

Practical Deep Raw Image Denoising on Mobile Devices

Yuzhi Wang^{1,2}, Haibin Huang², Qin Xu², Jiaming Liu², Yiqun Liu¹, and Jue Wang²

¹ Tsinghua University

² Megvii Technology

Abstract. Deep learning-based image denoising approaches have been extensively studied in recent years, prevailing in many public benchmark datasets. However, the state-of-the-art networks are computationally too expensive to be directly applied on mobile devices. In this work, we propose a light-weight, efficient neural network-based raw image denoiser that runs smoothly on mainstream mobile devices, and produces high quality denoising results. Our key insights are twofold: (1) by measuring and estimating sensor noise level, a smaller network trained on synthetic sensor-specific data can out-perform larger ones trained on general data; (2) the large noise level variation under different ISO settings can be removed by a novel *k-Sigma Transform*, allowing a small network to efficiently handle a wide range of noise levels. We conduct extensive experiments to demonstrate the efficiency and accuracy of our approach. Our proposed mobile-friendly denoising model runs at ~ 70 milliseconds per megapixel on Qualcomm Snapdragon 855 chipset, and it is the basis of the night shot feature of several flagship smartphones released in 2019.

1 Introduction

Smartphones have become the go-to devices for consumer photography in recent years. Compared with DSLR cameras, images captured with mobile devices are more easily contaminated with higher level of noise due to the use of relatively low-cost sensors and lenses, especially in low-light scenarios.

Despite decades of development in image denoising technologies, it remains challenging to restore high quality images from extremely noisy ones on mobile devices. Recently, deep neural network (DNN) based denoising methods [36, 10, 42, 21, 38, 30, 37] have achieved tremendous success and outperformed most traditional methods [43, 12, 18, 39, 40]. It is however not practical to directly deploy these heavy-weight DNNs on mobile devices due to the limited computational resources available on them.

In this work, we propose a simple yet efficient approach for deep raw image denoising. It can run efficiently on off-the-shelf smartphones with high quality

This work is supported by The National Key Research and Development Program of China under Grant 2018YFC0831700.



Fig. 1: Our proposed denoising method can run smoothly on smartphones with high quality noise reduction even in low-light conditions (see (b),(d)). In contrast, the default ISP image denoiser produces images with over-smoothed high texture regions (e.g. ground in (a)) and noisy smooth regions (e.g. sky in (c)).

noise reduction. Our key observation is that the noise characteristics for a specific sensor model are consistent and can be measured with sufficient accuracy. By capturing and modeling sensor noise, we can generate synthetic datasets with clean and noisy image pairs, and train a light-weight neural network on them. The trained model remains highly effective on real images captured by the same sensor (i.e. the same smartphone model). Furthermore, based on the parametric sensor noise model, we derive a unique linear transform in luminance space, dubbed as k-Sigma Transform, that maps noisy images captured under different ISO settings into an ISO-invariant signal-noise space, allowing a single network to handle different noise levels in different scenes. We show that this approach is not only theoretically elegant, but in practice is more efficient than training a separate model for each ISO setting, or having one large model trained on images with different noise levels.

To summarize, the main contributions of this work are as follows:

- A systematic approach to estimate sensor noise and train a sensor-specific denoising neural network using properly constructed synthetic data.
- A novel k-Sigma Transform to map noisy images under different ISO settings into an ISO-invariant signal-noise space. Instead of training separate models per ISO or a larger model to cover the variations, the proposed transform allows a single small network trained in this space to handle images with different noise levels.
- A mobile-friendly network architecture for efficient image denoising. We provide in-depth analysis and comparison with different network architectures and denoising methods, demonstrating that our method has compatible performance with state-of-the-art approaches with significantly less computational resources.

To the best of our knowledge, our solution is the first practical deep-learning-based image denoising approach that has satisfactory efficiency and accuracy on mobile devices. In Fig. 1 we show exemplar images captured by an off-the-shelf smartphone that use our approach in low-light photography. Compared with the default ISP image denoising technique, our results contain much more fine details of the scene.

2 Related work

Image denoising is a fundamental task in image processing and computer vision. Classical methods often rely on using sparse image priors, such as non-local means (NLM) [6], sparse coding [13, 28, 2], 3D transform-domain filtering (BM3D) [12], and others [18, 32]. Among them BM3D is usually deemed as the leading method considering its accuracy and robustness. Most of these methods are designed for general noise and do not take advantage of known sensor noise characteristics. Their algorithmic complexity is usually high, making full-fledged implementation difficult on smartphones.

With the recent development of convolutional neural networks (CNNs), training end-to-end denoising CNNs has gained considerable attention. Earlier work that uses multi-layer perceptron (MLP) [7] has achieved comparable results with BM3D. Further improvements have been achieved with the introduction of more advanced network architectures, resulting in a large number of CNN-based denoising methods [36, 10, 21, 38, 30, 37, 41, 23]. These works are primarily focused on novel network structures for improving the accuracy, without paying much attention to their adaptability to mobile devices.

Our work focuses on denoising raw image, i.e., images read out from the sensor in the raw Bayer format before demosaicing and other ISP-processing. On the recently proposed public raw image denoising benchmark datasets [1, 8, 3], CNN-based methods [8, 20, 17] have achieved the best results. It is however a very tedious work to construct such high quality real datasets with clean and noisy image pairs. Thus, the problem of synthesizing realistic image noise for training has also been extensively studied, including Gaussian-Poisson noise [16, 27], Gaussian Mixture Model (GMM) [44], in-camera process simulation [25, 34], GAN-generated noises [9] and so on. It has been shown that networks properly trained from the synthetic data can generalize well to real data [42, 5].

The existing best practice for raw image denoising on mobile devices is to capture and merge multiple frames [19, 24, 31]. These methods generally require accurate and fast image alignment, which is hard to achieve when moving objects present in the scene. Furthermore, when noise level is high, averaging multiple frames can reduce, but not completely remove the noise, leading to unsatisfactory results. To the best of our knowledge, our proposed method is the first single frame, deep learning-based raw image denoiser specifically designed for mobile devices.

3 Method

In this section, we first revisit the general ISO-dependent noise model of camera sensor, and then describe how to estimate noise parameters given a new sensor. We further show how to synthesize ISO-independent training data using the proposed k-Sigma Transform, and use it to train a small neural network that can handle a wide range of noise levels.

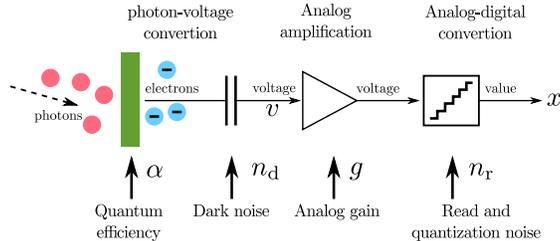


Fig. 2: Photon transfer pipeline: multiple noise sources like shot noise, read-out noise and thermal noise are involved along the camera pipeline. Check [14] for more details.

3.1 The Noise Model

A camera sensor converts the photons hitting the pixel area during the exposure time into a digitized luminance map. As shown in the photon transfer pipeline shown in Fig. 2, this process contains multiple stages, where each stage introduces specific noise. Let us first consider an ideal system with no noise. Under the linear camera model, at each pixel, the sensor conversion is a linear amplification as:

$$x^* = g\alpha u^*, \quad (1)$$

where u^* is the expected number of photons hitting the pixel area, α is the quantum efficiency factor and g is the analog gain. Now considering the system noise in each step of the pipeline in Fig. 2, we have:

$$x = g(\alpha u + n_d) + n_r, \quad (2)$$

where u denotes the actual collected amount of photons, and $n_d \sim \mathcal{N}(0, \sigma_d^2)$ and $n_r \sim \mathcal{N}(0, \sigma_r^2)$ are Gaussian noise before and after applying the analog gain. Furthermore, it is demonstrated in [14] that u obeys a Poisson distribution of u^* , given by

$$u \sim \mathcal{P}(u^*). \quad (3)$$

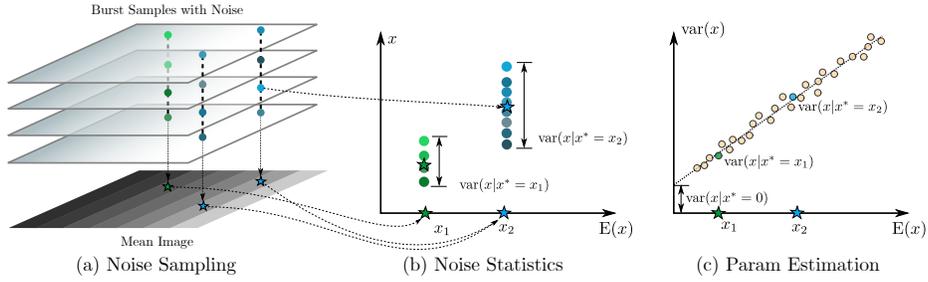


Fig. 3: Noise parameter estimation with a burst series of raw images of a static grayscale chart.

Combining Eqn. (1) to Eqn. (3), we have:

$$x \sim (g\alpha)\mathcal{P}\left(\frac{x^*}{g\alpha}\right) + \mathcal{N}(0, g^2\sigma_d^2 + \sigma_r^2). \quad (4)$$

This is consistent with the Poisson-Gaussian noise model that has been extensively studied in previous work [16, 27]. This formulation can be further simplified by replacing $k = g\alpha$ and $\sigma^2 = g^2\sigma_d^2 + \sigma_r^2$:

$$x \sim k\mathcal{P}\left(\frac{x^*}{k}\right) + \mathcal{N}(0, \sigma^2). \quad (5)$$

Note that both k and σ^2 are related to g , which is determined by the ISO setting of the camera.

3.2 Parameter Estimation

To sample the distribution described in Eqn. (5), we need an accurate estimation of k and σ under a specified ISO setting of a specific sensor. Luckily, as we check the mean and variance over x , shown in Eqn. (6), we can turn it into the following linear regression problem:

$$\begin{cases} E(x) &= x^*, \\ \text{Var}(x) &= kx^* + \sigma^2. \end{cases} \quad (6)$$

Similar to [15], we capture a series of raw images of a static grayscale chart in burst mode, depicted in Fig. 3a, and compute $E(x)$ from the series of luminance values at the same pixel location. Next, as shown in Fig. 3b, we bracket all pixels that have the same estimated luminance, and compute $\text{Var}(x)$ from them. A linear regression is then applied to find the optimal estimation of k and σ^2 , illustrated in Fig. 3c.

3.3 The k-Sigma Transform

In real applications the camera will automatically adjust the ISO settings according to the scene illumination, thus one has to consider different noise levels when training the denoising neural network. A straightforward solution is to train a single network to cover a wide range of ISO settings, but it puts extra burden on the network itself as the noise variation in the training data becomes quite large. Inspired by variance stabilizing transformations [4, 29], here we propose a k-Sigma Transform to avoid this problem.

Specifically, we define a linear transform

$$f(x) = \frac{x}{k} + \frac{\sigma^2}{k^2}. \quad (7)$$

According to our noise model of Eqn. (5),

$$f(x) \sim \mathcal{P}\left(\frac{x^*}{k}\right) + \mathcal{N}\left(\frac{\sigma^2}{k^2}, \frac{\sigma^2}{k^2}\right). \quad (8)$$

To analyze this distribution, a usual simplification is to treat the Poisson distribution $\mathcal{P}(\lambda)$ as a Gaussian distribution of $\mathcal{N}(\lambda, \lambda)$ [16]. Therefore:

$$\begin{aligned} & \mathcal{P}\left(\frac{x^*}{k}\right) + \mathcal{N}\left(\frac{\sigma^2}{k^2}, \frac{\sigma^2}{k^2}\right) \\ & \approx \mathcal{N}\left(\frac{x^*}{k}, \frac{x^*}{k}\right) + \mathcal{N}\left(\frac{\sigma^2}{k^2}, \frac{\sigma^2}{k^2}\right) \\ & = \mathcal{N}\left(\frac{x^*}{k} + \frac{\sigma^2}{k^2}, \frac{x^*}{k} + \frac{\sigma^2}{k^2}\right) \\ & = \mathcal{N}[f(x^*), f(x^*)]. \end{aligned} \quad (9)$$

Combining Eqn. (8) and Eqn. (9), the approximate distribution of $f(x)$ is:

$$f(x) \sim \mathcal{N}[f(x^*), f(x^*)]. \quad (10)$$

Eqn. (10) indicates that the distribution of $f(x)$ only depends on $f(x^*)$. As shown in Fig. 4, we can train a single network that takes $f(x)$ as input and outputs $f(\hat{x}^*)$ as an estimation of $f(x^*)$. The estimated true image value x^* can then be computed by applying the inverted k-Sigma Transform $f^{-1}(\cdot)$ to $f(\hat{x}^*)$. In other words, we apply ISO-dependent transforms to the input and output of the neural network, so that the network can be trained using normalized data without considering the ISO setting.

4 Learning to Denoise

4.1 Mobile-friendly Network Architecture

We further introduce a mobile-friendly convolutional neural network for image denoising, as shown in Fig. 5. We use a U-Net-like [33] architecture with 4 encoder and 4 decoder stages with skip connections, illustrated in Fig. 5a.

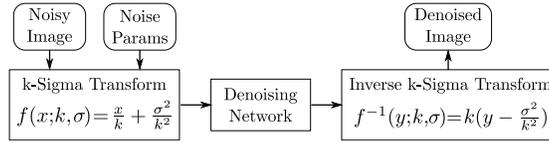
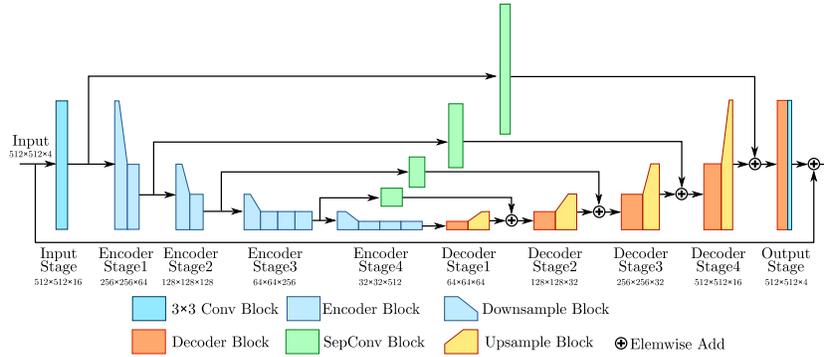
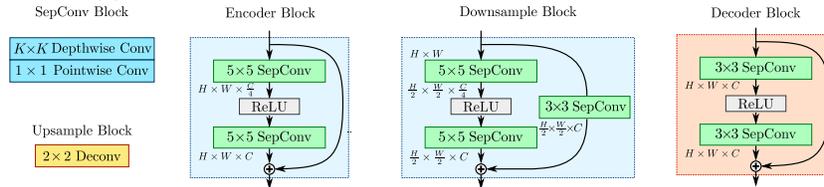


Fig. 4: The pipeline of running ISO-independent denoising network with k-Sigma Transform.



(a) The U-Net-like overall structure of the denoising network.



(b) Detailed structure of network blocks.

Fig. 5: The architecture of the proposed denoising network.

Fig. 5b depicts the detailed structures of the network blocks. Specifically, in order to run on mobile devices, we use separable-conv [11] in all encoder and decoder stages to reduce the computation cost, and normal dense convolution layers are only used in the input and output stage. In encoders, we use 5×5 kernel size to increase receptive field and decrease network depth, and downsample feature maps with stride-2 convolutions. In decoder, we only use 3×3 -speconv and upsample feature maps with 2×2 deconvolutions. The inputs of each encoder stages are combined into its corresponding decoder stage by element-wise adding, a 3×3 -speconv is adopted in the skip connect to match the channel shape. Finally, the last convolution layer outputs a residual added to the input image as the denoising result.

4.2 Training Dataset

To train our denoising network, we need pairs of noisy and clean RAW images. In this paper, we use a subset of See-in-the-Dark (SID) dataset proposed in [8] as the ground truth clean images. The SID dataset contains RAW images captured from a Sony $\alpha 7s$ II and a Fujifilm X-T2 camera, we choose the 10s and 30s long-exposure subset captured by the Sony $\alpha 7s$ II camera, and manually take out those with visible noise, leaving 214 high quality RAW images.

According to our noise model described in Section 3, if clean RAW images were available, we can synthesize noisy images by sampling from a Poisson-Gaussian distribution with estimated noise parameters measured from the target sensor.

4.3 Settings and Details

To generate training samples, we randomly crop 1024×1024 -sized bayer patches from the original dataset. We adopt the bayer-aug method described in [26] with random horizontal and vertical flipping and ensure the input bayer pattern is in R-G-G-B order. We then pack the bayer image to $512 \times 512 \times 4$ -shaped RGGGB tensor. We also randomly adjust the brightness and contrast of the cropped images for data augmentation. Noisy images are then synthesized according to the noise model, with noise parameters of randomly selected ISO value. Finally, we apply k-Sigma Transform to both the noisy and clean images so that the denoising network is trained in the ISO-independent space.

We use ℓ_1 distance between the noisy and clean images as the loss function, and train the network using Adam [22] optimizer. We adopt the triangular cyclical learning rate scheduling [35] with the maximum learning rate of 1e-3, cycle step of 50×214 iterations, and the base learning rate linearly decays to 1e-5 after 4000×214 iterations. The batch size is set to 1, and the training converges at 8000×214 iterations.

5 Experiments

In this section, we evaluate our denoising method with a real world dataset collected with an OPPO Reno-10x smartphone. This smartphone has three rear cameras, and we use the most commonly used main camera for our test. The sensor of this camera is Sony IMX586 sized at $1/2''$ with 48 megapixels, and the pixel size is $0.8 \mu\text{m}$. This sensor is widely used in many smartphones of 2019, including OPPO Reno series, Xiaomi 9, etc.

5.1 Noise Parameters Estimation

We first measure and estimate the noise parameters with the method described in Section 3.2. We write a simple app to collect RAW images with Android Camera-2 API, which allows us manually control the camera’s ISO and exposure time.

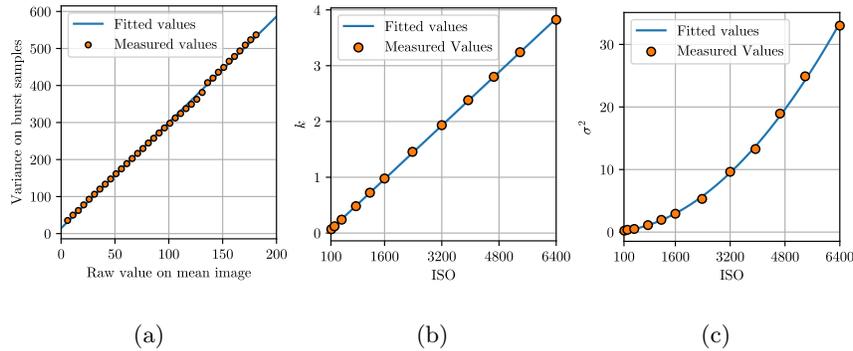


Fig. 6: Noise param estimation of Reno-10x smartphone: (a) parameter estimation at ISO-4800 (b) k values at different ISOs (c) σ^2 at different ISOs.

To keep a stable light condition, we use an Xrite SpectraLight QC light booth in a dark room. At each ISO and exposure time setting, we adjust the luminance of the light source to avoid over or under exposure, and the final values of the captured image are kept in an appropriate range.

At each ISO setting, 64 RAW images are captured in burst mode, and the mean image is considered as the clean image. With the method described in Section 3.2, we can estimate the noise params k and σ^2 at each specified ISO setting. Fig. 6a plots the value-variance curve of the test phone under ISO 4800, where the scattered dots represent the measured variances corresponding to each raw value on the mean image, and the blue line plots the linear regression result of Eqn. (6). From the figure we can see that our theoretical noise model can well fit the measurement results. The slope of the fitted line is the estimated noise parameter \hat{k} and the y-intercept value is the estimated noise parameter $\hat{\sigma}^2$.

The adjustable range of IMX586 sensor analog gain is [1.0, 64.0], corresponding to the ISO value of OPPO Reno-10x camera as [100, 6400]. According to our noise model Eqn. (4), the params k and σ^2 are linearly, and quadratically correlated to the ISO value, respectively. We measure and estimate the noise params at each ISO setting, and plot the ISO- k and ISO- σ^2 curve in Fig. 6. The scattered dots represent the estimated noise params under each ISO setting, and the blue lines in Fig. 6b and Fig. 6c respectively represent the linearly and quadratically fitted curves, which demonstrate that our theoretical model matches the measurements well.

With the ISO- k and ISO- σ^2 curves well fitted, the noise params under any ISO setting can be easily calculated, and thus satisfying the requirements of both synthesizing training data and applying the k-Sigma transform.

5.2 Test Dataset and Metrics

Since our proposed denoising network needs to be trained for specific sensors, we cannot directly use public benchmarks such as SIDD [1] due to the mismatch-

ing of sensors. Therefore, we build a testing dataset to evaluate our denoising method.

We use an Xrite SpectraLight QC light booth in a dark room to build a stable light condition. For each capturing process, 64 RAW images are captured and the mean image can be used as the estimated ground truth. As shown in Fig. 7, we capture 4 static scenes as the content of testing images, and 2 luminance conditions are set for each scene. When capturing one scene, the camera position and scene contents are kept fixed. We set 5 exposure combinations at each scene and luminance, which are ISO-800@160ms, ISO-1600@80ms, ISO-3200@40ms, ISO-4800@30ms and ISO-6400@20ms, respectively. These settings share an identical value of the product of ISO and exposure time, so that the captured images have similar brightness but different noise levels.

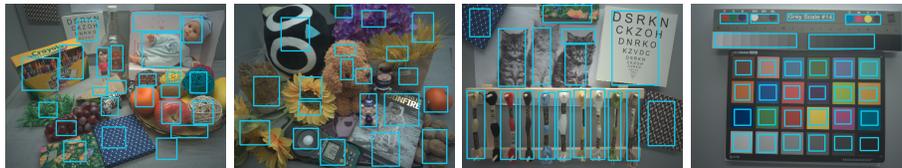


Fig. 7: The 4 scenes from the test dataset: the blue boxes represent the regions of interest where image quality metrics are calculated.

We use peak signal-noise-ratio (PSNR) and structural similarity (SSIM) to measure the performance of denoising methods. The PSNR and SSIM between the denoising results and the clean images are measured in sRGB domain with a simple post processing pipeline, including (1) white balance correction, (2) demosaicking, (3) color correction, and (4) gamma correction. The parameters for white balance correction and color correction are obtained from the metadata of the RAW image. The demosaicking algorithm is pixel-grouping (PPG) and the gamma value is 2.2.

5.3 Results

We first show the comparison between our method and the previous state-of-the-art raw image denoising method proposed in [26], which is trained and tested using the SIDD dataset. An extra-large UNet-like network which costs 1T multiply-and-cumulate operations (MACs) per megapixel is proposed in [26] and achieved the state-of-the-art performance in NTIRE 2019 denoising challenge. In addition to the UNet-1T architecture, we modify the network by reducing the channel width and layer depth to fit various computation complexities. We train these models with two different data sources: SIDD dataset and the training data generated by our method.

As shown in Fig. 8, because of our accurate noise modeling, the models trained with our synthetic data outperform those trained with the SIDD dataset

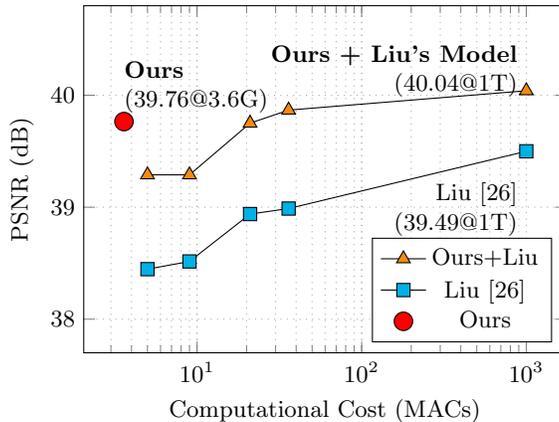


Fig. 8: PSNRs at different computation costs of our method, our method using Liu’s model [26], and Liu’s method. Note that the MACs are in the log space.

Table 1: Running time of denoising 1 MP on Qualcomm Snapdragon 855 GPU

Method	Ours	UNet-5G	UNet-21G	UNet-36G
ms/MP	70.70	79.19	292.0	383.9

by a large margin. Moreover, our mobile-friendly network trained on our synthetic data achieves comparable performance with the UNet-36G with only 10% of its computational complexity (3.6G vs 36.3G). More visual comparisons are provided in Fig. 9.

We further test the actual running time of different models on mobile devices, listed in Table 1. Our mobile-friendly model can process a 1024×1024 Bayer input with 70.70 ms (aka ~ 850 ms for a 12MP full-sized image) on a Qualcomm Snapdragon 855 GPU, while other models with comparable performance require significantly longer time, 292 ms for 21G network and 383 ms for 36G network (aka 4.45 s for 12MP), making them impractical to be deployed on mobile devices.

5.4 Ablation Studies

Data synthesis method To verify the effectiveness of our data synthesis method, we train our denoising network with four different training datasets, including

- testset overfitting: directly use the inputs and ground truth from the testset as the training data;
- testset with synthetic: use the ground truth of the testset and add synthetic noise as the input;
- testset with noise scaling: use our noise synthesis method but scale the noise parameters to purposely mismatch with the target sensor;

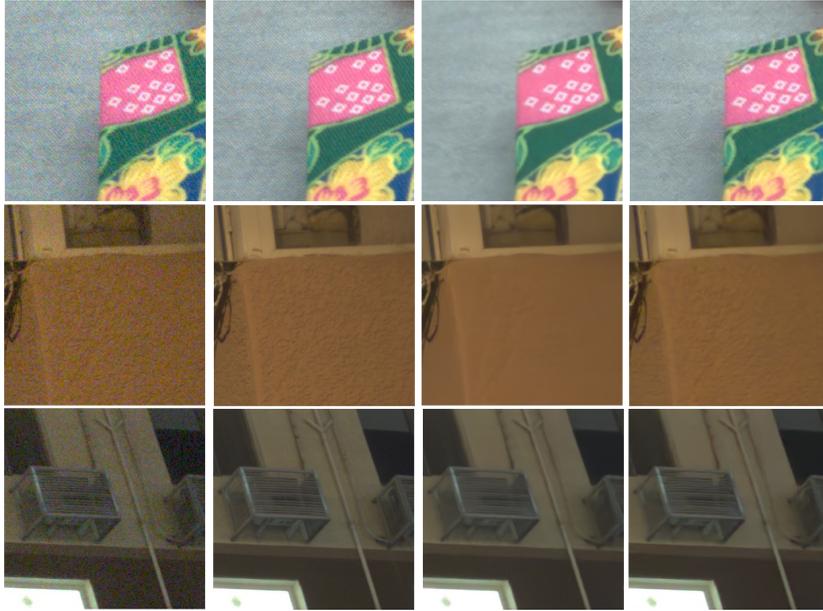


Fig. 9: More visual results on our real dataset. From left to right: input image; ground truth; result based on method in [26]; our result. Compared with [26] which generates blurred areas, our method can efficiently reduce the noise as well maintain underlying details. Moreover, our method utilizes a significantly smaller network (3.6G vs 1T).

- SID: synthesize the training set with SID dataset described in Section 4.2.

The test results are listed in Table 2. Not surprisingly, the overfitting experiment obtains the highest PSNR of 40.06 dB and SSIM of 0.9335, which sets an upper-bound for this comparison. The testset with our synthetic noise achieves the second best result of 39.97 dB PSNR. When using inaccurate noise parameters by scaling k and σ^2 by 0.5 or 2.0, the network results in the lowest PSNRs in this experiment. Our method of using the SID dataset achieves 39.76 dB PSNR.

This experiment shows that our noise model and synthetic noise can well match the characteristics of the real input noise, and the testing results can be close to the upper bound performance. Inaccurate noise parameters, even in the overfitting experiment, lead to noticeable performance degradation.

Robustness to ISO We compare several strategies of denoising images of different ISO settings to verify the effectiveness of the k-Sigma Transform. We compare two ways of handling multiple ISO settings: (1) iso-augmentation: randomly choose ISO settings when synthesizing training samples and directly feed them into the denoising network; and (2) concat-variance: the method proposed in [5], where the estimated noise variance are concatenated to the input as 4

Table 2: Comparison of training datasets and noise parameters

Dataset	PSNR (dB)	SSIM
testset overfitting	40.06	0.9335
testset+synthetic	39.97	0.9337
testset+synthetic scale 0.5	36.87	0.8850
testset+synthetic scale 2.0	39.68	0.9301
SID	39.76	0.9310

Table 3: Strategies of denoising for multiple ISOs

Method	PSNR (dB)	SSIM
k-Sigma Transform	39.76	0.9310
concat-variance [5]	39.65	0.9307
iso-augmentation	39.57	0.9299
single-iso-1600	35.74	0.8072
single-iso-3200	38.56	0.9089
single-iso-6400	38.16	0.9167

additional channels. In addition, we also test the performance of the single-ISO method, where the training data is synthesized using noise parameters of a single ISO.

The results are listed in Table 3. The concat-variance strategy achieves the PSNR of 39.65 dB, which is about 0.09 dB higher than the iso-augmentation strategy. This means that explicit noise level information can help the model achieve better results than blind denoising. With the proposed k-Sigma Transform, our network achieves the highest PSNR in this experiment. In comparison, all single-ISO methods perform much worse than multi-ISO ones.

Table 4 gives a more detailed analysis of single-ISO methods and our approach, where the PSNRs measured in the testset are grouped into different ISO settings. From the table we can see that when the ISO setting of the testset matches with the single-ISO model, it can produce competitive denoising results. Our method base on the k-Sigma Transform performs consistently well under all ISO settings.

Table 4: PSNRs under different ISO settings

	ISO-800	ISO-1600	ISO-3200	ISO-4800	ISO-6400
k-Sigma Transform	43.21	41.48	39.49	38.17	36.94
single-iso-1600	42.96	41.48	35.01	31.33	28.86
single-iso-3200	41.79	40.87	39.51	36.97	34.08
single-iso-6400	39.59	38.59	38.11	37.80	36.91

6 Conclusion

We have presented a new raw image denoiser designed for mobile devices. By accurate sensor noise estimation, we can utilize a light-weight network trained on sensor-specific synthetic data that generalizes well to real noise. We also propose a k-Sigma Transform to process the input and output data, so that denoising can be learned in an ISO-independent space. This allows the network to handle a wide range of noise level without increasing network complexity. Our results show that the proposed method can achieve compatible performance with state-of-the-art methods, which typically employ much larger networks that cannot be directly applied for mobile applications.

In applications, our method can be integrated into existing camera pipeline and replace its denoising component. Since our method can produce high quality denoised raw images, it gives a strong base for the ISP to apply more aggressive post-processing. Our methods have been featured in the night shot mode of several flagship phones released in 2019, with stable and outstanding performance on mobile devices.

In the future, we would like to explore how to further reduce the computational complexity of the proposed method, so that we can apply it on video streams in real-time. Also, we believe it will be interesting and promising to explore deep learning based approaches for raw image processing that can improve or even replace camera’s ISP pipeline.

References

1. Abdelhamed, A., Lin, S., Brown, M.S.: A high-quality denoising dataset for smart-phone cameras. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
2. Aharon, M., Elad, M., Bruckstein, A., et al.: K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on signal processing* **54**(11), 4311 (2006)
3. Anaya, J., Barbu, A.: Renoir—a dataset for real low-light image noise reduction. *Journal of Visual Communication and Image Representation* **51**, 144–154 (2018)
4. Anscombe, F.J.: The transformation of poisson, binomial and negative-binomial data. *Biometrika* **35**(3/4), 246–254 (1948)
5. Brooks, T., Mildenhall, B., Xue, T., Chen, J., Sharlet, D., Barron, J.T.: Unprocessing images for learned raw denoising. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 11036–11045 (2019)
6. Buades, A., Coll, B., Morel, J.M.: A non-local algorithm for image denoising. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). vol. 2, pp. 60–65. IEEE (2005)
7. Burger, H.C., Schuler, C.J., Harmeling, S.: Image denoising: Can plain neural networks compete with bm3d? In: CVPR (2012)
8. Chen, C., Chen, Q., Xu, J., Koltun, V.: Learning to see in the dark. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3291–3300 (2018)
9. Chen, J., Chen, J., Chao, H., Yang, M.: Image blind denoising with generative adversarial network based noise modeling. In: CVPR (2018)
10. Chen, Y., Pock, T.: Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE transactions on pattern analysis and machine intelligence* **39**(6), 1256–1272 (2017)
11. Chollet, F.: Xception: Deep Learning with Depthwise Separable Convolutions (oct 2016), <http://arxiv.org/abs/1610.02357>
12. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image restoration by sparse 3d transform-domain collaborative filtering. In: Image Processing: Algorithms and Systems VI. vol. 6812, p. 681207. International Society for Optics and Photonics (2008)
13. Elad, M., Aharon, M.: Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing* **15**(12), 3736–3745 (2006)
14. European Machine Vision Association.: Standard for Characterization of Image Sensors and Cameras (2010). <https://doi.org/10.1063/1.1518010>
15. Foi, A., Alenius, S., Katkovnik, V., Egiazarian, K.: Noise measurement for raw-data of digital imaging sensors by automatic segmentation of nonuniform targets. *IEEE Sensors Journal* **7**(10), 1456–1461 (2007)
16. Foi, A., Trimeche, M., Katkovnik, V., Egiazarian, K.: Practical poissonian-gaussian noise modeling and fitting for single-image raw-data. *IEEE Transactions on Image Processing* **17**(10), 1737–1754 (2008)
17. Gharbi, M., Chaurasia, G., Paris, S., Durand, F.: Deep joint demosaicking and denoising. *ACM Transactions on Graphics (TOG)* **35**(6), 191 (2016)
18. Gu, S., Zhang, L., Zuo, W., Feng, X.: Weighted nuclear norm minimization with application to image denoising. In: CVPR (2014)

19. Hasinoff, S.W., Sharlet, D., Geiss, R., Adams, A., Barron, J.T., Kainz, F., Chen, J., Levoy, M.: Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics* **35**(6), 1–12 (2016). <https://doi.org/10.1145/2980179.2980254>, <http://dl.acm.org/citation.cfm?doid=2980179.2980254>
20. Hirakawa, K., Parks, T.W.: Joint demosaicing and denoising. *IEEE Transactions on Image Processing* **15**(8), 2146–2157 (2006)
21. Jain, V., Seung, S.: Natural image denoising with convolutional networks. In: *Advances in neural information processing systems*. pp. 769–776 (2009)
22. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
23. Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T.: Noise2noise: Learning image restoration without clean data. *arXiv preprint arXiv:1803.04189* (2018)
24. Liba, O., Murthy, K., Tsai, Y.T., Brooks, T., Xue, T., Karnad, N., He, Q., Barron, J.T., Sharlet, D., Geiss, R., Hasinoff, S.W., Pritch, Y., Levoy, M.: Handheld mobile photography in very low light. *ACM Transactions on Graphics* **38**(6) (2019). <https://doi.org/10.1145/3355089.3356508>
25. Liu, C., Szeliski, R., Kang, S.B., Zitnick, C.L., Freeman, W.T.: Automatic estimation and removal of noise from a single image. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(2), 299–314 (2008)
26. Liu, J., Wu, C.H., Wang, Y., Xu, Q., Zhou, Y., Huang, H., Wang, C., Cai, S., Ding, Y., Fan, H., Wang, J.: Learning Raw Image Denoising with Bayer Pattern Unification and Bayer Preserving Augmentation (apr 2019), <http://arxiv.org/abs/1904.12945>
27. Liu, X., Tanaka, M., Okutomi, M.: Practical signal-dependent noise parameter estimation from a single noisy image. *IEEE Transactions on Image Processing* **23**(10), 4361–4371 (2014)
28. Mairal, J., Bach, F.R., Ponce, J., Sapiro, G., Zisserman, A.: Non-local sparse models for image restoration. In: *ICCV*. vol. 29, pp. 54–62. Citeseer (2009)
29. Makitalo, M., Foi, A.: Optimal inversion of the anscombe transformation in low-count poisson image denoising. *IEEE transactions on Image Processing* **20**(1), 99–109 (2010)
30. Mao, X., Shen, C., Yang, Y.B.: Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In: *NeurIPS* (2016)
31. Mildenhall, B., Barron, J.T., Chen, J., Sharlet, D., Ng, R., Carroll, R.: Burst Denoising with Kernel Prediction Networks (dec 2017), <https://arxiv.org/abs/1712.02327>
32. Portilla, J., Strela, V., Wainwright, M.J., Simoncelli, E.P.: Image denoising using scale mixtures of gaussians in the wavelet domain. *IEEE Trans Image Processing* **12**(11) (2003)
33. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 234–241. Springer (2015)
34. Shi, G., Zifei, Y., Kai, Z., Wangmeng, Z., Lei, Z.: Toward convolutional blind denoising of real photographs. In: *arXiv preprint arXiv:1807.04686* (2018)
35. Smith, L.N.: Cyclical learning rates for training neural networks. In: *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*. pp. 464–472. IEEE (2017)

36. Tai, Y., Yang, J., Liu, X., Xu, C.: Memnet: A persistent memory network for image restoration. In: Proceedings of the IEEE international conference on computer vision. pp. 4539–4547 (2017)
37. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9446–9454 (2018)
38. Xie, J., Xu, L., Chen, E.: Image denoising and inpainting with deep neural networks. In: Advances in neural information processing systems. pp. 341–349 (2012)
39. Xu, J., Zhang, L., Zhang, D., Feng, X.: Multi-channel weighted nuclear norm minimization for real color image denoising. In: ICCV (2017)
40. Yair, N., Michaeli, T.: Multi-scale weighted nuclear norm image restoration. In: CVPR (2018)
41. Zhang, K., Zuo, W., Zhang, L.: Ffdnet: Toward a fast and flexible solution for cnn based image denoising. IEEE Transactions on Image Processing (2018)
42. Zhou, Y., Jiao, J., Huang, H., Wang, Y., Wang, J., Shi, H., Huang, T.: When awgn-based denoiser meets real noises. arXiv preprint arXiv:1904.03485 (2019)
43. Zhou, Y., Liu, D., Huang, T.: Survey of face detection on low-quality images. In: 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018). pp. 769–773. IEEE (2018)
44. Zhu, F., Chen, G., Heng, P.A.: From noise modeling to blind image denoising. In: CVPR (2016)