

Score Priors Guided Deep Variational Inference for Unsupervised Real-World Single Image Denoising

Jun Cheng, Tao Liu, Shan Tan*

Huazhong University of Science and Technology, Wuhan, China

{jcheng24, hust_liutao, shantan}@hust.edu.cn

Abstract

*Real-world single image denoising is crucial and practical in computer vision. Bayesian inversions combined with score priors now have proven effective for single image denoising but are limited to white Gaussian noise. Moreover, applying existing score-based methods for real-world denoising requires not only the explicit train of score priors on the target domain but also the careful design of sampling procedures for posterior inference, which is complicated and impractical. To address these limitations, we propose a score priors-guided deep variational inference, namely ScoreDVI, for practical real-world denoising. By considering the deep variational image posterior with a Gaussian form, score priors are extracted based on easily accessible minimum MSE Non-*i.i.d* Gaussian denoisers and variational samples, which in turn facilitate optimizing the variational image posterior. Such a procedure adaptively applies cheap score priors to denoising. Additionally, we exploit a Non-*i.i.d* Gaussian mixture model and variational noise posterior to model the real-world noise. This scheme also enables the pixel-wise fusion of multiple image priors and variational image posteriors. Besides, we develop a noise-aware prior assignment strategy that dynamically adjusts the weight of image priors in the optimization. Our method outperforms other single image-based real-world denoising methods and achieves comparable performance to dataset-based unsupervised methods.*

1. Introduction

Image denoising is a fundamental task in low-level computer vision. Different from the additive white Gaussian noise (AWGN) that has been widely studied, real-world noise is typically signal-dependent and spatially correlated (or structured) and is prevalent in various image processing applications, including photography with sRGB noise [24, 29] and bioimaging with microscopy image noise [44].

The removal of real-world structured noise is therefore critical for the subsequent image analysis and understanding.

While recent deep learning-based supervised methods [62, 31, 60, 6] and unsupervised/self-supervised methods [40, 29, 66, 24] have shown promising results on real-world image denoising, they suffer from data collection difficulties and model generalization issues [65]. As an alternative, single image-based real-world denoising methods that rely on only a single image are gaining increasing attention as a more practical and promising solution. However, existing single image-based methods, e.g., [55, 4, 45, 67, 5, 65, 30], show unsatisfactory performance in real-world structured noise removal, which highlights the need for new and effective approaches.

Nowadays, deep generative models, represented by score-based methods [52, 15], have enabled powerful image priors to be learned from large-scale clean datasets. The learned score priors from score networks have demonstrated remarkable capacity in image generation and various image restorations [21, 52, 51, 47]. However, current score-based methods related to denoising, e.g., [17, 23, 21, 22, 13, 9], typically consider AWGN and cannot handle real-world noise. Furthermore, applying current score-based methods for real-world denoising not only requires the explicit train of score priors on the target domain but also has to devise sampling procedures elaborately for posterior inference. The accurate noise distribution is also demanded in the conditional sampling. Above restrictions makes them cumbersome and impractical for real-world noise removal.

To address these limitations, we propose an effective and practical score priors guided deep variational inference, namely ScoreDVI, for unsupervised real-world single image denoising. Rather than explicitly training score priors on the target domain, we relate the general score functions to the easily available minimum MSE (MMSE) Non-*i.i.d* Gaussian denoisers. We extract score priors by employing well-trained Non-*i.i.d* Gaussian denoisers to denoise samples from the variational image posterior parameterized by deep neural networks and then use the extracted scores to facilitate optimizing this variational image posterior, which in

*Corresponding author.

turn allows adaptive extraction of scores in the next round. This optimization procedure adaptively applies cheap score priors to denoising and frees the elaborate design of sampling strategies.

Instead of relying on the accurate noise model, we adopt the Non-*i.i.d* Gaussian mixture model (GMM) to roughly represent real-world noise, and a variational noise posterior is built and optimized to improve the modeling accuracy by means of hierarchical Bayesian inference. This modeling strategy also allows for the pixel-wise fusion of multiple image priors and variational image posteriors, further enhancing the denoising capability. Moreover, traditional variational inference treats image priors and likelihoods equally during optimization, which may not be optimal for inputs with varying levels of noise. Since noisier images require more priors, we propose a noise-aware prior assignment strategy that dynamically assigns different weights of image priors during optimization based on the estimated noise levels of noisy images. By doing so, our method can maximize performance in individual instances.

In total, our contributions are summarized as follows:

1. We propose an adaptive extraction and utilization of score priors based on Non-*i.i.d* Gaussian denoisers and deep variational inference for variational optimization.
2. We incorporate Non-*i.i.d* GMM likelihood model and the variational noise posterior to model the real-world noise, which enables pixel-wise fusion of multiple image priors and variational image posteriors.
3. We develop a noise-aware prior assignment strategy, which adaptively assigns the degree of image priors to the optimization objective for different noisy images.
4. To the best of our knowledge, we are the first to successfully remove real-world structured noise by applying score priors. Our method outperforms other single image-based methods on a variety of benchmark datasets and achieves competitive results against dataset-based unsupervised methods.

2. Related works

2.1. Real-world image denoising

Dataset-based unsupervised learning methods. To address the data collection difficulties of supervised approaches [61, 31, 60, 6], several unsupervised and self-supervised real-world denoising methods have been proposed. Blind-spot (BS) based strategies, e.g., Noise2self [4], Nei2nei [16], and BS networks [26, 56] as well as their zero-shot versions, are limited to handling pixel-independent noise. Noiser2noise [38], R2R [42], and IDR [63] can address structured noise, but they require accurate noise distribution, making them impractical.

Recently, [67] proposed using pixel-shuffle downsampling (PD) to break down the spatial correlation of structured noise. AP-BSN [29] used two PDs with different strides for training and testing. CVF-SID [40] attempted to separate the latent image and noise from the noisy input by a cyclic module and self-supervised losses. LUD-VAE [66] designed two hidden variables for the latent image domain and noise domain, respectively, and derived an optimizable loss for unpaired data. [24, 18, 49] explicitly learned the real-world noise distribution. These methods require noisy or unpaired data and generalization problems remain.

Single image-based real-world denoising. Deep image prior (DIP) [55] achieved single image denoising by regularizing the solution via CNN architectures. Self2Self [45] built training pairs from a single noisy image via Bernoulli sampling to train denoisers. [67] proposed a pixel-shuffle downsampling (PD) to break down the spatial correlation of real-world structured noise. StructN2V [5] estimated the noise correlation from a single image and developed a structured mask to apply N2V [25]. NN+denoiser [65] used conditional VAE to model the likelihood of real-world noise and combined plug-and-play priors to denoise. However, these methods show poor performance when removing severe structured noise. Instead, our ScoreDVI can denoise noisier images effectively and efficiently.

2.2. Image priors modeling

Classical image priors [27, 28, 35, 12] utilize low-level statistics of natural images to denoise images. Recently, complex image priors can be learned by using deep generative models and are then employed in Bayesian inversion to solve various image restoration tasks [14, 41, 21]. Particularly, score-based methods have achieved remarkable performance [21, 52, 47]. However, current research on image restoration with score priors [23, 22, 19, 51, 46, 8, 21, 13, 9] generally consider AWGN or noiseless inputs. And the direct application of these methods for real-world noise is complicated and impractical. Recently, [53] used the diffusion model to remove structured noise, while their method simply extended the existing paradigm and can only handle simple cases such as removing digits from face images. Instead, we propose to combine score priors with MMSE Gaussian denoisers and variational inference, which allows us to effectively and efficiently remove real-world noise.

2.3. Variational inference

Traditional variational inference-based methods [32, 2, 3] employ analytical image priors and cyclic optimization to obtain the approximate posterior, but they often suffer from poor performance. Recently, deep learning-based variational inference methods for real-world image denoising such as VDN [58], VNIR [59], and VDIR [50] construct variational posteriors via DNNs and utilize ground truths as

strong priors. However, these methods require paired training data. Instead, our ScoreDVI is unsupervised and can deal with structured noise by employing score priors.

3. Method

In this section, we present our unsupervised deep variational inference method for real-world single image denoising. We first introduce the variational inference framework in section 3.1, followed by the idea of adaptively applying score priors to facilitate the optimization of variational objective in section 3.2. The Non-*i.i.d* Gaussian mixture likelihood model and hyperpriors are presented in section 3.3. Section 3.4 shows the deep variational posteriors and the complete optimization objective. The noise-aware prior assignment is finally given in section 3.5.

3.1. Variational inference framework

In Bayesian statistics, image denoising is viewed as finding the posterior distribution $p(\mathbf{x}|\mathbf{y})$ and its estimators, e.g., posterior mean and MAP solution, where $\mathbf{x}, \mathbf{y} \in \mathbb{R}^N$ are latent clean image and noisy observation, respectively. $N = C \times H \times W$ is the image dimension. In practice, the posterior distribution $p(\mathbf{x}|\mathbf{y})$ is non-trivial and has no closed-form solution. The idea of variational inference is to construct a tractable distribution $q(\mathbf{x})$ and minimize its ‘distance’ to the true $p(\mathbf{x}|\mathbf{y})$. KL divergence is usually employed as the distance metric,

$$\begin{aligned} \text{KL}(q(\mathbf{x})||p(\mathbf{x}|\mathbf{y})) &= \int \log \frac{q(\mathbf{x})}{p(\mathbf{x}|\mathbf{y})} q(\mathbf{x}) d\mathbf{x} \\ &= \log p(\mathbf{y}) - \underbrace{(\mathbb{E}_{q(\mathbf{x})}(\log p(\mathbf{y}|\mathbf{x})) - \text{KL}(q(\mathbf{x})||p(\mathbf{x})))}_{\text{ELBO}} \end{aligned} \quad (1)$$

where $\log p(\mathbf{y})$ is the marginal likelihood (also called evidence), and $\mathbb{E}_{q(\mathbf{x})}(\log p(\mathbf{y}|\mathbf{x})) - \text{KL}(q(\mathbf{x})||p(\mathbf{x}))$ is the Evidence Low Bound (ELBO) as $\text{KL}(q(\mathbf{x})||p(\mathbf{x})) \geq 0$.

As $\log p(\mathbf{y})$ is generally not computable, the ELBO instead acts as the objective of variational inference. Hence, modeling the image prior $p(\mathbf{x})$, likelihood $p(\mathbf{y}|\mathbf{x})$, and variational posterior $q(\mathbf{x})$ involved in the ELBO is important.

3.2. Adaptive score priors for variational inference

In this section, we propose an adaptive extraction and application of score priors in variational inference.

Scores and MMSE Gaussian denoisers. [20] derived the relationship between scores and MMSE denoisers for *i.i.d* Gaussian noise. Here, we extend such a relationship to the case of general structured Gaussian.

Theorem 1 *Suppose a noisy observation $\tilde{\mathbf{x}} = \mathbf{x} + \mathbf{n} \in \mathbb{R}^N$ where $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \Sigma)$ has the zero mean and covariance Σ , and $\mathbf{x} \sim p(\mathbf{x})$. Then the score of $p(\tilde{\mathbf{x}})$ is*

$$\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}}) = \Sigma^{-1} (E_{p(\mathbf{x}|\tilde{\mathbf{x}})}(\mathbf{x}) - \tilde{\mathbf{x}}) \quad (2)$$

where $E_{p(\mathbf{x}|\tilde{\mathbf{x}})}(\mathbf{x})$ is the conditional mean for input $\tilde{\mathbf{x}}$ corrupted by general Gaussian noise.

The proof of Theorem 1 is given in Supplementary Material S1.1. If we further consider the noise covariance has diagonal form $\Sigma = \text{diag}(\sigma^2)$, Eq. (2) becomes

$$\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}}) = \frac{1}{\sigma^2} \odot (E_{p(\mathbf{x}|\tilde{\mathbf{x}})}(\mathbf{x}) - \tilde{\mathbf{x}}) \quad (3)$$

where \odot denotes element-wise multiplication.

In practice, the posterior mean $E_{p(\mathbf{x}|\tilde{\mathbf{x}})}(\mathbf{x})$ in Eq. (3) can be approximated by the output of a deep MMSE Non-*i.i.d* (independent but not identical) Gaussian denoiser $\mathcal{G}(\tilde{\mathbf{x}})$. Moreover, a CNN-based blind MMSE Gaussian denoiser trained over AWGN with a wide range of noise variance has the ability to handle such Non-*i.i.d* noise as the local connectivity property of CNNs. Therefore, score priors are directly available from well-trained blind MMSE Gaussian denoisers, freeing the explicit training of score networks.

Scores facilitate ELBO optimization. Within the ELBO objective in Eq. (1), the image prior $p(\mathbf{x})$ is only involved in the KL divergence term of the ELBO, which can be approximated as

$$\begin{aligned} \text{KL}(q(\mathbf{x})||p(\mathbf{x})) &= \mathbb{E}_{q(\mathbf{x})}(\log q(\mathbf{x}) - \log p(\mathbf{x})) \\ &\approx \mathbb{E}_{q(\mathbf{x})} \log q(\mathbf{x}) - \frac{1}{M} \sum_{m=1}^M \log p(\mathbf{x}_m) \end{aligned} \quad (4)$$

where the Monte Carlo (MC) integration with M samples is used to approximate $\mathbb{E}_{q(\mathbf{x})} \log p(\mathbf{x})$, and $\mathbf{x}_m \sim q(\mathbf{x})$.

If we take the derivative of $\text{KL}(q(\mathbf{x})||p(\mathbf{x}))$ in Eq. (4) with regard to \mathbf{x} , then the score function appears. In order to incorporate the score prior from Gaussian denoisers, we consider that $q(\mathbf{x})$ follows diagonal Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \text{diag}(\boldsymbol{\sigma}^2))$, and then the derivatives of $\mathbb{E}_{q(\mathbf{x})} \log p(\mathbf{x})$ with regard to variational posterior parameters $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$ can be computed as

$$\begin{aligned} \nabla_{\boldsymbol{\mu}} \mathbb{E}_{q(\mathbf{x})} \log p(\mathbf{x}) &\approx \frac{1}{M} \sum_{m=1}^M (\mathcal{G}(\mathbf{x}_m) - \mathbf{x}_m) \odot \frac{1}{\boldsymbol{\sigma}^2} \\ \nabla_{\boldsymbol{\sigma}^2} \mathbb{E}_{q(\mathbf{x})} \log p(\mathbf{x}) &\approx \frac{1}{M} \sum_{m=1}^M (\mathcal{G}(\mathbf{x}_m) - \mathbf{x}_m) \odot \frac{1}{\boldsymbol{\sigma}^2} \odot \boldsymbol{\epsilon} \end{aligned} \quad (5)$$

$$\mathbf{x}_m = \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\epsilon}, \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, I)$$

The derivation of Eq. (5) is given in Supplementary Material S1.2. Eq. (5) reveals that there is no need to compute $\mathbb{E}_{q(\mathbf{x})} \log p(\mathbf{x})$ explicitly, instead, we can utilize the score prior from MMSE Non-*i.i.d* Gaussian denoisers to perform the gradient-based optimization of $\text{KL}(q(\mathbf{x})||p(\mathbf{x}))$ in the ELBO objective. Moreover, the extraction of scores from MMSE Gaussian denoisers is adaptive, and we do

not need to explicitly schedule the noise variance and design the sampling strategy as done in score-based generative models [52, 22, 21]. That is, the score for $\mathbf{x}_m(t) \sim \mathcal{N}(\boldsymbol{\mu}(t), \boldsymbol{\sigma}^2(t))$ in the t -th iteration assists in the update of $\boldsymbol{\mu}(t+1)$ and $\boldsymbol{\sigma}^2(t+1)$, which in turn helps extract the score for $\mathbf{x}_m(t+1)$ dynamically in the $(t+1)$ -th iteration.

3.3. Non-*i.i.d* Gaussian mixture likelihood model

The real-world structured noise is generally signal-dependent and spatially correlated, which differs greatly from AWGN. In order to incorporate our knowledge about real-world noise, we utilize the Non-*i.i.d* Gaussian mixture model (GMM) to model the likelihood function, which is

$$p(\mathbf{y}|X, \Phi, \Omega) = \prod_{i=1}^N \sum_{k=1}^K \omega_{ik} \mathcal{N}(x_{ik}, \phi_{ik}^{-1}), \sum_{k=1}^K \omega_{ik} = 1 \quad (6)$$

where $X = [\mathbf{x}_1, \dots, \mathbf{x}_K]$, $\Phi = [\phi_1, \dots, \phi_K]$, $\Omega = [\omega_1, \dots, \omega_K] \in \mathbb{R}^{N \times K}$ are matrices with K columns, respectively. x_{ik} in X , ϕ_{ik} in Φ , and ω_{ik} in Ω denote the mean, precision, and mixing parameter of the k -th Gaussian component, respectively, for modeling i -th pixel y_i of \mathbf{y} . K is the total number of mixtures.

In Eq. (6), each noisy pixel y_i is represented by a GMM, so as to better model the spatially variant characteristic of real-world noise. Besides, such consideration allows the ensemble of different image priors and variational image posteriors in a pixel-wise manner, which will be introduced in the following section. To facilitate optimization, we refer to [7] and reparameterize the likelihood function in Eq. (6) as

$$\begin{aligned} p(\mathbf{y}|X, Z, \Phi) &= \prod_{i=1}^N \prod_{k=1}^K \mathcal{N}(x_{ik}, \phi_{ik}^{-1})^{z_{ik}} \\ p(Z|\Omega) &= \prod_{i=1}^N \prod_{k=1}^K \omega_{ik}^{z_{ik}}, \sum_{k=1}^K z_{ik} = 1 \end{aligned} \quad (7)$$

where $Z = [z_1, \dots, z_K] \in \mathbb{R}^{N \times K}$ is an auxiliary matrix variable and each row of Z follows a categorical distribution with one item one and the others zero.

Hyperpriors for hyperparameters. Different from existing score-based methods which carry out posterior sampling based on accurate likelihood functions with fixed distribution parameters, we view Φ and Ω in Eq. (7) as random variables and infer them by introducing hyperpriors based on hierarchical Bayesian inference. In order to simplify the ELBO objective, we choose the following conjugate priors for Φ and Ω , respectively

$$\begin{aligned} p(\Phi|Z) &= \prod_{i=1}^N \prod_{k=1}^K \text{Gamma}(\phi_{ik}; \alpha, \beta)^{z_{ik}} \\ p(\Omega) &= \prod_{i=1}^N Z_{\text{Dir}}(\mathbf{d}) \prod_{k=1}^K \omega_{ik}^{d_k-1}, \mathbf{d} \in \mathbb{R}^K \end{aligned} \quad (8)$$

where Gamma denotes the Gamma distribution with parameters α and β . Each row of Ω follows the Dirichlet distribution with parameter \mathbf{d} and normalization coefficient $Z_{\text{Dir}}(\cdot)$.

3.4. Deep modeling of variational posteriors

The modeling of likelihood function in Eq. (7) enables us to build the factorized image prior conditioned on Z as

$$p(X|Z) = \prod_{i=1}^N \prod_{k=1}^K p(x_{ik})^{z_{ik}} \quad (9)$$

which enables the image prior $p(X)$ to be modeled by fusing many priors in a pixel-wise manner.

Considering Eqs. (7), (8) and (9), the posterior distribution we want to infer now is $p(X, \Phi, Z, \Omega|\mathbf{y})$, while it is non-trivial. Therefore, we propose a powerful and trivial deep variational posterior $q(X, \Phi, Z, \Omega)$ that approximates the true $p(X, \Phi, Z, \Omega|\mathbf{y})$.

Specifically, we factorize the variational posterior $q(X, \Phi, Z, \Omega)$ as $q(X|Z)q(\Phi|Z)q(Z)q(\Omega)$ with each item

$$\begin{aligned} q(X|Z) &= \prod_{i=1}^N \prod_{k=1}^K \mathcal{N}(\mu_{ik}, \sigma_{ik}^2)^{z_{ik}} \\ q(\Phi|Z) &= \prod_{i=1}^N \prod_{k=1}^K \text{Gamma}(\phi_{ik}; \hat{\alpha}_{ik}, \hat{\beta}_{ik})^{z_{ik}} \\ q(Z) &= \prod_{i=1}^N \prod_{k=1}^K \pi_{ik}^{z_{ik}}, q(\Omega) = \prod_{i=1}^N Z_{\text{Dir}}(\hat{\mathbf{d}}_i) \prod_{k=1}^K \omega_{ik}^{\hat{d}_{ik}-1} \end{aligned} \quad (10)$$

All involved parameters of variational posteriors in Eq. (10), i.e., $\Theta = \{\mu_{ik}, \sigma_{ik}^2, \hat{\alpha}_{ik}, \hat{\beta}_{ik}, \pi_{ik}, \hat{d}_{ik} | i \in \{1, \dots, N\}, k \in \{1, \dots, K\}\}$ are modeled by CNNs as they provide good regularization for Θ . Note that $q(\Phi)$ is variational noise posterior. As shown in Figure 1, we build four CNNs, i.e., X -net, Φ -net, Z -net and Ω -net, which have \mathbf{y} as the input and output variational posterior parameters Θ .

Combing Eq. (7), Eq. (9), Eq. (8) and Eq. (10), we can derive the complete optimization objective, which is

$$\mathcal{L} = -\text{ELBO} \quad (11a)$$

$$= -\mathbb{E}_{q(X, \Phi, Z)} (\log p(\mathbf{y}|X, \Phi, Z)) \rightarrow \mathcal{L}_1 \quad (11b)$$

$$+ \mathbb{E}_{q(Z)} (\text{KL}(q(X|Z)||p(X|Z))) \rightarrow \mathcal{L}_2 \quad (11c)$$

$$+ \mathbb{E}_{q(Z)} (\text{KL}(q(\Phi|Z)||p(\Phi|Z))) \rightarrow \mathcal{L}_3 \quad (11d)$$

$$+ \mathbb{E}_{q(\Omega)} (\text{KL}(q(Z)||p(Z|\Omega))) \rightarrow \mathcal{L}_4 \quad (11e)$$

$$+ \text{KL}(q(\Omega)||p(\Omega)) \rightarrow \mathcal{L}_5 \quad (11f)$$

with each component as follows:

$$\begin{aligned} \mathcal{L}_1 &= -\sum_{i=1}^N \sum_{k=1}^K \pi_{ik} \left(-\frac{((\mu_{ik} - y_i)^2 + \sigma_{ik}^2) \hat{\alpha}_{ik}}{2\hat{\beta}_{ik}} \right. \\ &\quad \left. + \frac{\psi(\hat{\alpha}_{ik}) - \log \hat{\beta}_{ik}}{2} \right) \end{aligned} \quad (12)$$

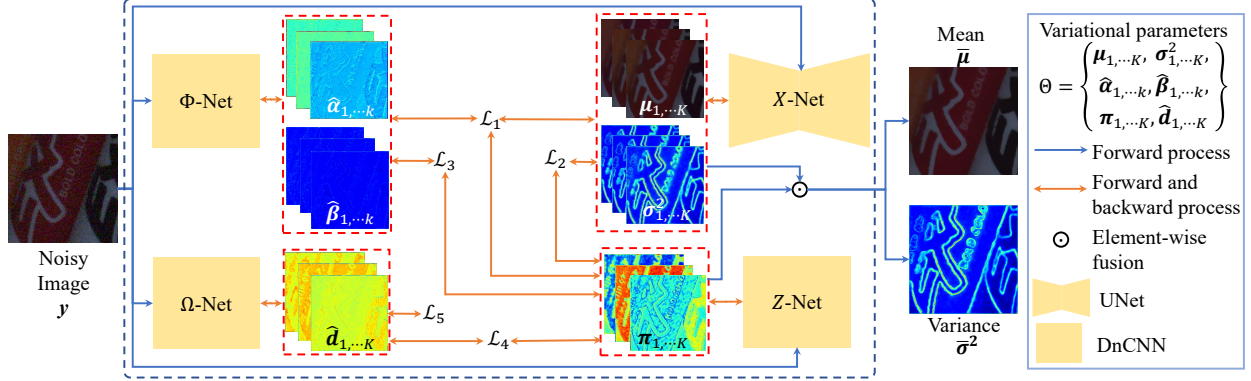


Figure 1: Overview of the proposed ScoreDVI. Our method receives a noisy image \mathbf{y} and outputs the denoised image $\bar{\boldsymbol{\mu}}$ as well as its variance $\bar{\boldsymbol{\sigma}}^2$ when the optimization is done. The four networks are used to output the variational parameters Θ .

$$\mathcal{L}_2 = \sum_{i=1}^N \sum_{k=1}^K \left(-\frac{1}{2} \pi_{ik} \log \sigma_{ik}^2 - \pi_{ik} \mathbb{E}_{q(x_{ik})} \log p(x_{ik}) \right) \quad (13)$$

$$\mathcal{L}_3 = \sum_{i=1}^N \sum_{k=1}^K \pi_{ik} \left(\log \frac{\Gamma(\alpha)}{\Gamma(\hat{\alpha}_{ik})} + (\hat{\alpha}_{ik} - \alpha) \psi(\hat{\alpha}_{ik}) \right. \\ \left. + \alpha \log \frac{\hat{\beta}_{ik}}{\beta} + (\beta - \hat{\beta}_{ik}) \frac{\hat{\alpha}_{ik}}{\hat{\beta}_{ik}} \right) \quad (14)$$

$$\mathcal{L}_4 = \sum_{i=1}^N \sum_{k=1}^K \pi_{ik} \left(\log \pi_{ik} - \psi(\hat{d}_{ik}) + \psi\left(\sum_{k=1}^K \hat{d}_{ik}\right) \right) \quad (15)$$

$$\mathcal{L}_5 = \sum_{i=1}^N \sum_{k=1}^K (\hat{d}_{ik} - d_k) (\psi(\hat{d}_{ik}) - \psi\left(\sum_{k=1}^K \hat{d}_{ik}\right)) \\ + \sum_{i=1}^N \log \frac{Z_{\text{Dir}}(\hat{\mathbf{d}}_i)}{Z_{\text{Dir}}(\mathbf{d})} \quad (16)$$

where $\Gamma(\cdot)$ and $\psi(\cdot)$ denote gamma function and digamma function, respectively; For detailed derivations, please refer to Supplementary Material S1.3.

Note that $\mathcal{L}_1, \mathcal{L}_3, \mathcal{L}_4, \mathcal{L}_5$ and $-\frac{1}{2} \pi_{ik} \log \sigma_{ik}^2$ in \mathcal{L}_2 are all directly computable and analytically derivable. Regarding $\mathbb{E}_{q(x_{ik})} \log p(x_{ik})$ in \mathcal{L}_2 , due to the factorized $p(X|Z)$ and $q(X|Z)$, we have

$$\nabla_{\mu_{ik}} \mathbb{E}_{q(x_{ik})} \log p(x_{ik}) = \left(\nabla_{\boldsymbol{\mu}_k} \mathbb{E}_{q(\mathbf{x}_k)} \log p(\mathbf{x}_k) \right)_i \\ \nabla_{\sigma_{ik}^2} \mathbb{E}_{q(x_{ik})} \log p(x_{ik}) = \left(\nabla_{\boldsymbol{\sigma}_k^2} \mathbb{E}_{q(\mathbf{x}_k)} \log p(\mathbf{x}_k) \right)_i \quad (17)$$

which can then be directly computed by using Eq. (5).

The diagram and the whole pipeline of the proposed ScoreDVI are illustrated in Figure 1 and Algorithm 1, respectively. Once the optimization of Eq. (11) is done, we

Algorithm 1 Algorithm of the Proposed ScoreDVI

Input: GMM number K , MC samples M , MMSE Gaussian denoisers $\mathcal{G}_1, \dots, \mathcal{G}_K$, X-net, Phi-net, Omega-net, total iterations T , noisy image \mathbf{y} , $t = 1$
Output: denoised image and its variance

- 1: **while** $t \leq T$ **do**
- 2: Compute Θ from X-net, Phi-net, Z-net and Omega-net
- 3: **for** $k \leq K$ **do**
- 4: Compute the derivatives of $\mathbb{E}_{q(x_{ik})} \log p(x_{ik})$ in \mathcal{L}_2 to $\boldsymbol{\mu}_k(t)$ and $\boldsymbol{\sigma}_k^2(t)$ based on Eq. (5), Eq. (17) and \mathcal{G}_k
- 5: Compute $\mathcal{L}_1, \mathcal{L}_3, \mathcal{L}_4, \mathcal{L}_5$ and the first item in \mathcal{L}_2
- 6: Compute the derivatives of Θ with regard to weights of four CNNs ▷ This is done in autograd
- 7: Update model weights with the optimizer
- 8: **end for**
- 9: **end while**
- 10: **return** $\bar{\boldsymbol{\mu}}(T)$ and $\bar{\boldsymbol{\sigma}}^2(T)$ based on Eq. (18)

can obtain the mean and variance of $q(X)$ from $q(X|Z)$ and $q(Z)$:

$$\bar{\boldsymbol{\mu}} = \sum_{k=1}^K \pi_k \odot \boldsymbol{\mu}_k, \bar{\boldsymbol{\sigma}}^2 = \sum_{k=1}^K \pi_k^2 \odot \boldsymbol{\sigma}_k^2 \quad (18)$$

which turn out to be the pixel-wise fusion of different variational image posteriors \mathbf{x}_k . $\bar{\boldsymbol{\mu}}$ is then utilized as the final denoised image and $\bar{\boldsymbol{\sigma}}^2$ represents the uncertainty of $\bar{\boldsymbol{\mu}}$.

3.5. Noise-aware prior assignment

Observe that in the loss function \mathcal{L} derived from standard VI, the image prior $p(\mathbf{x})$ and the likelihood are given equal weight. In practice, however, we expect the weighting between these two terms to depend on the specific input \mathbf{y} . For images with less noise, we would like the reconstruction results to be more biased toward the actual observations. On the other hand, for noisier images, we expect the restorations to be more dependent on the prior information. Therefore, we propose to adaptively adjust the weighting between

the image priors and likelihoods based on the level of noise present in the input image. Specifically, we scale \mathcal{L}_2 to

$$\mathcal{L}_2 = \lambda(\mathbf{y}) \mathbb{E}_{q(Z)} (\text{KL}(q(X|Z)||p(X|Z))) \quad (19)$$

where $\lambda(\mathbf{y})$ is a unary weight function depending on \mathbf{y} .

To roughly estimate the noise level of \mathbf{y} , we use the noise estimation method presented in [33] for generalized signal-dependent noise. Since real-world noise is spatially correlated, we first split \mathbf{y} into down-scaled sub-images using a PD operation with a stride of 4. We then estimate the noise variance of local patches in these sub-images, following [33], and retain the patches with weak texture. The average noise standard deviation δ is then calculated from the remaining patches, providing a rough indication of the noise level presented in \mathbf{y} .

Based on δ , we make a simple schedule of $\lambda(\mathbf{y})$ to assign different beliefs of image priors to \mathcal{L} in Eq. (11), that is

$$\lambda(\mathbf{y}) = \begin{cases} \frac{1}{\gamma} & \text{if } \delta < l_1 \\ 1 & \text{else if } l_1 \leq \delta < l_2 \\ \gamma & \text{else if } \delta \geq l_2 \end{cases} \quad (20)$$

4. Experiments

In this section, we first introduce the experimental settings. Then we present qualitative and quantitative results of our method as well as comparisons with other methods. The ablation study is carried out in the last.

4.1. Experimental settings

Dataset. We evaluate our ScoreDVI on four widely-used real-world noise datasets, namely SIDD [1], PolyU [57], CC [39], and FMDD [64]. The SIDD dataset contains natural sRGB images from smartphones, and we evaluate our method using the SIDD validation and benchmark datasets, each of which consists of 1280 patches with size $3 \times 256 \times 256$. The PolyU and CC consist of 100 and 15 natural images, respectively, taken from diverse commercial camera brands. Each image in these datasets has a size of $3 \times 512 \times 512$. FMDD contains fluorescence microscopy images, and we evaluate our method on the mixed test set with raw images, which have an image size of 512×512 .

Implementation of MMSE Non-*i.i.d* Gaussian denoisers. We obtain MMSE Non-*i.i.d* Gaussian denoisers using the bias-free (BF) [37] version of DnCNN [62] as the network architecture. We train multiple blind BF-DnCNNs over different clean datasets, including CBS300 [36], DIV2K [54], ImageNet [11] validation dataset, and Waterloo Exploration Database [34], to implicitly learn diverse score priors for sRGB denoising. For each dataset, we add AWGN with noise standard deviation of $[0, 100]$ to clean patches to construct noisy/clean pairs. The BF-DnCNN is then optimized with the MMSE objective using these paired

data. We train BF-DnCNNs on corresponding grayscale images to capture scores for microscopy image denoising.

Implementation of variational posteriors. We construct four CNNs to represent the parameters Θ in Eq. (10). For the X -net, we use the Unet [48] architecture with skip connections, which outputs two maps, each with a shape of (K, C, H, W) , to denote $\mu_{1, \dots, K}$ and $\sigma_{1, \dots, K}^2$, respectively. For $\hat{\alpha}_{1, \dots, K}$ and $\hat{\beta}_{1, \dots, K}$, we build upon the DnCNN with 5 convolution layers as done in VDN [58] and have the Φ -net output two maps with the same shape as μ_k and σ_k^2 . Regarding Z -net and Ω -net, we use the same architectures as Φ -net to output one map each to represent $\hat{d}_{1, \dots, K}$ and $\omega_{1, \dots, K}$. Details of these networks can be found in Supplementary Material S2.

Selection of hyperprior parameters. We set $\alpha = 1, \beta = 0.02$ for the SIDD and FMDD datasets, which are known to contain highly noisy images. For CC and PolyU datasets, which have medium and low noise levels respectively, we set $\alpha = 1, \beta = 0.01$, and $\alpha = 1, \beta = 0.005$, respectively. We set \mathbf{d} as an all-one vector in Eq. (6), to assign the same prior weight to each component.

Optimization details. During optimization, all losses in \mathcal{L} except \mathcal{L}_2 can be expressed as scalar functions in Pytorch [43]. As for \mathcal{L}_2 , we implement its gradient backpropagation by using a customized *autograd* unit inherited from *torch.autograd.Function*. Its backward function computes and returns the gradients based on Eqs. (5) and (17). Incorporating this into the total loss \mathcal{L} enables end-to-end training of ScoreDVI in PyTorch. We use the ADAM optimizer with a fixed learning rate of 10^{-3} for all four networks. In Eq. (20), we set $l_1 = 10, l_2 = 25$, and $\gamma = 2$. We choose $K = 3, M = 5$, and perform the optimization with a total of 400 iterations for each noisy input. We evaluate the denoising quality using peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) metrics. All denoising experiments are conducted on an Nvidia 2080 GPU.

4.2. Evaluation on real-world noise

Comparisons with single image-based methods. We compare our method against several single image-based methods, including BM3D [10], DIP [55], self2self [45], PD-denoising [67], NN+denoiser [65], and APBSN-single [29]. 1) For DIP, we choose the same Unet architecture as our X -net. Due to its problem of overfitting, we choose the optimal denoising results with the reference to ground truth (GT) images. Hence, DIP in our experiments only serves as a reference. 2) For APBSN-single, we adapt APBSN proposed in [29] to directly denoise a single image. The strides of PD in training and testing are 5 and 2, respectively. 3) For NN+denoiser [65], we choose its best version, i.e. NN+BM3D for single image denoising. We re-evaluate the performance of NN+BM3D on FMDD, PolyU, and CC without the reference of GTs based on their source code. 4)

Table 1: Quantitative comparisons (PSNR(dB)/SSIM) of our ScoreDVI and other real-world denoising methods including single image-based methods and dataset-based methods on SIDD, FMDD, PolyU, and CC datasets. The best and second-best PSNR/SSIM results are marked in **bold** and underlined in each denoising category.

Category	Method	SIDD validation	SIDD benchmark	FMDD	PolyU	CC
Single image	BM3D [10]	25.65/0.475	25.65/0.685	30.06/0.771	37.40/0.953	35.19/0.858
	DIP [55]	32.11/0.740	–	32.90/0.854	37.17/0.912	35.61/0.912
	Self2Self [45]	29.46/0.595	29.51/0.651	30.76/0.695	38.33/0.962	37.44/0.948
	PD-denoising [67]	<u>33.97/0.820</u>	<u>33.61/0.894</u>	<u>33.01/0.856</u>	37.04/0.940	35.85/0.923
	NN+denoiser [65]	–	33.18/0.895	32.21/0.831	37.66/0.956	36.52/0.943
	APBSN-single [29]	30.90/0.818	30.71/0.869	28.43/0.804	29.61/0.897	27.72/0.891
	ScoreDVI (Ours)	34.75/0.856	34.60/0.920	33.10/0.865	<u>37.77/0.959</u>	<u>37.09/0.945</u>
Noisy or Unpaired images	APBSN [29]	–	36.91/0.931	<u>31.99/0.836</u>	<u>37.03/0.951</u>	<u>34.88/0.925</u>
	CVF-SID [40]	<u>34.81/0.944</u>	34.71/0.917	32.73/0.843	35.86/0.937	33.29/0.913
	LUD-VAE [66]	34.91/0.892	<u>34.82/0.926</u>	–	<u>36.99/0.955</u>	35.48/0.941

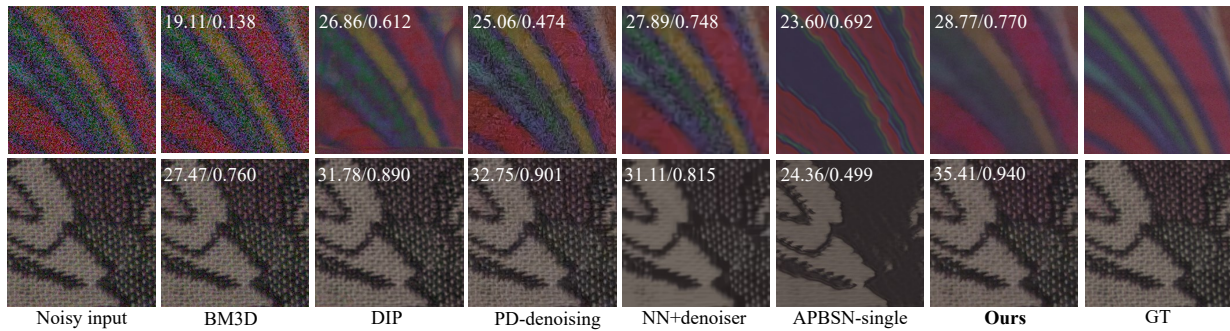


Figure 2: Visual comparison of our method against other single image-based denoising methods in SIDD validation dataset. The PSNR/SSIM results are listed within the images.




Figure 3: Visual comparison of our method against other single image-based denoising methods in PolyU dataset.

For the remaining methods, we either use the authors’ code or directly adopt their published results if available. We summarize the quantitative comparisons in Table 1. Qualitative comparisons between different single image-based methods on SIDD, FMDD, and PolyU are shown in Figures 2, 4 and 3, respectively. More visual comparisons can be found in Supplementary Material S4.

Our ScoreDVI achieves the best quantitative and qualitative performance among other single image-based methods on the SIDD and FMDD, which contain noisier images. PD-denoising can effectively remove noise, but weedy ar-

tifacts are frequently observed in restored images, as indicated in Figures 2, 4, and 3. Due to the large stride of PD used in training, APBSN-single severely destroys the details of original images and leaves clear color artifacts, as shown in Figures 2 and 3, or severe blur as shown in Figure 4. NN+denoiser introduces over-smoothing while removing noise. Self2Self performs badly on SIDD and FMDD, and cannot completely eliminate structured noise. In contrast, our method can both denoise noisy images and retain more details, thus yielding better results. Regarding the CC and PolyU datasets, Self2Self achieves the best PSNR with the




Figure 4: Visual comparison of single image-based denoising methods in FMDD. PSNR/SSIM are evaluated on the whole image. The noisy patch is from Confocal_BPAE_B_4.

cost of hundreds of thousands of training iterations. It also brings over-smoothed reconstructions, as shown in Figure 3. Instead, our ScoreDVI is quite efficient while also preserving fine details. Note that our method surpasses the optimal DIP among all datasets, as shown in Table 1.

Comparisons with dataset-based methods. Our ScoreDVI is also compared against several dataset-based unsupervised real-world denoising methods, including CVF-SID [40], APBSN [29], and LUD-VAE [66]. APBSN and CVF-SID are directly trained on the test datasets, while LUD-VAE is trained on unpaired data where the clean data is from the SIDD small dataset and the noisy data are from the test datasets. We present the quantitative results in Table 1, and the visual comparisons can be found in Supplementary Material. Our method shows competitive results compared to dataset-based methods among these datasets. While the results of dataset-based methods outperform ours on SIDD validation and benchmark, their performance degrades with the decrease of available noisy images in datasets such as PolyU, CC, and FMDD. In this case, they can only utilize fewer available images to train their networks. In contrast, our method requires only a single noisy image and performs consistently over all datasets, making it more practical for real-world applications.

Evaluating the existing score-based method for real-world denoising. We evaluate one typical score-based method by combining techniques from [52] and [46] for real-world denoising and present the result in Table 2. We first employ the NCSN++ network and VESDE strategy proposed in [52] to train scores on ImageNet validation. To obtain posterior samples, we use the predictor-corrector sampler with 1000 steps, where the unconditional score in the original corrector is replaced with the posterior score for Gaussian denoising presented in [46]. For each input, we obtain eight posterior samples and evaluate their means. By comparing the results in Table 1 and 2, we observe that our ScoreDVI outperforms this simple adoption of score-based

Table 2: Performance of the score-based method by combining [52] and [46] on SIDD validation and CC.

Dataset	SIDD validation	CC
PSNR/SSIM	32.30/0.769	36.18/0.936

Table 3: Effects of Gaussian mixture numbers K and prior assignment coefficient γ on SIDD validation dataset.

(a) $\gamma = 1$		(b) $K = 3$	
K	PSNR/SSIM	γ	PSNR/SSIM
1	33.20/0.789	1	33.79/0.807
2	33.68/0.808	4/3	34.55/0.852
3	33.79/0.807	2	34.75/0.856
4	33.87/0.815	4	33.83/0.831

methods for real-world structured noise by large margins.

4.3. Ablation study

Here, we conduct ablation studies on SIDD validation dataset to better evaluate the performance of our method. **Ablation on Gaussian numbers K .** Table 3a presents the results of our method with different K on the SIDD validation. As shown in Table 3a, the denoising performance consistently improves as we increase K . Additionally, Figure 5 shows that the ensemble of multiple means, μ_k , of the variational image posteriors, achieves better performance compared to using each single μ_k . Considering both performance and efficiency, we set $K = 3$ in our method.

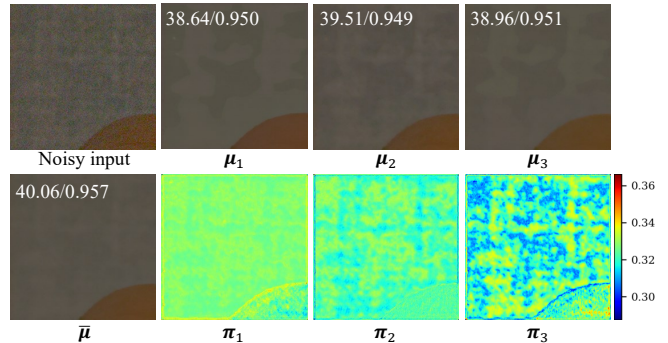


Figure 5: Visualization of different means μ_k and π_k for a noisy sRGB image of SIDD validation dataset.

Ablation on prior assignment coefficient γ . Table 3b presents the results of our method with different γ . From Table 3b and Figure 6, we have the following observations: when $\gamma = 1$, it means no adaptive prior assignment, and the corresponding denoised images either contain much noise for the noisier image (the first row of Figure 6) or are over-smooth for the less noisy image (the second row of Figure 6); with the increase of γ , the noise is gradually removed, and details emerge; when γ becomes too large, either image




Figure 6: Visualization of our method under different γ for noisier (first row) and less noisy (second row) images.

Table 4: Effects of deep modeling of Θ (A), image-wise fusion (B), and pixel-wise fusion (C) on SIDD validation.

A	B	C	PSNR/SSIM
		✓	31.76/0.767
✓	✓		34.34/0.850
✓		✓	34.75/0.856

priors or actual observations dominate the optimization, and the denoised images show artifacts or become noisy. Therefore, we set $\gamma = 2$ for both clean and sharp results.

Ablation on the modeling of parameters Θ . We evaluate the effect of using CNNs to model Θ on our method’s performance. We compare our method and the one that directly uses optimizable Θ initialized from *rand* Pytorch tensors to minimize \mathcal{L} . The results in the first and third row of Table 4 demonstrate the effectiveness of deep modeling of Θ .

Ablation on the ensemble strategy of image priors and posteriors. We evaluate the effectiveness of the pixel-wise fusion in Eq. (18) by comparing it with an image-level fusion strategy that bases on the image-level mixing of GMM for the likelihood. The final denoised image is hence the image-wise fusion of variational image posteriors. Details of this strategy are presented in Supplementary Material S3. The second and third row in Table 4 shows that the pixel-wise fusion achieves better results.

4.4. Efficiency comparisons and empirical convergence analysis

We present the inference time, model parameters, and FLOPs of the compared deep learning-based methods in Table 5. By comparing Table 1 and Table 5, our method achieves a balance between effectiveness and efficiency.

We ablate different iteration numbers and show the corresponding results (PSNR/SSIM) in Table 6. As shown, increasing the iteration number results in the convergence of performance and does not show overfitting issues. We find that 400 iterations are sufficient to achieve a balance between performance and efficiency *on average*, but it may be not the optimal number for each input.

Table 5: Efficiency comparisons of deep learning-based methods under the input size $256 \times 256 \times 3$

Method	Infer. time (s)	Params (M)	FLOPs (G)
DIP	146.2	13.4	31.06
Self2Self	3546.5	1.0	9.55
PD-denoising	0.36	0.7	46.94
NN+denoiser	897.6	13.4	31.06
APBSN-single	121.4	3.66	234.63
ScoreDVI	81.2	13.5	37.87

Table 6: Denoising performance of the proposed ScoreDVI with regard to different iteration numbers

Iters.	SIDD Val	CC
300	34.46/0.847	36.91/0.943
400	34.75/0.856	37.09/0.945
500	34.78/0.861	37.13/0.946
600	34.75/0.861	37.12/0.946

5. Conclusion

In this paper, we propose a novel unsupervised deep variational inference method that incorporates score priors embedded in MMSE Non-*i.i.d* denoisers and the Non-*i.i.d* Gaussian mixture likelihood model for real-world single image denoising. A noise-aware priors assignment strategy is developed to further improve denoising performance. Our proposed method shows superior performance compared with other single image-based real-world denoising methods on a variety of benchmark datasets. Moreover, our approach only requires a single noisy image and is more generalizable than dataset-based un/self-supervised real-world denoising methods.

6. Acknowledgment

This work was supported in part by the National Natural Science Foundation of China (NNSFC), under Grant Nos. 61672253 and 62071197.

References

- [1] Abdelrahman Abdelhamed, Stephen Lin, and Michael S Brown. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1692–1700, 2018. 6
- [2] S. Derin Babacan, Rafael Molina, and Aggelos K. Katsaggelos. Variational bayesian blind deconvolution using a total variation prior. *IEEE Transactions on Image Processing*, 18(1):12–26, 2009. 2
- [3] S. Derin Babacan, Rafael Molina, and Aggelos K. Katsaggelos. Variational bayesian super resolution. *IEEE Transactions on Image Processing*, 20(4):984–999, 2011. 2

- [4] Joshua Batson and Loic Royer. Noise2self: Blind denoising by self-supervision. In *International Conference on Machine Learning*, pages 524–533. PMLR, 2019. 1, 2
- [5] Coleman Broaddus, Alexander Krull, Martin Weigert, Uwe Schmidt, and Gene Myers. Removing structured noise with self-supervised blind-spot networks. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 159–163. IEEE, 2020. 1, 2
- [6] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*, pages 17–33. Springer, 2022. 1, 2
- [7] Yang Chen, Xiangyong Cao, Qian Zhao, Deyu Meng, and Zongben Xu. Denoising hyperspectral image with non-iid noise structure. *IEEE transactions on cybernetics*, 48(3):1054–1066, 2017. 4
- [8] Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated mri. *Medical Image Analysis*, page 102479, 2022. 2
- [9] Zhuo-Xu Cui, Chentao Cao, Shaonan Liu, Qingyong Zhu, Jing Cheng, Haifeng Wang, Yanjie Zhu, and Dong Liang. Self-score: Self-supervised learning on score-based models for mri reconstruction. *arXiv preprint arXiv:2209.00835*, 2022. 1, 2
- [10] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007. 6, 7
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [12] Weisheng Dong, Guangming Shi, and Xin Li. Nonlocal image restoration with bilateral variance estimation: a low-rank approach. *IEEE transactions on image processing*, 22(2):700–711, 2012. 2
- [13] Kuang Gong, Keith A Johnson, Georges El Fakhri, Quanzheng Li, and Tinsu Pan. Pet image denoising based on denoising diffusion probabilistic models. *arXiv preprint arXiv:2209.06167*, 2022. 1, 2
- [14] Mario González, Andrés Almansa, and Pauline Tan. Solving inverse problems by joint posterior maximization with autoencoding prior. *SIAM Journal on Imaging Sciences*, 15(2):822–859, 2022. 2
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 1
- [16] Tao Huang, Songjiang Li, Xu Jia, Huchuan Lu, and Jianzhuang Liu. Neighbor2neighbor: Self-supervised denoising from single noisy images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14781–14790, 2021. 2
- [17] Ajil Jalal, Marius Arvinte, Giannis Daras, Eric Price, Alexandros G Dimakis, and Jon Tamir. Robust compressed sensing mri with deep generative priors. *Advances in Neural Information Processing Systems*, 34:14938–14954, 2021. 1
- [18] Geonwoon Jang, Wooseok Lee, Sanghyun Son, and Kyoung Mu Lee. C2n: Practical generative noise modeling for real-world denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2350–2359, 2021. 2
- [19] Zahra Kadkhodaie and Eero Simoncelli. Stochastic solutions for linear inverse problems using the prior implicit in a denoiser. *Advances in Neural Information Processing Systems*, 34:13242–13254, 2021. 2
- [20] Zahra Kadkhodaie and Eero Simoncelli. Stochastic solutions for linear inverse problems using the prior implicit in a denoiser. *Advances in Neural Information Processing Systems*, 34:13242–13254, 2021. 3
- [21] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. 2022. 1, 2, 4
- [22] Bahjat Kawar, Gregory Vaksman, and Michael Elad. Snips: Solving noisy inverse problems stochastically. *Advances in Neural Information Processing Systems*, 34:21757–21769, 2021. 1, 2, 4
- [23] Bahjat Kawar, Gregory Vaksman, and Michael Elad. Stochastic image denoising by sampling from the posterior distribution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1866–1875, 2021. 1, 2
- [24] Shayan Kousha, Ali Maleky, Michael S Brown, and Marcus A Brubaker. Modeling srgb camera noise with normalizing flows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17463–17471, 2022. 1, 2
- [25] Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void-learning denoising from single noisy images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2129–2137, 2019. 2
- [26] Samuli Laine, Tero Karras, Jaakko Lehtinen, and Timo Aila. High-quality self-supervised deep image denoising. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [27] Matti Lassas and Samuli Siltanen. Can one use total variation prior for edge-preserving bayesian inversion? *Inverse Problems*, 20(5):1537–1563, aug 2004. 2
- [28] Marc Lebrun, Antoni Buades, and Jean-Michel Morel. A nonlocal bayesian image denoising algorithm. *SIAM Journal on Imaging Sciences*, 6(3):1665–1688, 2013. 2
- [29] Wooseok Lee, Sanghyun Son, and Kyoung Mu Lee. Apbsn: Self-supervised denoising for real-world images via asymmetric pd and blind-spot network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17725–17734, 2022. 1, 2, 6, 7, 8
- [30] Jason Lequyer, Reuben Philip, Amit Sharma, Wen-Hsin Hsu, and Laurence Pelletier. A fast blind zero-shot denoiser. *Nature Machine Intelligence*, pages 1–11, 2022. 1
- [31] Jingyun Liang, Jiezhong Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 1833–1844. IEEE Computer Society, 2021. 1, 2

- [32] A.C. Likas and N.P. Galatsanos. A variational approach for bayesian blind image deconvolution. *IEEE Transactions on Signal Processing*, 52(8):2222–2233, 2004. 2
- [33] Xinhao Liu, Masayuki Tanaka, and Masatoshi Okutomi. Practical signal-dependent noise parameter estimation from a single noisy image. *IEEE Transactions on Image Processing*, 23(10):4361–4371, 2014. 6
- [34] Kede Ma, Zhengfang Duanmu, Qingbo Wu, Zhou Wang, Hongwei Yong, Hongliang Li, and Lei Zhang. Waterloo exploration database: New challenges for image quality assessment models. *IEEE Transactions on Image Processing*, 26(2):1004–1016, 2016. 6
- [35] Julien Mairal, Michael Elad, and Guillermo Sapiro. Sparse representation for color image restoration. *IEEE Transactions on image processing*, 17(1):53–69, 2007. 2
- [36] 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 *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 416–423. IEEE, 2001. 6
- [37] Sreyas Mohan, Zahra Kadkhodaie, Eero P Simoncelli, and Carlos Fernandez-Granda. Robust and interpretable blind image denoising via bias-free convolutional neural networks. In *International Conference on Learning Representations*. 6
- [38] Nick Moran, Dan Schmidt, Yu Zhong, and Patrick Coady. Noisier2noise: Learning to denoise from unpaired noisy data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12064–12072, 2020. 2
- [39] Seonghyeon Nam, Youngbae Hwang, Yasuyuki Matsushita, and Seon Joo Kim. A holistic approach to cross-channel image noise modeling and its application to image denoising. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1683–1691, 2016. 6
- [40] 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. 1, 2, 7, 8
- [41] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting deep generative prior for versatile image restoration and manipulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [42] Tongyao Pang, Huan Zheng, Yuhui Quan, and Hui Ji. Recorruped-to-recorruped: unsupervised deep learning for image denoising. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2043–2052, 2021. 2
- [43] 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. *Advances in neural information processing systems*, 32, 2019. 6
- [44] Mangal Prakash, Mauricio Delbracio, Peyman Milanfar, and Florian Jug. Interpretable unsupervised diversity denoising and artefact removal. In *International Conference on Learning Representations*, 2022. 1
- [45] Yuhui Quan, Mingqin Chen, Tongyao Pang, and Hui Ji. Self2self with dropout: Learning self-supervised denoising from single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1890–1898, 2020. 1, 2, 6, 7
- [46] Zaccharie Ramzi, Benjamin Remy, Francois Lanusse, Jean-Luc Starck, and Philippe Ciuciu. Denoising score-matching for uncertainty quantification in inverse problems. In *NeurIPS 2020 Workshop on Deep Learning and Inverse Problems*. 2, 8
- [47] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 1, 2
- [48] 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. Springer, 2015. 6
- [49] Benjamin Salmon and Alexander Krull. Towards structured noise models for unsupervised denoising. In *Computer Vision—ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IV*, pages 379–394. Springer, 2023. 2
- [50] Jae Woong Soh and Nam Ik Cho. Variational deep image restoration. *IEEE Transactions on Image Processing*, 31:4363–4376, 2022. 2
- [51] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *International Conference on Learning Representations*, 2021. 1, 2
- [52] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2020. 1, 2, 4, 8
- [53] Tristan SW Stevens, Jean-Luc Robert, Faik C Yu, Jun Seob Shin, and Ruud JG van Sloun. Removing structured noise with diffusion models. *arXiv preprint arXiv:2302.05290*, 2023. 2
- [54] Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, and Lei Zhang. Ntire 2017 challenge on single image super-resolution: Methods and results. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 114–125, 2017. 6
- [55] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018. 1, 2, 6, 7
- [56] Xiaohe Wu, Ming Liu, Yue Cao, Dongwei Ren, and Wangmeng Zuo. Unpaired learning of deep image denoising. In *European conference on computer vision*, pages 352–368. Springer, 2020. 2

- [57] Jun Xu, Hui Li, Zhetong Liang, David Zhang, and Lei Zhang. Real-world noisy image denoising: A new benchmark. *arXiv preprint arXiv:1804.02603*, 2018. [6](#)
- [58] Zongsheng Yue, Hongwei Yong, Qian Zhao, Deyu Meng, and Lei Zhang. Variational denoising network: Toward blind noise modeling and removal. *Advances in neural information processing systems*, 32, 2019. [2](#), [6](#)
- [59] Zongsheng Yue, Hongwei Yong, Qian Zhao, Lei Zhang, Deyu Meng, and Kwan-Yen K. Wong. Deep variational network toward blind image restoration. 2020. [2](#)
- [60] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. *arXiv preprint arXiv:2111.09881*, 2021. [1](#), [2](#)
- [61] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14821–14831, 2021. [2](#)
- [62] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017. [1](#), [6](#)
- [63] Yi Zhang, Dasong Li, Ka Lung Law, Xiaogang Wang, Hongwei Qin, and Hongsheng Li. Idr: Self-supervised image denoising via iterative data refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2098–2107, 2022. [2](#)
- [64] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11710–11718, 2019. [6](#)
- [65] Dihan Zheng, Sia Huat Tan, Xiaowen Zhang, Zuoqiang Shi, Kaisheng Ma, and Chenglong Bao. An unsupervised deep learning approach for real-world image denoising. In *International Conference on Learning Representations*, 2021. [1](#), [2](#), [6](#), [7](#)
- [66] Dihan Zheng, Xiaowen Zhang, Kaisheng Ma, and Chenglong Bao. Learn from unpaired data for image restoration: A variational bayes approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [1](#), [2](#), [7](#), [8](#)
- [67] Yuqian Zhou, Jianbo Jiao, Haibin Huang, Yang Wang, Jue Wang, Honghui Shi, and Thomas Huang. When awgn-based denoiser meets real noises. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13074–13081, 2020. [1](#), [2](#), [6](#), [7](#)