

IDR: Self-Supervised Image Denoising via Iterative Data Refinement

Yi Zhang¹ Dasong Li¹ Ka Lung Law² Xiaogang Wang¹
 Hongwei Qin² Hongsheng Li¹

¹The Chinese University of Hong Kong ²SenseTime Research

zhangyi@link.cuhk.edu.hk

Abstract

The lack of large-scale noisy-clean image pairs restricts supervised denoising methods' deployment in actual applications. While existing unsupervised methods are able to learn image denoising without ground-truth clean images, they either show poor performance or work under impractical settings (e.g., paired noisy images). In this paper, we present a practical unsupervised image denoising method to achieve state-of-the-art denoising performance. Our method only requires single noisy images and a noise model, which is easily accessible in practical raw image denoising. It performs two steps iteratively: (1) Constructing a noisier-noisy dataset with random noise from the noise model; (2) training a model on the noisier-noisy dataset and using the trained model to refine noisy images to obtain the targets used in the next round. We further approximate our full iterative method with a fast algorithm for more efficient training while keeping its original high performance. Experiments on real-world, synthetic, and correlated noise show that our proposed unsupervised denoising approach has superior performances over existing unsupervised methods and competitive performance with supervised methods. In addition, we argue that existing denoising datasets are of low quality and contain only a small number of scenes. To evaluate raw image denoising performance in real-world applications, we build a high-quality raw image dataset SenseNoise-500 that contains 500 real-life scenes. The dataset can serve as a strong benchmark for better evaluating raw image denoising. Code and dataset will be released¹

1. Introduction

In recent years, deep neural networks have made great progress in image denoising. The mainstream of the advances is powered by supervised learning-based approaches [6, 11, 13, 22, 29]. However, large-scale datasets with high-quality ground-truth are still unavailable due to high human




Figure 1. **Top row:** The training targets created at different epochs. They have been refined progressively in our approach. **Bottom row:** Comparison with previous works: DBSN [31] (ECCV 2020), Noisier2Noise [23] (CVPR 2020).

and time costs, which limit the learning-based approaches' applications in real-world scenarios.

A few unsupervised learning-based approaches have been proposed to learn image denoising without ground truth. However, existing unsupervised denoising methods either show poor performance or work under impractical settings. One research direction focuses on unsupervised learning of denoising from single noisy images [5, 14, 16, 28, 31]. Although such training data are convenient to collect, learning from single noisy images shows relatively poor performance since the task is highly ill-posed. Another research direction focuses on learning from paired noisy images [18]. It utilizes two perfectly aligned images of the same scene as the training pairs. Such a method significantly improves the denoising performance and can sometimes even outperform supervised learning-based methods. However, multiple perfectly aligned noisy images of the same scene are sometimes challenging to capture in practice, especially for moving scenes. One recent approach applies noise models to add synthetic noise on the already noisy images to create *noisier-noisy* image pairs for training a denoising network to partially capture the noise prior [23]. However, such synthetic noisier-noisy training pairs' distributions cannot well fit that of the actual noisier-

¹<https://github.com/zhangyi-3/IDR>

clean image pairs (we call such train-inference distribution mismatch the *data bias* between the two datasets), which limits the method’s denoising performance on actual noisy images. There are also methods [26,32] that train a separate network for denoising each image separately and show better results. However, such methods inevitably have higher computational costs.

In this paper, we propose an unsupervised denoising method that adopts a practical setting for preparing training data and achieves state-of-the-art performances. It only requires single noisy images and a noise model for training a denoising neural network. Note that the noise model is convenient to obtain in actual raw image denoising (*e.g.*, noise profile in DNG files [15]), and can be both point-wise and correlated types (*e.g.*, raw image noise, stripe noise, multiplicative noise).

To start with, we add noise from the known noise model to actual noisy images to create noisier-noisy dataset, which treats the created noisier images as inputs and actual noisy images as learning targets. Then, we conduct a study on learning denoising from such noisier-noisy datasets to empirically validate two findings: (1) the denoising model trained on the noisier-noisy dataset has the generalization capability of denoising actual noisy images; (2) reducing the mismatch (data bias) between the created noisier-noisy dataset and the ideal noisy-clean dataset results in a better denoising model to process the actual noisy images.

The two findings inspire us to propose the iterative data refinement (IDR) scheme, which iteratively reduces the data bias between the created noisier-noisy dataset and the ideal noisy-clean dataset to train better denoising models. In particular, we first train a denoising model based on the noisier-noisy dataset. According to our finding 1, the network trained on the noisier-noisy dataset can be generalized to denoise the actual noisy images. The denoised images together with the noise model can then be used to create a new intermediate dataset for a new round of training, which regards the denoised images from the previous round as the targets. As the data bias between the noisier-noisy and the ideal noisy-clean dataset is reduced (based on our finding 2), we can train a better denoising model on the newly constructed intermediate dataset in the next round. This newly trained model can denoise the original noisy images again, based on which, better learning targets can be created for the next round. This process can be performed iteratively to gradually train better denoising models. To alleviate the heavy training time cost of this iterative method, we further propose a fast version of our IDR by reducing the training epochs in each data refinement iteration.

To evaluate the denoising models’ performances in real-world raw-image denoising, we collect a low-light dataset *SenseNoise-500*, which contains 500 diverse scenes with perfectly aligned and high-quality ground truth. Our method outperforms existing unsupervised methods on both real-world and synthetic noise denoising tasks with a wide

range of noise types and noise levels. Other than pointwise noises, we also show our algorithm’s denoising capability on typical correlated noises. All results demonstrate the effectiveness and superiority of our proposed method.

Our contributions are summarized as follows:

- A self-supervised image denoising scheme, iterative data refinement (IDR), is proposed. It works under a more practical setting that only requires single noisy images and a noise model.
- We also propose a fast approximation of our full method IDR, which costs the same training time as training an ordinary denoising model while maintaining the similar performance as our full-version model.
- We collect a high-quality raw image denoising dataset for training denoising models and test their performance in real-world scenarios. It contains 500 diverse scenes, with images being well-aligned and the clean images are created with $> 60s$ exposure.
- Our method outperforms existing unsupervised methods on a wide range of noise types and noise levels. In some cases, it even shows comparable performance with fully supervised methods.

2. Related Work

Image denoising has been extensively studied in the literature. In this review, we divide them into three categories: Non-learning denoising methods, learning from clean data, and learning from noisy data. In the end, we also review some popular real-world denoising benchmarks which are widely used in other papers.

Non-learning denoising methods. Representative non-learning methods usually are based on different assumptions. This category of methods include anisotropic diffusion, total variation denoising [27], wavelet domain denoising [25], sparse coding [20], etc. Another important category of methods is driven by image self-similarity. Classical methods such as NLM [7, 8], BM3D [12, 21] utilize similar patches to remove the noise.

Learning from clean data. Learning from clean data is the mainstream in image denoising. Many works [1, 3, 11] have explored learning from natural noisy and clean image pairs directly. While the ideal training pairs can be used, the amount of those pairs is quite limited. Another way is to use external clean images and create synthetic noisy images based on the noise model [13, 22, 35, 36]. As a result, there exists a natural domain gap between the real dataset and the synthetic dataset. To reduce the domain gap between synthetic images and natural noisy images, some progress [4, 6, 33] has been made to process the external data to the target domain so that the external data becomes similar to the testing data. Other works [2, 10, 13, 30, 37] focus on building a more accurate noise model for the real-word

noise. Generating more realistic noise for training helps the model generalize better in real applications.

Learning from noisy data. One line in unsupervised learning is single-image-based [26, 28, 32]. They require re-training models for different images, which usually is not practical. As reported in self2self [26], it costs around 1.2 GPU (RTX 2080Ti) hours to retrain and test for a 256×256 image. On the other hand, batch-based unsupervised approaches start from Noise2Noise (N2N) [18], which demonstrates how to learn image denoising with noisy image pairs. But noisy image pairs generally are not available in the real world and N2N also requires designing different loss functions according to the noise type. The followed-up works try to learn image denoising by using individual noisy images. Mask-based unsupervised approaches [5, 16, 17] design different mask schemas to blind pixels in the input noisy images. Then, they train the network to predict the masked pixels according to noisy pixels in the input receptive field. More recent work DBSN [31] uses unpaired clean images to improve unsupervised denoising results through a two-stage scheme. Noisier2noise [23] introduces the noise model into unsupervised learning and trains the network directly for the single noise level denoising. So there still exists the dataset bias.

Real-world denoising datasets. Several datasets have been proposed in the literature. Darmstadt Noise Dataset (DND) [24] collects 50 image pairs which can only be used for testing. And most of the images in DND have relatively low levels of noise. Smartphone Image Denoising Dataset (SID) [3] was proposed to evaluate denoising on smartphones. It captures multiple times for each scene and contains 10 indoor scenes, but occasionally exhibits obvious misalignments. SID dataset [11] was captured by DSLRs with the under-exposure strategy. It uses a single long-exposure frame as the ground truth.

3. Method

In this section, we first introduce the creation of noisier-noisy dataset, based on which, we conduct a study on learning denoising with the noisier-noisy dataset and empirically validate two findings. We then present our self-supervised denoising framework: iterative data refinement (IDR). To accelerate the training process of this iterative method, we further introduce a fast version of IDR to reduce the training time while leading to similar performance as the full version of our IDR.

3.1. Pilot Study on Data Bias of Learning Denoising

For training denoising models with a known noise model, a typical type of approaches [35] is to create synthetic noisy images via adding noise n onto the ground-truth clean images y , which is denoted as $y + n$ and used as network input. We name such synthetic training data *noisy-clean dataset* for learning the mapping $(y + n) \rightarrow y$. How-

ever, the ground-truth clean images y are generally difficult to obtain. Another plausible way is to add noise n onto the actual noisy images x to create the even noisier images $x + n$, where the later noisier $x + n$ are treated as inputs and the former noisy images x are used as learning targets. We call such training data *noisier-noisy dataset*. Although the later dataset is much easier to obtain, there is inevitable data bias between the above two types of datasets and the denoising network trained on the later noisier-noisy dataset cannot well handle the actual noisy images x . We conduct a pilot study as follows and empirically validate two important findings. Formally, the datasets are denoted as

$$\begin{aligned} \text{Noisy-clean Dataset: } & \{(y_i + n_i, y_i)\}_{i=1}^N, \\ \text{Noisier-noisy Dataset: } & \{(x_i + n_i, x_i)\}_{i=1}^N, \end{aligned} \quad (1)$$

where $x_i = y_i + n_i$.

Here, y_i is a clean image (which is very difficult to obtain in practice), $x_i = y_i + n_i$ is its corresponding noisy image, n_i is the random noise generated from the noise model, and $y_i + n_i^2$ denotes applying the sampled noise n_i to the actual noisy image y_i to create a synthetic but even noisier image. For simplicity, we use $\{y_i + n_i, y_i\}$ and $\{x_i + n_i, x_i\}$ to denote $\{(y_i + n_i, y_i)\}_{i=0}^N$ and $\{(x_i + n_i, x_i)\}_{i=0}^N$, respectively.

To analyze the data bias of training a network with noisier-noisy datasets, we train a U-Net from scratch on both the noisy-clean and noisier-noisy datasets³ and investigate three representative noise types, raw image, Gaussian, and correlated noise, with various noise levels. They cover common noise types *i.e.*, signal-dependent/signal-independent noises and point-wise/correlated noises. More setups of the noise models can be found in the caption of Fig. 2. For testing, we sample 4 discrete noise levels uniformly from the training noise levels and test the denoising models on actual noisy-clean image pairs, which have real ground truth.

As shown in Fig. 2, training on various noise types and noise levels show consistent results, which lead to the following two findings: (1) The denoising networks trained with noisier-noisy datasets can denoise the actual noisy images (red vs. blue dots in Fig. 2) with the same noise model. (2) With the same noise model, training on the less biased dataset towards the ideal noisy-clean dataset helps the denoising networks achieve better denoising performance (Table 1).

For the first finding, the models trained with noisier-noisy datasets (red dots in Fig. 2) always show better PSNR than the actual noisy images without any denoising (blue dots in Fig. 2) on all noise types, which indicate that the model can denoise actual noisy images $\{x_i\}$ to some extent

²Note that “+” here represents applying general noise to the images, which is **not limited to only additive noise**.

³We randomly select 70% images of the ImageNet validation dataset for training and the rest for testing.

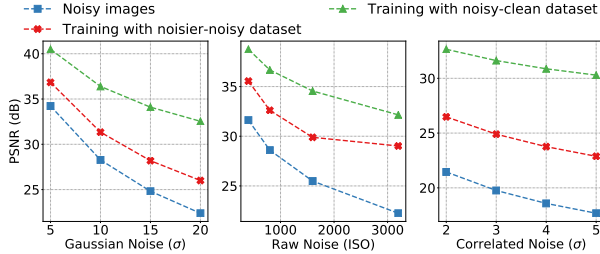


Figure 2. The denoising performance of denoising networks trained with the noisy-clean dataset and the noisier-noisy datasets of three representative noises and a range of noise levels. For Gaussian noise and correlated noise, we uniformly sample their noise level from [5, 20] and [2, 5] respectively. As for the real-world raw image noise*, we use the noise profile of a HUAWEI Mate20, which can be modeled as a Poisson-Gaussian noise model, and sample the ISO from [800, 3200] uniformly.

*Here, the “real-world raw image noise” is under the assumption of the Poisson-Gaussian noise model.

even with the existence of data bias between the noisier-noisy and noisy-clean datasets. Recently, both single-image-based [32] and batch-based [23] approaches have made similar conclusions on Gaussian noise but they only study a single noise level. Our investigation extends their observations to cover a wider range of noise types and noise levels.

For the second finding, we first remove all data bias in the noisier-noisy dataset, it becomes a noisy-clean dataset (zero data bias). The model trained based on it is considered to have the optimal performances (green dots in Fig. 2) on all three noise types and varying noise levels. We then synthesize different types and strengths of data bias to further verify the 2nd finding. Specifically, we apply Gaussian noise or Gaussian blur of different levels (denoted as σ in Table 1) to contaminate clean images y to create biased targets (denoted as y_{gn} and y_{gb}). The same three noise models (Gaussian, raw, correlated) with different noise levels as above are further added to the biased targets for creating noisier images ($y_{gn} + n$ and $y_{gb} + n$) as inputs for network training. Separate denoising networks are trained with biased datasets of different strengths and are tested on actual noisy images of the same noise model. The results in Table 1 show that the denoising networks trained on datasets with less data bias (smaller σ in Table 1) show better denoising performance and the second finding is consistent for two types of data bias, three different noise types, and three noise levels ($3 \times 2 \times 3 = 18$ separate denoising networks are trained and evaluated).

3.2. Iterative Data Refinement

The two findings in Section 3.1 inspire us to iteratively reduce the data bias between the noisier-noisy dataset and the ideal noisy-clean dataset. Gradually creating less biased datasets and training on them leads to promising denoising

Table 1. Denoising performance of models trained with biased datasets with different strengths. We corrupt the clean targets by applying Gaussian noise or Gaussian blur with different strengths (σ) to create biased targets (y_{gn} & y_{gb}) and add three different type of noise to create input noisier images ($y_{gn} + n$ & $y_{gb} + n$). The last six rows show the **PSNR drop** when models trained with different biased datasets to the ideal noisy-clean dataset (first row).

Training Datasets	σ	Gaussian Denoising	Raw Denoising	Correlated Denoising
$\{y + n, y\}$	-	30.896	35.263	30.484
$\{y_{gn} + n, y_{gn}\}$	5	-0.355	-3.100	-0.172
	3	-0.093	-1.620	-0.082
	1	-0.008	-0.310	-0.019
$\{y_{gb} + n, y_{gb}\}$	5	-0.074	-0.101	-0.944
	3	-0.069	-0.097	-0.595
	1	-0.067	-0.081	-0.150

performance.

We first train a model \mathcal{F}_0 based on the initial noisier-noisy dataset $\{x_i + n_i, x_i\}$. According to our finding 1 in Section 3.1, it can denoise the actual noisy images $\{x_i\}$ to some extent. Therefore, we can use it to denoise the noisy images and obtain a new set of refined training dataset:

$$\{\mathcal{F}_0(x_i) + n_i, \mathcal{F}_0(x_i)\}. \quad (2)$$

The newly constructed dataset shares the same number of images and scene contents as the original noisier-noisy dataset $\{x_i + n_i, x_i\}$. However, this new dataset has less data bias to the noisy-clean dataset since the L_2 distance of their targets have been reduced by our trained denoising model \mathcal{F}_0 ⁴ (according to our finding 1 in Section 3.1).

Then, using the less biased dataset construct above, we can train a new denoising model \mathcal{F}_1 from scratch in the next round:

$$\mathcal{F}_1 \leftarrow \{\mathcal{F}_0(x_i) + n_i, \mathcal{F}_0(x_i)\}. \quad (3)$$

Since the data bias between the new dataset $\{\mathcal{F}_0(x_i) + n_i, \mathcal{F}_0(x_i)\}$ and the ideal noisy-clean dataset $\{y_i + n_i, y_i\}$ is mitigated, according to our finding 2 in Section 3.1, the trained model \mathcal{F}_1 can better generalize to actual noisy images $\{x_i\}$ than \mathcal{F}_0 trained on the first-round noisier-noisy dataset.

Based on the improved model \mathcal{F}_1 , the above process can be performed iteratively and alternatively by refining the training targets and training better models as shown in Fig. 3. From the model’s perspective, once we have a denoising model \mathcal{F}_m , a better model \mathcal{F}_{m+1} can be trained with the constructed dataset:

$$\mathcal{F}_{m+1} \leftarrow \{\mathcal{F}_m(x_i) + n_i, \mathcal{F}_m(x_i)\}. \quad (4)$$

The previous model \mathcal{F}_m helps to close the gap between the new noisier-noisy dataset $\{\mathcal{F}_m(x_i) + n_i, \mathcal{F}_m(x_i)\}$ and the

⁴The PSNR improvements after denoised by \mathcal{F}_0 indicates the L_2 distance to clean targets has been reduced.

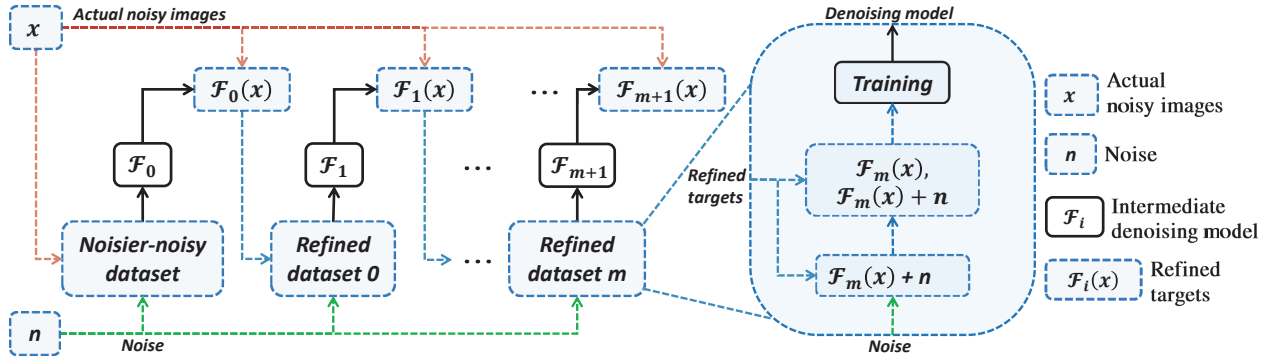


Figure 3. An overview of Iterative Data Refinement (IDR).

Algorithm 1: Fast Iterative Data Refinement

Input: Noisy images $\{x_i\}$; Noise model n ;
 Total epochs M ;
Output: Final model \mathcal{F}_M

- 1 Initialize the model \mathcal{F}_0 randomly ;
 - 2 Optimize the model \mathcal{F}_0 on the noisier-noisy dataset $\{x_i + n_i, x_i\}$ for one epoch ;
 - 3 **for** $m \leftarrow 1$ **to** M **do**
 - 4 Initialize the model $\mathcal{F}_m = \mathcal{F}_{m-1}$;
 - 5 Create “cleaner” targets $\{x_i^{(m)}\} \leftarrow \{\mathcal{F}_{m-1}(x_i)\}$;
 - 6 Construct the new training dataset $\{x_i^{(m)} + n_i, x_i^{(m)}\}$;
 - 7 Optimize current model \mathcal{F}_m by minimizing the L_1 loss for one epoch ;
 - 8 **end**
-

ideal noisy-clean dataset $\{y_i + n_i, y_i\}$. In the next round of training, the improvement of the newly created noisier-noisy dataset leads to the advance of our newly trained model \mathcal{F}_{m+1} . On the other hand, from the target’s perspective, a series of intermediate refined targets $\{\mathcal{F}_m(x_i)\}$ (for $m = 0, 1, \dots$) are produced by the intermediate denoising models. As shown in Fig. 1, the noise on training targets is removed and more textures are restored progressively. Moreover, unlike traditional iterative methods that denoise one noisy image several times and lead to losing textures heavily, our method refines the training dataset and denoises the noisy image once during inference.

3.3. Fast Iterative Data Refinement

While the above iterative data refinement does not increase time cost in inference, its training time is proportional to the number of rounds of the proposed self-supervised data refinement. To alleviate the heavy training time cost, we further propose a fast approximated training scheme. The whole algorithm is shown in the Algorithm 1.

We approximate our full method in two aspects. First, we introduce one-epoch training, which performs one round of data refinement after each epoch. Specifically, for each newly constructed dataset, we use it to train our model for only one epoch instead of training it until full convergence as before. In the next epoch, we train the denoising network based on a newly constructed dataset. We sacrifice the time needed for full model optimization but increase the iteration rounds of data refinement. The overhead we introduced is to produce refined targets before each training epoch. It costs less than 5% of the overall training time⁵. As a result, the total training time is reduced to almost the same as training a denoising model for only one round.

Second, when training on the new dataset at each epoch, our models are initialized by the model from the previous epoch instead of from scratch. This cumulative training strategy helps the denoising network converge faster with the proposed fast data refinement scheme and ensures that the final denoising model has been optimized continuously during the whole training process. The approximation algorithm also converges fast and keeps a similar performance as our full algorithm.

4. SenseNoise-500 dataset

There exists an obvious gap between existing raw image denoising benchmarks and real-world scenarios. For example, SID [11] was created with under-exposure photos which are rare in daily photography. The mobile benchmark SIDD [3] captures multiple images for each scene (10 indoor scenes). And, it suffers from global misalignments occasionally (e.g., see image Nos. 101 & 52 in SIDD) due to the different autofocus or OIS. To evaluate the denoising performance in real-world scenarios, we propose a new high-quality dataset under the common low-light conditions using a popular mobile CMOS sensor: IMX586 with 3000×4000 pixels. The dataset contains 500 scenes and

⁵The forward process of producing refined targets can be performed in parallel. We use batch size 32 for the inference. For training, our batch size is 4.



Figure 4. The qualitative results of compared methods on denoising images with Gaussian noise ($\sigma = 50$).

is larger than the commonly used denoising benchmarks: SIDD (10 scenes) [3], DND dataset (50 scenes) [24] and SID Sony (231 scenes) [11]. It was collected across 3 cities and covers common indoor and outdoor scenes in daily life. When we construct our dataset, there are some important considerations:

High-quality ground truth. The ground truths in the existing denoising benchmarks [3, 11] are still more or less noisy, especially at high noise levels. To obtain high-quality ground truths, we capture a sequence of 64 frames for each scene, among which, 4 frames of them are captured in normal exposure with the default camera setting and used as noisy testing images. The other 60 frames are long-exposure for creating the ground truth. For each pixel location, we collect all pixel values at the same location in the 60 long-exposure images and use the median value as the ground truth. The exposure time for each long-exposure frame lies in $[1s, 2s]$, which means that the total exposure time for creating each ground truth is between 60s and 120s. It guarantees that all ground truths are noise-free even at the highest noise level.

Alignment and noise calibration. To obtain well-aligned noisy/clean image pairs, we capture each image sequence on the tripod and control the mobile camera with a wireless shutter. All camera parameters are set and fixed in advance through camera2 API [19]. In this way, both normal-exposure and long-exposure frames are captured successively without any movement to the device. On the other hand, we try to avoid large object motions and dynamic illuminations in both indoor and outdoor scenes. Under this setting, all frames (long-exposure and normal-exposure) captured for one scene are perfectly aligned. While some companies may provide the noise profiles in raw files (*e.g.*, DNG [15]), they are usually not accurate. We re-calibrate

the noise parameters of IMX586 under the noise formulation in ELD [30]. However, we choose the Gaussian distribution to model the read noise, since we observe that it is more robust for mobile sensors. More details and examples of *SenseNoise-500* can be found in the supplementary materials.

5. Experiments

To evaluate our method, we follow previous works to start with synthetic noises on popular datasets. Then, we move to tackle more challenging real-world raw image denoising, including both DSLR and mobile datasets. Moreover, we conduct ablation studies to show the generalization of our method and make further comparisons. **More experiments can be found in supplementary materials.**

5.1. Implementation details

Training details. In all experiments, we use our fast IDR for comparisons. The network architecture and training settings are kept the same as N2N [18]. Specifically, we use a shallow U-Net architecture without batch normalization layers as our backbone. All models are trained from scratch with batch size 4 for 50k iterations. We use the Adam optimizer and L_1 loss function. The initial learning rate is fixed to 3×10^{-4} , and then halves twice at iteration 25k and 40k. Its input and output channels are differently set to adapt to different denoising tasks (*e.g.*, gray-image denoising has 1 input channel).

Datasets for Synthetic Experiments. For synthetic noises, we randomly extract patches of size 256×256 from the ImageNet validation dataset for training. For testing datasets, we select three largest of them for evaluation: Kodak (24 images) and BSDS300 (300 images) for color images, and BSD68 (68 images) for gray-scale images. Fol-

Table 2. The PSNR/SSIM results of Gaussian denoising on both sRGB datasets (Kodak and BSDS300) and the gray dataset (BSD64). One superior supervised method (MPRnet [34]) is also provided for comparison. Best and second best results are **highlighted** and underlined.

Dataset	σ	BM3D [21]	N2V [16]	Nr2n [23]	DBSN [31]	N2N [18]	Supervised	Ours	MPRnet [34]
Kodak	25	31.88 / 0.869	31.63 / 0.869	31.96 / 0.869	32.07 / 0.875	32.39 / 0.886	32.39 / 0.885	32.36 / 0.884	32.85 / 0.893
	50	28.64 / 0.772	28.57 / 0.776	28.73 / 0.770	28.81 / 0.783	29.23 / 0.803	<u>29.24 / 0.803</u>	29.27 / 0.803	29.78 / 0.817
BSDS300	25	30.47 / 0.863	30.72 / 0.874	29.57 / 0.815	31.12 / 0.881	31.39 / 0.889	<u>31.40 / 0.889</u>	31.48 / 0.890	31.87 / 0.897
	50	27.14 / 0.745	27.60 / 0.775	26.18 / 0.684	27.87 / 0.782	<u>28.17 / 0.799</u>	<u>28.17 / 0.799</u>	28.25 / 0.802	28.65 / 0.814
BSD64	25	28.55 / 0.782	27.64 / 0.781	NA	28.81 / 0.818	29.15 / 0.831	<u>29.20 / 0.835</u>	29.20 / 0.835	29.40 / 0.840
	50	25.59 / 0.670	25.46 / 0.681	NA	25.95 / 0.703	26.23 / 0.725	<u>26.24 / 0.727</u>	26.25 / 0.726	26.48 / 0.736
Average		28.71 / 0.784	28.60 / 0.793	29.11 / 0.785	29.11 / 0.807	29.43 / 0.822	<u>29.44 / 0.823</u>	29.47 / 0.823	29.84 / 0.833

Table 3. Denoising performances (PSNR/SSIM) of different compared methods on raw image denoising. For the SID dataset [11], we follow ELD [30] to show the results on ISO 100, 250, and 300. For our SenseNoise-500 dataset, we divide it into normal scenes (ISO ≤ 12800) and extreme scenes (ISO > 12800). Best and second best results are **highlighted** and underlined.

Dataset	ISO	BM3D [21]	N2S [5]	N2V [16]	N2N [18]	Ours	Supervised	MIT5K+unprocess
SID	100	32.92 / 0.758	35.61 / 0.830	35.64 / 0.827	37.83 / 0.898	<u>38.02 / 0.906</u>	38.60 / 0.912	37.67 / 0.884
	250	29.56 / 0.686	31.39 / 0.698	30.58 / 0.655	33.44 / 0.814	<u>33.88 / 0.819</u>	37.08 / 0.886	33.14 / 0.767
	300	28.88 / 0.674	29.78 / 0.649	28.88 / 0.595	31.61 / 0.769	<u>32.14 / 0.779</u>	36.29 / 0.874	31.36 / 0.727
SenseNoise-500	≤ 12800	41.05 / 0.958	31.05 / 0.867	41.86 / 0.961	41.22 / 0.956	<u>44.01 / 0.978</u>	44.80 / 0.981	43.36 / 0.978
	> 12800	37.54 / 0.940	27.95 / 0.871	38.28 / 0.948	39.53 / 0.965	<u>40.17 / 0.966</u>	40.23 / 0.966	39.92 / 0.965

lowing the previous work [18], we test on four diverse noise types: (1) Gaussian noise with training on the continuous noise levels $\sigma \in [0, 50]$ and testing on $\sigma = 25$ and 50 , (2) binomial and impulse noise with training on $p \in [0, 0.95]$ and testing on $p = 0.5$, and (3) correlated noise with $\sigma = 1$. **Datasets for Real-world Experiments.** For real-world raw image denoising, we train and test on the SID dataset (DSLR) [11] and the proposed *SenseNoise-500* dataset (smartphones). Those two datasets cover different camera sensors and low-light conditions. The training and testing sets for SID dataset are the same as SID [11]. For *SenseNoise-500* dataset, we randomly split 400 images for training and the other 100 images for testing. The patch size has been raised to 512×512 for better handling the low-light condition. As for the noise model, we recalibrate the noise profile using the recent ELD noise model [30] for both SID dataset and our *SenseNoise-500* dataset.

5.2. Main results

Due to the space limitation, we only put the results of Gaussian noise and real-world noise in the main text. Please find the results of other noises in supplementary materials.

Details of Gaussian noise experiments. We compare our method with six learning-based methods (N2N [18], N2V [16], N2S [5], DBSN [31] and Noisier2noise [23].) and also the representative traditional method BM3D [21]. For BM3D, we use the authors’ refined implementation⁶. For learning-based methods, all of them use the same U-Net architecture except for DBSN [31] provided by the authors, which is $10\times$ larger (97.4G FLOPS) than U-Net

(8.2G FLOPS). For DBSN, we use the author’s released models. All other methods are retrained using the author’s codes and keep the same training settings as N2N.

Results of Gaussian noise experiments. The PSNR and SSIM results are shown in Table 2. Our method outperforms all unsupervised methods on both color and gray images. And, surprisingly, our method even surpasses the supervised method slightly in most cases. Both our method and N2N introduce new training targets to unsupervised denoising, and show clear gains compared with other methods that only use single noisy images. But, the well-aligned paired noisy images required by N2N are difficult to capture in practice. On the contrary, the noise model used in our method can be easily calibrated and already saved in raw images. They can be easily obtained or recalibrated [30, 37] in read-world raw image denoising. Some qualitative results are shown in Fig. 4. Our method restores more textures and colors while keeping the high-quality denoising results.

Details of real-world raw noise experiments. For BM3D, we use the mixture Gaussian noise model to estimate the noise level. For learning-based methods, we re-train all methods by adopting the authors’ code.

Results of real-world raw noise experiments. The results are shown in Table 3. Some qualitative results are shown in Fig. 5. Our method performs consistently better than other unsupervised methods on both DSLR and smartphones datasets. Both the traditional method BM3D [21] and the unsupervised methods N2V [16] and N2S [5] with single noisy images cannot work well and show more than

⁶<https://pypi.org/project/bm3d/>

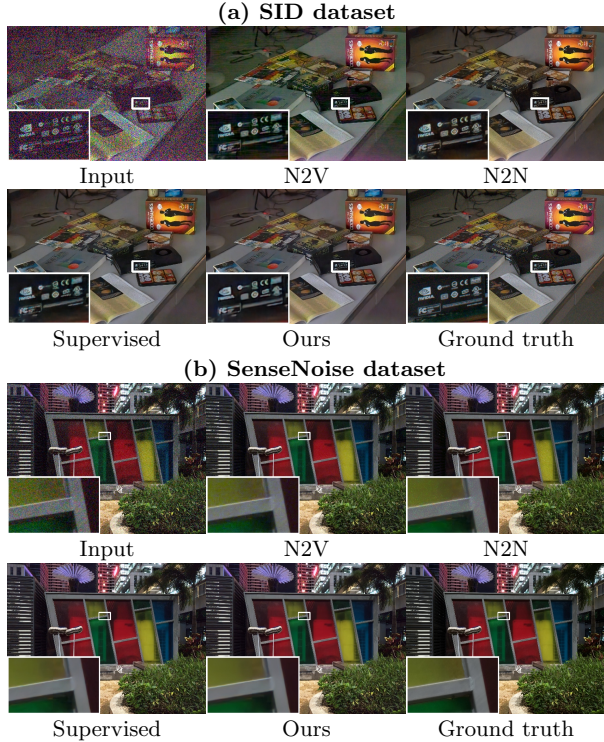


Figure 5. Qualitative results of raw image denoising on SID and *SenseNoise-500* datasets. (Zoom in for details)

Table 4. Comparing our method with supervised learning on raw image denoising. The supervised model is trained on the unprocessed [6] external clean sRGB data (MIT5K [9] and SID sRGB).

Dataset	ISO	MIT5K sRGB	SID sRGB	Ours
SID	100	37.67 / 0.884	37.76 / 0.887	38.02 / 0.906
	250	33.14 / 0.767	33.24 / 0.777	33.88 / 0.819
	300	31.36 / 0.727	31.50 / 0.741	32.14 / 0.779

2 dB performance drop compared with our method. Since the real-world noises are more complicated, they may not satisfy the zero-mean assumption in N2N [18]. As a result, N2N with paired noisy images shows 0.4 dB performance drop compared with our method.

There exists an obvious performance gap between unsupervised and supervised learning methods, especially at higher noise levels. We observe that the gap derives from two aspects. First, in extremely ill-posed conditions, some information might never be recovered without seeing the clean targets. The higher the ISO, the more obvious this phenomenon is. The second aspect might be the inaccurate noise profile. The noise distributions might be slightly different between the DSLRs of the same camera model (*e.g.* fix-pattern noise).

6. Comparison with supervised learning from external clean sRGB images.

Once knowing the noise model, we can synthesize data for supervised training. So, an interesting comparison of our method would be the supervised training with the synthetic data, especially for the real-world raw image denoising. But, the large-scale clean raw image dataset for each sensor is hard and expensive to collect. Following the common practice, we unprocess the external clean sRGB image to obtain clean raw images [6].

We conduct the study on the SID raw dataset. For the supervised method, we collect the clean sRGB images from MIT5K dataset [9] and unprocess them to obtain raw images [6] with proper parameters. For a fair comparison, we keep the same amount of data and the noise model for both methods. The results are shown in Tab. 4. Our methods could still outperform supervised learning with external clean ground-truths due to the dataset content bias. Then, we remove the image content bias by unprocessing the SID sRGB images for the supervised method, its performance improves over the training on MIT5K, but the performance gap with our method still exists as shown in Tab. 4. While the supervised method (“clean image+simulated noise”) uses the same noise model as ours, our method outperforms it over 0.5dB (averaged PSNR 34.17dB vs. 34.68dB) on SID raw dataset. The reason is that there still exists a clear gap (image content, raw image characters) between the synthetic and real clean raw images (unprocessed MIT5K/SID sRGB vs. SID raw).

7. Limitations and conclusions

In this paper, we propose a practical and iterative denoising method (IDR) that outperforms the existing unsupervised denoising methods. Our method is based on empirical findings on data bias of learning denoising. While lacking theoretical support, experiments on three synthetic noises, real-world raw image noise, and even spatially-correlated noise demonstrate the consistent robustness and effectiveness of our method. Moreover, we build a high-quality low-light denoising dataset for evaluating the denoising performance in real-world scenarios. The dataset contains 500 diverse scenes, each with high-quality ground truth. We expect the dataset can serve as a real-world benchmark for image denoising.

8. Acknowledgement

This work is supported in part by Centre for Perceptual and Interactive Intelligence Limited, in part by the General Research Fund through the Research Grants Council of Hong Kong under Grants (Nos. 14204021, 14207319, 14203118, 14208619), in part by Research Impact Fund Grant No. R5001-18, in part by CUHK Strategic Fund.

References

- [1] Abdelrahman Abdelhamed, Mahmoud Afifi, Radu Timofte, Michael S. Brown, et al. NTIRE 2020 challenge on real image denoising: Dataset, methods and results. In *CVPR Workshops*, pages 2077–2088. IEEE, 2020. [2](#)
- [2] Abdelrahman Abdelhamed, Marcus A Brubaker, and Michael S Brown. Noise Flow: Noise Modeling with Conditional Normalizing Flows. In *International Conference on Computer Vision (ICCV)*, 2019. [2](#)
- [3] Abdelrahman Abdelhamed, Stephen Lin, and Michael S. Brown. A high-quality denoising dataset for smartphone cameras. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [2](#), [3](#), [5](#), [6](#)
- [4] Mahmoud Afifi, Abdelrahman Abdelhamed, Abdullah Abuolaim, Abhijith Punnappurath, and Michael S Brown. Cie xyz net: Unprocessing images for low-level computer vision tasks. *arXiv preprint*, 2020. [2](#)
- [5] Joshua Batson and Loïc Royer. Noise2self: Blind denoising by self-supervision. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 524–533. PMLR, 2019. [1](#), [3](#), [7](#)
- [6] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T Barron. Unprocessing images for learned raw denoising. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [1](#), [2](#), [8](#)
- [7] Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. A non-local algorithm for image denoising. In *CVPR (2)*, pages 60–65. IEEE Computer Society, 2005. [2](#)
- [8] Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. Non-local means denoising. *Image Process. Line*, 1, 2011. [2](#)
- [9] Vladimir Bychkovsky, Sylvain Paris, Eric Chan, and Frédo Durand. Learning photographic global tonal adjustment with a database of input / output image pairs. In *The Twenty-Fourth IEEE Conference on Computer Vision and Pattern Recognition*, 2011. [8](#)
- [10] Ke-Chi Chang, Ren Wang, Hung-Jin Lin, Yu-Lun Liu, Chia-Ping Chen, Yu-Lin Chang, and Hwann-Tzong Chen. Learning camera-aware noise models. In *European Conference on Computer Vision (ECCV)*, 2020. [2](#)
- [11] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [12] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen O. Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Trans. Image Processing*, 16(8):2080–2095, 2007. [2](#)
- [13] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward convolutional blind denoising of real photographs. *2019 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [1](#), [2](#)
- [14] Tao Huang, Songjiang Li, Xu Jia, Huchuan Lu, and Jianzhuang Liu. Neighbor2neighbor: Self-supervised denoising from single noisy images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14781–14790, June 2021. [1](#)
- [15] ADOBE SYSTEMS INCORPORATED. Digital negative (dng) specification. https://www.adobe.com/content/dam/acom/en/products/photoshop/pdfs/dng_spec_1.4.0.0.pdf, 2012. [2](#), [6](#)
- [16] Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void-learning denoising from single noisy images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2129–2137, 2019. [1](#), [3](#), [7](#)
- [17] Samuli Laine, Tero Karras, Jaakko Lehtinen, and Timo Aila. High-quality self-supervised deep image denoising. In *NeurIPS*, pages 6968–6978, 2019. [3](#)
- [18] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2noise: Learning image restoration without clean data. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*, pages 2971–2980. PMLR, 2018. [1](#), [3](#), [6](#), [7](#), [8](#)
- [19] Google LLC. Android camera2 api. <http://developer.android.com/reference/android/hardware/camera2/package-summary.html>, 2016. [6](#)
- [20] Julien Mairal, Francis R. Bach, Jean Ponce, Guillermo Sapiro, and Andrew Zisserman. Non-local sparse models for image restoration. In *ICCV*, pages 2272–2279. IEEE Computer Society, 2009. [2](#)
- [21] Ymir Mäkinen, Lucio Azzari, and Alessandro Foi. Exact transform-domain noise variance for collaborative filtering of stationary correlated noise. In *ICIP*, pages 185–189. IEEE, 2019. [2](#), [7](#)
- [22] Ben Mildenhall, Jonathan T Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. Burst denoising with kernel prediction networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [1](#), [2](#)
- [23] Nick Moran, Dan Schmidt, Yu Zhong, and Patrick Coady. Noisier2noise: Learning to denoise from unpaired noisy data. In *CVPR*, pages 12061–12069. IEEE, 2020. [1](#), [3](#), [4](#), [7](#)
- [24] Tobias Plotz and Stefan Roth. Benchmarking denoising algorithms with real photographs. In *CVPR*, pages 2750–2759. IEEE Computer Society, 2017. [3](#), [6](#)
- [25] Javier Portilla, Vasily Strela, Martin J. Wainwright, and Eero P. Simoncelli. Image denoising using scale mixtures of gaussians in the wavelet domain. *IEEE Trans. Image Processing*, 12(11):1338–1351, 2003. [2](#)
- [26] Yuhui Quan, Mingqin Chen, Tongyao Pang, and Hui Ji. Self2self with dropout: Learning self-supervised denoising from single image. In *CVPR*, pages 1887–1895. IEEE, 2020. [2](#), [3](#)
- [27] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992. [2](#)
- [28] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. Deep image prior. In *CVPR*, pages 9446–9454. IEEE Computer Society, 2018. [1](#), [3](#)
- [29] Thijs Vogels, Fabrice Rousselle, Brian McWilliams, Gerhard Röhlin, Alex Harvill, David Adler, Mark Meyer, and Jan Novák. Denoising with kernel prediction and asymmetric loss functions. *ACM Trans. Graph.*, 37(4):124:1–124:15, 2018. [1](#)
- [30] Kaixuan Wei, Ying Fu, Jiaolong Yang, and Hua Huang. A physics-based noise formation model for extreme low-light raw denoising. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. [2](#), [6](#), [7](#)

- [31] Xiaohe Wu, Ming Liu, Yue Cao, Dongwei Ren, and Wangmeng Zuo. Unpaired learning of deep image denoising. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 3, 7
- [32] Jun Xu, Yuan Huang, Li Liu, Fan Zhu, Xingsong Hou, and Ling Shao. Noisy-as-clean: Learning unsupervised denoising from the corrupted image. *CoRR*, abs/1906.06878, 2019. 2, 3, 4
- [33] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Cycleisp: Real image restoration via improved data synthesis. In *CVPR*, 2020. 2
- [34] Syed Waqas Zamir et al. Multi-stage progressive image restoration. In *CVPR*, 2021. 7
- [35] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017. 2, 3
- [36] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for CNN based image denoising. *IEEE Transactions on Image Processing*, 2018. 2
- [37] Yi Zhang, Hongwei Qin, Xiaogang Wang, and Hongsheng Li. Rethinking noise synthesis and modeling in raw denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4593–4601, October 2021. 2, 7