

# Noise2NoiseFlow: Realistic Camera Noise Modeling without Clean Images

Ali Maleky<sup>1,3,\*</sup>, Shayan Kousha<sup>1,3,\*</sup>, Michael S. Brown<sup>3</sup>, Marcus A. Brubaker<sup>1,2,3</sup>  
<sup>1</sup>York University <sup>2</sup>Vector Institute <sup>3</sup>Samsung AI Center–Toronto

## Abstract

Image noise modeling is a long-standing problem with many applications in computer vision. Early attempts that propose simple models, such as signal-independent additive white Gaussian noise or the heteroscedastic Gaussian noise model (a.k.a., camera noise level function) are not sufficient to learn the complex behavior of the camera sensor noise. Recently, more complex learning-based models have been proposed that yield better results in noise synthesis and downstream tasks, such as denoising. However, their dependence on supervised data (i.e., paired clean images) is a limiting factor given the challenges in producing ground-truth images. This paper proposes a framework for training a noise model and a denoiser simultaneously while relying only on pairs of noisy images rather than noisy/clean paired image data. We apply this framework to the training of the Noise Flow architecture. The noise synthesis and density estimation results show that our framework outperforms previous signal-processing-based noise models and is on par with its supervised counterpart. The trained denoiser is also shown to significantly improve upon both supervised and weakly supervised baseline denoising approaches. The results indicate that the joint training of a denoiser and a noise model yields significant improvements in the denoiser.

## 1. Introduction

Image noise modeling is a long-standing problem in computer vision that has relevance for many applications [6, 9, 10, 21, 22, 29]. Recently, data-driven noise models based on deep learning have been proposed [2, 12, 23]. Unfortunately, these models generally require clean (i.e., noise-free) images, which are practically challenging to collect in real scenarios [1]. In this work we propose a new approach, Noise2NoiseFlow, which can accurately learn noise models without the need for clean images. Instead, only pairs of noisy images of a fixed scene are required.

While efforts are made to reduce noise during capture, post-capture modeling is a critical piece of many

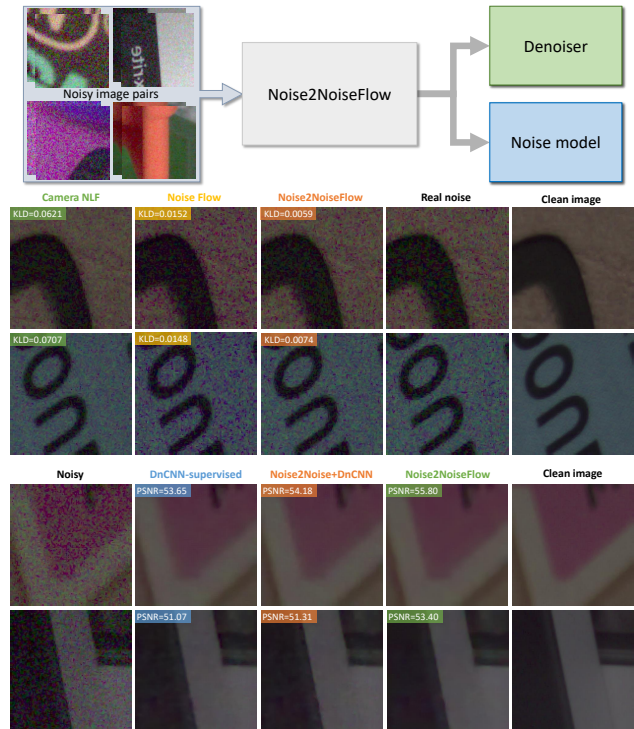


Figure 1. Overview of Noise2NoiseFlow. (top) Given pairs of noisy images of the same scene, Noise2NoiseFlow simultaneously trains both a noise model and a denoiser. (middle) Noise synthesis results from Camera NLF, Noise Flow, and Noise2NoiseFlow compared to the real noise in the SIDD dataset. Noise generated by Noise2NoiseFlow is the most similar to the real noise both visually and in KL divergence but without requiring clean, noise-free images. (bottom) Example denoising results from the jointly trained denoiser compared to its supervised DnCNN baseline, and a DnCNN trained with Noise2Noise loss.

downstream tasks and in many domains large amounts of noise are intrinsic to the problem—for example, astrophotography and medical imaging. As a result, noise is an integral and significant part of signal capture in many imaging domains, and modeling it accurately is critical. For instance, noise model estimation is necessary for removing fixed pattern effects from CMOS sensors [11] and enhancing video in extreme low-light conditions [32]. Noise

\*Work performed while interns at the Samsung AI Center–Toronto.

models can also be used to train downstream tasks to be robust in the presence of realistic input noise. Most naturally, they can also be used to train noise reduction algorithms without the need to collect pairs of clean and noisy images [2, 26, 38]. However, as mentioned in [3, 31, 39] denoisers trained with unrealistic noise models—for example, simple Gaussian noise—may not perform well on real data.

Early attempts at noise modeling were limited and failed to fully capture the characteristics of real noise. Simple IID Gaussian noise (also called a homoscedastic Gaussian noise) ignores the fact that photon noise is signal-dependent. Heteroscedastic Gaussian noise (*e.g.*, [9]) captures this by modeling noise variance as a linear function of clean image intensity but does not take into account the spatial non-uniformity of noise power, amplification noise, quantization effects, and more. More recently, Noise Flow [2] was proposed as a new parametric structure that uses conditional normalizing flows to model noise in the camera imaging pipeline. This model is a combination of unconditional and conditional transformations that map simple Gaussian noise into a more complex, signal-, camera-, and ISO-dependent noise distribution and outperformed previous baselines by a large margin in the normalizing flows [17] framework. However, it required supervised noise data—namely, pairs of clean and noisy images—in order to learn the noise model. Unfortunately gathering supervised data consisting of corresponding clean and noisy images can be challenging [1, 3, 28, 34] and is a limiting factor in the realistic characterization of noise. This is even worse for other downstream tasks, which typically require large amounts of data for training.

In the context of image denoising specifically, there has been significant recent interest in methods that avoid the need for supervised data, either from careful collection or synthesis. The well-known BM3D method [8] proposed a denoising scheme based on transform domain representation without clean image correspondence. However, the similar patch search step makes the inference time complexity inefficient for large-scale datasets. Recently, Lehtinen et al. [20] introduced the Noise2Noise framework, which allowed for training of a denoiser given pairs of noisy images of the same underlying image signal. Following this work, several others were proposed aiming to further reduce the data requirements; in particular Noise2Void [18] and Noise2Self [4] allow training of a denoiser with only individual noisy images by forcing the denoiser to predict the intensity of each pixel using only its neighbours. Other methods attempted to add additional noise to noisy input images [25, 27, 35] or use unpaired images in a GAN framework [5, 7, 13–15]. However, in all cases these methods are aimed primarily at denoising instead of noise modeling.

In this work, we aim to leverage these recent advances in training denoisers without direct supervision in the context

of noise modeling. Specifically, we extend the Noise2Noise framework to train a noise model with pairs of independently sampled noisy images rather than clean data. The resulting approach, called Noise2NoiseFlow and illustrated in Figure 1, produces both a denoiser and an explicit noise model, both of which are competitive with or out-perform fully supervised training of either model individually.

## 2. Background

Image noise can be described as an undesirable corruption added to an underlying clean signal. Formally,

$$\tilde{\mathbf{I}} = \mathbf{I} + \mathbf{N}, \quad (1)$$

where  $\mathbf{I}$  is the underlying and mostly unobserved clean image and  $\mathbf{N}$  is the unwanted noise corrupting the signal, and their addition results in the noisy observation  $\tilde{\mathbf{I}}$ . Different noise models are then defined by the choice of distribution assumed for  $\mathbf{N}$ . A widely used noise model assumes that  $\mathbf{N}(x, y) \sim \mathcal{N}(0, \sigma^2)$ —namely, that the noise at each pixel is drawn from a zero-mean Gaussian distribution with some fixed variance. This model has commonly been used to train and test denoisers; however, it fails to capture significant aspects of real noise, most prominently the signal-dependent variance, which is a result of the inherent Poisson shot noise [22, 24]. A significant improvement over this is heteroscedastic Gaussian noise (HGN) [9, 10, 22] which assumes that the variance of the noise at each pixel is a linear function of the clean image intensity. That is  $\mathbf{N}(x, y) \sim \mathcal{N}(0, \sigma^2(\mathbf{I}(x, y)))$ , where  $\sigma^2(\mathbf{I}) = \beta_1 \mathbf{I} + \beta_2$  and  $\beta_1, \beta_2$  are parameters. This model is also sometimes referred to as the “noise level function” (NLF). Recent work has shown that NLF parameters from camera manufacturers are often poorly calibrated [38]; however, the NLF neglects important noise characteristics, including spatial correlation, defective pixels, clipping, quantization, and more.

To address the limitations of these pixel-independent, Gaussian-based noise models, Abdelhamed et al. [2] proposed the Noise Flow model, a parametric noise model based on conditional normalizing flows specifically designed to capture different noise components in a camera imaging pipeline. In particular, Noise Flow can be seen as a strict generalization of HGN due to its use of a signal-dependent transformation layer. However, unlike HGN, Noise Flow is capable of capturing non-Gaussian distributions and complex spatial correlations.

More recently, the DeFlow model [33] was proposed to handle a broader range of image degradations beyond traditional noise. Other approaches consider mixture models or Generative Adversarial Networks (GAN) to simulate noisy and clean images in the context of denoiser training [5, 7, 12–15, 40]. However, these models are typically focused on denoising as opposed to noise modeling. Further,

GANs do not have tractable likelihoods, making the quality of the synthesized noise difficult to assess. Most importantly, the above methods require clean images, and potentially pairs of noisy and corresponding clean images for training. In this work we construct a formulation that explicitly trains a noise model without the need for clean images. Because of the flexibility and generality of the normalizing flow framework and quality of its results, we will focus on the Noise Flow model [2] here, though, as we will discuss, other choices are possible.

## 2.1. Image Denoising

Image noise reduction has been a long-standing topic of study in computer vision [6, 8, 19, 21, 29, 37]. Here we focus on recent methods that have found success by leveraging large training sets and deep learning architectures [36]. These methods are characterized by regressing, typically with a convolutional neural network, from a noisy image observation to its clean counterpart. Given a training set  $\mathcal{D} = \{(\tilde{\mathbf{I}}^{(i)}, \mathbf{I}^{(i)})\}_{i=1}^N$  of noisy images  $\tilde{\mathbf{I}}$  and their corresponding clean images  $\mathbf{I}$ , learning of a denoiser  $\mathbf{D}$  is then formulated as minimizing

$$\sum_{i=1}^N \mathcal{L}(\mathbf{D}(\tilde{\mathbf{I}}^{(i)}; \theta), \mathbf{I}^{(i)}), \quad (2)$$

where  $\mathcal{L}$  is typically an  $L_1$  or  $L_2$  norm and  $\mathbf{D}$  is a deep neural network with parameters  $\theta$ .

This approach is limited by the need to have access to the corresponding clean image  $\mathbf{I}$ , and several notable approaches have recently been explored to remove this requirement. Most relevant to this work is the Noise2Noise framework, proposed by Lehtinen et al. [20]. Rather than requiring clean/noisy pairs of images, it simply requires two noisy observations of the same underlying clean signal. Given a dataset of noisy image pairs  $\{(\tilde{\mathbf{I}}_1^{(i)}, \tilde{\mathbf{I}}_2^{(i)})\}_{i=1}^N$ , the Noise2Noise framework optimizes the loss function

$$\sum_{i=1}^N \mathcal{L}(\mathbf{D}(\tilde{\mathbf{I}}_1^{(i)}; \theta), \tilde{\mathbf{I}}_2^{(i)}) + \mathcal{L}(\mathbf{D}(\tilde{\mathbf{I}}_2^{(i)}; \theta), \tilde{\mathbf{I}}_1^{(i)}). \quad (3)$$

That is, the the second noisy image is used as the target for the denoiser of the first and vice versa. Perhaps surprisingly, training with this objective is still able to produce high-quality denoising results, despite the lack of access to clean images [20]. In this work, we aim to explore the generalization of this approach to noise modeling.

## 3. Noise2NoiseFlow

In this section, we define our approach to learning a noise model with weak supervision—namely, through the use of only pairs of noisy images. There are two main components, a denoiser  $\mathbf{D}(\cdot; \theta)$ , which learns to predict the clean

image  $\mathbf{I}$  given a noisy image,  $\tilde{\mathbf{I}}$ , as input, and a model of a noisy image given the clean image  $\mathbf{I}$ ,  $p_{\tilde{\mathbf{I}}}(\cdot|\mathbf{I}; \phi)$ . The denoiser and noise model have parameters  $\theta$  and  $\phi$  respectively. Our goal is to learn the distribution  $p_{\tilde{\mathbf{I}}}(\tilde{\mathbf{I}}|\mathbf{I})$ —namely, the distribution of noisy image conditioned on the clean image—without explicitly requiring  $\mathbf{I}$ .<sup>1</sup> To do this, we propose to use the output of the denoiser as an estimate of the clean image—That is,  $\mathbf{I} \approx \hat{\mathbf{I}} = \mathbf{D}(\tilde{\mathbf{I}}; \theta)$ . We could in principle then learn  $p_{\tilde{\mathbf{I}}}$  by minimizing  $-\log p_{\tilde{\mathbf{I}}}(\hat{\mathbf{I}}|\hat{\mathbf{I}}; \phi)$  with respect to the noise model parameters  $\phi$ . However, this requires a well-trained denoiser, which, in turn, typically requires access to clean images to train. Further, if we tried to simultaneously train the denoiser and noise model, there is a trivial singular optimum where the denoiser converges to the identity and the noise model converges to a Dirac delta at zero.

Drawing inspiration from the Noise2Noise framework [20], we instead assume we have access to pairs of noisy observations  $\tilde{\mathbf{I}}_1, \tilde{\mathbf{I}}_2$  which both have the same underlying clean signal,  $\mathbf{I}$ . That is,  $\tilde{\mathbf{I}}_1 = \mathbf{I} + \mathbf{N}_1$  and  $\tilde{\mathbf{I}}_2^{(i)} = \mathbf{I} + \mathbf{N}_2$ , where  $\mathbf{N}_1$  and  $\mathbf{N}_2$  are independent samples of noise. Then, given the pairs of noisy images, we can use the denoiser applied to one image to estimate the clean image for the other image in the pair. That is, we propose to optimize the loss

$$\mathcal{L}_{nm}(\tilde{\mathbf{I}}_1, \tilde{\mathbf{I}}_2) = -\log p_{\tilde{\mathbf{I}}_1}(\tilde{\mathbf{I}}_1|\mathbf{D}(\tilde{\mathbf{I}}_2; \theta); \phi) - \log p_{\tilde{\mathbf{I}}_2}(\tilde{\mathbf{I}}_2|\mathbf{D}(\tilde{\mathbf{I}}_1; \theta); \phi) \quad (4)$$

for both the noise model parameters  $\phi$  and the denoiser parameters  $\theta$ . Because the two images are of the same underlying scene, the output of the denoiser should ideally be the same for both noisy images. However, because the two images have independent samples of noise, the denoiser cannot simply collapse to the identity. This is analogous to the Noise2Noise objective, where the output of the denoiser on one image is used as the target for the other image in the pair. In practice, we find it beneficial to include the Noise2Noise objective function to stabilize the training of the denoiser together with the noise model objective. That is, we propose to train the denoiser and noise model jointly with the loss  $\mathcal{L}(\tilde{\mathbf{I}}_1, \tilde{\mathbf{I}}_2) = \mathcal{L}_{nm}(\tilde{\mathbf{I}}_1, \tilde{\mathbf{I}}_2) + \lambda \mathcal{L}_{dn}(\tilde{\mathbf{I}}_1, \tilde{\mathbf{I}}_2)$ , where

$$\mathcal{L}_{dn}(\tilde{\mathbf{I}}_1, \tilde{\mathbf{I}}_2) = \|\mathbf{D}(\tilde{\mathbf{I}}_1; \theta) - \tilde{\mathbf{I}}_2\|_2^2 + \|\mathbf{D}(\tilde{\mathbf{I}}_2; \theta) - \tilde{\mathbf{I}}_1\|_2^2 \quad (5)$$

is the Noise2Noise loss. Given a dataset of pairs of noisy images,  $\mathcal{D} = \{(\tilde{\mathbf{I}}_1^{(i)}, \tilde{\mathbf{I}}_2^{(i)})\}_{i=1}^N$ , we optimize the loss over the set of pairs

$$\sum_{i=1}^N \mathcal{L}(\tilde{\mathbf{I}}_1^{(i)}, \tilde{\mathbf{I}}_2^{(i)}),$$

where the optimization can be done with a stochastic optimizer. In this work we use Adam [16].

<sup>1</sup>Note that this is equivalent to learning the distribution of the noise conditioned on the clean image by simply shifting the distribution of noise by the clean image.

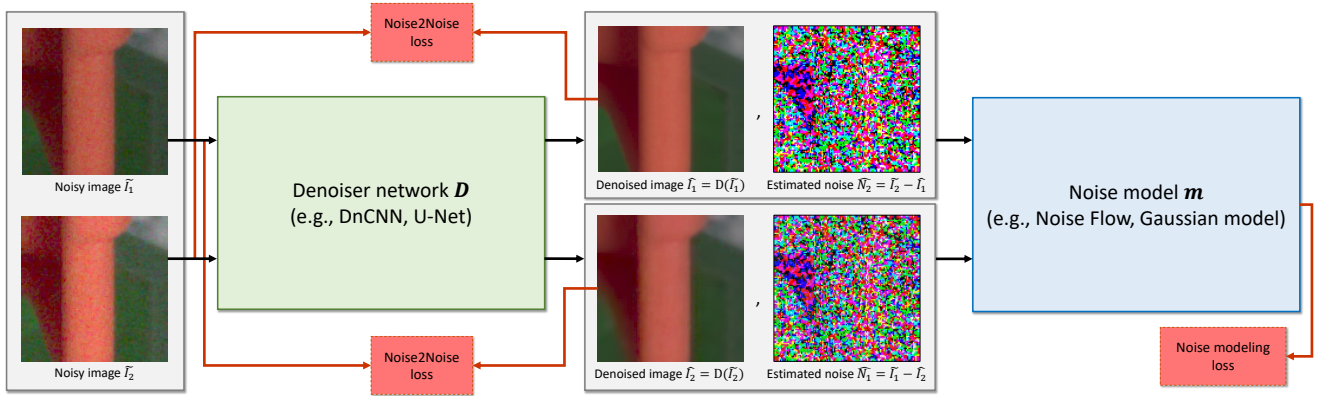


Figure 2. An overview of the training loss for the proposed Noise2NoiseFlow framework. Given a pair of independent noisy samples of the same underlying signal, it runs both noisy samples through a denoiser network  $\mathbf{D}$ , which outputs the estimated clean signal. We then use the estimated clean image from the first image in place of the true clean signal for the second noisy observation and vice versa. This prevents the denoiser from collapsing into a degenerate solution of an identity transformation. Note that in the paper we formulate the noise model as a distribution over the noisy image  $\tilde{\mathbf{I}}$ , though, as shown here, it is common for noise models to be expressed as a distribution of the residual noise  $\mathbf{N} = \tilde{\mathbf{I}} - \mathbf{I}$ . These two formulations are equivalent.

Figure 2 shows an overview of the proposed approach. We note that the formulation is generic to the choice of denoiser and noise model, requiring only that the noise model’s density function can be evaluated and that both the noise model and denoiser can be differentiated as needed. In the experiments that follow we primarily use the DnCNN architecture [36] for the denoiser, as it is a standard denoiser architecture based on residual connections and convolutional layers. For the noise model we primarily focus on Noise Flow [2] due to its flexibility and tractability and, consequently, dub our proposed method **Noise2NoiseFlow**. However, we also explore other choices for these components, such as a U-Net architecture for the denoiser and the heteroscedastic Gaussian noise model.

## 4. Experiments

Here we explore the performance of the proposed Noise2NoiseFlow approach. To do this we make use of Smartphone Image Denoising Dataset (SIDD) [1] to assess the accuracy of both our learned noise model and the image denoiser. SIDD contains images of 10 different scenes consisting of a range of objects and lighting conditions, which were captured with five different smartphone cameras at a range of different ISO levels. Multiple captures of each scene instance were taken and carefully aligned in order to produce a corresponding “clean” image for each noisy image. While our proposed method does not require the clean images for training, we do make use of them for a quantitative evaluation against a range of baselines, including methods that require clean image supervision. Here we use two different subsets of SIDD—namely SIDD-Full and SIDD-Medium. While SIDD provides both sRGB and rawRGB

images, here we only consider the rawRGB images. SIDD-Full provides 150 different noisy captures for each corresponding clean image. In contrast, SIDD-Medium contains only a single noisy image for each clean image. To extract the noisy/noisy image pairs of the same clean signal from SIDD-Full that are required by our method for training, we select pairs of noisy images corresponding to the same clean image. In order to maximize alignment between the selected two images, we select consecutive images from the 150 available for each scene in SIDD-Full.

We use SIDD-Medium to evaluate the performance of our method. Specifically, while we use noisy/noisy pairs of images extracted from SIDD-Full for training as described above, we evaluate the performance of both the denoiser  $\mathbf{D}(\cdot)$  and the noise model  $p_{\tilde{\mathbf{I}}}(\cdot|\mathbf{I})$  using the noisy/clean image pairs in SIDD-Medium. In order to test Noise2NoiseFlow against our baselines, we use supervised noisy/clean pairs from SIDD-Medium. Denoting  $(\tilde{\mathbf{I}}, \mathbf{I})$  as a noisy/clean image pair, we evaluate the noise modeling using the negative log-likelihood per dimension  $D^{-1} \log p_{\tilde{\mathbf{I}}}(\tilde{\mathbf{I}}|\mathbf{I}; \phi)$ , where  $D$  is the total number of dimensions (both pixels and channels) in the input. Negative log likelihood is a standard evaluation metric for generative models and density estimation, but it is known to be less sensitive to distributions that overestimate the variance of a distribution. To account for this we also evaluate the model using the Kullback-Leibler (KL) divergence metric introduced in [2]. Both NLL and KL divergence are reported in nats. Specifically, given a noisy and clean image, we compute a histogram of both real noise and noise generated by a model by subtracting the clean image and computing the KL divergence between the two histograms. See [2] for



more details on this metric. To evaluate the denoiser, we compute peak signal-to-noise ratio (PSNR) and the structural similarity index measure (SSIM).

SIDD contains scenes with ISO levels ranging from 50 to 10,000; however, many of those ISO levels have only a small number of images available. To be consistent with other methods that use SIDD for noise modeling—for example, [2]—we remove images with rare ISO levels, keeping only ISO levels 100, 400, 800, 1600, and 3200. After filtering, approximately 500,000 patches of size  $32 \times 32$  pixels are extracted. The extracted patches are separated into training and test sets using the same training and testing split of SIDD scenes that was used in [2]. Approximately 70% of the extracted patches were used for training and the remaining were used as testing. We trained all models using the Adam optimizer [16] for 2,000 epochs. We used a value of  $\lambda = 2^{18}$  in all experiments, unless otherwise noted. To speed up convergence and avoid early training instabilities we pre-trained the denoiser  $\mathbf{D}$  on the training set using  $\mathcal{L}_{dn}$  alone for all of the experiments. The architecture of the Noise Flow noise model and DnCNN denoiser was the same as in [2], but both were reimplemented in PyTorch and verified to produce equivalent results as the original Noise Flow implementation.

#### 4.1. Noise Modeling

We first compare our proposed approach quantitatively to traditional noise models which have been calibrated using supervised, clean images. Table 1 compares the results of our model against the camera noise level function (Cam-NLF), a simple additive white Gaussian noise model (AWGN), and Noise Flow [2]. Despite only having access to pairs of noisy images, the proposed Noise2NoiseFlow has effectively identical performance to the state-of-the-art Noise Flow model which is trained on clean/noisy image pairs. To demonstrate the benefit of joint training, we trained a Noise2Noise denoiser [20] on noisy/noisy paired data and use this to denoise images to train Noise Flow. We refer to this as “N2N+NF.”

We also compared our results to the recently released “calibrated Poisson-Gaussian” noise model described in [38]. The results for this comparison in terms of KL divergence can be found in Table 2 for the three cameras reported in the paper [38], as the Calibrated P-G model included noise parameters only for three different sensors: iPhone 7, Samsung Galaxy S6 Edge, and Google Pixel. It is clear that while the Calibrated P-G model improves over the in-camera noise level function, it still lags behind both Noise Flow and Noise2NoiseFlow. We again see that the proposed Noise2NoiseFlow outperforms this very recent method.

Figure 3 shows qualitative noise samples generated by Noise2NoiseFlow, as well as other baselines compared to the real noise. The samples are generated for different cam-

Model	$NLL$	$D_{KL}$
AWGN	-2.874	0.4815
Cam. NLF	-3.282	0.0578
N2N+NF	-3.459	0.0363
Noise Flow	-3.502	0.0267
Noise2NoiseFlow	-3.501	0.0265

Table 1. Negative log-likelihood per dimension and  $D_{KL}$  results on test data for baseline models and our proposed Noise2NoiseFlow model. Noise2NoiseFlow significantly improves over AWGN and Camera NLF and is on par with the Noise Flow model, while requiring no clean images. It also improves over separately training a Noise2Noise denoiser and NoiseFlow (N2N+NF), demonstrating the value of joint training.

	IP	S6	GP	Agg
AWGN	0.4353	0.4863	0.5865	0.4934
Cam. NLF	0.0513	0.1014	0.0212	0.0596
Calibrated P-G	0.0188	0.0981	0.0332	0.0492
Noise Flow	0.0112	0.0469	0.0180	0.0250
Noise2NoiseFlow	0.0125	0.0444	0.0190	0.0249

Table 2. Per camera KL divergence performance of our model Noise2NoiseFlow compared to the baselines on three camera sensors for which the Calibrated P-G model is defined as well as the aggregate results on the test data for these three sensors.

era sensors, ISO levels, and scenes. The suffix N corresponds to normal light and L corresponds to the low-light conditions. As evidenced by these images, the results from Noise2NoiseFlow are both visually and quantitatively better than other baselines, especially in low-light/high-ISO settings, where other baselines underperform.

#### 4.2. Noise Reduction

While the primary goal of this work was noise modeling, it also includes a denoiser as a key component. Here we investigate the performance of the denoiser by evaluating its performance in terms of PSNR on the held-out test set. We compared against three scenarios, which are reported in Table 3. In all cases the exact same DnCNN architecture is used. First, we trained the same denoiser architecture  $\mathbf{D}$  using the Noise2Noise [20] loss alone. This is shown in Table 4 as “Noise2Noise+DnCNN” and shows that, indeed, the joint noise model training improves the denoising performance by over 1.2dB, a significant margin in PSNR. Second, we trained a supervised DnCNN model using the corresponding clean image patches for the training set; this is indicated in the table as “DnCNN-supervised”. Noise2NoiseFlow outperforms this by nearly 1.5dB, despite not having access to clean images. In fact, both Noise2Noise+DnCNN and Noise2NoiseFlow outperform this clean-image supervised baseline, suggesting that




Figure 3. Noise synthesis samples from (a) the AWGN model, (b) Camera NLF, (c) Calibrated P-G [38], (d) Noise Flow [2], and our proposed method, Noise2NoiseFlow, compared to the (f) real noise in SIDD. Our samples are closer to real noise in terms of both visual perception and the KL divergence metric. Codes on the left indicate [camera]-[ISO]-[brightness]. For Calibrated P-G [38] calibrated parameters for SIDD only include three camera sensors (IP, GP, and S6) so some synthesis results are not available.

the increased variety of data available with noisy image pairs appears to be more valuable than access to clean images. We also trained a supervised Noise Flow model and used samples generated from the model to train a DnCNN denoiser. We refer to this baseline as “DnCNN - NF synthesized”. The “DnCNN - NF synthesized” outperforms the “DnCNN-supervised” baseline which is consistent with the results reported in the Noise Flow paper [2]. However, it still significantly underperforms Noise2NoiseFlow.

Figure 4 shows qualitative denoising results from Noise2NoiseFlow and the aforementioned baselines. The results show that our model performs better in denoising, especially in more severe situations (high ISO and low brightness). The estimated clean signal tends to be much

smoother and cleaner for Noise2NoiseFlow than both of its baselines in terms of visual perception and PSNR in almost all the cases. Taken together, our results demonstrate that the joint training of both an explicit noise model and a denoiser not only allows for weakly supervised training, but also improves the resulting estimated denoiser.

### 4.3. Ablation Studies

We next investigate the design choices for our framework and their impact on the results. First, we conduct an ablation on the value of  $\lambda$ , the weighting factor for the Noise2Noise loss. We explored a wide range of values, from  $\lambda = 0$  to  $\lambda = 2^{18}$ . For each value, we computed the negative log-likelihood per dimension and the PSNR of the denoiser.






Figure 4. Denoising results from (b) DnCNN-supervised, (c) Noise2Noise trained with DnCNN, and (d) Noise2NoiseFlow on samples from SIDD-Medium testing data. The codes on the left indicate the ISO level as well as the lighting condition.

Model	PSNR	SSIM
Noise2Noise+DnCNN	51.57	0.977
DnCNN-supervised	51.32	0.980
DnCNN - NF synthesized	51.71	0.980
Noise2NoiseFlow(ours)	<b>52.80</b>	<b>0.984</b>

Table 3. Denoising results from a DnCNN trained with supervised noisy/clean paired data from SIDD-Medium, Noise2Noise with a DnCNN trained on noisy/noisy image pairs, a DnCNN trained with noise samples generated from a supervised Noise Flow model, and our denoiser model trained on the same noisy/noisy data measured by PSNR and SSIM on the test set.

The results are plotted in Fig. 5 and show that our results are relatively robust to the choice of  $\lambda$ . While a value of  $\lambda = 0$  produces reasonable results, better results are generally obtained with larger values of  $\lambda$ . This indicates that the Noise2Noise loss in Eq. 5 plays an important role in stabilizing the training and ensuring consistency of the denoiser.

Next, we consider a different form of the loss function




Figure 5. Negative log-likelihood per dimension and PSNR results on test data as a function of the regularization term  $\lambda$ . Cross-sample loss refers to training with Eq. 4 and self-sample loss refers to training with Eq. 6. Cross marks indicate the loss at the last epoch in cases where training failed as evidenced by spikes in NLL and significant drops in PSNR. All experiments with the self-sample loss and  $\lambda \leq 2^{16}$  ultimately diverge.

where we use the estimated clean image based on  $\tilde{\mathbf{I}}_1$  for the noise model loss function with  $\tilde{\mathbf{I}}_1$ . Formally, we use the noise model objective

$$\mathcal{L}_{nm}(\tilde{\mathbf{I}}_1, \tilde{\mathbf{I}}_2) = -\log p_{\tilde{\mathbf{I}}_1}(\tilde{\mathbf{I}}_1 | \mathbf{D}(\tilde{\mathbf{I}}_1; \theta); \phi) - \log p_{\tilde{\mathbf{I}}_2}(\tilde{\mathbf{I}}_2 | \mathbf{D}(\tilde{\mathbf{I}}_2; \theta); \phi) \quad (6)$$

instead of the one proposed in Equation 4. We refer to training based on this model as the self-sample loss, in comparison to the cross-sample loss. While a seemingly innocuous change, training based on Equation 6 becomes extremely unstable. In this case, the denoiser can converge to a degenerate solution of the identity function—namely,  $\mathbf{D}(\tilde{\mathbf{I}}) = \tilde{\mathbf{I}}$ —which allows the noise model  $p_{\tilde{\mathbf{I}}}$  to converge to a Dirac delta and the value of  $\mathcal{L}_{nm}(\tilde{\mathbf{I}}_1, \tilde{\mathbf{I}}_2)$  goes to negative infinity. This behaviour can be alleviated with large values of  $\lambda$ , which can be seen in Figure 5, where settings of  $\lambda$  that resulted in diverged training are indicated with a cross at the value of the epoch before the divergence occurred. As the figure shows, values  $\lambda$  less than  $2^{17}$  resulted in this behaviour. In contrast, the proposed loss function in Equation 4 is robust to the choice of  $\lambda$ , even allowing training with a value of  $\lambda = 0$ , which disables the  $\mathcal{L}_{dn}$  term from Equation 5 entirely. We also explored higher values for  $\lambda$  (e.g.,  $2^{19}$ ) but did not observe significant changes in behaviour.

We also explored different choices of denoiser architecture and noise model as our framework is agnostic to these specific choices. For the denoiser, beyond the DnCNN architecture, we also considered the U-Net [30] denoiser architecture used in [20]. For the noise model, beyond the Noise Flow-based model, we also considered the heteroscedastic Gaussian noise model, or noise level function (NLF), due to its ubiquity. We implemented the NLF as a variation on a Noise Flow architecture. Specifically, taking the signal-dependent and gain layers of the Noise Flow model, without any of the other flow layers, results in a model that is equivalent to the NLF.

Denoiser	Noise model	$NLL$	$D_{KL}$	PSNR
DnCNN	Noise Flow	-3.501	0.0265	52.80
U-Net	Noise Flow	-3.500	0.0255	52.64
DnCNN	NLF	-3.461	0.0288	52.69
U-Net	NLF	-3.463	0.0332	52.13

Table 4. Performance on test data from our ablated models. Each row corresponds to a specific choice for the noise model and the denoiser architecture. NLF corresponds to our modified version of heteroscedastic Gaussian model implemented as a bijective normalizing flow transformation.

The results of this experiment can be found in Table 4, which reports the negative log likelihood per dimension, KL divergence metric, and PSNR of the resulting noise model and denoiser for all combinations of these choices. The results indicate that the choice of denoiser architecture is not particularly important. Both U-Net and DnCNN produce similar results to one another, for both choices of noise model. However, we see that the use of the Noise Flow model over the heteroscedastic Gaussian noise model does provide a boost in performance for both noise modeling and denoising. Further, and consistent with results reported recently elsewhere [38], we see that a retrained heteroscedastic Gaussian noise model can outperform the parameters provided by camera manufacturers.

#### 4.4. Training with Individual Noisy Images

Here we have proposed a novel approach to noise model training by coupling the training of a noise model with a denoiser and based on the Noise2Noise framework. This naturally raises the question of whether a noise model could be trained with only individual noisy images, particularly given the success of such approaches for denoisers. All of these approaches aim to prevent the denoiser from collapsing into the degenerate solution of an identity transformation, similar to the behaviour identified above with the alternative loss formulation in Equation 6, by either using a blind-spot network architecture (*e.g.*, Noise2Void [18] and Noise2Self [4]), or adding additional noise to the input images (*e.g.*, Noisier2Noise [25], Noisy-as-Clean [35], and R2R [27]). To investigate this idea we considered using the R2R [27] framework, which, given a single noisy image  $\tilde{\mathbf{I}}$ , generates two new noisy images as

$$\tilde{\mathbf{I}}_{\text{input}} = \tilde{\mathbf{I}} + \mathbf{D}^T \mathbf{z}, \quad \tilde{\mathbf{I}}_{\text{target}} = \tilde{\mathbf{I}} - \mathbf{D}^{-1} \mathbf{z}, \quad (7)$$

where  $\mathbf{z}$  is drawn from  $\mathcal{N}(0, I)$ , and  $\mathbf{D} = \alpha I$  is an invertible matrix with scale parameter  $\alpha$ . We modify our loss functions to utilize these new images so that  $\mathcal{L}_{nm} = -\log p_{\tilde{\mathbf{I}}}(\tilde{\mathbf{I}}|\mathbf{D}(\tilde{\mathbf{I}}_{\text{input}}))$  and  $\mathcal{L}_{dn} = \|\tilde{\mathbf{I}}_{\text{target}} - \mathbf{D}(\tilde{\mathbf{I}}_{\text{input}})\|_2^2$  and train by optimizing  $\mathcal{L} = \mathcal{L}_{nm} + \lambda\mathcal{L}_{dn}$  as described above. We use the same DnCNN architecture for  $\mathbf{D}$  and

Model	$NLL$	$D_{KL}$	PSNR
Noise Flow	-3.502	0.0267	N/A
Noise2NoiseFlow	-3.501	0.0265	52.80
R2RFlow	-3.443	0.0983	50.08

Table 5. Performance on test data from our proposed R2RFlow formulation that requires only single noisy samples and no supervision in any forms compared to Noise2NoiseFlow (weak supervision) and the Noise Flow model (supervised).

the Noise Flow model for  $p_{\tilde{\mathbf{I}}}$  and report the results in Table 5, with this variation labelled as R2RFlow and compared against a clean-image supervised Noise Flow model and the noisy-pair supervised Noise2NoiseFlow. The results indicate that the R2RFlow approach yields a reasonable noise model, though significantly below the performance of Noise2NoiseFlow, particularly in terms of denoising. However, the experiment is enticing and suggests that this is a promising direction for future work.

## 5. Conclusions and Future Work

We introduced a novel framework for jointly training a noise model and denoiser that does not require clean image data. Our experimental results showed that, even without the corresponding clean images, the noise modeling performance is largely the same when training only with pairs of noisy images. We believe this approach can improve the practicality of the existing noise models in real-world scenarios by reducing the need to collect clean image data, which can be a challenging, tedious, and time-consuming process and may not be possible in some settings, *e.g.*, medical imaging. Further, joint training was shown to improve denoising performance when compared with a denoiser trained alone. The learned denoiser can even surpass supervised baselines, which we hypothesize is due to the increased number of noisy images and indicating that noise modeling can provide useful feedback for denoising.

While training a noise model without clean image data is a significant step towards more practical noise models, our proposed approach still required paired noisy images. We believe that it may be possible to go further still and train a noise model in a purely unsupervised manner, *i.e.*, without clean images or pairs of noisy images. Our preliminary experiments with the R2R framework [27] suggest that this may indeed be feasible, but more work remains to be done. Code for this paper is available at: <https://yorkucvil.github.io/Noise2NoiseFlow/>.

**Acknowledgements** This work was done during an internship at the Samsung AI Center in Toronto, Canada. AM’s internship was funded by a Mitacs Accelerate. SK’s and AM’s student funding came in part from the Canada First Research Excellence Fund for the Vision: Science to Applications (VISTA) programme and an NSERC Discovery Grant.



## References

- [1] A. Abdelhamed, S. Lin, and M. S. Brown, “A high-quality denoising dataset for smartphone cameras,” in *CVPR*, 2018. 1, 2, 4
- [2] A. Abdelhamed, M. A. Brubaker, and M. S. Brown, “Noise flow: Noise modeling with conditional normalizing flows,” in *ICCV*, 2019. 1, 2, 3, 4, 5, 6
- [3] J. Anaya and A. Barbu, “Renoir—a dataset for real low-light image noise reduction,” *Journal of Visual Communication and Image Representation*, vol. 51, pp. 144–154, 2018. 2
- [4] J. Batson and L. Royer, “Noise2self: Blind denoising by self-supervision,” in *ICML*, 2019. 2, 8
- [5] S. Cha, T. Park, and T. Moon, “Gan2gan: Generative noise learning for blind image denoising with single noisy images,” *ICLR*, 2021. 2
- [6] S. G. Chang, B. Yu, and M. Vetterli, “Adaptive wavelet thresholding for image denoising and compression,” *TIP*, vol. 9, no. 9, pp. 1532–1546, 2000. 1, 3
- [7] J. Chen, J. Chen, H. Chao, and M. Yang, “Image blind denoising with generative adversarial network based noise modeling,” in *CVPR*, 2018. 2
- [8] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, “Image denoising by sparse 3-d transform-domain collaborative filtering,” *TIP*, vol. 16, no. 8, pp. 2080–2095, 2007. 2, 3
- [9] A. Foi, “Clipped noisy images: Heteroskedastic modeling and practical denoising,” *Signal Processing*, vol. 89, no. 12, pp. 2609–2629, 2009. 1, 2
- [10] A. Foi, M. Trimeche, V. Katkovnik, and K. Egiazarian, “Practical poissonian-gaussian noise modeling and fitting for single-image raw-data,” *TIP*, vol. 17, no. 10, pp. 1737–1754, 2008. 1, 2
- [11] G. E. Healey and R. Kondepudy, “Radiometric ccd camera calibration and noise estimation,” *TPAMI*, vol. 16, no. 3, pp. 267–276, 1994. 1
- [12] B. Henz, E. S. Gastal, and M. M. Oliveira, “Synthesizing camera noise using generative adversarial networks,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 3, pp. 2123–2135, 2020. 1, 2
- [13] Z. Hong, X. Fan, T. Jiang, and J. Feng, “End-to-end unpaired image denoising with conditional adversarial networks,” in *AAAI*, 2020. 2
- [14] G. Jang, W. Lee, S. Son, and K. M. Lee, “C2n: Practical generative noise modeling for real-world denoising,” in *ICCV*, 2021.
- [15] D.-W. Kim, J. Ryun Chung, and S.-W. Jung, “Grdn: Grouped residual dense network for real image denoising and gan-based real-world noise modeling,” in *CVPR*, 2019. 2
- [16] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *ICLR*, 2015. 3, 5
- [17] I. Kobyzev, S. J. Prince, and M. A. Brubaker, “Normalizing flows: An introduction and review of current methods,” *TPAMI*, vol. 43, no. 11, p. 3964–3979, 2021. 2
- [18] A. Krull, T.-O. Buchholz, and F. Jug, “Noise2void-learning denoising from single noisy images,” in *CVPR*, 2019. 2, 8
- [19] D. T. Kuan, A. A. Sawchuk, T. C. Strand, and P. Chavel, “Adaptive noise smoothing filter for images with signal-dependent noise,” *TPAMI*, no. 2, pp. 165–177, 1985. 3
- [20] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila, “Noise2noise: Learning image restoration without clean data,” in *ICML*, 2018. 2, 3, 5, 7
- [21] C. Liu, R. Szeliski, S. B. Kang, C. L. Zitnick, and W. T. Freeman, “Automatic estimation and removal of noise from a single image,” *TPAMI*, vol. 30, no. 2, pp. 299–314, 2007. 1, 3
- [22] X. Liu, M. Tanaka, and M. Okutomi, “Practical signal-dependent noise parameter estimation from a single noisy image,” *TIP*, vol. 23, no. 10, pp. 4361–4371, 2014. 1, 2
- [23] Y. Liu, S. Anwar, Z. Qin, P. Ji, S. Caldwell, and T. Gedeon, “Disentangling noise from images: A flow-based image denoising neural network,” *arXiv preprint arXiv:2105.04746*, 2021. 1
- [24] A. M. Mohsen, M. F. Tompsett, and C. H. Sèquin, “Noise measurements in charge-coupled devices,” *IEEE Transactions on Electron Devices*, vol. 22, no. 5, pp. 209–218, 1975. 2
- [25] N. Moran, D. Schmidt, Y. Zhong, and P. Coady, “Noisier2noise: Learning to denoise from unpaired noisy data,” in *CVPR*, 2020. 2, 8
- [26] S. Nam, Y. Hwang, Y. Matsushita, and S. J. Kim, “A holistic approach to cross-channel image noise modeling and its application to image denoising,” in *CVPR*, 2016. 2
- [27] T. Pang, H. Zheng, Y. Quan, and H. Ji, “Recorrupted-to-recorrupted: Unsupervised deep learning for image denoising,” in *CVPR*, 2021. 2, 8
- [28] T. Plotz and S. Roth, “Benchmarking denoising algorithms with real photographs,” in *CVPR*, 2017. 2
- [29] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, “Image denoising using scale mixtures of gaussians in the wavelet domain,” *TIP*, vol. 12, no. 11, pp. 1338–1351, 2003. 1, 3
- [30] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015. 7
- [31] T. Seybold, C. Keimel, M. Knopp, and W. Stechele, “Towards an evaluation of denoising algorithms with respect to realistic camera noise,” in *2013 IEEE International Symposium on Multimedia*. IEEE, 2013. 2
- [32] W. Wang, X. Chen, C. Yang, X. Li, X. Hu, and T. Yue, “Enhancing low light videos by exploring high sensitivity camera noise,” in *ICCV*, 2019. 1
- [33] V. Wolf, A. Lugmayr, M. Danelljan, L. Van Gool, and R. Timofte, “Deflow: Learning complex image degradations from unpaired data with conditional flows,” in *CVPR*, 2021. 2
- [34] J. Xu, H. Li, Z. Liang, D. Zhang, and L. Zhang, “Real-world noisy image denoising: A new benchmark,” *arXiv preprint arXiv:1804.02603*, 2018. 2
- [35] J. Xu, Y. Huang, M.-M. Cheng, L. Liu, F. Zhu, Z. Xu, and L. Shao, “Noisy-as-clean: learning self-supervised denoising from corrupted image,” *TIP*, vol. 29, pp. 9316–9329, 2020. 2, 8
- [36] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, “Beyond a gaussian denoiser: Residual learning of deep cnn for

- image denoising,” *TIP*, vol. 26, no. 7, pp. 3142–3155, 2017. [3](#), [4](#)
- [37] M. Zhang and B. K. Gunturk, “Multiresolution bilateral filtering for image denoising,” *TIP*, vol. 17, no. 12, pp. 2324–2333, 2008. [3](#)
- [38] Y. Zhang, H. Qin, X. Wang, and H. Li, “Rethinking noise synthesis and modeling in raw denoising,” in *ICCV*, 2021. [2](#), [5](#), [6](#), [8](#)
- [39] Y. Zhou, J. Jiao, H. Huang, Y. Wang, J. Wang, H. Shi, and T. Huang, “When awgn-based denoiser meets real noises,” in *AAAI*, 2020. [2](#)
- [40] F. Zhu, G. Chen, and P.-A. Heng, “From noise modeling to blind image denoising,” in *CVPR*, 2016. [2](#)