

LIGHTWEIGHT NETWORK TOWARDS REAL-TIME IMAGE DENOISING ON MOBILE DEVICES

Zhuoqun Liu^{†*} Meiguang Jin[†] Ying Chen[†] Huaida Liu[†] Canqian Yang[†] Hongkai Xiong^{*}

[†] Alibaba Group

^{*} Shanghai Jiao Tong University

ABSTRACT

Deep convolutional neural networks have achieved great progress in image denoising tasks. However, their complicated architectures and heavy computational cost hinder their deployments on mobile devices. Some recent efforts in designing lightweight denoising networks focus on reducing either FLOPs (floating-point operations) or the number of parameters. However, these metrics are not directly correlated with the on-device latency. In this paper, we identify the real bottlenecks that affect the CNN-based models’ runtime performance on mobile devices: memory access cost and NPU-incompatible operations, and build the model based on these. To further improve the denoising performance, the mobile-friendly attention module *MFA* and the model reparameterization module *RepConv* are proposed, which enjoy both low latency and excellent denoising performance. To this end, we propose a mobile-friendly denoising network, namely *MFDNet*. The experiments show that *MFDNet* achieves state-of-the-art performance on real-world denoising benchmarks SIDD and DND under real-time latency on mobile devices. The code and pre-trained models will be released.

Index Terms— Image Denoising, Mobile-friendly Network Design

1. INTRODUCTION

With the rapid development of deep learning techniques, the performance of image denoising is improved significantly in recent years [1, 2, 3, 4]. However, deploying a state-of-the-art (SOTA) denoising model on resource constrained devices, such as mobile devices, remains challenging. On the one hand, although NPUs specifically optimized for deep neural networks are ubiquitous on mobile devices, most SOTA network architectures do not consider NPUs’ compatibility. Hence they cannot fully utilize the powerful NPUs. On the other hand, the requirement of high resolution processing (720p/1080p or even higher) for real applications exponentially increases the computational and memory access cost, which are key efficiency bottlenecks on mobile devices. There have been some attempts to design lightweight models

to promote mobile deployments [5, 2, 1]. These lightweight models reduce either the FLOPs or the number of parameters of the model. However, recent works show that reducing the FLOPs or the number of parameters does not necessarily lead to a low latency on mobile devices [6, 7]. For example, skip connections and multi-branch structures are commonly used design choices for low-level vision tasks [8, 4, 3, 2]. However, these operations can incur a high memory access cost, hindering fast inference on mobile devices. In addition to the model architectures, whether the operations are well optimized by NPUs is also essential when improving runtime performance [7]. Benefiting from the powerful parallel computing capability and specialized optimization for common operations, NPUs show great advantages over other processors when processing neural networks. However, operations well optimized by the NPUs are quite limited. Architectures containing NPU-incompatible operations will be partially processed by CPUs or GPUs. This introduces the additional data transfer cost between processors and increases the synchronization cost, leading to a severe overhead. For example, the ESA module [9], which is adopted by the winner in the runtime track of the NTIRE 2022 efficient super-resolution challenge [10], contains NPU-incompatible operations like max-pooling with a kernel size of 7 and a stride of 3 and bilinear interpolation with a scale larger than 2, which will be processed by GPUs. In this case, the feature maps have to be moved from NPUs to GPUs until the process is finished, which significantly slows down the inference speed of the model, as shown in Table 2.

In this paper, by conducting extensive experiments on an iPhone 11, we identify the preference of the Apple Neural Engine (ANE), which is a typical type of NPUs, for different network architectures and operations. Based on this, a mobile-friendly denoising network is proposed to significantly improve the model’s runtime performance on mobile devices while achieving enhanced denoising performance than other lightweight denoising networks on the real-world denoising benchmarks SIDD [11] and DND [12]. The experiments show that our model only takes around 20ms when processing a 720p image on an iPhone 11, which offers the possibility to process 720p images or videos in real-time on mobile devices.




Fig. 1. The architecture of the baseline model.

Table 1. Quantitative performance of different downsampling factors.

Model	Setting	Factor	MACs/G	Memory/M	PSNR/dB	Latency/ms
Baseline	C16_N8	1	12.61	2271	36.74	19.87
	C32_N12	↓ 2	24.66	1448	38.52	18.96
	C48_N16	↓ 4	20.30	822	38.52	16.87

2. NETWORK ARCHITECTURE

In this section, we build a mobile-friendly denoising network from scratch. To ensure that our model can have efficient runtime performance on mobile devices, we only use operations compatible with the ANE. The results of models with different capacities are shown in the experiment section.

2.1. Baseline Model

As mentioned before, to ensure real-time on-device runtime performance, the memory access cost has to be strictly controlled. Therefore, we start from a DnCNN-like [13] plain topology, as shown in Fig.1 left, which contains only the well-optimized 3×3 convolutions and ReLU activation functions. To reduce the memory access cost, only one residual connection is adopted. We start from this architecture because the DnCNN-like architectures have demonstrated their effectiveness in the image denoising task [13]. And we remove batch normalization to reduce potential artifacts [14, 15].

Based on this plain topology, we adopt multiple 3×3 convolutions with stride 2 to downsample the input, followed by N Conv- 3×3 +ReLU blocks with a local skip connection. At the end of the network is a PixelShuffle [16] for reconstruction. The modified model is shown in Fig. 1 right.

Downsampling in the beginning brings two benefits. First, most of the operations in the model are executed at a low resolution, which significantly reduces the total memory access cost, as shown in Table 1. Second, the overall increased receptive field of the network enables the network to capture more contextual information and improves denoising performance.

We compare the denoising performance of the models with different downsampling factors on the SIDD validation dataset and the runtime performance on an iPhone 11. Both the runtime performance and the total memory access cost (read and write) are evaluated with a 720p input. The results are summarized in Table 1, where C and N represent the number of channels and the number of Conv- 3×3 +ReLU blocks, respectively. Results show that the pre-downsampling enables a wider and deeper model under a similar latency by




Fig. 2. The architecture of MFA.

Table 2. Performance of different attention mechanisms.

Model	Attention	MACs/G	Memory/M	PSNR/dB	Latency/ms
Baseline	ESA [9]	12.78	562	38.59	43.23*
	SCA [1]	12.52	400	38.41	15.20
	HFAB [17]	14.60	574	38.51	18.94
	MFA	12.85	448	38.60	15.61

significantly reducing the memory access cost, thus leading to an enhanced denoising performance. So we choose the $4 \times$ downsampling model as our final baseline model.

2.2. Attention

Attention mechanisms have been extensively studied in low-level vision tasks [9, 17, 1, 2, 4]. In previous works, attention modules often adopt complex topology or use operations that are not optimized by ANE [9], which affects the runtime performance severely, as shown in Table 2. In this paper, we propose a simple yet effective mobile-friendly attention module, MFA. Its architecture is shown in Fig.2. Our proposed architecture has three advantages. First, in order to reduce the memory access cost and boost the runtime performance of the module, we discard the complex topology and retain only the necessary residual connection for multiplying the learned attention maps with the input feature maps as a spatial attention mechanism. Second, we further downsample the feature maps to reduce the latency and enlarge the receptive field of the attention module. Finally, we only use operations well-optimized by ANE, like the common 3×3 convolution, the ReLU activation function and bilinear interpolation with a scale of 2 to further boost the efficiency.

In practice, we propose a mobile-friendly denoising block, MFDB, which contains K Conv- 3×3 +ReLU blocks followed by one MFA. The width of the MFA is set to $\frac{1}{4}$ of the width of the model. In Table 2, we compare MFA with the attention modules commonly used in low-level vision tasks, including ESA [9], SCA [1], and HFAB [17]. It can be seen from the table that MFA achieves the best denoising performance. Meanwhile, MFA is also close to SCA in terms of latency, with a difference of only 0.4ms. Note that although ESA has a comparable memory access cost compared to other methods, its on-device latency is still significantly higher due to incompatible operations mentioned in section 1.

2.3. Activation

Although Rectified Linear Unit (ReLU) has been extensively used in low-level vision tasks, many SOTA methods tend to

Table 3. Quantitative performance of different activation functions.

Model	Activation	PSNR/dB	Latency/ms
Baseline+ MFA	ReLU	38.60	15.61
	GELU	38.71	16.11
	PReLU	38.71	15.84
	LReLU	38.73	15.74

replace ReLU with other activation functions, such as GELU, LReLU, PReLU [1, 17] for better performance. In order to test the compatibility of the ANE for different activation functions and the potential performance gains, we replace ReLU with several different activation functions, as shown in Table 3. The results in the table show that these activation functions are all well-optimized by the ANE. Replacing ReLU with LReLU results in a performance gain of 0.13dB on the validation dataset of SIDD, while the inference time on an iPhone 11 remains nearly unchanged. We, therefore, replace all the ReLU in our baseline model with LReLU.

2.4. Lightweight Feature Extraction

To further improve the feature extraction and representation capabilities of the downsampling module, we replace the stride-2 convolution with the Haar transform, which can be efficiently implemented in a convolution form with a kernel size of 2. Compared to a stride-2 convolution, Haar transform is invertible, which ensures that the frequency information of the input can be effectively captured in a lossless manner. This is very helpful in preserving textures [18]. In addition, Haar transform is a more compact feature representation, meaning that fewer hidden channels between the downsampling blocks can lead to a better denoising performance, which also reduces the latency. Table 5 shows that using the Haar transform to extract features significantly improves the model’s denoising and on-device runtime performance. It brings 0.22dB performance gain on SIDD and 1.32ms latency reduction on an iPhone 11.

2.5. Reparameterization

The idea of reparameterization was first proposed by RepVGG [19]. The core idea behind reparameterization is to parameterize a plain topology with the parameters transformed from a more complex topology (e.g., multi-branch topology). Model reparameterization is very beneficial in the design of




Fig. 3. The structure of RepConv.




Fig. 4. The overall architecture of MFDNet.

Table 4. Comparisons between different reparameterization methods.

Model	Reparameterization Method	PSNR/dB
MFDNet w/o reparameterization	/	38.95
	ECB [7]	38.97
	RRRB [17]	38.95
	RepConv	38.99

mobile-friendly models. Complex topologies, such as multi-branch structures, can significantly increase the memory access cost and slow down the inference. However, the plain topology lags in the feature extraction capability, resulting in compromised model performance. Model reparameterization can be used to address this issue. In the training phase, the model takes advantage of the complex topology to enrich the feature representation and bring performance gains. In the testing phase, the model reparameterization method is used to simplify the topology and improve the inference speed of the model without a performance drop.

In this paper, we use the expand-and-squeeze topology for training, named RepConv, since the wider features result in better feature representation. As shown in Fig.3, the topology consists of two 1×1 convolutions, a 3×3 convolution, and a skip connection in the training phase. In the testing phase, we merge three convolutions and a skip connection into a single 3×3 convolution by reparameterization, thus eliminating the cascaded and multi-branch structures. In Table 4, we compare the RepConv with other reparameterization methods commonly used in low-level vision tasks, including ECB [7] and RRRB [17]. Results show that RepConv achieves the best denoising performance with an improvement of 0.04dB on SIDD.

2.6. Summary

At this point, we build our mobile-friendly denoising network step by step from the baseline model. The architecture of MFDNet is shown in Fig.4. For MFDNet, we set both the number of MFDBs (M) and RepConv+LReLU blocks in each MFDB (K) to 3.

3. EXPERIMENT

In this section, we first analyze the role of the different model design choices mentioned in the previous sections in terms of both denoising and runtime performance. We then apply our model to the real-world denoising benchmarks SIDD and

Table 5. Ablation study of different components in MFDNet.

	MFA	ReLU → LReLU	Stride-2 Conv → Haar Transform	RepConv	PSNR/dB	Latency/ms
Baseline					38.54	12.82
	✓				38.60	15.61
	✓	✓			38.73	15.74
	✓	✓	✓		38.95	14.42
	✓	✓	✓	✓	38.99	14.42

DND. To test the models’ on-device runtime performance, we execute them 300 times on an iPhone 11 to get the average elapsed time. Note that in all experiments, latency with the * notation indicates that the model contains ANE incompatible operations and is processed by CPUs/GPUs. The computational complexity, latency and total memory access cost in all experiments are evaluated with a 720p input.

3.1. Ablation Study


We conduct our ablation study on the validation dataset of SIDD and measure the models’ latency on an iPhone 11. We train our model using the Adam optimizer with the learning rate initialized with $4e-4$ and halved every 100k iterations. We train the model for a total of 1M iterations with a batch size of 32 and use cropped patches of 256×256 from SIDD-Medium as the training dataset.

We start from the baseline model with a downsampling factor of 4, as mentioned in section 3.1, and build the MFDNet step by step. We set the depth of the baseline model (N) to 9 and the width to 48. Table 5 shows the effectiveness of different components.

3.2. Application

We evaluate our model on two real-world denoising benchmarks, SIDD and DND. All training settings are identical to those used in the ablation experiments. To see the generalization ability of our model on different NPUs, we also test the on-device latency on iPhone 14 Pro. The results are summarized in Table 6. For MFDNet, we also design two models with different scales: MFDNet-S and MFDNet-L. For MFDNet-S, we reduce M and K to 1. For MFDNet-L, we keep $K = 3$ unchanged and increase M to 6. We compare our model with models proposed in [13, 20, 1, 3, 21]. Note that as these models are not designed specifically for mobile devices, we prune them in terms of depth and width to ensure that they can run on mobile devices. For DnCNN [13], we set both the width and depth to 12, and for DnCNN-S, we set the width to 16 and the depth to 3. For RIDNet [20], we trim the number of channels to 8 and set the channel reduction to 2. For NAFNet [1], the number of channels is also trimmed to 8 and the number of blocks is reduced to 7. For HINet [3], we adjust the number of channels and the depth of the model to 8 and 2, respectively.

We can see from Table 6 that, MFDNet and MFDNet-L achieve the best denoising performance with the lowest latency compared to models with comparable computational

**Fig. 5.** Qualitative comparisons on SIDD.**Table 6.** Quantitative performance of different methods.

Model	MACs /G	Memory /M	Latency/ms		SIDD		DND	
			iPhone 11	iPhone 14 Pro	PSNR	SSIM	PSNR	SSIM
DnCNN-S [13]	2.77	583	11.7	5.7	33.84	0.877	36.41	0.910
NAFNet [1]	3.81	2974	1384.5*	25.9	38.66	0.951	38.74	0.945
MFDNet-S	2.34	142	8.3	4.0	36.93	0.932	38.09	0.939
DnCNN [13]	12.88	2720	23.4	12.5	36.24	0.921	38.13	0.936
HINet [3]	11.68	2178	84.7*	23.3	37.83	0.943	36.82	0.931
MFDNet	11.46	384	14.2	6.2	38.90	0.952	39.06	0.947
RIDNet [20]	20.39	4224	75.9	30.0	38.01	0.945	38.88	0.947
PMRID [21]	15.07	4448	80.8	40.1	38.96	0.953	38.82	0.948
MFDNet-L	21.81	684	24.1	8.9	39.01	0.954	39.15	0.948

complexity. For models with computational complexity under 5GMACs, MFDNet-S achieves the best runtime performance with the denoising performance far exceeds that of DnCNN. Note that the proposed MFDNet is fully compatible with both iPhone 11 and iPhone 14 Pro, which indicates its excellent generalization ability. In contrast, NAFNet and HINet are only compatible with iPhone 14 Pro. The layer normalization operation in NAFNet and the instance normalization operation in HINet are incompatible with iPhone 11, which severely affect the inference speed of these models. By removing the incompatible operations from NAFNet and HINet, the latency on iPhone 11 will reduce to 71.3ms and 47.3ms, respectively.

4. CONCLUSION

In this paper, we identify the network architectures and operations that can run on NPUs with low latency and excellent denoising performance through extensive analysis and experiments. Based on that, we build a mobile-friendly denoising network from scratch. Experiments show advances of our method in terms of both denoising and runtime performance. We hope this work will promote the application of CNN-based denoising models on mobile devices.

5. REFERENCES

- [1] L. Chen, X. Chu, X. Zhang, and J. Sun, “Simple baselines for image restoration,” *arXiv preprint arXiv:2204.04676*, 2022.
- [2] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, and M.-H. Yang, “Restormer: Efficient transformer for high-resolution image restoration,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5728–5739.
- [3] L. Chen, X. Lu, J. Zhang, X. Chu, and C. Chen, “Hinet: Half instance normalization network for image restoration,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021, pp. 182–192.
- [4] Z. Wang, X. Cun, J. Bao, W. Zhou, J. Liu, and H. Li, “Uformer: A general u-shaped transformer for image restoration,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17683–17693.
- [5] L. Xu, J. Zhang, X. Cheng, F. Zhang, X. Wei, and J. Ren, “Efficient deep image denoising via class specific convolution,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021, vol. 35, pp. 3039–3046.
- [6] P. K. A. Vasu, J. Gabriel, J. Zhu, O. Tuzel, and A. Rajan, “An improved one millisecond mobile backbone,” *arXiv preprint arXiv:2206.04040*, 2022.
- [7] X. Zhang, H. Zeng, and L. Zhang, “Edge-oriented convolution block for real-time super resolution on mobile devices,” in *Proceedings of the 29th ACM International Conference on Multimedia*, 2021, pp. 4034–4043.
- [8] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, M.-H. Yang, and L. Shao, “Multi-stage progressive image restoration,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2021, pp. 14821–14831.
- [9] J. Liu, W. Zhang, Y. Tang, J. Tang, and G. Wu, “Residual feature aggregation network for image super-resolution,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2020, pp. 2359–2368.
- [10] F. Kong, M. Li, S. Liu, D. Liu, J. He, Y. Bai, F. Chen, and L. Fu, “Residual local feature network for efficient super-resolution,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2022, pp. 766–776.
- [11] A. Abdelhamed, S. Lin, and M. S. Brown, “A high-quality denoising dataset for smartphone cameras,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1692–1700.
- [12] T. Plotz and S. Roth, “Benchmarking denoising algorithms with real photographs,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1586–1595.
- [13] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, “Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising,” *IEEE transactions on image processing*, vol. 26, no. 7, pp. 3142–3155, 2017.
- [14] X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y. Qiao, and C. Change Loy, “Esrgan: Enhanced super-resolution generative adversarial networks,” in *Proceedings of the European conference on computer vision (ECCV) workshops*, 2018, pp. 0–0.
- [15] B. Lim, S. Son, H. Kim, S. Nah, and K. Mu Lee, “Enhanced deep residual networks for single image super-resolution,” in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017, pp. 136–144.
- [16] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, “Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1874–1883.
- [17] Z. Du, D. Liu, J. Liu, J. Tang, G. Wu, and L. Fu, “Fast and memory-efficient network towards efficient image super-resolution,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 853–862.
- [18] P. Liu, H. Zhang, K. Zhang, L. Lin, and W. Zuo, “Multi-level wavelet-cnn for image restoration,” in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2018, pp. 773–782.
- [19] X. Ding, X. Zhang, N. Ma, J. Han, G. Ding, and J. Sun, “Repvgg: Making vgg-style convnets great again,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021, pp. 13733–13742.
- [20] S. Zhuo, Z. Jin, W. Zou, and X. Li, “Ridnet: recursive information distillation network for color image denoising,” in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2019, pp. 0–0.
- [21] Y. Wang, H. Huang, Q. Xu, J. Liu, Y. Liu, and J. Wang, “Practical deep raw image denoising on mobile devices,” in *European Conference on Computer Vision*. Springer, 2020, pp. 1–16.