

# Exploring Sparsity in Image Super-Resolution for Efficient Inference

Longguang Wang<sup>1</sup>, Xiaoyu Dong<sup>2,3</sup>, Yingqian Wang<sup>1</sup>, Xinyi Ying<sup>1</sup>, Zaiping Lin<sup>1</sup>, Wei An<sup>1</sup>, Yulan Guo<sup>1\*</sup>

<sup>1</sup>National University of Defense Technology <sup>2</sup>The University of Tokyo <sup>3</sup>RIKEN AIP

{wanglongguang15, yulan.guo}@nudt.edu.cn

## Abstract

Current CNN-based super-resolution (SR) methods process all locations equally with computational resources being uniformly assigned in space. However, since missing details in low-resolution (LR) images mainly exist in regions of edges and textures, less computational resources are required for those flat regions. Therefore, existing CNN-based methods involve redundant computation in flat regions, which increases their computational cost and limits their applications on mobile devices. In this paper, we explore the sparsity in image SR to improve inference efficiency of SR networks. Specifically, we develop a Sparse Mask SR (SMSR) network to learn sparse masks to prune redundant computation. Within our SMSR, spatial masks learn to identify “important” regions while channel masks learn to mark redundant channels in those “unimportant” regions. Consequently, redundant computation can be accurately localized and skipped while maintaining comparable performance. It is demonstrated that our SMSR achieves state-of-the-art performance with 41%/33%/27% FLOPs being reduced for  $\times 2/3/4$  SR. Code is available at: <https://github.com/LongguangWang/SMSR>.

## 1. Introduction

The goal of single image super-resolution (SR) is to recover a high-resolution (HR) image from a single low-resolution (LR) observation. Due to the powerful feature representation and model fitting capabilities of deep neural networks, CNN-based SR methods have achieved significant performance improvements over traditional ones. Recently, many efforts have been made towards real-world applications, including few-shot SR [38, 39], blind SR [12, 49, 42], and scale-arbitrary SR [15, 43]. With the popularity of intelligent edge devices (such as smartphones and VR glasses), performing SR on these devices is highly demanded. Due to the limited resources of edge devices<sup>1</sup>, efficient SR is crucial to the applications on these devices.

<sup>1</sup>For example, the computational performance of Kirin 990 and RTX 2080Ti are 0.9 and 13.4 tFLOPs, respectively.

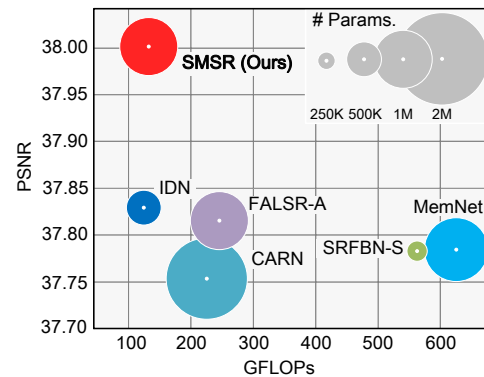


Figure 1. Trade-off between PSNR performance, number of parameters and FLOPs. Results are achieved on Set5 for  $\times 2$  SR.

Since the pioneering work of SRCNN [8], deeper networks have been extensively studied for image SR. In VDSR [19], SR network is first deepened to 20 layers. Then, a very deep and wide architecture with over 60 layers is introduced in EDSR [29]. Later, Zhang *et al.* further increased the network depth to over 100 and 400 in RDN [51] and RCAN [50], respectively. Although a deep network usually improves SR performance, it also leads to high computational cost and limits the applications on mobile devices. To address this problem, several efforts have been made to reduce model size through information distillation [17] and efficient feature reuse [2]. Nevertheless, these networks still involve redundant computation. Compared to an HR image, missing details in its LR image mainly exist in regions of edges and textures. Consequently, less computational resources are required in those flat regions. However, these CNN-based SR methods process all locations equally, resulting in redundant computation within flat regions.

In this paper, we explore the sparsity in image SR to improve inference efficiency of SR networks. We first study the intrinsic sparsity of the image SR task and then investigate the feature sparsity in existing SR networks. To fully exploit the sparsity for efficient inference, we propose a sparse mask SR (SMSR) network to dynamically skip redundant computation at a fine-grained level. Our SMSR learns spatial masks to identify “important” regions (*e.g.*, edge and texture regions) and uses channel masks to mark redundant channels in those “unimportant” regions. These

two kinds of masks work jointly to accurately localize redundant computation. During network training, we soften these binary masks using the Gumbel softmax trick to make them differentiable. During inference, we use sparse convolution to skip redundant computation. It is demonstrated that our SMSR can effectively localize and prune redundant computation to achieve better efficiency while producing promising results (Fig. 1).

Our main contributions can be summarized as: 1) We develop an SMSR network to dynamically skip redundant computation for efficient image SR. In contrast to existing works that focus on lightweight network designs, we explore a different route by pruning redundant computation to improve inference efficiency. 2) We propose to localize redundant computation by learning spatial and channel masks. These two kinds of masks work jointly for fine-grained localization of redundant computation. 3) Experimental results show that our SMSR achieves state-of-the-art performance with better inference efficiency. For example, our SMSR outperforms previous methods on Set14 for  $\times 2$  SR with a significant speedup on mobile devices (Table 2).

## 2. Related Work

In this section, we first review several major works for CNN-based single image SR. Then, we discuss CNN acceleration techniques related to our work, including adaptive inference and network pruning.

**Single Image SR.** CNN-based methods have dominated the research of single image SR due to their strong representation and fitting capabilities. Dong *et al.* [8] first introduced a three-layer network to learn an LR-to-HR mapping for single image SR. Then, a deep network with 20 layers was proposed in VDSR [19]. Recently, deeper networks are extensively studied for image SR. Lim *et al.* [29] proposed a very deep and wide network (namely, EDSR) by cascading modified residual blocks. Zhang *et al.* [51] further combined residual learning and dense connection to build RDN with over 100 layers. Although these networks achieve state-of-the-art performance, the high computational cost and memory footprint limit their applications on mobile devices.

To address this problem, several lightweight networks were developed [22, 17, 2]. Specifically, distillation blocks were proposed for feature learning in IDN [17], while a cascading mechanism was introduced to encourage efficient feature reuse in CARN [2]. Different from these manually designed networks, Chu *et al.* [6] developed a compact architecture using neural architecture search (NAS). Recently, Lee *et al.* [24] introduced a distillation framework to leverage knowledge learned by powerful teacher SR networks to boost the performance of lightweight student SR networks. Although these lightweight SR networks successfully reduce the model size, redundant computation is still involved and hinders them to achieve better computational efficiency.

In contrast to many existing works that focus on compact architecture designs, few efforts have been made to exploit the redundancy in SR networks for efficient inference.

**Adaptive Inference.** Adaptive inference techniques [44, 37, 36, 11, 26] have attracted increasing interests since they can adapt the network structure according to the input. One active branch of adaptive inference techniques is to dynamically select an inference path at the levels of layers. Specifically, Wu *et al.* [45] proposed a BlockDrop approach for ResNets to dynamically drop several residual blocks for efficiency. Mullapudi *et al.* [36] proposed an HydraNet with multiple branches and used a gating approach to dynamically choose a set of them at test time. Another popular branch is early stopping techniques that skip the computation at a location whenever it is deemed to be unnecessary [46]. On top of ResNets, Figurnov *et al.* [9] proposed a spatially adaptive computation time (SACT) mechanism to stop computation for a spatial position when the features become “good enough”. Liu *et al.* [31] introduced adaptive inference for SR by producing a map of local network depth to adapt the number of convolutional layers implemented at different locations. However, these adaptive inference methods only focus on spatial redundancy without considering redundancy in channel dimension.

**Network Pruning.** Network pruning [13, 32, 33] is widely used to remove a set of redundant parameters for network acceleration. As a popular branch of network pruning methods, structured pruning approaches are usually used to prune the network at the level of channels and even layers [25, 32, 33, 14]. Specifically, Li *et al.* [25] used  $L_1$  norm to measure the importance of different filters and then pruned less important ones. Liu *et al.* [32] imposed a sparsity constraint on scaling factors of the batch normalization layers and identified channels with lower scaling factors as less informative ones. Different from these static structured pruning methods, Lin *et al.* [30] conducted runtime neural network pruning according to the input image. Recently, Gao *et al.* [10] introduced a feature boosting and suppression method to dynamically prune unimportant channels at inference time. Nevertheless, these network pruning methods treat all spatial locations equally without taking their different importance into consideration.

## 3. Sparsity in Image Super-Resolution

In this section, we first illustrate the intrinsic sparsity of the single image SR task and then investigate the feature sparsity in state-of-the-art SR networks.

Given an HR image  $I^{HR}$  and its LR version  $I^{LR}$  (e.g.,  $\times 4$  downsampled), we super-resolve  $I^{LR}$  using Bicubic and RCAN to obtain  $I_{Bicubic}^{SR}$  and  $I_{RCAN}^{SR}$ , respectively. Figure 2 shows the absolute difference between  $I_{Bicubic}^{SR}$ ,  $I_{RCAN}^{SR}$  and  $I^{HR}$  in the luminance channel. It can be observed from Fig. 2(b) that  $I_{Bicubic}^{SR}$  is “good enough” for flat

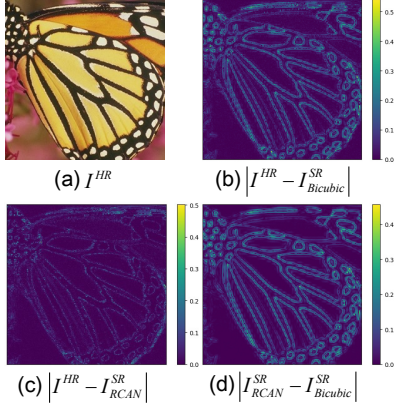


Figure 2. Absolute difference between  $I_{Bicubic}^{SR}$ ,  $I_{RCAN}^{SR}$  and  $I^{HR}$  in the luminance channel.

regions, with noticeable missing details in only a small proportion of regions ( $\sim 17\%$  pixels with  $|I^{HR} - I_{Bicubic}^{SR}| > 0.1$ ). That is, the SR task is intrinsically sparse in spatial domain. Compared to Bicubic, RCAN performs better in edge regions while achieving comparable performance in flat regions (Fig. 2(c)). Although RCAN focuses on recovering high-frequency details in edge regions (Fig. 2(d)), those flat regions are equally processed at the same time. Consequently, redundant computation is involved.

Figure 3 illustrates the feature maps after the ReLU layer in a backbone block of RCAN. It can be observed that the spatial sparsity varies significantly for different channels. Moreover, a considerable number of channels are quite sparse (sparsity  $\geq 0.8$ ), with only edge and texture regions being activated. That is, computation in those flat regions is redundant since these regions are not activated after the ReLU layer. In summary, RCAN activates only a few channels for “unimportant” regions (e.g., flat regions) and more channels for “important” regions (e.g., edge regions). More results achieved with different SR networks and backbone blocks are provided in the supplemental material.

Motivated by these observations, we learn sparse masks to localize and skip redundant computation for efficient inference. Specifically, our spatial masks dynamically identify “important” regions while the channel masks mark redundant channels in those “unimportant” regions. Compared to network pruning methods [10, 30, 14], we take region redundancy into consideration and only prune channels for “unimportant” regions. Different from adaptive inference networks [37, 27], we further investigate the redundancy in channel dimension to localize redundant computation at a finer-grained level.

## 4. Our SMSR Network

Our SMSR network uses sparse mask modules (SMM) to prune redundant computation for efficient image SR.

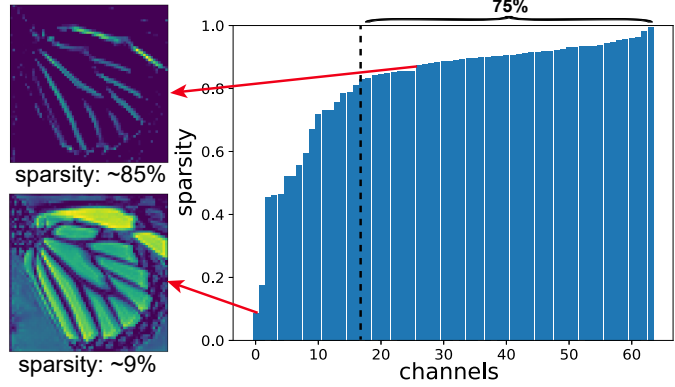


Figure 3. Visualization of feature maps after the ReLU layer in the first backbone block of RCAN. Note that, sparsity is defined as the ratio of zeros in the corresponding channels.

Within each SMM, spatial and channel masks are first generated to localize redundant computation, as shown in Fig. 4. Then, the redundant computation is dynamically skipped using  $L$  densely-connected sparse mask convolutions. Since only necessary computation is performed, our SMSR can achieve better efficiency while maintaining comparable performance.

### 4.1. Sparse Mask Generation

#### 1) Training Phase

**Spatial Mask.** The goal of spatial mask is to identify “important” regions in feature maps (i.e., 0 for “unimportant” regions and 1 for “important” ones). To make the binary spatial mask learnable, we use Gumbel softmax distribution to approximate the one-hot distribution [18]. Specifically, input feature  $F \in \mathbb{R}^{C \times H \times W}$  is first fed to an hourglass block to produce  $F^{spa} \in \mathbb{R}^{2 \times H \times W}$ , as shown in Fig. 5(a). Then, the Gumbel softmax trick is used to obtain a softened spatial mask  $M_k^{spa} \in \mathbb{R}^{H \times W}$ :

$$M_k^{spa}[x, y] = \frac{\exp\left(\left(F^{spa}[1, x, y] + G_k^{spa}[1, x, y]\right)/\tau\right)}{\sum_{i=1}^2 \exp\left(\left(F^{spa}[i, x, y] + G_k^{spa}[i, x, y]\right)/\tau\right)}, \quad (1)$$

where  $x, y$  are vertical and horizontal indices,  $G_k^{spa} \in \mathbb{R}^{2 \times H \times W}$  is a Gumbel noise tensor with all elements following Gumbel(0, 1) distribution and  $\tau$  is a temperature parameter. When  $\tau \rightarrow \infty$ , samples from Gumbel softmax distribution become uniform. That is, all elements in  $M_k^{spa}$  are 0.5. When  $\tau \rightarrow 0$ , samples from Gumbel softmax distribution become one-hot. That is,  $M_k^{spa}$  becomes binary. In practice, we start at a high temperature and anneal to a small one to obtain binary spatial masks.

**Channel Mask.** In addition to spatial masks, channel masks are used to mark redundant channels in those “unimportant” regions (i.e., 0 for redundant channels and 1 for preserved ones). Here, we also use Gumbel softmax trick

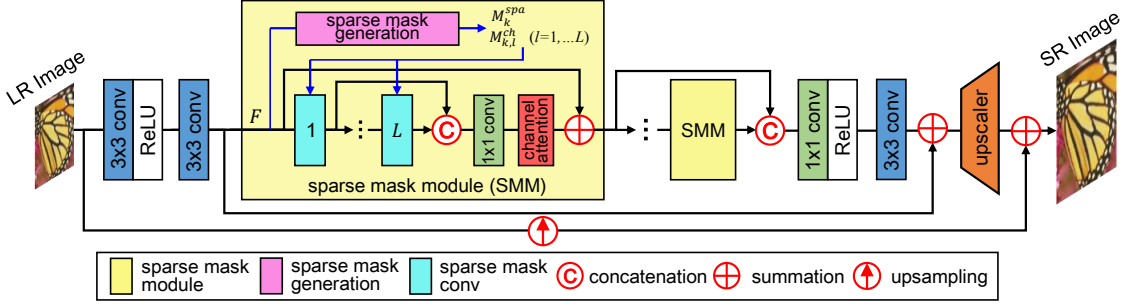


Figure 4. An overview of our SMSR network.

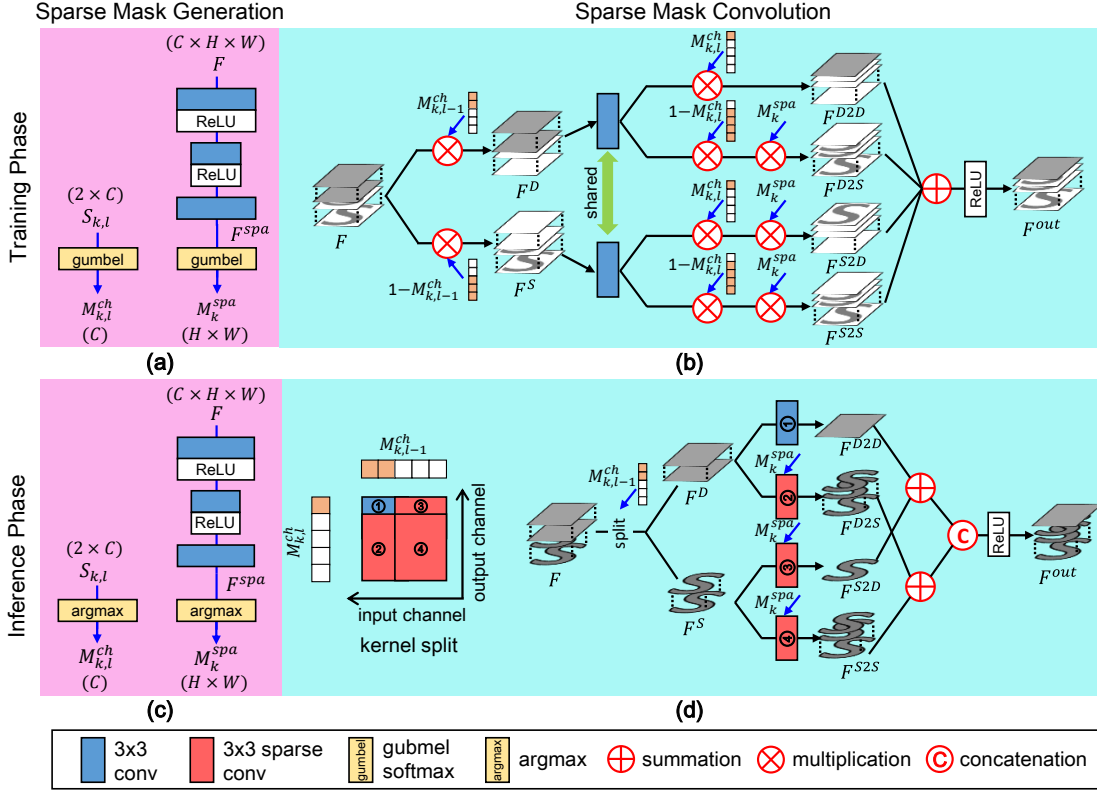


Figure 5. An illustration of sparse mask generation and sparse mask convolution.

to produce binary channel masks. For the  $l^{\text{th}}$  convolutional layer in the  $k^{\text{th}}$  SMM, we feed auxiliary parameter  $S_{k,l} \in \mathbb{R}^{2 \times C}$  to a Gumbel softmax layer to generate softened channel masks  $M_{k,l}^{ch} \in \mathbb{R}^C$ :

$$M_{k,l}^{ch}[c] = \frac{\exp\left((S_{k,l}[1, c] + G_{k,l}^{ch}[1, c])/\tau\right)}{\sum_{i=1}^2 \exp\left((S_{k,l}[i, c] + G_{k,l}^{ch}[i, c])/\tau\right)}, \quad (2)$$

where  $c$  is the channel index and  $G_{k,l}^{ch} \in \mathbb{R}^{2 \times C}$  is a Gumbel noise tensor. In our experiments,  $S_{k,l}$  is initialized using random values drawn from a Gaussian distribution  $\mathcal{N}(0, 1)$ .

**Sparsity Regularization.** Based on spatial and channel masks, we define a sparsity term  $\eta_{k,l}$ :

$$\eta_{k,l} = \frac{1}{C \times H \times W} \sum_{c,x,y} \left( (1 - M_{k,l}^{ch}[c]) \times M_k^{spa}[x, y] + M_{k,l}^{ch}[c] \times I[x, y] \right), \quad (3)$$

where  $I \in \mathbb{R}^{H \times W}$  is a tensor with all ones. Note that,  $\eta_{k,l}$  represents the ratio of activated locations in the output feature maps. To encourage the output features to be more sparse with fewer locations being activated, we further introduce a sparsity regularization loss:

$$L_{reg} = \frac{1}{K \times L} \sum_{k,l} \eta_{k,l}, \quad (4)$$

where  $K$  is the number of SMMs and  $L$  is the number of sparse mask convolutional layers within each SMM.

**Training Strategy.** During the training phase, the temperature parameter  $\tau$  in Gumbel softmax layers is annealed us-

Table 1. Comparative results achieved on Set14 by our SMSR with different settings for  $\times 2$  SR.

Model	Spatial Mask	Channel Mask	Conv	#Params.	Sparsity	FLOPs	PSNR	SSIM
1	✗	✗	Vanilla	926K	0	1.00×	33.65	0.9180
2	✗	✓	Vanilla	587K	0.46	0.60×	33.53	0.9169
3	✓	✗	Sparse	985K	0.42	0.65×	33.60	0.9176
4 (Ours)	✓	✓	Sparse	985K	0.46	0.61×	33.64	0.9179

ing the schedule  $\tau = \max(0.4, 1 - \frac{t}{T_{temp}})$ , where  $t$  is the number of epochs and  $T_{temp}$  is empirically set to 500 in our experiments. As  $\tau$  gradually decreases, Gumbel softmax distribution is forced to approach an one-hot distribution to produce binary spatial and channel masks.

## 2) Inference Phase

During training, Gumbel softmax distributions are forced to approach one-hot distributions as  $\tau$  decreases. Therefore, we replace the Gumbel softmax layers with argmax layers after training to obtain binary spatial and channel masks, as shown in Fig. 5(c).

## 4.2. Sparse Mask Convolution

### 1) Training Phase

To enable backpropagation of gradients at all locations, we do not explicitly perform sparse convolution during training. Instead, we multiply the results of a vanilla “dense” convolution with predicted spatial and channel masks, as shown in Fig. 5(b). Specifically, input feature  $F$  is first multiplied with  $M_{k,l-1}^{ch}$  and  $(1 - M_{k,l-1}^{ch})$  to obtain  $F^D$  and  $F^S$ , respectively. That is, channels with “dense” and “sparse” feature maps in  $F$  are separated. Next,  $F^D$  and  $F^S$  are passed to two convolutions with shared weights. The resulting features are then multiplied with different combinations of  $(1 - M_k^{ch})$ ,  $M_k^{ch}$  and  $M_k^{spa}$  to activate different parts of the features. Finally, all these features are summed up to generate the output feature  $F^{out}$ . Thanks to Gumbel softmax trick used in mask generation, gradients at all locations can be preserved to optimize the kernel weights of convolutional layers.

### 2) Inference Phase

During the inference phase, sparse convolution is performed based on the predicted spatial and channel masks, as shown in Fig. 5(d). Take the  $l^{\text{th}}$  layer in the  $k^{\text{th}}$  SMM as an example, its kernel is first splitted into four sub-kernels according to  $M_{k,l-1}^{ch}$  and  $M_{k,l}^{ch}$  to obtain four convolutions. Meanwhile, input feature  $F$  is splitted into  $F^D$  and  $F^S$  based on  $M_{k,l-1}^{ch}$ . Then,  $F^D$  is fed to convolutions ① and ② to produce  $F^{D2D}$  and  $F^{D2S}$ , while  $F^S$  is fed to convolutions ③ and ④ to produce  $F^{S2D}$  and  $F^{S2S}$ . Note that,  $F^{D2D}$  is produced by a vanilla “dense” convolution while  $F^{D2S}$ ,  $F^{S2D}$  and  $F^{S2S}$  are generated by sparse convolutions with only “important” regions (marked by  $M_k^{spa}$ ) being computed. Finally, features obtained from these four branches are summed and concatenated to produce the output feature  $F^{out}$ . Using sparse mask convolution, compu-

tation for redundant channels within those “unimportant” regions can be skipped for efficient inference.

## 4.3. Discussion

Different from many recent works that use lightweight network designs [17, 2, 6] or knowledge distillation [24] for efficient SR, we speedup SR networks by pruning redundant computation. Previous adaptive inference and network pruning methods focus on redundant computation in spatial and channel dimensions independently. Directly applying these approaches cannot fully exploit the redundancy in SR networks and suffers notable performance drop, as demonstrated in Sec. 5.2. In contrast, our SMSR provides a unified framework to consider redundancy in both spatial and channel dimensions. It is demonstrated that our spatial and channel masks are well compatible to each other and facilitate our SMSR to obtain fine-grained localization of redundant computation.

## 5. Experiments

### 5.1. Implementation Details

We used 800 training images and 100 validation images from the DIV2K dataset [1] as training and validation sets. For evaluation, we used five benchmark datasets including Set5 [4], Set14 [48], B100 [34], Urban100 [16], and Manga109 [35]. Peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) were used as evaluation metrics to measure SR performance. Following the evaluation protocol in [50, 51], we cropped borders and calculated the metrics in the luminance channel.

During training, 16 LR patches of size  $96 \times 96$  and their corresponding HR patches were randomly cropped. Data augmentation was then performed through random rotation and flipping. We set  $C = 64$ ,  $L = 4$ ,  $K = 5$  for our SMSR. We used the Adam method [21] with  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$  for optimization. The initial learning rate was set to  $2 \times 10^{-4}$  and reduced to half after every 200 epochs. The training was stopped after 1000 epochs. The overall loss for training is defined as  $L = L_{SR} + \lambda L_{reg}$ , where  $L_{SR}$  is the  $L_1$  loss between SR results and HR images,  $L_{reg}$  is defined in Eq. 4. To maintain training stability, we used a warmup strategy  $\lambda = \lambda_0 \times \min(\frac{t}{T_{warm}}, 1)$ , where  $t$  is the number of epochs,  $T_{warm}$  is empirically set to 50 and  $\lambda_0$  is set to 0.1.

### 5.2. Model Analysis

We first conduct experiments to demonstrate the effectiveness of sparse masks. Then, we investigate the effect of

Table 2. Comparative results achieved on Set14 by our SMSR with different sparsities for  $\times 2$  SR.

Model	Conv	$\lambda_0$	Sparsity	#Params.	FLOPs	Memory	Time		PSNR	SSIM
							GPU	CPU		
baseline	Vanilla	0	0	926K	1.00 $\times$	1.00 $\times$	1.00 $\times$	1.00 $\times$	33.65	0.9180
5	Sparse	0.1	0.46	985K	0.61 $\times$	0.89 $\times$	1.22 $\times$	0.64 $\times$	33.64	0.9179
6	Sparse	0.2	0.64	985K	0.46 $\times$	0.87 $\times$	1.11 $\times$	0.55 $\times$	33.61	0.9174
7	Sparse	0.3	0.73	985K	0.38 $\times$	0.85 $\times$	1.04 $\times$	0.54 $\times$	33.52	0.9169
IDN [17]	-	-	-	553K	0.57 $\times$	0.91 $\times$	1.04 $\times$	0.71 $\times$	33.30	0.9148
CARN [2]	-	-	-	1592K	0.99 $\times$	1.01 $\times$	1.00 $\times$	0.96 $\times$	33.52	0.9166
FALSR-A [6]	-	-	-	1021K	1.04 $\times$	2.02 $\times$	1.11 $\times$	1.05 $\times$	33.55	0.9168

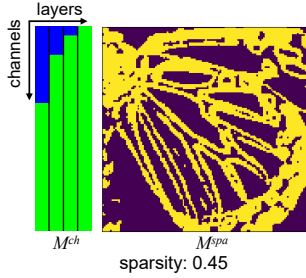


Figure 6. Visualization of sparse masks. Blue and green regions in  $M^{ch}$  represent channels with “dense” and “sparse” feature maps, respectively. In  $M^{spa}$ , “important” locations are shown in yellow.

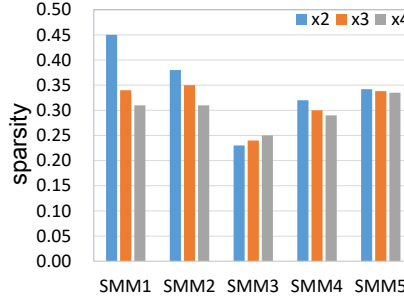


Figure 7. Comparison of sparsities achieved in different SMMs on butterfly for different scale factors.

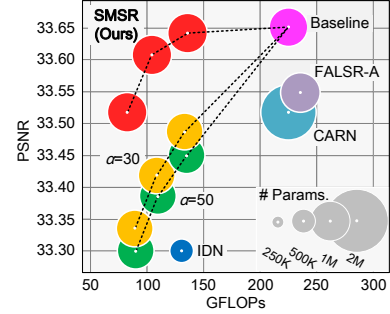


Figure 8. Comparison between learning-based masks (red circles) and gradient-based masks (yellow and green circles) on Set14.

sparsity and visualize sparse masks for discussion. Finally, we compare our learning-based masks with heuristic ones.

**Effectiveness of Sparse Masks.** To demonstrate the effectiveness of our sparse masks, we first introduced variant 1 by removing both spatial and channel masks. Then, we developed variants 2 and 3 by adding channel masks and spatial masks, respectively. Comparative results are shown in Table 1. Without spatial and channel masks, all locations and all channels are processed equally. Therefore, variant 1 has a high computational cost. Using channel masks, redundant channels are pruned at all spatial locations. Therefore, variant 2 can be considered as a pruned version of variant 1. Although variant 2 has fewer parameters and FLOPs, it suffers a notable performance drop (33.53 vs. 33.65) since beneficial information in “important” regions of these pruned channels are discarded. With only spatial masks, variant 3 suffers from a conflict between efficiency and performance since redundant computation in channel dimension cannot be well handled. Consequently, its FLOPs is reduced with a performance drop (33.60 vs. 33.65). Using both spatial and channel masks, our SMSR can effectively localize and skip redundant computation at a finer-grained level to reduce FLOPs by 39% while maintaining comparable performance (33.64 vs. 33.65).

**Effect of Sparsity.** To investigate the effect of sparsity, we retrained our SMSR with large  $\lambda_0$  to encourage high sparsity. Nvidia RTX 2080Ti, Intel I9-9900K and Kirin 990/810 were used as platforms of GPU, CPU and mobile processor for evaluation. For fair comparison of memory consumption and inference time, all convolutional layers in the backbone of different networks were implemented using im2col [5]

based convolutions since different implementation methods (e.g., Winograd [23] and FFT [41]) have different computational costs. Comparative results are presented in Table 2.

As  $\lambda_0$  increases, our SMSR produces higher sparsities with more FLOPs and memory consumption being reduced. Further, our network also achieves significant speedup on CPU and mobile processors. Due to the irregular and fragmented memory patterns, sparse convolution cannot make full use of the characteristics of general GPUs (e.g., memory coalescing) and relies on specialized designs to improve memory locality and cache hit rate for acceleration [47]. Therefore, the advantage of our SMSR cannot be fully exploited on GPUs without specific optimization. Compared to other state-of-the-art methods, our SMSR (variant 5) obtains better performance with lower memory consumption and shorter inference time on mobile processors. This clearly demonstrates the great potential of our SMSR for applications on mobile devices.

**Visualization of Sparse Masks.** We visualize the sparse masks generated in the first SMM for  $\times 2$  SR in Fig. 6. More results are provided in the supplemental material. It can be seen that locations around edges and textures in  $M^{spa}$  are considered as “important” ones, which is consistent with our observations in Sec. 3. Moreover, we can see that there are more sparse channels (i.e., green regions in  $M^{ch}$ ) in deep layers than shallow layers. This means that a subset of channels in shallow layers are informative enough for “unimportant” regions and our network progressively focuses more on “important” regions as the depth increases. Overall, our spatial and channel masks work jointly for fine-grained localization of redundant computation.

Table 4. Comparative results achieved for  $\times 2/3/4$  SR. PSNR/SSIM results of previous works are directly copied from corresponding papers. FLOPs is computed based on HR images with a resolution of 720p ( $1280 \times 720$ ). For SMSR, average sparsities on all datasets (0.49/0.39/0.33 for  $\times 2/3/4$  SR) are used to calculate FLOPs, with full FLOPs being shown in brackets. Best and second best results are **highlighted** and underlined.

Model	Scale	#Params	FLOPs	Set5	Set14	B100	Urban100	Manga109
Bicubic	$\times 2$	-	-	33.66/0.9299	30.24/0.8688	29.56/0.8431	26.88/0.8403	30.80/0.9339
SRCNN [8]	$\times 2$	57K	52.7G	36.66/0.9542	32.45/0.9067	31.36/0.8879	29.50/0.8946	35.60/0.9663
VDSR [19]	$\times 2$	665K	612.6G	37.53/0.9590	33.05/0.9130	31.90/0.8960	30.77/0.9140	37.22/0.9750
DRCN [20]	$\times 2$	1774K	9788.7G	37.63/0.9588	33.04/0.9118	31.85/0.8942	30.75/0.9133	37.55/0.9732
LapSRN [22]	$\times 2$	813K	29.9G	37.52/0.9591	33.08/0.9130	31.08/0.8950	30.41/0.9101	37.27/0.9740
MemNet [40]	$\times 2$	677K	623.9G	37.78/0.9597	33.28/0.9142	32.08/0.8978	31.31/0.9195	37.72/0.9740
SRFBN-S [28]	$\times 2$	282K	574.4G	37.78/0.9597	33.35/0.9156	32.00/0.8970	31.41/0.9207	38.06/0.9757
IDN [17]	$\times 2$	553K	127.7G	<u>37.83/0.9600</u>	33.30/0.9148	32.08/0.8985	31.27/0.9196	38.01/0.9749
CARN [2]	$\times 2$	1592K	222.8G	37.76/0.9590	33.52/0.9166	32.09/0.8978	31.92/0.9256	<u>38.36/0.9765</u>
FALSR-A [6]	$\times 2$	1021K	234.7G	37.82/0.9595	<u>33.55/0.9168</u>	<u>32.12/0.8987</u>	<u>31.93/0.9256</u>	-/-
SMSR	$\times 2$	985K	(224.1G)131.6G	<b>38.00/0.9601</b>	<b>33.64/0.9179</b>	<b>32.17/0.8990</b>	<b>32.19/0.9284</b>	<b>38.76/0.9771</b>
Bicubic	$\times 3$	-	-	30.39/0.8682	27.55/0.7742	27.21/0.7385	24.46/0.7349	26.95/0.8556
SRCNN [8]	$\times 3$	57K	52.7G	32.75/0.9090	29.30/0.8215	28.41/0.7863	26.24/0.7989	30.48/0.9117
VDSR [19]	$\times 3$	665K	612.6G	33.67/0.9210	29.78/0.8320	28.83/0.7990	27.14/0.8290	32.01/0.9340
DRCN [20]	$\times 3$	1774K	9788.7G	33.82/0.9226	29.76/0.8311	28.80/0.7963	27.14/0.8279	32.24/0.9343
MemNet [40]	$\times 3$	677K	623.9G	34.09/0.9248	30.01/0.8350	28.96/0.8001	27.56/0.8376	32.51/0.9369
SRFBN-S [28]	$\times 3$	375K	686.4G	34.20/0.9255	30.10/0.8372	28.96/0.8010	27.66/0.8415	33.02/0.9404
IDN [17]	$\times 3$	553K	57.0G	34.11/0.9253	29.99/0.8354	28.95/0.8013	27.42/0.8359	32.71/0.9381
CARN [2]	$\times 3$	1592K	118.8G	34.29/0.9255	30.29/0.8407	29.06/0.8034	28.06/0.8493	33.50/0.9440
SMSR	$\times 3$	993K	(100.5G)67.8G	<b>34.40/0.9270</b>	<b>30.33/0.8412</b>	<b>29.10/0.8050</b>	<b>28.25/0.8536</b>	<b>33.68/0.9445</b>
Bicubic	$\times 4$	-	-	28.42/0.8104	26.00/0.7027	25.96/0.6675	23.14/0.6577	24.89/0.7866
SRCNN [8]	$\times 4$	57K	52.7G	30.48/0.8628	27.50/0.7513	26.90/0.7101	24.52/0.7221	27.58/0.8555
VDSR [19]	$\times 4$	665K	612.6G	31.35/0.8830	28.02/0.7680	27.29/0.7260	25.18/0.7540	28.83/0.8870
DRCN [20]	$\times 4$	1774K	9788.7G	31.53/0.8854	28.02/0.7670	27.23/0.7233	25.18/0.7524	28.93/0.8854
LapSRN [22]	$\times 4$	813K	149.4G	31.54/0.8850	28.19/0.7720	27.32/0.7270	25.21/0.7560	29.09/0.8900
MemNet [40]	$\times 4$	677K	623.9G	31.74/0.8893	28.26/0.7723	27.40/0.7281	25.50/0.7630	29.42/0.8942
SRFBN-S [28]	$\times 4$	483K	852.9G	31.98/0.8923	28.45/0.7779	27.44/0.7313	25.71/0.7719	29.91/0.9008
IDN [17]	$\times 4$	553K	32.3G	31.82/0.8903	28.25/0.7730	27.41/0.7297	25.41/0.7632	29.41/0.8942
CARN [2]	$\times 4$	1592K	90.9G	<b>32.13/0.8937</b>	<b>28.60/0.7806</b>	<b>27.58/0.7349</b>	<b>26.07/0.7837</b>	<b>30.47/0.9084</b>
SMSR	$\times 4$	1006K	(57.2G)41.6G	<u>32.12/0.8932</u>	<b>28.55/0.7808</b>	<b>27.55/0.7351</b>	<b>26.11/0.7868</b>	<b>30.54/0.9085</b>

Table 3. Comparison between learning-based masks and gradient-based masks. Results are achieved on Set14 for  $\times 2$  SR.

$M^{spa}$	#Params.	$\alpha$	Sparsity	Set14	
				PSNR	SSIM
Gradient-based	926K	30	0.51	33.48	0.9163
	926K	30	0.62	33.42	0.9155
	926K	30	0.72	33.33	0.9151
	926K	50	0.50	33.45	0.9162
	926K	50	0.61	33.39	0.9153
Learning-based	926K	50	0.71	33.30	0.9150
	985K	-	0.46	33.64	0.9179
	985K	-	0.64	33.61	0.9174
	985K	-	0.73	33.52	0.9169

We further investigate the sparsities achieved by our SMMs for different scale factors. Specifically, we feed an LR image ( $\times 2$  downsampled) to  $\times 2/3/4$  SMSR networks and compare the sparsities in their SMMs. As shown in Fig. 7, the sparsities decrease for larger scale factors in most SMMs. Since more details need to be reconstructed for larger scale factors, more locations are marked as ‘‘important’’ ones (with sparsities being decreased).

**Learning-based Masks vs. Heuristic Masks.** As regions of edges are usually identified as important ones in our spatial masks (Fig. 6), another straightforward choice is to use heuristic masks. KernelGAN [3] follows this idea to identify regions with large gradients as important ones when applying ZSSR [38] and uses a masked loss to focus on these regions. To demonstrate the effectiveness of learning-based

masks in our SMSR, we introduced a variant with gradient-induced masks. Specifically, we consider locations with gradients larger than a threshold  $\alpha$  as important ones and keep the spatial mask fixed within the network. The performance of this variant is compared to our SMSR in Table 3. Compared to learning-based masks, the variant with gradient-based masks suffers a notable performance drop with comparable sparsity (e.g., 33.52 vs. 33.33/33.30). Further, we can see from Fig. 8 that learning-based masks facilitate our SMSR to achieve better trade-off between SR performance and computational efficiency. Using fixed heuristic masks, it is difficult to obtain fine-grained localization of redundant computation. In contrast, learning-based masks enable our SMSR to accurately localize redundant computation to produce better results.

### 5.3. Comparison with State-of-the-art Methods

We compare our SMSR with nine state-of-the-art methods, including SRCNN [8], VDSR [19], DRCN [20], LapSRN [22], MemNet [40], SRFBN-S [28], IDN [17], CARN [2], and FALSR-A [6]. As this paper focuses on lightweight SR networks ( $< 2M$ ), several recent works with large models (e.g., EDSR [29] ( $\sim 40M$ ), RCAN [50] ( $\sim 15M$ ) and SAN [7] ( $\sim 15M$ )) are not included for comparison. Quantitative results are presented in Table 4 and visualization re-

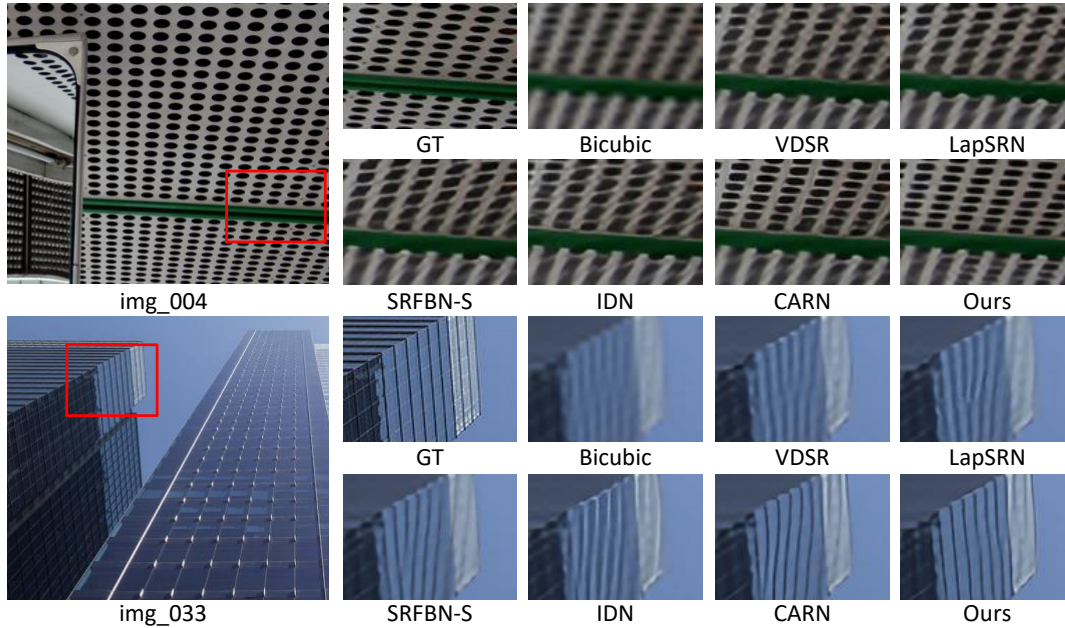


Figure 9. Visual comparison on the Urban100 dataset for  $\times 4$  SR.

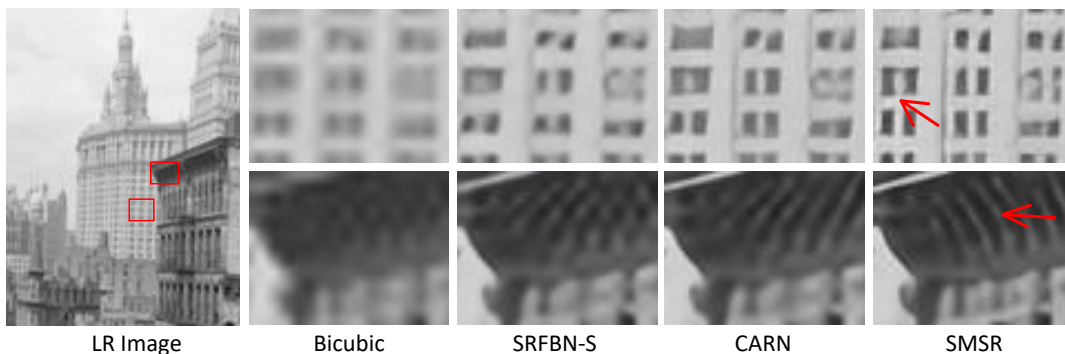


Figure 10. Visual comparison on a real-world image.

sults are shown in Figs. 9 and 10.

**Quantitative Results.** As shown in Table 4, our SMSR outperforms the state-of-the-art methods on most datasets. For example, our SMSR achieves much better performance than CARN for  $\times 2$  SR, with the number of parameters and FLOPs being reduced by 38% and 41%, respectively. With a comparable model size, our SMSR performs favorably against FALSAR-A and achieves better inference efficiency in terms of FLOPs (131.6G vs. 234.7G). With comparable computational complexity in terms of FLOPs (131.6G vs. 127.7G), our SMSR achieves much higher PSNR values than IDN. Using sparse masks to skip redundant computation, our SMSR reduces 41%/33%/27% FLOPs for  $\times 2/3/4$  SR while maintaining the state-of-the-art performance. We further show the trade-off between performance, number of parameters and FLOPs in Fig. 1. We can see that our SMSR achieves the best PSNR performance with low computational cost.

**Qualitative Results.** Figure 9 compares the qualitative re-

sults achieved on Urban100. Compared to other methods, our SMSR produces better visual results with fewer artifacts, such as the lattices in *img\_004* and the stripes on the building in *img\_033*. We further tested our SMSR on a real-world image to demonstrate its effectiveness. As shown in Fig. 10, our SMSR achieves better perceptual quality while other methods suffer notable artifacts.

## 6. Conclusion

In this paper, we explore the sparsity in image SR to improve inference efficiency of SR networks. Specifically, we develop a sparse mask SR network to prune redundant computation. Our spatial and channel masks work jointly to localize redundant computation at a fine-grained level such that our network can effectively reduce computational cost while maintaining comparable performance. Extensive experiments demonstrate that our network achieves the state-of-the-art performance with significant FLOPs reduction and a speedup on mobile devices.



## References

- [1] Eirikur Agustsson and Radu Timofte. NTIRE 2017 challenge on single image super-resolution: Dataset and study. In *CVPRW*, pages 1122–1131, 2017. [5](#)
- [2] Namhyuk Ahn, Byungkon Kang, and Kyung-Ah Sohn. Fast, accurate, and lightweight super-resolution with cascading residual network. In *ECCV*, pages 252–268, 2018. [1](#), [2](#), [5](#), [6](#), [7](#)
- [3] Sefi Bell-Kligler, Assaf Shocher, and Michal Irani. Blind super-resolution kernel estimation using an internal-gan. In *NeurIPS*, pages 284–293, 2019. [7](#)
- [4] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie-Line Alberi-Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *BMVC*, pages 1–10, 2012. [5](#)
- [5] Kumar Chellapilla, Sidd Puri, and Patrice Simard. High performance convolutional neural networks for document processing. In *IWFHR*, 2006. [6](#)
- [6] Xiangxiang Chu, Bo Zhang, Hailong Ma, Ruijun Xu, Jixiang Li, and Qingyuan Li. Fast, accurate and lightweight super-resolution with neural architecture search. In *ICPR*, 2020. [2](#), [5](#), [6](#), [7](#)
- [7] Tao Dai, Jianrui Cai, Yongbing Zhang, Shu-Tao Xia, and Lei Zhang. Second-order attention network for single image super-resolution. In *CVPR*, 2019. [7](#)
- [8] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *ECCV*, pages 184–199, 2014. [1](#), [2](#), [7](#)
- [9] Michael Figurnov, Maxwell D. Collins, Yukun Zhu, Li Zhang, Jonathan Huang, Dmitry P. Vetrov, and Ruslan Salakhutdinov. Spatially adaptive computation time for residual networks. In *CVPR*, pages 1790–1799, 2017. [2](#)
- [10] Xitong Gao, Yiren Zhao, Lukasz Dudziak, Robert D. Mullins, and Cheng-Zhong Xu. Dynamic channel pruning: Feature boosting and suppression. In *ICLR*, 2019. [2](#), [3](#)
- [11] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, pages 9224–9232, 2018. [2](#)
- [12] Jinjin Gu, Hannan Lu, Wangmeng Zuo, and Chao Dong. Blind super-resolution with iterative kernel correction. In *CVPR*, 2019. [1](#)
- [13] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *NeurIPS*, pages 1135–1143, 2015. [2](#)
- [14] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *CVPR*, pages 4340–4349, 2019. [2](#), [3](#)
- [15] Xuecai Hu, Haoyuan Mu, Xiangyu Zhang, Zilei Wang, Jian Sun, and Tieniu Tan. Meta-SR: A magnification-arbitrary network for super-resolution. In *CVPR*, 2019. [1](#)
- [16] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *CVPR*, pages 5197–5206, 2015. [5](#)
- [17] Zheng Hui, Xiumei Wang, and Xinbo Gao. Fast and accurate single image super-resolution via information distillation network. In *CVPR*, 2018. [1](#), [2](#), [5](#), [6](#), [7](#)
- [18] Eric Jang, Shixiang Gu, and Ben Poole. Categorical parameterization with gumbel-softmax. In *ICLR*, 2017. [3](#)
- [19] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *CVPR*, pages 1646–1654, 2016. [1](#), [2](#), [7](#)
- [20] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeply-recursive convolutional network for image super-resolution. In *CVPR*, pages 1637–1645, 2016. [7](#)
- [21] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. [5](#)
- [22] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *CVPR*, pages 5835–5843, 2017. [2](#), [7](#)
- [23] Andrew Lavin and Scott Gray. Fast algorithms for convolutional neural networks. In *CVPR*, pages 4013–4021, 2016. [6](#)
- [24] Wonkyung Lee, Junghyup Lee, Dohyung Kim, and Bumsub Ham. Learning with privileged information for efficient image super-resolution. In *ECCV*, 2020. [2](#), [5](#)
- [25] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. In *ICLR*, 2017. [2](#)
- [26] Hao Li, Hong Zhang, Xiaojuan Qi, Ruigang Yang, and Gao Huang. Improved techniques for training adaptive deep networks. In *ICCV*, pages 1891–1900, 2019. [2](#)
- [27] Xiaoxiao Li, Ziwei Liu, Ping Luo, Chen Change Loy, and Xiaoou Tang. Not all pixels are equal: Difficulty-aware semantic segmentation via deep layer cascade. In *CVPR*, pages 6459–6468, 2017. [3](#)
- [28] Zhen Li, Jinglei Yang, Zheng Liu, Xiaomin Yang, Gwanggil Jeon, and Wei Wu. Feedback network for image super-resolution. In *CVPR*, pages 3867–3876, 2018. [7](#)
- [29] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPR*, 2017. [1](#), [2](#), [7](#)
- [30] Ji Lin, Yongming Rao, Jiwen Lu, and Jie Zhou. Runtime neural pruning. In *NeurIPS*, pages 2181–2191, 2017. [2](#), [3](#)
- [31] Ming Liu, Zhilu Zhang, Liya Hou, Wangmeng Zuo, and Lei Zhang. Deep adaptive inference networks for single image super-resolution. In *ECCVW*, 2020. [2](#)
- [32] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *ICCV*, pages 2755–2763, 2017. [2](#)
- [33] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *ICCV*, pages 5068–5076, 2017. [2](#)
- [34] David Martin, Charless Fowlkes, Doron Tal, Jitendra Malik, et al. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, 2001. [5](#)
- [35] Yusuke Matsui, Kota Ito, Yuji Aramaki, Azuma Fujimoto, Toru Ogawa, Toshihiko Yamasaki, and Kiyoharu Aizawa. Sketch-based manga retrieval using manga109 dataset. *Multimedia Tools Appl.*, 76(20):21811–21838, 2017. [5](#)
- [36] Ravi Teja Mullapudi, William R. Mark, Noam Shazeer, and Kayvon Fatahalian. Hydranets: Specialized dynamic architectures for efficient inference. In *CVPR*, pages 8080–8089, 2018. [2](#)
- [37] Mengye Ren, Andrei Pokrovsky, Bin Yang, and Raquel Ur-

- tasun. Sbnnet: Sparse blocks network for fast inference. In *CVPR*, pages 8711–8720, 2018. 2, 3
- [38] Assaf Shocher, Nadav Cohen, and Michal Irani. “Zero-shot” super-resolution using deep internal learning. In *CVPR*, 2018. 1, 7
- [39] Jae Woong Soh, Sunwoo Cho, and Nam Ik Cho. Meta-transfer learning for zero-shot super-resolution. In *CVPR*, 2020. 1
- [40] Ying Tai, Jian Yang, Xiaoming Liu, and Chunyan Xu. Memnet: A persistent memory network for image restoration. In *ICCV*, pages 4549–4557, 2017. 7
- [41] Nicolas Vasilache, Jeff Johnson, Michaël Mathieu, Soumith Chintala, Serkan Piantino, and Yann LeCun. Fast convolutional nets with fbfft: A GPU performance evaluation. In *ICLR*, 2015. 6
- [42] Longguang Wang, Yingqian Wang, Xiaoyu Dong, Qingyu Xu, Jungang Yang, Wei An, and Yulan Guo. Unsupervised degradation representation learning for blind super-resolution. In *CVPR*, 2021. 1
- [43] Longguang Wang, Yingqian Wang, Zaiping Lin, Jungang Yang, Wei An, and Yulan Guo. Learning for scale-arbitrary super-resolution from scale-specific networks. *arXiv*, 2020. 1
- [44] Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E. Gonzalez. Skipnet: Learning dynamic routing in convolutional networks. In *ECCV*, volume 11217, pages 420–436, 2018. 2
- [45] Zuxuan Wu, Tushar Nagarajan, Abhishek Kumar, Steven Rennie, Larry S. Davis, Kristen Grauman, and Rogério Schmidt Feris. Blockdrop: Dynamic inference paths in residual networks. In *CVPR*, pages 8817–8826, 2018. 2
- [46] Zhenda Xie, Zheng Zhang, Xizhou Zhu, Gao Huang, and Stephen Lin. Spatially adaptive inference with stochastic feature sampling and interpolation. In *ECCV*, 2020. 2
- [47] Jiecao Yu, Andrew Lukefahr, David Palframan, Ganesh Dasika, Reetuparna Das, and Scott Mahlke. Scalpel: Customizing dnn pruning to the underlying hardware parallelism. In *ISCA*, 2017. 6
- [48] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *International Conference on Curves and Surfaces*, volume 6920, pages 711–730, 2010. 5
- [49] Kai Zhang, Luc Van Gool, and Radu Timofte. Deep unfolding network for image super-resolution. In *CVPR*, 2020. 1
- [50] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *ECCV*, pages 1646–1654, 2018. 1, 5, 7
- [51] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *CVPR*, pages 2472–2481, 2018. 1, 2, 5