

Blind2Sound: Self-Supervised Image Denoising without Residual Noise

Jiazheng Liu^{1,†} Zejin Wang^{1,†} Bohao Chen^{2,1} Hua Han^{1,3,*}

¹State Key Laboratory of Brain Cognition and Brain-inspired Intelligence Technology, Institute of Automation, Chinese Academy of Sciences, Beijing, China

²School of Advanced Interdisciplinary Sciences, University of Chinese Academy of Sciences, Beijing, China

³School of Future Technology, University of Chinese Academy of Sciences, Beijing, China

{liujiazheng2018, wangzejin2018, chenbohao2024, hua.han}@ia.ac.cn

Abstract

Self-supervised blind denoising for Poisson-Gaussian noise remains a challenging task. Pseudo-supervised pairs constructed from single noisy images re-corrupt the signal and degrade the performance. The visible blindspots solve the information loss in masked inputs. However, without explicitly noise sensing, mean square error as an objective function cannot adjust denoising intensities for dynamic noise levels, leading to noticeable residual noise. In this paper, we propose Blind2Sound, a simple yet effective approach to overcome residual noise in denoised images. The proposed adaptive re-visible loss senses noise levels and performs personalized denoising without noise residues while retaining the signal lossless. The theoretical analysis of intermediate medium gradients guarantees stable training, while the Cramer Gaussian loss acts as a regularization to facilitate the accurate perception of noise levels and improve the performance of the denoiser. Experiments on synthetic and real-world datasets show the superior performance of our method, especially for single-channel images.

1. Introduction

Supervised image denoisers [38, 47, 48, 19, 4, 43, 45] have demonstrated impressive and superior performance by utilizing numerous noisy-clean pairs. Notably, several researchers [38, 47, 48] pioneered the removal of additive white Gaussian noise and achieved notable performance gains. These denoisers [19, 4, 43, 45] can also eliminate other noise patterns, especially signal-dependent noise in the real world. However, acquiring a large amount of paired data to train a denoiser can be costly and, in some cases, even impossible, such as CT and MRI. Moreover, the performance of these denoisers rapidly degrades once they en-

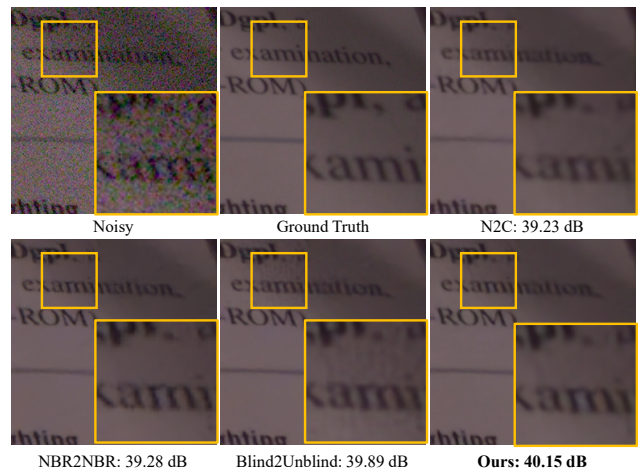


Figure 1: **Real-world image denoising on the SIDD validation.** Our approach performs thorough noise removal, while other methods have considerable noise artifacts.

counter unknown noise patterns due to the inherent prior embedded in the training data.

To tackle the challenge of acquiring paired data, researchers have conducted numerous studies [7, 44, 12, 41, 27, 24, 5] in two categories: data synthesis and self-supervised denoisers. Data synthesis involves generating paired data for supervised training via adding noise to clean images. For instance, UPI [7] and CycleISP [44] analyze the operations of the camera imaging pipeline on a signal-dependent noise and create an imaging framework capable of generating arbitrary real image pairs in both raw-RGB and sRGB space. Other researchers [12, 41] synthesize noisy-clean image pairs for training using pre-estimated noise parameters. However, while data synthesis approaches are practical with limited specific data, their poor generality limits their broader application.

Noise2Noise [27], a special case of supervised denoising that uses corrupted image pairs, is the foundation for self-supervised denoising. Self-supervised denoisers, which

* Corresponding author, † Equal contribution

learn only from a single raw noise image, have become the leading solution for denoising without data collection limitations. Blindspot schemes, such as manual masking [24, 5] and blindspot networks (BSN) [25, 41, 10, 9], enable single-image denoising by assuming that the signal is context-dependent and the noise is irrelevant. However, the information loss and sub-optimal masking strategies degrade the upper bound of denoising performance. Some works [42, 33, 35, 14] introduce additional noise to construct training data pairs. The pre-calibrated noise level restricts their application, and the re-corruption worsens the loss of valuable signals. NBR2NBR [22] sub-samples low-resolution training pairs from single noisy images for supervised training. The surrounding pixel approximation results in over-smoothing and block artifacts. FBI-Denoiser (FBI-D) [8] reimplements Gaussian estimation [11] as tensor operations combined with Generalized Anscombe Transformations (GAT) [3] to formulate Gaussian loss. However, Gaussian loss without fine-grained constraints zooms noise estimation errors. AP-BSN [26] breaks spatial correlation but remains in the scope of BSN. Blind2Unblind [40] achieves lossless denoising using the re-visible transition, but the greedy pixel-level fitting without noise sensing leads to sizable residual noise.

This paper presents a novel framework named Blind2Sound for self-supervised blind denoising, which achieves personalized denoising without noise residues while ensuring signal lossless. The framework includes an adaptive re-visible loss for lossless and personalized denoising and a Cramer Gaussian loss as a regularization to accurately sense noise knowledge. The framework assumes that the masked and visible branches are two independent generative processes, while denoising results following normal distributions and the respective noise following a Poisson-Gaussian distribution. This independence ensures that the masked results do not suppress the visible denoising. For optimal denoising, the adaptive loss adjusts noise removal via noise levels, while compatible with the lossless re-visible framework. The adaptive loss consists of two branches and their noise parameters, and the masked branch as an intermediate medium for gradient update. Due to signal dependence, the masked branch’s bottleneck significantly suppresses Poisson noise’s accuracy. Thus, the weighted Poisson noises in the mixed marginal likelihood is simplified to a single Poisson component. Moreover, the Cramer Gaussian loss introduces fine-grained noise knowledge, including sub-block and cross-channel constraints, to boost the estimation accuracy.

Our contribution can be summarized in three aspects:

1. The proposed self-supervised denoising framework adjusts denoising intensities based on sensed noise levels, achieving personalized denoising without residual noise while the signal is lossless.

2. We provide theoretical analysis for adaptive re-visible loss and the gradient of the intermediate medium.
3. Our approach shows excellent performance on various real-world datasets with dynamic noise patterns.

2. Related Work

Self-Supervised Blind Image Denoising Although Noise2Noise [27] makes collecting paired data much more manageable, it remains impractical as it requires noisy image pairs from the same source. Therefore, blindspot schemes such as manual masking [24, 5] and blindspot networks [25, 10] have been developed to achieve self-supervised denoising. However, these schemes suffer from the loss of valuable information, leading to severe artifacts. To address this issue, several methods, such as Laine [25], DBSN [41], FC-AIDE [10], BP-AIDE [9], AP-BSN [26] and ZS-N2N [31], have been proposed to enhance BSN. However, these methods either perform poorly or have a slow inference on Poisson-Gaussian noise. FBI-D [8] proposes a fast and effective blind image denoiser for Poisson-Gaussian noise, but its noise estimator has low accuracy in sRGB space without fine-grained noise constrains. Besides, the post-processing step can reinforce denoising errors accumulated in previous steps. CVF-SID [34] disentangles clean images and noise components in a cyclic manner, but relies on assumptions about noise distribution, limiting its applicability. DCD-Net [50] iteratively trains a denoiser and noise estimator, yet still requires an initial noise estimate, which can be challenging for real-world scenarios. ScoreDVI [13] uses score priors and deep variational inference for real-world denoising, but its computational complexity may limit practicality compared to simpler methods [50, 40]. Blind2Unblind [40] overcomes these limitations and achieves fast lossless denoising under blindspots. However, implicit denoising without sensing noise levels is unable to provide personalized noise removal for Poisson-Gaussian noise, resulting in residual noise artifacts.

Noise Estimation Methods Current noise estimation methods [28, 37, 11, 17, 29] are mainly designed for two noise models: additive white Gaussian noise (AWGN) and Poisson-Gaussian noise. For AWGN, a signal-independent noise, once the noise variance is known, its distribution can be determined. Previous methods [28, 37] for estimating AWGN variance used low-rank patch selection and principal component analysis. Chen *et al.* [11] addressed the bias issue for Gaussian parameter estimation using statistical decomposition of eigenvalues. In contrast, Poisson-Gaussian noise [17] is a mixture of Poisson and Gaussian components and is often considered as source-related noise in real-world scenarios. Foi *et al.* [17] introduced the Poisson-Gaussian model and provided fully automated noise parameter estimation for a single noisy image. Recently, Liu *et al.* [29]

improved local mean and noise variance estimation for selected low-rank patches.

3. Problem Setting and Preliminaries

3.1. Problem Formulation

Given a noisy observation \mathbf{y} , we aim to learn the clean image \mathbf{x} and its noise parameters directly from a single noisy image. We focus on Poisson-Gaussian noise [17], which is common in real-world imaging sensors. The noise-corrupted observation \mathbf{y} can be described as follows:

$$\mathbf{y} = \alpha P + N, \quad (1)$$

where $P \sim \text{Poisson}(\mathbf{x}/\alpha)$ is the signal-dependent Poisson noise caused by photon counting, and $N \sim \mathcal{N}(0, \sigma^2)$ is the signal-independent Gaussian noise resulting from electric and thermal noise. Here, α is a scaling factor that depends on the quantum efficiency and analog gain. To simplify the problem, the Poisson noise is approximated as signal-independent Gaussian noise [20], and the corruption can be reformulated as:

$$\mathbf{y} = \mathbf{x} + \mathcal{N}(0, \alpha \mathbf{x} + \sigma^2). \quad (2)$$

3.2. Generalized Anscombe Transformation (GAT)

GAT [3] is a popular variance-stabilized transform used for Poisson-Gaussian noise. It converts the mixture noise into stable Gaussian noise with unit variance:

$$G_{\alpha, \sigma}(\mathbf{y}) = \frac{2}{\alpha} \sqrt{\alpha \mathbf{y} + \frac{3}{8} \alpha^2 + \sigma^2}. \quad (3)$$

3.3. Gaussian Loss

As discussed in Section 1, Byun *et al.* [8] introduce a noise estimation network, denoted as $h_{\theta}(\cdot)$, that takes Poisson-Gaussian noise-corrupted images \mathbf{y} as input and produces the estimated noise parameters $(\hat{\alpha}, \hat{\sigma})$. To train this network, they reimplement the Gaussian parameter estimation proposed in [11] as tensor operations $\eta(\cdot)$ and leverage the property that the transformed noise level using GAT should be close to unit variance. Therefore, the Gaussian loss is defined for noise estimation as:

$$\arg \min_{\theta} \mathbb{E}_{\mathbf{y}} \left\| \eta \left(G_{\hat{\alpha}(\theta), \hat{\sigma}(\theta)}(\mathbf{y}) \right) - 1 \right\|_2^2, \quad (4)$$

the noise estimation network $h_{\theta}(\mathbf{y})$ predicts the noise parameters $\hat{\alpha}(\theta)$ and $\hat{\sigma}(\theta)$. Then, the transformed noisy image with unit variance is denoised, and the output of BSN is mapped back to the original image using the Inverse Anscombe transformation (IAT) [30]. Note that noise estimation and denoising are trained separately without joint optimization.

4. Main Method

As illustrated in Figure 2, the training process involves two modules: 1) The denoising network $f_{\omega}(\cdot)$ that outputs denoised results of the noisy masked volume $\Omega_{\mathbf{y}}$ and the raw noisy image \mathbf{y} , including their mean μ_m, μ_v , and covariance Σ_m, Σ_v . 2) The noise estimator $g_{\theta}(\cdot)$ that calculates the Poisson-Gaussian noise parameters $(\alpha, \sigma_1, \sigma_2)$. Since actual noise levels of denoised images may differ from raw noisy inputs, Cramer Gaussian loss only serves as a regularization to assist the denoiser in picking appropriate noise levels. Namely, the adaptive re-visible loss also determines the noise parameters for explicit and lossless personalized denoising. The training set D contains n training images denoted as $D = \{\mathbf{y}_i\}_{i=1}^n$, where \mathbf{y}_i denotes the i^{th} raw noisy image. The details of the proposed method are presented in subsequent sections.

4.1. Motivation

The re-visible framework overcomes the information loss in blindspot-driven methods. However, mean square error as objective function cannot adapt to varying noise levels, frequently resulting in residual noise in denoised images. Therefore, it is necessary to design an adaptive loss that adjusts denoising intensities based on sensed noise levels. For optimal noise removal, the adaptive loss should personalized denoising without noise residues, while ensuring valuable signal lossless. Lossless denoising requires that the adaptive loss is compatible with the re-visible framework. For this purpose, we regard the masked and visible branches as Gaussian processes and design an adaptive re-visible loss that satisfies personalized and lossless requirements.

4.2. Adaptive Re-Visible Loss

The goal of adaptive re-visible loss is to achieve optimal noise removal while retaining lossless denoising. We accomplish this by re-considering the re-visible scheme from Bayesian reasoning and developing a personalized noise removal method that does not require an auxiliary branch for noise estimation during inference.

First, we model $p(\mathbf{z}_1 | \Omega_{\mathbf{y}})$ as a multivariate Gaussian, which represents that the latent clean image \mathbf{z}_1 is generated from the masked noisy volume $\Omega_{\mathbf{y}}$ as follows:

$$\mathbf{z}_1 \sim \mathcal{N}(\mathbf{z}_1 | \mu_m, \Sigma_m), \quad (5)$$

where $\mathcal{N}(\cdot | \mu_m, \Sigma_m)$ denotes the multivariate Gaussian distribution with mean μ_m and variance Σ_m .

For the masked branch, Eq. (2) incorporates extra noise knowledge into the explicit corruption model, provided as the likelihood $p(\mathbf{y} | \mathbf{z}_1)$ given a clean value. Therefore, the marginal likelihood of the noisy training data can be constructed via the distribution of unobserved clean data \mathbf{z}_1 :




Figure 2: The architecture of the proposed Blind2Sound framework. The masked volume is represented by Ω_y , and the dotted lines indicate the forward process without updating the gradient. The Global Masker and Global Mask Mapper are detailed in the Supplementary Material (S.M.). Note that the noise estimator and the denoiser are trained through joint optimization. During inference, the auxiliary branches, namely the noise estimator and masked branch, are removed, and the denoiser directly generates denoised images from raw noisy inputs y .

$$p(\mathbf{y}_1) = \int p(\mathbf{y}|\mathbf{z}_1)p(\mathbf{z}_1|\Omega_y)d\mathbf{z}_1. \quad (6)$$

As illustrated in Eq. (6), when only noisy training data y are available, a known noise model is able to explicitly predict the masked prior $p(\mathbf{z}_1|\Omega_y)$. Specifically, for an approximate Gaussian noise model, the covariance of two mutually independent Gaussian convolutions is simply the sum of the components [6]. Hence, the marginal likelihood $p(\mathbf{y}_1)$ is calculated in closed form, allowing to obtain the distribution of \mathbf{z}_1 by maximizing Eq. (6). According to Eqs. (2) and (5) as well as the above analysis, the mean and variance of \mathbf{y}_1 become:

$$\mathbf{y}_1 \sim \mathcal{N}(\mathbf{y}_1|\mu_m, \Sigma_m + \text{diag}(\alpha_1\mu_m) + \sigma_1\mathbf{I}). \quad (7)$$

For the visible branch, we construct $p(\mathbf{z}_2|\mathbf{y})$ as the generation of the latent clean image \mathbf{z}_2 from the raw noise image y , which then becomes:

$$\mathbf{z}_2 \sim \mathcal{N}(\mathbf{z}_2|\mu_v, \Sigma_v), \quad (8)$$

where the mean μ_v and variance Σ_v are directly generated from the raw noise image y without gradient update. The marginal likelihood for the visible branch via the distribution of unobserved clean data \mathbf{z}_2 is then formulated as:

$$p(\mathbf{y}_2) = \int p(\mathbf{y}|\mathbf{z}_2)p(\mathbf{z}_2|\mathbf{y})d\mathbf{z}_2. \quad (9)$$

Similar to Eq. (7), the mean and variance of \mathbf{y}_2 become:

$$\mathbf{y}_2 \sim \mathcal{N}(\mathbf{y}_2|\mu_v, \Sigma_v + \text{diag}(\alpha_2\mu_v) + \sigma_2\mathbf{I}). \quad (10)$$

The marginal likelihood $p(\mathbf{y}_1)$ for the blind branch and $p(\mathbf{y}_2)$ for the visible branch are now available. Due to identity mapping, the mean and variance of \mathbf{z}_2 do not participate in backpropagation. However, maximizing only the blind distribution $p(\mathbf{z}_1|\Omega_y)$ via Eq. (6) has limited performance. To improve the performance, the loss errors of the visible branch $p(\mathbf{z}_2|\mathbf{y})$ are incorporated into the mask gradient. The decorrelation of the two branches enhances visible denoising without suppressing the masked results. Thus, \mathbf{y}_1 and \mathbf{y}_2 are modeled as i.i.d., and a Gaussian mixture is applied to boost their representation. Combining Eqs. (7) and (10), an enhanced target distribution y is obtained while retaining the independence of the two branches:

$$y \sim \sum_{i=1}^2 \pi_i \cdot \mathcal{N}(y_i|\mu_{y_i}, \Sigma_{y_i}), \quad (11)$$

where π_i is a hyper-parameter for the degree of re-visible. Besides, $0 \leq \pi_i \leq 1$ and $\pi_1 + \pi_2 = 1$. Set the blind factor π_1 to $1/(1+\lambda)$, the visible factor π_2 to $\lambda/(1+\lambda)$ and λ is a growing constant. Then, Eq. (11) is reformulated as:

$$y \sim \mathcal{N}(y|\frac{\mu_m + \lambda\mu_v}{1+\lambda}, \frac{\Sigma_{y_1} + \lambda^2\Sigma_{y_2}}{(1+\lambda)^2}). \quad (12)$$

We use $y \sim \mathcal{N}(y|\mu_y, \Sigma_y)$ to simplify notation and denote the clean target image as x . The mask mean μ_m is a lower bound of x and the signal-dependent factor α_1 magnifies this error. To improve accuracy, we replace the noise model in Eq. (12) with a more precise $p(y|x)$ that has zero mean

and variance $\text{diag}(\alpha\mu_y) + \pi_1^2\sigma_1\mathbf{I} + \pi_2^2\sigma_2\mathbf{I}$. The enhanced mixture marginal likelihood that bridges the blind and visible branches then become:

$$p(\mathbf{y}) = \int p(\mathbf{y}|\mathbf{x})p(\mathbf{x}|\mathbf{y}, \Omega_y)dx. \quad (13)$$

To fit the observed noisy training data, we minimize its negative log-likelihood loss in the training phase as follows:

$$\begin{aligned} \mathcal{L}_{arv} &= -\log p(\mathbf{y}) = -\log [\pi_1 p(\mathbf{y}_1) + \pi_2 p(\mathbf{y}_2)] \\ &= \frac{1}{2}[(\mathbf{y} - \mu_y)^T \Sigma_y^{-1} (\mathbf{y} - \mu_y)] \\ &\quad + \frac{1}{2} \log |\Sigma_y| + \text{const}, \end{aligned} \quad (14)$$

where const is an additive constant term that can be discarded, and \mathcal{L}_{arv} denotes the proposed adaptive re-visible loss. When the denoiser converges, the following is the optimal clean value \tilde{x} of Eq. (14):

$$\tilde{x} = \frac{\mu_m + \lambda\mu_v}{1 + \lambda}. \quad (15)$$

Assuming that $\mu_m = \mathbf{x} + \varepsilon_1$ and $\mu_v = \mathbf{x} + \varepsilon_2$. Empirically, $\|\varepsilon_1\|_1 > \|\varepsilon_2\|_1$ because valuable information is lost in Ω_y . Combined with Eqs. (11) and (15), it can be concluded that $\mu_m \leq \tilde{x} \leq \mu_v$ since \tilde{x} is the weighted average of μ_m and μ_v .

Using the analytic form of adaptive re-visible loss in Eq. (14), we explore and confirm the collaborative mechanism between the gradient update medium μ_m and the visible constant μ_v . Let $\mathbf{n}_y = \mathbf{y} - \mu_y$, the derivative of the medium μ_m gives its gradient ∇_{μ_m} (see S.M.):

$$\begin{aligned} \frac{\partial \mathcal{L}_{arv}}{\partial \mu_m} &= -\frac{1}{1 + \lambda} \Sigma_y^{-1} \mathbf{n}_y + \frac{\alpha}{(\lambda + 1)} (\Sigma_y^{-1} - \\ &\quad \Sigma_y^{-1} \mathbf{n}_y \mathbf{n}_y^T \Sigma_y^{-1}) \odot \mathbf{I} \odot (\sqrt{\mu_y \mu_y^T})' \mu_y. \end{aligned} \quad (16)$$

According to Eq. (16), it is observed that including the gradient update term μ_m in $\text{diag}(\alpha\mu_y)$ results in severe instability during training due to a complicated second term in the gradient. Therefore, disabling the gradient of $\text{diag}(\alpha\mu_y)$ is considered to stabilize the training process and improve the performance of the denoiser. Moreover, denoising can be performed directly from the raw noise image \mathbf{y} during inference. We provide more details in the supplementary material.

Laine *et al.* [25] utilize Bayesian reasoning to incorporate information from \mathbf{y} into maximum posterior probabilities (MAP) during test time. However, post-processing, which is not involved in training, performs poorly in practice. In contrast, additional MAP is redundant for adaptive re-visible denoising as $p(\mathbf{x}|\Omega_y, \mathbf{y})$ in Eq. (13) already includes information from \mathbf{y} . As a result, our approach outperforms other self-supervised methods.

4.3. Cramer Gaussian Loss

Gaussian loss [8] estimates noise directly from global images, which ignores the perceptual dimension of local noise knowledge. As a result, the estimated Poisson Gaussian parameters are less accurate. To overcome the coarse-scale limitation, the proposed Cramer Gaussian loss uses fine-grained local sub-block or cross-channel constraints to enrich noise perception dimensions, reducing the solution space to find the exact median of the noise level.

For single-channel images without channel correlation, each sub-block of the GAT-transformed images should have same noise levels as global ones. To ensure this, we introduce a fine-grained sub-block noise level constraint based on the coarse-scale constraint. The estimation incorporates overlapping sub-blocks at four corners as local noise knowledge for single-channel images. Denote $g_\theta(\mathbf{y})$ as the estimated noise parameter $(\hat{\alpha}, \hat{\sigma})$, $\eta(\cdot)$ as the Gaussian estimator [11] that estimates the Gaussian noise variance, and $G_{g_\theta}(\mathbf{y})$ as the GAT-transformed image $G_{g_\theta(\mathbf{y})}(\mathbf{y})$. The estimated Gaussian variance for the transformed noise and its sub-blocks should approximate unit variance. Cramer Gaussian loss for single-channel image \mathbf{y} becomes:

$$\mathcal{L}_{est} = \sum_{s=1}^4 \|\eta(G_{g_\theta}(\mathbf{y}^s)) - 1\|_2^2 + \|\eta(G_{g_\theta}(\mathbf{y})) - 1\|_2^2, \quad (17)$$

where \mathbf{y}_i^s denotes the s^{th} sub-block cropped from \mathbf{y}_i . We crop four identical sub-blocks from four corners. Each sub-block is three-quarters the size of the original image.

For multi-channel images, the implied noise level should be the same for different channels. To ensure a unique solution space, we introduce cross-channel noise level approximation, enabling cross-channel information exchange and overcoming the inherent limitation of Gaussian loss, *i.e.*, the problem of inter-channel estimated error offsetting. For each channel in the GAT-transformed image, the Gaussian noise variance of each channel should approximate the unit variance, and the noise level should be the same for each channel. Thus, the Cramer Gaussian loss for the multi-channel image \mathbf{y} becomes:

$$\begin{aligned} \mathcal{L}_{est} &= \sum_{j \neq k}^c \|\eta(G_{g_\theta}(\mathbf{y}_j)) - 1\|_2^2 \\ &\quad + \|\eta(G_{g_\theta}(\mathbf{y}_j)) - \eta(G_{g_\theta}(\mathbf{y}_k))\|_2^2, \end{aligned} \quad (18)$$

where c is the number of channels in image \mathbf{y}_i . j, k represent the j^{th} and k^{th} channel, respectively.

5. Experimental Results

5.1. Implementation Details

Training Details We use the same noise estimator as the FBI-D [8] and a modified U-Net [25, 22, 40] as the denoising network. Adam [23] with a weight decay of $1e^{-8}$ is used as the optimizer. The initial learning rate for the

noise estimator is 0.0001. For a small training set, the initial learning rate of the denoising network is 0.001, while 0.0003 and 0.0001 for ILSVRC2012 [16] validation set and SIDD [2], respectively. The learning rate for the noise estimator decreases by half every 10 epochs with 50 epochs trained, while the learning rate for the denoising network is halved every 20 epochs with 100 epochs trained. As for the hyper-parameter in the adaptive re-visible loss, we set $\lambda = 3$ as the initial value and progressively increase it to 11. The denoising network and noise estimation are jointly optimized during training. The patches of size 128×128 are randomly cropped for training. All models are trained using Python 3.10.4, Pytorch 1.11.0 [36], and an Nvidia Tesla V100 GPU.

Datasets We consider five types of Poisson-Gaussian noise in synthetic noise estimation: (1) $PG1 : \alpha = 0.1, \sigma = 0.02$, (2) $PG2 : \alpha = 0.1, \sigma = 0.0002$, (3) $PG3 : \alpha = 0.05, \sigma = 0.02$, (4) $PG4 : \alpha = 0.05, \sigma = 0.0002$, (5) $PG5 : \alpha = 0.01, \sigma = 0.02$. For grayscale images, we use BSD400 [32] for the training set, and for sRGB images, we use CBSD432 [47]. The noise levels are estimated on standard BSD68 [39] and CBSD68 [39] for grayscale and sRGB images, respectively. For synthetic denoising, we use ILSVRC2012 [16] validation set for sRGB image denoising and BSD400 [32] for grayscale image denoising. Specifically, following the setting in [25, 22, 40], we select 44328 images with sizes between 256×256 and 512×512 pixels from ILSVRC2012 validation set for training. The test sets used for grayscale image denoising are Set12, BSD68 [39], and Urbabn100 [21], while for sRGB denoising, we use Kodak [18], BSD300 [32], and Set14 [46]. For real noise experiments, we use SIDD [2] for real-world denoising in raw-RGB space and Fluorescence Microscopy Denoising (FMD) [49] for real-world grayscale denoising. We train SIDD using only the raw-RGB images in the SIDD Medium Dataset and validate and test using the SIDD Validation and Benchmark Datasets. Note that we receive the evaluation results for the SIDD Benchmark from the online public website [1]. For FMD, we use the Confocal Mice and Two-Photon Mice datasets and the 19th view is used for testing.

Baselines We compare the proposed Cramer Gaussian loss with three noise estimation methods, including the Gaussian loss in FBI-D [8], Foi [17], and Liu [29], and evaluate the denoising performance of Blind2Sound against two supervised methods (N2C [38] and N2N [27]), a traditional approach (GAT+BM3D [15]), and five self-supervised algorithms (N2V [24], NBR2NBR [22], FBI-D [8] and Blind2Unblind [40]). We adopt the experimental settings of FBI-D and Blind2Unblind, and reproduce the results using the official implementation for a fair comparison.

Noise Level	Grayscale				sRGB	
	Foi [17] ($\hat{\alpha}, \hat{\sigma}$)	Liu [29] ($\hat{\alpha}, \hat{\sigma}$)	FBI-D [8] ($\hat{\alpha}, \hat{\sigma}$)	Ours ($\hat{\alpha}, \hat{\sigma}$)	FBI-D [8] ($\hat{\alpha}, \hat{\sigma}$)	Ours ($\hat{\alpha}, \hat{\sigma}$)
$PG1$	0.096 /0.042	0.072/0.045	0.092/0.039	0.093/ 0.019	0.080/ 0.021	0.099 /0.001
$PG2$	0.097 /0.035	0.071/0.044	0.083/0.061	0.090/ 0.001	0.074/0.033	0.096 /0.001
$PG3$	0.049 /0.031	0.04/0.04	0.052/0.003	0.048/ 0.022	0.041/0.015	0.051 /0.020
$PG4$	0.051 /0.018	0.039/0.034	0.046/0.035	0.048/ 0.013	0.040/0.013	0.050 /0.0007
$PG5$	0.011/0.027	0.007/0.032	0.009/0.034	0.010 / 0.021	0.008/0.009	0.010 /0.028

Table 1: The average noise parameter estimates for grayscale images in BSD68 and sRGB images in CBSD68.

5.2. Results for Noise Estimation

We first evaluate the Cramer Gaussian loss for noise estimation. Table 1 shows the average noise parameters ($\hat{\alpha}, \hat{\sigma}$) predicted by Foi [17], Liu [29], FBI-D [8], and Cramer Gaussian loss for five noise patterns in grayscale and sRGB images. Cramer Gaussian loss shows superior performance in estimating Gaussian parameters for grayscale images, and its Poisson level estimation is comparable to Foi while requiring less inference time. Moreover, compared to FBI-D, Cramer Gaussian loss eliminates its severe Gaussian estimation error on grayscale images using a fine-grained fusion strategy. For sRGB space, cross-channel approximation enables our method to predict Poisson-Gaussian parameters close to their actual value. In contrast, FBI-D is inaccurate for estimating noise in sRGB images.

Noise Type	Method	BSD68	Set12	Urbabn100
$\alpha = 0.01$ $\sigma = 0.0002$	Baseline, N2C [38]	30.82/0.877	31.49/0.881	30.80/0.901
	Baseline, N2N [27]	30.76/0.876	31.47/0.880	30.84/0.900
	GAT+BM3D [15]	30.08/0.849	30.36/0.871	30.44/0.881
	N2V [24]	29.04/0.824	29.29/0.830	27.48/0.830
	FBI-D [8]	28.09/0.784	29.13/0.828	27.28/0.829
	NBR2NBR [22]	30.38/0.861	31.07/0.869	30.15/0.884
	Blind2Unblind [40]	30.61/0.869	31.45/0.880	30.70/0.900
	Ours	30.83/0.875	31.68/0.886	31.14/0.908
$\alpha = 0.01$ $\sigma = 0.02$	Baseline, N2C [38]	30.54/0.868	31.26/0.876	30.51/0.892
	Baseline, N2N [27]	30.62/0.871	31.34/0.878	30.70/0.899
	GAT+BM3D [15]	29.86/0.844	30.55/0.857	30.23/0.882
	N2V [24]	29.01/0.828	29.43/0.842	27.53/0.838
	FBI-D [8]	27.95/0.776	29.11/0.824	27.21/0.822
	NBR2NBR [22]	30.28/0.860	31.00/0.868	30.15/0.886
	Blind2Unblind [40]	30.41/0.865	31.27/0.877	30.49/0.896
	Ours	30.62/0.869	31.47/0.880	30.91/0.902
$\alpha = 0.05$ $\sigma = 0.02$	Baseline, N2C [38]	27.11/0.765	27.78/0.801	26.69/0.813
	Baseline, N2N [27]	27.07/0.762	27.71/0.799	26.58/0.808
	GAT+BM3D [15]	26.16/0.732	27.26/0.785	26.40/0.795
	N2V [24]	26.25/0.719	26.48/0.755	24.71/0.741
	FBI-D [8]	25.75/0.683	26.42/0.745	24.51/0.722
	NBR2NBR [22]	26.88/0.751	27.51/0.783	26.39/0.800
	Blind2Unblind [40]	27.02/0.757	27.65/0.796	26.54/0.805
	Ours	27.17/0.766	27.96/0.805	26.96/0.819

Table 2: Grayscale image denoising results. **Bold** and underlined are the highest and second without supervision.

5.3. Results for Synthetic Denoising

Grayscale Denoising The denoising results on synthetic grayscale images are presented in Table 2. Three Poisson-Gaussian noise levels are simulated to evaluate the performance of the proposed method under low, medium, and

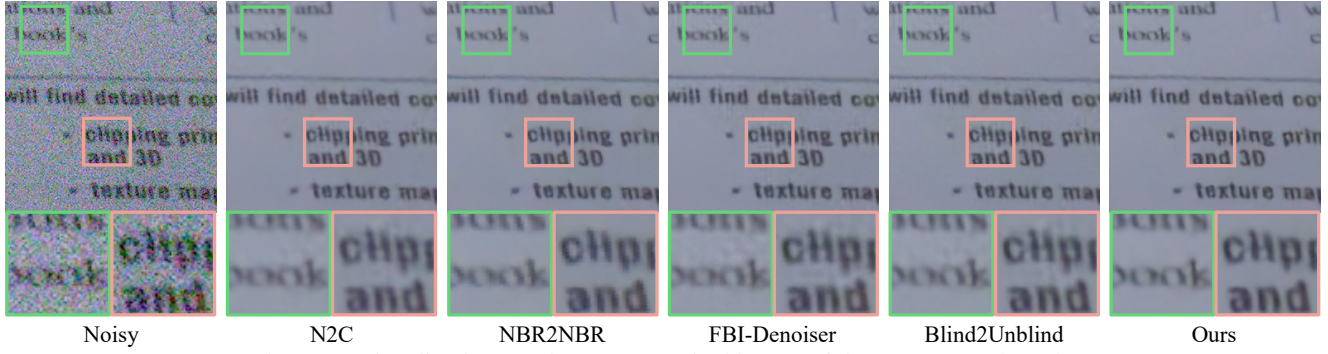


Figure 3: Visualization results on one typical image of the SIDD Benchmark.

high Poisson-Gaussian noise levels. The results show that our approach outperforms several denoising methods, including the traditional denoising method GAT+BM3D and four self-supervised denoising methods (N2V, NBR2NBR, FBI-D and Blind2Unblind) for both low and high Poisson-Gaussian noise levels. In addition, our method outperforms two supervised baselines (N2C and N2N) on Set12 and Urb100, with a maximum gain of 0.4 dB, and performs competitively on BSD68. These results highlight the superior generalization of our method on grayscale image denoising compared to supervised baselines. Comparing with other methods reveals the necessity of explicit personalized denoising.

Noise Type	Method	KODAK	SET14	BSD300
$\alpha = 0.01$ $\sigma = 0.0002$	Baseline, N2C [38]	34.67/0.925	33.16/0.904	33.61/0.929
	Baseline, N2N [27]	34.64/0.924	33.13/0.904	33.59/0.928
	GAT+BM3D [15]	33.63/0.913	31.80/0.883	32.47/0.909
	N2V [24]	31.68/0.871	30.72/0.848	29.71/0.844
	FBI-D [8]	31.66/0.871	30.66/0.848	29.69/0.843
	NBR2NBR [22]	34.10/0.918	32.69/0.896	32.89/0.919
	Blind2Unblind [40]	33.88/0.915	32.47/0.886	32.53/0.913
	Ours	34.23/0.920	32.75/0.896	33.00/0.921
$\alpha = 0.01$ $\sigma = 0.02$	Baseline, N2C [38]	34.39/0.920	32.93/0.899	33.28/0.923
	Baseline, N2N [27]	34.36/0.920	32.89/0.899	33.25/0.923
	GAT+BM3D [15]	33.39/0.909	31.58/0.876	32.21/0.904
	N2V [24]	31.51/0.867	30.56/0.845	29.55/0.838
	FBI-D [8]	31.54/0.867	30.56/0.846	29.58/0.838
	NBR2NBR [22]	<u>33.93/0.915</u>	<u>32.52/0.892</u>	<u>32.87/0.916</u>
	Blind2Unblind [40]	33.58/0.909	32.17/0.883	32.31/0.910
	Ours	34.12/0.917	32.60/0.893	32.95/0.917
$\alpha = 0.05$ $\sigma = 0.02$	Baseline, N2C [38]	30.80/0.854	29.69/0.837	29.45/0.843
	Baseline, N2N [27]	30.77/0.853	29.65/0.836	29.43/0.842
	GAT+BM3D [15]	29.19/0.824	27.58/0.801	27.87/0.799
	N2V [24]	29.03/0.793	28.14/0.785	27.42/0.759
	FBI-D [8]	29.15/0.800	28.29/0.791	27.48/0.763
	NBR2NBR [22]	30.49/0.848	29.46/0.832	29.20/0.837
	Blind2Unblind [40]	<u>30.58/0.849</u>	<u>29.52/0.832</u>	<u>29.27/0.837</u>
	Ours	30.69/0.850	29.57/0.833	29.33/0.838

Table 3: PSNR(dB)/SSIM on synthetic sRGB datasets.

sRGB Denoising The results of synthetic denoising for sRGB images are presented in Table 3. Our method outperforms GAT+BM3D and four self-supervised methods (N2V, NBR2NBR, FBI-D and Blind2Unblind) in both low and high noise levels. However, the gain over Blind2Unblind decreases at high noise levels due to numerous missing details that are valuable for restoring clean signals. Unlike grayscale denoising in Table 2, our method does not outper-

form supervised baselines in sRGB denoising but achieves competitive performance at high noise levels. Intuitively, cross-channel correlation makes noise removal in sRGB space more challenging than in grayscale images. As the noise level increases, the necessity for clean target supervision gradually decreases. Moreover, Blind2Unblind with greedy pixel-level objective outperforms NBR2NBR with neighboring approximation at high noise levels, highlighting the potential of lossless denoising for high noise levels.

Methods	SIDD			FMD	
	RAW Benchmark	RAW Validation	sRGB Benchmark	Confocal Mice	Two-Photon Mice
Baseline, N2C [38]	50.61/0.991	51.19/0.991	38.08/0.945	38.40/0.966	34.02/0.925
Baseline, N2N [27]	50.62/0.991	51.21/0.991	38.09/0.945	38.37/0.965	33.80/0.923
GAT+BM3D [15]	48.60/0.986	48.92/0.986	34.64/0.879	37.93/0.963	33.83/0.924
N2V [24]	48.01/0.983	48.55/0.984	34.21/0.864	37.49/0.960	33.38/0.916
DBSN [41]	49.56/0.987	50.13/0.988	36.77/0.917	30.61/0.730	26.24/0.423
BP-AIDE [9]	50.45/0.990	—	37.91/0.942	38.31/0.963	33.89/0.902
FBI-D [8]	50.57/0.990	—	38.07/0.942	38.32/0.963	33.95/0.908
NBR2NBR [22]	50.47/0.990	51.06/0.991	37.85/0.942	37.07/0.960	33.40/0.921
Blind2Unblind [40]	<u>50.79/0.991</u>	<u>51.36/0.992</u>	38.11/0.944	38.44/0.964	34.03/0.916
Ours	50.92/0.991	51.50/0.992	38.21/0.945	38.46/0.965	34.11/0.918

Table 4: PSNR(dB)/SSIM on SIDD and FMD dataset.

5.4. Results for Real-World Denoising

Table 4 shows the quantitative results for the SIDD benchmark and validation datasets in raw-RGB and sRGB space. The proposed Blind2Sound outperforms all self-supervised methods and supervised baselines, showcasing its strong generalization ability for real-world scenarios with various dynamic Poisson-Gaussian noise levels. Notably, our method achieves a gain of nearly 0.3 dB over FBI-D in raw-RGB space, validating the effectiveness of lossless denoising. Blind2Sound outperforms Blind2Unblind with a gain of 0.1 dB, confirming the importance of explicit personalized modeling. The visual results in Figure 3 demonstrate that Blind2Sound restores the highest level of texture details and pixel correlation while avoiding bad artifacts. In contrast, Blind2Unblind displays visible clumping shadows in multiple regions due to insufficient denoising. Even supervised baselines show notable artifacts under dynamic noise, further validating the superiority of noise-sensed personalized denoising. The denoising results on the

FMD dataset indicate that Blind2Sound outperforms self-supervised methods and has competitive performance compared to supervised baselines.




Figure 4: Ablation study on Cramer Gaussian loss with fixed weight 1 for L_{arv} and horizontal is the weight of L_{est} .

5.5. Ablation Study

This section presents ablation studies on various factors, including grain size, Cramer Gaussian loss, training scheme, noise model, re-visible correlation and visible factor. PSNR(dB)/SSIM is evaluated on CBSD68 while training on CBSD432 for 200 epochs.

Ablation study on grain size. Table 5 shows that fine-grained sub-block constraints improve the accuracy of the coarse-grained Gaussian loss, enhance robustness, and reduce solution space to provide more accurate noise parameters. However, smaller sub-block sizes result in limited noise context, leading to lower performance.

Ablation study on Cramer Gaussian loss. As shown in Figure 4, the denoiser with a weight of 0.01 performs best compared to 0 or 100. Joint optimization can restore more image details, and actual noise levels of optimal denoised images may differ from raw noisy images. Hence, Cramer Gaussian loss only serves as a regularization to assist the denoiser in sensing the most suitable noise level.

Ablation study on training scheme. As shown in Table 6, $T+J$ performs better than $T+P$ and $T+F$ at low noise, but they are comparable at high noise. At low noise, $T+P$ and $T+F$ mismatch the actual noise levels of restored images, thus degrading performance. At high noise, three training schemes estimate almost identical noise parameters.

Ablation study on noise model. Table 7 shows that \mathcal{M}_E performs slightly better than the other noise models at low noise but much better at high noise. \mathcal{M}_O amplifies Poisson noise error in masked branch due to signal dependence, while \mathcal{M}_S violates the independence of two branches.

Ablation study on re-visible correlation. Table 8 explores the impact of whether masked and visible branches are independent. The denoiser in the IID setting performs much better than in the non-IID setting. IID decouples the correlation between the blind and visible branches, thus achieving visible denoising without mask suppression.

Ablation study on visible factor. Table 9 shows the performance using different visible factors. The degree of visible

Grain Size	(0.01, 0.0002)	(0.01, 0.02)	(0.05, 0.02)	(0.1, 0.02)
CG	0.007/0.038	0.009/0.034	0.039/0.057	0.092/0.039
FG1	0.010/0.011	0.010/0.021	0.046/0.021	0.089/0.008
CG+FG1	0.010/0.012	0.010/0.021	0.048/0.022	0.093/0.019
CG+FG2	0.009/0.008	0.010/0.022	0.043/0.027	0.085/0.013
CG+FG1+FG2	0.009/0.009	0.010/0.022	0.046/0.024	0.092/0.017

Table 5: Ablation study on grain size for BSD68. CG is global coarse grain, FG1 is four sub-blocks each in three-quarter size and FG2 is nine sub-blocks in one-half size.

Training Scheme	(0.01, 0.0002)	(0.01, 0.02)	(0.05, 0.02)
$T+P$	32.78/0.914	32.69/0.913	29.24/0.838
$T+F$	33.02/0.917	32.70/0.913	29.26/0.839
$T+J$	33.14/0.920	32.89/0.916	29.29/0.840

Table 6: Ablation study on training scheme, $T+P$, $T+F$ and $T+J$ represent training using frozen pre-trained noise estimator, training with fixed true noise or joint training.

Loss Type	(0.01, 0.0002)	(0.01, 0.02)	(0.05, 0.02)
\mathcal{M}_O	33.11/0.919	32.84/0.913	29.12/0.833
\mathcal{M}_S	33.10/0.919	32.86/0.915	29.14/0.835
\mathcal{M}_E	33.14/0.920	32.89/0.916	29.29/0.840

Table 7: Ablation study on noise model. \mathcal{M}_O and \mathcal{M}_E denote Eq. (12) and Eq. (13). \mathcal{M}_S means $\text{diag}(\alpha\mu_y) + \sigma\mathbf{I}$.

Loss Type	(0.01, 0.0002)	(0.01, 0.02)	(0.05, 0.02)
non-IID	32.19/0.904	32.24/0.905	28.95/0.827
IID	33.14/0.920	32.89/0.916	29.29/0.840

Table 8: Ablation study on re-visible correlation. non-IID and IID indicate whether the masked and visible branches are independent.

Noise Type	$\lambda_f = 3$	$\lambda_f = 11$	$\lambda_f = 20$	$\lambda_f = 40$
(0.01, 0.0002)	33.05/0.919	33.14/0.920	33.11/0.920	33.06/0.919
(0.01, 0.02)	32.75/0.914	32.89/0.916	32.84/0.915	32.76/0.914
(0.05, 0.02)	29.24/0.838	29.29/0.840	29.27/0.839	29.23/0.837

Table 9: Ablation study on visible factor. Note that the initial value $\lambda_s = 3$ and λ_f is the final value.

is not proportional to the performance. Instead, the performance increases and decreases as the visible factor increases, reaching a peak when $\lambda_f = 11$.

6. Conclusion

We propose Blind2Sound, a self-supervised blind denoising framework that removes Poisson-Gaussian noise without residual noise and adapts to sensed noise levels. Adaptive re-visible loss associates noise parameters with re-visible transitions to achieve personalized and lossless denoising. Cramer Gaussian loss introduces fine-grained noise knowledge to improve estimation accuracy. The noise estimator is removed during inference. Extensive experiments show that our approach performs better, especially in real scenes with dynamic noise.

Acknowledgements

This work was supported by the grants from the STI 2030-Major Projects (2021ZD0204500, 2021ZD0204503), the National Natural Science Foundation of China (32171461).

References

- [1] Abdelrahman Abdelhamed. Website. <https://www.eecs.yorku.ca/~kamel/sidd/benchmark.php>. Accessed: 2018-11-01. 6
- [2] Abdelrahman Abdelhamed, Stephen Lin, and Michael S Brown. A high-quality denoising dataset for smartphone cameras. In *CVPR*, pages 1692–1700, 2018. 6
- [3] Francis J Anscombe. The transformation of poisson, binomial and negative-binomial data. *Biometrika*, 35(3/4):246–254, 1948. 2, 3
- [4] Saeed Anwar and Nick Barnes. Real image denoising with feature attention. In *ICCV*, pages 3155–3164, 2019. 1
- [5] Joshua Batson and Loic Royer. Noise2self: Blind denoising by self-supervision. In *ICML*, pages 524–533, 2019. 1, 2
- [6] Paul A Bromiley. Products and convolutions of gaussian distributions. *Medical School, Univ. Manchester, Manchester, UK, Tech. Rep*, 3:2003, 2003. 4
- [7] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T Barron. Unprocessing images for learned raw denoising. In *CVPR*, pages 11036–11045, 2019. 1
- [8] Jaeseok Byun, Sungmin Cha, and Taesup Moon. Fbi-denoiser: Fast blind image denoiser for poisson-gaussian noise. In *CVPR*, pages 5768–5777, 2021. 2, 3, 5, 6, 7
- [9] Jaeseok Byun and Taesup Moon. Learning blind pixel-wise affine image denoiser with single noisy images. *SPL*, 27:1105–1109, 2020. 2, 7
- [10] Sungmin Cha and Taesup Moon. Fully convolutional pixel adaptive image denoiser. In *ICCV*, pages 4160–4169, 2019. 2
- [11] Guangyong Chen, Fengyuan Zhu, and Pheng Ann Heng. An efficient statistical method for image noise level estimation. In *ICCV*, pages 477–485, 2015. 2, 3, 5
- [12] Jingwen Chen, Jiawei Chen, Hongyang Chao, and Ming Yang. Image blind denoising with generative adversarial network based noise modeling. In *CVPR*, pages 3155–3164, 2018. 1
- [13] Jun Cheng, Tao Liu, and Shan Tan. Score priors guided deep variational inference for unsupervised real-world single image denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12937–12948, 2023. 2
- [14] Hamadi Chihaoui and Paolo Favaro. Masked and shuffled blind spot denoising for real-world images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3025–3034, 2024. 2
- [15] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *TIP*, 16(8):2080–2095, 2007. 6, 7
- [16] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 6
- [17] Alessandro Foi, Mejdí Trimeche, Vladimir Katkovnik, and Karen Egiazarian. Practical poissonian-gaussian noise modeling and fitting for single-image raw-data. *TIP*, 17(10):1737–1754, 2008. 2, 3, 6
- [18] Rich Franzen. Kodak lossless true color image suite. *source: http://r0k.us/graphics/kodak*, 4(2), 1999. 6
- [19] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward convolutional blind denoising of real photographs. In *CVPR*, pages 1712–1722, 2019. 1
- [20] Samuel W Hasinoff. Photon, poisson noise. *Computer Vision: A Reference Guide*, pages 608–610, 2014. 3
- [21] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *CVPR*, pages 5197–5206, 2015. 6
- [22] Tao Huang, Songjiang Li, Xu Jia, Huchuan Lu, and Jianzhuang Liu. Neighbor2neighbor: Self-supervised denoising from single noisy images. In *CVPR*, pages 14781–14790, 2021. 2, 5, 6, 7
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [24] Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void-learning denoising from single noisy images. In *CVPR*, pages 2129–2137, 2019. 1, 2, 6, 7
- [25] Samuli Laine, Tero Karras, Jaakko Lehtinen, and Timo Aila. High-quality self-supervised deep image denoising. *NeurIPS*, 32:6970–6980, 2019. 2, 5, 6
- [26] Wooseok Lee, Sanghyun Son, and Kyoung Mu Lee. Ap-bsn: Self-supervised denoising for real-world images via asymmetric pd and blind-spot network. In *CVPR*, pages 17725–17734, 2022. 2
- [27] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2noise: Learning image restoration without clean data. In *ICML*, 2018. 1, 2, 6, 7
- [28] Xinhao Liu, Masayuki Tanaka, and Masatoshi Okutomi. Single-image noise level estimation for blind denoising. *TIP*, 22(12):5226–5237, 2013. 2
- [29] Xinhao Liu, Masayuki Tanaka, and Masatoshi Okutomi. Practical signal-dependent noise parameter estimation from a single noisy image. *TIP*, 23(10):4361–4371, 2014. 2, 6
- [30] Markku Makitalo and Alessandro Foi. Optimal inversion of the generalized anscombe transformation for poisson-gaussian noise. *TIP*, 22(1):91–103, 2012. 3
- [31] Youssef Mansour and Reinhard Heckel. Zero-shot noise2noise: Efficient image denoising without any data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14018–14027, 2023. 2
- [32] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, volume 2, pages 416–423, 2001. 6
- [33] Nick Moran, Dan Schmidt, Yu Zhong, and Patrick Coady. Noisier2noise: Learning to denoise from unpaired noisy data. In *CVPR*, pages 12064–12072, 2020. 2

- [34] Reyhaneh Neshatavar, Mohsen Yavartanoo, Sanghyun Son, and Kyoung Mu Lee. Cvf-sid: Cyclic multi-variate function for self-supervised image denoising by disentangling noise from image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17583–17591, 2022. 2
- [35] Tongyao Pang, Huan Zheng, Yuhui Quan, and Hui Ji. Recorrupted-to-recorrupted: Unsupervised deep learning for image denoising. In *CVPR*, pages 2043–2052, 2021. 2
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 32:8026–8037, 2019. 6
- [37] Stanislav Pyatykh, Jürgen Hesser, and Lei Zheng. Image noise level estimation by principal component analysis. *TIP*, 22(2):687–699, 2012. 2
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241, 2015. 1, 6, 7
- [39] Stefan Roth and Michael J Black. Fields of experts: A framework for learning image priors. In *CVPR*, volume 2, pages 860–867. IEEE, 2005. 6
- [40] Zejin Wang, Jiazheng Liu, Guoqing Li, and Hua Han. Blind2unblind: Self-supervised image denoising with visible blind spots. In *CVPR*, 2022. 2, 5, 6, 7
- [41] Xiaohe Wu, Ming Liu, Yue Cao, Dongwei Ren, and Wangmeng Zuo. Unpaired learning of deep image denoising. In *ECCV*, pages 352–368, 2020. 1, 2, 7
- [42] Jun Xu, Yuan Huang, Ming-Ming Cheng, Li Liu, Fan Zhu, Zhou Xu, and Ling Shao. Noisy-as-clean: learning self-supervised denoising from corrupted image. *TIP*, 29:9316–9329, 2020. 2
- [43] Zongsheng Yue, Hongwei Yong, Qian Zhao, Deyu Meng, and Lei Zhang. Variational denoising network: Toward blind noise modeling and removal. *NeurIPS*, 32:1690–1701, 2019. 1
- [44] 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*, pages 2696–2705, 2020. 1
- [45] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for real image restoration and enhancement. In *ECCV*, pages 492–511, 2020. 1
- [46] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *International conference on curves and surfaces*, pages 711–730, 2010. 6
- [47] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *TIP*, 26(7):3142–3155, 2017. 1, 6
- [48] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *TIP*, 27(9):4608–4622, 2018. 1
- [49] Yide Zhang, Yinhao Zhu, Evan Nichols, Qingfei Wang, Siyuan Zhang, Cody Smith, and Scott Howard. A poisson-gaussian denoising dataset with real fluorescence microscopy images. In *CVPR*, 2019. 6
- [50] Yunhao Zou, Chenggang Yan, and Ying Fu. Iterative denoiser and noise estimator for self-supervised image denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13265–13274, 2023. 2