive attention mechanism. We describe technical details about our attention model in this section.

3.1 Progressive Attentive Process

Let $\mathbf{f}^l \in \mathbb{R}^{H_l \times W_l \times C_l}$ be an output feature map of a layer $l \in \{0, \dots, L\}$ in CNN with width W_l , height H_l and number of channels C_l , and $f_{i,j}^l \in \mathbb{R}^{C_l}$ be a feature at (i, j) of feature map \mathbf{f}^l . In

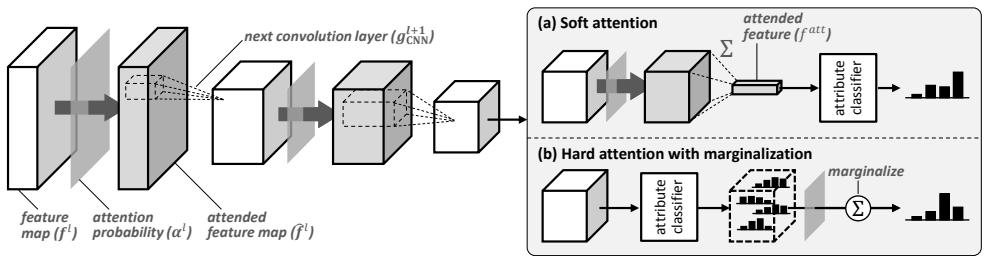


Figure 2: Overall procedure of our progressive attention model. Attentive processes are successively applied to feature maps at multiple layers and the resulting attended feature maps are used as input feature maps for the next convolution layers in the CNN. Attention probabilities α^l are estimated from feature maps and input query. In the last attention layer, (a) the attended feature maps are aggregated to a single feature vector (by sum pooling) as in soft attention and fed to the final attribute classifier, or (b) attention probabilities are used to marginalize attribute predictions from hard-selected feature at each spatial location as in hard attention. The bar graphs in dotted cube denote attribute prediction results from individual spatial locations and Σ in circle represents marginalization with attention probabilities.

the proposed model, an attentive process is applied to multiple layers of CNN and we obtain the following attended feature map $\hat{f}^l = [\hat{f}_{i,j}^l]$, where

$$\hat{f}_{i,j}^l = \alpha_{i,j}^l f_{i,j}^l. \quad (1)$$

In the above equation, attention probability $\alpha_{i,j}^l$ for a feature $f_{i,j}^l$ and a query q are given by

$$s_{i,j}^l = g_{\text{att}}^l(f_{i,j}^l, q; \theta_{\text{att}}^l), \quad \text{and} \quad \alpha_{i,j}^l = \begin{cases} \text{softmax}_{i,j}(s^l) & \text{if } l = L \\ \sigma(s_{i,j}^l) & \text{otherwise} \end{cases} \quad (2)$$

where $g_{\text{att}}^l(\cdot)$ denotes the attention function with a set of parameters θ_{att}^l , $s_{i,j}^l$ is the attention score at (i, j) , and $\sigma(\cdot)$ is a sigmoid function. The attention probability at each location is estimated independently in the same feature map; a sigmoid function is employed to constrain attention probabilities between 0 and 1. For the last layer of attention, we use a softmax function over the entire spatial region for later feature aggregation.

Unlike the soft attention model [23], the attended feature map \hat{f}^l in the intermediate attention layers is not summed up to generate a single vector representation of the attended regions. Instead, the attended feature map is forwarded to the next layer as an input to compute the next feature map, which is given by

$$f^{l+1} = g_{\text{CNN}}^{l+1}(\hat{f}^l; \theta_{\text{CNN}}^{l+1}) \quad (3)$$

where $g_{\text{CNN}}^{l+1}(\cdot)$ is the convolution operation at layer $l+1$ in CNN parameterized by $\theta_{\text{CNN}}^{l+1}$. This feedforward procedure with attentive processes in CNN is performed from the input layer, where $f^0 = I$, until \hat{f}^L is obtained. The attended feature f^{att} is finally retrieved by summing up all the features in the final attended feature map \hat{f}^L as in soft attention, which is given by

$$f^{\text{att}} = \sum_i^H \sum_j^W \hat{f}_{i,j}^L = \sum_i^H \sum_j^W \alpha_{i,j}^L f_{i,j}^L. \quad (4)$$

The attended feature f^{att} obtained by this process is then used as the input to visual attribute classifier as shown in Figure 2a.

In our models, we place the attention layers to the output of max pooling layers instead of




Figure 3: Attention estimation (a) without and (b) with local context. In (a), $\alpha_{i,j}^l$ is predicted from $f_{i,j}^l$ only while its spatially adjacent features are also employed to estimate $\alpha_{i,j}^l$ in (b).

every layer in CNN because the reduction of feature resolution within CNN mainly comes from pooling layers. In practice, we can also skip the first few pooling layers and attach the attention module only to the outputs of last K pooling layers.

3.2 Incorporating Hard Attention

The derivation of the final attended feature f^{att} shown in Eq. (4) is similar to soft attention mechanism. However, since f^{att} is obtained by a weighted sum of features at all locations, it loses semantic layout of the original feature map. To overcome this limitation, we extend our model to incorporate hard attention. Given the final feature map f^L , the hard attention model predicts an answer distribution using the selected feature $f_{i,j}^L$ as

$$p(a|f^L, i, j) = H(f_{i,j}^L) \quad (5)$$

where a is the predicted answer and $H(\cdot)$ is the visual attribute classifier. The final answer distribution given the feature map f^L is obtained by

$$p(a|f^L) = \sum_i^H \sum_j^W p(a|f^L, i, j) p(i, j|f^L) = \sum_i^H \sum_j^W H(f_{i,j}^L) \alpha_{i,j}^L \quad (6)$$

where $p(i, j|f^L)$ is modeled by $\alpha_{i,j}^L$ in Eq. (2). This marginalization eliminates the need for REINFORCE technique used in [23] by removing the hard selection process, and helps maintain the unique characteristics of feature maps.

3.3 Multi-Resolution Attention Estimation

As shown in Eq. (2), the resolution of attention probability map α^l depends on the size of the feature map in the corresponding layer. Since the resolution of α^l decreases inherently with increasing depth of CNNs and the attentive processes are performed over multiple layers recursively in our framework, we can attend to the regions with arbitrary resolutions even in higher layers. Hence, the proposed network exploits high-level semantics in deep representations for inference without losing attention resolution.

The progressive attention model is also effective in predicting fine attention shapes as attention information is aggregated over multiple layers to suppress irrelevant structures at different levels. In lower layers, features whose receptive fields contain small distractors are suppressed first while the features from a part of large distractors remain intact. In higher layers, features corresponding to these large distractors would have low attention probability as each feature contains information from larger receptive fields enabling the attention module to distinguish between distractors and target objects. This phenomenon is demonstrated well in the qualitative results in our experiments (Section 4). An additional benefit of our progressive attention is that inference is based only on feedforward procedure.




Figure 4: Example of the MREF datasets.

3.4 Local Context

PAN can improve the quality of attention estimation by allowing its attention layers to observe a local context of the target feature. The local context $\mathcal{F}_{i,j}^l$ of a feature $f_{i,j}^l$ is composed of its spatially adjacent features as illustrated in Figure 3, and is formally denoted by $\mathcal{F}_{i,j}^l = \{f_{s,t}^l | i - \delta \leq s \leq i + \delta, j - \delta \leq t \leq j + \delta\}$. The attention score is now predicted by the attention network with local context as

$$s_{i,j}^l = g_{\text{att}}^l(\mathcal{F}_{i,j}^l, q; \theta_{\text{att}}^l). \quad (7)$$

In our architecture, the area of the local context is given by the filter size corresponding to the composite operation of convolution followed by pooling in the next layer. The local context does not need to be considered in the last layer of attention since its activations are used to compute the final attended feature map. Local context improves attention prediction quality by comparing centroid features with their surroundings and making the estimated attention more discriminative.

3.5 Training Progressive Attention Networks


Training a PAN is as simple as training a soft attention network [23] because every operation within the network is differentiable. The entire network with both soft and hard attention is trained end-to-end by the standard backpropagation minimizing the cross entropy of the object-specific visual attributes. When we train it from a pretrained CNN, the CNN part should always be fine-tuned together since the intermediate attention maps may change the input distributions of their associated layers in the CNN.

4 Experiments


This section discusses experimental results on two datasets. Note that our experiments focus on the tasks directly related to visual attention to minimize any artifacts caused by irrelevant components. The codes are available at <https://github.com/hworang77/PAN>.

4.1 MNIST Reference

Datasets We conduct experiments on synthetic datasets created from MNIST [24]. The first synthetic dataset is referred to as MNIST Reference (MREF; Figure 4a), where each training example is a triple of an image, a query number, and its color label. The task on this dataset is to predict the color of the number given by a query. Five to nine distinct MNIST digits with different colors out of $\{\text{green, yellow, white, red, blue}\}$ and scales in $[0.5, 3.0]$ are randomly sampled and located within a 100×100 empty image with black background. When coloring numbers, Gaussian noise is added to the reference color value. To simulate more realistic situations, we made two variants of MREF by adding distractors (MDIST; Figure 4b) or replacing background with natural images (MBG; Figure 4c). Distractors in MDIST are constructed with randomly cropped 5×5 patches of MNIST images whereas backgrounds



(a) Network architectures of models on MREF. Arrows represents direct connection to next layer without attention.



(b) Architecture of attention function $g_{\text{att}}^l(\cdot)$. Local contexts $\mathcal{F}_{i,j}^l$ are used only in PAN[*]+CTX.

Figure 5: Detailed illustration of network architectures on MNIST Reference experiments.

of MBG are filled with natural scene images randomly chosen from the SUN Database [20]. The training, validation and test sets contain 30K, 10K and 10K images, respectively.

Experimental setting We implement two variants of the proposed network with soft and hard attention at its final layer, denoted by PAN[S] and PAN[H], respectively. Each of them is implemented with and without the local context observation; the networks with local context observation are denoted by ‘+CTX’ in our results. In addition, soft attention network (SAN), hard attention network (HAN) [23], and two variants of spatial transformer network (STN-S and STN-M) [10] are used as baseline models for comparisons. While STN-S is the model with a single transformer layer, STN-M contains multiple transformer layers in the network. We reimplemented SAN and STNs following the descriptions in [23] and [10], respectively, and trained HAN by optimizing the marginal log-likelihood. The architecture of image encoding networks in SAN and HAN, and localization networks in STNs are all identical for fair comparisons. The CNN in the proposed network also has the same architecture except for the additional layers for progressive attention. The CNN is composed of four stacks of 3×3 convolutions with 32 channels (stride 1) followed by a 2×2 max pooling layer (stride 2) as illustrated in Figure 5a. We used a single fc layer for attribute classifier because the task requires simple color prediction. The attention functions $g_{\text{att}}^l(\cdot)$ for all models are given by multi-layer perceptrons with two layers (Figure 5b). The function takes the concatenation of a query vector q and a feature vector $f_{i,j}^l$, and outputs an attention score $s_{i,j}^l$, where the query vector q is a one-hot vector representing a target object. In the proposed models, the intermediate attention functions additionally take local context $\mathcal{F}_{i,j}^l$ containing spatially adjacent features with $\delta = 1$. Every model is trained end-to-end from scratch by RMSprop until the models show no improvement for 50 epochs. We exponentially decay the initial learning rate by 0.9 at every 30 epoch after the 50th epoch, and grid-search the best initial learning rate for each model in [0.001, 0.005].

Results Table 1(left) presents color prediction accuracy of all compared algorithms. PAN[*] clearly outperforms other approaches with significant margins and PAN[*]+CTX further improves performance by exploiting local context for attention estimation. Our progressive attention model with the hard attention process and local context achieves the best performance in most cases. Note that the attention capability of the baseline models is restricted to either rhombic or coarsely shaped regions. In contrast, the proposed networks predict attention maps with arbitrary shapes by capturing spatial support of target area better.

Table 1: Results on MREF datasets: (left) color prediction accuracy [%] and true-positive ratio [%]. (right) accuracy [%] with different scales on MBG test subset.

	MREF		MDIST		MBG		Scale ranges					
	Acc.	TPR	Acc.	TPR	Acc.	TPR	0.5-1.0	1.0-1.5	1.5-2.0	2.0-2.5	2.5-3.0	
Uniform	20.00	2.34	20.00	2.35	20.00	2.39	31.1	35.3	35.8	37.7	45.0	
STN-S	39.10	0.94	38.32	0.52	32.27	0.74	51.8	58.3	55.3	58.5	61.7	
STN-M	93.89	0.66	85.09	0.74	52.25	0.65	49.1	68.9	63.6	61.7	54.6	
SAN	82.94	13.83	75.73	12.82	53.77	6.92	51.2	68.7	70.8	64.0	59.5	
HAN	81.84	13.95	78.49	13.81	55.84	7.64	67.9	76.6	73.9	69.6	61.3	
PAN[S]	95.92	44.25	91.65	52.85	69.46	32.07	70.8	79.6	76.8	70.2	67.3	
PAN[H]	95.32	62.31	90.46	52.07	72.50	38.32	79.6	86.8	85.1	81.9	72.1	
PAN[S]+CTX	98.51	60.45	96.02	59.60	81.01	43.87	84.2	87.1	87.5	81.6	79.6	
PAN[H]+CTX	98.53	62.36	95.82	61.39	84.62	51.49						




Figure 6: Qualitative results of SAN, HAN, PAN[S]+CTX and PAN[H]+CTX. (a) Magnitude of activations in feature maps $f_{i,j}^l$ before attention; the activations are mapped to original image space by spreading activations to their receptive fields. (b) Magnitude of activations in attended feature maps $\hat{f}_{i,j}^l$ showing the effect of attention in contrast to (a). For PAN[*]+CTX, we only show last two attentions, which accumulate the attentions of earlier layers. Every map is rescaled into $[0, 1]$ by $(x - \min)/(\max - \min)$.

We also present the attention quality of the models using true-positive ratio (TPR) in Table 1(left). TPR measures how strong attention is given to proper location by computing the ratio of the aggregated attention probability within the desired area (*i.e.*, ground-truth segmentation) to the attention probability in the whole image. To calculate TPR of STN baselines, we assigned the uniform attention probabilities to the attended rhombic regions. The models with progressive attention give the best results with significant margins compared to all the other methods. These results suggest that progressive attention constructs more accurate shapes of attended regions than all other attention models. Integrating hard attention and local context further improves overall performance in most cases.

To evaluate scale sensitivity, we divide test sets into five subsets based on target object scales with uniform intervals and computed accuracies of the models. In Table 1(right), the results on MBG show that the models with progressive attention are robust to scale variations due to their multi-scale attention mechanism especially when the hard attention and local contexts are incorporated. The tendencies on MREF and MDIST are also similar.

An important observation regarding to STNs is that these models actually do not attend to target objects. Instead, STNs generate almost identical transformation regardless of input images and pass the filtering process to next layers. As the results, the transformed images are padded ones containing the entire original image with different affine transformations. Consequently, these models show very low TPRs, even lower than the uniform attention.

Table 2: Weighted mAP of the attribute prediction and TPR of attentions measured with ground-truth bounding boxes on VG dataset (left). TPR of attentions measured with ground-truth segmentations in VOC 2007 (right).

	attention only		w/ prior		TPR (VOC 2007)	
	mAP	TPR	mAP	TPR		
STN-S	28.87	11.59	30.50	7.78	SAN	22.01
STN-M	29.12	1.99	31.17	2.28	HAN	24.91
					PAN[S]+CTX	27.16
HAN	27.72	17.24	31.84	17.65	PAN[H]+CTX	31.79
PAN[S]+CTX	29.38	18.01	32.50	20.17		
PAN[H]+CTX	30.00	22.23	34.19	24.37		

Figure 6 illustrates qualitative results of the two proposed methods and two baselines on MBG dataset. Our models yield accurate attended regions by gradually augmenting attention and suppressing irrelevant regions in the image. We observe that the proposed models maintain high attention resolution through the progressive attention process. In contrast, the baseline models attend to target regions only once at the top layer resulting in coarse attention.

4.2 Attribute prediction on Visual Genome

Dataset Visual Genome (VG) [10] is an image dataset containing several types of annotations: question/answer pairs, image captions, objects, object attributes and object relationship. We formulate object attribute prediction as a multi-label classification task with reference. Given an input image and a query (*i.e.*, an object category), we predict binary attributes of individual objects specified by the query. We used 827 object classes and 749 attribute classes that appear more than 100 times. A total of 86,674 images with 667,882 object attribute labels are used for our experiment, and they are split into training, validation and test sets, which contain 43,337, 8,667 and 34,670 images, respectively. The task is challenging because appearances and semantics of objects largely vary.

Experimental settings and results All networks share VGG-16 network [18] pretrained on ImageNet [5] and is further fine-tuned for attribute prediction. For SAN and HAN, an attention layer is attached to the last pooling layer of VGG-16 while PAN[S]+CTX and PAN[H]+CTX stack an additional attention layer with local contexts $\mathcal{F}_{i,j}^l$ with $\delta = 2$ on top of the last three pooling layers in VGG-16. We skip the first two pooling layers (pool1 and pool2) for placing attention because the features in those layers are not discriminative enough to filter out. We also test the models with object class conditional prior. For the purpose, the final attended feature is fused with the query once more using a fully connected layer, which allows the network to reflect the conditional distribution of the attributes given the query. We train the models by Adam with different module-wise initial learning rates and the initial learning rates are exponentially decayed by the factor of 0.9. Refer to the supplementary document and the code for more detailed descriptions on the network architectures and experimental settings, respectively.

All models are evaluated via mean average precision (mAP) weighted by the frequencies of the attribute labels in test set, where the computation of mAP follows PASCAL VOC protocol [1]. Our progressive attention process with local context and hard attention mechanism consistently achieves the best weighted mAP scores in both experimental settings as shown in Table 2(left). Table 2(left) also presents TPR of each model measured with ground-truth bounding boxes due to lack of the ground-truth segmentation labels for evaluating attention




Figure 7: Visualization of example attentions of models on VG dataset. Two variants of progressive attention models gradually attend to target objects (queried by *sign*) in fine resolution. For PAN[*]+CTX, we only show last two attentions, which accumulate attentions of earlier layers. More qualitative results are presented in supplementary document.

qualities. PAN[H]+CTX shows the best TPR although the computation of TPR with bounding boxes is more favorable to other methods.

To measure the attention quality with segmentation masks, we also measure TPRs of models on train/val set in PASCAL VOC 2007 [6] as shown in Table 2(right). We use class names as queries for images to obtain an attention map of each image and measure TPR with its corresponding object segmentation mask. We observe that our progressive attention models outperform the baselines and become stronger when incorporated with hard attention.

STNs have trouble to attend to target objects as in MREF and show very poor TPRs. Note that STNs show higher mAPs than the other baselines in ‘attention only’ setting. We believe that this is because STNs utilize object class conditional priors by encoding queries through a manipulation in the transformation process. Figure 7 presents qualitative results on VG dataset, which show that the progressive attention models gradually attend to target objects.

5 Conclusion

We proposed a novel attention network, which progressively attends to regions of interest through multiple layers in a CNN. As the model is recursively applied to multiple layers of a CNN with an inherent feature hierarchy, it accurately predicts regions of interest with variable sizes and shapes. We also incorporate local contexts into our attention network for more robust estimation. While progressive attention networks can be implemented with either soft or hard attention, we demonstrated that both versions of the model substantially outperform existing methods on both synthetic and real datasets.

Acknowledgement

This research is partly supported by Adobe Research and Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MIST) (No. 2017-0-01778, Development of Explainable Human-level Deep Machine Learning Inference Framework; No.2017-0-01780, The Technology Development for Event Recognition/Relational Reasoning and Learning Knowledge based System for Video Understanding).

References

- [1] Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Deep compositional question answering with neural module networks. In *CVPR*, 2016.

- [2] Jimmy Ba, Volodymyr Mnih, and Koray Kavukcuoglu. Multiple object recognition with visual attention. In *ICLR*, 2015.
- [3] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015.
- [4] Jan Chorowski, Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. End-to-end continuous speech recognition using attention-based recurrent nn: first results. *arXiv preprint arXiv:1412.1602*, 2014.
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [6] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>.
- [7] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [8] Alex Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.
- [9] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- [10] Karol Gregor, Ivo Danihelka, Alex Graves, and Daan Wierstra. Draw: A recurrent neural network for image generation. In *ICML*, pages 1462–1471, 2015.
- [11] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *NIPS*, pages 2008–2016, 2015.
- [12] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *arXiv preprint arXiv:1602.07332*, 2016.
- [13] Hugo Larochelle and Geoffrey E Hinton. Learning to combine foveal glimpses with a third-order boltzmann machine. In *NIPS*, pages 1243–1251, 2010.
- [14] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [15] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. Effective approaches to attention-based neural machine translation. In *EMNLP*, 2015.
- [16] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *NIPS*, pages 2204–2212, 2014.
- [17] Hyeonwoo Noh and Bohyung Han. Training recurrent answering units with joint loss minimization for vqa. *arXiv preprint arXiv:1606.03647*, 2016.

- [18] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2015.
- [19] Jason Weston, Sumit Chopra, and Antoine Bordes. Memory networks. In *ICLR*, 2015.
- [20] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [21] Jianxiong Xiao, Krista A Ehinger, James Hays, Antonio Torralba, and Aude Oliva. Sun database: Exploring a large collection of scene categories. *International Journal of Computer Vision*, pages 1–20, 2014.
- [22] Huijuan Xu and Kate Saenko. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. *arXiv preprint arXiv:1511.05234*, 2015.
- [23] Kelvin Xu, Jimmy Ba, Ryan Kiros, Aaron Courville, Ruslan Salakhutdinov, Richard Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.
- [24] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. *arXiv preprint arXiv:1511.02274*, 2015.
- [25] Dongfei Yu, Jianlong Fu, Tao Mei, and Yong Rui. Multi-level attention networks for visual question answering. In *CVPR*, pages 4709–4717, 2017.
- [26] Wojciech Zaremba and Ilya Sutskever. Reinforcement learning neural turing machines. *arXiv preprint arXiv:1505.00521*, 2015.
- [27] Heliang Zheng, Jianlong Fu, Tao Mei, and Jiebo Luo. Learning multi-attention convolutional neural network for fine-grained image recognition. In *CVPR*, pages 5209–5217, 2017.