

MOFA: A Model Simplification Roadmap for Image Restoration on Mobile Devices

Xiangyu Chen^{1,2†}, Ruiwen Zhen^{2†}, Shuai Li², Xiaotian Li², Guanghui Wang^{3*}

¹Department of EECS, University of Kansas, KS, USA

²SenseBrain Technology, San Jose, CA, USA

³Department of CS, Toronto Metropolitan University, Toronto, ON, Canada

xychen@ku.edu, {zhenruiwen, shuai.li, lixiaotian}@sensebrain.site, wangcs@torontomu.ca (* corresponding author)

Abstract

Image restoration aims to restore high-quality images from degraded counterparts and has seen significant advancements through deep learning techniques. The technique has been widely applied to mobile devices for tasks such as mobile photography. Given the resource limitations on mobile devices, such as memory constraints and runtime requirements, the efficiency of models during deployment becomes paramount. Nevertheless, most previous works have primarily concentrated on analyzing the efficiency of single modules and improving them individually. This paper examines the efficiency across different layers. We propose a roadmap that can be applied to further accelerate image restoration models prior to deployment while simultaneously increasing PSNR (Peak Signal-to-Noise Ratio) and SSIM (Structural Similarity Index). The roadmap first increases the model capacity by adding more parameters to partial convolutions on FLOPs non-sensitive layers. Then, it applies partial depthwise convolution coupled with decoupling upsampling/downsampling layers to accelerate the model speed. Extensive experiments demonstrate that our approach decreases runtime by up to 13% and reduces the number of parameters by up to 23%, while increasing PSNR and SSIM on several image restoration datasets. Source Code of our method is available at <https://github.com/xiangyu8/MOFA>.

1. Introduction

During the past decade, deep learning has been the dominant approach for various computer vision tasks, including image classification [29, 43], object detection [22, 41], segmentation [33, 32], and image restoration [39, 4, 46, 45,

10, 8], owing to its remarkable feature extraction capabilities derived from a substantial number of parameters. However, these large models often demand significant computational resources and result in longer inference latency, posing challenges for direct deployment on mobile devices. As a consequence, their application in mobile photography has been restricted.

To address this challenge, researchers have been exploring different approaches to simplify deep learning models by reducing the size of models directly, like channel dimensions and the number of layers [3, 21, 38]. Further, some works design efficient modules to replace inefficient modules [47, 52, 28, 15, 35, 37, 13], e.g. replacing the original convolution with depthwise convolution [3, 38]. However, most model compression methods focus on reducing the FLOPs or runtime only, at the cost of performance loss.

Partial Convolution (PConv) [2] is a recently proposed efficient module that simplifies original convolutions by performing calculations only on a portion of the channel dimension. In comparison to depthwise convolution, PConv demonstrates smaller runtime during deployment on mobile devices, even though it requires more parameters. Building on this observation, we can leverage the idea of adding more parameters during model compression to compensate for any potential performance loss, particularly in terms of PSNR (Peak Signal-to-Noise Ratio). By doing so, the model can still achieve faster runtime through model compression algorithms, while also benefiting from the larger capacity enabled by the additional parameters. As a result, the overall performance of the model is either maintained or even improved, considering the enhanced capacity brought about by the increased parameters.

Motivated by this concept, we propose a model simplification roadmap, which encompasses a set of tricks and techniques designed to optimize image restoration models for deployment on mobile devices. Through the application of

[†] These authors contributed equally to this work

these techniques, we aim to strike the right balance between model efficiency and performance, enabling the deployment of high-quality image restoration models on mobile devices without compromising on speed or output quality. The contributions of this paper are summarized below.

- We propose a model simplification roadmap, consisting of tricks to add parameters to expand model capacity first and then reduce runtime, *i.e.* **MOre parameters and FAster (MOFA)**. Given a small model, the model can be further accelerated before deployment with improved performance following the roadmap.
- Building upon the concept of partial convolution, we propose Partial Depthwise Convolution (PDWConv) to further enhance acceleration. PDWConv involves performing depthwise convolution across only a portion of the channel dimension for input features.
- We tested our roadmap on widely used small image restoration models, PMRID and NAFNet, on image restoration tasks including image denoising, image deblurring, and image deraining. Extensive experiments show that using our tricks, runtime for PMRID and NAFNet is decreased by up to 13% and parameters are reduced by 23% with better PSNR and SSIM.

2. Related work

This section presents related works on image restoration models on mobile devices as well as those modules designed to increase efficiency.

2.1. Image restoration on mobile devices

Image restoration aims to reverse degraded images into high-quality images. Tasks include image denoising, image deblurring, image draining, etc. These tasks are more commonly used in mobile devices like mobile photography. This encourages researchers to work on designing efficient deep learning models considering both runtime and image quality (IQ) [20, 14, 40, 16, 5, 53, 44, 18, 6, 1, 51]. PMRID [38] is one of the earliest efficient models. It carefully designs a U-Net-like model and increases efficiency by putting most parameters and calculations on the encoder stages. MicroISP [17] propose an attention-based multi-branch model so that it can be adapted to the different computational power of edge devices flexibly. MFDNet [23] designs a mobile-friendly attention module coupled with a reparameterization module to increase the denoising efficiency on iPhone 11. FusionNet [21] designs a U-Net-like model to take two images from dual cameras as input and output the deblurred images. Beyond image restoration, MobileOne [36] proposes re-parameterizable structures to decrease the runtime during inference for general tasks like

image classification. Orthogonally, we propose a roadmap to further accelerate U-Net-like backbones.

2.2. Efficient modules on computer vision

The first set of efficient modules is related to the convolution operation. Vanilla convolution performs well but it visits each pixel the number of output channel times, which takes much computation and is inefficient for mobile applications. To improve this, Xception [7] proposes to use separable convolution, a depthwise convolution followed by a pointwise convolution, which requires visiting each pixel only once by depthwise convolution. This reduces computation greatly even counting the following pointwise convolution. PConv [2] proposes to calculate partial convolution, *i.e.* a portion like 1/4 of the input feature across channel dimension as they observe most kernels are learning similar features.

In addition, some works propose a model simplification process by reducing unnecessary modules or replacing them with more efficient modules, *e.g.* NAFNet [3], ConvNeXt [24] and Efficient Unet [34]. Other efficient modules include dynamic convolution [12], fast nearest convolution [26], IRA block [9], and self-calibrated module [27].

However, all these consider only the efficiency of a single module and attempt to replace them with more efficient ones. Instead, we consider the efficiency of the entire network and manipulate inefficient layers according to the FLOPs distribution. PConv module provides us the flexibility to add parameters freely by setting the portion differently for different layers.

3. Roadmap for simplification

This section presents our model simplification roadmap, under the guideline that using more parameters wherever possible to compensate for the performance loss brought by rapid runtime decrease. We employ the widely used PMRID model from image denoising as an example for illustration. Following our roadmap, either the runtime gets reduced drastically or more parameters are added to the model in each step. At the end of optimization, the runtime is decreased by 11%, the number of parameters is reduced by 6% and the PSNR is increased by 0.02dB. Noticeably, this roadmap can be applied to any U-Net-like model as a model compression method before deployment on mobile devices.

Settings We train PMRID on SIDD dataset. Runtime is tested on one single v100 GPU with input size $3 \times 1024 \times 1024$, averaging 600 rounds. Each trick mentioned below from subsection 3.2 is built on top of the previous ones.

3.1. Starting point - PMRID

PMRID [38] is a U-Net-like model proposed for raw image denoising on mobile devices as in Figure 1 and 2. It is

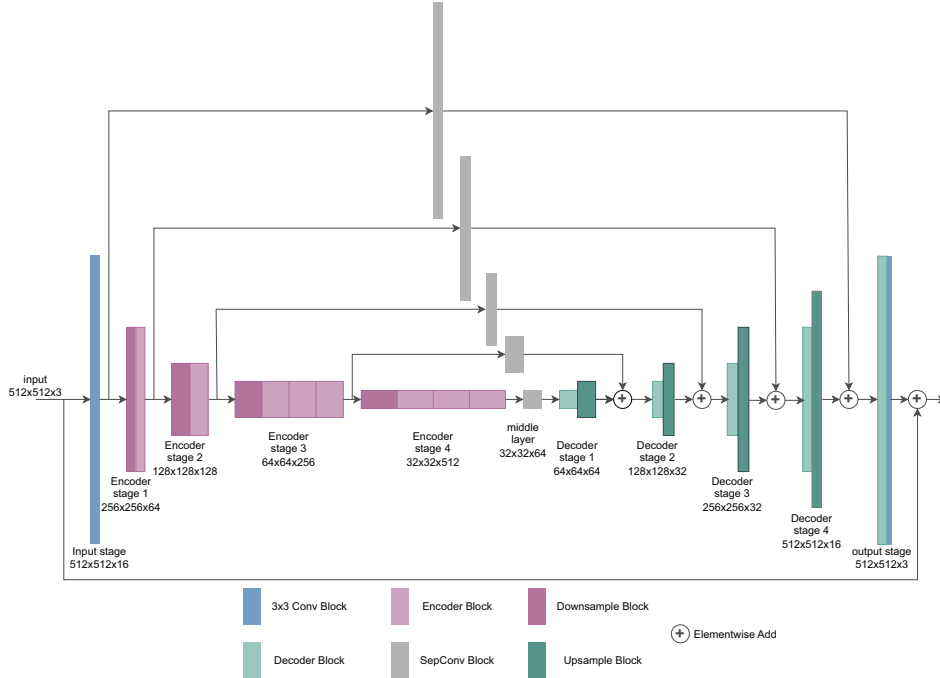


Figure 1. The overall network structure of PMRID

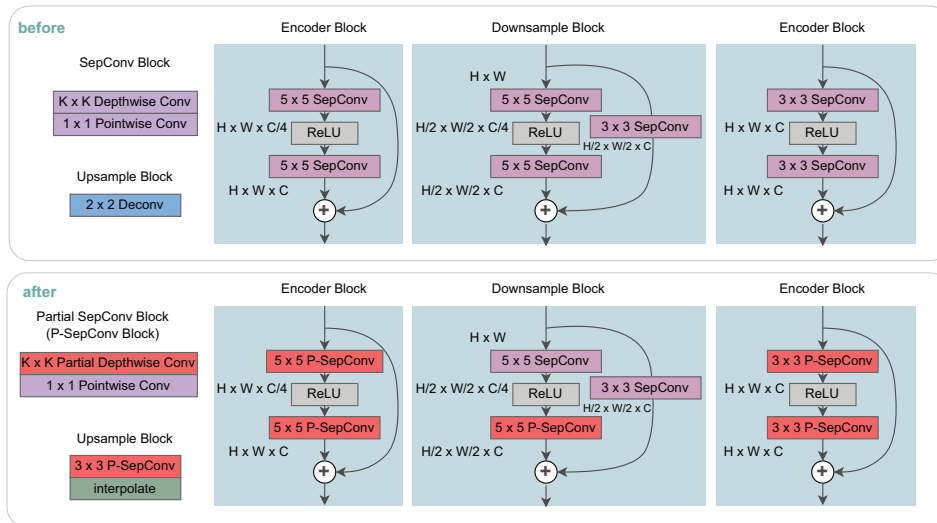


Figure 2. Detailed Blocks. We replace all Separable Convolutions (stride = 1) with one of the following convolutions: partial depth-wise convolution and partial convolution (PConv), with the portion 1/4 or 1/2.

worth noting that PMRID holds much more parameters than normal U-Net-like image restoration models on mobile devices like DnCNN [50] and NAFNet, but comparable runtime. This is because PMRID puts most of its parameters on its encoder only like half-UNet [25], while most U-Net-like models distribute equally on the encoder and decoder. This asymmetrical design gives more freedom to put even more parameters on high-level encoders, where the input feature size is the smallest and hence the FLOPs will not increase drastically along with the increase of parameter amounts.

For all convolutional layers, most layers use efficient depth-wise convolution followed by point-wise convolution, excluding the input and output layer, where vanilla convolutions are used as the channel dimension is relatively small hence it will not add too many calculations/FLOPs into the whole structure.

3.2. Depthwise separable convolution to PConv

Partial convolution (PConv) [2] is a newly proposed efficient replacement for convolution modules on mobile de-

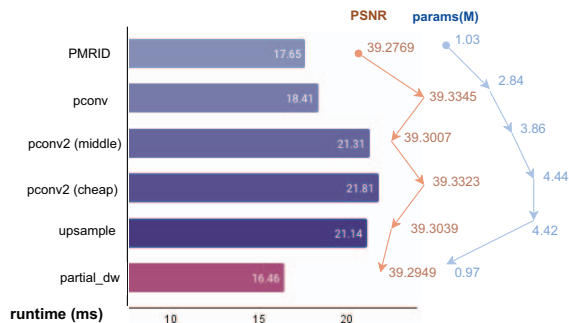


Figure 3. Roadmap. Starting from the baseline PMRID, we add each trick on top of all previous tricks, and it ends with partial_dw.

ices, aiming to reduce both the number of floating-point operations (FLOPs), related to the number of multiply-adds, and the floating-point operations per second (FLOPS), considering the hardware deployment efficiency like memory access time. Unlike depthwise convolution which requires much memory access as convolution among different depths does not share parameters, partial convolution calculates only part of the convolution in terms of channel dimension, *e.g.* 1/4 by default. In this way, the computation gets reduced as depthwise convolution compared with the original convolution while it also enjoys the deployment time reduction due to parameter sharing as the original convolution, reducing memory access time drastically compared with depthwise convolution. Besides this, partial convolutions provide us more flexibility to manipulate later, *e.g.* the portion. Hence, we first replace all depth-wise separable convolutions with partial convolutions. If there are vanilla convolutions exist (except the input and output layer), we also replace all those convolutions with partial convolutions directly. After this replacement, PSNR and SSIM increased to 39.3345 and 0.9559 respectively and runtime also increased by 0.76ms. This may be caused by the increase in parameters and FLOPs after the replacement. From now on, we use this partial convolution version as our model.

3.3. More parameters - middle layers

More parameters mean higher network capacity and better performance at most times. This motivates us to add more parameters to compensate for the PSNR and SSIM loss brought by the conversion from convolution to partial convolution. To achieve this, we can manipulate how large the “part” is in partial convolution, *e.g.* increasing it from default 1/4 to 1/2. Another problem is, where can we put those extra parameters to increase feature extraction ability without hurting runtime too much?

As for image restoration tasks, they target on reconstruct pixels, not the semantic understanding of the whole images as high-level vision tasks, *e.g.* image classification and detection. Thus we claim that more parameters should be

added in shallow layers to increase low-level feature extraction power. However, more parameters in shallow layers usually mean much larger FLOPs and runtime as the feature size for shallow layers is usually larger than deeper layers. Considering both the performance and runtime, we increase the portion from 1/4 to 1/2 in partial convolution to add parameters on middle layers for both encoder and decoder, *e.g.*, encoder 3 and decoder 2 for the PMRID model. As a result, PSNR and SSIM decreased a little bit, and runtime increased only 2.9ms even if the number of parameters increased from 2.84M to 3.86M and FLOPs from 2.12G to 3.25G. This increases model capacity.

3.4. More parameters - cheap layers

To add more parameters without adding too many FLOPs and runtime, we first examine the flop distribution for each layer. There are two ways to achieve this. The most intuitive way is to calculate the FLOPs layer by layer using Python tools. However, this will include some unnecessary FLOPs we are not interested in here, *e.g.* calculation in skip connections. Another way is to use a theoretical formula to estimate the FLOPs of convolution layers only. Here, we use the second method. The estimated FLOPs distribution can be found in Figure 4. It shows that all the decoders along with the middle layer take up relatively fewer FLOPs compared with encoders, hence we call them “cheap layers”. As a result, we increase the portion from 1/4 to 1/2 for partial convolutions in decoder layers. After adding this trick, PSNR, and SSIM increased to 39.3323 and 0.9560 respectively while runtime only increased 0.5ms. The amount of parameters now is 4.44M, which is 4× compared with the original PMRID model. After adding parameters to increase the model capacity using previous tricks, the following subsections focus on reducing the runtime.

3.5. Faster - upsampling/downsampling

Downsampling and upsampling modules are vital in U-Net-like networks. There are some choices for them, *e.g.* PMRID used depthwise separable convolution with stride 2 for downsampling and deconvolution for upsampling while NAFNet employed convolution coupled with pixel shuffle for upsampling. Instead, paper [34] argued that performing downsampling before the convolution and upsampling after the convolution can improve the efficiency, different from the normal order for downsampling and upsampling. Hence, we may first decouple downsampling from convolution with stride 2 to one average pooling layer followed by a convolution layer with stride 1. Upsampling deconvolutions are replaced with an interpolation layer following a stride-1 convolution layer. We select the most efficient downsampling/upsampling from the choices mentioned above and the original downsampling/upsampling strategies for differ-

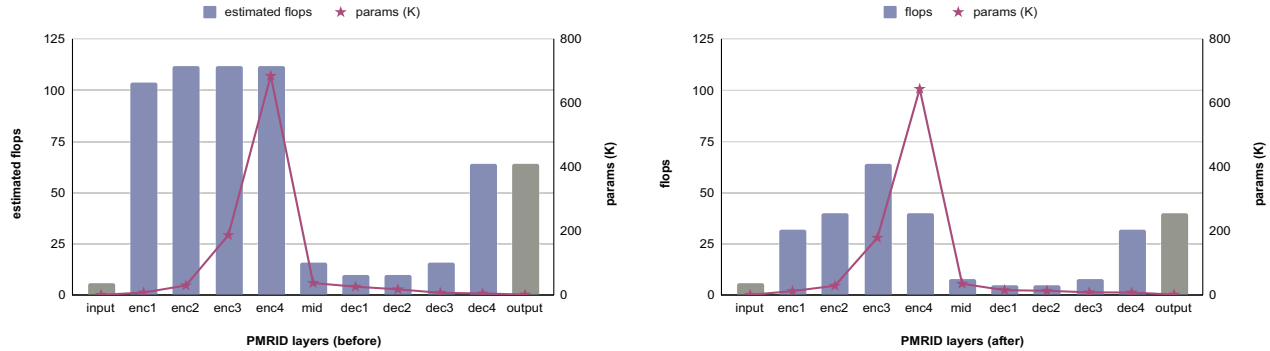


Figure 4. Estimated FLOPs distribution across layers before/after applying our roadmap tricks (dimension threshold $d * p$ is 0 here). Note that in PMRID structure, some layers (e.g. input and output layers) are vanilla convolutions and separable convolutions elsewhere. To get a clear structure-related FLOPs distribution (only related to the channel dimensions and feature size changes), we estimate the FLOPs by assuming all layers are vanilla convolutions and excluding FLOPs in skip connections.

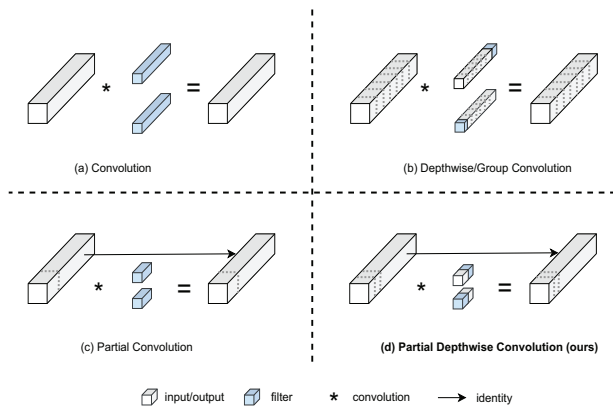


Figure 5. The structure of vanilla convolution, depthwise convolution, partial convolution, and partial depthwise convolution.

ent structures. For PMRID, we replace only the upsampling layers with the decoupled ones. After this change, the runtime gets reduced by 0.67 ms, PSNR dropped a little bit while SSIM remains the same.

3.6. Faster - PConv to PDWConv

To further accelerate it, we propose to replace partial convolution (PConv) with partial depthwise convolution (PDWConv). As illustrated in Figure 5, Partial depthwise convolution calculates only a portion across the channel dimension of the input features as in PConv and the rest untouched part uses identity mapping. Unlike PConv, we perform depthwise convolution for the selected portion while PConv uses original convolution.

Besides this, noticing that after taking the portion across the channel dimension, the input feature dimension will be decreased drastically. For features with small channel dimensions, e.g. 64 or 32, only 16 or 8 channels are used to learn the parameters, which may hurt the performance

worse than larger features. Thus, we replace only those PConv with PDWConv for larger channel dimensions while keeping PConv for smaller ones. For the PMRID model, we set the dimension threshold $d * p$ (d is the input channel dimension for the current layer and p is the portion) to 32 according to the experimental results.

Finally, the runtime gets decreased to 16.46ms while PSNR and SSIM become 39.2769 and 0.9556 respectively, still better than the original PMRID backbone.

4. Experiments

This section presents a detailed ablation study and analysis of some tricks in enhancing PMRID model on SIDD denoising dataset and both quantitative and qualitative results when we accelerate other models on RGB image denoising, image deblurring, raw image denoising, image deblurring with JPEG artifacts and image deraining tasks.

4.1. Setting

To verify the effectiveness of our pipeline, we apply it on two different models, NAFNet-tiny [3] which has only 7 blocks, and the backbone from PMRID [38]. For both models, our settings are mainly based on NAFNet [3] as PMRID lacks the settings and reported results on most image restoration benchmarks. Adam optimizer [19] is used with $\beta_1 = 0.9$, $\beta_2 = 0.9$ and weight decay 0. Total iteration is 200K during training. In addition, we report the FLOPs (input size $3 \times 256 \times 256$) and the number of parameters for comparison. For the performance, we use PSNR and SSIM for quantitative comparison. Qualitative visualization can be found in Section 4.5. We tested runtime using two ways. The first one is gpu based as mentioned above. We tested on a single v100 with $3 \times 1024 \times 1024$ as the input size, averaging from 600 rounds. Besides this, we also test our models on Android phone Tecno Camon 19 Pro based on Pytorch

Mobile framework, using image size $3 \times 256 \times 256$, averaging runtime from 3 forward passes. We use a larger input size on GPU runtime testing because when both our model and input are too small, the variance of runtime would be comparable to the actual runtime on v100 GPU, leading to biased averaged runtime.

4.2. Ablation study

In this subsection, we examine how each trick works in this acceleration process. Results can be found in Table 1. We can find the first 3 tricks mainly increased the PSNR and SSIM by adding more parameters efficiently without adding too much runtime. Specifically, replacing with partial convolution increased PSNR and SSIM to 39.3345 and 0.9559 respectively compared with the baseline (PSNR: 39.2769, SSIM: 0.9556). Adding more parameters for middle layers also boosts PSNR and SSIM to 39.3007 and 0.9557 separately. Adding more parameters on cheap layers coupled with middle layers together further increased PSNR to 39.3323 and SSIM to 0.9560. All these tricks bring much more parameters to the model to increase the model capacity without introducing too many FLOPs and runtime, making space for the runtime reduction tricks to keep PSNR and SSIM.

For the upsampling/downsampling modules, PMRID uses deconvolution layers to perform upsampling. When we decouple it into a convolution layer and an upsampling layer, the runtime gets reduced by 0.65ms while PSNR and SSIM are still comparable.

Finally, we replace all PConv with PDWConv, *i.e.* we use depth-wise convolution to replace the original convolution in PConv and it results in partial depth-wise convolution as illustrated in Figure 5. When we apply this directly on PConv skipping all steps in the middle, performance dropped from (PSNR: 39.3345, SSIM: 0.9559) to (PSNR: 39.2104, SSIM: 0.9552), even if runtime dropped significantly. This also demonstrates the necessity of expanding the model capacity to increase the model performance in the middle steps.

4.3. Analysis of different components

This subsection presents experimental results of how to determine alternatives for each trick.

Analysis of middle layers As mentioned above, we aim to add more parameters in middle-level layers to balance between performance and runtime, as we believe more parameters in shallower layers would be more appealing for low-level vision tasks as they focus more on extracting low-level features. However, more parameters in shallow layers also mean much more FLOPs as they take the largest feature size in the entire U-Net-like network. As a result, we decide to add parameters on the middle layers by increasing the portion of PConv from $1/4$ to $1/2$. The experiments can

be found in Table 2.

As adding more parameters on PConv layers in encoder stage 3 and decoder stage 2 increased PSNR from 39.2769 to 39.3007 compared with the PMRID backbone baseline while other encoder-decoder pairs even dropped the PSNR and SSIM, we use encoder stage 3 and decoder stage 2 as our final middle layers to add parameter in PMRID. Besides this, this pair also introduces much more parameters than others. This is because encoder stage 3 in this pair has 4 encoder blocks while other encoder stages have only 2.

Analysis of upsampling/downsampling layers For upsampling, we replace deconvolution layers with convolution and an upsampling layer in PMRID. This order can decrease the FLOPs compared with the reverse. For downsampling, the downsampling layer coupled with the convolution layer (stride 1) replaced the stride 2 convolution layers. Results can be found in Table 3. Here, the runtime is based on Pytorch Mobile tested on Android phones. We can find adding upsampling can reduce GPU runtime a little bit as in Table 1 while increasing inference time on the device as in Table 3. This may be caused by different acceleration implementations for GPU and Pytorch Mobile on devices. Besides this, downsampling increased runtime on Pytorch Mobile a lot for PMRID, from 255ms to 460ms. Hence, we replace upsampling only for PMRID.

Analysis of PDWConv threshold Noticing that PDWConv calculates only a portion of the entire depthwise convolution across channel dimension, and when the dimension is relatively small, *e.g.* 32, it may hurt performance drastically compared with the convolution in PConv which hold much more parameters. As a result, we replace only PConv with PDWConv with a larger channel dimension and set a threshold to determine. If the channel dimension of input features d divided by the portion p , *i.e.* $d * p$, $p = [1/2, 1/4]$, the actual channel dimension to be calculated, is larger than the threshold, replace PConv with PDWConv, otherwise, we keep them PConv. Considering both runtime and PSNR/SSIM, we choose 32 as the threshold for PMRID. It may differ for different backbones.

4.4. Applications

We apply our roadmap with two different backbones, PMRID [38] and NAFNet [3], and get their accelerated versions PMRID_{mofa} and NAFNet_{mofa} . As layer normalization is not deployable on mobile devices, we remove this layer directly to test runtime. We tested these models on various image restoration tasks, including denoising, deblurring, and deraining.

RGB Image Denoising. We test denoising performance on SIDD dataset. Quantitative results are shown in Table 4. This table shows for both PMRID and NAFNet, our accelerated versions decreased runtime on mobile devices by 11% and 13%, and PSNR and SSIM are also increased at

	PConv	PConv2 (middle)	PConv2 (cheap)	down/up	PDWConv	FLOPs (G)	Param (M)	PSNR(\uparrow)	SSIM(\uparrow)	runtime (ms) (\downarrow) GPU Mobile	
baseline						1.11	1.03	39.2769	0.9556	17.65	220
	✓					2.12	2.84	39.3345	0.9559	18.41	212
	✓	✓				3.25	3.86	39.3007	0.9557	21.31	250
	✓		✓			2.69	3.46	39.3211	0.9559	19.19	218
	✓	✓	✓			3.69	4.44	39.3323	0.9560	21.81	255
				✓		0.95	1.01	39.2741	0.9556	16.99	153
	✓				✓	1.03	0.95	39.2104	0.9552	14.46	162
	✓	✓	✓	✓	✓	1.11	0.97	39.2949	0.9558	16.46	196

Table 1. Ablation study based on PMRID model for image denoising on SIDD dataset.



Figure 6. Visualization. Image restoration results from PMRID, $PMRID_{mofa}$, NAFNet and $NAFNet_{mofa}$ on different tasks.

layers	FLOPs (G)	Param(M)	PSNR	SSIM
baseline	1.11	1.03	39.2769	0.9556
+PConv	2.12	2.84	39.3345	0.9559
enc3, dec2	3.25	3.86	39.3007	0.9557
enc2, dec3	2.58	2.93	39.2508	0.9556
enc1, dec4	2.54	2.86	39.2770	0.9556

Table 2. Adding more parameters to middle layers on PMRID.

replacement	Param(M)	runtime	PSNR	SSIM
no	4.44	255	39.3323	0.9560
upsampling	4.42	268	39.3039	0.9561
downsampling	4.61	460	39.3334	0.9559

Table 3. Replacement of upsampling/downsampling in PMRID.

the same time, from 39.2769dB to 39.2949dB for PMRID and from 39.2319dB to 39.3098dB for NAFNet. Note that

runtime reduction is applied to all the following tasks as we use the same models for comparison.

Image Deblurring. We tested image deblurring performance on GoPro [30] dataset. As shown in Table 4, we have increased PSNR by 0.14dB for PMRID and 0.05dB for NAFNet while reducing the runtime by more than 10% for both backbones.

Image Deblurring with JPEG artifacts. As in NAFNet [38], we also tested the image deblurring effect with JPEG artifacts. Following settings in NAFNet [38], we use REDS [31] dataset and evaluate on REDS-val-300. Table 4 shows $NAFNet_{mofa}$ increased PSNR by 0.07dB compared with NAFNet while the runtime gets decreased and $PMRID_{mofa}$ is comparable to PMRID.

Image deraining. Following settings in HINet [4], we also train our models on image deraining task with Rain13k dataset and test on Test100 [49], Rain100H [42], Rain100L [42], Test2800 [11] and Test1200 [48] datasets. Quantitative results are shown in Table 5. From this ta-

Method	runtime (ms)	# params	FLOPs (G)	denoising (SIDD)		deblurring (GoPro)		deblurring (REDS)	
	Mobile			PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
PMRID	220*	1.03M	1.15	39.2769	0.9556	29.0840	0.9153	27.6501	0.8348
PMRID _{mofa}	196(-11%)	0.97M(-6%)	1.11	39.2949	0.9558	29.2274	0.9176	27.6493	0.8347
NAFNet	69*	286.6K	1.03	39.2042	0.9555	29.4337	0.9237	27.6747	0.8381
NAFNet _{mofa}	60(-13%)	222.1K(-23%)	0.94	39.3098	0.9560	29.4810	0.9242	27.7450	0.8398

Table 4. Image denoising and deblurring results on different datasets. * Represents we also added *split_cat* operations for fair comparison.

Method	Test100		Rain100H		Rain100L		Test2800		Test1200	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
PMRID	28.3400	0.8690	27.9492	0.8343	32.5887	0.9316	32.6028	0.9260	31.4316	0.9001
PMRID _{mofa}	28.2158	0.8739	28.3177	0.8453	32.7912	0.9330	32.6106	0.9272	31.4936	0.9034
NAFNet	28.0719	0.8775	27.8863	0.8414	33.6056	0.9467	32.8942	0.9298	32.6602	0.9181
NAFNet _{mofa}	28.9186	0.8888	28.5133	0.8531	33.9560	0.9500	32.9450	0.9303	32.4535	0.9150

Table 5. Image deraining results on different test datasets.

thres $d * p$	runtime(ms)	Param(M)	PSNR	SSIM
0	14.10	0.94	39.2513	0.9554
16	15.8	0.95	39.2408	0.9557
32	16.46	0.97	39.2949	0.9558
64	16.56	1.21	39.2920	0.9559
128	16.71	1.25	39.2973	0.9561

Table 6. The effect of threshold $d * p$ to replace PConv with PDWConv.

ble, our accelerated versions for both backbones have improved the performance for most deraining datasets while reducing the runtime. Specifically, PMRID_{mofa} defeats PMRID on Rain100H, Rain100L, Test2800 and Test1200 datasets, and NAFNet_{mofa} improves over NAFNet on Test100, Rain100H, Rain100L and Test2800 datasets.

4.5. Visualization

To understand the qualitative results, we visualize image restoration results on different datasets, SIDD dataset for image denoising, GoPro dataset for image deblurring, REDS dataset for image deblurring with JPEG artifacts, and Rain100L dataset for deraining as in Figure 6.

For image denoising, we can find all these models, PMRID, PMRID_{mofa}, NAFNet and NAFNet_{mofa} have good denoising ability. Among them, we observe NAFNet_{mofa} is the best.

In GoPro dataset, the quality of deblurring results increases from PMRID to PMRID_{mofa} where the shape of the hand in the figure changes from blurred for PMRID to clear for NAFNet_{mofa}.

For deraining results on Rain100L dataset, we can see both PMRID_{mofa} and NAFNet_{mofa} show better deraining

results than their baselines.

5. Conclusion and limitation

This paper has presented a comprehensive roadmap to accelerate image restoration models before deploying them on mobile devices. The key idea is to enhance the model’s capacity without significantly increasing runtime by strategically adding more parameters to FLOP-insensitive layers and middle layers. Additionally, we propose the decoupling of upsampling/downsampling layers and apply partial depthwise convolution to further accelerate the models. The proposed roadmap is effectively applied to PMRID and NAFNet for various image restoration tasks. Extensive experimental results demonstrate the effectiveness of our strategy. Following our approach, we achieve a runtime reduction of up to 13% and reduce the number of parameters by 23% while simultaneously improving the image restoration performance.

Our approach offers a powerful solution for deploying image restoration models on mobile devices, striking a balance between increased model capacity and faster inference without compromising restoration quality. For deployment, the decrease of parameters is suitable for parameters-sensitive implementation like sensor programming. The introduction of *split_cat* operation in PDWConv and PConv might be less efficient in deployment based on Pytorch Mobile, ONNX and Caffe and need to be further accelerated.

Acknowledgement

The work was supported by and conducted at SenseBrain Technology. G. Wang was partly supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

- [1] Hanqing Chen, Yunhe Wang, Jianyuan Guo, and Dacheng Tao. Vanillanet: the power of minimalism in deep learning. *arXiv preprint arXiv:2305.12972*, 2023. 2
- [2] Jierun Chen, Shiu-hong Kao, Hao He, Weipeng Zhuo, Song Wen, Chul-Ho Lee, and S-H Gary Chan. Run, don't walk: Chasing higher flops for faster neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12021–12031, 2023. 1, 2, 3
- [3] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*, pages 17–33. Springer, 2022. 1, 2, 5, 6
- [4] Liangyu Chen, Xin Lu, Jie Zhang, Xiaojie Chu, and Chengpeng Chen. Hinet: Half instance normalization network for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 182–192, 2021. 1, 7
- [5] Haram Choi, Cheolwoong Na, Jinseop Kim, and Jihoon Yang. Exploration of lightweight single image denoising with transformers and truly fair training. In *Proceedings of the 2023 ACM International Conference on Multimedia Retrieval*, pages 452–461, 2023. 2
- [6] Haram Choi, Cheolwoong Na, Jihyeon Oh, Seungjae Lee, Jinseop Kim, Subeen Choe, Jeongmin Lee, Taehoon Kim, and Jihoon Yang. Ramit: Reciprocal attention mixing transformer for lightweight image restoration. *arXiv preprint arXiv:2305.11474*, 2023. 2
- [7] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017. 2
- [8] Xiaojie Chu, Liangyu Chen, Chengpeng Chen, and Xin Lu. Improving image restoration by revisiting global information aggregation. In *European Conference on Computer Vision*, pages 53–71. Springer, 2022. 1
- [9] Marcos V Conde, Florin Vasluiianu, Javier Vazquez-Corral, and Radu Timofte. Perceptual image enhancement for smartphone real-time applications. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1848–1858, 2023. 2
- [10] Akshay Dudhane, Syed Waqas Zamir, Salman Khan, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Burst image restoration and enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5759–5768, 2022. 1
- [11] Xueyang Fu, Jiabin Huang, Delu Zeng, Yue Huang, Xinghao Ding, and John Paisley. Removing rain from single images via a deep detail network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3855–3863, 2017. 7
- [12] Jiaming Guo, Xueyi Zou, Yuyi Chen, Yi Liu, Jianzhuang Liu, Youliang Yan, and Jia Hao. Asconvsr: Fast and lightweight super-resolution network with assembled convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1582–1592, 2023. 2
- [13] Meng-Hao Guo, Cheng-Ze Lu, Zheng-Ning Liu, Ming-Ming Cheng, and Shi-Min Hu. Visual attention network. *Computational Visual Media*, pages 1–20, 2023. 1
- [14] Yu Guo, Axel Davy, Gabriele Facciolo, Jean-Michel Morel, and Qiyu Jin. Fast, nonlocal and neural: a lightweight high quality solution to image denoising. *IEEE Signal Processing Letters*, 28:1515–1519, 2021. 2
- [15] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 1
- [16] Andrey Ignatov, Grigory Malivenko, Radu Timofte, Yu Tseng, Yu-Syuan Xu, Po-Hsiang Yu, Cheng-Ming Chiang, Hsien-Kai Kuo, Min-Hung Chen, Chia-Ming Cheng, et al. Pynet-v2 mobile: efficient on-device photo processing with neural networks. In *2022 26th International Conference on Pattern Recognition (ICPR)*, pages 677–684. IEEE, 2022. 2
- [17] Andrey Ignatov, Anastasia Sycheva, Radu Timofte, Yu Tseng, Yu-Syuan Xu, Po-Hsiang Yu, Cheng-Ming Chiang, Hsien-Kai Kuo, Min-Hung Chen, Chia-Ming Cheng, et al. Microisp: processing 32mp photos on mobile devices with deep learning. In *European Conference on Computer Vision*, pages 729–746. Springer, 2022. 2
- [18] Daniel Justus, John Brenner, Stephen Bonner, and Andrew Stephen McGough. Predicting the computational cost of deep learning models. In *2018 IEEE international conference on big data (Big Data)*, pages 3873–3882. IEEE, 2018. 2
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [20] Avisek Lahiri, Sourav Bairagya, Sutanu Bera, Siddhant Haldar, and Prabir Kumar Biswas. Lightweight modules for efficient deep learning based image restoration. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(4):1395–1410, 2020. 2
- [21] Wei-Sheng Lai, Yichang Shih, Lun-Cheng Chu, Xiaotong Wu, Sung-Fang Tsai, Michael Krainin, Deqing Sun, and Chia-Kai Liang. Face deblurring using dual camera fusion on mobile phones. *ACM Transactions on Graphics (TOG)*, 41(4):1–16, 2022. 1, 2
- [22] Kaidong Li, Nina Y Wang, Yiju Yang, and Guanghui Wang. Sgnet: A super-class guided network for image classification and object detection. In *2021 18th Conference on Robots and Vision (CRV)*, pages 127–134. IEEE, 2021. 1
- [23] Zhuoqun Liu, Meiguang Jin, Ying Chen, Huaida Liu, Canqian Yang, and Hongkai Xiong. Mfdnet: Towards real-time image denoising on mobile devices. *arXiv preprint arXiv:2211.04687*, 2022. 2
- [24] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022. 2

- [25] Haoran Lu, Yifei She, Jun Tie, and Shengzhou Xu. Half-unet: A simplified u-net architecture for medical image segmentation. *Frontiers in Neuroinformatics*, 16:911679, 2022. [3](#)
- [26] Ziwei Luo, Youwei Li, Lei Yu, Qi Wu, Zhihong Wen, Haoqiang Fan, and Shuaicheng Liu. Fast nearest convolution for real-time efficient image super-resolution. In *European Conference on Computer Vision*, pages 561–572. Springer, 2022. [2](#)
- [27] Long Ma, Tengyu Ma, Risheng Liu, Xin Fan, and Zhongxuan Luo. Toward fast, flexible, and robust low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5637–5646, 2022. [2](#)
- [28] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European conference on computer vision (ECCV)*, pages 116–131, 2018. [1](#)
- [29] Wenchu Ma, Xuemin Tu, Bo Luo, and Guanghui Wang. Semantic clustering based deduction learning for image recognition and classification. *Pattern Recognition*, 124:108440, 2022. [1](#)
- [30] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3883–3891, 2017. [7](#)
- [31] Seungjun Nah, Sanghyun Son, Suyoung Lee, Radu Timofte, and Kyoung Mu Lee. Ntire 2021 challenge on image deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 149–165, 2021. [7](#)
- [32] Krushi Bharatbhai Patel, Fengjun Li, and Guanghui Wang. Fuzzynet: A fuzzy attention module for polyp segmentation. In *NeurIPS’22 Workshop on All Things Attention: Bridging Different Perspectives on Attention*, 2022. [1](#)
- [33] Raiyan Rahman, Christopher Indris, Tianxiao Zhang, Kaidong Li, Brian McCornack, Daniel Flippo, Ajay Sharda, and Guanghui Wang. On the real-time semantic segmentation of aphid clusters in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6298–6305, 2023. [1](#)
- [34] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. [2](#), [4](#)
- [35] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. [1](#)
- [36] Pavan Kumar Anasosalu Vasu, James Gabriel, Jeff Zhu, Oncel Tuzel, and Anurag Ranjan. Mobileone: An improved one millisecond mobile backbone. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7907–7917, 2023. [2](#)
- [37] Wenhui Wang, Jifeng Dai, Zhe Chen, Zhenhang Huang, Zhiqi Li, Xizhou Zhu, Xiaowei Hu, Tong Lu, Lewei Lu, Hongsheng Li, et al. Internimage: Exploring large-scale vision foundation models with deformable convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14408–14419, 2023. [1](#)
- [38] Yuzhi Wang, Haibin Huang, Qin Xu, Jiaming Liu, Yiqun Liu, and Jue Wang. Practical deep raw image denoising on mobile devices. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI*, pages 1–16. Springer, 2020. [1](#), [2](#), [5](#), [6](#), [7](#)
- [39] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17683–17693, 2022. [1](#)
- [40] Lu Xu, Jiawei Zhang, Xuanye Cheng, Feng Zhang, Xing Wei, and Jimmy Ren. Efficient deep image denoising via class specific convolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3039–3046, 2021. [2](#)
- [41] Wenju Xu, Yuanwei Wu, Wenchu Ma, and Guanghui Wang. Adaptively denoising proposal collection for weakly supervised object localization. *Neural Processing Letters*, 51:993–1006, 2020. [1](#)
- [42] Wenhui Yang, Robby T Tan, Jiashi Feng, Jiaying Liu, Zongming Guo, and Shuicheng Yan. Deep joint rain detection and removal from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1357–1366, 2017. [7](#)
- [43] Yiyu Yang, Tianxiao Zhang, Guanyu Li, Taejoon Kim, and Guanghui Wang. An unsupervised domain adaptation model based on dual-module adversarial training. *Neurocomputing*, 475:102–111, 2022. [1](#)
- [44] Shuochao Yao, Yiran Zhao, Huajie Shao, ShengZhong Liu, Dongxin Liu, Lu Su, and Tarek Abdelzaher. Fastdeepiot: Towards understanding and optimizing neural network execution time on mobile and embedded devices. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*, pages 278–291, 2018. [2](#)
- [45] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5728–5739, 2022. [1](#)
- [46] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14821–14831, 2021. [1](#)
- [47] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for fast image restoration and enhancement. *IEEE transactions on pattern analysis and machine intelligence*, 45(2):1934–1948, 2022. [1](#)

- [48] He Zhang and Vishal M Patel. Density-aware single image de-raining using a multi-stream dense network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 695–704, 2018. 7
- [49] He Zhang, Vishwanath Sindagi, and Vishal M Patel. Image de-raining using a conditional generative adversarial network. *IEEE transactions on circuits and systems for video technology*, 30(11):3943–3956, 2019. 7
- [50] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017. 3
- [51] Wei Zhang, Wanshu Fan, Xin Yang, Qiang Zhang, and Dongsheng Zhou. Lightweight single-image super-resolution via multi-scale feature fusion cnn and multiple attention block. *The Visual Computer*, pages 1–13, 2023. 2
- [52] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6848–6856, 2018. 1
- [53] Yifeng Zhou, Xing Xu, Shuaicheng Liu, Guoqing Wang, Huimin Lu, and Heng Tao Shen. Thunder: Thumbnail based fast lightweight image denoising network. *arXiv preprint arXiv:2205.11823*, 2022. 2