

USER CONSTRAINED THUMBNAIL GENERATION USING ADAPTIVE CONVOLUTIONS

Perla Sai Raj Kishore¹, Ayan Kumar Bhunia^{2*}, Shuvojit Ghose¹, Partha Pratim Roy³

¹Institute of Engineering & Management, India ²Nanyang Technological University, Singapore

³Indian Institute of Technology Roorkee, India

²ayanbhunia@ntu.edu.sg

ABSTRACT

Thumbnails are widely used all over the world as a preview for digital images. In this work we propose a deep neural framework to generate thumbnails of any size and aspect ratio, even for unseen values during training, with high accuracy and precision. We use Global Context Aggregation (GCA) and a modified Region Proposal Network (RPN) with adaptive convolutions to generate thumbnails in real time. GCA is used to selectively attend and aggregate the global context information from the entire image while the RPN is used to generate candidate bounding boxes for the thumbnail image. Adaptive convolution eliminates the difficulty of generating thumbnails of various aspect ratios by using filter weights dynamically generated from the aspect ratio information. The experimental results indicate the superior performance of the proposed model¹ over existing state-of-the-art techniques.

Index Terms— Thumbnail generation, aspect ratio, Global Context Aggregation, Adaptive Convolution, Region Proposal Network.

1. INTRODUCTION

With an increasing number of digital images everyday, it is a need to display images in any given space efficiently for easy browsing. This is facilitated by the use of thumbnails which are a smaller version of the original images that effectively capture and represent the content of the original images. Thumbnails are widely used in social media platforms and applications such as *Image Gallery* where the images are shown within a fixed resolution display in an organized manner.

Earlier methods on thumbnail generation [1, 2, 3] employ a two-stage framework which first predict a salient map specifying the importance of various regions and objects of the input image and then use this for generating a crop for the thumbnail image through a region search algorithm. It is important for the generated crop to be of same aspect ratio as that of the target thumbnail, since any mismatch could lead to significant distortions when scaling down the crop to thumbnail



Fig. 1. An example of output generated by our model. The original image is shown on the left with the generated thumbnails of different sizes and aspect ratios on the right.

size. Although [3] could crop images with a limited set of aspect ratios, it was mentioned it could be infeasible for a given overall saliency threshold value. Other approaches like [4, 5] attempt to crop the most aesthetic part of the image. Huang et al. [6] were the first to address this problem with a direct solution. They used manually selected features and SVM to score a large set of candidate crops. The crop with largest score was finally selected as the thumbnail. But their method was slow and only considered thumbnails of fixed size. Esmaeili et al. [7], proposed a fast thumbnail generation algorithm inspired from an object detection framework (F-RCN [8]) to predict bounding box coordinates of the thumbnail image. They maintain a set of convolutional filter banks to generate thumbnails for a limited selection of aspect ratio values.

In this work, we propose a deep neural framework to generate thumbnails of any given size and aspect-ratio for an input image in real time. We use Global Context Aggregation (GCA) and a modified Region Proposal Network (RPN) [9] with adaptive convolutions to generate thumbnails of varying sizes with high accuracy and precision. GCA selectively attends and aggregates the contextual information from the entire image at each context location, increasing the receptive field of the model to the entire image and enhancing the feature representability of the CNNs. A RPN is then used to generate candidate bounding boxes for the thumbnail image. The RPN uses adaptive convolutions where the filter weights are dynamically generated from the aspect ratio of the thumb-

*Corresponding Author

¹The source code of the proposed system is publicly available at <https://github.com/sairajk/Thumbnail-Generation>

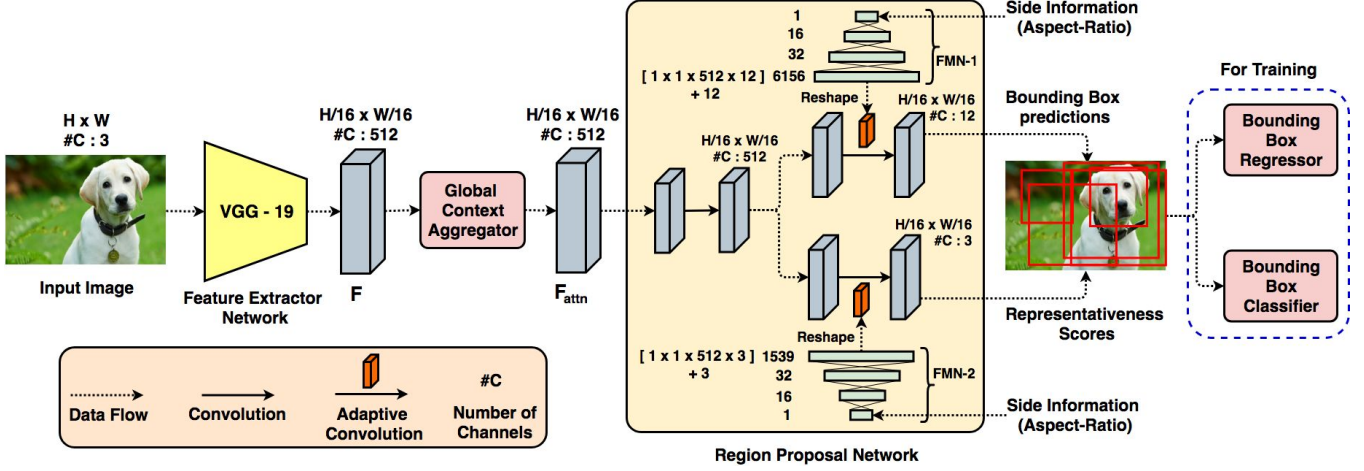


Fig. 2. Architecture of the proposed model. The input image is passed through a Feature Extraction Network to generate F . The global context is then selectively aggregated at each location to generate F_{attn} . A modified RPN is then used to generate candidate bounding boxes and representativeness scores. The network is finally trained using Regression and Classification losses.

nail image. The adaptive convolution layer provides a smooth manifold for the convolution filters to vary and is used to disentangle the variations in shape of the bounding box predictions with respect to aspect ratio. The smooth manifold for convolution filters also allows us to interpolate and use unseen values of aspect ratio for thumbnails during inference.

In summary, we make the following two contributions :

- (1) Our method can generate thumbnails of any size and aspect ratio efficiently, even for unseen values during training.
- (2) Our method is fast and can generate thumbnails in real time. It can process around 13 images per second on a GPU.

2. METHODOLOGY

2.1. Global Context Aggregation (GCA)

Given a convolutional feature map, $\mathbf{F} \in \mathbb{R}^{H \times W \times C}$, with height H , width W and number of channels C , GCA aims to generate an attention map at each location, over \mathbf{F} to selectively aggregate the global context information from the entire image. We employ a recurrent approach, inspired from [10] and [11], for this purpose. A pair of Bidirectional LSTMs is used to scan the given feature map in horizontal and vertical directions from the either ends to aggregate the global context at each spatial location $\langle h, w \rangle$. We then apply convolution to generate an intermediate output $Z \in \mathbb{R}^{H \times W \times D}$, where $D = H \times W$. Softmax operation is then applied to normalize each of these D -dimensional vectors of Z and can be represented as :

$$Z_{norm}^{(h,w,i)} = \frac{\exp(Z^{(h,w,i)})}{\sum_{i=1}^D \exp(Z^{(h,w,i)})} \quad (1)$$

where i is the index along the depth of the convolution output Z and Z_{norm} is the normalized feature map. The features

at all positions in \mathbf{F} are then weighted summed by $Z_{norm}^{(h,w)}$ to generate the attended contextual features $\mathbf{F}_{attn} \in \mathbb{R}^{H \times W \times C}$:

$$F_{attn}^{(h,w)} = \sum_{i=1}^D Z_{norm}^{(h,w,i)} f^i \quad (2)$$

where f^i is the C -dimensional feature vector at the i^{th} location (row/column-wise) of \mathbf{F} .

2.2. Adaptive Convolutional Layer

To address the difficulty of generating thumbnails of various aspect-ratios, we make use of adaptive convolutions [12] which model the filter weights as a low-dimensional manifold, parametrized by the side information, within the high-dimensional space of filter weights. In other words, the filter weights for adaptive convolution are generated dynamically based on some side information (aspect ratio, in this case). Hence, as the aspect ratio changes, the convolution filter weights also changes accordingly. This eliminates the need of aspect ratio-specific filter banks as used in [7], making the same convolutional layer adaptive to various values of aspect ratio for the thumbnail image. This allows us to disentangle aspect ratio-specific outputs and also use unseen but similar values of aspect ratio during inference.

Compared to a normal convolution layer where the filter weights remain constant after training, the filter weights for the adaptive convolution layer are dynamically generated by a set of fully connected layers, called the Filter Manifold Network (FMN), with progressively increasing number of neurons and side information as input. The convolution filter weights are generated by reshaping the output of the FMN into a $4D$ tensor for convolution kernel and a $1D$ vector for bias. The adaptive convolution layer is more formally defined

by, $F_{adap} = A(x * g(z; w))$, where x is an input image or feature maps from a convolution operation, z is the side information, $g(\cdot; w)$ is the filter manifold network with learnable weights w and $A(\cdot)$ is the activation function.

2.3. Thumbnail Generation

We employ a two-step process for generating the thumbnail, a bounding box with user defined aspect ratio is first predicted to crop the region of interest, the cropped image is then rescaled to generate the thumbnail of desired size.

The input image is first passed through a series of convolutional layers, which act as a feature extractor network. The extracted features, \mathbf{F} , are then passed through a GCA module to aggregate the global context information at each location. These attended contextual features, \mathbf{F}_{attn} , are then passed on to a Region Proposal Network (RPN) to predict candidate bounding boxes for the thumbnail image. The RPN generates two outputs, bounding box coordinates and the corresponding representativeness scores at each spatial location of the input feature map. The representativeness score depicts how well the corresponding bounding box represents the original image. The RPN consists of a convolution layer with 3×3 kernel and 512 channels as output, followed by two parallel convolutions using 1×1 kernel and whose number of output channels depend upon the number of bounding boxes (anchors), k , predicted at each location. One branch of the 1×1 convolution generates bounding box predictions with $4k$ channels as output ($\langle x_{center}, y_{center}, width, height \rangle$ for each bounding box), whereas the other branch generates the corresponding representativeness scores with k channels as output (probability of the corresponding bounding box representing the original image). We modify the RPN, proposed in [9], to generate the bounding box coordinates and the corresponding representativeness scores constrained by the user defined aspect ratio for the thumbnail image. For this we use adaptive convolutional layers where the filter weights are dynamically generated from the aspect ratio information. We make the 1×1 convolution layers of the RPN adaptive, as these are the layers which are finally responsible for generating the bounding boxes coordinates and representativeness scores for the target thumbnail image.

We use the attended contextual features, \mathbf{F}_{attn} , instead of directly using the features, \mathbf{F} , from the feature extractor network for bounding box predictions. This is vital for a task like thumbnail generation where the entire global context is important for generating the thumbnail, as opposed to object detection tasks where only the local neighborhood context is important. GCA allows the RPN to *look* at the entire image in a selective manner before making the bounding box predictions at each location and to assess how well the predicted bounding boxes represent the original image (through representativeness scores). GCA also allows us to use a lighter feature extractor network (in contrast to ResNet-101 [13] used

in [7]) by increasing the receptive field, filtering information (by attention) and enhancing the feature representability of the CNNs. As a result, GCA plays a crucial role in reducing the memory footprint while improving the overall accuracy and speed of the system.

We predict a total of 3 bounding boxes (i.e. $k = 3$), with same aspect ratio but different scales (with box areas of 60^2 , 120^2 and 180^2 pixels), at each context location of \mathbf{F}_{attn} . The aspect ratio is determined by the convolution kernel generated dynamically from the user input. We employ a similar approach as used in [9] for training the network. We use regression loss for bounding box predictions and classification loss for the probability scores. The candidate bounding boxes are first classified as ‘positive’ or ‘negative’ depending upon their Intersection-over-Union (IoU) scores with the ground truth bounding box according to the thresholds described in [9]. We then randomly sample 256 candidate bounding boxes to form a balanced mini-batch of positive and negative boxes. It is used to calculate the classification loss (L_{cls}) using binary cross entropy. We then only use the positive bounding boxes to calculate the regression loss (L_{reg}). For L_{reg} , we use smooth L_1 loss defined in [14]. The overall loss function is given by,

$$L(p, b) = \frac{1}{N_{cls}} \sum_i L_{cls}(p_i, \hat{p}_i) + \lambda \frac{1}{N_{reg}} \sum_i \hat{p}_i L_{reg}(b_i, \hat{b}_i) \quad (3)$$

where i is the index of the bounding boxes in the mini-batch and p_i is the predicted probability of bounding box being representative of the original image. The ground-truth label \hat{p}_i is 1 if the bounding box is positive, and is 0 if the bounding is negative. The predicted bounding box coordinates are represented by b_i , and that of the associated ground truth box by \hat{b}_i . The two losses are normalized by N_{cls} and N_{reg} and are weighted by λ . During inference, the bounding box with highest representativeness score is selected and rescaled to generate the target thumbnail.

3. EXPERIMENTS

Implementation and Training details: For our experiments we use the dataset released in [7]. The dataset contains a total of 70,048 thumbnail annotations over 28,064 images taken from the photo-quality dataset [15]. The dataset is split into 63043 annotations (from 24,154 images) for training and 7,005 annotations (from 3,910 images) for testing. All images are resized to have a maximum height of 650 and a maximum width of 800 pixels. The size of the annotated thumbnails vary from 32 to 200 pixels and the aspect ratio varies from 0.5 to 2.

The entire network is trained in an end-to-end manner supervised by the loss function stated in Equation 3 with $\lambda = 10$. We use VGG-19 [16] as the feature extractor network with weights initialized by pretraining the model for Ima-

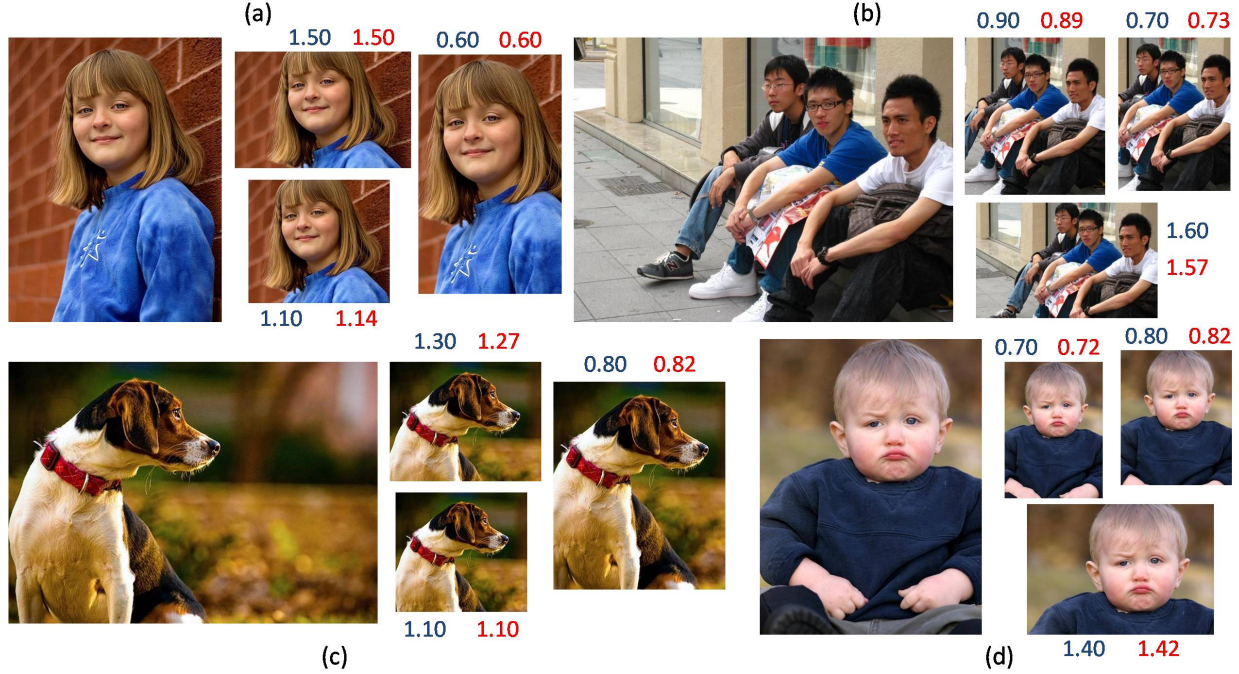


Fig. 3. Thumbnails generated by our model on the test set. The original image is shown on the left with the generated thumbnails on the right. The query aspect ratio is given in blue and the aspect ratio of the generated thumbnail is given in red.

geNet classification [17]. The features, \mathbf{F} , are extracted from the `conv5_4` layer of the VGG-19 network and all the layers upto `conv5_4` are fine-tuned during training. All other weights are initialized with values sampled from zero-mean Gaussian distribution with standard deviation 0.05. We use Adam Optimizer [18] for training with learning rate and momentum set to 0.001 and 0.9 respectively. The model is implemented using the TensorFlow library and all experiments are conducted on a single NVIDIA Tesla K80 GPU.

Baselines and Evaluation Metrics: Our model is compared against the following 2 state-of-the-art image cropping methods: Fast-AT [7] and Aesthetic Image Cropping (AIC) [4]. Fast-AT maintains a set of aspect ratio-specific convolutional filters to generate thumbnails of different aspect ratios. Whereas [4] crops the most aesthetic part of the image irrespective of the target image aspect ratio. We use the following evaluation metrics defined in [6] and [7] to compare our model to others: Center offset (CO), Rescaling factor (RF), Intersection over Union (IoU), Aspect ratio mismatch (ARM), Hit Ratio (h_r) and Background Ratio (b_r).

Performance Analysis: A quantitative comparison of the proposed model to the other methods for various evaluation metrics is shown in Table 1 and some qualitative results generated by our model are shown in Figure 1 and Figure 3. It is evident from the results in Table 1 that our model achieves the best results for all evaluation metrics and outperforms the existing state-of-the-art methods for thumbnail generation. We observe a significant drop in ARM score which is attributed

Model	CO	RF	IoU	ARM	h_r	b_r
AIC [4]	89.4	1.402	0.54	0.213	65.6%	43.1%
Fast-AT [7]	55.0	1.148	0.68	0.010	83.7%	37.1%
Ours	51.8	1.032	0.77	0.0013	87.6%	30.3%

Table 1. Evaluation Metrics evaluated on different thumbnail generation methods.

to adaptive convolutions, which help us generate bounding boxes with the given aspect ratio with high precision. Fast-AT [7] works moderately well due discretized aspect ratio-specific filters which approximate the target aspect ratio. Aesthetic Image Cropping [4] performs poorly as it aims to crop the most aesthetic region of the image irrespective of the target aspect ratio.

4. CONCLUSION

In this work, we proposed a deep neural framework to generate thumbnails of variable aspect ratios, even for unseen values during training, with high accuracy and precision. We used Global Context Aggregation and a modified Region Proposal Network with adaptive convolutions to generate state-of-the-art results for automatic thumbnail generation. The proposed model is fast, efficient and can be extended to other similar photo cropping tasks easily.

5. REFERENCES

- [1] Bongwon Suh, Haibin Ling, Benjamin B Bederson, and David W Jacobs, “Automatic thumbnail cropping and its effectiveness,” in *UIST*, 2003, pp. 95–104.
- [2] Gianluigi Ciocca, Claudio Cusano, Francesca Gasparini, and Raimondo Schettini, “Self-adaptive image cropping for small displays,” *IEEE Transactions on Consumer Electronics*, vol. 53, no. 4, 2007.
- [3] Jiansheng Chen, Gaocheng Bai, Shaoheng Liang, and Zhengqin Li, “Automatic image cropping: A computational complexity study,” in *CVPR*, 2016, pp. 507–515.
- [4] Wenguan Wang, Jianbing Shen, and Haibin Ling, “A deep network solution for attention and aesthetics aware photo cropping,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [5] Debang Li, Huikai Wu, Junge Zhang, and Kaiqi Huang, “A2-rl: Aesthetics aware reinforcement learning for image cropping,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8193–8201.
- [6] Jingwei Huang, Huarong Chen, Bin Wang, and Stephen Lin, “Automatic thumbnail generation based on visual representativeness and foreground recognizability,” in *ICCV*, 2015, pp. 253–261.
- [7] Seyed A Esmaeili, Bharat Singh, and Larry S Davis, “Fast-at: Fast automatic thumbnail generation using deep neural networks,” in *CVPR*, 2017, pp. 4178–4186.
- [8] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun, “R-fcn: Object detection via region-based fully convolutional networks,” in *Advances in Neural Information Processing Systems*, 2016, pp. 379–387.
- [9] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” in *Advances in neural information processing systems*, 2015, pp. 91–99.
- [10] Nian Liu, Junwei Han, and Ming-Hsuan Yang, “Picanet: Learning pixel-wise contextual attention for saliency detection,” in *CVPR*, 2018, pp. 3089–3098.
- [11] Francesco Visin, Kyle Kastner, Kyunghyun Cho, Matteo Matteucci, Aaron Courville, and Yoshua Bengio, “Renet: A recurrent neural network based alternative to convolutional networks,” *arXiv preprint arXiv:1505.00393*, 2015.
- [12] Di Kang, Debarun Dhar, and Antoni B Chan, “Crowd counting by adapting convolutional neural networks with side information,” *arXiv preprint arXiv:1611.06748*, 2016.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” *CVPR*, pp. 770–778, 2016.
- [14] Ross Girshick, “Fast r-cnn,” in *ICCV*, 2015, pp. 1440–1448.
- [15] Wei Luo, Xiaogang Wang, and Xiaoou Tang, “Content-based photo quality assessment,” in *ICCV*, 2011, pp. 2206–2213.
- [16] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [17] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al., “Imagenet large scale visual recognition challenge,” *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [18] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.