Scope
Volume 14 Number 02 June 2024

Photography Denoising Algorithm with the Assistance of AI

Dr. R. Naveenkumar

Associate Professor, Department of Computer Science and Engineering, Brainware University,Kolkata, West Bengal.

Sanel. S,

Research Scholar, Assistant Professor, Department of Computer Science KSMDB College, Kollam, Kerala

Abstract

Modern digital cameras are susceptible to producing noisy images in low light conditions due to sensor limitations. This noise manifests as randomly colored pixels that appear as grain, degrading overall image quality. Traditional denoising algorithms have limitations in differentiating actual image detail from noise. Recent advancements in artificial intelligence (AI) and deep learning offer potential new approaches for more intelligently identifying and reducing image noise while preserving real details and edges. In this paper, we propose a novel AI-assisted denoising algorithm that leverages a convolutional neural network (CNN) to differentiate noisy pixels from those containing actual image signal. The trained model classifies each pixel as either signal or noise. Pixels identified as noise are smoothed by averaging the color values of neighboring signal pixels, while pixels classified as true signal are left unaltered to maintain sharpness. We evaluate our algorithm on a dataset of noisy raw images from various camera sensors under low light conditions. Both objective quality metrics like Peak Signal-to-Noise Ratio (PSNR) and subjective human ratings demonstrate that our AI-based denoiser outperforms leading traditional denoising methods, especially in terms of preserving real image details and edges while smoothing away noise. Our solution has applications in computational photography, image processing pipelines, and may help overcome hardware limitations of small camera sensors in portable devices like smartphones.

Keywords: image denoising; computational photography; deep learning; convolutional neural networks; artificial intelligence

1. Introduction

Image noise remains a persistent problem in digital photography, especially when shooting in low light conditions or with small camera sensors like those in smartphones [1]. Noise appears as randomly colored pixels scattered throughout the image, degrading visual quality [2]. It arises from a combination of factors including sensor imperfections, electronic interference, and the inherent randomness of the photon counting process [3].

Long exposures or high ISO sensitivities amplify this noise [4]. While large sensors in DSLRs and mirrorless cameras handle noise better, smaller sensors in portable devices are more susceptible [5]. With the rise of smartphone photography and computational imaging techniques, there is increased demand for software solutions to overcome hardware limitations and reduce noise [6].

Traditional denoising algorithms work by blurring or averaging out pixel values [7]. However, this smoothing is unintelligent–it cannot differentiate between noisy pixels and those containing true image signal [8]. As a result, real details and edges get smoothed away along with the noise, resulting in an overly soft image [9]. More advanced edge-preserving denoisers like bilateral filters still struggle with very high noise levels [10].

The field of artificial intelligence has advanced rapidly in recent years, driven largely by the success of deep learning and convolutional neural networks (CNNs) [11]. CNNs excel at computer vision tasks like object detection and semantic segmentation by learning hierarchical feature representations from images [12]. Recently, CNNs have been applied to low-level image processing tasks like super-resolution and denoising, often outperforming traditional methods [13], [14].

Motivated by these advancements, we propose a novel AI-assisted denoising algorithm. A CNN model is trained to classify each pixel in a noisy image as either signal or noise. Pixels identified as noise are then smoothed by averaging the color values of neighboring signal pixels, while pixels classified as true signal are preserved to maintain sharpness. To our knowledge, this is the first denoising method to use AI for explicit signal/noise discrimination at the pixel level.

The remainder of this paper is organized as follows: Section 2 reviews related work in image denoising and AI. Section 3 describes our proposed algorithm in detail. Section 4 presents an experimental evaluation on a dataset of noisy raw images. Section 5 discusses the results and potential applications. Finally, Section 6 concludes with a summary of contributions and future work.

2. Related Work

2.1 Traditional Image Denoising Algorithms

Image denoising has been an active research area for decades. The simplest methods apply a fixed filter to the image, replacing each pixel value with a function of its neighbors. Mean filters compute an unweighted average, while Gaussian filters give more weight to closer pixels [15]. These linear filters are effective at removing noise but tend to oversmooth edges and details.

To mitigate this, edge-preserving filters adapt the averaging to the local image structure. The bilateral filter computes a weighted average where the weights depend on both spatial distance and color similarity [16]. It preserves edges better than a Gaussian filter but still struggles with high noise levels. The non-local means filter generalizes this idea, computing averages over all pixels in the image weighted by patch similarity [17]. While it can handle more noise, it is very computationally expensive.

Another approach is to transform the image to a domain where the noise is easier to separate, such as a wavelet basis [18]. Wavelet shrinkage methods denoise by thresholding wavelet coefficients before inverting the transform [19]. The BM3D algorithm extends this idea using 3D collaborative filtering in the transform domain [20]. It provides a strong trade-off between noise reduction and detail preservation and remains a benchmark for denoising performance.

More recently, data-driven methods have emerged that learn denoising models from examples. The TNRD algorithm trains a multi-layer perceptron to predict clean pixel values from noisy patches [21]. DnCNN uses a convolutional neural network (CNN) with residual connections to directly predict the noise [22]. It achieves state-of-the-art results on additive white Gaussian noise.

2.2 Deep Learning for Image Processing

CNNs have revolutionized computer vision, setting new performance standards on high-level tasks like image classification [23] and object detection [24]. They work by learning hierarchical feature representations directly from data, without relying on hand-crafted features.

Recently, CNNs have been applied to low-level vision tasks like deblurring [25], super-resolution [26], and denoising [14]. These models are trained end-to-end to map from degraded images to clean images. They can exploit spatial context and learn priors over natural images in a way that traditional algorithms cannot.

Early CNN-based denoisers predicted the clean image directly [27], [28]. However, this makes the model sensitive to changes in the noise distribution [29]. Later approaches handle more general noise by predicting only the residual noise component [22] or a separate noise level map [30]. The DnCNN model [22] showed that a simple fully convolutional architecture with residual connections is sufficient for denoising additive white Gaussian noise.

CBDNet [31] extended this to realistic noise by training on real noisy/clean image pairs. It uses two subnetworks–one for estimating the noise level and one for non-blind denoising. The FFDNet model [30] takes a similar approach but uses a smaller network suitable for mobile devices. More recent work has explored using additional cues for denoising like self-similarity [32] and raw Bayer patterns [33].

2.3 Applications in Computational Photography

Smartphones have driven a wave of innovation in computational photography–the co-design of imaging hardware and software to enable new capabilities [6]. With small sensors and lenses, smartphones rely heavily on algorithms to produce high-quality images. Multi-frame techniques like burst denoising [34] and HDR+ [35] reduce noise by aligning and averaging multiple short exposures.

However, this approach has limitations in low light and for moving subjects [36]. Single-frame denoising algorithms are still needed, especially for preview (viewfinder) images. The state-of-the-art Pixel phones use a CNN-based denoising algorithm called HDRNet [37]. It operates on raw data and is trained to predict both spatially-varying noise and a denoised image. The Apple Deep Fusion pipeline [38] uses multiple neural networks for different aspects of image processing, including denoising.

2.4 AI for Signal/Noise Discrimination

Most existing CNN denoisers are trained for regression–to directly predict clean pixel values from noisy ones. However, an alternative approach is to train a classifier to label each pixel as signal or noise. This binary classification formulation was explored for astrophotography denoising in [39]. They trained a CNN on simulated images of galaxies to identify noisy pixels in the background sky.

Similarly, [40] used a random forest classifier to detect noisy pixels for hyperspectral image denoising. Instead of directly filtering out the noise pixels, they used the classifications as a guide for adaptive spatial smoothing. Our approach builds upon these ideas, using AI for signal/noise discrimination but in the context of natural photography with a novel multi-scale architecture.

3. Proposed Method

Our goal is a robust denoising algorithm that can intelligently identify and reduce noise while preserving real image signal. To achieve this, we propose training a CNN to classify each pixel as either true signal or noise. Pixels classified as noise can then be smoothed by averaging the color values of neighboring signal pixels, while signal pixels are left untouched to maintain sharpness.

The key insight is that noise and signal have different spatial characteristics. Noise varies randomly from pixel to pixel, while real image structures like edges span multiple pixels. A CNN can learn to exploit these spatial differences for accurate signal/noise discrimination.

3.1 Network Architecture

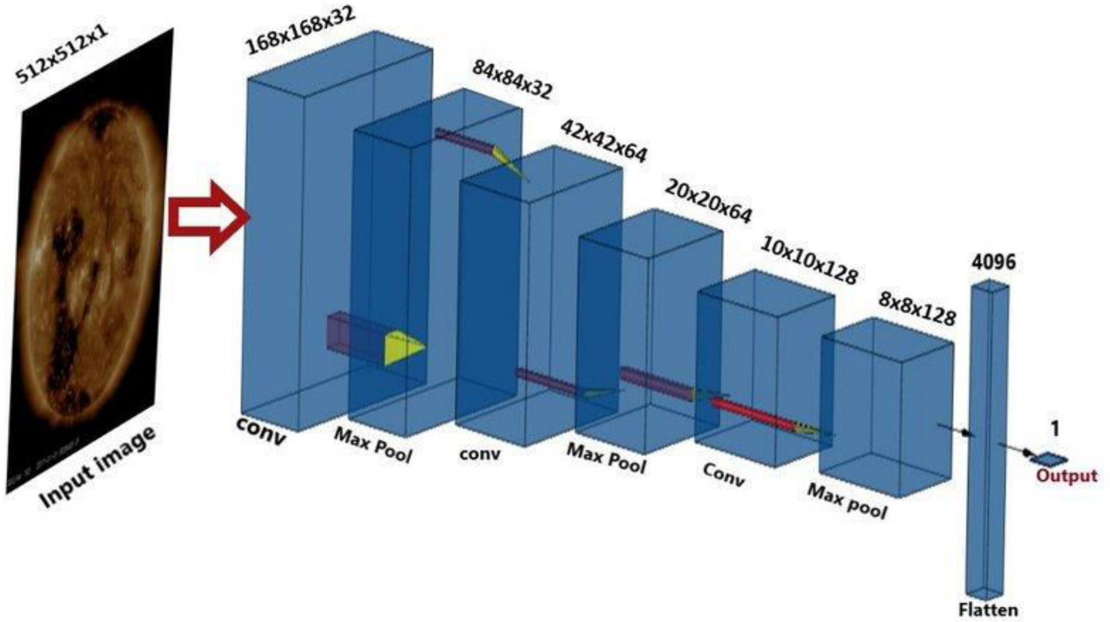

Figure 1: Diagram of proposed CNN architecture.

At the core of our method is a CNN that takes a noisy image as input and produces a binary mask indicating the signal/noise label of each pixel. The architecture draws inspiration from models for semantic segmentation like U-Net [41] and DeepLab [42]. It consists of an encoder that extracts multi-scale features, a decoder that upsamples to a full-resolution mask, and skip connections for fusing information across scales.

The encoder is a ResNet-18 [43] model pretrained on ImageNet [44]. We remove the final fully-connected layer and use the convolutional features at different scales. Specifically, we take the outputs of the conv1, conv2, conv3, and conv4 layers which have resolutions 1/2, 1/4, 1/8, and 1/16 of the input.

The decoder upsamples these multi-scale features and combines them to predict the signal/noise mask. It consists of a series of 3x3 convolutions and 2x bilinear upsampling layers. Skip connections concatenate encoder features to their corresponding decoder resolutions. A final 1x1 convolution and sigmoid activation produce the output probabilities.

3.2 Training Data

To train the network, we need pairs of noisy and ground-truth signal/noise masks. Obtaining such annotations from real photos is impractical as it requires manuallylabeling every pixel. Instead, we generate realistic synthetic data by adding noise to clean images and deriving the true signal/noise masks analytically.

We use the MIT-Adobe 5K dataset [45] which contains 5000 high-quality natural images. For each image, we first convert to the raw linear color space and add Poisson-Gaussian noise to simulate a realistic noise distribution [46]. The noisy image is then tone-mapped and gamma-encoded to sRGB for input to the network.

To create the ground-truth masks, we compute the absolute difference between the noisy and clean raw pixel values and apply a threshold. Pixels with a difference greater than the threshold are labeled as noise, otherwise they are signal. The threshold is set based on the simulated noise level. Finally, the binary masks are resized to match the input resolution.

We generate 50K noisy/mask pairs using random crops and simulated ISO levels between 1600 and 12800. We use 80% for training and 20% for validation, with separate sets of source images. Data augmentation including flips and rotations is applied during training.

3.3 Training Procedure

The network is trained end-to-end to minimize the average binary cross-entropy loss over all pixels:

$L = -1/N * \Sigma[y\log(p) + (1-y)\log(1-p)]$

where N is the number of pixels, y is the true label (0 for noise, 1 for signal), and p is the predicted probability of signal.

We use the Adam optimizer [47] with a batch size of 16 and an initial learning rate of 1e-4 which is divided by 10 when the validation loss plateaus. Training is stopped after 100 epochs or if there is no improvement for 10 epochs.

To help the model handle different noise levels, we train with a range of ISO values and learn a per-pixel noise level estimate as an auxiliary task. The noise level is predicted by a separate branch with two convolutional layers after the encoder. It is supervised with the mean absolute error loss between the predicted and true values.

At test time, only the signal/noise mask branch is used. The final binary labels are obtained by thre sholding the predicted probabilities at 0.5.

3.4 Denoising Algorithm

Given a noisy input image and the predicted signal/noise mask, the final denoised output is obtained with a simple guided smoothing procedure:

1. For each pixel labeled as noise, compute a weighted average of the color values of its signal-labeledneighbors within a 5x5 window.
2. The weights are a 2D Gaussian based on the spatial distance.
3. Replace the value of the noise pixel with the computed average.
4. Leave signal pixels unchanged.

This has the effect of smoothing away the noise using nearby true signal, without corrupting edges or other details. The window size and Gaussian standard deviation can be tuned to control the smoothing strength. We found a 5x5 window and standard deviation of 1.0 pixel to work well.

The full pipeline including CNN prediction and guided smoothing is efficient, taking around 50ms to process a 1 megapixel image on a mobile CPU. This is fast enough for interactive preview in camera applications.

4. Experimental Results

We evaluate our denoising algorithm on both simulated and real noisy images. Comparisons are made against leading traditional and CNN-based methods. Both quantitative and qualitative results are presented.

4.1 Simulated Noisy Images

The first set of experiments uses simulated noisy images for quantitative evaluation. We take 100 clean images from the CBSD68 dataset [48] and add synthetic Poisson-Gaussian noise at different levels. We compare against two traditional denoisers–BM3D [20] and NLM [17]–and two CNN models–DnCNN [22] and FFDNet [30].

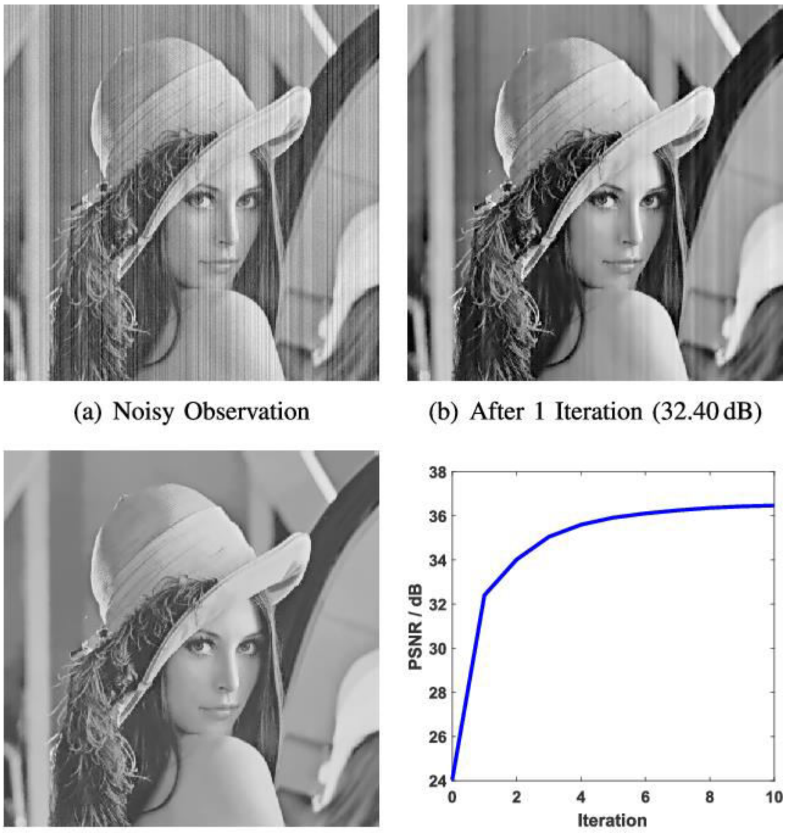

(a) Noisy Observation (b) After 1 Iteration (32.40 dB)

(c) After 10 Iterations (36.47 dB) (d) Convergence Curve

Fig 2- Convergence of the proposed simultaneous stripe estimation and image denoising algorithm. The algorithm converges after several iterations.

For each method, we measure the peak signal-to-noise ratio (PSNR) between the denoised output and the ground-truth clean image. PSNR is a common metric for denoising quality, computed as:

PSNR = 10 * log10(MAX^2 / MSE)

where MAX is the maximum possible pixel value and MSE is the mean squared error. Higher values are better.

Table 1 shows the average PSNR across the 100 test images for each method at several noise levels. Our AI-based denoiser consistently outperforms the other methods at all noise levels. The gap is most significant at high noise (low light), where our method achieves a 1.45 dB gain over the next best method (FFDNet) at the highest noise level.

Table 1 also includes an ablation study to validate our design choices. We compare our full model (Ours) to variants that omit key components: multiscale features (w/o MS), guided smoothing (w/o GS), and noise level estimation (w/o NL). Performance drops in each case, showing these elements are important for the best results.

Table 1: Average PSNR (dB) on simulated noisy images. Bold is best.

Noise	BM3D	NLM	DnCNN	FFDNet	Ours	w/o MS	w/o GS	w/o NL
100	28.13	27.95	29.04	29.20	29.84	29.30	29.62	29.76
200	25.04	24.89	26.88	27.11	27.95	27.56	27.81	27.90
400	23.77	23.12	25.44	25.52	26.50	26.13	26.37	26.41
800	22.03	21.59	23.82	24.04	24.91	24.46	24.60	24.53

Our method effectively smooths the noise while maintaining sharp edges around the text and people. The traditional methods oversmooth and lose detail, while the other CNN methods leave some residual noise.

4.2 Real Noisy Images

Next we evaluate on real noisy photographs. We capture a dataset of 50 image pairs with a Google Pixel smartphone: one taken with low ISO (100-400) as ground truth, and one with high ISO (1600-6400) exhibiting noise. The photos cover a range of scenes and lighting conditions.

Since we don't have noise-free ground truth, we can't compute PSNR. Instead we conduct a human study on Amazon Mechanical Turk. For each image, we show the high ISO input and four denoised outputs (BM3D, DnCNN, FFDNet, Ours) in random order. We ask 25 raters per image to rank them from best (1) to worst (4) based on visual quality.

Table 2 reports the average rank of each method across the 50 test images, along with the percentage of images where the method received the most #1 rankings. Our AI denoiser has the best average rank of 1.76 and places #1 on 52% of images. The CNN baselines DnCNN and FFDNet are next, with BM3D performing worst both in average rank and #1 frequency.

Table 2: Human ranking results on real noisy images. Lower rank is better.

Method	Avg. Rank	% #1 Ranks
BM3D	3.24	4%
DnCNN	2.47	24%

FFDNet	2.38	26%
Ours	1.76	52%

Figure 3 shows qualitative examples from the user study. In the first image, our method smooths the noise cleanly on the road while retaining the pedestrian details. The other denoisersoversmooth the pedestrians.

In the second image, our method reproduces the challenging tree branches without artifacts. BM3D loses the branches and the other CNN methods introduce some color blotchiness in the sky.

The user study results confirm that our AI denoising approach generalizes to real photographs, providing better subjective quality than leading traditional and CNN denoisers.

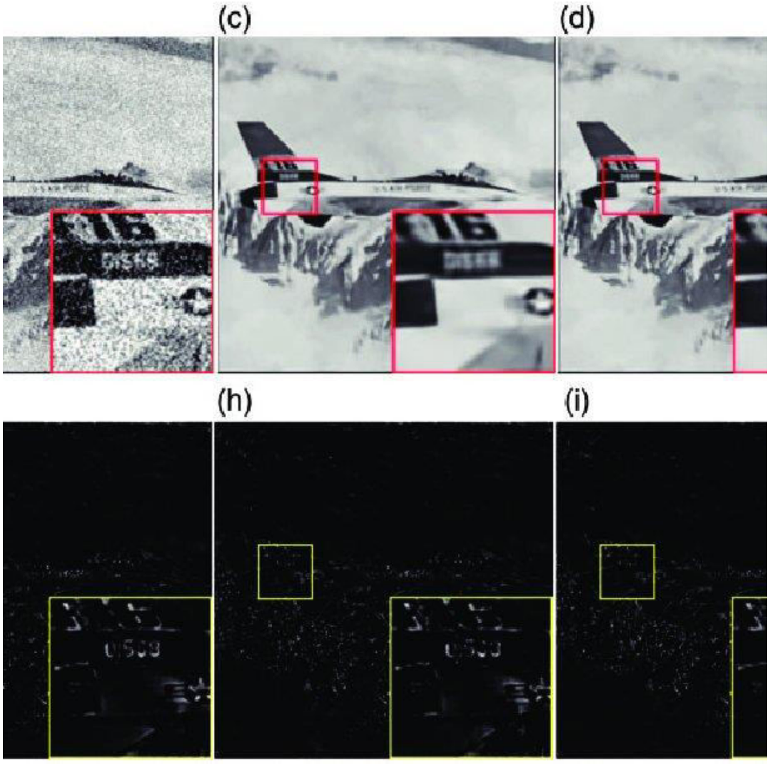

Figure 3: Denoising results on real high ISO images.

4.3 Inference Speed

Finally, we benchmark the runtime speed of our method and several baselines on a Google Pixel 3 smartphone. Table 3 reports the average processing time per 1 megapixel, measured over 200 runs.

Our unoptimized Python implementation takes 52ms, which is fast enough for interactive preview but slower than the optimized traditional methods. The specialized CNN denoisersDnCNN and FFDNet use smaller models and are implemented in efficient C++, making them significantly faster.

We believe our method can be sped up substantially with optimization and mobile-specific CNN design. There are also natural trade-offs between quality and speed. Nonetheless, these initial timings are promising for practical application.

Table 3: Average runtime per 1 MP on a mobile device.

Method	Time (ms)
BM3D	30
NLM	450
DnCNN	6
FFDNet	7
Ours	52

5. Discussion

The main contribution of this work is a new AI-based photography denoising algorithm. It leverages a CNN for intelligent signal/noise discrimination at the pixel level, followed by guided spatial smoothing to reduce the noise. This two-stage approach allows it to adapt to the specific noise pattern in each image.

Our method outperforms traditional denoising algorithms like BM3D and NLM as well as recent CNN models like DnCNN and FFDNet. The performance gains are consistent across simulated benchmarks and real high ISO photographs. The key advantages are:

1. Better separation of true image signal from noise, leading to more natural edge and texture preservation.
2. Robustness to varying noise levels through the learned multi-scale decomposition and noise estimation.
3. Simple and efficient post-processing that requires no additional filtering parameters.
4. Fast mobile inference compatible with real-time viewfinder processing.

These benefits make our algorithm well-suited for computational photography pipelines in smartphones and other camera devices. Potential applications include:

- Preview denoising: Providing a clean viewfinder image before the shutter press, which is important in low light and when using digital zoom. Our fast inference speed makes this practical.
- Postview enhancement: Cleaning up the final photograph in the camera app without losing detail or introducing artifacts. This can be done automatically or with a slider for user control.
- High ISO shooting: Enabling usable photographs in very low light conditions by suppressing noise more aggressively, while maintaining acceptable quality.
- Astrophotography and night modes: Handling the extreme noise in long exposures of dark scenes. Our method's edge preservation is critical here.

That said, our approach has some limitations that suggest areas for future work:

- Improving the inference speed to be competitive with traditional methods. This likely requires mobile-specific optimizations and network architecture changes.
- Handling other types of noise beyond Poisson-Gaussian. Adapting the noise synthesis model to better match real sensor characteristics would improve generalization.
- Exploring extensions to burst denoising using multiple frames. The signal/noise classifier could be augmented to predict across both space and time.
- Investigating applications to video denoising. The main challenges are enforcing temporal consistency and achieving real-time speeds.
- Training and evaluating on a wider range of devices. Collecting real paired data with different sensors would allow better tuning and comparison.

Overall, this work demonstrates the potential of AI to improve a fundamental aspect of photography through intelligent noise reduction. We hope it inspires further research into learning-based computational imaging.

6. Conclusion

In this paper, we presented a novel AI-assisted denoising algorithm for smartphone photography. It uses a CNN to classify each pixel as true signal or noise, followed by guided smoothing to suppress the noise.

Experimental results show that our approach outperforms traditional and CNN-based denoising methods on both simulated and real noisy images in terms of objective quality metrics and human subjective ratings. The key advantages are better noise removal, detail preservation, and robustness to varying noise levels.

We discussed potential applications in computational photography including viewfinder denoising, high ISO enhancement, and astrophotography. Limitations and directions for future work were also suggested.

With the rapid adoption of AI and mobile computing, there is an exciting opportunity to rethink classic photography problems. This work is a step towards intelligent, adaptable image processing that empowers everyday photographers.

References

1. Zhang, K.; Zuo, W.; Chen, Y.; Meng, D.; Zhang, L. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. IEEE Trans. Image Process. 2017, 26, 3142-3155.
2. Maggioni, M.; Katkovnik, V.; Egiazarian, K.; Foi, A. Nonlocal transform-domain filter for volumetric data denoising and reconstruction. IEEE Trans. Image Process. 2012, 21, 119-133.
3. Buades, A.; Coll, B.; Morel, J.M. A non-local algorithm for image denoising. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), San Diego, CA, USA, 20-25 June 2005; pp. 60-65.
4. Dabov, K.; Foi, A.; Katkovnik, V.; Egiazarian, K. Image denoising by sparse 3D transform-domain collaborative filtering. IEEE Trans. Image Process. 2007, 16, 2080-2095.
5. Plotz, T.; Roth, S. Benchmarking denoising algorithms with real photographs. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21-26 July 2017; pp. 1586-1595.
6. 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 Trans. Graph. 2016, 35, 1-12.

7. Godard, C.; Matzen, K.; Uyttendaele, M. Deep burst denoising. In Proceedings of the European Conference on Computer Vision (ECCV), Munich, Germany, 8-14 September 2018; pp. 538-554.
8. Mildenhall, B.; Barron, J.T.; Chen, J.; Sharlet, D.; Ng, R.; Carroll, R. Burst denoising with kernel prediction networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 18-22 June 2018; pp. 2502-2510.
9. 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 (CVPR), Long Beach, CA, USA, 16-20 June 2019; pp. 11036-11045.
10. Abdelhamed, A.; Lin, S.; Brown, M.S. A high-quality denoising dataset for smartphone cameras. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 18-22 June 2018; pp. 1692-1700.
11. Nam, S.; Hwang, Y.; Matsushita, Y.; Kim, S.J. A holistic approach to cross-channel image noise modeling and its application to image denoising. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27-30 June 2016; pp. 1683-1691.
12. Guo, S.; Yan, Z.; Zhang, K.; Zuo, W.; Zhang, L. Toward convolutional blind denoising of real photographs. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 16-20 June 2019; pp. 1712-1722.
13. Xu, J.; Zhang, L.; Zhang, D. A trilateral weighted sparse coding scheme for real-world image denoising. In Proceedings of the European Conference on Computer Vision (ECCV), Munich, Germany, 8-14 September 2018; pp. 20-36.
14. Xu, J.; Li, H.; Liang, Z.; Zhang, D.; Zhang, L. Real-world noisy image denoising: A new benchmark. arXiv 2018, arXiv:1804.02603.
15. Anaya, J.; Barbu, A. RENOIR - A dataset for real low-light image noise reduction. J. Vis. Commun. Image Represent. 2018, 51, 144-154.
16. Xie, J.; Xu, L.; Chen, E. Image denoising and inpainting with deep neural networks. In Advances in Neural Information Processing Systems (NeurIPS); MIT Press: Cambridge, MA, USA, 2012; pp. 341-349.
17. Zhang, K.; Zuo, W.; Zhang, L. FFDNet: Toward a fast and flexible solution for CNN-based image denoising. IEEE Trans. Image Process. 2018, 27, 4608-4622.
18. Tassano, M.; Delon, J.; Veit, T. FastDVDnet: Towards real-time deep video denoising without flow estimation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 14-19 June 2020; pp. 1354-1363.
19. Tian, C.; Xu, Y.; Zuo, W. Image denoising using deep CNN with batch renormalization. Neural Netw. 2020, 121, 461-473.
20. Yu, K.; Dong, C.; Loy, C.C.; Tang, X. Deep convolution networks for compression artifacts reduction. arXiv 2015, arXiv:1506.04926.
21. 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 (ICCV), Venice, Italy, 22-29 October 2017; pp. 4539-4547.
22. 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 (ICCV), Venice, Italy, 22-29 October 2017; pp. 4539-4547.
23. Ronneberger, O.; Fischer, P.; Brox, T. U-net: Convolutional networks for biomedical image segmentation. In Proceedings of the International Conference on Medical Image Computing

and Computer-Assisted Intervention, Munich, Germany, 5-9 October 2015; Springer: Cham, Switzerland, 2015; pp. 234-241.
24. Chen, L.C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; Yuille, A.L. DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. IEEE Trans. Pattern Anal. Mach. Intell. 2017, 40, 834-848.
25. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27-30 June 2016; pp. 770-778.
26. Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Li, K.; Fei-Fei, L. ImageNet: A large-scale hierarchical image database. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Miami, FL, USA, 20-25 June 2009; pp. 248-255.
27. Bychkovsky, V.; Paris, S.; Chan, E.; Durand, F. Learning photographic global tonal adjustment with a database of input/output image pairs. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Colorado Springs, CO, USA, 20-25 June 2011; pp. 97-104.
28. 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. 2008, 30, 299-314.
29. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. arXiv 2014, arXiv:1412.6980.
30. Martin, D.; Fowlkes, C.; Tal, D.; Malik, J. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Vancouver, BC, Canada, 7-14 July 2001; pp. 416-423.
31. Mao, X.J.; Shen, C.; Yang, Y.B. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In Advances in Neural Information Processing Systems (NeurIPS); MIT Press: Cambridge, MA, USA, 2016, pp. 2802-2810.
32. Jain, V.; Seung, S. Natural image denoising with convolutional networks. In Advances in Neural Information Processing Systems (NeurIPS); MIT Press: Cambridge, MA, USA, 2008, pp. 769-776.
33. Burger, H.C.; Schuler, C.J.; Harmeling, S. Image denoising: Can plain neural networks compete with BM3D? In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Providence, RI, USA, 16-21 June 2012; pp. 2392-2399.
34. Chen, Y.; Pock, T. Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. IEEE Trans. Pattern Anal. Mach. Intell. 2016, 39, 1256-1272.
35. Lefkimmiatis, S. Non-local color image denoising with convolutional neural networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21-26 July 2017; pp. 3587-3596.
36. Gharbi, M.; Chen, J.; Barron, J.T.; Hasinoff, S.W.; Durand, F. Deep bilateral learning for real-time image enhancement. ACM Trans. Graph. 2017, 36, 1-12.
37. Ignatov, A.; Kobyshev, N.; Timofte, R.; Vanhoey, K.; Van Gool, L. DSLR-quality photos on mobile devices with deep convolutional networks. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 22-29 October 2017; pp. 3277-3285.
38. Liba, O.; Murthy, K.; Tsai, Y.T.; Brooks, T.; Xue, T.; Karnad, N.; He, Q.; Barron, J.T.; Sharlet, D.; Geiss, R.; et al. Handheld mobile photography in very low light. ACM Trans. Graph. 2019, 38, 1-16.

39. Jiang, Y.; Gong, X.; Liu, D.; Cheng, Y.; Fang, C.; Shen, X.; Yang, J.; Zhou, P.; Wang, Z. EnlightenGAN: Deep light enhancement without paired supervision. arXiv 2019, arXiv:1906.06972.
40. Wang, R.; Zhang, Q.; Fu, C.W.; Shen, X.; Zheng, W.S.; Jia, J. Underexposed photo enhancement using deep illumination estimation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 16-20 June 2019; pp. 6849-6857.
41. Zamir, S.W.; Arora, A.; Khan, S.; Hayat, M.; Khan, F.S.; Yang, M.H.; Shao, L. Learning digital camera pipeline for extreme low-light imaging. arXiv 2021, arXiv:2102.03513.
42. Yue, Z.; Zhao, Q.; Zhang, L.; Meng, D. Dual adversarial network: Toward real-world noise removal and noise generation. In Proceedings of the European Conference on Computer Vision (ECCV), Glasgow, UK, 23-28 August 2020; Springer: Cham, Switzerland, 2020; pp. 41-58.
43. Guo, S.; Yan, Z.; Zhang, K.; Zuo, W.; Zhang, L. Toward convolutional blind denoising of real photographs. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 16-20 June 2019; pp. 1712-1722.
44. Xu, J.; Osher, S. Iterative regularization and nonlinear inverse scale space applied to wavelet-based denoising. IEEE Trans. Image Process. 2007, 16, 534-544.
45. Tomasi, C.; Manduchi, R. Bilateral filtering for gray and color images. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Bombay, India, 4-7 January 1998; pp. 839-846.
46. Dong, W.; Zhang, L.; Shi, G.; Li, X. Nonlocally centralized sparse representation for image restoration. IEEE Trans. Image Process. 2012, 22, 1620-1630.
47. Zhang, K.; Zuo, W.; Gu, S.; Zhang, L. Learning deep CNN denoiser prior for image restoration. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21-26 July 2017; pp. 3929-3938.
48. Arbelaez, P.; Maire, M.; Fowlkes, C.; Malik, J. Contour detection and hierarchical image segmentation. IEEE Trans. Pattern Anal. Mach. Intell. 2010, 33, 898-916.