

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

IMDeception: Grouped Information Distilling Super-Resolution Network

Mustafa Ayazoğlu Aselsan Research Ankara, Turkey

mayazoglu@aselsan.com.tr

Abstract

Single-Image-Super-Resolution (SISR) is a classical computer vision problem that has benefited from the recent advancements in deep learning methods, especially the advancements of convolutional neural networks (CNN). Although state-of-the-art methods improve the performance of SISR on several datasets, direct application of these networks for practical use is still an issue due to heavy computational load. For this purpose, recently, researchers have focused on more efficient and high-performing network structures. Information multi-distilling network (IMDN) is one of the highly efficient SISR networks with high performance and low computational load. IMDN achieves this efficiency with various mechanisms such as Intermediate Information Collection (IIC), working in a global setting, Progressive Refinement Module (PRM), and Contrast Aware Channel Attention (CCA), employed in a local setting. These mechanisms, however, do not equally contribute to the efficiency and performance of IMDN. In this work, we propose the Global Progressive Refinement Module (GPRM) as a less parameter-demanding alternative to the IIC module for feature aggregation. To further decrease the number of parameters and floating point operations per second (FLOPS), we also propose Grouped Information Distilling Blocks (GIDB). Using the proposed structures, we design an efficient SISR network called IMDeception. Experiments reveal that the proposed network performs on par with state-of-the-art models despite having a limited number of parameters and FLOPS. Furthermore, using grouped convolutions as a building block of GIDB increases room for further optimization during deployment. To show its potential, the proposed model was deployed on NVIDIA Jetson Xavier AGX and it has been shown that it can run in realtime on this edge device.

1. Introduction

Single-Image-Super-Resolution (SISR) is a well-studied computer vision problem. The problem's goal is to create



(a) Original: Div2K 0886



Figure 1. Results of our method compared with other methods

a high-resolution image from a single low-resolution image. Due to its nature, it is an ill-posed problem. Starting with the seminal work of Dong *et al.* [5] the problem is addressed by using deep-learning approaches. Dong's model used a CNN with only 3-layers and beat the traditional approaches. Later on, to decrease the computational load, FSRCNN [6] proposed postponing the upscaling to the end of the network while most of the computation and feature extracting done in low-resolution. Shi. *et al.* proposed ESPCN, [26] which replaced the transposed convolution layer with Depth2Space operator. Later, Kim *et al.* proposed [15], a 20-layer network, and showed that increasing the number of parameters can improve the network's performance. EDSR proposed by Lim et al. [21] further improved the state-of-the-art by increasing the number of layers and omitting BatchNorm layers from the network. Later on, Yu et al. proposed WDSR [32], a network with 75M parameters and improved super-resolution results. Indeed, increasing the number of parameters improves the performance of a network, but it also makes it harder to use it in many practical real-time scenarios. For these reasons, researches started working on efficient models which aim to maintain image reconstruction performance with those of millions-of-parameter-networks while still being applicable for real-time scenarios [20, 33]. To decrease the number of parameters, recursive networks are employed [14, 29], but the number of FLOPS is very high for these networks. Besides this work, there are some work incorporating the attention idea into the SISR domain, such as [24, 34], which increases the receptive field and hence the performance of the network while keeping the parameters low at the cost of an increased number of operations.

In this context, Hui *et al.* proposed IDN [12] which uses channel splitting method to separate the high-level features from the low-level ones while keeping the number of parameters low and maintaining acceptable performance. IMDN [11] further investigated the channel splitting idea at a granularity level and further improved the performance and inference speed. Besides channel splitting, IMDN employed Intermediate Information Collection (IIC) at the global level to accumulate the information from different information multi-distilling blocks (IMDB) and in the IMDBs it used Progressive Refinement Module (PRM) which splits the outputs of different convolution layers such that a portion of the information is directly flows to the end of the block while the rest is fed to the next Conv2d layer for further refinement.

Although IMDN is an efficient and well-performing network, global information fusion modules (IIC) and IMDB blocks are not ideal and there is still room for improvement. To this end, following the Network-in-Network [22] and Inception [28] spirit, we propose the Global Progressive Refinement Module (GPRM), which is an extension of the PRM in the global setting, in-place of the IIC module. Using GPRM gives us the flexibility to control the number of parameters while being able to integrate the midlevel information to the end of the network. To further reduce the number of parameters and operations, we proposed grouped information distilling blocks (GIDB) as the building blocks that employ grouped convolutions. Using grouped convolutions increases the room for further optimization during deployment. Furthermore, by incorporating the block-based non-local attention (NLA) blocks at the global level, [30] we further improved the performance of the proposed model.

Reconstruction efficiency of the model is shown in various different datasets, and inference efficiency is shown using NVIDIA TensorRT since it is training framework agnostic and optimizes the network for the hardware at the hand.

2. Related Works

As with many computer vision problems, SISR has benefited a lot from the recent advancements in deep learning. The first SISR model using deep learning started with the work of Dong [5]. Later on, by postponing the upscaling stage to the end and processing the input image at a lower resolution, FSRCNN [6] improved the inference speed. FS-RCNN also replaced ReLU activation with PReLU. Later on VDSR [15] introduced a deeper network and introduced a long upscaling skip connection. These showed that deeper networks improve the performance and long skip connection helps with the optimization. The same spirit continued with recursive architectures where a shared parameter sub-network is repeatedly applied at a cost of increased operations to solve SISR problem. LapSRN [16] aimed efficient super-resolution and used Laplacian pyramids to progressively extract features and reconstruct images at different scales with the same network. EDSR [21] improved the reconstruction results by eliminating Batch Normalization layers from the network and increasing the number of parameters to 43M. WDSR [32] further increased the parameters of the model to 75M and improved the results of EDSR. RDN [35] used DenseNet [10] style intermediate feature aggregation with residual blocks. More recently, researchers incorporated new ideas (such as grouped convolutions, attention layers etc.) into super resolution networks [3, 4, 24]. One obvious thing that can be deduced from these advancements is that as the number of parameters increases, the performance of the model increases as well. However, this comes at the cost of the model being practically not applicable. For these reasons, research interest in SISR has recently shifted towards building efficient models [20, 33]. IDN [12] follows this spirit; it uses channel splitting to distil features efficiently. IMDN further improves on this idea and uses channel splitting at granularity level and proposes information multi-distilling block (IMDB) which also includes a contrast-aware channel attention (CCA) layer. At the global level distilled information from the IMDBs are aggregated using Intermediate Information Collection (IIC). In this type of information collection, the information from the intermediate levels directly flows to the ends of the model. Indeed, this can be seen as a subset of the information collection type used in DenseNet and RDN, where DenseNet structure in RDN allows intermediate-to-intermediate flow as well.

The problem of a deep learning model not being practically applicable is indeed a problem with other deeplearning models from different fields as well. Because of this, researchers proposed different approaches that can make a model run in real-time, such as, Hand Picked Architectures / Blocks, Network Pruning/ Sparsification, Knowledge Distilling, Quantization, Network Architecture Search (NAS)

Hand-picked architectures focuses on manually designed architectures and blocks. Network sparsification and pruning, such as [8], follow a different approach and try to eliminate the redundancies in a larger network to come up with a more efficient network. Knowledge distilling [9] uses heavy teacher and lighter student networks in a setting where the teacher network guides the student network. Quantization, such as [25], focuses on the deployment side and tries to keep the network performance under lighter arithmetic operations. Network architecture search [36] goes beyond these ideas and tries to find the network architecture in an optimization setting.

Indeed, these ideas can be used to design superresolution networks as well. For this purpose, Li *et al.* [18] proposed a differentiable pruning model. Their method reduced the number of parameters, FLOPS, and run-time of EDSR Baseline [21] and several other networks by a significant amount. In [19], Li et al. proposed Layer-Wise Differentiable Network Architectures to adjust the channel sizes of predefined networks and successfully reduced the number of parameters of EDSR Baseline while improving its performance. Song et al. [27] proposed an evolutionary network search algorithm for efficiently searching residual dense blocks for super-resolution networks. Wu et al. [31] proposed a trilevel NAS algorithm for optimizing networks, cells, and kernels of super-resolution networks at the same time. In [17] Li et al. followed a different approach for reducing the number of parameters, and proposed a learning basis for convolutional layers. Their method compresses the number of parameters of EDSR Baseline by up to 93%.

While designing IMDeception we followed a manual approach since other approaches can still be applied to further push its limits.

3. Proposed Method

In this section, we describe the details of the proposed network. As it was mentioned before, the main motivation of this paper is efficiency while keeping performance at a comparable level with million-parameter networks. As a starting point, IMDN [11] is selected as the baseline of our work. The original IMDN architecture can be seen in Fig. 3a.

The variations of this network are already known to be high-performing [20,23,33] and improving it is challenging. This is due to the already employed network mechanisms such as Progressive Refinement Module (PRM) (Fig. 3c) Contrast Aware Channel Attention (CCA) (Fig. 3c) and Intermediate Information Collection(IIC) (Fig. 3a) are very efficient. The modules are indeed studied further in the original work, and each of these modules' contributions to the final model is noted. The individual contributions of each module can be seen in Tab. 1

Module-Dataset	Set5	Set14	Manga109
Basic	31.86	28.43	29.92
PRM Improvement	0.15	0.06	0.24
CCA Improvement	0.09	0.02	0.09
IIC Improvement	0.01	0.01	0.03

Table 1. PRM, CCA and IIC modules' contributions to the basic network, *The table is derived from Table 3 of [11]

Note that in IMDN, PRM and CCA are used locally inside the IMDBs, however IIC is used in a more global setting. Also note that the improvement provided by the PRM is much larger than the CCA and IIC. Furthermore the number of parameters drops with the PRM. Motivated by these facts and inspired by the Inception network's repeated structure [28] we created a network structure where PRM is repeated locally in the blocks and globally among the blocks to improve performance and reduce the number of parameters. This is done in such a way that IIC in the global setting is replaced with the proposed Global PRM (Fig. 2a). Furthermore, CCA layers are used in every IMDB, but the performance contributions from these layers are marginal compared to the number of operations and parameters that they add to the network. However, since attention layers are great at increasing the receptive field, we decided to use a limited number of block-based non-local attention blocks [30] in our proposed network's main path. To further reduce the number of parameters and number of operations of the network, every single Conv2D operation inside the IMDB is replaced with Gblocks (Fig. 2b) as in XLSR [3] which is based on grouped convolutions. We call these group convolution based structures as Grouped Information Distilling Blocks (GIDB). Although the grouped convolutions are not well optimized in training frameworks, [7] if utilised correctly within an inference-oriented framework, group convolutions can lead to speed ups as noted in [3, 7]especially in mobile devices where efficient network structures are usually employed.

Mathematically, the model can be described as follows; Given a low resolution image I^{lr} , super-resolved image, I^{sr} , can be obtained as follows:

$$I^{sr} = H_{IMDeception}(I^{lr}) \tag{1}$$

Here, $H_{IMDeception}()$ is our proposed optimized superresolution model. In the begining of the network a 64 channel 3x3 convolution is employed for feature extraction, as



(a) IMDeception Network



(c) GIDB Block

(b) Gblock: Here Red and Orange Stripes show ReLU and LeakyReLU activations respectively.

Figure 2. IMDeception: Proposed Architecture

in IMDN, let I_0^{feat} represent these features. These features are both transferred to the end of the network and processed in the Global Progressive Refinement Module as follows:

$$\begin{split} I_{1}^{dfeat}, I_{1}^{sfeat} &= Split(GIDB_{1}(I_{0}^{feat})) \\ I_{2}^{dfeat}, I_{2}^{sfeat} &= Split(GIDB_{2}(I_{1}^{dfeat})) \\ I_{2}^{dfeat}, I_{2}^{sfeat} &= Split(GIDB_{3}(I_{2}^{dfeat})) \\ I_{3}^{dfeat}, I_{3}^{sfeat} &= Split(GIDB_{4}(I_{3}^{dfeat})) \\ I_{4}^{dfeat}, I_{4}^{sfeat} &= Split(GIDB_{4}(I_{3}^{dfeat})) \\ I_{4}^{dfeat}, I_{5}^{sfeat} &= Split(GIDB_{5}(I_{4}^{dfeat})) \\ I_{5}^{dfeat}, I_{5}^{sfeat} &= Split(GIDB_{5}(I_{4}^{dfeat})) \\ I_{6}^{sfeat} &= GIDB_{6}(I_{5}^{dfeat}) \\ I_{7}^{feat} &= Concat([I_{1}^{sfeat}, ..., I_{6}^{sfeat}]) \end{split}$$

In the above equations, Split(), Concat(), NLA() represent 3/1 ratio channel splitting, channel concatenation, and block based non-local attention. I_n^{dfeat}, I_n^{sfeat} are

channel split features of $n^{th} GIDB$ block which is our proposed Grouped Information Distilling Block (GIDB). Note that here GPRM is used for global feature distilling and aggregation, and operating on the outputs of GIDBs. At the local level, features are further processed by GIDBs as follows:

$$\begin{split} I_{n-1}^{dfeat1}, I_{n-1}^{sfeat1} &= Split(GBlock_{1}(I_{n-1}^{dfeat})) \\ I_{n-1}^{dfeat2}, I_{n-1}^{sfeat2} &= Split(GBlock_{2}(I_{n-1}^{dfeat1})) \\ I_{n-1}^{dfeat3}, I_{n-1}^{sfeat3} &= Split(GBlock_{3}(I_{n-1}^{dfeat2})) \\ I_{n-1}^{dfeat4}, I_{n-1}^{sfeat4} &= Split(GBlock_{4}(I_{n-1}^{dfeat3})) \\ I_{n}^{dfeat} &= H_{fuse}(Concat([I_{n-1}^{sfeat1}, ..., I_{n-1}^{sfeat4}, I_{n-1}^{dfeat}])) \end{split}$$

Here H_{fuse} represents 1x1 convolution operation used for information fusion, Note that at the local level input features I_{n-1}^{dfeat} are processed and refined in grouped fashion using Gblock()s. Gblock() is implemented using 3x3 grouped convolution (groups=4) and cascaded 1x1 convolution to allow information flow between the groups. Using information grouping and processing the features in a



(a) IMDN Original Network: Note that here conv2d 1x1 includes a leaky relu activation.



Figure 3. IMDN Structure and the submodules PRM and IIC are included here for reference. CCA layer details are omitted and can be found in the original work.

grouped fashion reduces the number of parameters while almost at no cost of performance loss. The detailed Gblock() implementation along with the used activation functions can be seen in Fig. 2b.

The output of the GPRM module, I_7^{feat} , is further processed to construct the super-resolved image, I^{sr} as follows:

$$\begin{split} I_8^{feat} &= LeakyReLU(H_{1x1}(I_7^{feat}))\\ I_9^{feat} &= LeakyReLU(H_{3x3}(I_8^{feat})) + I_0^{feat} \qquad (4)\\ I^{sr} &= Upsample(I_9^{feat}) \end{split}$$

Here H_{1x1} , H_{3x3} represent 1x1 and 3x3 convolution layers respectively. We used Leaky ReLU (slope=0.05) activations. Upsample() is the upsampling layer implemented as shown in Fig. 3b

Our proposed network structure, which we call IMDeception, combining all of these ideas, can be seen in Fig. 2.

Note that we used global PRM among the GIDB and local PRM as in IMDB inside GIDB. Our proposed architecture defines a class of highly efficient architectures sharing the same structure with different channel numbers on the filters. As it can be seen from Fig. 2, we define the complexity of the models using *core* parameter. Depending on the needs, the *core* parameter can be used to adjust the complexity of the network. From our experiments, we have observed that even *core* = 4 with no attention blocks can still show high reconstruction performance with great inference timings.

The performance parameters of various IMDeception networks using *core* parameter and existence of attention blocks can be seen in Tab. 2

4. Experiments

4.1. Datasets

For the training, we used Div2K Dataset [1], and Flickr2K [21] combined (DF2K). In total, the training set includes, 3450 images, and for validation we used Div2K validation set, which includes 100 images.

4.2. Training Details

The proposed model (core = 16 + NLA) was trained in two different phases in all of the phases we used;

IMDeception	core = 16 + NLA	core = 12	core = 8	core = 4 + NLA	core = 4
#Params	316K	198K	113K	57K	57K
FLOPS[G]	20.7	12.9	7.4	3.7	3.7
Act.[M]	206	149	103	57	57
Runtime[ms]*	60	45	31	24	21
#Conv2d	62	58	58	62	58
#Div2K Val. (PSNR)	29.02	28.82	28.70	28.48	28.45

Table 2. IMDeception performance parameters, *Averaged on Div2K Validation Set on NVIDIA 2080 Super

- Adam optimizer with $\beta_1 = 0.9, \beta_2 = 0.999$.
- Mini-batch size of 8.
- Cropped HR Image Size of 512 x 512.
- Zero padding is used when necessary.
- Each epoch contains 800 mini batches.
- Knee Learning Rate Scheduling [13] with 10 epochs warm-up 400 epochs exploit and 400 epochs cooldown period with maximum learning rate $0.5e^{-3}$ (Fig. 4b).
- PyTorch model is trained within PyTorch-Lightning framework.

For the first phase, we used Charbonnier loss with $\epsilon = 0.1$ as in Eq. (5) and trained for 2000 epochs, which lasted 2 days and 7 hours on a single NVIDIA Tesla v100. See Fig. 4 for training curves and learning rate policy.

$$Charbonnier(x) = \sqrt{x^2 + \epsilon^2} \tag{5}$$

The second phase of the training started from the best checkpoint, and this time L2-norm was used as the loss function, and trained for 1300 epochs which lasted 1 day and 16 hours.

4.3. Results

In this section, the proposed architecture's PSNR results are given on various different datasets. The PSNR results of IMDeception and other state-of-the-art methods can be seen in Tab. 3. From the experiments it can be seen that, although IMDeception (core = 16 + NLA) has very limited number of parameters and FLOPS, it has on par performance with state-of-the-art algorithms. Especially, IMDeception's performance on Urban100, Manga109 datasets is well above E-RFDN, IMDN, CARN, LapSRN methods except EDSR (which has 43M parameters). An interesting result is IMDeception (core = 4)'s PSNR performance on Urban100 and Manga109 surpasses LapSRN although it has only 7% of number of parameters. The number-ofparameters and PSNR results of these methods can be best seen in Fig. 5



Figure 4. Training Curves

Another important property that IMDeception has, its precise output on the repeated structures and patterns which can be seen in Fig. 6.

In terms of run-time, our proposed method has great potential for optimization on edge devices [7], due to parallel grouped convolutions and a reduced number of parameters. As it can be seen from the Tab. 3, the proposed model defines a set of efficient architectures, which can be used in different devices with different inference run-times with good reconstruction performance. As a reference and as an indication of its potential, we run our proposed models on NVIDIA RTX 2080 Super and NVIDIA Jetson Xavier AGX 30W devices. To do this, we have converted trained models to ONNX format and used NVIDIA's TRT Engine application to convert them to an inference engine to use

Model	Set5	Set14	BSD100	Urban100	Manga109	Div2K (Val)
Bicubic	28.42	26.00	25.96	23.14	24.89	26.66
LapSRN (812K)	31.54	28.19	27.32	25.21	29.09	28.75
IMDN (779K)	32.21	28.58	27.56	26.04	30.45	28.94
EDSR (43M)	32.46	28.80	27.71	26.64	31.02	29.25
CARN [2] (1.5M)	32.14	28.61	27.58	26.07	30.46	28.96
E-RFDN (433K)	32.16	28.65	27.60	26.15	30.59	29.04
IMDeception (316K)	22.14	20 61	27.60	26.20	20.67	20.02*
core=16 + NLA	32.14	28.04	27.00	20.20	50.07	29.02
IMDeception(198K)	31.83	28 47	27 50	25.82	30.28	28.83
core=12	51.05	20.47	27.50	25.62	50.20	20.05
IMDeception(113K)	31.60	28 35	27 42	25.60	20.80	28.69
core=8	51.07	20.55	27.72	25.00	27.07	20.07
IMDeception(57K)	31 35	28.12	77 77	25.23	29.20	28.45
core=4	51.55	20.12	21.21	23.23	27.20	20.43
IMDeception(57K)	31 42	28.18	27.26	25.26	29.28	28.48
core=4 + NLA	51.42	20.10	27.20	23.20	27.20	20.40

Table 3. Experimental results of the proposed method were compared with various different methods for x4 scaling. Note that except for Div2K Validation result, all PSNR values are calculated on Luminance (Y) channel to be consistent with the literature. * Div2K Test Set PSNR is 28.73



Figure 5. PSNR vs Number of Parameters for Urban100 Dataset for Different Super-Resolution Models

the hardware's full potential. The run-times are listed in Tab. 4. Note that IMDeception can run on this edge device at up to 24fps while outputting high-resolution 2K images. An important conclusion that can be made from the run-time experiments although the number of parameters and FLOPS are lower for IMDeception core = 12 compared to core = 16 + NLA, the inference run-times are higher. This is due to the fact that GPUs are usually optimized for the channel sizes, which are powers of 2. Because 12 is not a power of 2, additional processing in the GPU is required, negating the benefits of the reduced number of parameters and FLOPS. This is an important conclusion to make since this phenomenon is not observable during inference with a training framework such as PyTorch.

Model Hardware	RTX 2080	Jetson Xavier	
Widder I laidware	TensorRT (ms)	TensorRT (ms)	
IMDeception	177	88.9	
core=16+NLA	17.7		
IMDeception*	10	08.0	
core=12	19	90.9	
IMDeception	0.0	51.1	
core=8	9.9	51.1	
IMDeception	0.8	117	
core=4+NLA	9.0	44./	
IMDeception	0.2	41.0	
core=4	9.2	41.9	

Table 4. Inference Results for an input image size of 512x256 * Note the increased inference time. This is because 12 is not a power of 2 and the GPUs are optimized and have kernels specific to sizes of power of 2.

5. Conclusion

We proposed an efficient model based on the IMDN network called IMDeception. IMDeception employs the proposed Global Progressive Refinement Module (GPRM), which is an extension of the Progressive Refinement Module (PRM). Unlike PRM that works only with the Conv2d layer at the local scale, GPRM can be used with any arbitrary block, as we did with the newly proposed Grouped Information Distilling Blocks (GIDB). Both of the proposed mechanisms/blocks can be used for different networks and in different structures. These structures are designed with



(a) Original: BSD100 8023.png



(b) Bicubic



(c) CARN



(e) IMDN



(f) Original: Urban100 img_048.png

(h) CARN

(j) IMDN

Figure 6. Example Images from Reference Datasets. IMDeception is compared with other methods

efficiency in mind, which reduces the number of parameters and FLOPS while maintaining high performance. GPRM is an efficient way of combining features and can be an alternative to Dense Networks style or IIC-style feature aggregating methods. One nice feature of it is that it separates the aggregated part from the distilled part, which helps controlling the network size while maintaining network performance. On the other hand, GIDB uses grouped convolutions, which, if implemented with efficiency in mind, can provide a speed boost during inference. We also showed that the proposed model is very high-performing on various different datasets and has great inference timings on different hardware, including NVIDIA Jetson Xavier AGX.

References

- Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, July 2017. 5
- [2] Namhyuk Ahn, Byungkon Kang, and Kyung-Ah Sohn. Fast, accurate, and lightweight super-resolution with cascading residual network. In *Eur. Conf. Comput. Vis.* 7
- [3] Mustafa Ayazoglu. Extremely lightweight quantization robust real-time single-image super resolution for mobile devices. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 2472–2479, 2021. 2, 3
- [4] Ming Zhuo Chen and Jun Ming Wu. Group feature information distillation network for single image super-resolution. In 2021 7th International Conference on Computer and Communications (ICCC), pages 1827–1831, 2021. 2
- [5] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. volume 8692, pages 184–199, 2014. 1, 2
- [6] Chao Dong, Chen Change Loy, and Xiaoou Tang. Accelerating the super-resolution convolutional neural network. *Eur. Conf. Comput. Vis.*, 9906:391–407, 2016. 1, 2
- [7] Perry Gibson, José Cano, Jack Turner, Elliot J. Crowley, Michael O'Boyle, and Amos Storkey. Optimizing grouped convolutions on edge devices, 2020. 3, 6
- [8] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Int. Conf. Comput. Vis.*, pages 1398–1406, 2017. 3
- [9] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In Adv. Neural Inform. Process. Syst., 2015. 3
- [10] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2261–2269, 2017. 2
- [11] Zheng Hui, Xinbo Gao, Yunchu Yang, and Xiumei Wang. Lightweight image super-resolution with information multidistillation network. In *Proceedings of the 27th ACM International Conference on Multimedia*. ACM, oct 2019. 2, 3
- [12] Zheng Hui, Xiumei Wang, and Xinbo Gao. Fast and accurate single image super-resolution via information distillation network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 723–731, 2018. 2
- [13] Nikhil Iyer, V. Thejas, Nipun Kwatra, Ramachandran Ramjee, and Muthian Sivathanu. Wide-minima density hypothesis and the explore-exploit learning rate schedule. *CoRR*, abs/2003.03977, 2020. 6
- [14] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeplyrecursive convolutional network for image super-resolution, 2015. 2
- [15] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1646–1654, 2016. 1, 2
- [16] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and

accurate super-resolution. In IEEE Conf. Comput. Vis. Pattern Recog., 2017. 2

- [17] Yawei Li, Shuhang Gu, Luc Van Gool, and Radu Timofte. Learning filter basis for convolutional neural network compression. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. 3
- [18] Yawei Li, Shuhang Gu, Kai Zhang, Luc Van Gool, and Radu Timofte. DHP: differentiable meta pruning via hypernetworks. In *Eur. Conf. Comput. Vis.*, 2020. 3
- [19] Yawei Li, Wen Li, Martin Danelljan, Kai Zhang, Shuhang Gu, Luc Van Gool, and Radu Timofte. The heterogeneity hypothesis: Finding layer-wise differentiated network architectures. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2144–2153, 2021. 3
- [20] Yawei Li, Kai Zhang, Luc Van Gool, Radu Timofte, et al. Ntire 2022 challenge on efficient super-resolution: Methods and results. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2022. 2, 3
- [21] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, July 2017. 2, 3, 5
- [22] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network, 2014. 2
- [23] Jie Liu, Jie Tang, and Gangshan Wu. Residual feature distillation network for lightweight image super-resolution. In *Eur. Conf. Comput. Vis.* 3
- [24] Ben Niu, Weilei Wen, Wenqi Ren, Xiangde Zhang, Lianping Yang, Shuzhen Wang, Kaihao Zhang, Xiaochun Cao, and Haifeng Shen. Single image super-resolution via a holistic attention network. *Int. Conf. Comput. Vis.*, 2020. 2
- [25] Tao Sheng, Chen Feng, Shaojie Zhuo, Xiaopeng Zhang, Liang Shen, and Mickey Aleksic. A quantizationfriendly separable convolution for mobilenets. *CoRR*, abs/1803.08607, 2018. 3
- [26] Wenzhe Shi, Jose Caballero, Ferenc Huszar, 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. *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1874–1883, 2016. 1
- [27] Dehua Song, Chang Xu, Xu Jia, Yiyi Chen, Chunjing Xu, and Yunhe Wang. Efficient residual dense block search for image super-resolution. *Proceedings of the AAAI Conference* on Artificial Intelligence, 34(07):12007–12014, Apr. 2020. 3
- [28] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1–9, 2015. 2, 3
- [29] Ying Tai, Jian Yang, and Xiaoming Liu. Image superresolution via deep recursive residual network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017. 2
- [30] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks, 2017. 2, 3
- [31] Yan Wu, Zhiwu Huang, Suryansh Kumar, Rhea Sanjay Sukthanker, Radu Timofte, and Luc Van Gool. Trilevel neural ar-

chitecture search for efficient single image super-resolution. *CoRR*, abs/2101.06658, 2021. **3**

- [32] Jiahui Yu, Yuchen Fan, Jianchao Yang, Ning Xu, Zhaowen Wang, Xinchao Wang, and Thomas S. Huang. Wide activation for efficient and accurate image super-resolution. *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, 2018. 2
- [33] Kai Zhang, Martin Danelljan, Yawei Li, Radu Timofte, et al. Aim 2020 challenge on efficient super-resolution: Methods and results. In *Computer Vision – ECCV 2020 Workshops*, pages 5–40, 2020. 2, 3
- [34] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *Eur. Conf. Comput. Vis.*, pages 294–310, 2018. 2
- [35] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. In *IEEE Conf. Comput. Vis. Pattern Recog.* 2
- [36] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition. *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8697–8710, 2018. 3