AIM 2019 Challenge on RAW to RGB Mapping: Methods and Results

Andrey Ignatov Radu Timofte Sung-Jea Ko Seung-Wook Kim Kwang-Hyun Uhm
Seo-Won Ji Sung-Jin Cho Jun-Pyo Hong Kangfu Mei Juncheng Li
Jiajie Zhang Haoyu Wu Jie Li Rui Huang Muhammad Haris
Greg Shakhnarovich Norimichi Ukita Yuzhi Zhao Lai-Man Po Tiantian Zhang
Zongbang Liao Xiang Shi Yujia Zhang Weifeng Ou Pengfei Xian
Jingjing Xiong Chang Zhou Wing Yin Yu Yubin Bingxin Hou Bumjun Park
Songhyun Yu Sangmin Kim Jechang Jeong

Abstract

This paper reviews the first AIM challenge on mapping camera RAW to RGB images with the focus on proposed solutions and results. The participating teams were solving a real-world photo enhancement problem, where the goal was to map the original low-quality RAW images from the Huawei P20 device to the same photos captured with the Canon 5D DSLR camera. The considered problem embraced a number of computer vision subtasks, such as image demosaicing, denoising, gamma correction, image resolution and sharpness enhancement, etc. The target metric used in this challenge combined fidelity scores (PSNR and SSIM) with solutions' perceptual results measured in a user study. The proposed solutions significantly improved baseline results, defining the state-of-the-art for RAW to RGB image restoration.

1. Introduction

Over the past years, the topic of image restoration and enhancement has witnessed an increased interest from the vision and graphics communities, which resulted in many works targeting the improvement of different image quality aspects [26, 3, 1, 4, 7, 2, 21, 8], including its perceptual quality, resolution, color rendition, etc. One of the key real-world problems in this area is the restoration of the low-quality images obtained with compact camera sensors present in many portable mobile devices [11, 14, 15, 5]. The first works dealing with a comprehensive image enhancement [11, 12] appeared back in 2017, followed by many subsequent papers that have substantially improved the baseline results [22, 6, 28, 10, 18]. A further development in this field was facilitated by the PIRM challenge on perceptual image enhancement on smartphones [15] and the NTIRE 2019 challenge on image enhancement [13] that were working with a diverse DPED dataset [11] and produced a large number of efficient solutions.

The AIM 2019 RAW to RGB mapping challenge is a step forward in benchmarking example-based single image enhancement. It uses a large-scale Zurich RAW to RGB (ZRR) dataset consisting of RAW photos captured with the Huawei P20 mobile camera and the Canon 5D DSLR, and is taking into account both quantitative and qualitative visual results of the proposed solutions. In the next sections we describe the challenge and the corresponding dataset, present and discuss the results and describe the proposed methods.

A. Ignatov and R. Timofte ({andrey,radu.timofte}@vision.ee.ethz.ch, ETH Zurich) are the challenge organizers, while the other authors participated in the challenge.
The Appendix A contains the authors' teams and affiliations.
AIM 2019 webpage:
http://vision.ee.ethz.ch/aim19

2. AIM 2019 RAW to RGB Mapping Challenge

The objectives of the AIM 2019 Challenge on mapping RAW to RGB images are to gauge and push the state-of-the-art in image enhancement, to compare different approaches and solutions, and to promote realistic image enhancement settings defined by the ZRR dataset described below.

2.1. Zurich RAW to RGB (ZRR) dataset

In order to tackle the problem of image translation from the original RAW photos captured by smartphone cameras to superior quality images achieved by a professional DSLR camera, a large-scale real-world dataset containing more than 20K images was collected. The dataset consists of photos taken in the wild synchronously by Canon 5D Mark IV DSLR camera and Huawei P20 phone capturing images in the RAW format. 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, and we used default settings for all cameras

2473-9944/19/$31.00 ©2019 IEEE
DOI 10.1109/ICCVW.2019.00443

Figure 1. Example images from the collected dataset. From left to right: original 4-channel RAW image (channels = [R, GR, B, GB]) saved in PNG format (transparency channel = GB), original visualized RAW image and Canon 5D Mark IV target image.

throughout the whole collection procedure. An example set of collected images can be seen in figure 1.

Since training deep learning models on the high-resolution images is infeasible, the patches of size 224×224 pixels were extracted from the P20-RAW / Canon image pairs preliminary aligned using SIFT keypoints and RANSAC method. Around 200 original images were reserved for testing, the rest of the photos were used for training and validation. This procedure resulted in 90 thousand training and 2.4 thousand test patches.

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 solutions submitted by the participants were evaluated based on three measures:

- PSNR measuring fidelity score,
- SSIM, a proxy for perceptual score,
- MOS (mean opinion scores) by a user study for explicit image quality assessment.

The AIM 2019 RAW to RGB mapping challenge consists of two tracks. In the first "Fidelity" track, the target is to obtain an output image with the highest pixel fidelity to the ground truth as measured by PSNR and SSIM metrics. Since SSIM and PSNR scores are not reflecting many aspects of real image quality, in the second, "Perceptual" track, we are evaluating the solutions based on their Mean Opinion Scores (MOS). For this, we conduct a user study involving several hundreds of participants (using Amazon's MTurk platform) evaluating the visual results of all proposed methods. The users were asked to rate the quality of each submitted solution (based on 10 full resolution enhanced test images) by selecting one of five quality levels (0 - almost identical, 1 - mostly similar, 2 - similar, 3 - somewhat similar, 4 - mostly different) for each method result in comparison with the original Canon images. The expressed preferences were averaged per each test image and then per each method to obtain the final MOS.

3. Challenge Results

From above 70 registered participants, 7 teams entered the final phase and submitted results, codes / executables, and factsheets. Table 1 summarizes the final challenge results and reports PSNR, SSIM and MOS scores for each submitted solution, as well as the self-reported runtimes and hardware / software configurations. The methods are briefly described in section 4, and the team members and affiliations are listed in Appendix A.

3.1. Architectures and Main Ideas

All the proposed methods are relying on end-to-end deep learning-based solutions. The best results are achieved by models that are combining the idea of multi-scale image processing with the residual learning, usually by using a modified U-net architecture [23] or Back-Projection Networks [9]. The most popular loss functions are L_1 (often leading to the highest fidelity scores) and VGG [25]-based (for better perceptual quality) losses, though several works also considered adversarial, color, MSE and style losses. Almost all teams are using Adam optimizer [17] to train deep learning models, and in all cases PyTorch was used to implement and train the models.

https://www.mturk.com/

Team	Author	Framework	Factsheet Info Hardware, GPU	Runtime, ms	Track 1: Fidelity PSNR↑	SSIM↑	Track 2: Perceptual PSNR↑	SSIM↑	MOS↓
CVIP	kh3005	PyTorch	NVIDIA Titan Xp	36	**22.59**	**0.81**	**22.21**	**0.81**	1.28
The First Team of Hogwarts School	MKFMIKU	PyTorch	NVIDIA Titan X	700	22.24	0.80	22.20	0.80	**1.24**
TTI	iim_lab	PyTorch	NVIDIA Titan X	20	21.94	0.79	21.83	0.77	1.46
CityU TI Lab	zyz987	PyTorch	GeForce GTX 1080 Ti	16	21.91	0.79	-	-	1.56
superTeam	wyb123568	PyTorch	n.a.	20	19.46	0.75	19.46	0.75	1.92
houbingxin	houbingxin	PyTorch	NVIDIA Titan	8	-	-	-	-	2.16
Team Eraser	sangmin_kim	PyTorch	n.a.	1280	16.86	0.53	16.86	0.53	n.a.*

Table 1. AIM 2019 RAW to RGB mapping challenge results and final rankings. * – The solution from "Team Eraser" was not able to process full-resolution images.

3.2. Performance

The runtime varies considerably among the proposed solutions. According to the self-reported runtime, the fastest approaches process an image in less than 36ms, while the slowest require around 1s on a common Nvidia GPU card. The fastest approach of houbingxin (8ms) is also the worst in terms of perceptual quality with a MOS score of 2.16. Also, the best MOS score is achieved by The First Team of Hogwarts School with one of the slowest approaches – 750ms. CVIP is only slightly below The First Team of Hogwarts School in terms of MOS. However, CVIP provides the best balance between runtime efficiency and overall fidelity and perceptual quality. CVIP provides the best fidelity scores in this challenge, for both fidelity and perceptual tracks. CVIP and The First Team of Hogwarts School are the winners of the AIM 2019 RAW to RGB mapping challenge. We note that Team Eraser which achieves the lowest PSNR and SSIM scores was not able to process full-resolution images and, consequently, was not included in the user study for perceptual ranking.

3.3. Discussion

The AIM 2019 RAW to RGB mapping challenge promoted realistic settings for the image enhancement task — instead of using synthetic datasets capturing only a very limited amount of image quality aspects, the participants were proposed to improve the quality of real unprocessed RAW photos captured with a standard mobile camera, and were provided with the ZRR dataset containing paired and aligned photos captured with the Huawei P20 phone and Canon 5D Mark IV, a high-end DSLR camera. A diversity of proposed approaches demonstrated significantly improved visual results compared to the original RAW images. We conclude that the challenge through the proposed solutions defined the state-of-the-art for the practical RAW to RGB image mapping task.

4. Challenge Methods and Teams

This section describes solutions submitted by all teams participating in the final stage of the AIM 2019 RAW to RGB Challenge.

4.1. CVIP

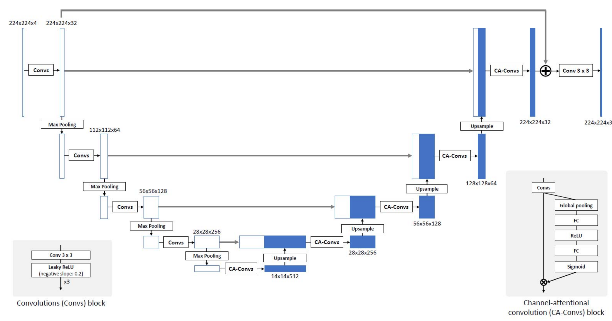

Figure 2. The U-Net architecture used by CVIP team.

The authors proposed a two-stage model consisting of two stacked U-Nets shown in Fig. 2. At the beginning, the first one is trained for 100 epochs and then is frozen, and the second U-Net is subsequently trained for 25 epochs with a learning rate of 10^{-4}. Finally, both networks are trained together for 1 epoch with a learning rate of 10^{-5}. The model was trained with a combination of the VGG-based (perceptual) [16], color (cosine distance between the RGB color vectors of the predicted and target labels) and L_1 losses. Additionally, an ensemble of three models which results were averaged was used in the fidelity track. We refer to [27] for more details on the solution proposed by CVIP team.

4.2. The First Team of Hogwarts School

This team proposed a Higher-Resolution Network for RGB-to-RAW image mapping [20] which architecture is visualized in Fig. 3. To learn both local and global features, two parallel paths and a pyramid full-image encoder are used: the first one (denoted by the green color) is processing downsampled by a factor of 4 images and has 16 Residual in Residual Blocks (RIR), where each one consists of 10 CNN layers without channel attention units. The second one, full-resolution path, consists of 8 Multi-Scale Residual Blocks (denoted by the pink color) and is processing the images in the original resolution. Besides, the pyramid full-image encoder is processing the images in a fixed resolution into an n-dimensional vector. The extracted global fea-

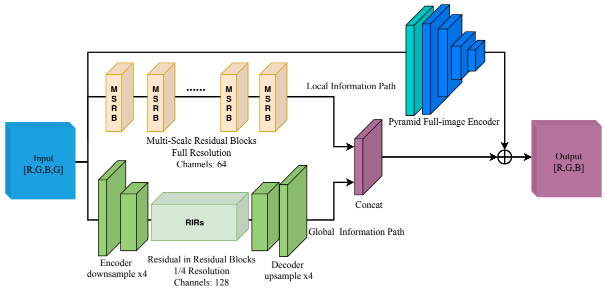

Figure 3. Higher-Resolution Network proposed by the First Team of Hogwarts School.

tures and local features are concatenated together, and the n-dimensional vector containing high-level features is then added, which enables the network to process images with arbitrary resolution. To accelerate the training speed, a progressive training method for high-resolution image patches was proposed. During the initial training period, small image patches (72×72 px) were used, and their resolution was increased gradually throughout the training while no other modifications were introduced to the network architecture. The model was trained for 48 hours on the four 4 Titan X GPUs to minimize the sum of L_1 and MSE losses.

4.3. TTI

Team TTI used the architecture illustrated in Fig. 4 that is inspired by the Deep Back-Projection Networks [9]. The model consists of iterative down-up projection units and is constructed based on the assumption that the down-projection unit can be used to remove the noise by down-scaling the feature-maps, while the up-projection unit is upscaling them back to the original resolution. The filter size in the projection units is 8×8 with striding by four and padding by two. All convolutional and deconvolutional layers are followed by parametric rectified linear units (PReLUs). The authors used error feedbacks from the up- and down-scaling steps to guide the network to achieve the optimal result.

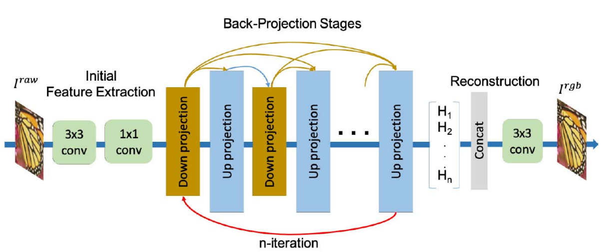

Figure 4. Deep Back-Projection Network used by TTI team.

In the fidelity track, only L_1 loss was used for training the model. In the perceptual track, a combination of the MSE, VGG (max pooling layers 2, 3, 4 and 5) [16], adversarial and style [16] losses was used. The network was trained with a batch size of 10 for 50 epochs with a learning rate initialized to $1e^{-4}$ and decreased by a factor of 10 after 25 epochs. For optimization, Adam algorithm with momentum 0.9 was used.

4.4. CityU TI Lab

The authors proposed a saliency map-aided generative adversarial network for RAW to RGB mapping [30] (Fig. 5) that contains a U-Net based generator with short-cut connections and a patch-based discriminator. The generator consists of one encoder and two decoders that jointly produce RGB images and saliency maps given a RAW input image. The saliency map decoder is attached to each layer of the RGB image decoder without sharing weights. 4 ResBlocks are utilized to preserve perceptual information so that mapping quality can be optimized. The discriminator receives the output or the ground truth RGB image and generates its feature embeddings evaluated by the LSGAN [19]. Instance normalization is used in each convolutional layer except the first and the last ones in order to address the artifact distortion problem.

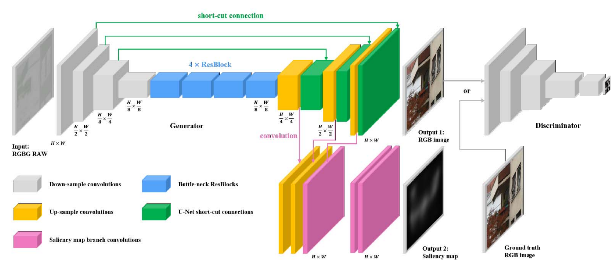

Figure 5. Saliency map-aided generative adversarial network proposed by CityU TI Lab.

The training process was divided into two stages. At the first stage, the generator is trained only with L_1 loss in order to achieve high pixel accuracy that is vital for stabilizing GAN training at second phase (since the generator already produces relatively good results). The input RAW and output RGB images were mapped to [-1, 1] interval, and the output saliency map was rescaled to [0, 1] range. At the seconds stage, a combination of the adversarial and VGG perceptual losses were added to the target loss function. The learning rate is initialized to $2e^{-4}$, and the whole system was trained for 10 epochs during the first stage and for 30 epochs during the second one. The model was trained using the Adam optimizer with β_1 = 0.5, β_2 = 0.999 and a batch size equal to 4.

4.5. superTeam

The authors adopted the model from [24] consisting of 20 convolutional layers. At each layer, 61 out of the 64 channels are standard feed-forward features, while the other 3 channels contain a correction for the RGB values of the previous block, *i.e.*, they contain a residual image that is added to the estimation of the previous layer. The model was trained for 1000 epochs using the Adam optimizer with a learning rate of $5e^{-6}$.

4.6. houbingxin

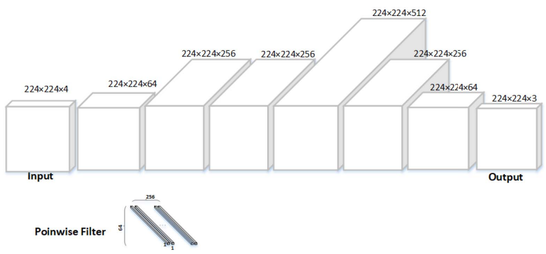

Figure 6. Houbingxin's network architecture.

The authors used a simple convolutional network illustrated in Fig. 6. The model consists of 7 pointwise convolutional layers followed by non-linear activation functions and contains in total 362,307 parameters.

4.7. Team Eraser

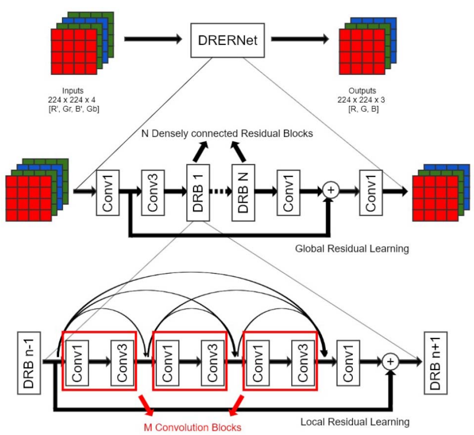

Figure 7. Densely connected residual network proposed by Team Eraser.

Team Eraser used an end-to-end network with densely connected residual blocks (Fig. 7) inspired by the RDN [29] architecture. The model contains 15 densely connected residual blocks, each one consisting of 3 convolutional blocks / 7 convolutional layers. The first and the last layers of the network are 1×1 convolutional layers with 12 and 3 feature maps, respectively, while in the second layer the number of feature maps is increased to 32. In total, the model contains 2.1M parameters and is trained for 6 epochs to minimize L_1 loss function using the Adam optimizer with a learning rate of $1e^{-4}$, momentum 0.9 and a batch size of 1.

Acknowledgments

We thank the AIM 2019 sponsors.

A. Appendix 1: Teams and affiliations

AIM 2019 RAW to RGB Mapping Team

Members:

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

Affiliation:

Computer Vision Lab, ETH Zurich, Switzerland

CVIP

Title:

W-Net: Two-stage U-Net with Misaligned Data for Raw-to-RGB Mapping

Members:

Sung-Jea Ko – sjko@korea.ac.kr,
Seung-Wook Kim, Kwang-Hyun Uhm, Seo-Won Ji, Sung-Jin Cho, Jun-Pyo Hong

Affiliations:

Korea University, Republic of Korea

The First Team of Hogwarts School

Title:

Higher-Resolution Network for Image Enhancement

Members:

Kangfu Mei – kangfumei@link.cuhk.edu.cn,
Juncheng Li, Jiajie Zhang, Haoyu Wu, Jie Li, Rui Huang

Affiliations:

The Chinese University of Hong Kong, China
Kuaishou Technology, China
East China Normal University, China

TTI

Title:

Deep Back-Projection Networks

Members:

Muhammad Haris – mharis@toyota-ti.ac.jp,
Greg Shakhnarovich, Norimichi Ukita

Affiliations:

Toyota Technological Institute (TTI), Japan
Toyota Technological Institute at Chicago (TTIC), USA

CityU TI Lab

Title:

Saliency Map-aided Generative Adversarial Network for RAW to RGB Mapping

Members:

Yuzhi Zhao – yzzhao2-c@my.cityu.edu.hk,
Lai-Man Po, Tiantian Zhang, Zongbang Liao, Xiang Shi, Yujia Zhang, Weifeng Ou, Pengfei Xian, Jingjing Xiong, Chang Zhou, Wing Yin Yu

Affiliations:

City University of Hong Kong, Hong Kong SAR, China
Huazhong University of Science and Technology, China

superTeam

Title:

Using Deep Convolution Networks for RAW to RGB Mapping

Members:

Yubin – 997597442wyb@gmail.com

Affiliations:

None

houbingxin

Title:

Pointwise Convolution for RAW to RGB Image Mapping

Members:

Bingxin Hou – houbingxin@gmail.com

Affiliations:

Santa Clara University, USA

Team Eraser

Title:

End-to-End RAW to RGB Mapping Network using Densely connected Residual Blocks

Members:

Bumjun Park – kkbbbj@gmail.com,
Songhyun Yu, Sangmin Kim, Jechang Jeong

Affiliations:

Image Communication and Signal Processing Laboratory, Hanyang University, Republic of Korea

References

[1] Abdelrahman Abdelhamed, Radu Timofte, Michael S. Brown, et al. Ntire 2019 challenge on real image denoising: Methods and results. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019.

[2] Codruta O. Ancuti, Cosmin Ancuti, Radu Timofte, 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.

[3] Yochai Blau, Roey Mechrez, Radu Timofte, Tomer Michaeli, and Lihi Zelnik-Manor. The 2018 pirm challenge on perceptual image super-resolution. In *The European Conference on Computer Vision (ECCV) Workshops*, September 2018.

[4] Jianrui Cai, Shuhang Gu, Radu Timofte, Lei Zhang, et al. Ntire 2019 challenge on real image super-resolution: Methods and results. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.

[5] Yu-Sheng Chen, Yu-Ching Wang, Man-Hsin Kao, and Yung-Yu Chuang. Deep photo enhancer: Unpaired learning for image enhancement from photographs with gans. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[6] Etienne de Stoutz, Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, and Luc Van Gool. Fast perceptual image enhancement. In *European Conference on Computer Vision Workshops*, 2018.

[7] Shuhang Gu and Radu Timofte. A brief review of image denoising algorithms and beyond. In *Springer series on Challenges in Machine Learning*, 2019.

[8] Shuhang Gu, Radu Timofte, Richard Zhang, et al. Ntire 2019 challenge on image colorization: Report. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019.

[9] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Deep backprojection networks for super-resolution. In *Conference on Computer Vision and Pattern Recognition*, 2018.

[10] Zheng Hui, Xiumei Wang, Lirui Deng, and Xinbo Gao. Perception-preserving convolutional networks for image enhancement on smartphones. In *European Conference on Computer Vision Workshops*, 2018.

[11] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Dslr-quality photos on mobile devices with deep convolutional networks. In *the IEEE Int. Conf. on Computer Vision (ICCV)*, 2017.

[12] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Wespe: weakly supervised photo enhancer for digital cameras. *arXiv preprint arXiv:1709.01118*, 2017.

[13] Andrey Ignatov and Radu Timofte. Ntire 2019 challenge on image enhancement: Methods and results. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.

[14] Andrey Ignatov, Radu Timofte, William Chou, Ke Wang, Max Wu, Tim Hartley, and Luc Van Gool. Ai benchmark: Running deep neural networks on android smartphones. In

Proceedings of the European Conference on Computer Vision (ECCV), pages 0–0, 2018.

[15] Andrey Ignatov, Radu Timofte, et al. Pirm challenge on perceptual image enhancement on smartphones: Report. In *European Conference on Computer Vision Workshops*, 2018.

[16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, pages 694–711. Springer, 2016.

[17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[18] Hanwen Liu, Pablo Navarrete Michelini, and Dan Zhu. Deep networks for image to image translation with mux and demux layers. In *European Conference on Computer Vision Workshops*, 2018.

[19] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2794–2802, 2017.

[20] Kangfu Mei, Juncheng Li, Jiajie Zhang, Haoyu Wu, Jie Li, and Rui Huang. Higher-resolution network for image demosaicing and enhancing. In *ICCV Workshops*, 2019.

[21] Seungjun Nah, Radu Timofte, Shuhang Gu, Sungyong Baik, Seokil Hong, Gyeongsik Moon, Sanghyun Son, and Kyoung Mu Lee. Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.

[22] Zhu Pengfei, Huang Jie, Geng Mingrui, Ran Jiewen, Zhou Xingguang, Xing Chen, Wan Pengfei, and Ji Xiangyang. Range scaling global u-net for perceptual image enhancement on mobile devices. In *European Conference on Computer Vision Workshops*, 2018.

[23] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[24] Eli Schwartz, Raja Giryes, and Alex M Bronstein. Deepisp: learning end-to-end image processing pipeline. *arXiv preprint arXiv:1801.06724*, 2018.

[25] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[26] Radu Timofte, Shuhang Gu, Jiqing Wu, and Luc Van Gool. 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.

[27] Kwang-Hyun Uhm, Seung-Wook Kim, Seo-Won Ji, Sung-Jin Cho, Jun-Pyo Hong, and Sung-Jea Ko. W-net: Two-stage u-net with misaligned data for raw-to-rgb mapping. In *ICCV Workshops*, 2019.

[28] Thang Van Vu, Cao Van Nguyen, Trung X Pham, Tung Minh Liu, and Chang D. Youu. Fast and efficient image quality enhancement via desubpixel convolutional neural networks. In *European Conference on Computer Vision Workshops*, 2018.

[29] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2472–2481, 2018.

[30] Yuzhi Zhao, Lai-Man Po, Tiantian Zhang, Zongbang Liao, Xiang Shi, Yujia Zhang, Weifeng Ou, Pengfei Xian, Jingjing Xiong, Chang Zhou, and Wing Yin Yu. Saliency map-aided generative adversarial network for raw to rgb mapping. In *ICCV Workshops*, 2019.