

# Joint Demosaicking and Denoising in the Wild: The Case of Training Under Ground Truth Uncertainty

Jierun Chen, Song Wen, S.-H. Gary Chan

Department of Computer Science and Engineering  
The Hong Kong University of Science and Technology, Hong Kong, China  
{jcheneh, swenaa, gchan}@cse.ust.hk

## Abstract

Image demosaicking and denoising are the two key fundamental steps in digital camera pipelines, aiming to reconstruct clean color images from noisy luminance readings. In this paper, we propose and study Wild-JDD, a novel learning framework for joint demosaicking and denoising in the wild. In contrast to previous works which generally assume the ground truth of training data is a perfect reflection of the reality, we consider here the more common imperfect case of ground truth uncertainty in the wild. We first illustrate its manifestation as various kinds of artifacts including zipper effect, color moire and residual noise. Then we formulate a two-stage data degradation process to capture such ground truth uncertainty, where a conjugate prior distribution is imposed upon a base distribution. After that, we derive an evidence lower bound (ELBO) loss to train a neural network that approximates the parameters of the conjugate prior distribution conditioned on the degraded input. Finally, to further enhance the performance for out-of-distribution input, we design a simple but effective fine-tuning strategy by taking the input as a weakly informative prior. Taking into account ground truth uncertainty, Wild-JDD enjoys good interpretability during optimization. Extensive experiments validate that it outperforms state-of-the-art schemes on joint demosaicking and denoising tasks on both synthetic and realistic raw datasets.

## Introduction

Modern digital cameras use a single sensor overlaid with a color filter array (CFA) to capture an image. This means that only one color channel’s value is recorded for each pixel location. Let  $N$  be the number of pixels in an image, the raw data acquisition process can be simply modeled as

$$\mathbf{x} = \mathbf{A}\mathbf{z} + \mathbf{n}, \quad (1)$$

where  $\mathbf{x} \in \mathbb{R}^N$  is a noisy raw data vector of luminance readings,  $\mathbf{A} \in \mathbb{R}^{N \times 3N}$  is a mosaicking operation,  $\mathbf{z} \in \mathbb{R}^{3N}$  is an unknown clean image with three color channels, and  $\mathbf{n} \in \mathbb{R}^N$  is a noise vector.

Before the final “cooked” image is ready for the users, the raw data undergoes a series of processing steps, known as the image processing pipeline. Among those, demosaicking and denoising (DM&DN) are two of the very early and

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



(a) Zipper effect. (b) Color moire. (c) Residual noise.

Figure 1: Imperfect ground truth examples (*electronic zoom-in recommended*): (a) A ground truth image from CBSD dataset (Arbeláez et al. 2011) suffering from zipper effect, an artificial jagged pattern around edges; (b) Color moire in an image from ImageNet dataset (Russakovsky et al. 2015). Such artifact appears as false coloring due to interpolation error; (c) Noticeable residual noise in the collected “clean” image from Renoir dataset (Anaya and Barbu 2018).

most crucial steps. Demosaicking aims to undo the mosaicking operation  $\mathbf{A}$  by interpolating the missing two-thirds of each pixel’s color channels, while denoising removes the inevitable noise  $\mathbf{n}$  from the measurement  $\mathbf{x}$ . Due to their modular property, substantial traditional literature takes them as independent tasks and executes them in a sequential manner. This yields potentially suboptimal performance, and inspires several works on jointly addressing the DM&DN tasks (Liu et al. 2020; Kokkinos and Lefkimmatis 2019; Tan et al. 2017a).

Among the joint DM&DN works, data-driven approaches (Liu et al. 2020; Tan et al. 2018; Kokkinos and Lefkimmatis 2018) have been shown more effective than applying handcrafted priors and filters. These approaches usually require a collection of paired data, which are the mosaicked noisy images  $\mathbf{x}$  and the demosaicked clean “ground truth” counterparts  $\mathbf{y}$ . However, it is often costly and tedious to collect a large amount of high quality real-life data. Furthermore, the collected  $\mathbf{y}$  is not perfect without artifacts or noise. We illustrate this in Figure 1. For demosaicking, many approaches (Syu, Chen, and Chuang 2018; Tan et al. 2017b) take the output from a camera pipeline as  $\mathbf{y}$ , possibly introducing artifacts like zipper effect or color moire in regions with rich textures and sharp edges. For denoising, the

“clean” images are often collected by either setting a low-ISO (Plotz and Roth 2017; Anaya and Barbu 2018) or averaging a set of repeated shots of the same scene (Abdelhamed, Lin, and Brown 2018), which still contain noticeable noise. Moreover, such denoising data collecting process usually assumes the captured objects to be perfectly still, or requires a precise spatial alignment and intensity calibration among a burst of images. Potential failure cases would introduce additional error into the collected dataset. Therefore, all these in-the-wild issues means that the “ground truth”  $\mathbf{y}$  deviates from the needed authentic  $\mathbf{z}$ , limiting the performance of DM&DN model.

To account for the fact that the collected ground truth  $\mathbf{y}$  is not a perfect reflection of  $\mathbf{z}$ , we propose Wild-JDD, a novel joint demosaicking and denoising learning framework to enable training under ground truth uncertainty. In Wild-JDD, we first formulate a two-stage data degradation process, where a conjugate prior distribution is imposed upon a base Gaussian distribution. Then, we derive an ELBO loss from a variational perspective. In this way, the optimization process is aware of the target uncertainty and prevents the trained neural network from over-fitting to those randomness errors. Beyond that, when the testing image falls outside of the training range, we further enhance the performance by regarding the input as a weakly informative prior.

Our main contributions are summarized as follows:

- We identify in existing DM&DN datasets the ground truth uncertainty issues, manifesting themselves as various artifacts in the wild, such as zipper effect, color moire and residual noise.
- We introduce a novel learning framework for joint demosaicking and denoising in the wild (Wild-JDD), where a two-stage data degradation and an ELBO loss are formulated for optimization. We also propose a simple but effective fine-tuning strategy for out-of-distribution input.
- Instead of simply generating a demosaicked clean image, networks instantiated from our framework are capable of estimating all the parameters involved in data degradation and reconstruction, which provides better interpretability of the optimization process.
- We conduct extensive experiments on both synthetic and realistic datasets. Quantitative and qualitative comparisons show that Wild-JDD substantially outperforms state-of-the-art works.

## Related Work

In this section, we review the most relevant DM&DN works from sequential processing to joint optimization, and from supervised learning to self-supervised fine-tuning.

Traditionally, demosaicking and denoising are performed sequentially in arbitrary order (Zhang et al. 2011; Akiyama, Tanaka, and Okutomi 2015; Zhang et al. 2014). However, demosaicking first would correlate the noise spatially, break the commonly used independent identically distributed (i.i.d.) assumption imposed on the noise modeling, and increase the difficulty of the following denoising step (Nam et al. 2016; Zhang et al. 2009). Another issue

arises if denoising is performed first, ending up with an over-smoothed result (Jin, Facciolo, and Morel 2020). To address the above problems, recent studies jointly consider DM&DN for better performance. Khashabi et al. proposed one of the very first joint approaches through a learned non-parametric random field, and published an MSR dataset for evaluation. Heide et al. embedded a non-local natural prior into a global primal-dual optimization. Klatzer et al. formulated a sequential energy minimization framework. However, these heuristics-based methods were outperformed by the deep-learning-based approaches. Gharbi et al. trained a neural network on millions of images to achieve better results and shorter running time. After that, new approaches were developed to extend the CNN’s capability in the field: first with more effective network blocks (Huang et al. 2018; Tan et al. 2018), or a cascade of refinement frameworks (Kokkinos and Lefkimmiatis 2019, 2018), then by relying on a variational deep image prior (Park et al. 2020), and exploiting density map guidance to adaptively recover regions with different frequencies (Liu et al. 2020). These learning-based methods have achieved state-of-art performances, but the assumption that the clean color image is perfect remains in doubt.

Data-driven approaches normally perform well if the testing image shares a similar distribution with the training data (Mohseni et al. 2020; Ehret et al. 2019). However, in practice, the noise type remains diversified and may fall outside the training range. For this reason, Lehtinen et al. introduced a pioneering “noise2noise” training strategy using pairs of noisy images. A similar “mosaic2mosaic” framework improved a joint DM&DN network by fine-tuning bursts of raw images (Ehret et al. 2019). Nonetheless, performance of these methods can be limited by inadequate shots of the same scene. Batson and Royer; Krull, Buchholz, and Jug tackled the problem by using only one realization for each image at the price of a performance drop. These insightful methods enable adaptive fine-tuning, but the quality or uncertainty level of the pseudo ground truth has not been taken into account.

Our method differs with previous works in that we do not ideally assume the collected “ground truth” as perfect data. Instead, we acknowledge the presence of various artifacts and consider the case of training a joint DM&DN network under the ground truth uncertainty.

## Wild-JDD Methodology

We start with a dataset  $\mathcal{D} = \{\mathbf{x}^{(i)}, \mathbf{y}^{(i)}\}_{i=1}^M$ , consisting of  $M$  pairs of images. Our goal is to learn a function that maps  $\mathbf{x}^{(i)}$  to its corresponding authentic ground truth  $\mathbf{z}^{(i)}$  (note that  $\mathbf{y}^{(i)}$  is an approximation of  $\mathbf{z}^{(i)}$ ). This goal is approached by two steps: (1) we formulate a two-stage data degradation to link up all the parameters involved; (2) we derive an expression of data likelihood for optimization. Moreover, a fine-tuning strategy is applied to deal with out-of-distribution inputs. To keep dimensional consistency across different variables, in the following, without otherwise stated,  $\mathbf{x}^{(i)}$  is first bilinearly interpolated to a color image  $\tilde{\mathbf{x}}^{(i)} \in \mathbb{R}^{3N}$ , e.g. interpolating a missing green channel

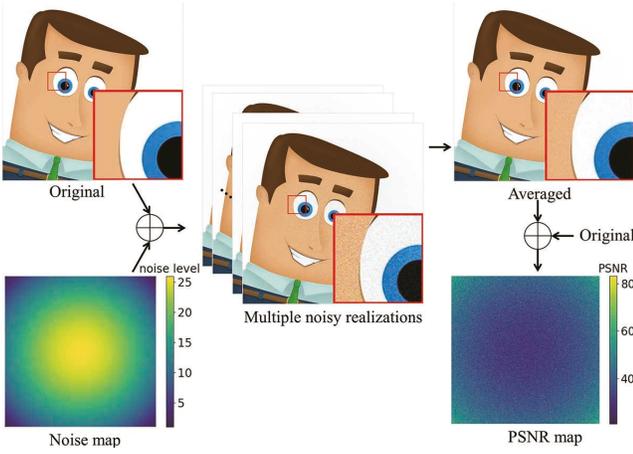


Figure 2: Visualization of the relation between the quality of collected clean image and the noise level. Firstly multiple (ten here) noisy realizations are generated given the original cartoon image and its spatial noise map. Then the relatively clean image is collected by averaging across those noisy realizations. The PSNR map between the averaged image and the original version indicates that those regions with higher noise levels correspond to lower PSNR values and higher uncertainties of the collected clean pixels.

by averaging its four green neighbors. The superscript ( $i$ ) is ignored in following subsections for simplicity.

### Two-stage Data Degradation

Conventionally, the pixel/channel-wise expression for data degradation is modelled with an additive Gaussian noise (Zhou et al. 2019; Jia et al. 2019) as

$$\tilde{x}_j | y_j, \sigma_j^2 \sim \mathcal{N}(y_j, \sigma_j^2), \quad (2)$$

where  $j = 1, 2, \dots, 3N$  specifies a dimension within an image. However,  $y_j$  is just a point estimator of the unknown authentic ground truth  $z_j$ , and training with  $y_j$  only achieves a suboptimal performance as variance of this estimator is not considered. As such, we adopt a brand-new degradation model as

$$\tilde{x}_j | z_j, \sigma_j^2 \sim \mathcal{N}(z_j, \sigma_j^2), \quad (3)$$

and seek to first parameterize the unknown authentic ground truth  $z_j$ . Suppose  $z_j$  follows a normal distribution with mean  $y_j$  and variance  $\sigma_j^2/\lambda$ :

$$z_j | y_j, \sigma_j^2, \lambda \sim \mathcal{N}(y_j, \sigma_j^2/\lambda), \quad (4)$$

where  $\sigma_j^2$  has an inverse gamma distribution parameterized by  $\alpha, \beta_j$ :

$$\sigma_j^2 | \alpha, \beta_j \sim \Gamma^{-1}(\alpha, \beta_j). \quad (5)$$

Then  $(z_j, \sigma_j^2)$  can be jointly denoted as a normal-inverse-gamma distribution:

$$(z_j, \sigma_j^2) | y_j, \lambda, \alpha, \beta_j \sim \text{N-}\Gamma^{-1}(y_j, \lambda, \alpha, \beta_j), \quad (6)$$

which serves as a conjugate prior distribution over the base distribution in Equation (3). The collected ground truth  $y_j$

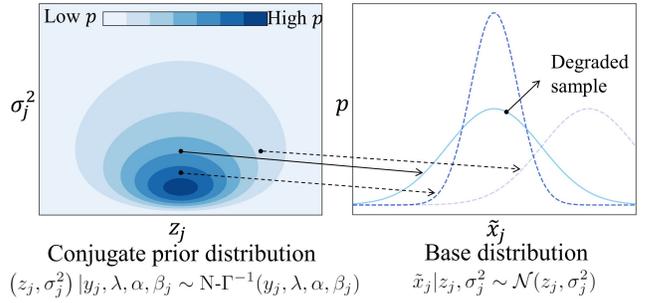


Figure 3: Two-stage data degradation corresponds to two sampling processes: first sampling from the conjugate prior distribution to obtain parameters of the base distribution; then sampling from the base distribution to obtain a degraded sample.

is an estimate of  $z_j$ , while parameter  $\lambda$  reflects the quality or uncertainty level of this estimation. For  $\alpha$  and  $\beta_j$ , they can be interpreted in a way that the variance  $\sigma_j^2$  is estimated from  $2\alpha$  observations with a sum of sample squared deviations  $2\beta_j$ . Similar to estimating the noise level by applying a Gaussian filter to the variance map (Yue et al. 2019), we parameterize  $\alpha$  and  $\beta_j$  as

$$\begin{aligned} \alpha &= \frac{w^2}{2}, \\ \beta_j &= \frac{w^2}{2} \mathcal{B}(\{(x_{j+t} - y_{j+t})^2\}_{t=-\lfloor w^2/2 \rfloor}^{\lfloor w^2/2 \rfloor}), \end{aligned} \quad (7)$$

where  $\mathcal{B}$  denotes a bilateral filtering operation on a variance map patch centered at pixel  $j$  with odd window size  $w$ . Note that  $\lambda$  and  $\alpha$  here are invariant to the dimension  $j$ . According to Equation (4), the variance of  $z_j$  or the uncertainty of  $y_j$  is proportional to the noise level  $\sigma_j^2$ . This can be reasonable, e.g., when collecting the clean ground truth by averaging multiple shots of the same scene, the noisier region should have lower quality and higher uncertainty (see Figure 2).

Therefore, based on equations (3) and (6), the data degradation process comprises two stages (see Figure 3): (1)  $z_j$  and  $\sigma_j^2$  take values by sampling from the prior distribution  $p(z_j, \sigma_j^2 | y_j, \lambda, \alpha, \beta_j)$ ; (2) then a degraded sample  $\tilde{x}_j$  is generated from the conditional distribution  $p(\tilde{x}_j | z_j, \sigma_j^2)$ .

### Maximizing ELBO

For the purpose of predicting latent variables  $\mathbf{z}, \boldsymbol{\sigma}^2$  given observed data point  $\tilde{\mathbf{x}}$ , or approximating the intractable posterior  $p(\mathbf{z}, \boldsymbol{\sigma}^2 | \tilde{\mathbf{x}})$ , an encoder  $q_w(\mathbf{z}, \boldsymbol{\sigma}^2 | \tilde{\mathbf{x}})$  is introduced with learnable weights  $w$ , as in the work (Kingma and Welling 2013). Notably, due to our use of a conjugate prior in Equation (6), this encoder  $q_w(\mathbf{z}, \boldsymbol{\sigma}^2 | \tilde{\mathbf{x}})$  is in the same probability distribution family as the prior, i.e.,

$$q_w(\mathbf{z}, \boldsymbol{\sigma}^2 | \tilde{\mathbf{x}}) = \prod_{i=1}^{3N} \text{N-}\Gamma^{-1}(\hat{y}_j, \hat{\lambda}_j, \hat{\alpha}_j, \hat{\beta}_j), \quad (8)$$

where  $\{\hat{y}_j, \hat{\lambda}_j, \hat{\alpha}_j, \hat{\beta}_j\}_{j=1}^{3N}$  is the output of a neural network. Note that here  $\hat{\lambda}_j$  is dimension-wise output learning the

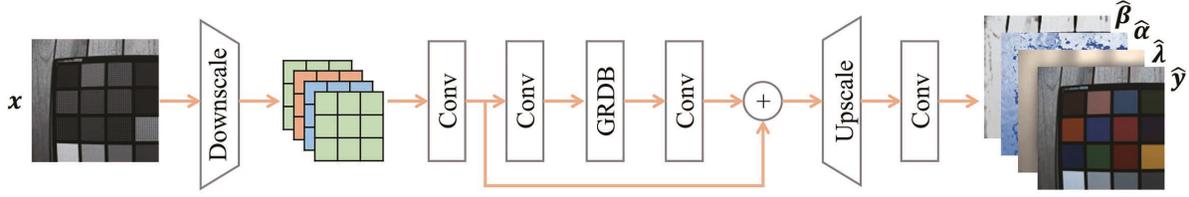


Figure 4: The overall architecture. The input was first downsampled into four color maps. Then a series of convolutional layers are applied to remove the noise and to interpolate missing color values. The output is obtained by applying an upscaling operation followed by an additional convolutional layer. This network learns a posterior distribution  $q_w(z, \sigma^2 | x) = \text{N-}\Gamma^{-1}(\hat{y}, \hat{\lambda}, \hat{\alpha}, \hat{\beta})$ .

same target  $\lambda$  for different  $j$ . The same is applicable to  $\hat{\alpha}_j$  and  $\alpha$ .

We next describe how we actually train such a neural network using a maximum likelihood estimation approach. As described in the work (Kingma and Welling 2013), the marginal log-likelihood can be calculated as

$$\log p(\tilde{x}) = D_{KL}(q_w(z, \sigma^2 | \tilde{x}) || p(z, \sigma^2 | \tilde{x})) + \mathcal{L}(w; \tilde{x}), \quad (9)$$

where the KL divergence term is a non-negative value. Therefore, maximizing the marginal log-likelihood  $\log p(\tilde{x})$  is converted to maximizing the second term, called the evidence lower bound (ELBO), and decomposed as

$$\mathcal{L}(w; \tilde{x}) = -D_{KL}(q_w(z, \sigma^2 | \tilde{x}) || p(z, \sigma^2)) + \mathbb{E}_{q_w(z, \sigma^2 | \tilde{x})} [\log p(\tilde{x} | z, \sigma^2)]. \quad (10)$$

This ELBO loss is maximized when: (1) the divergence term encourages the distribution returned by the encoder network close to the prior; (2) the expectation term guides the network predicting parameters with a high likelihood after seeing the corrupted image. In the work (Kingma and Welling 2013), the sampling process requires a reparameterization trick for gradient back-propagation. However, such kinds of tricks are unnecessary here because a closed form expression for  $\mathcal{L}(w; \tilde{x})$  can be derived analytically as follows:

$$\begin{aligned} & D_{KL}(q_w(z, \sigma^2 | \tilde{x}) || p(z, \sigma^2)) \\ &= \sum_{j=1}^{3N} \left\{ \frac{\lambda \hat{\alpha}_j}{2 \hat{\beta}_j} (y_j - \hat{y}_j)^2 + \frac{\lambda}{2 \hat{\lambda}_j} - \frac{1}{2} \log \frac{\lambda}{\hat{\lambda}_j} + \alpha_j \log \frac{\hat{\beta}_j}{\beta_j} \right. \\ & \quad \left. - \frac{1}{2} + \log \frac{\Gamma(\alpha_j)}{\Gamma(\hat{\alpha}_j)} + (\hat{\alpha}_j - \alpha_j) \psi(\hat{\alpha}_j) - (\hat{\beta}_j - \beta_j) \frac{\hat{\alpha}_j}{\hat{\beta}_j} \right\}, \quad (11) \end{aligned}$$

$$\begin{aligned} \mathbb{E}_{q(z, \sigma^2 | \tilde{x})} [\log p(\tilde{x} | z, \sigma^2)] &= \sum_{j=1}^{3N} \left\{ -\frac{\log 2\pi}{2} \right. \\ & \quad \left. - \frac{\log \hat{\beta}_j - \psi(\hat{\alpha}_j)}{2} - \frac{\hat{\beta}_j}{2 \hat{\lambda}_j^2 (\hat{\alpha}_j - 1)} - \frac{\hat{\alpha}_j (\tilde{x}_j - \hat{y}_j)^2}{2 \hat{\beta}_j} \right\}, \quad (12) \end{aligned}$$

where  $\Gamma(\cdot)$ ,  $\psi(\cdot)$  denote the Gamma and Di-gamma function respectively (detailed derivations are provided in the supplementary materials). Looking deeper into the term

$\frac{\lambda \hat{\alpha}_j}{2 \hat{\beta}_j} (y_j - \hat{y}_j)^2$  in Equation (11), we can notice that if the parameter  $\lambda_j$  is set to be large enough, the ELBO loss would degenerate to a mean squared error (MSE). When using an MSE loss, too much attention would be put into the restoration term  $(y_j - \hat{y}_j)^2$ , leaving the existence of ground truth uncertainty omitted and making the model biased to the training data. Therefore, from a variational point of view, WildJDD provides a sound interpretability of the reason why the restoration term and the rest regularization terms should co-exist in the training process.

After formulating the ELBO loss for each single image, the overall optimization objective is obtained by computing across the entire dataset:

$$\min_w \sum_{i=1}^M -\log p(\tilde{x}^{(i)}); \log p(\tilde{x}^{(i)}) \approx \mathcal{L}(w; \tilde{x}^{(i)}). \quad (13)$$

At test time, the desired demosaicked clean image is obtained by taking the expectation of  $z$ , i.e.  $\mathbb{E}[z] = \hat{y}$ , while the noise map is parameterized as  $\mathbb{E}[\sigma^2] = \frac{\hat{\beta}}{\hat{\alpha} - 1}$ , according to the definition of  $\text{N-}\Gamma^{-1}$  distribution.

### Corrupted Input as a Weakly Informative Prior

Data-driven approaches are generally promising when a test image shares similar characteristics with the training set. However, their performances are limited when the input is considerably different, e.g. having a different type of noise. Inspired by the ‘‘noise2noise’’ (Lehtinen et al. 2018) and ‘‘mosaic2mosaic’’ (Ehret et al. 2019) algorithms, we further improve our trained model by taking the corrupted input as a weakly informative prior, i.e. replacing  $y_j$  by  $\tilde{x}_j$  during fine-tuning and using a smaller  $\lambda$  to indicate the increased uncertainty. However, it comes with an underlying problem: the network may merely learn an identity mapping function, i.e. predicting  $\tilde{x}_j$  given  $\tilde{x}_j$ .

It can be observed that in smooth regions, pixels share a strong spatial correlation. Therefore, we tackle the above issue by an alternative scheme, replacing  $y_j$  by a random pixel  $\tilde{x}_{j+t}$  in a small patch  $\{\tilde{x}_{j+t}\}_{t=-\lfloor p^2/2 \rfloor}^{\lfloor p^2/2 \rfloor}$  centered at pixel  $\tilde{x}_j$ , where  $p$  denotes the patch size. It is possible that  $\tilde{x}_{j+t}$  differs a lot to  $\tilde{x}_j$  in texture-rich regions, requiring another step to exclude the outlying  $\tilde{x}_{j+t}$ . We achieve this by applying a simple filter: if the value  $\tilde{x}_{j+t}$  falls outside the confidence interval  $(\tilde{x}_j - 2\sigma_j, \tilde{x}_j + 2\sigma_j)$ , this informative prior  $\tilde{x}_{j+t}$  is

Method	$\sigma$	Kodak (24 images)		McMaster (18 images)		WED-CDM (100 images)		MIT moire (1000 images)		Urban100 (100 images)	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
FlexISP		31.31	0.8694	31.17	0.8627	31.08	0.8754	29.06	0.8206	30.37	0.8832
SEM		34.59	0.9269	32.36	0.8869	32.85	0.9234	27.46	0.8292	27.19	0.7813
ADMM		31.60	0.8787	32.63	0.8966	31.79	0.9003	28.58	0.7923	28.57	0.8578
DeepJoint	5	36.11	0.9455	35.47	0.9378	35.09	0.9485	31.82	0.9015	34.04	0.9510
Kokkinos		36.22	0.9426	34.74	0.9252	35.12	0.9410	31.94	0.8882	34.07	0.9358
SGNet		-	-	-	-	-	-	32.15	<b>0.9043</b>	34.54	0.9533
Wild-JDD		36.88	<i>0.9520</i>	35.85	<i>0.9425</i>	35.92	<i>0.9543</i>	32.29	0.8987	<i>34.70</i>	<i>0.9534</i>
Wild-JDD*		<b>36.97</b>	<b>0.9526</b>	<b>35.94</b>	<b>0.9435</b>	<b>36.01</b>	<b>0.9551</b>	<b>32.39</b>	<i>0.8999</i>	<b>34.83</b>	<b>0.9540</b>
FlexISP		28.64	0.7583	28.51	0.7534	28.24	0.7691	26.61	0.7491	27.51	0.8196
SEM		29.78	0.7681	28.68	0.7306	28.90	0.7563	25.45	0.7531	25.36	0.7094
ADMM		31.04	0.8595	31.72	0.8699	30.90	0.8758	28.26	0.7720	27.48	0.8388
DeepJoint	10	33.10	0.9018	33.18	0.9047	32.69	0.9156	29.75	0.8561	31.60	0.9152
Kokkinos		33.32	0.9022	32.75	0.8956	32.76	0.9066	30.01	0.8123	31.73	0.8912
SGNet		-	-	-	-	-	-	30.09	0.8619	32.14	0.9229
Wild-JDD		<i>33.81</i>	<i>0.9127</i>	<i>33.53</i>	<i>0.9123</i>	<i>33.44</i>	<i>0.9244</i>	<i>30.30</i>	<i>0.8645</i>	<i>32.42</i>	<i>0.9288</i>
Wild-JDD*		<b>33.88</b>	<b>0.9136</b>	<b>33.61</b>	<b>0.9137</b>	<b>33.51</b>	<b>0.9255</b>	<b>30.37</b>	<b>0.8657</b>	<b>32.54</b>	<b>0.9299</b>
FlexISP		26.67	0.6541	26.55	0.6572	26.24	0.6694	24.91	0.6851	25.55	0.7642
SEM		25.79	0.5954	25.45	0.5800	25.46	0.5799	23.23	0.6527	23.25	0.6156
ADMM		30.16	0.8384	30.50	0.8412	29.85	0.8497	27.58	0.7497	28.37	0.8440
DeepJoint	15	31.25	0.8603	31.49	0.8707	30.99	0.8823	28.22	0.8088	29.73	0.8802
Kokkinos		31.28	0.8674	30.98	0.8605	30.94	0.8710	28.28	0.7693	29.87	0.8451
SGNet		-	-	-	-	-	-	28.60	0.8188	30.37	0.8923
Wild-JDD		<i>31.92</i>	<i>0.8765</i>	<i>31.90</i>	<i>0.8846</i>	<i>31.75</i>	<i>0.8965</i>	<i>28.89</i>	<i>0.8310</i>	<i>30.79</i>	<i>0.9055</i>
Wild-JDD*		<b>31.99</b>	<b>0.8777</b>	<b>31.97</b>	<b>0.8863</b>	<b>31.82</b>	<b>0.8979</b>	<b>28.95</b>	<b>0.8325</b>	<b>30.89</b>	<b>0.9070</b>

Table 1: Comparison against state-of-the-art works on five datasets. The parameter  $\sigma$  indicates the noise level of inputs corrupted by additive white Gaussian noise. The best and second best results are in bold and Italic, respectively. Note that for SGNet, the code is not released publicly and their results on Kodak, McMaster and WED-CDM datasets are not reported in their paper.

masked from computing the fine-tuning ELBO loss, i.e.,

$$\mathcal{L}_{ft} = \sum_{j=1}^{3N} \mathbb{1}(\tilde{x}_{j+t} \in (\tilde{x}_j - 2\sigma_j, \tilde{x}_j + 2\sigma_j)) \mathcal{L}_j(\mathbf{w}; \tilde{\mathbf{x}}), \quad (14)$$

where  $\mathbb{1}(\cdot)$  denotes an indicator function, and  $\mathcal{L}_j(\mathbf{w}; \tilde{\mathbf{x}})$  can be computed with  $j$ -indexed components in Equation (10), (11), (12) after replacing  $y_j$  by  $\tilde{x}_{j+t}$ .

## Illustrative Experimental Results

To show the effectiveness of our framework, we conduct extensive experiments with both synthetic datasets and realistic raw data. We focus on the Bayer pattern, which has been the dominating choice among various CFA patterns.

### Network Architecture

In previous section,  $q_{\mathbf{w}}(\mathbf{z}, \sigma^2 | \tilde{\mathbf{x}})$  represents the network taking  $\tilde{\mathbf{x}}$  as the input. However, the original input is actually  $\mathbf{x}$ , and the bilinear interpolation process from  $\mathbf{x}$  to  $\tilde{\mathbf{x}}$  can be considered as part of the job done by the network. Therefore, our network is trained to learn a mapping function  $q_{\mathbf{w}}(\mathbf{z}, \sigma^2 | \mathbf{x})$ . We use a light-weight network architecture as shown in Figure 4. The GRDB building module refers to a grouped Residual Dense Block (Zhang et al. 2018), consisting of dense connected layers and a local feature fusion.

A downscaling layer is positioned at the first layer to rearrange a mosaicked input to four quarter-resolution color maps. This rearrangement helps to save memory and speed up the training. Each of the first three convolution layers has 64 filters with  $3 \times 3$  kernel size. After that, an upscaling layer is used to unpack the features back to full-resolution. The last convolution layer produces 12 feature maps with  $3 \times 3$  kernel size. These 12 feature maps correspond to four parameters  $\hat{\mathbf{y}}, \hat{\lambda}, \hat{\alpha}, \hat{\beta}$ , each of which has 3 maps. The network is implemented using PyTorch framework.

### Experiments on Synthetic Datasets

We first compare our method with previous works on the synthetic datasets in sRGB space, following the convention without an inverse ISP processing (Liu et al. 2020; Klatzer et al. 2016). In this experiment, 800 high-resolution images from DIV2K (Timofte et al. 2017) and 2650 high-resolution images from Flickr2K (Lim et al. 2017) are used for training. These images are randomly cropped into  $120 \times 120$  patches with a batch size of 128. After the augmentation by flipping and rotation, the noisy mosaicked inputs are generated by applying the Bayer pattern sampling and adding random Gaussian noise in the range of  $[0, 20]$ . Unlike most denoising works assuming the i.i.d. noise, which deviates from the practical application, we adopt non-i.i.d Gaussian modeling with spatially variant noise levels, following the work (Yue

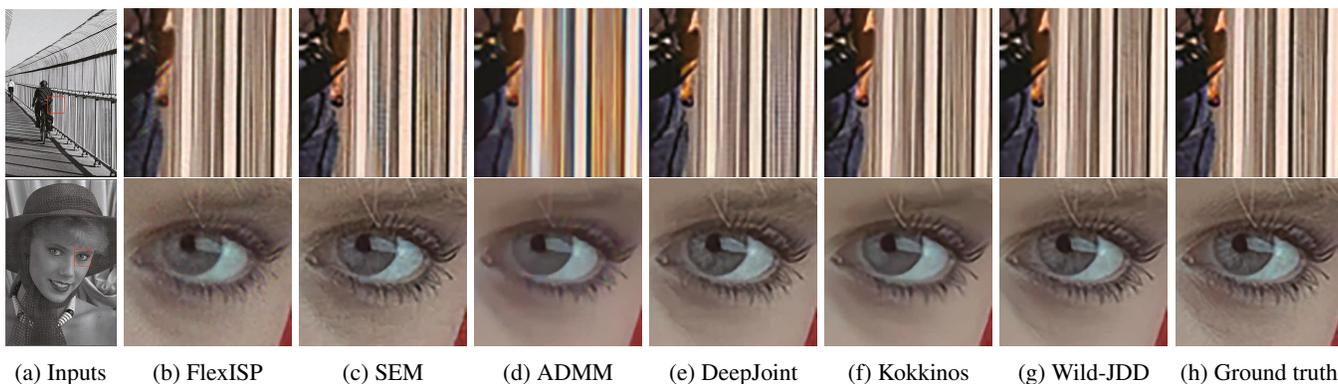


Figure 5: Visual comparison of our method against competing related works. Our reconstructions preserve texture details of high quality without introducing noticeable moire or zipper artifacts.

et al. 2019). During network training, we empirically set the parameter  $\lambda$  as  $2e3$  and  $\alpha$  as 180.5 (window size  $w$  as 19). The Adam optimizer (Kingma and Ba 2014) is used. The learning rate is initialized as  $5e-4$ , reduced by a factor of 0.8 when the training meets a plateau in PSNR, with a minimum value of  $1e-4$ . The whole training process takes around 5 days on a single RTX 2080Ti GPU.

For testing, five widely-used benchmark datasets are used, including: Kodak<sup>1</sup>, McMaster (Zhang et al. 2011), WED-CDM (Tan et al. 2017b), MIT moire (Gharbi et al. 2016) and Urban100 (Huang, Singh, and Ahuja 2015). These datasets are collected from various devices under different scenarios. Note that the ground truth in these datasets are not perfect either. However, if one method can consistently outperform others across various datasets, its effectiveness can still be validated and approved. For comparison, six existing state-of-the-art works on joint DM&DN task are adopted, including: FlexISP (Heide et al. 2014), SEM (Klatzer et al. 2016), ADMM (Tan et al. 2017a), DeepJoint (Gharbi et al. 2016), Kokkinos (Kokkinos and Lefkimmiatis 2018) and SNet (Liu et al. 2020). We run their source code for evaluation or directly cite their reported performance if the code is not available. The results are reported in both PSNR and SSIM matrices listed in Table 1.

Overall, our method outperforms all other works quantitatively, though our method is trained for non-i.i.d. noise cases. For both DeepJoint and SNet, they assume an accurate noise map as known input, which is not reasonable in practice. In contrast, our method is able to perform a truly-blind reconstruction without such a known noise map. To further improve the performance, we adopt a self-ensemble strategy by applying the flipping and rotation on the input to generate 8 augmented inputs. After being processed by the network, 8 outputs are obtained and transformed back to the original geometry, followed by an averaging to get a unified final output. Note that the augmentation on input would break its Bayer pattern, e.g. from RGGB to GRBG. Therefore, a Bayer Preserving Unification is utilized by padding and cropping the image borders (Liu et al. 2019). We denote

<sup>1</sup><http://r0k.us/graphics/kodak>

Method	linear		sRGB	
	PSNR	SSIM	PSNR	SSIM
JMCDM	37.44	0.971	31.35	0.942
RTF	37.77	0.976	31.77	0.951
FlexISP	38.28	0.974	31.76	0.941
SEM	38.93	0.980	32.93	<b>0.960</b>
DeepJoint	38.61	0.963	32.58	0.913
Kokkinos	39.29	0.975	33.37	0.930
MMNet20	40.07	0.979	34.24	0.942
DMCNN-VD	38.33	0.968	32.00	0.920
DMCNN-VD-Tr	40.07	<b>0.981</b>	34.08	0.957
Wild-JDD	40.16	0.980	34.34	0.945
Wild-JDD*	<b>40.36</b>	<b>0.981</b>	<b>34.59</b>	0.947

Table 2: Evaluation on realistic raw data. Our network is trained once using linear data and evaluated on both linear and sRGB space.

our method with self-ensemble as Wild-JDD\*. Qualitative comparison is also provided in Figure 5. Our reconstructions remove the noise and preserve details pretty well without introducing noticeable artifacts, while other works tend to produce color moire in high-frequency regions. Although Kokkinos (Kokkinos and Lefkimmiatis 2018) also has good immunity to those artifacts, the produced images are over-smoothed due to their iterative processing properties.

## Experiments on Realistic Raw Data

In the previous experiment, we trained and evaluated on sRGB data to enable more comparison with other works. However, Khashabi et al. suggested that the evaluation should also be conducted on the raw data and thus proposed a realistic MSR 16-bits benchmark dataset. We retrain our network on their Linear Bayer Panasonic set with 200 images, in the same parameter setting to previous experiment. Table 2 reports our overall better performance in both linear and sRGB space compared to other representative works, including JMCDM (Chang, Ding, and Li 2015), RTF (Khashabi et al. 2014), FlexISP (Heide et al. 2014), SEM (Klatzer et al. 2016), DeepJoint (Gharbi et al. 2016), Kokki-

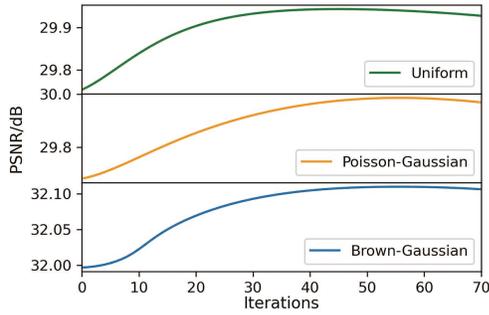


Figure 6: Increasing PSNR values when fine-tuning for different corrupted inputs. For each iteration, the updated PSNR values are obtained by averaging them across the McM dataset.

Method	Kodak		McMaster		WED-CDM	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
MSE	33.77	0.9124	33.45	0.9096	33.37	0.9233
ELBO	<b>33.81</b>	<b>0.9127</b>	<b>33.53</b>	<b>0.9123</b>	<b>33.44</b>	<b>0.9244</b>

Table 3: Comparison of ELBO against the MSE on synthetic datasets with noise level  $\sigma = 10$ .

nos (Kokkinos and Lefkimmiatis 2018), MMNet20 (Kokkinos and Lefkimmiatis 2019), DMCNN-VD and DMCNN-VD-Tr (Syu, Chen, and Chuang 2018).

### Fine-tuning Out-of-distribution Input

To examine the effectiveness of our fine-tuning strategy, three types of noise are considered, including Uniform, Poisson-Gaussian and Brown-Gaussian. Implementation details are similar to previous experiments except that parameter  $\lambda$  is set to be a smaller value 1, the learning rate decreased to  $2e-6$  and  $p$  is empirically set as 3. The results in Figure 6 show a 0.1~0.3 dB PSNR improvement as the number of iterations increases until roughly 50. Notably, with too many iterations, the performance would drop from its peak. This concern could be eased with the help of no-reference image quality assessment tools (Xu, Jiang, and Min 2017).

### Ablation Study

**ELBO versus MSE** When setting the parameter  $\lambda$  to be large enough, our ELBO loss would degrade to a commonly used MSE loss, which assumes the dataset ground truth to be a perfect target. The superiority of using ELBO against MSE loss is validated in Table 3, where a consistent PSNR improvement can be observed. This slight improvement comes from capturing the mild uncertainty embedded in the collected ground truth during training.

We also conduct an additional experiment to compare their optimization process. We take a single image from the Cartoon Set (Royer et al. 2020) as  $z$ . Then a Bayer pattern mosaicking and AWGN noise are applied to this image, followed by a bilinear interpolation to obtain a corrupted version  $\tilde{x}$ . This corrupted image  $\tilde{x}$  is set to be the learning objective of a neural work. As described in the work Deep Image Prior (Ulyanov, Vedaldi, and Lempitsky 2018), a neural

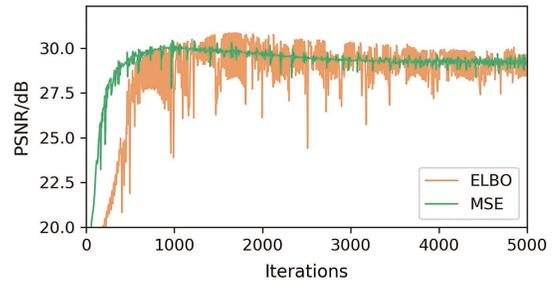


Figure 7: Comparison of ELBO against MSE on learning a cartoon image.

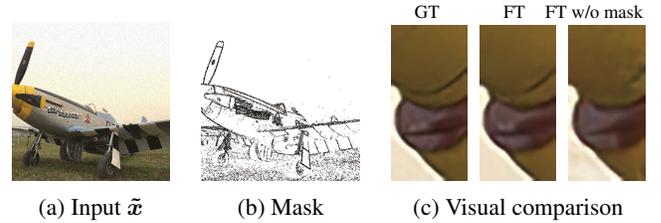


Figure 8: Fine-tuning with and without the mask.

network tends to learn the clean signal faster than learning the random noise. Therefore, we can see that in Figure 7, the PSNR curve of MSE increases and then decreases. When ELBO comes into play, its PSNR curve fluctuates roughly around the curve of MSE. This fluctuation results from the interaction between ELBO’s restoration term and ELBO’s regularization terms. This is consistent to our expectation, that the neural network is aware of the uncertainty affiliated with the target  $\tilde{x}$  instead of treating  $\tilde{x}$  as the absolute learning target. Therefore, training with our ELBO loss can achieve a higher intermediate peak of the PSNR curve.

**Effect of the mask** Using the mask during fine-tuning can effectively avoid edges getting blurred. As shown in Figure 8 (b), a mask map generated by our simple confidence interval scheme is able to outline those edges. In Figure 8 (c), fine-tuning with such a mask preserves the sharp edges more faithfully than fine-tuning without the mask.

### Conclusion

We have presented the Wild-JDD, a novel learning framework for joint demosaicking and denoising tasks. We identify the ground truth uncertainty issues, formulate a two-stage data degradation process and derive an ELBO loss for optimization. We also propose a simple but effective fine-tuning strategy for out-of-distribution input. Comprehensive experiments demonstrate the effectiveness of our method. Wild-JDD not only outperforms state-of-the-art solutions in terms of both statistical and perceptual quality by a clear margin, but also provides good interpretability, where the restoration term and the rest regularization terms coexist to account for the learning target uncertainty. We hope that Wild-JDD will inspire more future research to study the effective training under the ground truth uncertainty in image-to-image translation tasks.

## Acknowledgements

This work was supported, in part, by Hong Kong General Research Fund under grant number 16200120.

## References

- Abdelhamed, A.; Lin, S.; and Brown, M. S. 2018. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1692–1700.
- Akiyama, H.; Tanaka, M.; and Okutomi, M. 2015. Pseudo four-channel image denoising for noisy CFA raw data. In *2015 IEEE International Conference on Image Processing (ICIP)*, 4778–4782. IEEE.
- Anaya, J.; and Barbu, A. 2018. RENOIR-A benchmark dataset for real noise reduction evaluation. *Journal of Visual Communication and Image Representation* 144–154.
- Arbeláez, P.; Maire, M.; Fowlkes, C.; and Malik, J. 2011. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence* 33(5): 898–916.
- Batson, J.; and Royer, L. 2019. Noise2Self: Blind Denoising by Self-Supervision. In *International Conference on Machine Learning*, 524–533.
- Chang, K.; Ding, P. L. K.; and Li, B. 2015. Color image demosaicking using inter-channel correlation and nonlocal self-similarity. *Signal Processing: Image Communication* 39: 264–279.
- Ehret, T.; Davy, A.; Arias, P.; and Facciolo, G. 2019. Joint Demosaicking and Denoising by Fine-Tuning of Bursts of Raw Images. In *Proceedings of the IEEE International Conference on Computer Vision*, 8868–8877.
- Gharbi, M.; Chaurasia, G.; Paris, S.; and Durand, F. 2016. Deep joint demosaicking and denoising. *ACM Transactions on Graphics (TOG)* 35(6): 1–12.
- Heide, F.; Steinberger, M.; Tsai, Y.-T.; Rouf, M.; Pajak, D.; Reddy, D.; Gallo, O.; Liu, J.; Heidrich, W.; Egiazarian, K.; et al. 2014. Flexisp: A flexible camera image processing framework. *ACM Transactions on Graphics (TOG)* 33(6): 1–13.
- Huang, J.-B.; Singh, A.; and Ahuja, N. 2015. Single image super-resolution from transformed self-exemplars. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5197–5206.
- Huang, T.; Wu, F. F.; Dong, W.; Shi, G.; and Li, X. 2018. Lightweight Deep Residue Learning for Joint Color Image Demosaicking and Denoising. In *2018 24th International Conference on Pattern Recognition (ICPR)*, 127–132. IEEE.
- Jia, X.; Liu, S.; Feng, X.; and Zhang, L. 2019. FOCNet: A fractional optimal control network for image denoising. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6054–6063.
- Jin, Q.; Facciolo, G.; and Morel, J.-M. 2020. A Review of an Old Dilemma: Demosaicking First, or Denoising First? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 514–515.
- Khashabi, D.; Nowozin, S.; Jancsary, J.; and Fitzgibbon, A. W. 2014. Joint demosaicing and denoising via learned nonparametric random fields. *IEEE Transactions on Image Processing* 23(12): 4968–4981.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Klatzer, T.; Hammernik, K.; Knobelreiter, P.; and Pock, T. 2016. Learning joint demosaicing and denoising based on sequential energy minimization. In *2016 IEEE International Conference on Computational Photography (ICCP)*, 1–11. IEEE.
- Kokkinos, F.; and Lefkimmiatis, S. 2018. Deep image demosaicking using a cascade of convolutional residual denoising networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 303–319.
- Kokkinos, F.; and Lefkimmiatis, S. 2019. Iterative joint image demosaicking and denoising using a residual denoising network. *IEEE Transactions on Image Processing* 28(8): 4177–4188.
- Krull, A.; Buchholz, T.-O.; and Jug, F. 2019. Noise2void-learning denoising from single noisy images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2129–2137.
- Lehtinen, J.; Munkberg, J.; Hasselgren, J.; Laine, S.; Karras, T.; Aittala, M.; and Aila, T. 2018. Noise2noise: Learning image restoration without clean data. *arXiv preprint arXiv:1803.04189*.
- Lim, B.; Son, S.; Kim, H.; Nah, S.; and Mu Lee, K. 2017. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 136–144.
- Liu, J.; Wu, C.-H.; Wang, Y.; Xu, Q.; Zhou, Y.; Huang, H.; Wang, C.; Cai, S.; Ding, Y.; Fan, H.; et al. 2019. Learning raw image denoising with bayer pattern unification and bayer preserving augmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 0–0.
- Liu, L.; Jia, X.; Liu, J.; and Tian, Q. 2020. Joint Demosaicing and Denoising With Self Guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2240–2249.
- Mohseni, S.; Pitale, M.; Yadawa, J.; and Wang, Z. 2020. Self-Supervised Learning for Generalizable Out-of-Distribution Detection. In *AAAI*, 5216–5223.
- Nam, S.; Hwang, Y.; Matsushita, Y.; and Joo Kim, S. 2016. 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*, 1683–1691.
- Park, Y.; Lee, S.; Jeong, B.; and Yoon, J. 2020. Joint Demosaicing and Denoising Based on a Variational Deep Image Prior Neural Network. *Sensors* 20(10): 2970.

- Plotz, T.; and Roth, S. 2017. Benchmarking denoising algorithms with real photographs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1586–1595.
- Royer, A.; Bousmalis, K.; Gouws, S.; Bertsch, F.; Mosseri, I.; Cole, F.; and Murphy, K. 2020. Xgan: Unsupervised image-to-image translation for many-to-many mappings. In *Domain Adaptation for Visual Understanding*, 33–49. Springer.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115(3): 211–252.
- Syu, N.-S.; Chen, Y.-S.; and Chuang, Y.-Y. 2018. Learning deep convolutional networks for demosaicing. *arXiv preprint arXiv:1802.03769*.
- Tan, H.; Xiao, H.; Lai, S.; Liu, Y.; and Zhang, M. 2018. Deep residual learning for image demosaicing and blind denoising. *Pattern Recognit. Lett.*
- Tan, H.; Zeng, X.; Lai, S.; Liu, Y.; and Zhang, M. 2017a. Joint demosaicing and denoising of noisy bayer images with admm. In *2017 IEEE International Conference on Image Processing (ICIP)*, 2951–2955. IEEE.
- Tan, R.; Zhang, K.; Zuo, W.; and Zhang, L. 2017b. Color image demosaicking via deep residual learning. In *Proc. IEEE Int. Conf. Multimedia Expo (ICME)*, 793–798.
- Timofte, R.; Agustsson, E.; Van Gool, L.; Yang, M.-H.; and Zhang, L. 2017. Ntire 2017 challenge on single image super-resolution: Methods and results. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 114–125.
- Ulyanov, D.; Vedaldi, A.; and Lempitsky, V. 2018. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9446–9454.
- Xu, S.; Jiang, S.; and Min, W. 2017. No-reference/blind image quality assessment: a survey. *IETE Technical Review* 34(3): 223–245.
- Yue, Z.; Yong, H.; Zhao, Q.; Meng, D.; and Zhang, L. 2019. Variational denoising network: Toward blind noise modeling and removal. In *Advances in neural information processing systems*, 1690–1701.
- Zhang, L.; Lukac, R.; Wu, X.; and Zhang, D. 2009. PCA-based spatially adaptive denoising of CFA images for single-sensor digital cameras. *IEEE transactions on image processing* 18(4): 797–812.
- Zhang, L.; Wu, X.; Buades, A.; and Li, X. 2011. Color demosaicking by local directional interpolation and nonlocal adaptive thresholding. *Journal of Electronic imaging* 20(2): 023016.
- Zhang, X.; Sun, M.-T.; Fang, L.; and Au, O. C. 2014. Joint Denoising and demosaicking of noisy CFA images based on inter-color correlation. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5784–5788. IEEE.
- Zhang, Y.; Tian, Y.; Kong, Y.; Zhong, B.; and Fu, Y. 2018. Residual dense network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2472–2481.
- Zhou, Y.; Jiao, J.; Huang, H.; Wang, J.; and Huang, T. 2019. Adaptation strategies for applying awgn-based denoiser to realistic noise. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 10085–10086.