

Generating Training Data for Denoising Real RGB Images via Camera Pipeline Simulation

Ronnachai Jaroensri Camille Biscarrat Miika Aittala Frédo Durand
MIT CSAIL
{tiam,cjbisc,miika,fredo}@csail.mit.edu




Figure 1: (a) Real noise in cellphone-processed JPEG pictures is very different from uncorrelated Gaussian noise widely assumed (see Fig. 2). (b) A blind denoiser trained on additive white Gaussian noise (AWGN) is unable to recognize the noise pattern resulting and denoise the image. (c) In contrast, the network trained on realistic noise model generated by our pipeline was able to denoise properly, resulting in over 3dB improvement.

Abstract

Image reconstruction techniques such as denoising often need to be applied to the RGB output of cameras and cell-phones. Unfortunately, the commonly used additive white noise (AWGN) models do not accurately reproduce the noise and the degradation encountered on these inputs. This is particularly important for learning-based techniques, because the mismatch between training and real world data will hurt their generalization. This paper aims to accurately simulate the degradation and noise transformation performed by camera pipelines. This allows us to generate realistic degradation in RGB images that can be used to train machine learning models. We use our simulation to study the importance of noise modeling for learning-based denoising. Our study shows that a realistic noise model is required for learning to denoise real JPEG images. A neural network trained on realistic noise outperforms the one trained with AWGN by 3 dB. An ablation study of our pipeline shows that simulating denoising and demosaicking is important to this improvement and that realistic demosaicking algorithms, which have been rarely considered, is

needed. We believe this simulation will also be useful for other image reconstruction tasks, and we will distribute our code publicly.

1. Introduction

Most image reconstruction techniques such as denoising operate on RGB images, either JPEGs directly from a camera or RAW files that have been demosaicked later. In this paper, we show that the simple additive Gaussian noise (AWGN) usually used in the literature [33, 22, 32] does not accurately model the artifacts observed in real image. This is especially true when working from JPEG images, which undergo a long pipeline that includes operations such as demosaicking, denoising, and compression that dramatically transform the noise (see Fig. 2). The mismatch of noise profiles can have a strong adverse effect on performance, especially for learning-based approaches.

Several works have shown that image reconstruction tasks can benefit from better noise modeling [4, 9, 5, 13]. However, most noise and degradation models found in the



Figure 2: AWGN noise (left) and real noise from Pixel XL phones (right). The noise in real images is processed extensively by the camera pipeline. For this particular camera, the artifact is long-grained (right) and very different from the fine chroma pattern of AWGN (left).

literature remain simplistic. For example, most works do not consider demosaicking artifacts, and the ones that do typically use bilinear demosaicking [5, 13, 25], which is rarely used in real consumer cameras [15, 14].

In this paper, we propose a camera simulation pipeline that can be used to realistically simulate camera processing of images. We implement over 40 individual modules that can be custom-built into a camera pipeline. While not exhaustive, they cover a good range of typical camera modules such as tonemapping, demosaicking, and denoising. From these modules, we build a pipeline capable of processing RAW images, with some manual tuning, into visually similar RGB images that some cellphone cameras produce.

We believe this pipeline can be used to generate data for many low-level vision tasks. To demonstrate, we use our pipeline to study the importance of noise modeling in supervised denoising of JPEG images. We generate different versions of a dataset, where the pipeline processing parameters and, therefore, the noise characteristics vary. We train state-of-the-art denoising convolutional neural networks (CNN) [26] on these datasets and measure their performance on real processed images. We show that networks trained with our realistic pipeline outperformed ones that are trained on AWGN by roughly 3 dB on real test images (see Fig. 1 and Fig. 7).

To understand how our pipeline contributes to this improvement, we train additional networks with different combinations of simulation components. We find that performance drops markedly if denoising and/or demosaicking components are removed. Furthermore, the choice of demosaicking algorithm is also important. Using bilinear demosaicking in camera simulation pipelines [5, 13, 25] leads to less effective denoising compared to using edge-aware methods such as the Adaptive Homogeneity-Directed (AHD) algorithm [17].

While we are not the first to propose a camera simula-

tion [19, 11], our main contribution is to integrate it into the learning pipeline and use it to show the importance of realistic simulation for learning-based image restoration tasks. We summarize our contributions as follow:

1. We propose a camera simulator that is expressive enough to simulate processing of real cameras. We believe that this simulator will be useful for many learning-based image restoration tasks.
2. Using this simulator, we study the importance of realistic noise modeling for denoising real images. We show that:
 - (a) A realistic noise model is beneficial for denoising real JPEG images and leads to superior performance compared to AWGN.
 - (b) Denoising filters and demosaicking are the most important components for simulating realistic noise.
 - (c) Realistic demosaicking algorithm is important and leads to improvement over bilinear demosaicking commonly used in camera simulation pipelines [5, 13, 25].

All code and evaluation datasets will be released publicly.

2. Related Work

Classical denoising techniques often create probabilistic models of the noise and signal and use this model to derive a denoising algorithm. Wavelet coring is based on the observation that noise is usually smaller than the image signal, resulting in smaller wavelet coefficients that can be suppressed [29, 27, 10]. The current state-of-the-art classical denoising method remains BM3D [8]. The algorithm performs non-local matching within the image and average these matched blocks together. These methods typically assume an AWGN model in order to simplify their modeling effort.

With the growing popularity of CNNs [21], learning-based denoising is becoming prevalent. DnCNN [32] uses CNNs to predict a residual map that corrects noisy images. N3Net [26] formulates a differentiable version of nearest neighbor search to further improve DnCNN. FFDNet [33] attempts to address spatially varying noise by appending noise level maps to the input of DnCNN. Despite many improvements, these works perform very similarly, with often less than 0.5 dB difference. Moreover, they assume spatially uncorrelated noise, which is not true for real noisy JPEG images.

Many works in image processing are recognizing the important of noise modeling. [12, 15] jointly model noise with




Figure 3: Our camera pipeline consists of four main stages: artifact generation, demosaicking, denoising, and post-processing.

their demosaicking task and found it to improve their performance. [4] uses an adhoc noise model that simulates spatial correlation of noise. They found this to significantly boost the quality of their deblurring results.

Recent denoising work proposes simulating camera pipelines. [5] unprocesses JPEG images to get RAW representation and focus on RAW-to-RAW denoising. Very related to our work is [13] who also uses simulated camera pipeline to supplement real training data. However, these works tend to assume a limited camera pipeline and do not evaluate on real processed images. Our work follows in the same spirit, though we aim to accurately model realistic camera pipelines, and evaluate our results with real images.

3. Camera Simulation Pipeline

Our camera pipeline is designed to mirror typical camera processing stages. We build our pipeline from individual modules that are easily extensible. Additionally, we also include an *artifact generation* stage that simulates the degradation of the image signal, by introducing artifacts such as noise and motion blur.

Fig. 3 shows the 4 main stages of our pipeline: artifact simulation, demosaicking, denoising, and postprocessing. In sum, we have over 70 parameters that control the behavior of our pipeline. We describe each stage as follows.

Artifact Generation. The first stage of our pipeline is the physical artifact simulators. It aims to simulate the physical degradation process that happens before the sensor. It includes motion blur, chromatic aberration, multiplicative exposure adjustment, and noise.

Noise at the sensor is largely uncorrelated and zero-mean. So we only simulate spatially uncorrelated noise here. Because photon noise is Poisson in nature and the sensor read noise is Gaussian, we provide both additive and multiplicative noise to simulate the two effects. We optionally mosaick input images before adding noise if the user wishes to simulate the Bayer pattern, which will then be demosaicked and processed in the subsequent stages.

Demosaicking. If the input is mosaicked, we demosaick the input at this stage. To our knowledge, most cameras use more advanced algorithms, such as the Adaptive Homogeneity-Directed (AHD) algorithm [17]. We provide a Python adaptation of the reference AHD algorithm and an algorithm developed by Kodak Inc. [16] implemented in high performance language Halide [23]. We provide

bilinear demosaicking as well because it has been widely used in the recent camera simulation literature [5, 25, 13]. Hot/Dead pixel correction and white balance also occur here.

Cellphone Denoising. Demosaicking noisy images tends to result in long-grained artifacts (as Fig. 2 shows). Our third stage applies denoising to the image. We include three denoising algorithms—bilateral filters [30], median filters [18], and wavelet coring [10]—with the option to turn each algorithm on/off as well as reorder them. Performing tonemapping prior to denoising can be beneficial for denoising different intensity ranges because it can non-linearly compress a particular range of the intensity leading to a more smoothing effect. We include a pre-tonemapping operator, which can be a gamma or an s-shape tone curve.

Tonemapping and Post-processing. The last stage performs postprocessing that aims to generally improve the aesthetics of the image. We include saturation adjustment, tonemapping for additional tone/contrast enhancement, unsharp mask for detail enhancement, and JPEG compression for JPEG compression artifacts simulation.

We build our pipeline largely on top of the PyTorch package [24]. This allows us to readily integrate it into learning frameworks. Because some of the software used does not support differentiation [31], we do not utilize the differentiability of our pipeline. While we believe that our pipeline is realistic and rich in features, it is by no means a comprehensive set of operations implemented by camera manufacturers. In particular, we do not consider automatic adjustments such as auto-exposure and auto white-balance. These modules will become crucial in automatic processing of cellphone images. Nonetheless, we demonstrate in section 3.1 that our pipeline can emulate cellphone processing, given an appropriate set of parameters.

3.1. Camera Simulation Evaluation

We show that our pipeline is expressive enough to perform the same image processing as a camera’s image signal processing (ISP) unit. Fortunately, modern cellphones allow RAW and JPEG captures from the same exposure. This means that if we are able to process the RAW image into the same, or similar, JPEG image as captured, we will have successfully emulated the camera’s ISP.

Because our pipeline is missing automatic adjustments commonly found in a camera ISP, we allow adjustments of parameters to individual images. In particular, tones and color balance are adaptively adjusted per image/scene. Denoising parameters, on the other hand, are held fixed per camera to reduce the risk of overfitting.

We captured RAW + JPEG images with an iPhone 7, an iPhone 8, and a Samsung Galaxy S7. We chose these phones because they are recent enough to allow the capture of RAW but not too recent as to have superior imaging




Figure 4: Comparison of our processed RAW and camera JPEG for iPhone 7, iPhone 8, and Samsung Galaxy S7. The tone and noise pattern match well. For more results, please refer to the supplementary material.

sensors. Including both iOS and Android phones demonstrates the versatility of our pipeline because they are likely to have different processings and imaging sensors. We captured approximately 10 scenes on each phone. We focus on low-light scenes so that the noise pattern is visible, allowing us to evaluate the similarity of the processing results of our pipeline.

To find the best parameters for each image, we performed grid search of tone and color parameters, using L2 loss on the luminance and chrominance channels respectively. We then hand-tuned each parameter to obtain the final result.

3.2. Evaluation Result

Fig. 4 shows the comparison of our pipeline processing and the camera JPEG. Our simulation obtains an average PSNR of 28.9dB - 30.8dB and SSIM of 0.873 - 0.888 across the three phones. In addition to these metrics, we visually




Figure 5: Some examples where we do not match the appearance well. The PSNR for the top picture is high because while we are able to match the tone well, but the noise pattern is over-smoothed.

inspect the noise pattern in both the camera JPEG and our processed RAW, and we find them to be subjectively similar.

While our pipeline can achieve good PSNR and SSIM numbers, these metrics tend to over-emphasize tones and low-frequency image content. We find some visible differences in the level of smoothing across intensity levels that may require per-image denoising parameter tuning to remove (see Fig. 5). Nonetheless, the level of smoothing is satisfactory overall, and we show that this pipeline can be used to improve end-to-end denoising task in Section 4.

4. Denoising Experiment

We demonstrate that our pipeline can be used for generating training data for real image denoising. Denoising is a well-studied topic, yet, few works have attempted to model realistic noise correlation. We show that the lack of realistic noise can be detrimental to denoising performance.

We synthetically generate our datasets using the camera simulation pipeline described in Section 3. Using different sets of parameters, we seek to answer two important questions: does having a realistic noise model matter, and if so, how realistic does the noise model have to be?

Many works on denoising are shifting towards denoising RAW images, where noise is easier to model [12, 5, 7]. We focus on denoising JPEG images for two reasons. First, most image reconstruction algorithms deal primarily with JPEG images. But for these methods, using an additive

white Gaussian noise model with JPEG images can lead to inferior results [4]. Second, many photographs taken are in JPEG format because it is often easier to work with and uses less storage. Therefore, any algorithm that aims to be widely adopted must be able to deal with the degradation present in the JPEG images.

We primarily focus on learning-based approaches, for which synthetic data generation is useful for training. Methods that do not require training data may still find it useful to generate realistic test data as an alternative to collecting their own dataset.

4.1. Training Data and Architecture

Our denoising setup aims to denoise RGB images that have been processed by the camera. Fig. 6 shows our training setup. It starts from an input JPEG image with gamma compression. We undo the gamma compression to obtain a linear image to be degraded by our camera simulation pipeline (Section 3). The degraded output is then fed into a denoising CNN. Finally, the denoised image is compared to the original linear, clean RGB image to provide training signal to the denoising network.

Choice of Modules in the Camera Pipeline Simulation. Since we focus only on the noise pattern, we turn off all tonemapping and color operations. This way the denoising network does not have to learn to adjust tones, simplifying the learning problem considerably.

We observe that real cellphone denoising is often a combination of bilateral and median filters, so we use these two algorithms in our cellphone simulation pipeline. We find that the Kodak algorithm [16] and AHD [17] perform roughly the same, so we choose the Kodak algorithm for which we have a more efficient implementation.

Parameters of the Pipeline. We set the configuration of our processing pipeline based on the range of values observed during our experiment (see Section 3.2). For simplicity, we randomize each parameter independently. We choose noise strengths based on measurement data from the iPhone 7 and the Samsung Galaxy S6 at various ISO [1, 2]. We exaggerate the noise strength to ensure that the network sees very noisy samples in the training set. Table 1 lists noise strengths and processing performed on each of our datasets.

Source Dataset. We use the MIT-Adobe5k dataset [6] as our input images because it has high-quality photographs. We use their expert-C retouched images so that the input and target tones are representative of JPEG images. We downsample the images by 4x to reduce any remaining noise and artifacts. We extract 5 patches randomly from each image in the dataset, resulting in a total of 25k patches available for training.

Denoising Network. Since the focus of this work is not the network architecture, we used the author’s implementa-

Table 1: Configuration of Training Datasets

Training Data	Gaussian STD	Poisson Mult Factor	Additional Processing
AWGN	0 - 0.2	0	None
Add-Mult WGN	0 - 0.1	0 - 0.02	None
Ours	0 - 0.1	0 - 0.02	Demosaicking, Denoising, Post-processing
Samsung S7 Measurement @ ISO800	0.007	0.02	N/A

tion of the Neural Nearest Neighbor network [26], which has been shown to achieve state of the art result in denoising. We follow the author’s training method, using the Adam optimizer [20] with learning rate of 0.001. The author also noted that increasing learning rate decay is beneficial, so we decay the learning rate by 10^{-5} over 100 epochs (instead of 10^{-5} over 50 epochs in the original paper) and train for 100 epochs.

Performance Consideration. Because our dataset is synthetic, we are able to generate it on-the-fly. This allows us to rapidly prototype and change configurations without pre-generating the entire dataset. Additionally, each input patch receives different randomized processing parameters in each epoch, which increases the complexity of our dataset. Our pipeline implementation is based largely on PyTorch modules [24] and uses the high performance Halide language [23]. While performance varies with the system and configuration, we are able to largely saturate a machine with a Tesla P100 GPU and 32 CPU cores (80x80 patch, batch size=32). Training takes roughly 9 hours.

4.2. Testing Data

Because we focus on denoising real JPEG photographs, real JPEG images are required to measure the denoising performance. This is challenging because we do not have access to the blackbox camera processing, and our pipeline cannot process large amounts of images automatically. Furthermore, some artifacts in the JPEG images cannot be removed by averaging.

Existing datasets do not provide the required clean JPEG images. [3] and [7] provide only RAW images, while [25] uses simple processing which may not be realistic. [28] provides short- and long-exposure image pairs, but they do not keep exposure levels constant, resulting in large tone shifts between the ground truth and noisy images. Furthermore, we find the noise in their long-exposure ground truth images to still be significant.

Because of these limitations, we use averaged RAW images from bursts as the target. Noise in RAW images is




Figure 6: Our training setup. We use our pipeline described in Section 3 to generate the data for the denoising network. By varying configuration of the pipeline, datasets that simulate different noise profiles (AWGN, processed JPEG from cellphones, etc) can be generated, allowing a comparative study of these noise profiles and their effectiveness.

Table 2: Quantitative comparison of training data between AWGN vs our pipeline for different cellphone cameras.

Metric	Input vs Ground truth	Training Data		
		AWGN	AMWGN	Our Pipeline
iPhone 8				
PSNR	32.4	32.3	32.6	35.2
SSIM	0.788	0.788	0.799	0.892
Pixel XL				
PSNR	31.2	31.3	31.6	35.1
SSIM	0.760	0.761	0.775	0.881
Samsung Galaxy S7				
PSNR	41.1	41.0	41.4	42.5
SSIM	0.933	0.931	0.939	0.954

zero-mean and can be reduced by averaging. However, using RAW images as the target requires demosaicking and normalizing the tone. Because PSNR and SSIM are very sensitive to tone change, we normalize each ground truth image to the output image at test time by matching their means and standard deviations per color channel.

We collected test images using the iPhone 8, Pixel XL, and Samsung Galaxy S7 to test generalization across camera models. For each phone, roughly 20-25 scenes were captured, and for each scene, one high-ISO image and a burst of 10 low-ISO images were taken. All images were captured in the RAW + JPEG format and the exposure were kept roughly constant. We used sturdy tripods and avoided moving objects and reflections as much as possible. We also set a timer and used a shutter cable to avoid any movement that resulted from interacting with the phones.

5. Denoising Results

In this Section, we report the findings of our denoising experiments.

5.1. Additive/Multiplicative White Gaussian Noise (AMWGN) vs Realistic Noise

The network trained on our dataset significantly outperforms ones that were trained with additive/multiplicative

Gaussian noise. Table 2 shows denoising results of N3Net [26] trained with different datasets. On the iPhone 8 and Pixel XL test sets, the model trained on our dataset achieved a 3 dB higher PSNR and nearly 0.1 higher SSIM. On the Samsung Galaxy S7, the improvement is approximately 1.5 dB in PSNR and 0.015 in SSIM. These are significant margins because many recent denoising works often report improvements that are less than 0.5-1 dB [26, 32].

Visual inspection of the resulting denoised patch reveals that the AMWGN models seem to ignore noise entirely—the output patch is almost identical to the input patch, as Fig. 7 shows. On the iPhone 8 test data, the PSNR between the input and output patches are over 50dB, and the SSIM is over 0.996 (vs 35.7 dB and 0.856 for our model).

To show that our AWGN model works correctly, we pass the patches with additive Gaussian noise with STD of 0.1 (on a 0-1 scale) to the model. Fig. 8 shows the denoising results. The AWGN networks are able to properly denoise the patches with PSNRs of 36.6 dB and 36.0 dB for the additive and additive-multiplicative models, respectively. This suggests that their performance on real images is likely the result of a mismatch between real test JPEG image and the additive Gaussian noise training data, and not the faulty implementation of our models.

Denoising Demosaicked RAW. Most image reconstruction algorithms are designed for RGB images, so when working with RAW images, demosaicking is often applied (except for a few works [12, 15]). We demosaick our real RAW noisy images and use them as test input. We find that our data outperforms AWGN by 7-9 dB in PSNR and 0.2-0.3 on SSIM, depending on the demosaicking algorithms applied (the training always uses Kodak [16]).

5.2. Ablation Study

In order to understand the essential features of our pipeline, we train additional networks with different components of our pipeline turned off. We group the components based on stages outlined in Section 3: demosaicking, denoising, and post-processing.

Denoising and Demosaicking. We find demosaicking




Figure 7: Sample output patches from denoising networks trained with AWGN vs our dataset, on iPhone 8, Pixel XL, and Samsung Galaxy S7 test data. More results in the supplementary material.

and denoising to be important to the smoothing of the image. Fig. 9 shows a sample patch from three different networks: ones that are trained without post-processing, without denoising, and without demosaicking. The network

trained without post-processing produces the smoothest outputs, while the other two retain long-grained artifacts present in the input image. Table 3 shows the quantitative result for these networks. While the PSNRs are compara-




Figure 8: Our AWGN-trained and Add-Mult WGN models are able to denoise patches with AWGN, suggesting that the models are working correctly. More result is in the supplementary material.

Table 3: Quantitative comparison of handicapped data generation by turning one component off at a time (iPhone 8 test data).

Metric	Full-Pipeline	No Post-processing	No Denoising	No Demosaicking
PSNR	35.2	34.0	34.7	33.9
SSIM	0.892	0.866	0.870	0.846




Figure 9: Qualitative evaluation of our ablation datasets (on the iPhone 8 test data). The network trained without post-processing is able to smooth the real image, while the network without denoising and demosaicking shows less smoothing of the noise.

ble, removing demosaicking suffers the largest reduction in SSIM, confirming our qualitative observation. For brevity, we only show results on iPhone 8 test data, but we observe similar trends on both Pixel XL and Samsung Galaxy S7 test data.

Table 4: Quantitative comparison on the choice of demosaicking algorithm (iPhone 8 test data).

Metric	Full Pipeline	Kodak [16]	AHD[17]	Bilinear
PSNR	35.2	33.6	32.9	31.1
SSIM	0.892	0.840	0.821	0.746




Figure 10: Qualitative evaluation of simulation with different demosaicking algorithms (iPhone 8 test data). The networks trained with edge-aware demosaicking (Kodak [16], AHD [17]) were able to smooth images well, while one with the commonly used bilinear demosaicking retains the most artifacts. More result is in the supplementary material.

Choice of Demosaicking Algorithm. We further investigate the choice of demosaicking algorithm used, because most works that simulate the camera processing pipeline use bilinear demosaicking [5, 13, 25].

As Fig. 10 shows, network trained with bilinear interpolation retains the most JPEG artifacts in their denoising result. On the other hand, the networks trained with the other two edge-aware demosaicking algorithms are able to remove more of these artifacts. Table 4 shows the quantitative results. The AHD [17] and the Kodak algorithm [16] outperforms bilinear demosaicking by more than 2dB on PSNR and over 0.07 on SSIM.

6. Conclusion and Future Work

We have proposed a realistic camera pipeline simulation that is expressive enough to process RAW inputs into JPEG images that is visually similar to the ones cameras produce. We use this simulation to generate realistic datasets for training denoising CNNs and show that it improves the performance of such networks on real JPEG images by over 3dB. Demosaicking and denoising seem to be the most important components of our pipeline that enable such improvement. Removing either of them leads to a significant drop in the quality of the denoised output. Using correct algorithms for these components is also important. The bilinear demosaicking algorithm commonly used in previous camera simulation work [5, 13, 25] leads to a significant performance drop, while edge-aware algorithms such as AHD [17] do not.

While we have shown our pipeline is useful and realistic, it still requires significant manual tuning in order to match the appearance of the processed JPEG. The ability to automatically match the appearance is an interesting future direction. This will help ensure realism, so that the generated data can be used for any arbitrary camera models.

Acknowledgments

The authors would like to thank the Toyota Research Institute for their generous support of the projects. We thank Tzu-Mao Li for his helpful comments, and Luke Anderson for his help revising this draft.

References

- [1] Read noise in dns versus iso setting. [http://www.photonstophotos.net/Charts/RN_ADU.htm#Apple%20iPhone%207_12,Samsung%20Galaxy%20S6\(S5K2P2\)_10](http://www.photonstophotos.net/Charts/RN_ADU.htm#Apple%20iPhone%207_12,Samsung%20Galaxy%20S6(S5K2P2)_10). Accessed: 2018-10-30. **5**
- [2] Samsung galaxy s6 edge : Measurements - dxomark. <https://www.dxomark.com/Cameras/Samsung/Galaxy-S6-Edge---Measurements>. Accessed: 2018-10-30. **5**
- [3] A. Abdelhamed, S. Lin, and M. S. Brown. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1692–1700, 2018. **5**
- [4] M. Aittala and F. Durand. Burst image deblurring using permutation invariant convolutional neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 731–747, 2018. **1, 3, 5**
- [5] T. Brooks, B. Mildenhall, T. Xue, J. Chen, D. Sharlet, and J. T. Barron. Unprocessing images for learned raw denoising. *arXiv preprint arXiv:1811.11127*, 2018. **1, 2, 3, 4, 8**
- [6] V. Bychkovsky, S. Paris, E. Chan, and F. Durand. Learning photographic global tonal adjustment with a database of input/output image pairs. In *CVPR 2011*, pages 97–104. IEEE, 2011. **5**
- [7] C. Chen, Q. Chen, J. Xu, and V. Koltun. Learning to see in the dark. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3291–3300, 2018. **4, 5**
- [8] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising with block-matching and 3d filtering. In *Image Processing: Algorithms and Systems, Neural Networks, and Machine Learning*, volume 6064, page 606414. International Society for Optics and Photonics, 2006. **2**
- [9] J. Dong, J. Pan, D. Sun, Z. Su, and M.-H. Yang. Learning data terms for non-blind deblurring. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 748–763, 2018. **1**
- [10] D. L. Donoho and J. M. Johnstone. Ideal spatial adaptation by wavelet shrinkage. *biometrika*, 81(3):425–455, 1994. **2, 3**
- [11] J. E. Farrell, F. Xiao, P. B. Catrysse, and B. A. Wandell. A simulation tool for evaluating digital camera image quality. In *Image Quality and System Performance*, volume 5294, pages 124–132. International Society for Optics and Photonics, 2003. **2**
- [12] M. Gharbi, G. Chaurasia, S. Paris, and F. Durand. Deep joint demosaicking and denoising. *ACM Transactions on Graphics (TOG)*, 35(6):191, 2016. **2, 4, 6**
- [13] S. Guo, Z. Yan, K. Zhang, W. Zuo, and L. Zhang. Toward convolutional blind denoising of real photographs. *arXiv preprint arXiv:1807.04686*, 2018. **1, 2, 3, 8**
- [14] S. W. Hasinoff, D. Sharlet, R. Geiss, A. Adams, J. T. Barron, F. Kainz, J. Chen, and M. Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics (TOG)*, 35(6):192, 2016. **2**
- [15] F. Heide, M. Steinberger, Y.-T. Tsai, M. Rouf, D. Pajak, D. Reddy, O. Gallo, J. Liu, W. Heidrich, K. Egiazarian, et al. Flexisp: A flexible camera image processing framework. *ACM Transactions on Graphics (TOG)*, 33(6):231, 2014. **2, 6**
- [16] R. H. Hibbard. Apparatus and method for adaptively interpolating a full color image utilizing luminance gradients, Jan. 17 1995. US Patent 5,382,976. **3, 5, 6, 8**
- [17] K. Hirakawa and T. W. Parks. Adaptive homogeneity-directed demosaicing algorithm. *IEEE Transactions on Image Processing*, 14(3):360–369, 2005. **2, 3, 5, 8**
- [18] T. Huang, G. Yang, and G. Tang. A fast two-dimensional median filtering algorithm. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 27(1):13–18, 1979. **3**
- [19] H. C. Karaimer and M. S. Brown. A software platform for manipulating the camera imaging pipeline. In *European Conference on Computer Vision*, pages 429–444. Springer, 2016. **2**
- [20] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. **5**
- [21] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. **2**
- [22] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila. Noise2noise: Learning image restoration without clean data. *arXiv preprint arXiv:1803.04189*, 2018. **1**
- [23] T.-M. Li, M. Gharbi, A. Adams, F. Durand, and J. Ragan-Kelley. Differentiable programming for image processing and deep learning in halide. *ACM Transactions on Graphics (TOG)*, 37(4):139, 2018. **3, 5**
- [24] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017. **3, 5**
- [25] T. Plotz and S. Roth. Benchmarking denoising algorithms with real photographs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1586–1595, 2017. **2, 3, 5, 8**
- [26] T. Plötz and S. Roth. Neural nearest neighbors networks. In *Advances in Neural Information Processing Systems*, pages 1095–1106, 2018. **1, 2, 5, 6**
- [27] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli. Image denoising using scale mixtures of gaussians in the wavelet domain. *IEEE Trans Image Processing*, 12(11), 2003. **2**
- [28] E. Schwartz, R. Giryes, and A. M. Bronstein. Deepisp: Toward learning an end-to-end image processing pipeline. *IEEE Transactions on Image Processing*, 28(2):912–923, 2019. **5**

- [29] E. P. Simoncelli and E. H. Adelson. Noise removal via bayesian wavelet coring. In *Proceedings of 3rd IEEE International Conference on Image Processing*, volume 1, pages 379–382. IEEE, 1996. [2](#)
- [30] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *null*, page 839. IEEE, 1998. [3](#)
- [31] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, and the scikit-image contributors. scikit-image: image processing in Python. *PeerJ*, 2:e453, 6 2014. [3](#)
- [32] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017. [1](#), [2](#), [6](#)
- [33] K. Zhang, W. Zuo, and L. Zhang. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018. [1](#), [2](#)