

Reconstructing the Noise Variance Manifold for Image Denoising

Ioannis Marras¹, Grigorios G. Chrysos², Ioannis Alexiou¹, Gregory Slabaugh¹,
and Stefanos Zafeiriou²

¹ Huawei Noah's Ark

{ioannis.marras,ioannis.alexiou,gregory.slabaugh}@huawei.com

² Imperial College London

{g.chrysos,s.zafeiriou}@imperial.ac.uk

Abstract. Deep Convolutional Neural Networks (CNNs) have been successfully used in many low-level vision problems like image denoising. Although the conditional image generation techniques have led to large improvements in this task, there has been little effort in providing conditional generative adversarial networks (cGANs) with an explicit way of understanding the image noise for object-independent denoising reliable for real-world applications. The task of leveraging structures in the target space is unstable due to the complexity of patterns in natural scenes, so the presence of unnatural artifacts or over-smoothed image areas cannot be avoided. To fill the gap, in this work we introduce the idea of a cGAN which explicitly leverages structure in the image noise variance space. By learning directly a low dimensional manifold of the image noise variance, the generator promotes the removal from the noisy image only that information which spans this manifold. This idea brings many advantages while it can be appended at the end of any denoiser to significantly improve its performance. Based on our experiments, our model substantially outperforms existing state-of-the-art architectures, resulting in denoised images with less over-smoothing and better detail.

1 Introduction

During image acquisition, due to the presence of noise some image corruption is inevitable and can degrade the visual quality considerably. Therefore, noise removal is essential for many digital imaging and computer vision applications [22] and remains an important and active research topic.

Denoising algorithms can be grouped in two categories: learning-based and model-based. Modelling the image prior from a set of noisy and ground-truth image sets is the goal of discriminative learning. The performance of the current learning models is limited by their inadequacy of handling all possible levels of noise in a single model. In this category are methods such as brute force learning like MLP [11], CNNs [58, 59] or truncated inference [15]. On the other hand, the model-based algorithms are computationally expensive, and unable to characterize complex image textures. In the this category fall algorithms including

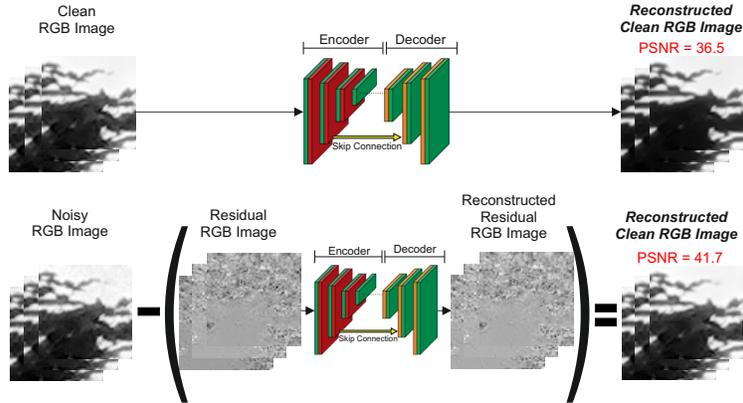


Fig. 1: **Motivation of our method:** By characterizing directly the image signal dependent noise, the reconstruction of the clean image is much more accurate. Instead of constraining the output of a generator to span the target space, is better to constrain it to remove from the noisy image only that information which spans the manifold of the residual image variance.

external priors [7], Markov random field models [49, 52], gradient methods [55, 54], non-local self-similarity [33] and sparsity (e.g. MCWNNM [24]).

A denoising algorithm should be efficient, perform denoising using a single model and handle spatially variant noise when the noise standard-deviation is known or unknown. The physics of digital sensors and the steps of an imaging pipeline are well-understood and can be leveraged to generate training data from almost any image using only basic information about the target camera sensor. Recent work has shifted to sophisticated signal-dependent single source noise models [27] that better match the physics of image formation [38, 43, 9]. Also, adapting a learned denoising algorithm to a new camera sensor may require capturing a new dataset. However, capturing noisy and noise-free image pairs is difficult, requiring long exposures or large bursts of images, and post-processing to combat camera motion and lighting changes.

In this paper, we introduce the idea of a cGAN [44] which directly constrains the image spatially variant noise for image denoising (Fig. 1). In this way, we avoid the direct characterization of the space of clean images, since the complexity of natural image patterns is extremely high. To do so, a combination of supervised (regression) and unsupervised (autoencoder) ‘*encoder-decoder*’ type subnets applies implicit constraints in the residual image (the difference between the noisy observation and the clean image) variance latent subspace. By adopting the idea of residual learning [58] in the regression subnet and using a shared decoder, the unsupervised subnet is explicitly constrained to generate residual image samples that span only the image noise variance manifold. Intuitively, this can be thought of as constraining the regression subnet to *subtract from the*

noisy image only the residual image that looks like realistic image noise coming from a specific camera sensor. The proposed idea: a) allows the direct association of one or more camera sensors with their corresponding noise statistics and b) introduces also the idea of a discriminator operating directly in the residual image domain. Our system: a) increases significantly the robustness of the image denoising task, b) makes easier the model adaptation to a new camera sensor, c) allows multi-camera noise reduction during one inference step, d) allows multi-source noise removal during one inference step, e) utilizes all the samples in the residual image domain even in the absence of the corresponding noisy input samples, f) can be applied at the end of any residual learning based denoiser improving its performance and g) deals with a wide range of noise levels.

2 Related Work

2.1 Image Prior Based Methods

Image prior based methods, e.g. NSCR [21], TWSC [56], WNNM [24], can be employed to solve the denoising problem of unknown noise because they do not require training data since they model the image prior over the noisy image directly. The classic BM3D [19] method is based on the idea that natural images usually contain repeated patterns (non-local self-similarity model). In non-local means (NLM) [10], the pixel values are predicted based on their noisy surroundings. Many variants of NLM and BM3D seeking self-similar patches in different transform domains were proposed, e.g. SAPCA [33], NLB [35]. Sparsity is enforced by dictionary-based methods [20] by employing self-similar patches and learning over-complete dictionaries from clean images. In contrast, Noise2Void (N2V) [34] and Noise2Noise (N2N) [37] do not require training noisy image pairs, nor clean target images. N2N attempts to learn a mapping between pairs of independently degraded versions of the same training image. For image patch restoration, maximum likelihood algorithms like Gaussian Mixture Models (GMMs), were employed to learn statistical priors from image patch groups [13, 57]. Dictionary learning based and basis-pursuit based algorithms such as KSVD [4], Fields-of-Experts or TNRD [16] operated by finding image representations where sparsity holds or statistical regularities are well-modeled [61]. In [36], an extension of non-local Bayes approach, named NC, was proposed to model the noise of each patch group to be zero-mean correlated and Gaussian distributed. The disadvantage of this category of methods is that external information from possible many other images taken under the same condition with the image to be denoised cannot be used. Furthermore, the generalization capabilities are limited because these methods are defined mostly based on human knowledge.

2.2 Discriminative Deep Learning Methods

In recent years, CNNs have achieved great success in image denoising. The first attempt of employing CNNs for this task was made in [30]. Discriminative deep

learning methods are trained offline, extracting information from ground truth annotated training sets before they are applied to test data. In DnCNN [58] and IrCNN [59] networks, stacked convolution, batch normalization and ReLU layers were used to estimate the residual image [26]. By adding symmetric skip connections, an improved encoder-decoder network for image denoising based on residual learning was proposed in [41]. A densely connected denoising network, named Memnet, constructed in [51] to enable memory of the network. A multi-level wavelet CNN (MWCNN) model based on a U-Net architecture used in [40] to incorporate large receptive field for image denoising. By incorporating non-local operations into a recurrent neural network (RNN), a non-local recurrent network (NLRN) for image restoration presented in [39]. A network named N³Net [47] employed the k-nearest neighbor matching in the denoising network to exploit the non-local property of the image features. A fast and flexible network (FFDNet) which can process images with non-uniform noise corruption proposed in [60]. A residual in the residual structure (RIDNet) used in [6] to ease the flow of low-frequency information and apply feature attention to exploit the channel dependencies. Recently, a blind denoising model for real photographs named CBDNet [25] is composed of two subnetworks: noise estimation and non-blind denoising. A self-guided network (SGN), which adopts a top-down self-guidance architecture to better exploit image multi-scale information presented in [23]. FOCNet network [31] solved a fractional optimal control problem in a multi-scale approach. Although the methods in this category achieved high denoising quality, they cannot work in the absence of paired training data.

2.3 Generative Models

GANs were recently trained to synthesize noise [14], thus pairs of corresponding clean and noisy images were obtained for training CNNs. Any further filtering of the RAW image changes the real noise statistics making that task very difficult [14]. Also, is not realistic to create noisy images by adding random generated noise to the clean images since in real images the noise variance is data dependent. Noise Flow method [1] combined well-established basic parametric noise models (e.g. signal-dependent noise) with the flexibility and expressiveness of normalizing flow architectures to model noise distributions observed from large datasets of real noisy images. However, it is not clear how to quantitatively assess the quality of the generated samples.

3 Our Method

In this section, we introduce our system for the task of image denoising. The goal is to produce a single clean (RGB or RAW) image from a corresponding single noisy (RGB or RAW) image. Firstly, we give a brief overview of the noise signal in real images (Section 3.1). Our method falls in the category of conditional image generation methods, thus to make the paper self-contained we briefly describe this category (Section 3.2) before introducing our method (Section 3.3).

3.1 Image Noise Modeling in Real-World Images

Camera sensors output RAW data in a linear color space where pixel measurements are proportional to the number of photo-electrons collected. The primary sources of noise are shot noise, a Poisson process with variance equal to the signal level, and read noise, an approximately Gaussian process caused by a variety of sensor readout effects. The noise is spatially variant; hence, the assumption that noise is spatially invariant does not hold for real images. The noise is well-modeled by a signal-dependent Gaussian distribution [27]:

$$x_p \sim \mathcal{N}(y_p, \sigma_r^2 + \sigma_s y_p) \quad (1)$$

where x_p is a noisy measurement of the true intensity y_p at pixel p . The parameters σ_r and σ_s : a) are fixed, given a specific camera sensor, for each sensor gain (ISO) value and varies as ISO changes and b) are different for different camera sensors. Since the noise is structured (not random) a low-dimensional manifold for noise variance exists. A realistic noise model is important aspect in training CNN-based denoising methods for real photographs [25, 6].

3.2 Conditional Image Generation

In computer vision, the task of conditional image generation is dominated by approaches similar to a GAN. The GAN consists of a generator and a discriminator module commonly optimized with alternating gradient descent methods. cGAN extend the formulation by providing the generator with additional labels. The generator G takes the form of an encoder-decoder network where the encoder projects the label into a low-dimensional latent subspace and the decoder performs the opposite mapping.

cGAN and its variants like Robust cGAN [17], were applied in the past for the task of object-dependent image denoising. The encoder-decoder generator of Robust cGAN performs a similar regression as its counterpart in cGAN. It accepts a sample from the source domain and maps it to the target domain by using a second CNN in the target domain which promotes more realistic regression outputs. Recently in non GAN-based methods, generators adopting a similar architecture were proposed for object-dependent image denoising. In [29], a two-tailed CNN is employed for inferring the clean image and the noise separately. The input noisy image is decoupled to the signal and noise in a latent space, while a decoder used to generate the signal and noise in the spatial domain. There are two major drawbacks of all these methods: i) in the absence of skip connections, these methods perform well only in the case of object-dependent image denoising (i.e. face denoising [17, 53]). The need of having different models for different objects makes them unsuitable for digital devices with limited resources (e.g. smartphones) where the run-time performance is of importance. ii) the purpose of their unsupervised learning sub-networks, whose (hidden) layers contain representations of the input data, is to be sufficiently powerful for compressing (and decompressing) the data while losing as little information as possible. However, even in the presence of skip connections, this procedure of

defining a nonlinear representation which can accurately reconstruct image patterns from a variety of real complex objects/scenes is not realistic. As a result, these methods very often hallucinate complex image structures by introducing severe blurry effects or unusual image patterns/artifacts.

3.3 Image Denoising Based on Noise Variance Manifold Reconstruction

To tackle the problems mentioned in Section 3.2, the proposed method introduces the general idea of explicitly constraining the residual image removed by a denoiser to lie in the low-dimensional manifold of the signal dependent image noise variance (Section 3.1). Like cGAN, our method consists of a generator and a discriminator. The generator includes two subnets: the first regression (*Reg*) subnet performs regression while the second reconstruction (*Rec*) is an autoencoder in the residual image domain. Both subnets consist of similar encoder-decoder networks, while a backbone network is used prior to the encoder-decoder network of the *Reg subnet*. By sharing the weights of their decoders, the generator adopts the residual learning strategy to remove from the noisy observation that information which spans the image noise variance manifold. A schematic of the proposed generator is illustrated in Fig. 2. Rather than directly outputting the denoised image, the supervised *Reg subnet* is designed to predict the ground-truth residual image $\mathbf{v} = \mathbf{s} - \mathbf{y}$, where \mathbf{s} and \mathbf{y} stand for the noisy and the clean (ground-truth) image, respectively. Thus, the unsupervised *Rec subnet* works as a conditional auto-encoder in the domain of \mathbf{v} . The *Rec subnet* during inference is no longer required, therefore the testing complexity remains the same as in standard cGAN. Two Unet style skip connections from the encoder to the decoder used in both subnets improving the learning of the residual between the features corresponding to the image and to the residual image structures. Also, a BEGAN style decoder skip connection [8], which creates a skip connection between the first decoder layer and each successive upsampling layer of the decoder, used to help gradient propagation.

The reconstruction of \mathbf{v} is an easier task compared to the reconstruction of \mathbf{y} as in cGAN. Thus, we can learn how to turn bad images into good images by only looking at the structure of the residual image. This property: i) makes our cGAN an object-independent image denoiser and ii) helps the denoiser largely avoid image over-smoothing/artifacts, something essential for image denoising.

Noising is a challenging process to be reversed by the few convolutional layers of the encoder in *Reg subnet* especially in an object-independent scenario. This is why a backbone network used to extract complex feature representations, φ , useful to preserve for later the low and high image frequencies. Different state-of-the-art denoisers could be used as backbone networks. Thus, the proposed idea could be applied at the end of any denoiser constraining its output improving in that way its performance as it is experimentally verified in Section 4.

In addition, the *Rec subnet* enables utilization of all the samples in the domain of the residual image even in the absence of the corresponding noisy input samples. In the case of a well defined image noise source, like the one described

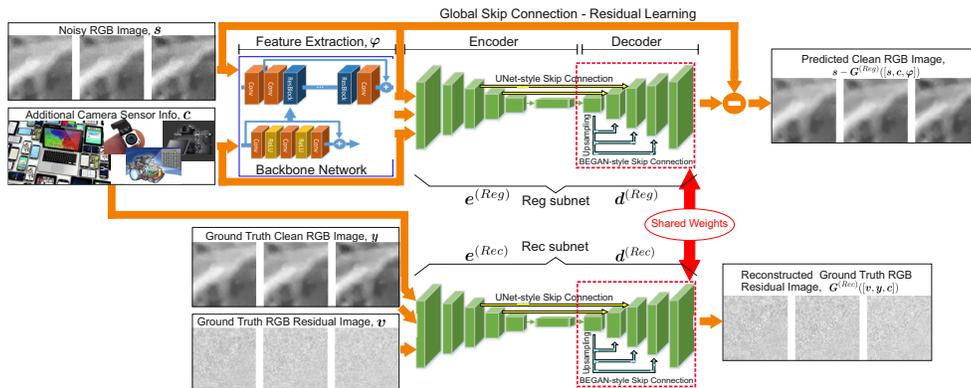


Fig. 2: Schematic of the proposed generator.

in Section 3.1, a huge amount of different residual image realizations (e.g. for different ISOs) could be generated and used to train that subnet.

The adaptation of an existing model to a new camera sensor is an easier task for our method. To do so, only the *Rec subnet* must be retrained from scratch while the *Reg subnet* needs only to be fine-tuned using a small number of paired training samples obtained using the new sensor. Also, our method can remove more than one noise source during one inference step. To do so, a different noise variance manifold for each noise source is obtained, thus a different *Rec subnet* per noise source constrains the denoiser in a sequential manner (Fig. 4(a)).

The task of learning directly the image noise variance manifold can be greatly benefit by any conditional information, \mathbf{c} , related to the camera sensor. This information varies and it is provided to both subnets. \mathbf{c} could contain the two noise parameters σ_r and σ_s (if available) associating in that way a camera sensor with its corresponding noise statistics. In the case of multi-camera noise reduction, \mathbf{c} could additionally contain a one hot vector per pixel defining the camera id used to take each picture, thus one or more noise sources are explicitly associated with the corresponding camera sensor. More specifically, *Reg subnet* gets as input \mathbf{c} in concatenation (denoted as $[\cdot]$) with \mathbf{s} and φ and outputs $\mathbf{y} - \mathbf{G}^{(Reg)}([\mathbf{s}, \mathbf{c}, \varphi])$, where $\mathbf{G}^{(Reg)}([\mathbf{s}, \mathbf{c}, \varphi])$ is the predicted residual image. The superscript ‘Reg’ abbreviates modules of the *Reg subnet*. Based on eq. 1, the noise variance for a pixel p depends, except for the camera sensor-based parameters, on y_p . Thus, the input to *Rec subnet* should be \mathbf{v} in concatenation with \mathbf{y} and \mathbf{c} . By giving explicitly \mathbf{y} as additional input, *the task of the Rec subnet is not to learn the underlying structure of a huge variety of complex image patterns, but to learn how clean image structures are affected by the presence of structured noise.*

The proposed idea deals with a wide range of noise levels in contrast to a standard cGAN or its variants. According to [26], when the original mapping $F(\mathbf{s})$ (as in cGAN) is more like an identity mapping, the residual mapping will be much easier to optimize. Note that \mathbf{s} is much more like $\mathbf{s} - \mathbf{G}^{(Reg)}([\mathbf{s}, \mathbf{c}, \varphi])$ than $\mathbf{G}^{(Reg)}([\mathbf{s}, \mathbf{c}, \varphi])$ (especially when the noise level is low). Thus, $F(\mathbf{s})$ would

be closer to an identity mapping than $\mathbf{G}^{(Reg)}([\mathbf{s}, \mathbf{c}, \boldsymbol{\varphi}])$, and the residual learning formulation is more suitable for image denoising [58].

In the case of image denoising in the RGB domain, \mathbf{s} represents 3-channel image based tensors. Regarding \mathbf{s} in the RAW domain, each pixel in a conventional camera (linear Bayer) sensor is covered by a single red, green, or blue color filter, arranged in a 4-channel Bayer pattern (i.e. R-G-G-B). The content loss consists of two terms that compute the per-pixel difference between the predicted clean image, and the clean (ground-truth) image. The two terms are i) the ℓ_1 loss between the ground-truth image and the output of the generator, ii) the ℓ_1 of their gradients; mathematically expressed as:

$$\mathcal{L}_c = \lambda_c \cdot \sum_{n=1}^N \|(\mathbf{s}^{(n)} - \mathbf{G}^{(Reg)}([\mathbf{s}^{(n)}, \mathbf{c}^{(n)}, \boldsymbol{\varphi}^{(n)}])) - \mathbf{y}^{(n)}\| + \lambda_{cg} \cdot \sum_{n=1}^N \|\nabla(\mathbf{s}^{(n)} - \mathbf{G}^{(Reg)}([\mathbf{s}^{(n)}, \mathbf{c}^{(n)}, \boldsymbol{\varphi}^{(n)}])) - \nabla \mathbf{y}^{(n)}\|, \quad (2)$$

where $\mathbf{G}^{(Reg)}([\mathbf{s}^{(n)}, \mathbf{c}^{(n)}, \boldsymbol{\varphi}^{(n)}]) = \mathbf{d}^{(Reg)}(\mathbf{e}^{(Reg)}([\mathbf{s}^{(n)}, \mathbf{c}^{(n)}, \boldsymbol{\varphi}^{(n)}]))$, N stands for the total number of training samples, \mathbf{e} stands for encoder, \mathbf{d} stands for decoder and $\lambda_c, \lambda_{cg} = 0.5 \cdot \lambda_{ae}$ are hyper-parameters to balance the loss terms. The unsupervised *Rec subnet* contributes the following loss term:

$$\mathcal{L}_{Rec} = \sum_{n=1}^N [f_d(\mathbf{v}^{(n)}, \mathbf{G}^{(Rec)}([\mathbf{v}^{(n)}, \mathbf{y}^{(n)}, \mathbf{c}^{(n)}]))] \quad (3)$$

where $\mathbf{G}^{(Rec)}([\mathbf{v}^{(n)}, \mathbf{y}^{(n)}, \mathbf{c}^{(n)}]) = \mathbf{d}^{(Rec)}(\mathbf{e}^{(Rec)}([\mathbf{v}^{(n)}, \mathbf{y}^{(n)}, \mathbf{c}^{(n)}]))$ is the *Rec subnet*'s output, f_d is a divergence metric (ℓ_2 loss due to the auto-encoder in the noise domain) and the superscript ‘Rec’ abbreviates modules of the *Rec subnet*.

Despite sharing the weights of the decoders, the latent representations of the two subnets are forced to span the same space using a latent loss term \mathcal{L}_{lat} . This term minimizes the distance between the encoders’ outputs, i.e. the two residual noise variance representations are spatially close. The latent loss term is:

$$\mathcal{L}_{lat} = \sum_{n=1}^N \|\mathbf{e}^{(Reg)}([\mathbf{s}^{(n)}, \mathbf{c}^{(n)}, \boldsymbol{\varphi}^{(n)}]) - \mathbf{e}^{(Rec)}([\mathbf{v}^{(n)}, \mathbf{y}^{(n)}, \mathbf{c}^{(n)}])\|. \quad (4)$$

As a part of the vanilla cGAN, the feature matching loss [50, 28] enables the network to match the data and the model’s distribution faster. The intuition is that to match the high-dimensional distribution of the data with *Reg subnet*, their projections in lower-dimensional spaces are encouraged to be similar. The feature matching loss is:

$$\mathcal{L}_f = \sum_{n=1}^N \|\pi(\mathbf{s}^{(n)} - \mathbf{G}^{(Reg)}([\mathbf{s}^{(n)}, \mathbf{c}^{(n)}, \boldsymbol{\varphi}^{(n)}])) - \pi(\mathbf{y}^{(n)})\|. \quad (5)$$

where $\pi()$ extracts the features from the penultimate layer of the discriminator.

Skip connections enable deeper layers to capture more abstract representations without the need of memorizing all the information. The lower-level representations are propagated directly to the decoder through the shortcut, which makes it harder to train the longer path [48]. The Decov loss term [18] used: a) to penalize the correlations in the representations of one or more layers, b) to implicitly encourage the representations to capture diverse and useful information and c) to maximize the variance captured by the longer path representations. For the j^{th} layer this loss is defined as:

$$\mathcal{L}_{decov}^j = \frac{1}{2} (\|\mathbf{C}^j\|_F^2 - \|\text{diag}(\mathbf{C}^j)\|_2^2), \quad (6)$$

where $diag()$ computes the diagonal elements of a matrix and \mathbf{C}^j is the covariance matrix of the layer representations. The loss is minimized when the covariance matrix is diagonal, i.e. it imposes a cost to minimize the covariance of hidden units without restricting the diagonal elements that include the variance of the hidden representations.

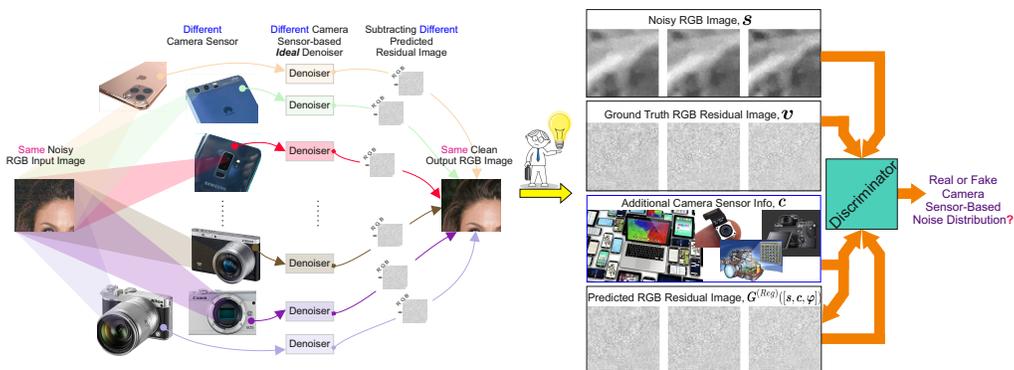


Fig. 3: The proposed discriminator for image denoising operates directly in the residual image domain.

In the case of multi-camera noise reduction, let's assume that the same scene is captured under the same lighting conditions by different camera sensors. Let's also assume that an ideal denoiser per camera sensor exists. In that case, the output of all the denoisers should be the same underlying clean image although the noise statistics of each camera can be very different. This leads to the idea of a discriminator which operates directly in the residual image domain (Fig. 3) thus trying to distinguish between the residual image samples generated by the denoiser and the ground-truth residual image distributions given a specific camera sensor. This is feasible in the proposed method because the *Rec subnet* constraints directly the denoiser to remove only that information which spans the learned noise variance manifold of each camera sensor. The generator samples \mathbf{z} from a prior distribution $p_{\mathbf{z}}$, e.g. uniform, and tries to model the target distribution p_d ; the discriminator D tries to distinguish between the samples generated from the model and the target image noise distributions. More specifically, the discriminator accepts as input $\mathbf{G}^{(Reg)}([\mathbf{s}, \mathbf{c}, \boldsymbol{\varphi}])$ along with \mathbf{v} , \mathbf{c} and \mathbf{s} , while the standard adversarial loss of cGAN is modified to:

$$\mathcal{L}_{adv}^*(\mathbf{G}^{(Reg)}, D) = \mathbb{E}_{\mathbf{s}, \mathbf{v} \sim p_d(\mathbf{s}, \mathbf{v})} [\log D(\mathbf{v} | \mathbf{s}, \mathbf{c})] + \mathbb{E}_{\mathbf{s} \sim p_d(\mathbf{s}), \mathbf{z} \sim p_z(\mathbf{z})} [\log(1 - D(\mathbf{G}^{(Reg)}([\mathbf{s}, \mathbf{c}, \boldsymbol{\varphi}] | \mathbf{s}, \mathbf{c}))]. \quad (7)$$

by solving the following min-max problem:

$$\min_{\mathbf{w}_G} \max_{\mathbf{w}_D} \mathcal{L}_{adv}^*(\mathbf{G}^{(Reg)}, D) = \min_{\mathbf{w}_G} \max_{\mathbf{w}_D} \mathbb{E}_{\mathbf{s}, \mathbf{v} \sim p_d(\mathbf{s}, \mathbf{v})} [\log D(\mathbf{v} | \mathbf{s}, \mathbf{c}, \mathbf{w}_D)] + \mathbb{E}_{\mathbf{s} \sim p_d(\mathbf{s}), \mathbf{z} \sim p_z(\mathbf{z})} [\log(1 - D(\mathbf{G}^{(Reg)}([\mathbf{s}, \mathbf{c}, \boldsymbol{\varphi}] | \mathbf{w}_G) | \mathbf{s}, \mathbf{c}, \mathbf{w}_D))]$$

where $\mathbf{w}_G, \mathbf{w}_D$ denote the generator’s and the discriminator’s parameters respectively. The final loss function of our method is:

$$\mathcal{L}_{total} = \mathcal{L}_{adv}^* + \mathcal{L}_c + \lambda_\pi \cdot \mathcal{L}_f + \lambda_{ae} \cdot \mathcal{L}_{Rec} + \lambda_l \cdot \mathcal{L}_{lat} + \lambda_d \cdot \sum^j \mathcal{L}_{decov}^j, \quad (8)$$

where $\lambda_\pi, \lambda_{ae}, \lambda_l$ and λ_d are extra hyper-parameters to balance the loss terms.

(a)

Method	Type	RAW		sRGB		Runtime (ms)
		PSNR	SSIM	PSNR	SSIM	
GCBD [14]	Blind	-	-	35.58	0.9217	-
FoE [49]	Non-blind	45.78	0.9666	35.99	0.9042	-
TNRD [16] + VST	Non-blind	45.70	0.9609	36.09	0.8883	5,200
MLP [11] + VST	Non-blind	45.71	0.9629	36.72	0.9122	~60,000
MCWNNM [24]	Non-blind	-	-	37.38	0.9294	208,100
EPLL [61] + VST	Non-blind	46.86	0.9730	37.46	0.9245	-
KSVD [4] + VST	Non-blind	46.87	0.9723	37.63	0.9287	>60,000
WNNM [24] + VST	Non-blind	47.05	0.9722	37.69	0.9260	-
NCSR [21] + VST	Non-blind	47.07	0.9688	37.79	0.9233	-
BM3D [19] + VST	Non-blind	47.15	0.9737	37.86	0.9296	6,900
Whitener [29]	Blind	47.16	0.9737	37.88	0.9307	48
RoCGAN [17]	Blind	47.17	0.9738	37.90	0.9310	49
TWSC [56]	Blind	-	-	37.94	0.9403	195,200
CBDNet [25]	Blind	-	-	38.06	0.9421	400
DnCNN [58]	Blind	47.37	0.9760	38.08	0.9357	60
N ³ Net [47]	Blind	47.56	0.9767	38.32	0.9384	210
RIDNet [6]	Blind	-	-	39.23	0.9526	215
UPI [9]	Blind	48.89	0.9824	40.35	0.9641	22
Ours (empty c. ResNet [26])	Blind	49.90 (+1.01)	0.9861	41.50 (+1.15)	0.9759	52
Ours (empty c. UPI)	Blind	50.05 (+1.16)	0.9866	41.59 (+1.24)	0.9760	64
Ours (Non empty c. ResNet)	Non-Blind	50.91 (+3.76)	0.9873	42.11 (+4.25)	0.9775	63

Ablations of Ours (empty c. ResNet [26])

Standard Discriminator [44]	Blind	49.50 (+0.61)	0.9835	41.0 (+0.65)	0.9714	52
No Rec subnet	Blind	47.51 (-1.38)	0.9766	38.54 (-1.81)	0.9417	52
No Rec subnet, No res. learning	Blind	46.92 (-1.97)	0.9725	37.73 (-2.62)	0.9316	52

(b)

Fig. 4: (a) In the case of camera multi-source image noise, more than one *Rec subnet* can be employed. Each subnet is responsible for removing noise structure that comes from a specific noise source, and (b) the quantitative results on the DnD benchmark of our method and its ablations. Regarding our method, in parentheses we define the type of denoising plus the used backbone network.

4 Experimental Results

4.1 Training Settings

Synthetic noisy images were combined with real noisy data to improve the generalization ability of our method to real photographs. To generate them, we followed the pipeline in [25]. To do so, we employed BSD500 [42], DIV2K [3], and MIT-Adobe FiveK [12], resulting in 3.5K images while for real noisy images, we extracted cropped patches from SSID [2] and RENOIR [5]. Finally, the data augmentation procedure results in 64×64 image patches. In our ‘*encoder-decoder*’ architecture (same for both subnets) 11 layers were used with an latent

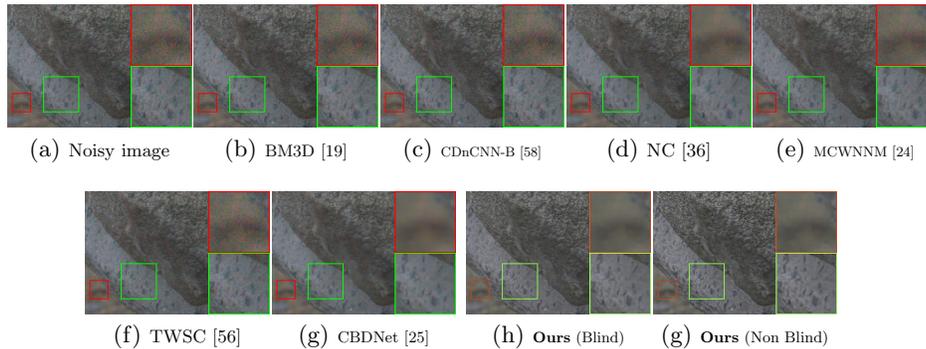


Fig. 5: Example of image denoising of a DnD image. Results of the proposed method shown when a standard ResNet used as backbone network.

space of dimensions $MB \times 2 \times 2 \times 1024$, where MB stands for the mini-batch size. The values of the additional hyper-parameters are $\lambda_{ae} = 2.5 * 10^3$, $\lambda_l = 0.5$ and $\lambda_d = 10^{-6}$. The common hyper-parameters λ_π and λ_c with the vanilla cGAN remain the same. In the beginning, the two subnets were trained separately while afterwards they were jointly trained. For both subnets: the kernel size used was 3×3 ; Adam [32] was used as the optimizer with default parameters; the learning rate was initially set to 10^{-3} and then halved after 10^6 iterations; ReLU activation used; the network ran for 50 epochs.

4.2 Comparisons on Real-World Images

The most three challenging public datasets that significantly improve upon earlier (and often unrealistic) benchmarks for denoising, were used to evaluate the performance of our method: the Darmstadt Noise Dataset (DnD) [46], the Nam Dataset [45] and the Smartphone Image Denoising Dataset (SID) [2]. These datasets are multi-camera datasets (camera id is provided), thus can be used for performing multi-camera noise reduction. To highlight the contribution of the proposed idea, as the backbone network in our method we used: a) a standard residual network (ResNet) [26] created by stacking three building blocks, and b) the best deep learning-based method in the literature according to each benchmark, if existing, excluding the last network layer since this network acts as a feature extractor. The pre-trained weights reported in the literature, if available, used as initialization of the backbone network.

Evaluation on DnD: DnD is a novel benchmark dataset which consists of realistic uncompressed photos from 50 scenes taken by 4 different standard consumer cameras of natural “in the wild” scene content. In DnD: the camera metadata have been captured; the noise properties have been carefully calibrated; and the image intensities are presented as RAW unprocessed linear intensities. For each real high-resolution image, the noisy high-ISO image is paired with the corresponding (nearly) noise-free low-ISO ground-truth image.

The evaluation of DnD is separated in two categories: algorithms that use linear Bayer sensor readings or algorithms that use bilinearly demosaiced sRGB images as input. Thus, PSNR and SSIM for each technique are reported for both categories. The quantitative results with respect to prior work of our method and its ablations are shown in Fig. 4(b). For algorithms which have been evaluated with and without a variance stabilizing transformation (VST), the version which performs better is reported. The evaluation of algorithms that only operate on sRGB inputs is also reported. The proposed idea was tested for both categories. A blind and a non-blind version of our method had been tested for each category based on the info that \mathbf{c} represents. The blind version uses no extra conditional information along with the noisy input image (empty \mathbf{c}). As described in Section 3.3, in the non-blind version, \mathbf{c} could contain information regarding the camera noise model and/or the camera id. As a backbone network for blind image denoising, two variants used: a) the standard ResNet and b) the best method in the literature named UPI [9]. In the case of RAW image domain, the first variant produced significantly higher PSNR (+1.01dB) and SSIM than UPI, while the second one impressively boosted the performance of UPI by 1.16dB. In the case of sRGB image domain, the first variant produced significantly higher PSNR (+1.15dB) and SSIM than UPI, while the second one impressively boosted the performance of UPI by 1.24dB. As a backbone network for non-blind image denoising, only the standard ResNet used since the best methods in the literature are not deep learning techniques. In the case of image RAW domain, our system produced significantly higher PSNR (+3.76dB) and SSIM compared to the second best method named BM3D [19]+VST. In case of sRGB image domain, the improvement over BM3D+VST was 4.25dB. Also, runtimes (mean over 100 runs) reported in the literature are presented as well in Fig. 4(b). The runtime (excluding data transferring to GPU) of our blind model with standard ResNet as backbone network is 52ms while for the non-blind one is 63ms given as input 512×512 images. Some qualitative results are given in Fig. 5.

Evaluation on Nam: The Nam dataset consists of 11 static scenes captured by 3 consumer cameras. For each scene, 500 JPEG noisy temporal images were captured to compute the temporal nearly noise-free mean image and covariance matrix for each pixel. The quantitative results with respect to prior work are shown in Fig. 7(a). Both the blind and non-blind version of our method were evaluated. As a backbone network for blind image denoising, two variants used: a) the standard ResNet and b) the best method in the literature named CBDNet [25]. The first variant produced significantly higher PSNR (+1.03dB) and SSIM than CBDNet, while the second one impressively boosted the performance of CBDNet by 1.18dB. CBDNet-JPEG [25] is a version of CBDNet which specifically deals with the JPEG compression. For fair comparison, we have retrained both variants by adopting this data augmentation technique. In that case, the first variant produced significantly higher PSNR (+0.89dB) and SSIM than CBDNet-JPEG, while the second one impressively boosted the performance of CBDNet-JPEG by 1.07dB. As a backbone network for non-blind image denoising, only the standard ResNet used since the best methods in bib-

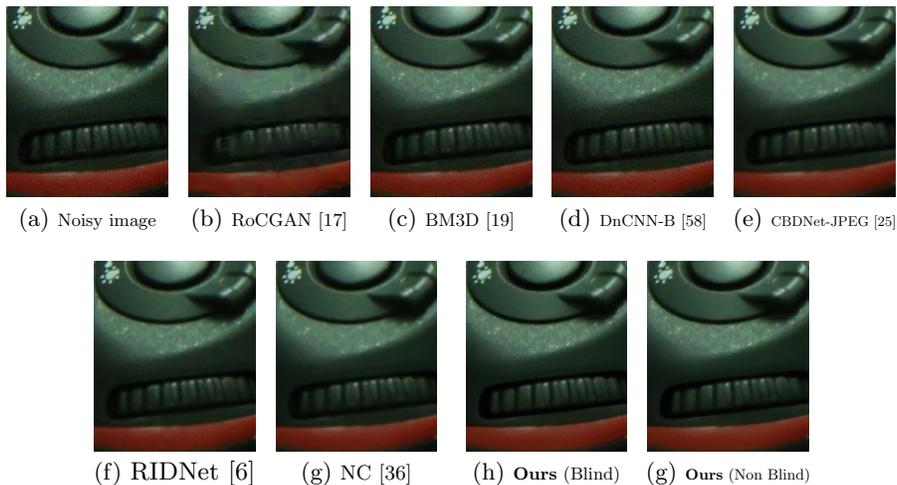


Fig. 6: Example of image denoising of a Nam image. Results of the proposed method shown when a standard ResNet used as backbone network.

liography are not deep learning techniques. Our system produced significantly higher PSNR (+0.97dB) and SSIM compared to the second best method named WNNM [24]. Some qualitative results are given in Fig. 6.

Evaluation on SIDD: SIDD is real noise dataset with a large number of available test (validation) images. The quantitative results on the SIDD benchmark with respect to prior work are shown in Fig. 7(b). Both the blind and non-blind version of our method were evaluated. As a backbone network for blind image denoising, two variants used: a) the standard ResNet and b) the best method in the literature named RIDNet [6]. The first variant produced significantly higher PSNR (+1.11dB) than RIDNet, while the second one impressively boosted the performance of RIDNet by 1.14dB. As a backbone network for non-blind image denoising, only the standard ResNet was used since the best method in the literature, named BM3D [19], is not a deep learning technique. Our system produced significantly higher PSNR (+8.93dB) compared to BM3D. Some qualitative results are given in Fig. 8.

Since the idea behind our method favours the multi-camera noise reduction task, there is a significant improvement in terms of performance across all datasets. Based on all our experiments, the proposed idea is general and can be appended at the end of existing image denoising methods to significantly improve their performance. In addition, the proposed idea better restores the true colors than the competing methods. Also, by directly characterizing the image noise, our method avoids in great degree the image over-smoothing.

Method	Type	PSNR	SSIM
CDnCNN-B [58]	Blind	37.49	0.9272
TWSC [56]	Blind	37.52	0.9292
MCWNNM [24]	Blind	37.91	0.9322
RoCGAN [17]	Blind	38.52	0.9517
Whitener [29]	Blind	38.62	0.9527
RIDNet [6]	Blind	39.09	0.9591
BM3D [19]	Non-blind	39.84	0.9657
CBDNet [25]	Blind	40.02	0.9687
NC [36]	Blind	40.41	0.9731
WNNM [24]	Non-blind	41.04	0.9768
Ours (empty c, ResNet [26])	Blind	41.05 (+1.03)	0.9772
Ours (empty c, CBDNet)	Blind	41.20 (+1.18)	0.9783
Ours (Non empty c, ResNet)	Non-Blind	42.01 (+0.97)	0.9830

Method	Type	PSNR
DnCNN-B [58]	Blind	26.21
FFDNet [60]	Blind	29.20
CBDNet-JPEG [25]	Blind	30.78
BM3D [19]	Non-blind	30.88
Whitener [29]	Blind	37.57
RoCGAN [17]	Blind	37.72
RIDNet [6]	Blind	38.71
Ours (empty c, RIDNet)	Blind	39.82 (+1.11)
Ours (empty c, ResNet [26])	Blind	39.85 (+1.14)
Ours (Non empty c, ResNet)	Non-Blind	39.81 (+8.93)

(a) (b)

Fig. 7: (a) The quantitative results on the Nam benchmark and (b) the quantitative results on the SIDD benchmark. Regarding our method, in parentheses we define the type of denosing plus the used backbone network.

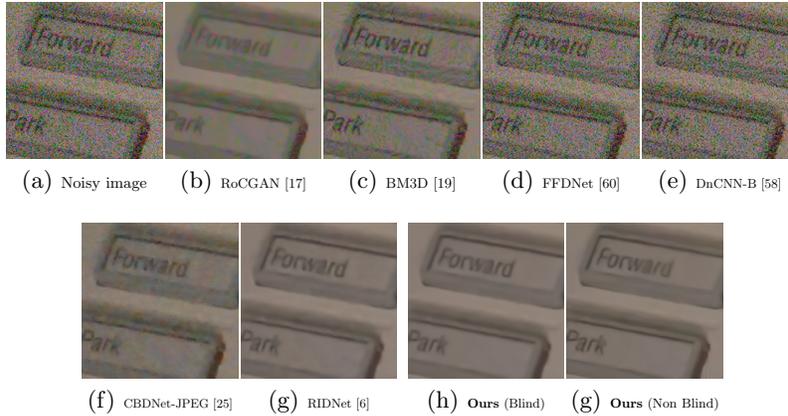


Fig. 8: Example of image denoising of a SIDD image. Results of the proposed method shown when ResNet used as backbone network.

5 Conclusions

In this work, we show that is easier to turn noisy images into clean images only by looking at the structure of the residual image. We introduce the idea of a cGAN that explicitly leverages structure in the image noise variance space. By adopting the residual learning, the generator promotes the removal from the noisy image only that information which spans the manifold of the image noise variance. Our method significantly outperforms existing state-of-the-art architectures.

References

1. Abdelhamed, A., Brubaker, M.A., Brown, M.S.: Noise flow: Noise modeling with conditional normalizing flows. ArXiv (2019)
2. Abdelhamed, A., Lin, S., Brown, M.S.: A high-quality denoising dataset for smart-phone cameras. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
3. Agustsson, E., Timofte, R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. pp. 1122–1131 (2017)
4. Aharon, M., Elad, M., Bruckstein, A.: K-SVD: An algorithm for designing over-complete dictionaries for sparse representation. IEEE Transactions on Signal Processing (TSP) **54**, 4311 – 4322 (2006)
5. Anaya, J., Barbu, A.: Renoir - a dataset for real low-light noise image reduction. Journal of Visual Communication and Image Representation **51**, 144–154 (2018)
6. Anwar, S., Barnes, N.: Real image denoising with feature attention. In: IEEE International Conference on Computer Vision (ICCV) (2019)
7. Anwar, S., Porikli, F., Huynh, C.P.: Category-specific object image denoising. IEEE Transactions on Image Processing **26**(11), 5506–5518 (2017)
8. Berthelot, D., Schumm, T., Metz, L.: Began: Boundary equilibrium generative adversarial networks. ArXiv [abs/1703.10717](https://arxiv.org/abs/1703.10717) (2017)
9. Brooks, T., Mildenhall, B., Xue, T., Chen, J., Sharlet, D., Barron, J.T.: Unprocessing images for learned raw denoising. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
10. Buades, A., Coll, B., Morel, J.M.: A non-local algorithm for image denoising. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 60–65 (2005)
11. Burger, H.C., Schuler, C.J., Harmeling, S.: Image denoising: Can plain neural networks compete with bm3d? pp. 2392–2399 (2012)
12. Bychkovsky, V., Paris, S., Chan, E., Durand, F.: Learning photographic global tonal adjustment with a database of input / output image pairs. In: The Twenty-Fourth IEEE Conference on Computer Vision and Pattern Recognition (2011)
13. Chen, F., Zhang, L., Yu, H.: External patch prior guided internal clustering for image denoising. In: IEEE International Conference on Computer Vision (ICCV). pp. 603–611 (2015)
14. Chen, J., Chen, J., Chao, H., Yang, M.: Image blind denoising with generative adversarial network based noise modeling. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
15. Chen, Y., Pock, T.: Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. IEEE Transactions on Pattern Analysis and Machine Intelligence **39**, 1–1 (08 2016)
16. Chen, Y., Yu, W., Pock, T.: On learning optimized reaction diffusion processes for effective image restoration. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015)
17. Chrysos, G.G., Kossaifi, J., Zafeiriou, S.: Robust conditional generative adversarial networks. In: International Conference on Learning Representations (ICLR) (2019)
18. Cogswell, M., Ahmed, F., Girshick, R., Zitnick, L., Batra, D.: Reducing overfitting in deep networks by decorrelating representations. In: International Conference on Learning Representations (ICLR) (2016)
19. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-D transform-domain collaborative filtering. IEEE Transactions on image processing **16**(8), 2080–2095 (2007)

20. Dong, W., Li, X., Zhang, L., Shi, G.: Sparsity-based image denoising via dictionary learning and structural clustering. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 674–697 (2011)
21. Dong, W., Zhang, L., Shi, G., Li, X.: Nonlocally centralized sparse representation for image restoration. *IEEE Transactions on Image Processing* **22**, 1620–1630 (2013)
22. Gonzalez, R.: *Digital Image Processing 2Nd Ed.* Prentice-Hall Of India Pvt. Limited (2002), <https://books.google.co.uk/books?id=iyJOPgAACAAJ>
23. Gu, S., Li, Y., Gool, L.V., Timofte, R.: Self-guided network for fast image denoising. In: IEEE International Conference on Computer Vision (ICCV) (2019)
24. Gu, S., Zhang, L., Zuo, W., Feng, X.: Weighted nuclear norm minimization with application to image denoising. *IEEE Conference on Computer Vision and Pattern Recognition* pp. 2862–2869 (2014)
25. Guo, S., Yan, Z., Zhang, K., Zuo, W., Zhang, L.: Toward convolutional blind denoising of real photographs. 2019 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
26. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
27. Healey, G., Kondepudy, R.: Radiometric ccd camera calibration and noise estimation. *IEEE Trans. Pattern Anal. Mach. Intell.* **16**, 267–276 (1994)
28. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017)
29. Izadi, S., Mirikharaji, Z., Zhao, M., Hamarneh, G.: Whitenner - blind image denoising via noise whiteness priors. In: International Conference on Computer Vision workshop on Visual Recognition for Medical Images (ICCV VRMI) (2019)
30. Jain, V., Sebastian, S.: Natural image denoising with convolutional networks. In: *Advances in Neural Information Processing Systems*. pp. 769–776 (2009)
31. Jia, X., Liu, S., Feng, X., Zhang, L.: Focnet: A fractional optimal control network for image denoising. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
32. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
33. Kostadin, K., Foi, A., Katkovnik, V., Egiazarian, K.: BM3D image denoising with shape-adaptive principal component analysis (2009)
34. Krull, A., Buchholz, T.O., Jug, F.: Noise2void - learning denoising from single noisy images. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
35. Lebrun, M., Buades, A., Morel, J.M.: A nonlocal bayesian image denoising algorithm. *SIAM Journal on Imaging Sciences* **6**, 1665–1688 (2013)
36. Lebrun, M., Colom, M., Morel, J.M.: The noise clinic: a blind image denoising algorithm. *Image Processing On Line (IPOL)* **5**, 1–54 (2015)
37. Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T.: Noise2Noise: Learning image restoration without clean data. In: International Conference on Machine Learning (ICML). pp. 2965–2974 (2018)
38. Liu, C., Szeliski, R., Kang, S.B., Zitnick, C.L., Freeman, W.T.: Automatic estimation and removal of noise from a single image. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **30**, 299–314 (2008)
39. Liu, D., Wen, B., Fan, Y., Loy, C., Huang, T.S.: Non-local recurrent network for image restoration. In: International Conference on Neural Information Processing Systems (NIPS). pp. 1680–1689 (2018)

40. Liu, P., Zhang, H., Zhang, K., Lin, L., , Zuog, W.: Multi-level wavelet-cnn for image restoration. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). p. 773782 (2018)
41. Mao, X., Shen, C., Yang, Y.B.: Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In: Advances in Neural Information Processing Systems. pp. 2802–2810 (2016)
42. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: IEEE International Conference on Computer Vision (ICCV). vol. 2, pp. 416–423 (July 2001)
43. Mildenhall, B., Barron, J.T., Chen, J., Sharlet, D., Ng, R., Carroll, R.: Burst Denoising With Kernel Prediction Networks. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
44. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014)
45. Nam, S., Hwang, Y., Matsushita, Y., Kim, S.J.: A holistic approach to cross-channel image noise modeling and its application to image denoising. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 1683–1691 (2016)
46. Plotz, T., Roth, S.: Benchmarking denoising algorithms with real photographs. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 2750–2759 (2017)
47. Plötz, T., Roth, S.: Neural nearest neighbors networks. In: Advances in Neural Information Processing Systems (NeurIPS) (2018)
48. Rasmus, A., Berglund, M., Honkala, M., Valpola, H., Raiko, T.: Semi-supervised learning with ladder networks. In: International Conference on Neural Information Processing Systems (NIPS). pp. 3546–3554 (2015)
49. Roth, S., Black, M.J.: Fields of experts. International Journal of Computer Vision **82**(2), 205–229 (2009)
50. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. In: International Conference on Neural Information Processing Systems (NIPS). pp. 2234–2242 (2016)
51. Tai, Y., Yang, J., Liu, X., Xu, C.: Memnet: A persistent memory network for image restoration. In: IEEE International Conference on Computer Vision (ICCV). pp. 4549–4557 (2017)
52. Tappen, M.F., Liu, C., Adelson, E.H., Freeman, W.T.: Learning gaussian conditional random fields for low-level vision. pp. 1–8 (07 2007)
53. Tripathi, S., Lipton, Z.C., Nguyen, T.Q.: Correction by projection: Denoising images with generative adversarial networks. ArXiv [abs/1803.04477](https://arxiv.org/abs/1803.04477) (2018)
54. Weiss, Y., Freeman, W.: What makes a good model of natural images? pp. 1 – 8 (07 2007)
55. Xu, J., Osher, S.: Iterative regularization and nonlinear inverse scale space applied to wavelet-based denoising. IEEE Transactions on Image Processing **16**, 534–544 (2007)
56. Xu, J., Zhang, L., Zhang, D.: A trilateral weighted sparse coding scheme for real-world image denoising. In: European Conference on Computer Vision (ECCV). pp. 21–38 (2018)
57. Xu, J., Zhang, L., Zuo, W., Zhang, D., Feng, X.: Patch group based nonlocal self-similarity prior learning for image denoising. IEEE International Conference on Computer Vision (ICCV) pp. 244–252 (2015)

58. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing* **26**(7), 3142–3155 (2017)
59. Zhang, K., Zuo, W., Gu, S., Zhang, L.: Learning deep cnn denoiser prior for image restoration. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 2808–2817 (2017)
60. Zhang, K., Zuo, W., Zhang, L.: Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing (TIP)* **27**(9), 4608–4622 (2018)
61. Zoran, D., Weiss, Y.: From learning models of natural image patches to whole image restoration. In: *IEEE International Conference on Computer Vision (ICCV)*. pp. 479–486 (2011)