

Real-Time Video Super-Resolution on Smartphones with Deep Learning, Mobile AI 2021 Challenge: Report

Andrey Ignatov	Andres Romero	Heewon Kim	Radu Timofte	Chiu Man Ho
Zibo Meng	Kyoung Mu Lee	Yuxiang Chen	Yutong Wang	Zeyu Long
Chenhao Wang	Yifei Chen	Boshen Xu	Shuhang Gu	Lixin Duan
Wang Bofei	Zhang Diankai	Zheng Chengjian	Liu Shaoli	Gao Si
Zhang xiaofeng	Lu Kaidi	Xu Tianyu	Zheng Hui	Xinbo Gao
Jiaming Guo	Xueyi Zhou	Hao Jia	Youliang Yan	

Abstract

Video super-resolution has recently become one of the most important mobile-related problems due to the rise of video communication and streaming services. While many solutions have been proposed for this task, the majority of them are too computationally expensive to run on portable devices with limited hardware resources. To address this problem, we introduce the first Mobile AI challenge, where the target is to develop an end-to-end deep learning-based video super-resolution solutions that can achieve a real-time performance on mobile GPUs. The participants were provided with the REDS dataset and trained their models to do an efficient 4X video upscaling. The runtime of all models was evaluated on the OPPO Find X2 smartphone with the Snapdragon 865 SoC capable of accelerating floating-point networks on its Adreno GPU. The proposed solutions are fully compatible with any mobile GPU and can upscale videos to HD resolution at up to 80 FPS while demonstrating high fidelity results. A detailed description of all models developed in the challenge is provided in this paper.

1. Introduction

An increased popularity of various video streaming services and a widespread of mobile devices have created a strong need for efficient and mobile-friendly video super-resolution approaches. Over the past years, many accurate deep learning-based solutions have been proposed for this

problem [46, 48, 57, 37, 51, 50, 12]. The biggest limitation of these methods is that they were primarily targeted at achieving high fidelity scores while not optimized for computational efficiency and mobile-related constraints, which is essential for tasks related to image [18, 19, 32] and video [47] enhancement on mobile devices. In this challenge, we take one step further in solving this problem by using a popular REDS [46] video super-resolution dataset and by imposing additional efficiency-related constraints on the developed solutions.

When it comes to the deployment of AI-based solutions on mobile devices, one needs to take care of the particularities of mobile NPUs and DSPs to design an efficient model. An extensive overview of smartphone AI acceleration hardware and its performance is provided in [28, 25]. According to the results reported in these papers, the latest mobile NPUs are already approaching the results of mid-range desktop GPUs released not long ago. However, there are still two major issues that prevent a straightforward deployment of neural networks on mobile devices: a restricted amount of RAM, and a limited and not always efficient support for many common deep learning layers and operators. These two problems make it impossible to process high resolution data with standard NN models, thus requiring a careful adaptation of each architecture to the restrictions of mobile AI hardware. Such optimizations can include network pruning and compression [8, 22, 41, 43, 49], 16-bit / 8-bit [8, 36, 35, 59] and low-bit [6, 55, 33, 44] quantization, device- or NPU-specific adaptations, platform-aware neural architecture search [14, 52, 58, 56], etc.

While many challenges and works targeted at efficient deep learning models have been proposed recently, the evaluation of the obtained solutions is generally performed on desktop CPUs and GPUs, making the developed solutions not practical due to the above mentioned issues. To address this problem, we introduce the first *Mobile AI Workshop and*

*Andrey Ignatov, Andres Romero, Heewon Kim and Radu Timofte (andrey@vision.ee.ethz.ch, roandres@ethz.ch, ghimhw@gmail.com, radu.timofte@vision.ee.ethz.ch) are the main Mobile AI 2021 challenge organizers. The other authors participated in the challenge. Appendix A contains the authors' team names and affiliations.




Figure 1. Sample crops from the original video frames, results obtained by challenge participants and the target high-resolution frames.

Challenges, where all deep learning solutions are developed for and evaluated on real mobile devices. In this competition, the participating teams were provided with the original high-quality and downsampled by a factor of 4 videos from the REDS [46] dataset to train their networks. Within the challenge, the participants were evaluating the runtime and tuning their models on the OPPO Find X2 smartphone featuring the Qualcomm Adreno 650 GPU that can efficiently accelerate floating-point neural networks. The final score of each submitted solution was based on the runtime and fidelity results, thus balancing between the image reconstruction quality and efficiency of the proposed model. Finally, all developed solutions are fully compatible with the TensorFlow Lite framework [53], thus can be deployed and accelerated on any mobile platform providing AI acceleration through the Android Neural Networks API (NNAPI) [3] or custom TFLite delegates [10].

This challenge is a part of the *MAI 2021 Workshop and Challenges* consisting of the following competitions:

- Learned Smartphone ISP on Mobile NPUs [17]
- Real Image Denoising on Mobile GPUs [16]

- Quantized Image Super-Resolution on Mobile NPUs [26]
- Real-Time Video Super-Resolution on Mobile GPUs
- Single-Image Depth Estimation on Mobile Devices [20]
- Quantized Camera Scene Detection on Smartphones [21]
- High Dynamic Range Image Processing on Mobile NPUs

The results obtained in the other competitions and the description of the proposed solutions can be found in the corresponding challenge papers.

2. Challenge

To develop an efficient and practical solution for mobile-related tasks, one needs the following major components:

1. A high-quality and large-scale dataset that can be used to train and evaluate the solution;
2. An efficient way to check the runtime and debug the model locally without any constraints;
3. An ability to regularly test the runtime of the designed neural network on the target mobile platform or device.

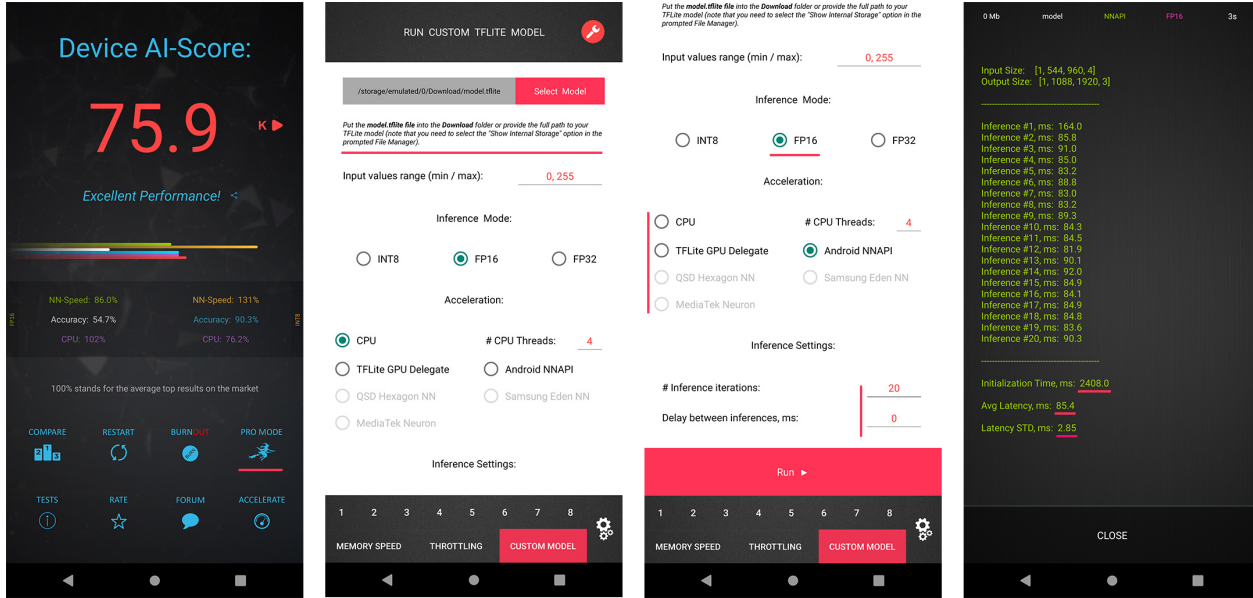


Figure 2. Loading and running custom TensorFlow Lite models with AI Benchmark application. The currently supported acceleration options include Android NNAPI, TFLite GPU, Hexagon NN, Samsung Eden and MediaTek Neuron delegates as well as CPU inference through TFLite or XNNPACK backends. The latest app version can be downloaded at <https://ai-benchmark.com/download>

This challenge addresses all the above issues. Real training data, tools, and runtime evaluation options provided to the challenge participants are described in the next sections.

2.1. Dataset

In this challenge, we use the REDS [46] dataset that serves as a benchmark for traditional video super-resolution task as it contains a large diversity of content and dynamic scenes. Following the standard procedure, we use 240 videos for training, 30 videos for validation, and 30 videos for testing. Each video has sequences of length 100, where every sequence contains video frames of 1280×720 resolution at 24 fps. To generate low-resolution data, the videos were bicubically downsampled with a factor of 4. The low-resolution video data is then considered as input, and the high-resolution — are the target.

2.2. Local Runtime Evaluation

When developing AI solutions for mobile devices, it is vital to be able to test the designed models and debug all emerging issues locally on available devices. For this, the participants were provided with the *AI Benchmark* application [25, 28] that allows to load any custom TensorFlow Lite model and run it on any Android device with all supported acceleration options. This tool contains the latest versions of *Android NNAPI*, *TFLite GPU*, *Hexagon NN*, *Samsung Eden* and *MediaTek Neuron* delegates, therefore supporting all current mobile platforms and providing the users with the ability to execute neural networks on smartphone NPUs, APUs, DSPs, GPUs and CPUs.

To load and run a custom TensorFlow Lite model, one needs to follow the next steps:

1. Download AI Benchmark from the official website¹ or from the Google Play² and run its standard tests.
2. After the end of the tests, enter the *PRO Mode* and select the *Custom Model* tab there.
3. Rename the exported TFLite model to *model.tflite* and put it into the *Download* folder of the device.
4. Select mode type (*INT8*, *FP16*, or *FP32*), the desired acceleration/inference options and run the model.

These steps are also illustrated in Fig. 2.

2.3. Runtime Evaluation on the Target Platform

In this challenge, we use the *OPPO Find X2* smartphone with the *Qualcomm Snapdragon 865* mobile SoC as our target runtime evaluation platform. The considered chipset demonstrates very decent AI Benchmark scores [4] and can be found in the majority of flagship Android smartphones released in 2020. It can efficiently accelerate floating-point networks on its Adreno 650 GPU with a theoretical FP16 performance of 2.4 TFLOPS. Within the challenge, the participants were able to upload their TFLite models to an external server and get a feedback regarding the speed of their model: the runtime of their solution on the above mentioned OPPO device or an error log if the network contains

¹<https://ai-benchmark.com/download>

²<https://play.google.com/store/apps/details?id=org.benchmark.demo>

Team	Author	Framework	Model Size, KB	PSNR \uparrow	SSIM \uparrow	Runtime per 10 frames \downarrow		Speed-Up	Final Score
						CPU, ms	GPU, ms		
Diggers	chenyuxiang	Keras / TensorFlow	230	28.33	0.8112	916	199	4.6	8.13
ZTE VIP	jieson.zheng	PyTorch / TensorFlow	50	27.85	0.7983	163	113	1.4	7.36
Rainbow	Zheng222	TensorFlow	204	27.99	0.8021	429	180	2.4	5.61
Noah_TerminalVision	JeremieG	TensorFlow	30	27.97	0.8017	448	Error *	-	2.19
Bicubic Upscaling	Baseline			26.50	0.7508				-

Table 1. Mobile AI 2021 Real-Time Video Super-Resolution challenge results and final rankings. During the runtime measurements, the models were upscaling 10 subsequent video frames from 180×320 to 1280×720 pixels on the OPPO Find X2 smartphone. Team *Diggers* is the challenge winner. * The solution from *Noah_TerminalVision* was not parsed correctly by the TFLite GPU delegate.

some incompatible operations. The models were parsed and accelerated using the TensorFlow Lite GPU delegate [40] demonstrating the best performance on this platform according to AI Benchmark results. The same setup was also used for the final runtime evaluation.

2.4. Challenge Phases

The challenge consisted of the following phases:

- I. *Development*: the participants get access to the data and AI Benchmark app, and are able to train the models and evaluate their runtime locally;
- II. *Validation*: the participants can upload their models to the remote server to check the fidelity scores on the validation dataset, to get the runtime on the target platform, and to compare their results on the validation leaderboard;
- III. *Testing*: the participants submit their final results, codes, TensorFlow Lite models, and factsheets.

2.5. Scoring System

All solutions were evaluated using the following metrics:

- Peak Signal-to-Noise Ratio (PSNR) measuring fidelity score,
- Structural Similarity Index Measure (SSIM), a proxy for perceptual score,
- The runtime on the target OPPO Find X2 smartphone.

The goal of this challenge was to produce an efficient solution balancing between the fidelity scores and latency. For the fidelity evaluation, we compute the PSNR and SSIM measures between the target sharp high-resolution and the produced super-resolved videos, both scores were averaged over the entire sequence of frames. Different to common VSR methods [57, 7] where the input of the model is a 5-dimensional tensor including the video sequence information, this challenge encouraged the participants to build models that receive mobile-friendly 4-dimensional tensors. The input model tensor should accept 10 subsequent video frames and have a size of $[1 \times 180 \times 320 \times 30]$, where the

first dimension is the batch size, the second and third dimensions are the height and width of the input frames, and the last dimension is the number of channels (3 color channels \times 10 frames). The size of the output tensor of the model should be $[1 \times 720 \times 1280 \times 30]$.

The score of each final submission was evaluated based on the next formula (C is a constant normalization factor):

$$\text{Final Score} = \frac{2^{2 \cdot \text{PSNR}}}{C \cdot \text{runtime}},$$

During the final challenge phase, the participants did not have access to the test dataset. Instead, they had to submit their final TensorFlow Lite models that were subsequently used by the challenge organizers to check both the runtime and the fidelity results of each submission under identical conditions. This approach solved all the issues related to model overfitting, reproducibility of the results, and consistency of the obtained runtime/accuracy values.

3. Challenge Results

From above 125 registered participants, 4 teams entered the final phase and submitted valid results, TFLite models, codes, executables and factsheets. Table 1 summarizes the final challenge results and reports PSNR, SSIM and runtime numbers for each submitted solution on the final test dataset and on the target evaluation platform, while Fig. 1 shows the obtained qualitative results. The proposed methods are described in section 4, and the team members and affiliations are listed in Appendix A.

3.1. Results and Discussion

All submitted solutions demonstrated a very high efficiency: the first three models can upscale video frames from 180×320 to 1280×720 resolution at more than 50 FPS on the target Snapdragon 865 chipset. The solution proposed by team *Noah_TerminalVision* can potentially achieve even higher frame rates, though it is currently not compatible with the TFLite GPU delegate due to *split* operations and thus was tested on Snapdragon’s CPU only. Team *Diggers* is the challenge winner — the model proposed by

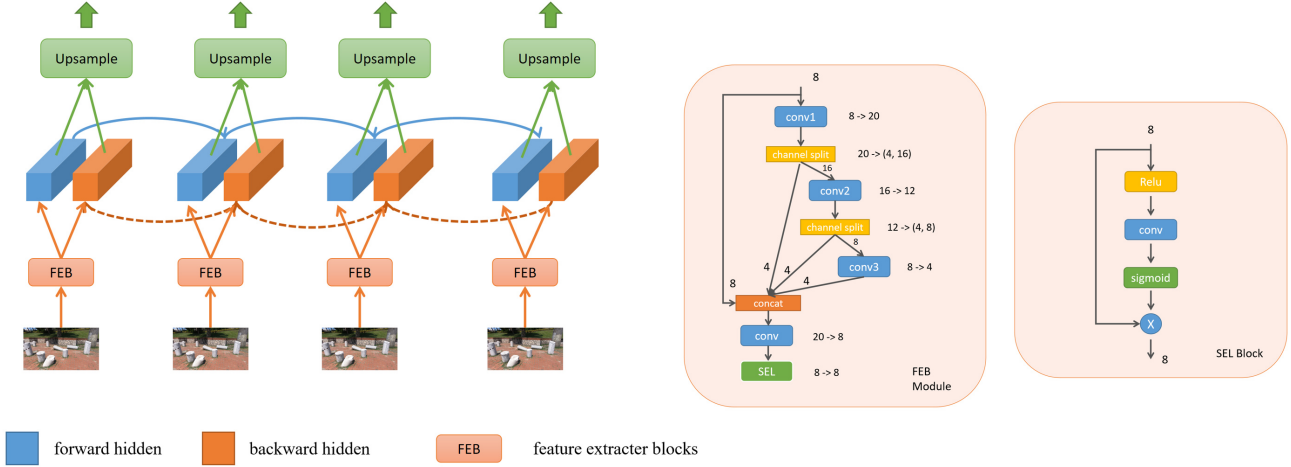


Figure 3. Team Diggers proposes a bidirectional RNN with efficient feature extractors (FEB) to exploit the temporal dependencies.

the authors achieves the best fidelity results while demonstrating good runtime values. This is the only solution in this challenge that applied recurrent connections to make use of inter-frame dependencies for getting better reconstruction results. While the other methods were performing only a standard per-frame upscaling, the image crops shown in Fig. 1 demonstrate that the visual quality of the reconstructed video frames obtained in their cases is just slightly behind the one of the winning solution, while all results are significantly better compared to the simple bicubic video interpolation. Therefore, we can conclude that, from the practical aspect, all proposed solutions can be applied for video super-resolution task on real mobile devices — the final choice will depend on the set of supported ops (*e.g.*, the first model might not be compatible with some mobile NPUs [26]) and the target FPS and runtime values.

4. Challenge Methods

This section describes solutions submitted by all teams participating in the final stage of the MAI 2021 Real-Time Video Super-Resolution challenge.

4.1. Diggers

Team Diggers proposed a bi-directional recurrent model for the considered video super-resolution task that uses feature maps computed for the previous and future video frames as an additional information while super-resolving the current frame (Fig. 3). The model architecture is generally based on the ideas proposed in [34] and [15]: for each input frame, two feature extraction blocks (FEBs) are applied to generate the corresponding feature maps: forward [blue] and backward [orange]. The forward feature maps of the current and previous frames are then combined and passed to another feature extraction block to generate the

final forward feature map for the current frame. As for the backward frames, the procedure is exactly the same, though the sequence is reversed. The obtained final forward and backward features are fed to the selection units layer (SEL) module [9], one IMDB [15] module and two convolutional and image resizing layers performing final frame upscaling.

During the training process, sequences of 21 subsequent low- and high-resolution video frames were used as model's inputs and targets. First, the model was trained for 31 epochs with a batch size of 16 and an initial learning rate of $4e-3$ multiplied by 0.7 each 2 epochs starting from the 7th one. Next, it was trained for another 31 epochs with a batch size of 32 and the same learning rate policy. L_2 loss was used as a target loss function, model parameters were optimized with Adam [39]. The images were additionally flipped randomly during the training for data augmentation.

4.2. ZTE VIP




Figure 4. Feed-forward CNN proposed by ZTE VIP team. Each basic module (BM) consists of a residual block with two convolutional layers.

Team ZTE VIP proposed a model performing per-frame upscaling without taking into account any inter-frame dependencies (Fig. 4) which can significantly speed-up the inference. In its first layer, the input is resized so that the batch size is equal to the number of video frames, and then they are processed separately by several residual blocks [13] and a *depth-to-space* layer performing final frame upsampling. The number of residual blocks and their size was found us-

ing the Neural Architecture Search (NAS) [38], where the target metric was composed of the fidelity loss and the number of model FLOPS. The final model contains five residual blocks, each one consisting of two 3×3 convolutions with eight feature maps. The network was trained to minimize L_1 loss with a batch size of 4 for 1000 epochs using Adam optimizer with a learning rate of $2e - 4$ down-scaled by a factor of 0.5 till 400 epochs. A more detailed description of the model, design choices and training procedure is provided in [42].

4.3. Rainbow




Figure 5. The overall model architecture proposed by Rainbow team (left). The structure of the information multi-distillation blocks (IMDB_s) is presented on the right side; 12, 9 and 3 denote the size of the output channels of each convolutional layer.

Similarly to the previous solution, Rainbow team has developed a pure CNN model performing per-frame video upscaling (Fig. 5). The authors presented a network consisting of three information multi-distillation blocks (IMDB_s) [15] followed by a *depth-to-space* upsampling layer processing each video frame separately. A global skip connection is used to improve the fidelity results of the model. L_1 loss was used as a target fidelity measure, network parameters were optimized using Adam with an initial learning rate of $2e - 4$ halved every 50K iterations. A batch size of 8 was used during the training, horizontal and vertical flipping was used to augment the training data.




Figure 6. The solution proposed by Noah.TerminalVision consists of a light-weight architecture with three residual blocks and asymmetric convolutions.

4.4. Noah.TerminalVision

Team Noah.TerminalVision presented a TinyVSRNet architecture that contains three residual blocks (each consisting of 2 convolutions with 16 channels) followed by *depth-to-space* upsampling layer and one global skip connection performing bilinear image upscaling (Fig. 6). The authors have also proposed a “single-frame” solution by converting 10 video frames from the channel dimension to the batch dimension with *split* and *concat* layers. To boost the model performance, they adopted the approach developed in [11], where three asymmetric convolution kernels (of size 3×3 , 1×3 , and 3×1) are used during the training and then fused into one single convolution op during the inference. With this modification, the results of the TinyVSRNet were improved by around 0.05 dB. The network was trained to minimize L_1 loss, its parameters were optimized using Adam for one million iterations with a cyclic learning rate starting from $5e - 4$ and decreased to $1e - 6$ each 200K iterations.

5. Additional Literature

An overview of the past challenges on mobile-related tasks together with the proposed solutions can be found in the following papers:

- Video Super-Resolution: [48, 47]
- Image Super-Resolution: [30, 45, 5, 54]
- Learned End-to-End ISP: [27, 31]
- Perceptual Image Enhancement: [30, 24]
- Bokeh Effect Rendering: [23, 29]
- Image Denoising: [1, 2]

Acknowledgements

We thank OPPO Mobile Co., Ltd., AI Witchlabs, ETH Zurich (Computer Vision Lab) and Seoul National University (SNU), the organizers and sponsors of this Mobile AI 2021 challenge.

A. Teams and Affiliations

Mobile AI 2021 Team

Title:

Mobile AI 2021 Image Super-Resolution Challenge

Members:

Andrey Ignatov^{1,3} (andrey@vision.ee.ethz.ch), Andres Romero¹ (roandres@ethz.ch), Heewon Kim⁴ (ghimhw@gmail.com), Radu Timofte^{1,3} (radu.timofte@vision.ee.ethz.ch), Chiu Man Ho² (chiuman@oppo.com), Zibo Meng² (zibo.meng@oppo.com), Kyoung Mu Lee⁴

Affiliations:

¹ Computer Vision Lab, ETH Zurich, Switzerland

² OPPO Mobile Co., Ltd, China

³ AI Witchlabs, Switzerland

⁴ Seoul National University, South Korea

Diggers

Title:

Real-Time Video Super-Resolution based on Bidirectional RNNs

Members:

Yuxiang Chen (chenyx.cs@gmail.com), Yutong Wang, Zeyu Long, Chenhao Wang, Yifei Chen, Boshen Xu, Shuhang Gu, Lixin Duan, Wen Li

Affiliations:

University of Electronic Science and Technology of China (UESTC), China

ZTE VIP

Title:

Evsrnet: Efficient video super-resolution with neural architecture search. [42]

Members:

Wang Bofei (wang.bofei@zte.com.cn), Zhang Diankai, Zheng Chengjian, Liu Shaoli, Gao Si, Zhang xiaofeng, Lu Kaidi, Xu Tianyu

Affiliations:

Audio & Video Technology Platform Department, ZTE Co., Ltd, China

Rainbow

Title:

Video Super-Resolution with Multi-Information Distillation Network on Mobile Device

Members:

Zheng Hui (zheng_hui@aliyun.com), Xinbo Gao, Xiumei Wang

Affiliations:

School of Electronic Engineering, Xidian University, China

Noah.TerminalVision

Title:

TinyVSRNet for Real-Time Video Super-Resolution

Members:

Jiaming Guo (guojiaming5@huawei.com), Xueyi Zhou, Hao Jia, Youliang Yan

Affiliations:

Huawei Technologies Co., Ltd, China

References

- [1] Abdelrahman Abdelhamed, Mahmoud Afifi, Radu Timofte, and Michael S Brown. Ntire 2020 challenge on real image denoising: Dataset, methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 496–497, 2020. 6
- [2] Abdelrahman Abdelhamed, Radu Timofte, and Michael S Brown. Ntire 2019 challenge on real image denoising: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 6
- [3] Android Neural Networks API. <https://developer.android.com/ndk/guides/neuralnetworks>. 2
- [4] AI Benchmark Archive. http://web.archive.org/web/20210425131428/https://ai-benchmark.com/ranking_processors.html. 3
- [5] Jianrui Cai, Shuhang Gu, Radu Timofte, and Lei Zhang. Ntire 2019 challenge on real image super-resolution: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 6
- [6] Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Zeroq: A novel zero shot quantization framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13169–13178, 2020. 1
- [7] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Basicvsr: The search for essential components in video super-resolution and beyond. *arXiv preprint arXiv:2012.02181*, 2020. 4
- [8] Cheng-Ming Chiang, Yu Tseng, Yu-Syuan Xu, Hsien-Kai Kuo, Yi-Min Tsai, Guan-Yu Chen, Koan-Sin Tan, Wei-Ting Wang, Yu-Chieh Lin, Shou-Yao Roy Tseng, et al. Deploying image deblurring across mobile devices: A perspective of quality and latency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 502–503, 2020. 1
- [9] Jae-Seok Choi and Munchurl Kim. A deep convolutional neural network with selection units for super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 154–160, 2017. 5

- [10] TensorFlow Lite delegates. <https://www.tensorflow.org/lite/performance/delegates>. 2
- [11] Xiaohan Ding, Yuchen Guo, Guiguang Ding, and Jungong Han. Acnet: Strengthening the kernel skeletons for powerful cnn via asymmetric convolution blocks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1911–1920, 2019. 6
- [12] Dario Fuoli, Shuhang Gu, and Radu Timofte. Efficient video super-resolution through recurrent latent space propagation. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3476–3485. IEEE, 2019. 1
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [14] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1314–1324, 2019. 1
- [15] Zheng Hui, Xinbo Gao, Yunchu Yang, and Xiumei Wang. Lightweight image super-resolution with information multi-distillation network. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 2024–2032, 2019. 5, 6
- [16] Andrey Ignatov, Kim Byeoung-su, and Radu Timofte. Fast camera image denoising on mobile gpus with deep learning, mobile ai 2021 challenge: Report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2021. 2
- [17] Andrey Ignatov, Jimmy Chiang, Hsien-Kai Kuo, Anastasia Sycheva, and Radu Timofte. Learned smartphone isp on mobile npus with deep learning, mobile ai 2021 challenge: Report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2021. 2
- [18] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Dslr-quality photos on mobile devices with deep convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3277–3285, 2017. 1
- [19] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Wespe: weakly supervised photo enhancer for digital cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 691–700, 2018. 1
- [20] Andrey Ignatov, Grigory Malivenko, David Plowman, Samarth Shukla, and Radu Timofte. Fast and accurate single-image depth estimation on mobile devices, mobile ai 2021 challenge: Report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2021. 2
- [21] Andrey Ignatov, Grigory Malivenko, and Radu Timofte. Fast and accurate quantized camera scene detection on smartphones, mobile ai 2021 challenge: Report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2021. 2
- [22] Andrey Ignatov, Jagruti Patel, and Radu Timofte. Rendering natural camera bokeh effect with deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 418–419, 2020. 1
- [23] Andrey Ignatov, Jagruti Patel, Radu Timofte, Bolun Zheng, Xin Ye, Li Huang, Xiang Tian, Saikat Dutta, Kuldeep Purohit, Praveen Kandula, et al. Aim 2019 challenge on bokeh effect synthesis: Methods and results. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3591–3598. IEEE, 2019. 6
- [24] Andrey Ignatov and Radu Timofte. Ntire 2019 challenge on image enhancement: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 6
- [25] 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) Workshops*, pages 0–0, 2018. 1, 3
- [26] Andrey Ignatov, Radu Timofte, Maurizio Denna, and Abdel Younes. Real-time quantized image super-resolution on mobile npus, mobile ai 2021 challenge: Report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2021. 2, 5
- [27] 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, et al. Aim 2019 challenge on raw to rgb mapping: Methods and results. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3584–3590. IEEE, 2019. 6
- [28] Andrey Ignatov, Radu Timofte, Andrei Kulik, Seungsoo Yang, Ke Wang, Felix Baum, Max Wu, Lirong Xu, and Luc Van Gool. Ai benchmark: All about deep learning on smartphones in 2019. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3617–3635. IEEE, 2019. 1, 3
- [29] Andrey Ignatov, Radu Timofte, Ming Qian, Congyu Qiao, Jiamin Lin, Zhenyu Guo, Chenghua Li, Cong Leng, Jian Cheng, Juewen Peng, et al. Aim 2020 challenge on rendering realistic bokeh. In *European Conference on Computer Vision*, pages 213–228. Springer, 2020. 6
- [30] Andrey Ignatov, Radu Timofte, Thang Van Vu, Tung Minh Luu, Trung X Pham, Cao Van Nguyen, Yongwoo Kim, Jae-Seok Choi, Munchurl Kim, Jie Huang, et al. Pirm challenge on perceptual image enhancement on smartphones: Report. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 6
- [31] Andrey Ignatov, Radu Timofte, Zhilu Zhang, Ming Liu, Haolin Wang, Wangmeng Zuo, Jiawei Zhang, Ruimao Zhang, Zhanglin Peng, Sijie Ren, et al. Aim 2020 challenge on learned image signal processing pipeline. *arXiv preprint arXiv:2011.04994*, 2020. 6
- [32] Andrey Ignatov, Luc Van Gool, and Radu Timofte. Replacing mobile camera isp with a single deep learning model. In *Proceedings of the IEEE/CVF Conference on Computer*

- Vision and Pattern Recognition Workshops*, pages 536–537, 2020. 1
- [33] Dmitry Ignatov and Andrey Ignatov. Controlling information capacity of binary neural network. *Pattern Recognition Letters*, 138:276–281, 2020. 1
 - [34] Takashi Isobe, Fang Zhu, Xu Jia, and Shengjin Wang. Re-visiting temporal modeling for video super-resolution. *arXiv preprint arXiv:2008.05765*, 2020. 5
 - [35] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2704–2713, 2018. 1
 - [36] Sambhav R Jain, Albert Gural, Michael Wu, and Chris H Dick. Trained quantization thresholds for accurate and efficient fixed-point inference of deep neural networks. *arXiv preprint arXiv:1903.08066*, 2019. 1
 - [37] Armin Kappeler, Seunghwan Yoo, Qiqin Dai, and Aggelos K Katsaggelos. Video super-resolution with convolutional neural networks. *IEEE Transactions on Computational Imaging*, 2(2):109–122, 2016. 1
 - [38] Heewon Kim, Seokil Hong, Bohyung Han, Heesoo Myeong, and Kyoung Mu Lee. Fine-grained neural architecture search. *arXiv preprint arXiv:1911.07478*, 2019. 6
 - [39] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
 - [40] Juhyun Lee, Nikolay Chirkov, Ekaterina Ignasheva, Yury Pisarchyk, Mogan Shieh, Fabio Riccardi, Raman Sarokin, Andrei Kulik, and Matthias Grundmann. On-device neural net inference with mobile gpus. *arXiv preprint arXiv:1907.01989*, 2019. 4
 - [41] Yawei Li, Shuhang Gu, Luc Van Gool, and Radu Timofte. Learning filter basis for convolutional neural network compression. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5623–5632, 2019. 1
 - [42] Shaoli Liu, Chengjian Zheng, Kaidi Lu, Si Gao, Ning Wang, Bofei Wang, Diankai Zhang, Xiaofeng Zhang, and Tianyu Xu. Evsrnet: Efficient video super-resolution with neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2021. 6, 7
 - [43] Zechun Liu, Haoyuan Mu, Xiangyu Zhang, Zichao Guo, Xin Yang, Kwang-Ting Cheng, and Jian Sun. Metapruning: Meta learning for automatic neural network channel pruning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3296–3305, 2019. 1
 - [44] Zechun Liu, Baoyuan Wu, Wenhan Luo, Xin Yang, Wei Liu, and Kwang-Ting Cheng. Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. In *Proceedings of the European conference on computer vision (ECCV)*, pages 722–737, 2018. 1
 - [45] Andreas Lugmayr, Martin Danelljan, and Radu Timofte. Ntire 2020 challenge on real-world image super-resolution: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 494–495, 2020. 6
 - [46] Seungjun Nah, Sungyong Baik, Seokil Hong, Gyeongsik Moon, Sanghyun Son, Radu Timofte, and Kyoung Mu Lee. Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 1, 2, 3
 - [47] Seungjun Nah, Sanghyun Son, Radu Timofte, and Kyoung Mu Lee. Ntire 2020 challenge on image and video deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 416–417, 2020. 1, 6
 - [48] Seungjun Nah, Radu Timofte, Shuhang Gu, Sungyong Baik, Seokil Hong, Gyeongsik Moon, Sanghyun Son, and Kyoung Mu Lee. Ntire 2019 challenge on video super-resolution: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 1, 6
 - [49] Anton Obukhov, Maxim Rakhuba, Stamatis Georgoulis, Menelaos Kanakis, Dengxin Dai, and Luc Van Gool. T-basis: a compact representation for neural networks. In *International Conference on Machine Learning*, pages 7392–7404. PMLR, 2020. 1
 - [50] Mehdi SM Sajjadi, Raviteja Vemulapalli, and Matthew Brown. Frame-recurrent video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6626–6634, 2018. 1
 - [51] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016. 1
 - [52] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2820–2828, 2019. 1
 - [53] TensorFlow-Lite. <https://www.tensorflow.org/lite>. 2
 - [54] Radu Timofte, Shuhang Gu, Jiqing Wu, and Luc Van Gool. Ntire 2018 challenge on single image super-resolution: Methods and results. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 852–863, 2018. 6
 - [55] Stefan Uhlich, Lukas Mauch, Fabien Cardinaux, Kazuki Yoshiyama, Javier Alonso Garcia, Stephen Tiedemann, Thomas Kemp, and Akira Nakamura. Mixed precision dnns: All you need is a good parametrization. *arXiv preprint arXiv:1905.11452*, 2019. 1
 - [56] Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuan-dong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, et al. Fbnetv2: Differentiable neural architecture search for spatial and channel dimensions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12965–12974, 2020. 1

- [57] Xintao Wang, Kelvin CK Chan, Ke Yu, Chao Dong, and Chen Change Loy. Edvr: Video restoration with enhanced deformable convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. [1](#), [4](#)
- [58] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10734–10742, 2019. [1](#)
- [59] Jiwei Yang, Xu Shen, Jun Xing, Xinmei Tian, Houqiang Li, Bing Deng, Jianqiang Huang, and Xian-sheng Hua. Quantization networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7308–7316, 2019. [1](#)