

Feature-Align Network with Knowledge Distillation for Efficient Denoising

Lucas D. Young*
Jon Morton
Xiaoyu Xiang

Fitsum A. Reda*[†]
Jun Hu
David Liu
Facebook Inc.

Rakesh Ranjan
Yazhu Ling
Vikas Chandra

Abstract

We propose an efficient neural network for RAW image denoising. Although neural network-based denoising has been extensively studied for image restoration, little attention has been given to efficient denoising for compute limited and power sensitive devices, such as smartphones and wearables. In this paper, we present a novel architecture and a suite of training techniques for high quality denoising in mobile devices. Our work is distinguished by three main contributions. (1) The Feature-Align layer that modulates the activations of an encoder-decoder architecture with the input noisy images. The auto modulation layer enforces attention to spatially varying noise that tends to be "washed away" by successive application of convolutions and non-linearity. (2) A novel Feature Matching Loss that allows knowledge distillation from large denoising networks in the form of a perceptual content loss. (3) Empirical analysis of our efficient model trained to specialize on different noise subranges. This opens an additional avenue for model size reduction by sacrificing memory for compute. Extensive experimental validation shows that our efficient model produces high quality denoising results that compete with state-of-the-art large networks, while using significantly fewer parameters and MACs. On the Darmstadt Noise Dataset benchmark, we achieve a PSNR of 48.28dB, while using 263× fewer MACs and 17.6× fewer parameters than the state-of-the-art network, which achieves 49.12dB.

1. Introduction

Image acquisition is inevitably contaminated by noise due to various environmental effects. Noisy images are especially more prevalent in small devices such as smartphones and wearables. Often these devices are characterized by small sensors and limited light intake ability, leading to low perceptual quality images. Thus, efficient image

*Equal contribution.

[†]Affiliated with Facebook at the time of this work.



(a) Noisy Input, PSNR = 18.76

(b) Ground Truth



(c) CycleISP [44], PSNR = 40.44, 702.73 GMACs/MP
(d) Our Model, PSNR = 36.33, 2.67 GMACs/MP

Figure 1. Our efficient denoising method compared to the existing state-of-the-art denoising method, CycleISP when applied to an image in the Darmstadt Noise Dataset [30]. PSNR is calculated in the RAW domain. Our method produces similar results with 263x fewer MACs.

denoising algorithms are highly desirable to restore image quality in mobile devices.

Image denoising is a classical yet actively studied topic

in image restoration [9, 6, 10, 13, 42, 43]. Recent attention for the problem focuses on applying deep neural networks to RAW sensor data or to images obtained after post-processing with the device’s image signal processor (ISP) [5, 37, 28, 45, 21, 48, 29, 44]. Several of these networks are based on the U-Net [31] architecture. They produce high quality denoising and leverage large quantities of training data. They are, however, computationally too expensive to run at real-time on compute and power sensitive edge devices.

The success of the U-Net architecture [31] has been celebrated in learning complex mappings for several dense prediction tasks, such as semantic segmentation [38], depth estimation [25] or image synthesis [40]. However, its direct deployment on mobile devices with limited compute-budget is not optimal. Trivial reduction of U-Net parameters is not sufficient to learn robust models. As such, efforts have been spent on efficient building blocks. MobileNets [18, 36, 17] proposed a suite of efficient blocks for efficient learning, starting with a depth-wise convolution, inverted residuals, and later with optimized blocks searched via neural architecture search algorithms. Similarly, [19] introduced squeeze-excitation-networks, and [47] proposed Shuffle Nets that proved to be effective for mobile applications.

Another recent phenomenon observed in U-Net is that successive application of convolutions and non-linearity tend to “wash away” important image cues in deeper layers. This phenomenon was first studied in [27] for the task of conditional image synthesis and shown to lead to sub-optimal image synthesis outcomes. The same work addresses this issue with a spatially-adaptive normalization layer. Specifically, it modulates the U-Net features with transformations learned from the input semantic layout, and shows improved learning outcomes in image synthesis. In our experiments, we observed this phenomenon to be even more pronounced in light-weight image denoising U-Net architectures.

In this work, we build upon the U-Net [31] architecture and introduce an efficient neural network for high quality RAW image denoising. We employ a variant of the MobileNetV2 [36] efficient module in place of normal convolutions. Specifically, our variant MobileNetV2 [36] uses group convolutions in place of depth convolutions, which further improves memory efficiency. To alleviate the “washing away” of features in U-Net, we take inspiration from the SPADE architecture [27] and propose to modulate the convolutional features at each layer with the input noisy images. This auto modulation of features with transformations learned directly from the input noisy images allows our models to attend to spatially varying noise. Experimental validations suggest the effectiveness of this mechanism in our efficient U-Net architecture. We refer to our modula-

tion layer as “Feature-Align”.

Student-teacher learning has proven to be effective in learning efficient models [39]. Knowledge distillation allows the transfer of richer knowledge to light-weight models. Our experiments also prove the effectiveness of knowledge distillation in image denoising, especially in settings with limited training datasets. In this work, we build up on the student-teacher training mechanism and propose a new Feature Matching Loss that performs knowledge distillation in the form of a perceptual content loss. Specifically, we extract deep multi-scale features from our efficient model’s output image and optimize them to match the features extracted from the clean ground-truth image, in a method similar to the VGG perceptual loss [22]. One distinction of our features matching loss is that we use a large pre-trained image denoising network, for instance [5], to perform the feature extraction, instead of ImageNet [11] trained image classification networks [34]. Using a task-specific feature extractor, in our case an image denoiser, allows the transfer of knowledge at deeper representation levels. As we will show in experiments, our Feature Matching Loss leads to more crisp and realistically denoised images than that of training with standard student-teacher distillation or regular perceptual losses.

Our work also explores overparameterization. Specifically, we train multiple networks targeting for specific noise subranges. Use of an array of light-weight specialized models allows us to improve output image quality solely by increasing additional parameter count without increasing inference time. Experimental validations suggest the effectiveness of this technique.

In summary, the primary contributions of this paper are:

- The Feature-Align layer that modulates the activations of an encoder-decoder architecture with the input noisy images. The auto modulation layer enforces attention to spatially varying noise that tend to be “washed away” by successive application of convolutions and non-linearity.
- A novel feature matching loss that allows knowledge distillation from large denoising networks in the form of a perceptual content loss.
- Empirical analysis of our efficient model trained to specialize on different noise subranges, which opens the additional avenue for model size reduction by sacrificing memory for compute.

2. Related Work

Single-frame image denoising is a fundamental problem in image processing and computer vision. Prior methods such as BM3D [9] and non-local means [6] rely on hand-engineered algorithms [10, 13, 42, 43]. With the

introduction of data-driven neural network-based methods [37, 28, 45, 21, 48, 29], datasets with paired noisy and noise-free images became sought after [1, 8, 2]. To resolve the scarcity of paired noisy and clean RAW images, other works proposed to simulate noisy images, in which clean images are post-processed by a simple addition of white gaussian noise (AWGN). Recent works observe that true sensor noise exhibits characteristics unlike AWGN, and thus they propose robust noise modeling mathematics by considering the physical properties of sensors [5, 30, 41, 4, 15, 14, 12]. The work of [5] and CycleISP [44] further synthesize RAW images from real-life RGB images by inverting the ISP pipeline. Public real-life image datasets, such as MIRFLICKR [20], are leveraged with artificial noise modeling to scale-up the training dataset. This approach has been demonstrated in achieving state-of-the-art performance on public benchmarks, suggesting the generalizability of the approach.

While these works introduce valuable techniques for denoising, they do not study less computationally expensive variants of their methods. PMRID [41] is the first to address this disparity in neural denoising. Specifically, it introduces an efficient model architecture using separable convolutions. It also builds upon the variance stabilizing transform [3, 26] by introducing the k-sigma transform, which normalizes data such that per-pixel variance is not dependent on the ISO of the exposure. PMRID also describes a noise modeling method similar to [5]. The work we propose here shares similar goals to PMRID. However, a direct quantitative comparison could not be performed at the time of this paper’s writing because PMRID’s trained model or test datasets are not publicly available.

Our proposed technique builds upon previous works and introduces an efficient denoising neural network and a suite of training techniques. In particular, it incorporates the following prior techniques:

- Noise modeling of our target camera’s sensor, Sony IMX258, with a method similar to PMRID and UPI, to add realistic artificial noise to ground truths.
- Unprocessing [5] the MIRFLICKR-1m [20] dataset to create a large RAW training dataset.
- Incorporation of Bayer Augmentation [23] in our training pipeline.
- Use of Bayer Unification [23] to use ground truths from the sensors in the SIDD [1] and SID [8] dataset.
- Incorporation of the k-sigma transform [41]

Lastly, we combine inspiration from a method traditionally applied to super-resolution, perceptual loss [22], and knowledge distillation commonly applied for student-teacher knowledge distillation to propose a novel Feature

Matching Loss function. As we will show in experiments, our new loss produces results that are better than using either of these established techniques.

3. Method

Given a noisy RAW image I_n , we aim to learn efficient neural networks to denoise and produce clean RAW image \tilde{I} , in a data-driven way. Mathematically, the denoising problem is formulated as,

$$\tilde{I} = \mathcal{F}(I_n), \tag{1}$$

where \mathcal{F} is an efficient neural network we aim to learn. We adopt the techniques introduced in [5] to create synthetic RAW training pairs. Denote the noise sampled from a noise model n and the unprocessed “clean” RAW image I . Note that, the RGB images we present here to visually motivate or validate our methods are obtained after processing the denoised RAW image \tilde{I} with the camera ISP. We use I_n^f and \tilde{I}^f to denote the final noisy and final denoised RGB image, respectively.

3.1. Model Architecture

We realize our neural network \mathcal{F} with an efficient U-Net [31] architecture. We propose techniques to boost the learning ability light-weight U-Net models. In particular, we make the following architectural changes: 1) To achieve efficient learning, we replace the convolutional layers with a variant of the MobileNet-V2 block [32]. 2) To reduce the memory footprint of skip layers, we introduce shrunken skip-connection, implemented with a point-wise convolution. 3) To minimize the “washing away” phenomenon of important image features in deep convolutional layers, as pointed out in the SPADE architecture [27], we introduce a new Feature-Align layer. The following subsections describe each design choice. Figure 4 illustrates our proposed efficient architecture.

A variant of MobileNet-V2 block: In this work, we use a variant of the MobileNet-V2 efficient block to realize each contracting and expanding layer in our U-Net. Specifically, as illustrated in Figure 2, we replace MobileNet-V2’s depth-wise convolution with a group convolution. This choice allows efficient memory access.

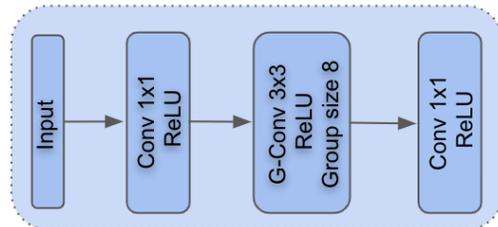


Figure 2. ARNet-block, a variant of MobileNet-V2 [32]

Shrunked Skip Connections: Skip-connections in a U-Net architecture allow joint processing of multi-scale features. High-resolution features from the contracting layers are combined with low-resolution features from the expanding layers in the decoder layer. In memory constrained devices, storing encoder activations until their processing by decoder layers could be memory inefficient. To alleviate this, as illustrated in Figure 4, we apply a point-wise convolution to shrink the encoder features. At the decoder layer, we replicate the contracted features to match the size of the expanding features to fusion the features.

Input Feature-Align Layer: Prior neural network-based denoising techniques take in noisy inputs and directly process them with a stack of convolutions and non-linearity. While this is also a common technique for various dense prediction tasks, in our experiments, we observed such direct processing to bias our light-weight models to exhibit global over- or under-denoising. As shown in our experiments, we attribute this to the recently studied phenomenon of "washing away" of input details during processing with a stack of multi-scale layers. In our case, we tightly integrate the input noisy information at each layer of our network with a Feature-Align layer, to enable our models to learn to denoise images in a spatially adaptive manner.

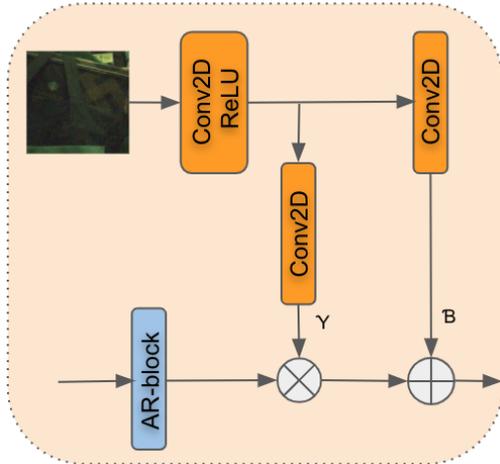


Figure 3. Input Feature-Align Layer, inspired by the SPADE architecture [27]

Let $\mathbf{F}^i \in \mathbb{R}^{N \times C^i \times H^i \times W^i}$ denote the feature map of layer i , with N, C^i, H^i, W^i being the batch size, channel count, height and width, respectively, of the feature map. Our Feature-Align layer, in a similar way as a Batch Norm layer, or the Spatially-Adaptive De-normalization (SPADE) layer, applies a scale and a bias to affinely transform the inputs, guided by the noisy input image. Specifically, we compute the pixel-wise scale and bias parameters based on the noisy input image, as illustrated in Figure 3. Mathematically, the affine transformation of a feature map is given by

$$\mathbf{G}^i = \gamma^i \cdot \mathbf{F}^i + \beta^i, \quad (2)$$

with $\gamma^i \in \mathbb{R}^{N \times C^i \times H^i \times W^i}$ and $\beta^i \in \mathbb{R}^{N \times C^i \times H^i \times W^i}$ being the scale and bias parameters that are a function of the input noisy image.

3.2. Loss functions

Charbonnier Loss: Our primary loss function is the Charbonnier Loss [7] over the denoised image, given by

$$\mathcal{L}_{charb} = \sqrt{(\tilde{\mathbf{I}} - \mathbf{I})^2 + c^2} \quad (3)$$

where c is a constant parameter ($1e - 6$ in our implementation).

RGB Perceptual Loss: UPI [5] describes a variant of their network in which training loss is computed after postprocessing the predicted image and ground truth with a minimalist ISP implementation. This incentivizes the network to predict RAW images that after postprocessing, yield RGB images that closely match the ground truth after postprocessing, rather than to optimize for RAW signal fidelity. We expand upon this idea with the RGB Perceptual Loss, in which a perceptual loss such as that described in [22] is computed upon the predicted image and ground truth after postprocessing. After the postprocessing step, the implementation is identical to that of [22]; the loss is defined as the difference between $\tilde{\mathbf{I}}^f$ and \mathbf{I}^f in content and style in activations of a VGG16 classification network [35] pretrained on ImageNet.

Simple Knowledge Distillation: The additional layers of conventional larger networks can enable the model to encode image features at a higher, more perceptual level and make imaginative judgments to "fill in the blank" of a missing texture obscured by noise. Since this capability is diminished with smaller networks, image outputs from efficient networks are prone to appear oversmoothed. As a result, knowledge distillation [16] is applicable to the challenge of creating an efficient model with perceptually pleasing image outputs.

Another observation from experimenting with a large denoising model (the same U-Net architecture used in [5]) is that a denoising network can learn to perform the identity transformation as desired - that is, if a noise-free image is fed to the network as an input, the output is virtually unaltered. This observation enables a simple implementation of knowledge distillation in the form of a loss function defined as the difference between the student network's predicted image, $\tilde{\mathbf{I}}_{student}$ and the large teacher network's predicted image, $\tilde{\mathbf{I}}_{teacher}$:

$$\mathcal{L}_{k.d.} = |\tilde{\mathbf{I}}_{student} - \tilde{\mathbf{I}}_{teacher}| \quad (4)$$

Feature Matching Loss for Knowledge Distillation: Expanding on Simple Knowledge Distillation and perceptual

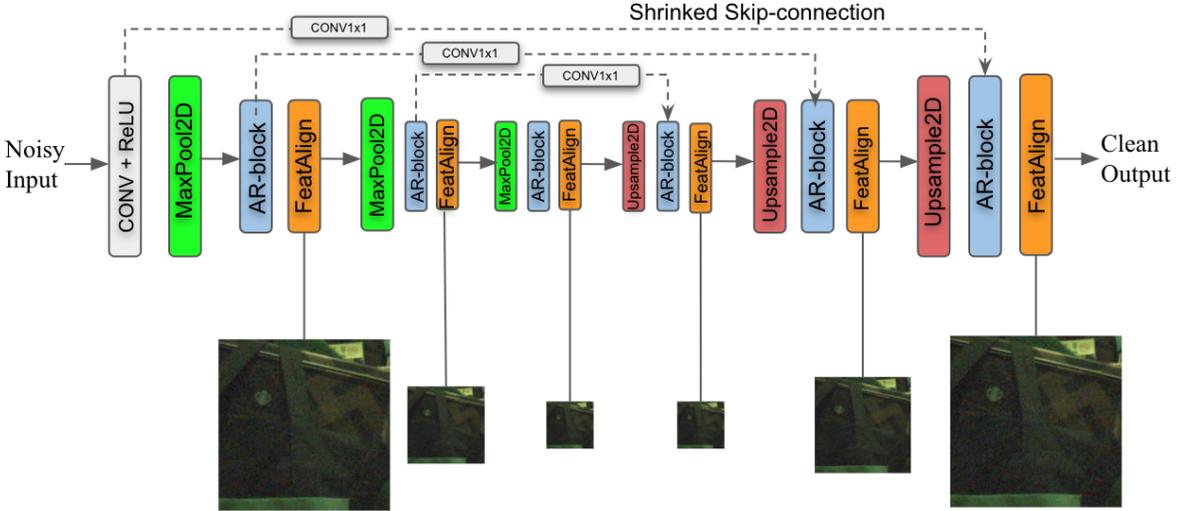


Figure 4. Our U-Net-based model architecture. Shrunked skip-connections combine high-resolution features in encoding layers with low-resolution features in decoder layers. A Feature-Align layer combines the original input features at different resolutions with the output for each efficient ARNet block.

loss, our novel Feature Matching Loss implements knowledge distillation in the form of a perceptual loss function. The Feature Matching Loss is defined as the difference in content and style between the activations of the teacher network on the predicted image and the ground truth image. The implementation is identical to [22], except the classifier model is substituted for the teacher network. By combining knowledge distillation with perceptual loss, the capability of large networks to recover fine details obscured by noise is specifically targeted for distillation.

3.3. Noise Subrange Model Array

The number of parameters in our efficient model architecture is limited in comparison to large model architectures traditionally used for neural denoising. In contrast, the problem space encountered in photography on mobile devices that typically have cameras with small aperture sizes widely ranges from minimal noise to extreme noise, requiring vastly different denoising strategies depending on the noise level. With a limited number of parameters, it is challenging to create an efficient network that optimally responds to every possible noise level. Thus, we offer the method of a Noise Subrange Model Array (NSMA), in which the range of noise levels is partitioned into n subranges where an individual model is trained for each noise subrange. To partition the noise levels, we refer to the regression between $\log a$ and $\log b$ in Figure 2 of the supplemental material and partition the x-axis into n parts. Given the global minimum signal dependent noise parameter, a_{min} , and global maximum signal dependent noise parameter a_{max} , Equation 5 and 6 describe the minimum and maximum a parameters for artificial noise in training for the

zero-indexed i th model in the array of n models.

$$\log a_{min_i} = \log a_{min} + i/n \cdot (\log a_{max} - \log a_{min}) \quad (5)$$

$$\log a_{max_i} = \log a_{min} + (i + 1)/n \cdot (\log a_{max} - \log a_{min}) \quad (6)$$

In testing, we observe the annotated a noise parameter of the image as described in supplemental Section 1 and select the corresponding model from the model array.

3.4. Training and Implementation Details

Our models are trained via adding synthetic noise to ground truths sourced from [1], [8], and the unprocessing method described in [5]. Bayer Augmentation and Unification [23] and the k-sigma transform [41] are incorporated into our method. A full description of training and implementation details is provided in supplemental Section 2.

3.5. Test Dataset

While our denoising network is trained with artificial noise, we test our network with real noise. We contribute the publicly available Feature-Align Paired Test Dataset of noisy RAW images paired with corresponding noise-free ground truths. This dataset consists of carefully aligned pairs of noisy short exposures that correspond to noise-free long exposures. The method used to collect this dataset is provided in supplemental Section 3.

4. Results

4.1. Experiments

In this section we evaluate denoising results of various experiments using the paired noisy and ground truth test set described in Section 3.5. We evaluate our results using traditional metrics for image quality, PSNR and SSIM, both before and after processing the images using the minimalist ISP pipeline from [5] with modified white balance gains for our sensor. We also evaluate our results qualitatively against the ground truth image. To display texture preservation, all images are zoomed in to the same highly textured region of an image in the test set. We quantify the computational efficiency of networks by the number of multiply-accumulate operations (MACs) required by the network.

Table 1 compares our Feature-Align method against a similarly sized model that uses the wavelet transform in up-sampling and down-sampling as a baseline to approximate the prior state-of-the-art [24, 33]. To simplify the comparison, the Noise Subrange Model Array is not used in this experiment. A plain U-Net that does not include the wavelet transform or the Feature-Align layer, along with the Feature-Align model with shrunk skip connections removed are also included in the comparison. With traditional metrics, the Feature-Align model performs virtually equivalently to the wavelets model. Qualitatively, the Feature-Align model shows sharper edges and less blurring of fine details. To show additional generalization of the Feature-Align layer, we compare the wavelet transform model to the Feature-Align model on collected noisy RAW Google Pixel-4 inputs in Figure 5. We note that the qualitative benefits of the Feature-Align layer are more easily observed on our Pixel-4 inputs than our paired test dataset. Furthermore, the Feature-Align model without shrunk skip connections performs the worst among the 4 networks. This indicates that the Feature-Align layer is not redundant to the shrunk skip connections despite the similarity of these techniques.

In Table 2 we examine the effect that the NSMA has on the Feature-Align model. The NSMA, which is comprised of 4 models, improves performance in every metric and demonstrates superior recovery of texture.

Using the Feature-Align model and NSMA, in Table 3 we compare the effect of altering the loss function. Since RGB Perceptual Loss, Simple Knowledge Distillation, and Feature Matching Loss are targeted at perceptual image quality rather than pure signal fidelity, for a fair comparison we include the perceptual metric, LPIPS [46]. A lower LPIPS score indicates greater perceptual similarity. Simple Knowledge Distillation performs the best in the traditional PSNR and SSIM metrics, while Feature Matching Loss performs the best in the LPIPS metric. Qualitatively, the outputs from the Feature Matching Loss models are sharper and have the most texture present. Charbonnier Loss out-



Figure 5. Left: Source image from Google Pixel-4. Red box highlights selected crop for comparison. Middle: Model output with wavelet transform. Right: Model output with our Feature-Align layer. Middle and right images are brightened for clarity.

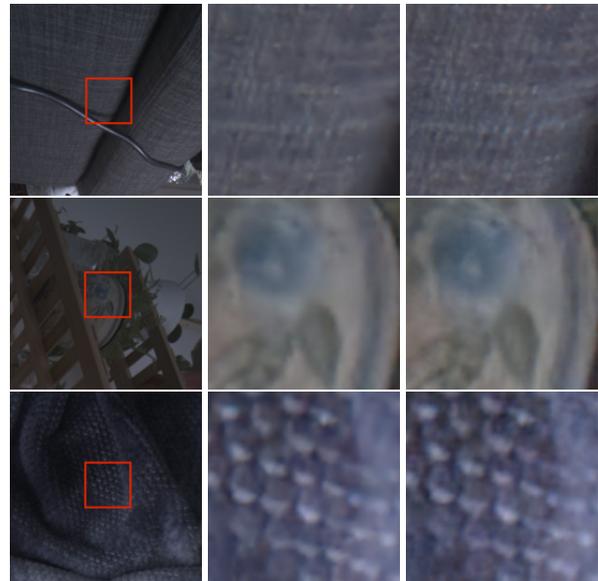


Figure 6. Summary of our improvements. Left: Source image. Red box highlights selected crop for comparison. Middle: Model output with wavelet transform, no NSMA, and Charbonnier Loss. Right: Model output with our Feature-Align layer, NSMA ($n=4$), and Feature Matching Loss. Middle and right images are brightened for clarity. Our method shows improved texture restoration.

performs Feature Matching Loss in the traditional metrics, further demonstrating that the traditional metrics tend to favor oversmoothing over texture reconstruction.

Figures 6 and 7 display a summary of our denoising improvements.

Model Architecture Comparison					
Method	RAW PSNR	RAW SSIM	RGB PSNR	RGB SSIM	GMACs/MP
Input	36.616	0.80509	27.581	0.51557	-
Plain U-Net	45.485	0.97711	34.519	0.93481	2.2840
Wavelet	45.576	0.97892	34.591	0.93157	1.9800
Feature-Align	45.520	0.97780	34.565	0.93491	2.6688
FA w/o connections	45.395	0.97769	34.487	0.92534	2.5688

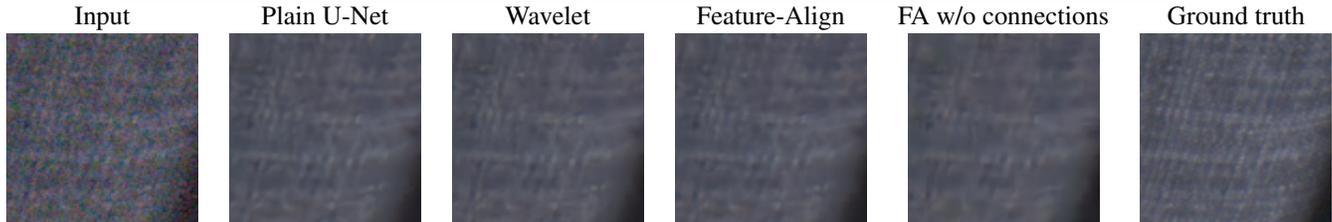


Table 1. Model with a Feature-Align layer shows a sharper image compared to a baseline wavelet transform model.

Effect of Noise Subrange Model Array				
Method	RAW PSNR	RAW SSIM	RGB PSNR	RGB SSIM
Input	36.616	0.80509	27.581	0.51557
No NSMA	45.520	0.97780	34.565	0.93491
NSMA (n=4)	45.668	0.97936	34.869	0.94014

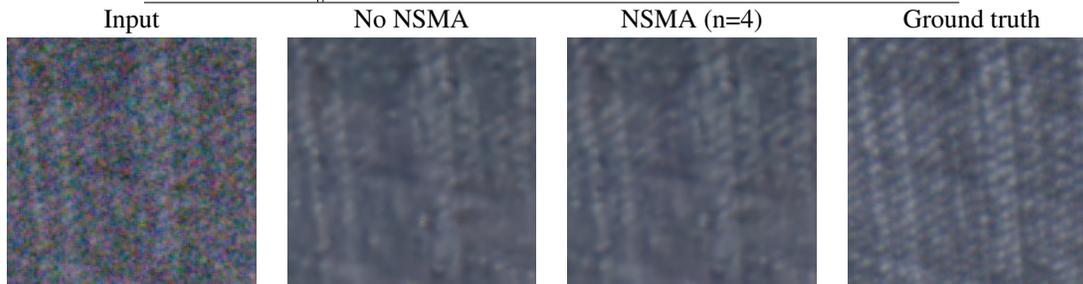


Table 2. NSMA substantially improves the image quality of our efficient model.

Effect of Loss Function					
Method	RAW PSNR	RAW SSIM	RGB PSNR	RGB SSIM	RGB LPIPS
Input	36.616	0.80509	27.581	0.51557	0.53680
Charbonnier	45.668	0.97936	34.869	0.94014	0.28882
RGB Perceptual	45.519	0.97685	34.665	0.92168	0.32448
Knowledge Dist.	45.813	0.98217	35.499	0.94079	0.27009
Feature Matching	45.483	0.97828	34.607	0.93685	0.24603

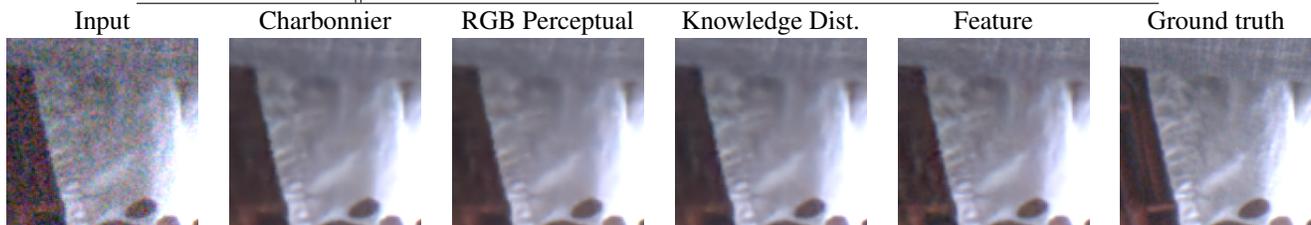


Table 3. Charbonnier Loss is outperformed by both Simple Knowledge Distillation and Feature Matching Loss. Simple Knowledge Distillation has the best performance in traditional metrics, while Feature Matching Loss has the best performance in the perceptual LPIPS metric. The model output with Feature Matching Loss appears to have the most detail.

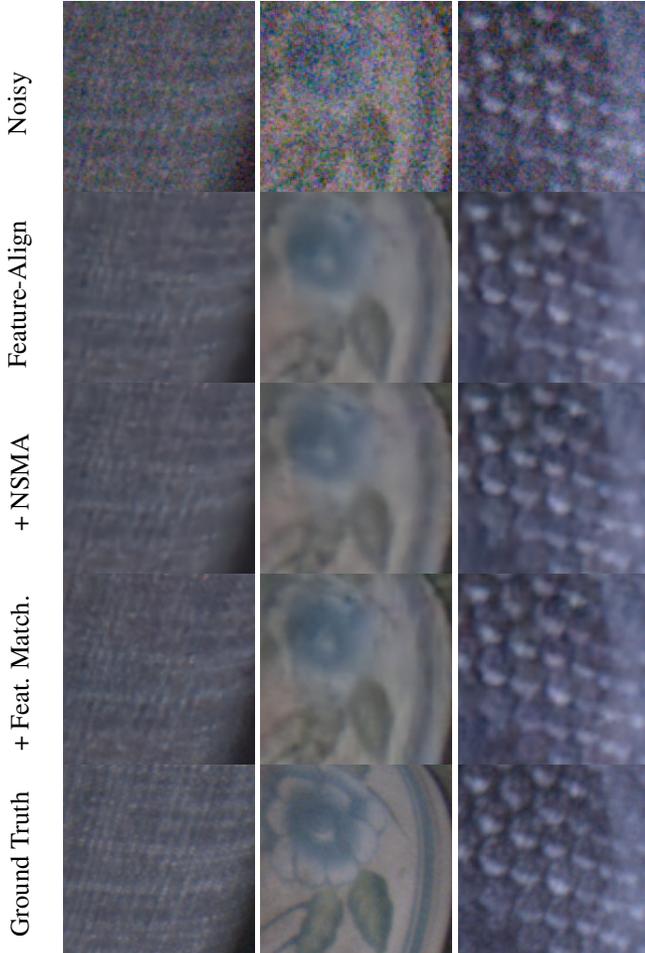


Figure 7. Our improvements in order: the effect of using the NSMA ($n=4$) on our Feature-Align method, followed by the effect of replacing Charbonnier Loss with Feature Matching Loss in the NSMA of Feature-Align models. Images are brightened for clarity.

4.2. Comparison to State-of-the-Art

For a standardized comparison of our method against existing state-of-the-art methods, we apply our method to the Darmstadt Noise Dataset public benchmark [30]. In training the noise parameters are tweaked to fit the DND test dataset and the white balance gains used to create the unprocessed MIRFLICKR dataset are tweaked to match [5] to be representative of DND. Table 4 shows our method compared to the top 3 submissions on the benchmark. Our method achieves a RAW PSNR of 48.2756, a comparable metric to that of CycleISP [44] 49.1251 RAW PSNR, while using 263.32 times fewer MACs per megapixel. Our network uses 89,696 bytes of learned parameters, and the NSMA of 4 of our networks takes 4 times 89,696 bytes, which is 358,784 bytes. This is 22.7% the size of learned parameters used by CycleISP, 1,578,624 bytes.

Darmstadt Noise Evaluation					
Method	RAW PSNR	RAW SSIM	RGB PSNR	RGB SSIM	GMACs/MP
Noisy	-	-	29.836	0.7018	-
CycleISP	49.1251	0.983	40.4987	0.9655	702.73
UPI Raw	48.8905	0.9824	40.1728	0.9623	74.305
UPI RGB	48.8824	0.9821	40.3545	0.9641	74.233
Ours	48.2756	0.9808	39.5061	0.9572	2.6688

Table 4. Our method achieves comparable results to the computationally expensive state-of-the-art methods.

Reduced Baselines vs. Our Method						
Method	RAW PSNR	RAW SSIM	RGB PSNR	RGB SSIM	GMACs/MP	Params
CycleISP-Lite	44.992	0.97868	34.728	0.92853	2.8568	11,063
UPI-Raw-Lite	44.222	0.94095	29.954	0.82964	3.1894	347,733
Ours	45.813	0.98217	35.499	0.94079	2.6688	239,304

Table 5. Scaled baselines versus the simple knowledge distillation variant of our method on our test set

To compare against UPI and CycleISP when the computational cost of methods are similar, in Table 5 we report the performance of UPI and CycleISP when these methods are scaled down by trivial means. UPI-Lite’s convolution blocks contain 1 convolution rather than 3, and the number of hidden channels is multiplied by a factor of 0.3. CycleISP-Lite contains 1 RRG rather than 8, 1 DAB rather than 4, and 16 hidden channels rather than 64. When all methods are normalized to the same compute, our method demonstrates superior performance in the conventional metrics. While CycleISP-Lite has much fewer parameters, this is not practically relevant when running the network on a typical CNN accelerator. In int8 format, our model consumes 239KB which is within the SRAM size of a typical CNN accelerator (usually in megabytes). A network’s latency and power-consumption arise from the amount of compute (MACs) and the activation memory footprint rather than number of parameters. Our method reduces MACs the most while retaining the highest image quality.

5. Conclusion

In this work we propose three innovations that enable high quality image denoising in an efficient model architecture: a Feature-Align layer, the use of an array of models tuned to different regions of the problem space, and a loss function for knowledge distillation that maximizes texture recovery. We combine these innovations along with existing state-of-the-art techniques and prove their efficacy by evaluating the method on a public benchmark. In addition, we propose a new public dataset of carefully constructed pairs of noisy and ground truth images with noise level annotations. The low computational cost of our method in comparison to existing state-of-the-art methods enables a variety of new applications for learning-based denoising in edge devices.

References

- [1] Abdelrahman Abdelhamed, Stephen Lin, and Michael S. Brown. A high-quality denoising dataset for smartphone cameras. In *CVPR*, June 2018.
- [2] Josue Anaya and Adrian Barbu. Renoir – a dataset for real low-light image noise reduction. *Journal of Visual Communication and Image Representation*, 51:144–154, Feb 2018.
- [3] F. J. Anscombe. The transformation of poisson, binomial and negative-binomial data. *Biometrika*, 35(3/4):246–254, 1948.
- [4] L. Azzari and A. Foi. Gaussian-cauchy mixture modeling for robust signal-dependent noise estimation. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5357–5361, 2014.
- [5] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T Barron. Unprocessing images for learned raw denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11036–11045, 2019.
- [6] H. C. Burger, C. J. Schuler, and S. Harmeling. Image denoising: Can plain neural networks compete with bm3d? In *CVPR*, pages 2392–2399, 2012.
- [7] P. Charbonnier, L. Blanc-Feraud, G. Aubert, and M. Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proceedings of 1st International Conference on Image Processing*, volume 2, pages 168–172 vol.2, 1994.
- [8] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *CVPR*, 2018.
- [9] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8):2080–2095, 2007.
- [10] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image restoration by sparse 3D transform-domain collaborative filtering. In Jaakko T. Astola, Karen O. Egiazarian, and Edward R. Dougherty, editors, *Image Processing: Algorithms and Systems VI*, volume 6812, pages 62 – 73. International Society for Optics and Photonics, SPIE, 2008.
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [12] Alessandro Foi. Clipped noisy images: Heteroskedastic modeling and practical denoising. *Signal Processing*, 89:2609–2629, 12 2009.
- [13] S. Gu, L. Zhang, W. Zuo, and X. Feng. Weighted nuclear norm minimization with application to image denoising. In *CVPR*, pages 2862–2869, 2014.
- [14] Samuel W. Hasinoff. *Photon, Poisson Noise*, pages 608–610. Springer US, Boston, MA, 2014.
- [15] S. W. Hasinoff, F. Durand, and W. Freeman. Noise-optimal capture for high dynamic range photography. *CVPR*, pages 553–560, 2010.
- [16] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015.
- [17] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1314–1324, 2019.
- [18] 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.
- [19] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.
- [20] Mark J. Huiskes and Michael S. Lew. The mir flickr retrieval evaluation. In *MIR '08: Proceedings of the 2008 ACM International Conference on Multimedia Information Retrieval*, New York, NY, USA, 2008. ACM.
- [21] Y. I. Jang, Y. Kim, and N. I. Cho. Dual path denoising network for real photographic noise. *IEEE Signal Processing Letters*, 27:860–864, 2020.
- [22] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [23] Jiaming Liu, Chi-Hao Wu, Yuzhi Wang, Qin Xu, Yuqian Zhou, Haibin Huang, Chuan Wang, Shaofan Cai, Yifan Ding, Haoqiang Fan, and Jue Wang. Learning raw image denoising with bayer pattern unification and bayer preserving augmentation, 2019.
- [24] Wei Liu, Qiong Yan, and Yuzhi Zhao. Densely self-guided wavelet network for image denoising. In *CVPRW*, June 2020.
- [25] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Transactions on Graphics (TOG)*, 39(4):71–1, 2020.
- [26] M. Makitalo and A. Foi. Optimal inversion of the anscombe transformation in low-count poisson image denoising. *IEEE Transactions on Image Processing*, 20(1):99–109, 2011.
- [27] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, pages 2337–2346, 2019.
- [28] Yali Peng, Yue Cao, Shigang Liu, Jian Yang, and Wangmeng Zuo. Progressive training of multi-level wavelet residual networks for image denoising. *arXiv preprint*, 2020.
- [29] Tobias Plötz and Stefan Roth. Neural nearest neighbors networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [30] Tobias Plötz and Stefan Roth. Benchmarking denoising algorithms with real photographs, 2017.
- [31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *CVPR*, 2015.
- [32] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, pages 4510–4520, 2018.
- [33] S. Sardy, P. Tseng, and A. Bruce. Robust wavelet denoising. *IEEE Transactions on Signal Processing*, 49(6):1146–1152, 2001.

- [34] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [35] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015.
- [36] Brijraj Singh, Durga Toshniwal, and Sharan Kumar Alur. Shunt connection: An intelligent skipping of contiguous blocks for optimizing mobilenet-v2. *Neural Networks*, 118:192–203, 2019.
- [37] Y. Song, Y. Zhu, and X. Du. Grouped multi-scale network for real-world image denoising. *IEEE Signal Processing Letters*, pages 1–1, 2020.
- [38] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *TPAMI*, 2019.
- [39] Lin Wang and Kuk-Jin Yoon. Knowledge distillation and student-teacher learning for visual intelligence: A review and new outlooks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [40] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.
- [41] Yuzhi Wang, Haibin Huang, Qin Xu, Jiaming Liu, Yiqun Liu, and Jue Wang. Practical deep raw image denoising on mobile devices, 2020.
- [42] Jun Xu, Lei Zhang, David Zhang, and Xiangchu Feng. Multi-channel weighted nuclear norm minimization for real color image denoising, 2018.
- [43] N. Yair and T. Michaeli. Multi-scale weighted nuclear norm image restoration. In *CVPR*, pages 3165–3174, 2018.
- [44] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Cycleisp: Real image restoration via improved data synthesis, 2020.
- [45] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for real image restoration and enhancement, 2020.
- [46] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [47] 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.
- [48] Yiyun Zhao, Zhuqing Jiang, Aidong Men, and Guodong Ju. Pyramid real image denoising network, 2019.