



AIM 2020 Challenge on Rendering Realistic Bokeh

Andrey Ignatov¹(✉), Radu Timofte¹, Ming Qian², Congyu Qiao², Jiamin Lin²,
Zhenyu Guo², Chenghua Li², Cong Leng², Jian Cheng², Juewen Peng³,
Xianrui Luo³, Ke Xian³, Zijin Wu³, Zhiguo Cao³, Densen Puthussery⁴,
C. V. Jiji⁴, P. S. Hrishikesh⁴, Melvin Kuriakose⁴, Saikat Dutta^{5,6,7},
Sourya Dipta Das^{5,6,7}, Nisarg A. Shah^{5,6,7}, Kuldeep Purohit⁵,
Praveen Kandula⁵, Maitreya Suin⁵, A. N. Rajagopalan⁵, M. B. Saagara⁴,
A. L. Minnu⁴, A. R. Sanjana⁴, S. Praseeda⁴, Ge Wu⁸, Xueqin Chen⁸,
Tengyao Wang⁸, Max Zheng⁸, Hulk Wong⁸, and Jay Zou⁸

¹ Computer Vision Lab, ETH Zurich, Zürich, Switzerland
{andrey,radu.timofte}@vision.ee.ethz.ch

² Nanjing Artificial Intelligence Chip Research, Institute of Automation,
Chinese Academy of Sciences (AiRiA), MAICRO, Beijing, China
20181223053@nuist.edu.cn

³ Huazhong University of Science and Technology, Wuhan, China
im.pengjw@gmail.com

⁴ College of Engineering Trivandrum, Trivandrum, India
puthusserydensen@gmail.com, {hrishikeshps,saagara}@cet.ac.in,
ar.sanjanar@gmail.com

⁵ Indian Institute of Technology Madras, Chennai, India
cs18s016@smail.iitm.ac.in, kuldeepurohit3@gmail.com

⁶ Jadavpur University, Kolkata, India

⁷ Indian Institute of Technology Jodhpur, Jodhpur, India
⁸ Beijing, China

1047670389@qq.com, 1843639867@qq.com

Abstract. This paper reviews the second AIM realistic bokeh effect rendering challenge and provides the description of the proposed solutions and results. The participating teams were solving a real-world bokeh simulation problem, where the goal was to learn a realistic shallow focus technique using a large-scale EBB! bokeh dataset consisting of 5K shallow/wide depth-of-field image pairs captured using the Canon 7D DSLR camera. The participants had to render bokeh effect based on only one single frame without any additional data from other cameras or sensors. The target metric used in this challenge combined the runtime and the perceptual quality of the solutions measured in the user study. To ensure the efficiency of the submitted models, we measured their runtime on standard desktop CPUs as well as were running the models on smartphone GPUs. The proposed solutions significantly improved the baseline results, defining the state-of-the-art for practical bokeh effect rendering problem.

A. Ignatov and R. Timofte are the challenge organizers, while the other authors participated in the challenge.

The Appendix A contains the authors' teams and affiliations.

AIM 2020 webpage: <https://data.vision.ee.ethz.ch/cvl/aim20/>

© Springer Nature Switzerland AG 2020

A. Bartoli and A. Fusiello (Eds.): ECCV 2020 Workshops, LNCS 12537, pp. 213–228, 2020.
https://doi.org/10.1007/978-3-030-67070-2_13

1 Introduction

The advances in image manipulation tasks are impressive. In particular, the image manipulation related to portable devices such as smartphone cameras has recently faced an interest boost from the research community to match the users' demands. Multiple novel solutions were proposed in the literature for various tasks, such as image quality enhancement [1, 12, 17, 35], style transfer [8, 22, 27], learning of an image signal processor (ISP) [20], photo segmentation and blurring [2, 4, 33, 36], etc. Moreover, modern mobile devices got powerful GPUs and NPUs that are well suitable for running the proposed deep learning models [15, 16].

Rendering an automatic bokeh effect has been one of the most popular topics over past few years, with many solutions that are now included within the majority of smartphone camera applications. In 2014, a seminal work on portrait segmentation [9] was published, and substantial improvements in segmentation accuracy were reported in many subsequent papers [33, 38]. Wadhwa *et al.* [36] provided a detailed description of the synthetic depth-of-field rendering method found in the Google Pixel phones and inspired further development in this field.

The AIM 2020 challenge on rendering realistic bokeh builds upon the success of the previous AIM 2019 challenge [14], and advances the benchmarking of example-based single image bokeh effect rendering by introducing two tracks with evaluation on several recent-generation desktop CPUs and smartphone GPUs. The AIM 2020 challenge uses the large-scale EBB! [13] dataset consisting of photo pairs with shallow and wide depth-of-field captured using the Canon 70D DSLR camera. Quantitative and qualitative visual results as well as the inference time and efficiency are used for ranking the proposed solutions. The challenge, the corresponding dataset, the results and the proposed methods are described and discussed in the next sections.

This challenge is one of the AIM 2020 associated challenges on: scene relighting and illumination estimation [6], image extreme inpainting [29], learned image signal processing pipeline [18], rendering realistic bokeh [19], real image super-resolution [37], efficient super-resolution [40], video temporal super-resolution [34] and video extreme super-resolution [7].

2 AIM 2020 Challenge on Realistic Bokeh

The objectives of the AIM 2020 challenge on rendering realistic bokeh effect is to promote realistic settings as defined by the *EBB!* Bokeh dataset, to push the state-of-the-art in synthetic shallow depth-of-field rendering, and to ensure that the final solutions are efficient enough to run both on desktop and mobile hardware.




Fig. 1. Sample wide and shallow depth-of-field image pairs from the EBB! dataset.

2.1 *Everything is Better with Bokeh!* Dataset

One of the biggest challenges in the bokeh rendering task is to get high-quality real data that can be used for training deep models. To tackle this problem, we used a large-scale *Everything is Better with Bokeh!* (EBB!) dataset presented in [13] that is containing more than 10 thousand images collected in the wild during several months. By controlling the aperture size of the lens, images with shallow and wide depth-of-field were taken. In each photo pair, the first image was captured with a narrow aperture ($f/16$) that results in a normal sharp photo, whereas the second one was shot using the highest aperture ($f/1.8$) leading to a strong bokeh effect. The photos were taken during the daytime in a wide variety of places and in various illumination and weather conditions. The photos were captured in automatic mode, the default settings were used throughout the entire collection procedure. An example set of collected images is presented in Fig. 1.

The captured image pairs are not aligned exactly, therefore they were first matched using SIFT keypoints and RANSAC method same as in [12]. The resulting images were then cropped to their intersection part and downscaled so that their final height is equal to 1024 pixels. From the resulting 10 thousand images, 200 image pairs were reserved for testing, while the other 4.8 thousand photo pairs can be used for training and validation.

2.2 Tracks and Competitions

The challenge consists of the following phases:

- i *development*: the participants get access to the data;
- ii *validation*: the participants have the opportunity to validate their solutions on the server and compare the results on the validation leaderboard;
- iii *test*: the participants submit their final results, models, and factsheets.

All submitted solutions were evaluated based on the following measures:

- PSNR measuring fidelity score,
- SSIM, a proxy for perceptual score,
- The runtime of the submitted models on desktop CPUs and mobile GPUs,
- MOS scores measured in the user study for explicit image quality assessment.

The AIM 2020 challenge on realistic bokeh consists of two tracks. In the first “CPU” track, the target was to produce a model which runtime is optimized for standard desktop CPUs. In the second, “Smartphone GPU” track, the goal was to develop a TensorFlow Lite [25] compatible solution that was tested on several mobile GPUs using a publicly available¹ *AI Benchmark* application [16] and an OpenCL-based TFLite GPU delegate [5]. During the development and validation phases, the quantitative performance of the solutions was measured by PSNR and SSIM metric. Since SSIM and PSNR scores are not reflecting many aspects of real quality of the resulted images, during the final test phase we evaluated the solutions based on their Mean Opinion Scores (MOS). For this, we conducted a user study evaluating the visual results of all proposed methods. The users were asked to rate the quality of each submitted solution by selecting one of the five quality levels (5 - comparable perceptual quality, 4 - slightly worse, 3 - notably worse, 2 - poor perceptual quality, 1 - completely corrupted image) for each method result in comparison with the original Canon images exhibiting bokeh effect. The expressed preferences were averaged per each test image and then per each method to obtain the final MOS.

3 Challenge Results

The Track 1 of the challenge attracted more than 110 registered participants and the Track 2 more than 80. However, only 9 teams provided results in the final phase together with factsheets and codes for reproducibility. Tables 1 and 2 summarize the final challenge results in terms of PSNR, SSIM and MOS scores for each submitted solution in addition to self-reported hardware/software configurations and runtimes. Short descriptions of the proposed solutions are provided in Sect. 4, and the team details (contact email, members and affiliations) are listed in Appendix A.

¹ <http://ai-benchmark.com>.

Table 1. AIM 2020 realistic bokeh rendering challenge, CPU Track: results and final rankings. The results are sorted based on the MOS scores. * - These teams submitted solutions that are using pre-computed depth maps and therefore were excluded from the final evaluation phase.

Team	Author	Factsheet Info		Track 1: Desktop CPU			
		Framework	Training Hardware, GPU	Avg. Runtime, s	PSNR↑	SSIM↑	MOS↑
Airia-bokeh	MingQian	TensorFlow	Nvidia TITAN RTX	5.52	23.58	0.8770	4.2
AIA-Smart	JuewenPeng	PyTorch	GeForce GTX 1080	1.71	23.56	0.8829	3.8
CET_SP	memelvin99	TensorFlow	Nvidia Tesla P100	1.17	21.91	0.8201	3.3
CET_CVLab	Densen	TensorFlow	Nvidia Tesla P100	1.17	23.05	0.8591	3.2
Team Horizon	tensorcat	PyTorch	GeForce GTX 1080 Ti	19.27	23.27	0.8818	3.2
IPCVLITM	ms_ipcv	PyTorch	NVIDIA Titan X	27.24	23.77	0.8866	2.5
CET21_CV	SaagaraMB	TensorFlow	Nvidia Tesla P100	0.74	22.80	0.8628	1.3
CET_ECE	Sanjana.A.R	TensorFlow	Nvidia Tesla P100	0.74	22.85	0.8629	1.2
xuehuapiaopiao-team	xuehuapiaopiao	TensorFlow	GeForce GTX 1080 Ti	-	22.98	0.8758	- *
Terminator	Max.zheng	TensorFlow	GeForce GTX 1080 Ti	-	23.04	0.8756	- *

Table 2. AIM 2020 realistic bokeh rendering challenge, GPU Track: results and final rankings. The results are sorted based on the MOS scores. The model submitted by the Team Horizon was unable to run on mobile GPUs due to NCHW channel order that is currently not supported by the TensorFlow Lite GPU delegate.

Team	Author	Factsheet Info		Track 2: Smartphone GPU			
		Framework	Training Hardware, GPU	Avg. Runtime, s	PSNR↑	SSIM↑	MOS↑
Airia-bokeh	MingQian	TensorFlow	Nvidia TITAN RTX	1.52	23.58	0.8770	4.2
AIA-Smart	JuewenPeng	PyTorch	GeForce GTX 1080	15.2	22.94	0.8842	4.0
CET_CVLab	Densen	TensorFlow	Nvidia Tesla P100	2.75	23.05	0.8591	3.2
Team Horizon	tensorcat	PyTorch	GeForce GTX 1080 Ti	- *	23.27	0.8818	3.2

3.1 Architectures and Main Ideas

All the proposed methods are relying on end-to-end deep learning-based solutions. Almost all submitted models have a multi-scale encoder-decoder architecture and are processing the images at several scales. This allows to achieve a significantly faster runtime as all heavy image processing is done on images of low resolution, as well as adds the possibility of introducing heavy global image manipulations. The majority of teams were using the $L1$, SSIM/MS-SSIM, VGG-based, Sobel and Charbonnier loss functions, while team Airia-bokeh demonstrated that a proper adversarial loss can significantly boost the quality of the resulting bokeh effect. Almost all teams were using the Adam optimizer [23] to train deep learning models and TensorFlow or PyTorch frameworks to implement and train the networks.

3.2 Performance

Quality. Airia-bokeh is the winner of the AIM 2020 challenge on rendering realistic bokeh. Airia-bokeh ranks the best in perceptual quality in both track 1 and track 2 with the same solution (deep model). Only one team – AIA-Smart, – submitted different models/solutions for the two tracks of the challenge. Surprisingly, the solution submitted for evaluation on smartphone GPU (Track 2) obtained better SSIM and MOS results than the one for CPU (Track 1). They are coming second to Airia-bokeh in the MOS score while reporting the best SSIM (0.8842) in Track 2. As expected, the perceptual ranking according to the MOS does not strongly correlate with the fidelity measures such as PSNR and SSIM. In particular, IPCV_IITM team ranks first in terms of SSIM and PSNR but only sixth in terms of perceptual quality (Track 1). Interestingly, the CET_SP team has the lowest fidelity (PSNR) and SSIM results, though comes third in perceptual quality (MOS).

Runtime. The measured average runtimes of the proposed solutions on standard Nvidia GPU cards (CPU Track 1) vary from 0.7s to more than 27s per single image. The fastest solutions (~ 0.7 s) are also among the worst performing in perceptual ranking, while the top fidelity method proposed by IPCV_IITM requires 27 s, and the top perceptual methods of Airia-bokeh and of AIA-Smart require 5.52s and 1.71s, respectively. When it comes to the solutions proposed in Track 2 evaluated on smartphone GPUs, the best perceptual quality solution of Airia-bokeh is also with the lowest inference time, 1.52s. We conclude that the proposed solutions do not meet the requirements for real-time applications on the current generation of smartphones, thus all processing should be done in the background after the image is obtained/captured.

3.3 Discussion

With the AIM 2020 challenge, we went further compared to the previously held challenges and aimed at solutions meant to run efficiently on desktop and smartphone hardware. The challenge employed the EBB! [13], a novel large dataset containing paired and aligned low- and high-aperture photos captured with a high-end Canon 70D DSLR camera. Several of the proposed approaches produced results with good perceptual quality and runtime suitable for on-device image processing. These methods are gauging the state-of-the-art for the practical bokeh synthesis task learned from pairs of real exemplars.

4 Challenge Methods and Teams

This section describes solutions submitted by all teams participating in the final stage of the AIM 2020 realistic bokeh rendering challenge.

4.1 Airia-Bokeh

Team Airia-bokeh proposed a Bokeh-Glass Network (BG-Net) [31] model for rendering realistic bokeh that is illustrated in Fig. 2. The model consists of two stacked U-Net based networks that were first trained separately using a combination of the L_1 and SSIM losses (with weights 0.5 and 1, respectively). During the second stage, two PatchGAN [42] discriminators with different receptive fields were added to improve the quality of the produced images. The generator and the discriminator were trained together using the WGAN-GP algorithm with a batch size of 1. The authors have additionally enhanced the EBB! dataset by removing some image pairs that did not correspond in color or were not in focus.

4.2 AIA-Smart

The solution of AIA-Smart team is based on defocus map estimation [28]. The proposed architecture consists of 4 modules (Fig. 3): defocus estimation, radiance, rendering and upsampling modules. Defocus estimation module is used to predict a defocus map, which works as a guidance for defocus rendering. Radiance module calculates the weight map used for estimating the weight of each pixel in the rendering process. In the rendering module, the low resolution bokeh result can be obtained with the input of the radiance map, weight map and defocus map using the refocusing pipeline proposed in [3]. In the upsampling module, the low resolution bokeh result and the high resolution original image are combined to generate the final full-resolution bokeh image.




Fig. 2. Bokeh-Glass Network network (top) and PatchGAN-based discriminators (bottom) proposed by Airia-bokeh team.

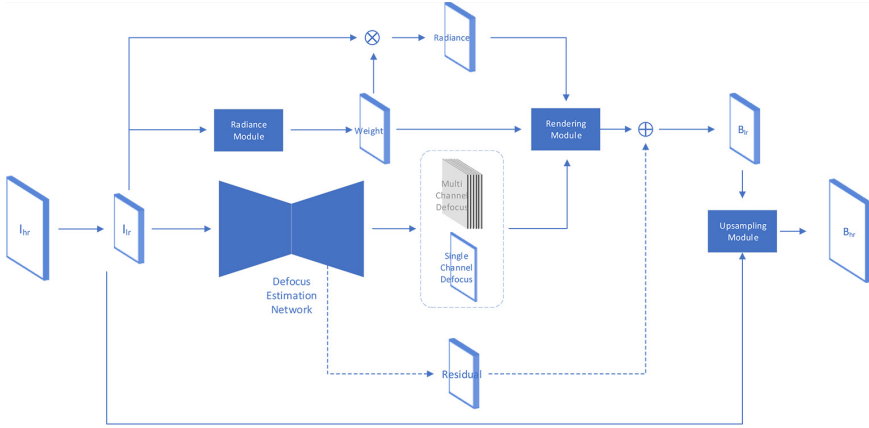


Fig. 3. AIA-Smart network consisting of the defocus estimation, radiance, rendering and upsampling modules.

The training of the network can be divided into 4 stages: 1) predicting the layered defocus maps and render the bokeh result on 1/4 of the original resolution, 2) rendering bokeh effect at 1/2 resolution while using the pretrained network from the first stage to refine the details around foreground boundaries, 3) replacing the multi-channel classification layer with a single-channel regression layer in the defocus estimation module to generate the pleasing ‘Circle of Confusion’, and 4) rendering the image at 1/2 resolution, upsampling the result by bilinear interpolation and calculating a soft foreground mask from the predicted single-channel defocus map. Finally, the foreground objects of the original image are covered on the rendering result to make the foreground more clear.

During the first stage, the model is trained with a combination of the L_1 , perceptual, SSIM and gradient loss functions using the images of resolution 256×256 pixels. The initial learning rate is set to $1e-4$ with a decay-cycle of 30 epochs. At the second and the third stages, the model is fine-tuned on 512×512 pixel images using the same set of loss functions.

4.3 CET_CVLab and CET_SP




Fig. 4. Dilated Wavelet CNN model used by CET_CVLab and CET_SP teams.

Both CET_CVLab and CET_SP teams used the same U-Net based Dilated Wavelet CNN model (Fig. 4) for generating bokeh images. In this network, the standard downsampling and upsampling operations are replaced with a discrete wavelet transform based (DWT) decomposition to minimize the information loss in these layers. The proposed methodology is computationally efficient and is based on the multi-level wavelet-CNN (MWCNN) proposed in [26].

CET_SP trained the model with a combination of the Charbonnier and perceptual VGG loss, while CET_CVLab additionally used Sobel and Grayscale (L_1 distance between the grayscale images) loss functions. Both models were optimized using the Adam algorithm with a batch size of 10 for 600 and 500 epoch, respectively.

4.4 Team Horizon

The authors proposed an encoder-decoder based model shown in Fig. 5 that is trained at several scales. At each level, the encoder-decoder module is producing the weight maps that are used together with the input image by the bokeh generation module to render the bokeh image. Generated weight maps and bokeh images are then upscaled and concatenated with the input image in the next level, while the upscaled encoded features are added to the corresponding encoded features used in the next level. The model is trained with a combination of the MS-SSIM and SSIM loss functions.

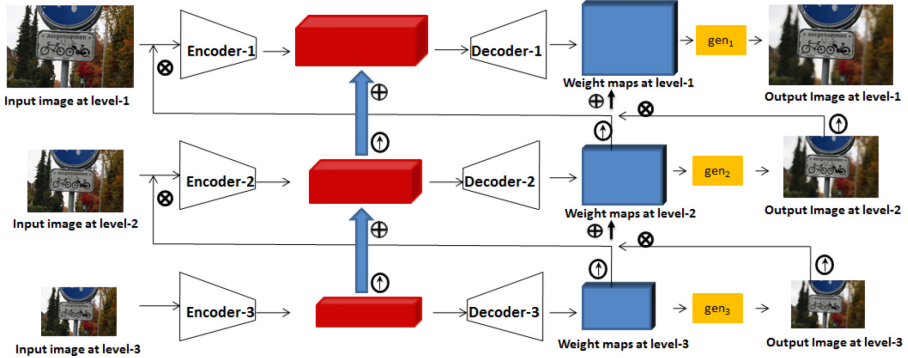


Fig. 5. A multiscale encoder-decoder based model proposed by Team Horizon.

4.5 IPCV_IITM

The authors proposed a depth-guided dynamic filtering dense network for rendering shallow depth-of-field (Fig. 6). At the onset, the network uses a space-to-depth module that divides each input channel into a number of blocks concatenated along the channel dimension. The output of this layer is concatenated with

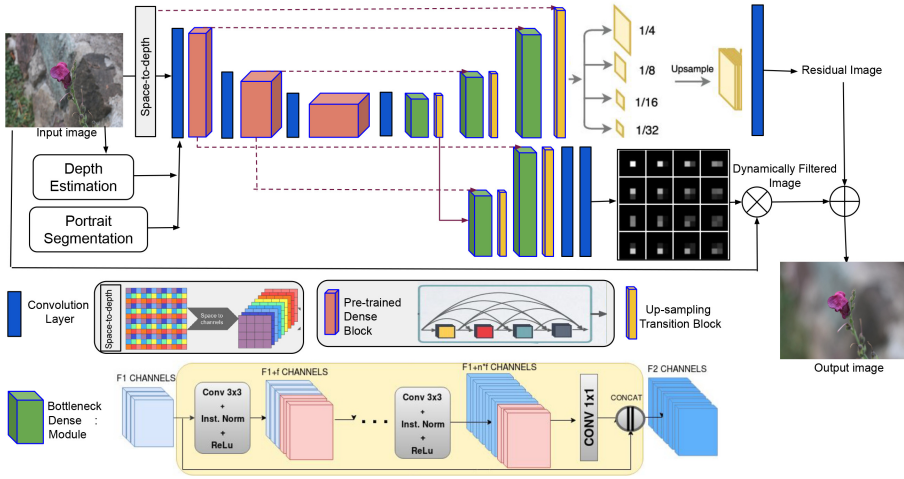


Fig. 6. Depth-guided Dynamic Filtering Dense Network proposed by IPCV_IITM.

the outputs of the pre-trained depth estimation [24] and salient object segmentation [10] networks to achieve more accurate rendering results. The resulting feature maps are passed to a U-net [32] based encoder consisting of densely connected modules. The first dense-block contains 12 densely-connected layers, the second block – 16, and the third one – 24 densely-connected layers. The weights of each block are initialized using the DenseNet-121 network [11] trained on the ImageNet dataset. The decoder has two difference branches which outputs are summed to produce the final result. The first branch has a U-net architecture with skip-connections and also consists of densely connected blocks. Its output is enhanced through multi-scale context aggregation through pooling and upsampling at 4 scales. The second branch uses the idea of dynamic filtering [21] and generates dynamic blurring filters conditioned on the encoded feature map. These filters are produced locally and on-the-fly depending on the input, the parameters of the filter-generating network are updated during the training. We refer to [30] for more details.

4.6 CET21_CV

CET21_CV proposed a modified U-Net model depicted in Fig. 7. Compared to the original U-Net implementation, the authors replaced the max-pooling down-sampling operation with a strided convolution layer, and the feature maps from shortcut connections are concatenated before applying the activation functions in the decoder module. *Leaky ReLU* activations are used in the convolutional layers, and the entire model is trained to minimize the mean absolute error loss using the Adam algorithm.

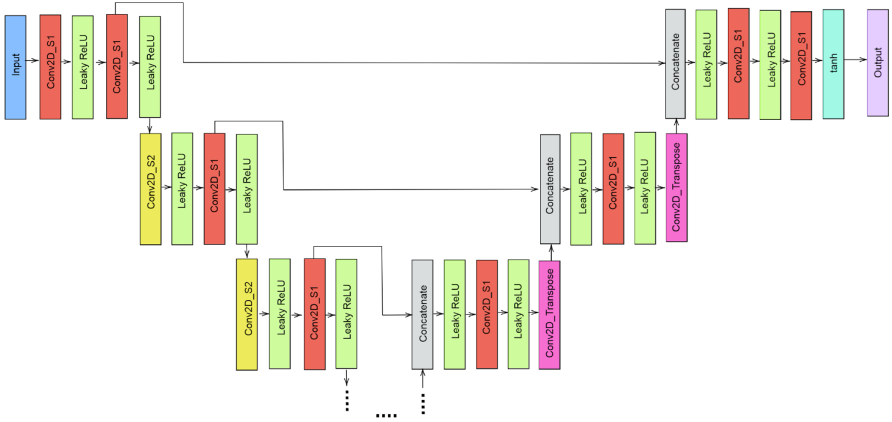


Fig. 7. A modified U-Net model used by CET21.CV team.

4.7 CET_ECE

The model architecture proposed by CET_ECE team was inspired from the wide activation [39] and channel attention [41] based networks. The proposed network (Fig. 8) is generally consisting of 2 block types: a feature extraction block and a series of wide activation residual blocks. To reduce the model complexity and information loss, a space-to-depth layer with a scale factor of 4 is used before the initial feature extraction block, and a depth-to-space operation is used as the last layer of the network. The Charbonnier loss function is used for training the network as it better captures the edge information compared to the MSE loss.

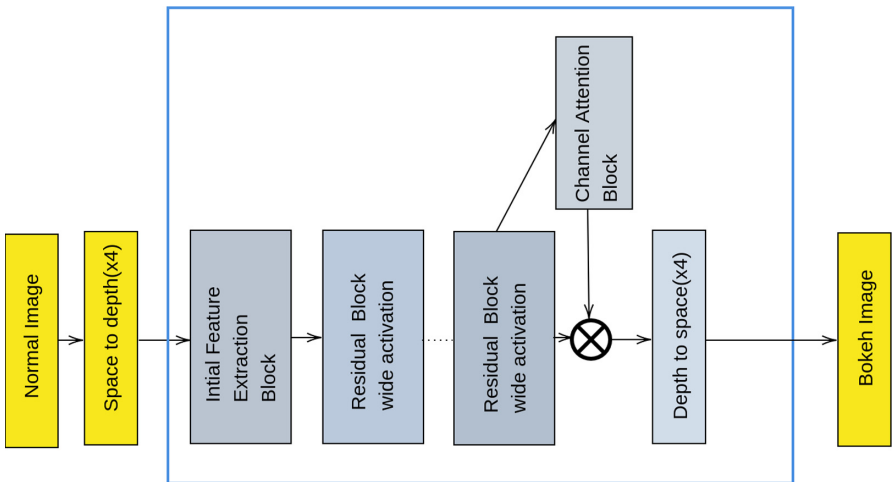


Fig. 8. CET_ECE's network architecture.

4.8 Xuehuapiaopiao-Team and Terminator

Both teams used a slightly modified PyNET [13] model for generating bokeh images. While the visual results of the proposed solutions were looking fine, they were relying on depth estimation modules that were not included in the submissions, and therefore were not ranked in the final phase of the challenge.

Acknowledgment. We thank the AIM 2020 sponsors: Huawei, MediaTek, Qualcomm, NVIDIA, Google and Computer Vision Lab/ETH Zürich.

A Appendix 1: Teams and affiliations

AIM 2020 Realistic Bokeh Rendering Challenge Team

Title: AIM 2020 Challenge on Rendering Realistic Bokeh

Members: Andrey Ignatov – andrey@vision.ee.ethz.ch,
Radu Timofte – radu.timofte@vision.ee.ethz.ch

Affiliations: Computer Vision Lab, ETH Zurich, Switzerland

Airia-bokeh

Title: BGNet: Bokeh-Glass Network for Rendering Realistic Bokeh

Members: Ming Qian – 20181223053@nuist.edu.cn,
Congyu Qiao, Jiamin Lin, Zhenyu Guo, Chenghua Li,
Cong Leng, Jian Cheng

Affiliations: Nanjing Artificial Intelligence Chip Research, Institute of Automation
Chinese Academy of Sciences (AiRiA), MAICRO, China

AIA-Smart

Title: Bokeh Rendering from Defocus Estimation [28]

Members: Juewen Peng – im.pengjw@gmail.com,
Xianrui Luo, Ke Xian, Zijin Wu, Zhiguo Cao

Affiliations: Huazhong University of Science and Technology, China

CET_CVLab

Title: Photorealistic Bokeh Effect Rendering with Dilated Wavelet CNN

Members: Densen Puthussery – puthusserydensen@gmail.com,
Jiji C V

Affiliations: College of Engineering Trivandrum, India

CET_SP**Title:** Bokeh Effect using VGG based Wavelet CNN**Members:** Hrishikesh P S – hrishikeshps@cet.ac.in,
Melvin Kuriakose**Affiliations:** College of Engineering Trivandrum, India**Team Horizon****Title:** Deep Multi-scale Hierarchical Network for Bokeh Effect Rendering**Members:** Saikat Dutta – cs18s016@smail.iitm.ac.in,
Sourya Dipta Das, Nisarg A. Shah**Affiliations:** Indian Institute of Technology Madras, India
Jadavpur University, India
Indian Institute of Technology Jodhpur, India**IPCV_IITM****Title:** Dense Dynamic Filtering Network for Rendering Synthetic Depth-of-Field Effect**Members:** Kuldeep Purohit – kuldeppurohit3@gmail.com,
Praveen Kandula, Maitreya Sun, A. N. Rajagopalan**Affiliations:** Indian Institute of Technology Madras, India**CET21_CV****Title:** Synthetic Bokeh Effect with Modified UNet**Members:** Saagara M B – saagara@cet.ac.in,
Minnu A L**Affiliations:** College of Engineering Trivandrum, India**CET_ECE****Title:** Bokeh Effect Rendering with Deep Convolutional Neural Network**Members:** Sanjana A R – ar.sanjanaar@gmail.com,
Praseeda S**Affiliations:** College of Engineering Trivandrum, India**Xuehuapiaopiao-team****Title:** Multi-scale Bokeh Rendering Network**Members:** Ge Wu – 1047670389@qq.com,
Xueqin Chen, Tengyao Wang**Affiliations:** None

Terminator

Title: Simulating Realistic Bokeh Rendering with an Improved Dataset and Robust Network

Members: Max Zheng – 1843639867@qq.com,
Hulk Wong, Jay Zou

Affiliations: None

References

1. Ancuti, C.O., Ancuti, C., Timofte, R., et al.: Ntire 2019 challenge on image dehazing: Methods and results. In: The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (2019)
2. Badrinarayanan, V., Kendall, A., Cipolla, R.: SegNet: a deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(12), 2481–2495 (2017)
3. Busam, B., Hog, M., McDonagh, S., Slabaugh, G.: SteReFo: efficient image refocusing with stereo vision. In: Proceedings of the IEEE International Conference on Computer Vision Workshops (2019)
4. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: DeepLab: Semantic image segmentation with deep convolutional nets, ATRous convolution, and fully connected CRFs. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**(4), 834–848 (2017)
5. delegate, T.L.G. <https://www.tensorflow.org/lite/performance/gpu>
6. El Helou, M., et al.: AIM 2020: scene relighting and illumination estimation challenge. In: Bartoli, A., Fusiello, A. (eds.) *ECCV 2020*, LNCS 12537, pp. 499–518. Springer, Cham (2020)
7. Fuoli, D., et al.: AIM 2020 challenge on video extreme super-resolution: Methods and results. In: Bartoli, A., Fusiello, A. (eds.) *ECCV 2020*, LNCS 12538, pp. 57–81. Springer, Cham (2020)
8. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2414–2423 (2016)
9. L.B. in the new Google Camera app. <https://ai.googleblog.com/2014/04/lens-blur-in-new-google-camera-app.html>
10. Hou, Q., Cheng, M.M., Hu, X., Borji, A., Tu, Z., Torr, P.H.: Deeply supervised salient object detection with short connections. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3203–3212 (2017)
11. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4700–4708 (2017)
12. Ignatov, A., Kobyshev, N., Timofte, R., Vanhoey, K., Van Gool, L.: DSLR-quality photos on mobile devices with deep convolutional networks. In: the IEEE International Conference on Computer Vision (ICCV) (2017)
13. Ignatov, A., Patel, J., Timofte, R.: Rendering natural camera bokeh effect with deep learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pp. 418–419 (2020)

14. Ignatov, A., et al.: AIM 2019 challenge on bokeh effect synthesis: methods and results. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), pp. 3591–3598. IEEE (2019)
15. Ignatov, A., et al.: AI Benchmark: running deep neural networks on android smartphones. In: Leal-Taixé, L., Roth, S. (eds.) ECCV 2018. LNCS, vol. 11133, pp. 288–314. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-11021-5_19
16. Ignatov, A., et al.: AI Benchmark: all about deep learning on smartphones in 2019. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), pp. 3617–3635. IEEE (2019)
17. Ignatov, A., Timofte, R., et al.: PIRM challenge on perceptual image enhancement on smartphones: report. In: European Conference on Computer Vision Workshops (2018)
18. Ignatov, A., Timofte, R., et al.: AIM 2020 challenge on learned image signal processing pipeline. In: Bartoli, A., Fusiello, A. (eds.) ECCV 2020, LNCS 12537, pp. 152–170. Springer, Cham (2020)
19. Ignatov, A., Timofte, R., et al.: AIM 2020 challenge on rendering realistic bokeh. In: Bartoli, A., Fusiello, A. (eds.) ECCV 2020, LNCS 12537, pp. 213–228. Springer, Cham (2020)
20. Ignatov, A., Van Gool, L., Timofte, R.: Replacing mobile camera ISP with a single deep learning model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pp. 536–537 (2020)
21. Jia, X., De Brabandere, B., Tuytelaars, T., Gool, L.V.: Dynamic filter networks. In: Advances in Neural Information Processing Systems, pp. 667–675 (2016)
22. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9906, pp. 694–711. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_43
23. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
24. Li, Z., Snavely, N.: Megadepth: learning single-view depth prediction from internet photos. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2041–2050 (2018)
25. Lite, T.: <https://www.tensorflow.org/lite>
26. Liu, P., Zhang, H., Zhang, K., Lin, L., Zuo, W.: Multi-level wavelet-CNN for image restoration. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 773–782 (2018)
27. Luan, F., Paris, S., Shechtman, E., Bala, K.: Deep photo style transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4990–4998 (2017)
28. Luo, X., Peng, J., Xian, K., Wu, Z., Cao, Z.: Bokeh rendering from defocus estimation. In: Bartoli, A., Fusiello, A. (eds.) ECCV 2020, LNCS 12537, pp. 245–261. Springer, Cham (2020)
29. Ntavelis, E., et al.: AIM 2020 challenge on image extreme inpainting. In: Bartoli, A., Fusiello, A. (eds.) ECCV 2020, LNCS 12537, pp. 716–741. Springer, Cham (2020)
30. Purohit, K., Suin, M., Kandula, P., Ambasamudram, R.: Depth-guided dense dynamic filtering network for bokeh effect rendering. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), pp. 3417–3426. IEEE (2019)

31. Qian, M., Qiao, C., Lin, J., Guo, Z., Li, C., Leng, C., Cheng, J.: Bggan: Bokeh-glass generative adversarial network for rendering realistic bokeh. In: Bartoli, A., Fusiello, A. (eds.) ECCV 2020, LNCS 12537, pp. 229–244. Springer, Cham (2020)
32. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
33. Shen, X., et al.: Automatic portrait segmentation for image stylization. In: Computer Graphics Forum, vol. 35, pp. 93–102. Wiley Online Library (2016)
34. Son, S., Lee, J., Nah, S., Timofte, R., Lee, K.M., et al.: AIM 2020 challenge on video temporal super-resolution. In: European Conference on Computer Vision Workshops (2020)
35. Timofte, R., Gu, S., Wu, J., Van Gool, L.: Ntire 2018 challenge on single image super-resolution: methods and results. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, June 2018
36. Wadhwa, N., Garg, R., Jacobs, D.E., Feldman, B.E., Kanazawa, N., Carroll, R., Movshovitz-Attias, Y., Barron, J.T., Pritch, Y., Levoy, M.: Synthetic depth-of-field with a single-camera mobile phone. *ACM Trans. Graphics (TOG)* **37**(4), 64 (2018)
37. Wei, P., Lu, H., Timofte, R., Lin, L., Zuo, W., et al.: AIM 2020 challenge on real image super-resolution: methods and results. In: Bartoli, A., Fusiello, A. (eds.) ECCV 2020, LNCS 12537, pp. 392–422. Springer, Cham (2020)
38. Xu, N., Price, B., Cohen, S., Huang, T.: Deep image matting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2970–2979 (2017)
39. Yu, J., et al.: Wide activation for efficient and accurate image super-resolution. arXiv preprint [arXiv:1808.08718](https://arxiv.org/abs/1808.08718) (2018)
40. Zhang, K., et al.: AIM 2020 challenge on efficient super-resolution: methods and results. In: Bartoli, A., Fusiello, A. (eds.) ECCV 2020, LNCS 12537, pp. 5–40. Springer, Cham (2020)
41. Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., Fu, Y.: Image super-resolution using very deep residual channel attention networks. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11211, pp. 294–310. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01234-2_18
42. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2223–2232 (2017)