

BSRAW: Improving Blind RAW Image Super-Resolution

Marcos V. Conde, Florin Vasluianu, Radu Timofte
 Computer Vision Lab, CAIDAS & IFI, University of Würzburg
 {marcos.conde, radu.timofte}@uni-wuerzburg.de
<https://github.com/mv-lab/AISP>

Abstract

In smartphones and compact cameras, the Image Signal Processor (ISP) transforms the RAW sensor image into a human-readable sRGB image. Most popular super-resolution methods depart from a sRGB image and upscale it further, improving its quality. However, modeling the degradations in the sRGB domain is complicated because of the non-linear ISP transformations. Despite this known issue, only a few methods work directly with RAW images and tackle real-world sensor degradations.

We tackle blind image super-resolution in the RAW domain. We design a realistic degradation pipeline tailored specifically for training models with raw sensor data. Our approach considers sensor noise, defocus, exposure, and other common issues. Our BSRAW models trained with our pipeline can upscale real-scene RAW images and improve their quality. As part of this effort, we also present a new DSLM dataset and benchmark for this task.

1. Introduction

The specialized on-board Image Signal Processor (ISP) transforms RAW sensor readings into more refined and perceptually meaningful RGB representations [20, 45, 52, 68]. Photographers often opt for RAW over RGB processing to yield perceptually superior images, attributing this preference to two key benefits of RAW data: (i) RAW information carries a broader data range, generally ranging from 12-14 bits, compared to the 8-bit RGB output from the ISP. (ii) RAW data linearly correlates with scene radiance, which simplifies the correction of degradations such as noise. Conversely, ISP operations are non-linear and undergo an irrevocable loss of information across various stages [6, 27], complicating subsequent image restoration [6, 54]. For these considerations, RAW imagery is a more advantageous starting point than RGB for a multitude of low-level vision applications like image denoising [1, 6, 36, 38, 59, 62], color balance [21], deblurring [30], exposure adjustment [11, 19, 69], and image super-resolution [42, 54, 60, 68].

Nevertheless, only a small fraction of low-vision research directly engages with RAW data, mainly due to the greater abundance and accessibility of general-purpose sRGB images. Consequently, the most popular deep learning architectures for image restoration predominantly operate on RGB images [13, 51, 61, 62]. The majority of cutting-edge Single Image Super-Resolution (SISR) techniques [4, 15, 26, 31, 63] rely on deep convolutional networks or Transformer [47] architectures, and operate in the RGB color space. Yet, these approaches have notable drawbacks. Primarily, they are often trained on artificially generated data, leading to poor generalization on real-world scenes. Secondly, modeling accurately the degradations in RGBs is challenging due to the non-linear transformations and information loss incurred during ISP processing (RAW to RGB).

The classical single-image super-resolution model [17, 54, 66] is formulated as:

$$\mathbf{y} = (\mathbf{x} \otimes \mathbf{k}) \downarrow_s + \mathbf{n} \quad (1)$$


It assumes the observed -or captured- low resolution (LR) image \mathbf{y} is obtained from an underlying high resolution (HR) image \mathbf{x} by applying a degradation kernel (or PSF) \mathbf{k} [16], followed by a downsampling operation \downarrow_s with scale factor s (e.g. Bicubic [46, 57]), and the addition of the noise \mathbf{n} [18]. The more realistic these factors are modelled and applied, the better the restoration and SISR models perform and generalize in real-world scenarios [54, 66].

This introduces the need for more complex degradation pipelines, to reduce the generalization gap observed in the real world domain. Recently, Zhang *et al.* [66] designed a practical degradation pipeline for the RGB domain SISR, achieving promising results in real scene super-resolution.


RAWSR from Xu *et al.* [54, 55] is one of the first works that apply a simple degradation pipeline (Eq. 1) for synthesizing low-resolution RAW images and train SISR models that successfully generalize on real scenes. Their approach is illustrated in Figure 1c, ultimately they aim at enhancing RGB images by leveraging both the sensor RAW data and the RGB produced by the ISP. For this reason, their studies are performed on the resultant RGBs, and their



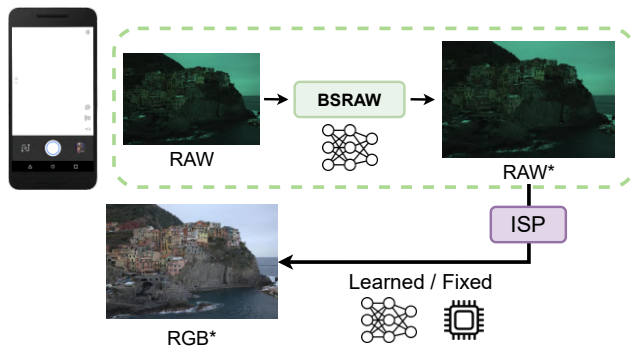
(a) Professional photography image editing [8, 27]. The software (e.g. Adobe Lightroom) usually applies restoration on the RAW image, and further ISP and enhancement steps.



(b) Image Restoration and Super-Resolution in the RGB domain as offline post-processing [13, 31, 51, 61, 66] — most popular.



(c) Xu *et al.* [54, 55] approach leverages RAW and RGB images for real scene super-resolution using a double-branch network.



(d) **Our approach BSRaw**. We focus on **Blind Super-Resolution of RAW** images. Our model operates before the ISP as a plug & play block, not as post-processing (b,c), thus it can benefit other downstream tasks.

Figure 1. Comparison of image restoration and SISR pipelines. Note that previous approaches (b) and (c) depend on the in-camera ISP, meanwhile our BSRaw method can complement any (learned) ISP, as a pre-processing block. We refer as RAW* and RGB* to the enhanced images. Sample image from the Adobe MIT5K dataset [8].

method can be seen as an offline *post-processing* (always *after the ISP*) for image enhancement and photo-finishing. This is similar to most popular image restoration and SISR techniques [13, 31, 61, 66] -see Figure 1b-.

In this work we study **RAW SISR** -see Figure 1d-, using exclusively RAW sensor data, thus, the signal depends only on the scene and the imaging system (e.g. light sources, user-camera interaction). Moreover, the degradations such as noise [1, 6] are not altered by the ISP stages [14, 27]. We build a realistic RAW degradation pipeline that extends and improves previous methods from Zhang *et al.* [66] and Xu *et al.* [54, 55]. Our baseline model, BSRaw, is an efficient blind RAW SISR method, trained using our degradation pipeline, and real data from DSLR and DSLM cameras. The *main contributions* of this paper are:

- 1) A controllable degradation pipeline to train deep neural models using RAW images and considering realistic downsampling, noise and blur.
- 2) A new dataset for RAW SISR including digital single-lens mirrorless (DSLM) cameras.
- 3) We study BSRaw and other methods for this task, and we provide a new benchmark for RAW SISR.

2. Related Work

Image Signal Processing Digital camera systems incorporate a specialized in-camera Image Signal Processor (ISP) [14, 20, 27, 45, 52], which converts the captured RAW sensor data into images that are both less deteriorated and visually appealing to human perception.

In the recent years many deep learning approaches have been proposed to learn the ISP transformation [12, 34]. Some popular methods in this direction are CycleISP [62], CameraNET [32], and PyNET [25].

In this work, we use ISP pipelines [27] with known parameters and operations, however, our proposed method can help to improve learned ISPs for real scene processing.

Degradation Models Creating authentic degradation models for synthesizing data is a critical endeavor in the field of low-level vision. Standard degradation models frequently adopt a sequential approach involving blur, downsampling, and noise injection (as described in Eq. 1). Since most degradation models are far from the real-world complexity, some works aimed at solving such problem by designing complex degradation pipelines [7, 35, 54, 66].

Zhang *et al.* [64, 66] devised a pragmatic degradation methodology specifically tailored for training sRGB image restoration models, achieving encouraging outcomes in Single Image Super-Resolution (SISR). Extensions of this approach to video denoising have also been reported [10].

Xu *et al.* [54, 55] implement the degradation equation (Eq. 1) to synthesize low-resolution RAW images and then

train models that exhibit strong generalization capabilities in real-world scenarios. In their work, they use heteroscedastic Gaussian noise, varying disk kernels for defocus blur, and minimal motion blur.

Little research exists on elaborating more intricate degradation phenomena, apart from noise, in the RAW image space [1, 18, 30]. In response, we introduce improvements such as enhanced noise profiles, anisotropic defocusing, motion blur variation, exposure inconsistencies, and compression artifacts. Details of our enhanced degradation model are elaborated in Section 3.

2.1. Image Super-Resolution and Restoration

In the recent years blind image restoration and super-resolution using deep learning has become a popular task [13, 15, 31, 49, 65, 66, 71], this is because blind methods do not require prior knowledge about the degradation factors or the camera sensor, therefore they are general purpose models *i.e.* practically, such methods can process any degraded RGB image. SwinIR [31] is an example of general purpose SISR and restoration method.

The vast majority of blind restoration and SISR methods operate in the RGB color space. We also focus on the blind super-resolution problem, yet, in the RAW domain.

There are far fewer works that solve complex inverse problems in the RAW domain. Most methods focus on RAW image denoising, since estimating and removing the noise on linear raw data is well-studied [1, 6, 12, 38]. SIDD [1] and DND [40] are the most popular RAW denoising benchmarks. These benchmarks show that *state-of-the-art* methods can achieve high reconstruction ($> 50\text{dB}$ PSNR) given a noisy RAW image, indicating that new challenging data should be used. Liang *et al.* studied RAW image deblurring on a simple setup using images from one DSLR camera [30]. We also find works focused on exposure correction for low-light image enhancement [11, 12, 19, 23]. Xu *et al.* introduced RAWSR [54, 55] for super-resolution and enhancement using both RAW and RGB images as inputs.

We aim to study the RAW SISR problem assuming (i) unknown real-world degradations, and (ii) cross-device generalization *i.e.* our method can be applied to images from different sensors, as the most popular RGB restoration models [13, 31, 61].

3. Controllable RAW Degradation Pipeline

We improve previous degradation models to synthesize realistic degraded RAW images [54] for training SR models. We consider the most common limiting factors related to the acquisition and the processing of high quality RAW images: (i) noise and low resolution – related to the size of the sensor [27]. (ii) exposure. (iii) motion and defo-

cus blur – related to stabilization issues. Note that previous methods [54, 55] did not consider multiple noise profiles, or complex blur and PSFs. We provide samples from our dataset in Figure 2.

3.1. Noise

Noise (**n**) is an omnipresent element across all stages of image capture and processing [1, 18]. Its occurrence is influenced by a multitude of variables, such as exposure duration and lighting conditions. In the domain of photography, noise removal—or denoising—commonly takes place in the RAW stage, where its linear characteristics have been thoroughly investigated [1, 6, 18, 19, 37]. This makes it more manageable to address prior to the non-linear transformations introduced by the ISP [27]. Most methods use Homoscedastic Gaussian noise [65–67] in their degradation models. Xu *et al.* [54] use sharper Heteroscedastic Gaussian distribution [1, 40].

We adopt a more practical shot-read noise [6, 18, 70]. In Equation 2, we can observe the intensity y as a sample of a Gaussian distribution having as mean the input signal x and the variance depending on the λ_s (shot) and λ_r (read) parameters [6]. This is derived from a Poisson-Gaussian noise model [50]:

$$y \sim \mathcal{N}(\mu = x, \sigma^2 = \lambda_r + \lambda_s x) \quad (2)$$

Utilizing Zhang *et al.* methodology [70], we derive noise profiles specific to our imaging devices, in addition to adopting established profiles from SIDD [1, 2] and DND [40] for smartphone and DSLR sensors, respectively. This diverges from earlier works that relied on a singular noise distribution [6, 54]. Our approach incorporates **varied real-world noise profiles**, which manifest stochastically. We also emulate the interaction effects (*e.g.*, attenuation, amplification) between introduced noise and other image degradations (*e.g.*, blur).

3.2. Blur

Blur (\mathcal{B}_k) is also one of the most common degradations appearing in image capturing *e.g.*, camera shake in hand-held photography, motion blur, and defocus blur [3, 22, 58]. Capturing aligned blurry-clean pairs in real scenarios is extremely difficult, for this reason, the most popular deblurring datasets are synthetic. Most approaches adopt a uniform blur by convolving the image with iso/anisotropic Gaussian kernels [66]. Xu *et al.* [54, 55] implemented defocus blur as a disk kernel, and a modest motion blur [44].

We create our blur degradation by convolving the image with a **diverse pool of kernels**: classical isotropic and anisotropic gaussian blur kernels [29, 66], real estimated motion blur kernels [39, 43, 56], and real estimated PSFs (point-spread-function) from [53, 63, 72].

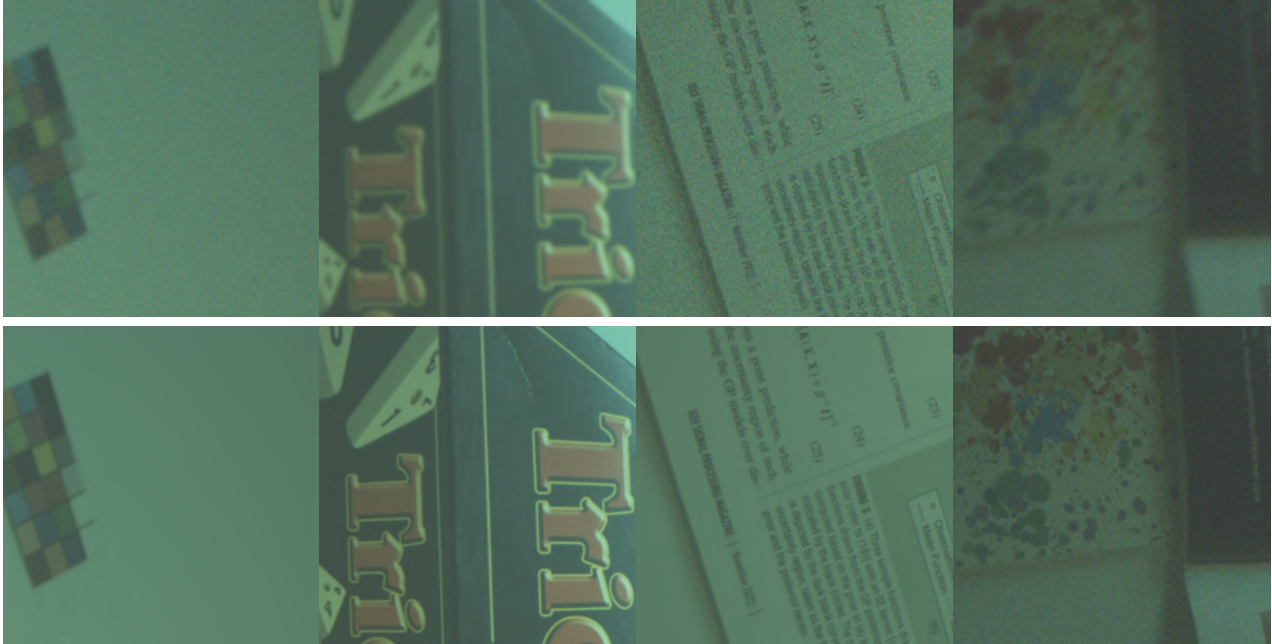


Figure 2. Different variants of our dataset **test images**, after applying our degradation model. (Up) Degraded images, (Bot.) Clean images. The RAW images were acquired from multiple sensors (*e.g.* Nikon D700 and Nikon Z7). RAW images are visualized after demosaicing, gamma correction and tone mapping. Best viewed in electronic version.

This task is fundamental since the current trend in reducing size and weight observed for the most popular DSLR or mirrorless cameras, and for the smartphone cameras, makes it challenging to implement sensor stabilization systems handling extreme levels of sensor instability during point-and-shoot acquisition.

3.3. Exposure

Exposure levels (EE) can introduce undesirable artifacts in photographic images. Both underexposure and overexposure can result in signal loss or increased noise in the imagery [11, 12, 19]. While our inclusion of this degradation isn't intended to address HDR issues [19], it does allow the model to acquire exposure correction skills. Additionally, the interplay between simulated exposure levels and noise augments the complexity of the degradation environment. We adopt a low exposure model inspired by the work of Punnappurath *et al.* [41], employing a linear ISO scaling technique to generate the resultant darkened image. It is worth noting that our real-world testing dataset comprises images with exposure-related challenges, along with their respective, carefully-selected ground truth.

3.4. Downsampling

One of the most important differences between cameras is the size of the sensor [14, 25, 27]. Photographers can use telephoto lenses to obtain HR images, but the resolution of

the scene is limited by the size of the sensor. Thus, it is more common to capture a wide-range scene as a LR image using a wide-angle lens, and then apply SISR [54, 55]. Considering this, it is desirable to be able to upscale RAW images [54, 55]. However, this is not a trivial task -as for RGB images- since most traditional methods do not preserve the Bayer pattern (RGGB).

Since obtaining real LR-HR pairs is extremely expensive and time-consuming, we follow [54, 66] to synthesize LR RAW images from the assumed HR real captures (*e.g.* 60MP RAWs). We downsample the high-quality RAW images so that each pixel could have its ground truth red, green and blue values. We denote this operation as \downarrow_s .

3.5. Complete RAW Degradation Model

The degradation pipeline covers most of the factors affecting the quality of RAW images. One important aspect is the noise sampling from multiple noise profiles, as well as the PSF sampling (and even application of two PSFs). The order of the degradations is defined by the optics theory [17] *i.e.* first we apply the kernel (\mathbf{k}) to the RAW signal, then we simulate a small sensor with downsampling and noise injection -that can be influenced by the exposure- and later we apply the other degradations including additional motion blur and downsampling.

We define four different degradation levels in our experiments: I) classical image restoration [17, 31]. Only blur



Figure 3. Real world samples representing usual RAW image defects. Best in the electronic format. Images captured manually with our Nikon Z7 on a tripod. Illumination and stabilization are controlled to ensure proper alignment.

Device	Sensor	# Train	# Test	Res.
DSLRs [8, 54]	Diverse	1000	150	10MP
Sony α 7R4	IMX551	90	10	2MP
Nikon Z7	IMX309BQJ	90	10	2MP

Table 1. Dataset split. We use diverse DSLR images from the MIT5K as previous works [8, 54] (Canon EOS 5D and Nikon D700). The sensors highlighted in gray correspond to our new dataset of real-world captures. The train/test split ensures no overlap and proper evaluation.

and noise are considered. II) SISR degradation model (see Eq. 1) III) SISR considering additional exposure degradations IV) Finally we apply our complete realistic degradation pipeline, including multiple noise profiles, the application of multiple kernels, exposure, and downsampling. We provide more details about the pipeline in the appendix.

4. Experiments

We conduct several experiments to prove the benefits of our approach for RAW image restoration and SISR. We propose BSRRAW as an efficient baseline method for this task. Due to the pages limit, we provide additional implementation details, dataset information, and results in the supplementary material.

4.1. BSRRAW Dataset

Our goal is to solve blind super-resolution for RAW images, and apart from a realistic degradation prior (Sec. 3), an image prior also contributes to the success of the deep model [66]. Our sources of real RAW data and our train/test split are summarized in Table 1. The quality of the DSLR RAW images is substantially better than in smartphones due to the larger optics and sensor, this implies that the RAW images show less noise and minimal optical aberrations.

Following previous work [52, 54, 55], we use images from the Adobe MIT5K dataset [8], which includes two DSLR cameras. The DSLR images are manually filtered

to ensure diversity and natural properties (*i.e.* remove extremely dark or overexposed images), we also remove the blurry images (*i.e.* we only consider all-in-focus images).

The **pre-processing** is as follows: (i) we normalize all RAW images depending on their black level and bit-depth. (ii) we convert (“pack”) the images into the well-known RGGB Bayer pattern (4-channels), which allows to apply the transformations and degradations without damaging the original color pattern information [33]. (iii) Following previous work on RAW processing and learned ISPs [24, 54], we train using image patches of dimension $248 \times 248 \times 4$.

Synthetic Dataset We apply our complete degradation pipeline (see Sec. 3.5) to the training images to generate aligned degraded-clean pairs. The synthetic test dataset is generated by applying our degradation pipeline, at different levels, to the corresponding test images – see Table 1.

Real-World Captures We believe the models trained using our pipeline have notable generalization capabilities (also concluded in [54, 66]), therefore, we collect a novel dataset with challenging real scenes. As indicated in Table 1, we capture images using two mirrorless full-frame cameras: Sony α 7R4 (60MP) and Nikon Z7 (45.7MP). In Figure 3 we show samples from our data capturing setup.

This is a paired dataset consisting of variations over 100 real-world indoor and outdoor scenes. All the devices used for data capturing allow for manual operation of the camera ISP parameters and optics, and are fixed on a tripod. For each “clean” reference image, we capture several variations, where each variation represents a degradation with a different degree of intensity. The scene variants represented there are degraded in terms of defocusing blur, exposure time, high ISO noise, and motion blur. Since interaction with the camera device was needed during the acquisition, we perform a post-processing alignment (similar to [5]) to further improve the quality and usability of the data.

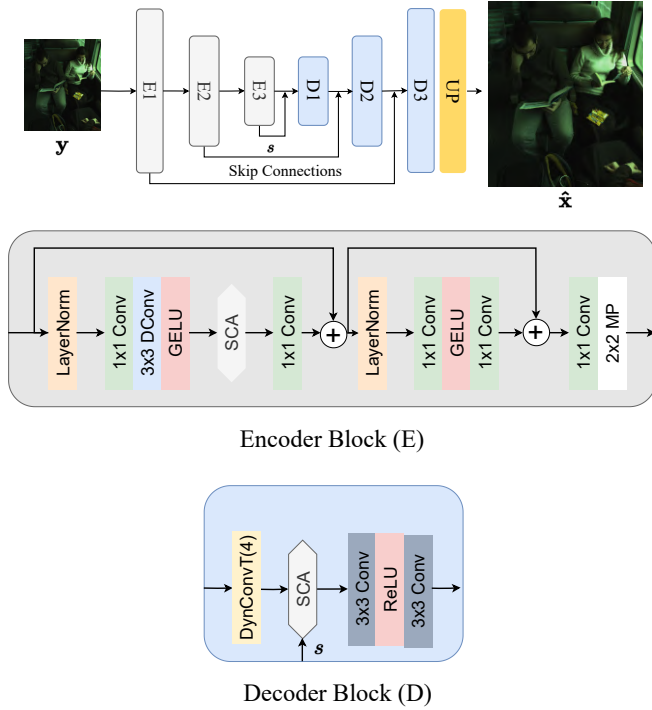


Figure 4. Architecture of the proposed BSRaw inspired in NAFNet [13]. We simplify the encoder blocks and use Dynamic convolutions to process images more efficiently.

4.2. Experimental Results

We propose BSRaw as an efficient baseline, able to process 4K resolution images in standard GPUs (120 ms for a 4K image on a Nvidia 3090Ti 24GB GPU), without patching or tiling strategies. The architecture is illustrated in Figure 4. BSRaw combines NAFNet [13] baseline blocks with dynamic convolutions and simplified stereo channel attention in the decoder part of the model. In comparison to NAFNet [13] -already quite efficient- our method has 10× less parameters and $\approx 20\%$ less MACs. We upscale the images using an adapted pixel-shuffle head.

SISR results We compare with previous method RAWSR [54] for SISR. Our realistic degradation pipeline helps our efficient model BSRaw to generalize better on complex degradations. We upscale the images using BSRaw, and later process them using a basic ISP [27] *e.g.* Ddraw as [54] to obtain the resultant RGB. Note that the training dataset derived from MIT5K [8] from Xu *et al.* RAWSR [54] is not publicly available. Therefore we train using our own curated DSLR dataset (see Table 1); this, together with BSRaw being trained using our degradation pipeline, allow us to achieve notably better results. We also compare with TENet [42] (based on RDN [71]) pipeline for “Denoising (DN) \rightarrow SR” for RAWs.

Methods	PSNR \uparrow	SSIM \uparrow	Domain
DeepISP	21.71	0.7323	RAW
RDN [71]	29.93	0.7804	RGB
TENet [42]	30.20	0.800	RAW
BSRGAN [66]	30.01	0.794	RGB
RAWSR [54]	<u>30.79</u>	<u>0.804</u>	RAW+RGB
BSRAW II (<i>Ours</i>)	31.02	0.812	RAW
BSRAW IV (<i>Ours</i>)	31.40	0.843	RAW
BSRAW \dagger (<i>Ours</i>)	31.86	0.860	RAW

Table 2. Comparison for RAW SISR $\times 2$. We use the DSLR test-set from [54], which consists on 150 images from MIT5K [8] degraded with a simple SISR model [54]. Our model BSRaw achieves *state-of-the-art* results thanks to our degradation pipeline and our dataset. Metrics calculated on RGBs. We process the up-scaled RAWs using a fixed ISP and metadata [54] *e.g.* Ddraw.

We present the results in Table 2. We distinguish three variants of our model, the **ablation study** is as follows:

(i) BSRaw II is trained using a simple SISR degradation pipeline, similar to [54] -our level II-. (ii) BSRaw IV is trained using our complete degradation pipeline. We can appreciate how this helps to improve results notably. Finally, (iii) BSRaw \dagger extends the previous variant, and we include training data from our dataset. This leads to the best results in terms of fidelity metrics.

In Table 3 we show the results from our model on the synthetic dataset at each degradation level.

Effectiveness on Real Images As we show in Figures 6 and Table 2 for RAW restoration and SISR, our model can generalize and provide reasonable results on real scenes. Figures 5 and 6 qualitative results show that BSRaw is able to produce noise-free sharp images with proper textures and color distributions, even using RAW images from unseen sensors such as Sony Quad sensors or SIDD phones [1].

Super-Resolution, RAW or RGB? We also compare the following two strategies for image processing:

(I) SwinIR [31] and BSRGAN [66], general-purpose models trained for real-world blind SISR on RGBs. Their input is a degraded RGB image obtained after applying an ISP to the original degraded RAW image. We use images from the SIDD dataset [1] with noise and low exposure.

(II) Our approach. BSRaw takes directly the degraded RAW image, applies restoration and upsampling, then we apply the ISP with known parameters and operations.

The ISP [27] is fixed for a fair comparison. The images are real-world captures from “unknown” sensors to the model.

In Figure 6 we compare these two possible variants of super-resolution. The quantitative results from Table 2, and qualitative results (Figure 6) show the advantages and potential of our approach. We provide further analysis and samples in the supplementary material.

4.3. Implementation Details

The method was implemented in PyTorch. We use the basic \mathcal{L}_1 loss between the ground-truth and the restored image. We set the minibatch size to 32. We use Adam [28] optimizer with standard hyper-parameters. The initial learning rate is set to $1e^{-4}$, we use a linear LR decay scheduler, until a minimum of $1e^{-5}$. We perform experiments using multiple NVIDIA RTX 3090Ti and 4090Ti.

The degradation pipeline can be integrated into the dataloader as an augmentation. All the degradations are applied on 4-channel RAW images (RGGB).

Note that the training dataset derived from MIT5K [8] from Xu *et al.* RAWSR [54] is not publicly available, and therefore cannot be used for comparison. We cannot perform a qualitative comparison because RAWSR weights and results are not publicly available *i.e.* the material at ¹ is obsoleted and links do not work.

In the *RAW domain*, we can compare the HR RAW ($1460 \times 2192 \times 4$) and the LR RAW ($365 \times 548 \times 4$) — which correspond to RGBs of resolution ($2920 \times 4384 \times 3$) and ($730 \times 1096 \times 3$) respectively. Upscaling the LR RAW using Bicubic interpolation implies a reconstruction of 27.13dB (and 0.81 SSIM) *w.r.t* the HR RAW.

4.4. Limitations

Downsampling RAW images is not trivial, this operation can cause color artifacts besides blurriness, and irreversible structure loss. Also, the MIT5K [8] images were captured with “old” cameras, for instance, 12MP DSLRs, in comparison with our 60MP DSLM cameras. We need further analysis using modern cameras, specially smartphones.

5. Conclusion

In this paper we discuss a controllable degradation pipeline to synthesize realistic degraded RAW images for training deep blind super-resolution models. We have curated a dataset with different DSLR cameras. As part of this effort, we provide a new dataset of diverse real scenes captured using DSLM cameras. Our experiments demonstrate that models trained with our degradation pipeline can restore real-world degraded RAWs. This represents a powerful alternative in image signal processing, and can be of benefit to other low-level downstream tasks.

¹<https://github.com/xuxy09/RAWSR>




Figure 5. Our BSRAW model, thanks to the proposed degradation pipeline, can perform $2\times$ (red) and $4\times$ (yellow) SISR and restore RAW images in the wild. We use images from a smartphone Sony IMX586 Quad Bayer sensor [24], which was not used for training (neither similar sensors). This proves certain generalization capabilities, even on unseen sensors.

Future Work We plan to extend this analysis to smartphone images, where the ISP is more challenging, and the possible degradations in the RAW images might be more complex due to the limited optics *i.e.* more noise and blur.

Acknowledgments This work was partly supported by the The Alexander von Humboldt Foundation (AvH).

Degradation	Super-Resolution $\times 2$				Super-Resolution $\times 4$			
	BSRAW-light		BSRAW		BSRAW-light		BSRAW	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
I	42.57	0.949	43.30	0.952	40.08	0.929	40.51	0.937
II	43.26	0.961	43.98	0.965	40.42	0.941	41.28	0.949
III	43.53	0.962	43.96	0.966	40.51	0.939	41.34	0.951
IV	43.78	0.964	44.25	0.971	39.71	0.927	41.45	0.957

Table 3. Ablation study for BSRAW on different degradation levels using our *synthetic MIT5K testset*. The metrics are calculated in the RGB domain, thus, the ISP setting was fixed for all the compared configurations, to avoid possible biases, we use a simple ISP with the canonical steps (demosaicing, white balance, color correction, gamma and tone mapping correction.) [27]. BSRAW-light is a variant with less encoder blocks and channels, described in the supplementary material.

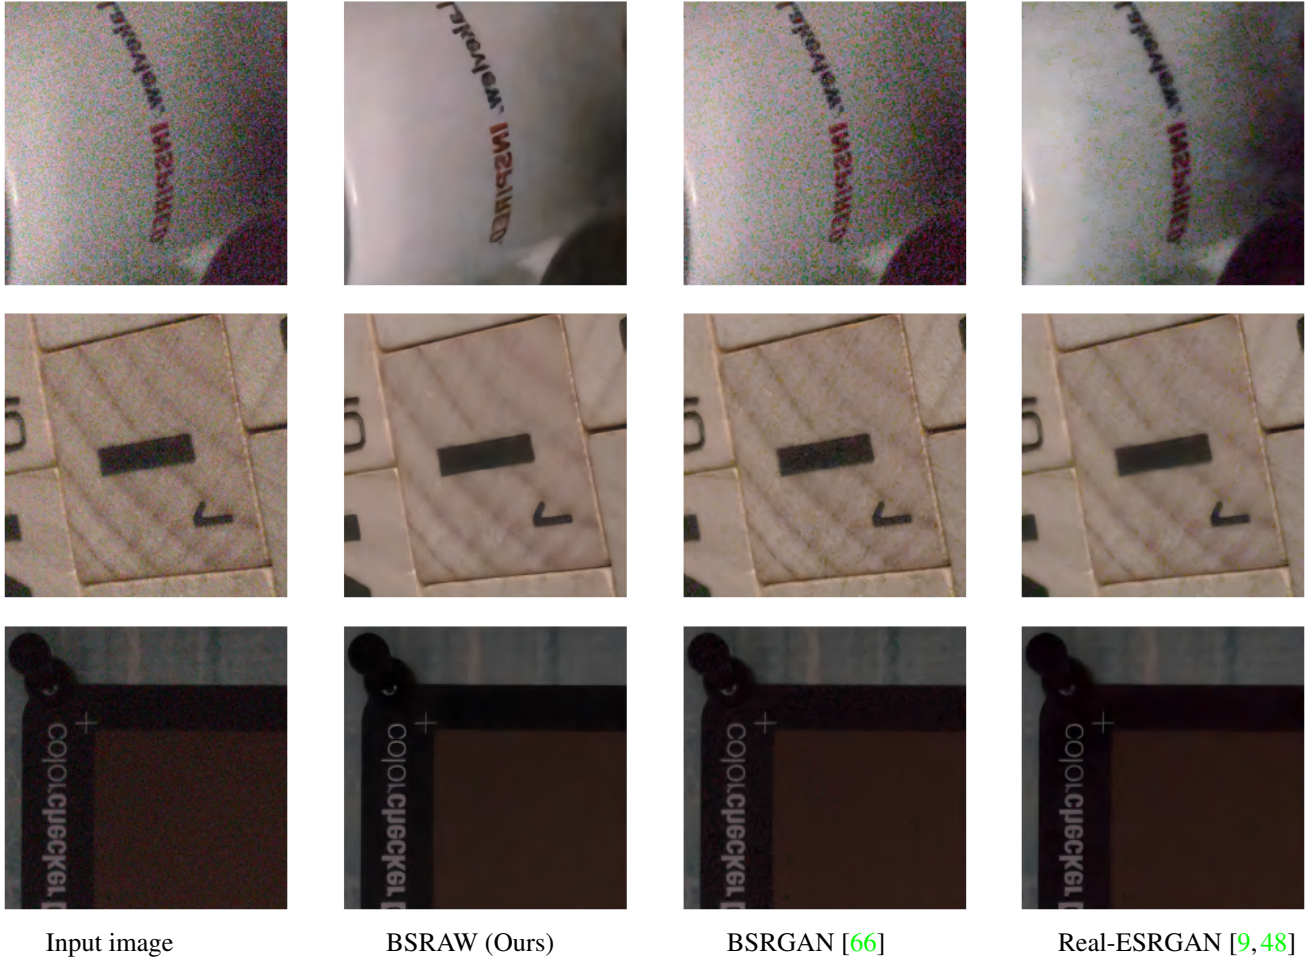


Figure 6. From left to right: sRGB rendered using the input RAW and a simple ISP [1], our restored RAW (we apply the ISP after restoring and upsampling), BSRGAN [66] and Real-ESRGAN [9,48] restored images. The approaches are illustrated in Figure 1. We believe both blind single image super-resolution (SISR) approaches might be complementary and represent future research.

References

- [1] Abdelrahman Abdelhamed, Stephen Lin, and Michael S Brown. A high-quality denoising dataset for smartphone cameras. In *CVPR*, pages 1692–1700, 2018. 1, 2, 3, 6, 8
- [2] Abdelrahman Abdelhamed, Radu Timofte, Michael S. Brown, et al. Ntire 2019 challenge on real image denoising: Methods and results. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, June 2019. 3
- [3] Abdullah Abuolaim, Mahmoud Afifi, and Michael S Brown. Improving single-image defocus deblurring: How dual-pixel images help through multi-task learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1231–1239, 2022. 3
- [4] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 126–135, 2017. 1
- [5] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Deep burst super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9209–9218, 2021. 5
- [6] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T Barron. Unprocessing images for learned raw denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11036–11045, 2019. 1, 2, 3
- [7] Adrian Bulat, Jing Yang, and Georgios Tzimiropoulos. To learn image super-resolution, use a gan to learn how to do image degradation first. In *Proceedings of the European conference on computer vision (ECCV)*, pages 185–200, 2018. 2
- [8] Vladimir Bychkovsky, Sylvain Paris, Eric Chan, and Frédo Durand. Learning photographic global tonal adjustment with a database of input / output image pairs. In *CVPR*, 2011. 2, 5, 6, 7
- [9] Jianrui Cai, Hui Zeng, Hongwei Yong, Zisheng Cao, and Lei Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *IEEE Conference on International Conference on Computer Vision*, pages 3086–3095, 2019. 8
- [10] Jiezhong Cao, Qin Wang, Jingyun Liang, Yulun Zhang, Kai Zhang, and Luc Van Gool. Practical real video denoising with realistic degradation model. *arXiv preprint arXiv:2208.11803*, 2022. 2
- [11] Meng Chang, Huajun Feng, Zhihai Xu, and Qi Li. Low-light image restoration with short-and long-exposure raw pairs. *IEEE Transactions on Multimedia*, 24:702–714, 2021. 1, 3, 4
- [12] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3291–3300, 2018. 2, 3, 4
- [13] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. *arXiv preprint arXiv:2204.04676*, 2022. 1, 2, 3, 6
- [14] Mauricio Delbracio, Damien Kelly, Michael S Brown, and Peyman Milanfar. Mobile computational photography: A tour. *arXiv preprint arXiv:2102.09000*, 2021. 2, 4
- [15] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *European Conference on Computer Vision*, pages 184–199, 2014. 1, 3
- [16] Netalee Efrat, Daniel Glasner, Alexander Apartsin, Boaz Nadler, and Anat Levin. Accurate blur models vs. image priors in single image super-resolution. In *IEEE Conference on International Conference on Computer Vision*, pages 2832–2839, 2013. 1
- [17] Michael Elad and Arie Feuer. Restoration of a single super-resolution image from several blurred, noisy, and undersampled measured images. *IEEE Transactions on Image Processing*, 6(12):1646–1658, 1997. 1, 4
- [18] Samuel W Hasinoff. Photon, poisson noise. *Computer Vision, A Reference Guide*, 4, 2014. 1, 3
- [19] Samuel W. Hasinoff, Dillon Sharlet, Ryan Geiss, Andrew Adams, Jonathan T. Barron, Florian Kainz, Jiawen Chen, and Marc Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 35(6), 2016. 1, 3, 4
- [20] Felix Heide, Markus Steinberger, Yun-Ta Tsai, Mushfiquir Rouf, Dawid Pajak, Dikpal Reddy, Orazio Gallo, Jing Liu, Wolfgang Heidrich, Karen Egiazarian, et al. Flexisp: A flexible camera image processing framework. *ACM Transactions on Graphics (ToG)*, 33(6):1–13, 2014. 1, 2
- [21] Daniel Hernandez-Juarez, Sarah Parisot, Benjamin Busam, Ales Leonardis, Gregory Slabaugh, and Steven McDonagh. A multi-hypothesis approach to color constancy. In *VPR*, 2020. 1
- [22] Mahdi S Hosseini and Konstantinos N Plataniotis. Convolutional deblurring for natural imaging. *IEEE Transactions on Image Processing*, 29:250–264, 2019. 3
- [23] Haofeng Huang, Wenhan Yang, Yueyu Hu, Jiaying Liu, and Ling-Yu Duan. Towards low light enhancement with raw images. *IEEE Transactions on Image Processing*, 31:1391–1405, 2022. 3
- [24] Andrey Ignatov, Cheng-Ming Chiang, Hsien-Kai Kuo, Anastasia Sycheva, and Radu Timofte. Learned smartphone isp on mobile npus with deep learning, mobile ai 2021 challenge: Report. In *CVPR Workshops*, pages 2503–2514, 2021. 5, 7
- [25] Andrey Ignatov, Luc Van Gool, and Radu Timofte. Replacing mobile camera isp with a single deep learning model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 536–537, 2020. 2, 4
- [26] Xiaozhong Ji, Yun Cao, Ying Tai, Chengjie Wang, Jilin Li, and Feiyue Huang. Real-world super-resolution via kernel estimation and noise injection. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 466–467, 2020. 1
- [27] Hakki Can Karaimer and Michael S Brown. A software platform for manipulating the camera imaging pipeline. In *ECCV*, pages 429–444, 2016. 1, 2, 3, 4, 6, 8

- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [29] Anat Levin, Yair Weiss, Fredo Durand, and William T Freeman. Understanding and evaluating blind deconvolution algorithms. In *2009 IEEE conference on computer vision and pattern recognition*, pages 1964–1971. IEEE, 2009. 3
- [30] Chih-Hung Liang, Yu-An Chen, Yueh-Cheng Liu, and Winston Hsu. Raw image deblurring. *IEEE Transactions on Multimedia*, 2020. 1, 3
- [31] Jingyun Liang, Jiezhong Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1833–1844, 2021. 1, 2, 3, 4, 6
- [32] Zhetong Liang, Jianrui Cai, Zisheng Cao, and Lei Zhang. Cameranet: A two-stage framework for effective camera isp learning. *IEEE Transactions on Image Processing*, 30:2248–2262, 2021. 2
- [33] Jiaming Liu, Chi-Hao Wu, Yuzhi Wang, Qin Xu, Yuqian Zhou, Haibin Huang, Chuan Wang, Shaofan Cai, Yifan Ding, Haoqiang Fan, et al. Learning raw image denoising with bayer pattern unification and bayer preserving augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 5
- [34] Shuai Liu, Chaoyu Feng, Xiaotao Wang, Hao Wang, Ran Zhu, Yongqiang Li, and Lei Lei. Deep-flexisp: A three-stage framework for night photography rendering. In *CVPR*, pages 1211–1220, 2022. 2
- [35] Andreas Lugmayr, Martin Danelljan, and Radu Timofte. Un-supervised learning for real-world super-resolution. In *IEEE Conference on International Conference on Computer Vision Workshops*, pages 3408–3416. IEEE, 2019. 2
- [36] Jiaqi Ma, Shengyuan Yan, Lefei Zhang, Guoli Wang, and Qian Zhang. Elmformer: Efficient raw image restoration with a locally multiplicative transformer. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 5842–5852, 2022. 1
- [37] Ali Maleky, Shayan Kousha, Michael S Brown, and Marcus A Brubaker. Noise2noiseflow: Realistic camera noise modeling without clean images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17632–17641, 2022. 3
- [38] Ben Mildenhall, Jonathan T Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. Burst denoising with kernel prediction networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2502–2510, 2018. 1, 3
- [39] Jinshan Pan, Deqing Sun, Hanspeter Pfister, and Ming-Hsuan Yang. Blind image deblurring using dark channel prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1628–1636, 2016. 3
- [40] Tobias Plotz and Stefan Roth. Benchmarking denoising algorithms with real photographs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1586–1595, 2017. 3
- [41] Abhijith Punnappurath, Abdullah Abuolaim, Abdelrahman Abdelhamed, Alex Levinshtein, and Michael S Brown. Day-to-night image synthesis for training nighttime neural isps. In *CVPR*, pages 10769–10778, 2022. 4
- [42] Guocheng Qian, Yuanhao Wang, Chao Dong, Jimmy S Ren, Wolfgang Heidrich, Bernard Ghanem, and Jinjin Gu. Rethinking the pipeline of demosaicing, denoising and super-resolution. *arXiv preprint arXiv:1905.02538*, 2019. 1, 6
- [43] Dongwei Ren, Kai Zhang, Qilong Wang, Qinghua Hu, and Wangmeng Zuo. Neural blind deconvolution using deep priors. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3341–3350, 2020. 3
- [44] Christian J Schuler, Harold Christopher Burger, Stefan Harmeling, and Bernhard Scholkopf. A machine learning approach for non-blind image deconvolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1067–1074, 2013. 3
- [45] Eli Schwartz, Raja Giryes, and Alex M Bronstein. Deepisp: Toward learning an end-to-end image processing pipeline. *IEEE Transactions on Image Processing*, 28(2):912–923, 2018. 1, 2
- [46] Radu Timofte, Vincent De Smet, and Luc Van Gool. A+: Adjusted anchored neighborhood regression for fast super-resolution. In *ACCV*, 2014. 1
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017. 1
- [48] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. *arXiv preprint arXiv:2107.10833*, 2021. 8
- [49] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *European Conference on Computer Vision Workshops*, pages 701–710, 2018. 3
- [50] Yuzhi Wang, Haibin Huang, Qin Xu, Jiaming Liu, Yiqun Liu, and Jue Wang. Practical deep raw image denoising on mobile devices. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI*, pages 1–16. Springer, 2020. 3
- [51] Zhendong Wang, Xiaodong Cun, Jianmin Bao, and Jianzhuang Liu. Uformer: A general u-shaped transformer for image restoration. *arXiv preprint arXiv:2106.03106*, 2021. 1, 2
- [52] Yazhou Xing, Zian Qian, and Qifeng Chen. Invertible image signal processing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6287–6296, 2021. 1, 2, 5
- [53] Li Xu and Jiaya Jia. Two-phase kernel estimation for robust motion deblurring. In *Computer Vision–ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part I 11*, pages 157–170. Springer, 2010. 3
- [54] Xiangyu Xu, Yongrui Ma, and Wenxiu Sun. Towards real scene super-resolution with raw images. In *CVPR*, 2019. 1, 2, 3, 4, 5, 6, 7

- [55] Xiangyu Xu, Yongrui Ma, Wenxiu Sun, and Ming-Hsuan Yang. Exploiting raw images for real-scene super-resolution. *IEEE transactions on pattern analysis and machine intelligence*, 44(4):1905–1921, 2020. 1, 2, 3, 4, 5
- [56] Xiangyu Xu, Jinshan Pan, Yu-Jin Zhang, and Ming-Hsuan Yang. Motion blur kernel estimation via deep learning. *IEEE Transactions on Image Processing*, 27(1):194–205, 2017. 3
- [57] Jianchao Yang, John Wright, Thomas S Huang, and Yi Ma. Image super-resolution via sparse representation. *IEEE transactions on image processing*, 19(11):2861–2873, 2010. 1
- [58] Lu Yuan, Jian Sun, Long Quan, and Heung-Yeung Shum. Image deblurring with blurred/noisy image pairs. In *ACM SIGGRAPH 2007 papers*, pages 1–es. 2007. 3
- [59] Huanjing Yue, Cong Cao, Lei Liao, Ronghe Chu, and Jingyu Yang. Supervised raw video denoising with a benchmark dataset on dynamic scenes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2301–2310, 2020. 1
- [60] Huanjing Yue, Zhiming Zhang, and Jingyu Yang. Real-rawvsr: Real-world raw video super-resolution with a benchmark dataset. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VI*, pages 608–624. Springer, 2022. 1
- [61] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5728–5739, 2022. 1, 2, 3
- [62] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Cycleisp: Real image restoration via improved data synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2696–2705, 2020. 1, 2
- [63] Kai Zhang, Luc Van Gool, and Radu Timofte. Deep unfolding network for image super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3217–3226, 2020. 1, 3
- [64] Kai Zhang, Yawei Li, Jingyun Liang, Jiezhang Cao, Yulun Zhang, Hao Tang, Radu Timofte, and Luc Van Gool. Practical blind denoising via swin-conv-unet and data synthesis. *arXiv preprint arXiv:2203.13278*, 2022. 2
- [65] Kai Zhang, Yawei Li, Wangmeng Zuo, Lei Zhang, Luc Van Gool, and Radu Timofte. Plug-and-play image restoration with deep denoiser prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 3
- [66] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *IEEE Conference on International Conference on Computer Vision*, 2021. 1, 2, 3, 4, 5, 6, 8
- [67] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017. 3
- [68] Xuaner Zhang, Qifeng Chen, Ren Ng, and Vladlen Koltun. Zoom to learn, learn to zoom. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3762–3770, 2019. 1
- [69] Yichi Zhang, Hengyu Liu, Dandan Ding, and Zhan Ma. Low light raw image enhancement using paired fast fourier convolution and transformer. In *2022 IEEE International Conference on Visual Communications and Image Processing (VCIP)*, pages 1–5. IEEE, 2022. 1
- [70] Yi Zhang, Hongwei Qin, Xiaogang Wang, and Hongsheng Li. Rethinking noise synthesis and modeling in raw denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4593–4601, 2021. 3
- [71] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2472–2481, 2018. 3, 6
- [72] Xiang Zhu, Filip Šroubek, and Peyman Milanfar. Deconvolving psfs for a better motion deblurring using multiple images. In *Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part V 12*, pages 636–647. Springer, 2012. 3