

DiffuseRAW: End-to-End Generative RAW Image Processing for Low-Light Images

Rishit Dagli
University of Toronto
rishit@cs.toronto.edu




Figure 1. Given a RAW image, DiffuseRAW can perform the entire image-processing pipeline for low-light enhancement. We show the camera output visualized in the sRGB space, the scaled-up RAW images, the output of a traditional pipeline, and compare them with our approach trained on the RAW data. These images are taken with Sony α 7SII and for these input images, we have that the illuminance at the camera is < 0.1 lux, and we show the ratio of exposure times between reference ground-truth and input images ($\times 250, \times 300$). (**Best viewed in color and with zoom.**)

Abstract

Imaging under extremely low-light conditions presents a significant challenge and is an ill-posed problem due to the low signal-to-noise ratio (SNR) caused by minimal photon capture. Previously, diffusion models have been used for multiple kinds of generative tasks and image-to-image tasks, however, these models work as a post-processing step. These diffusion models are trained on processed images and

learn on processed images. However, such approaches are often not well-suited for extremely low-light tasks. Unlike the task of low-light image enhancement or image-to-image enhancement, we tackle the task of learning the entire image-processing pipeline, from the RAW image to a processed image. For this task, a traditional image processing pipeline often consists of multiple specialized parts that are overly reliant on the downstream tasks. Unlike these, we develop a new generative ISP that relies on fine-tuning latent diffu-

sion models on RAW images and generating processed long-exposure images which allows for the apt use of the priors from large text-to-image generation models. We evaluate our approach on popular end-to-end low-light datasets for which we see promising results and set a new SoTA on the See-in-Dark (SID) dataset. Furthermore, with this work, we hope to pave the way for more generative and diffusion-based image processing and other problems on RAW data.

1. Introduction

All imaging systems need to capture light to form an image. When very little light is present in the scene, imaging systems can only capture a few photons from the scene; when few photons are collected from the scene, the captured image is corrupted by noise sources. The main challenge in these low-light scenes is that the signal measured by the sensor is low relative to the noise inherent in the measurement process.

While low-light imaging is a longstanding challenge, traditional methods still struggle to restore images corrupted by noise. For instance, amplifying the measured analog or digital signals can result in the noise being amplified, and significantly degrades the quality of the image. Applying post-processing techniques such as scaling or histogram stretching may improve the image quality to some extent, but they cannot resolve the issue of low signal-to-noise ratio (SNR) due to the low photon counts in the captured image. Physical imaging techniques to increase SNR, such as opening the aperture or using flash, also have their own drawbacks, including motion blur and camera shake that make downstream applications with such images very challenging, and in many cases not feasible. We can also use burst-imaging techniques [24], however, these often fail in very low-light scenarios since these methods are reliant on robust alignment techniques to account for motion in the scene which is immensely challenging in the extreme noise setting in which we are interested in.

Alternatively, denoising algorithms can be applied to noisy images to improve image quality. These range from conventional techniques such as spatial filters like [30, 65, 76] to other modern approaches like BM3D [13], Deep Learning approaches based on CNNs [8, 22, 43, 46, 79, 80], and generative models based on: GANs [31, 86], normalizing flows [68], and very recently on denoising diffusion [48, 50, 52]. Although these methods have been very successful in enhancing low-light images, these approaches do not target very low-light scenarios that we tackle or some of these approaches only tackle image enhancement as opposed

to the entire image processing pipeline like this work. Though there have been approaches that tackle very low-light scenarios (< 0.1 lux) like [83], however, they use a simple CNN to solve this problem and we show how we propose a new generative image processing pipeline. Applying our approach to captured RAW data outperforms other techniques that operate on pre-processed Low Dynamic Range (LDR) camera inputs, and is able to correctly generate minute details of the image as well demonstrated by Figure 1.

Our approach specifically solves this problem by proposing a new image-processing pipeline based on a diffusion model. We first use a standard pre-trained latent diffusion model and fine-tune this model on patches of up-sampled and corrected Bayer RAW images. The fine-tuned model then learns the entire image processing pipeline for low-light RAW data including but not limited to demosaicing, color correction, enhancement, and transformations. The fine-tuning process itself is supervised by long-exposure ground-truth images and the model is fine-tuned to produce the same resolution enhanced images. This greatly improves the overall image-processing pipeline for low-light RAW data and reduces the use of handcrafted processes in the image-processing pipeline, over other traditional and learned image-processing pipelines. Furthermore, this approach also easily allows combining our learned model with other forms of popular text conditioning like Dreambooth [61] and Instruct fine-tuning [6] among others and can potentially be extended to support other high-level image editing operations like stylization, inpainting, accessorization, and property modification. This allows our approach also enables a new type of end-to-end generative camera Image Signal Processor (ISP) wherein regions of an image can be inpainted or stylized using models that are applied end-to-end, starting from the RAW data.

Our approach not only demonstrates improvements throughout the image processing pipeline for low-light RAW data but also is the first approach using a diffusion-based model for working with RAW images and makes apt use of the vast information a RAW image has for the downstream task, here enhancing and denoising very low-light images. With this paper, we hope to pave the way for more diffusion-based approaches for a generative end-to-end camera processing pipeline leading to significant performance improvements.

Contributions. A traditional image processing pipeline often consists of multiple specialized parts that are overly reliant on the downstream tasks. The key novelty of our approach stems from modifying the

pipeline with a generalizable data-driven approach.

- We propose a first-of-its-kind generative camera ISP for extreme low-light image enhancement based on a diffusion model that improves performance on low-light image enhancement in most cases and is the first diffusion-based model that is modeled over RAW image.
- We develop an implementation of this technique using a latent diffusion model fine-tuned on low-light RAW images allowing for great extensibility and making use of the vast amount of priors that diffusion models have.
- We set a new state-of-the-art for image denoising on the popular See-in-the-Dark (SID) dataset [8] demonstrating that a fine-tuned diffusion model can enhance extremely low-light RAW images and perform the whole image processing pipeline on such data.

2. Related Work

Improving low-light images is a problem that has been very well explored throughout the years, and many deep-learning-based methods have recently performed very well in improving low-light images which has led to it being practical now. Images taken in low light have a low signal-to-noise ratio in the captured sensor data and thus are very noisy, to improve such low-light images, these images need to be denoised. We provide a brief overview of image-denoising algorithms. Our approach also makes use of latent diffusion models and we also provide a brief overview of diffusion models. A popular class of algorithms for improving low-light images are ones that mainly handle processed images or image enhancement algorithms as well as algorithms that handle the end-to-end image processing pipelines for this problem. We also provide a brief overview of both of these classes of algorithms.

2.1. Image Denoising

Image denoising is a very well-explored topic in vision and there have been many approaches proposed for denoising images. Many traditional image-denoising approaches draw inspiration from various mathematical and machine learning techniques, utilizing image priors, statistical insights, and the power of deep learning architectures. Within classical image denoising techniques, methods such as Total Variation (TV) denoising [60], wavelet-domain processing [55], sparse coding [16], nuclear norm minimization [20], and 3D Transform-Domain Filtering (BM3D) [14] perform well across a range of noise levels and exploit specific priors

such as sparsity intensity changes, and transformation domain.

The advent of deep learning approaches has significantly reshaped this landscape. There have been multiple convolution-based denoising approaches [42, 77]. These approaches learn convolutional networks to approximate the noise maps and are often based on strong priors. Furthermore, there have also been other specialized approaches for denoising like Stacked Sparse Denoising Auto-Encoders (SSDA) [70], Trainable Nonlinear Reaction-Diffusion (TNRD) [9], and Deep Autoencoders [44] demonstrating performance increases over traditional methods in capturing intricate noise patterns and restoring images. [48] proposed using physics-inspired GANs with stochastic corruptions during training, to enhance the appearance of night sky images.

2.2. Diffusion Models

Diffusion Models [25] are a type of probabilistic generative model that has gained popularity in recent years due to its ability to generate high-resolution images with diverse features. The underlying principle behind diffusion models involves a forward process that gradually adds noise to clean samples drawn from a prior distribution and a reverse process that reverses the corruption process to recover plausible samples from the noise [45]. Diffusion models have demonstrated their effectiveness in various image-based tasks, such as unconditional image generation, inpainting, colorization, image segmentation, and medical imaging. Compared to other generative models such as GANs and VAEs, diffusion models offer a stable training process and the ability to learn strong priors [25]. A Diffusion Model generates samples by gradually removing noise from a signal and learning a network to do so, while their training objective is a reweighted variational lower bound, have shown particularly great results for denoising images as well. In this work, we work on fine-tuning the broad class of latent diffusion models [59].

Another recent approach to low-light image enhancement using deep learning is based on probabilistic models, such as conditional denoising diffusion probabilistic models (DDPMs). Diffusion in the Dark (DiD) [50] is a recent diffusion model for low-light image reconstruction that provides qualitatively competitive reconstructions, especially for low-light text recognition. In contrast, our approach relies on fine-tuning a latent Diffusion model directly on RAW sensor data and aims at learning the entire image processing pipeline.

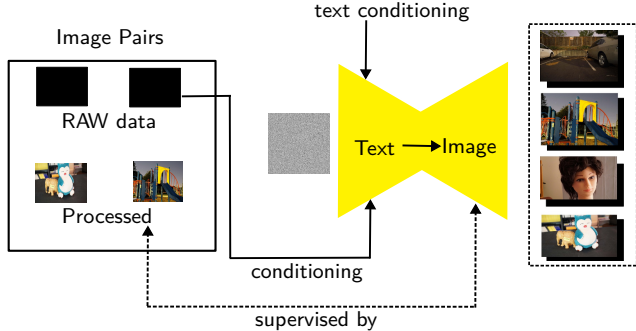


Figure 2. We show how our model is fine-tuned on text-to-image diffusion models which requires paired RAW sensor data and processed long-exposure image.

2.3. Low-light Photography.

Classical low-light reconstruction methods have been widely used to enhance the quality of low-light images, and they remain an active area of research. These methods can be broadly classified into histogram equalization-based methods [2, 10, 27, 33], Retinex-based methods [4, 19, 23, 32, 51, 58], and burst averaging techniques [3, 15, 73].

Histogram equalization, as described in [10], aims to enhance image contrast by redistributing pixel intensities. However, it often suffers from over-enhancing noise in darker regions. In contrast, Retinex-based methods separate images into reflectance and illumination components, inspired by Land’s theory of color vision [36]. These methods include single-scale Retinex (SSR) [11, 33], multi-scale Retinex (MSR) [56], and adaptive MS-Retinex (AMS-R) [37, 40], each with its unique advantages.

To address the limitations of classical techniques, recent research has turned to deep learning approaches. Convolutional neural networks (CNNs) have been employed in various papers such as [29] and [79], autoencoders in [5] and [63], and generative adversarial networks (GANs) in [28], and [48]. Additionally, diffusion-based methods have been explored, as seen in the work by [50]. These modern techniques enhance classical methods by decoupling illumination and reflectance, utilizing attention mechanisms, or applying normalizing flows to achieve improved performance and noise robustness. Some deep learning methods even operate without the need for paired low-light/well-lit image data, enhancing their practicality for real-world applications [35, 39, 75, 78].

However, image enhancement algorithms are often not suitable to be used as an entire image processing pipeline for enhancing low-light images. This is often because most image enhancement algorithms are only

trained with processed image pairs, which causes such methods to be unable to perform the multitude of steps required for the image processing pipeline as well as causes these algorithms to not be able to handle very low-light images.

2.4. Learning-based image processing pipeline

Recently there have been multiple data-driven or learned image-processing pipelines developed to improve low-light images. These approaches learn image-processing pipelines using hand-crafted features [72] or deep-learning approaches [8, 49, 81]. These algorithms aim to propose a new image processing pipeline directly on RAW images and the corresponding produced images are enhanced.

There have been approaches that train models to learn the entire image-processing pipeline [8, 17, 69, 71, 81, 82]. These works usually train a model that takes a RAW image or a lightly processed RAW image and handles the entire pipeline of creating an image while performing some enhancement like producing denoised images from low-light settings. Our work is closest to LSID [8] which tackles the same problem, of learning an image processing pipeline for extremely low-light images. LSID tackles this problem by learning a convolutional network and handcrafting some features like the brightness of the generated image or the black level, this is trained on image pairs of scenes with high exposure time and the same scene with varying levels of lower exposure time, which generates noisy images. The convolution network not only learns the entire image processing pipeline but also denoises this image well. However, our work develops a new generative camera-processing pipeline contrary to LSID. Our approach is based on fine-tuning existing diffusion models to learn this image-processing pipeline which improves upon the performance of LSID and other subsequent methods and is also the first work to present a diffusion model fine-tuned on RAW images. Our approach is also different from other Diffusion-based approaches [50, 52] which only perform image enhancement on post-processed images, unlike our approach which is a new end-to-end generative camera processing pipeline.

3. Method

Given a casually captured RAW image in very low-light settings, without any description or any other information about what the scene contains, our objective is to generate a processed final image with the scene as if it was taken in higher-light settings with more exposure. We do not impose any restrictions on input image capture settings and the captured images can have




Figure 3. Our fine-tuning approach consists of first naively packing the RAW images, subtracting the black level, scaling the image by an externally set amplification factor, upscaling the packed RAW images, and then fine-tuning a text-to-image latent diffusion model conditioned on these packed RAW images.

varying numbers of photons captured per pixel. Our approach being based on fine-tuning, in the future can easily be modified to be guided by more diverse text and support other high-level image editing operations like stylization, inpainting, accessorization, and property modification. We next provide some background on conditional latent diffusion models (Sec. 3.1), and then present our fine-tuning technique to develop a generative ISP (Sec. 3.2).

3.1. Conditional Latent Diffusion Models

Latent Diffusion models [59] project the input into a lower-dimensional latent space and train the diffusion model on this latent space instead of applying the diffusion process on a high-dimensional input pixel space. Specifically, the forward and backward diffusion processes occur in a lower-dimensional latent space and an encoder-decoder architecture is trained on a large image dataset to translate images into latent codes. At inference time, a random noise latent code goes through the backward diffusion process, and the pre-trained decoder is used to generate the final image. We have a conditional latent diffusion model ϵ_θ , which can be interpreted as a sequence of denoising autoencoder $\epsilon_\theta(\mathbf{z}_t, t, \dots)$; $t \sim \mathcal{U}(\{1, \dots, T\})$, that is trained using a squared error loss (or represent the denoising term in the ELBO) to denoise some variably noisy images as follows:

$$\mathbb{E}_{\mathcal{E}(\mathbf{x}), \mathbf{c}, \epsilon \sim \mathcal{N}(0, I), t} [\|\epsilon - \epsilon_\theta(\mathbf{z}_t, t, \tau_\theta(\mathbf{c}))\|_2^2] \quad (1)$$

where \mathbf{x}_t is some noisy image, \mathbf{c} is a conditioning vector, τ_θ is an encoder that projects \mathbf{c} to an intermediate representation, \mathcal{E} encodes \mathbf{x}_t into a latent representation \mathbf{z}_t , and we jointly optimize ϵ_θ and τ_θ .

3.2. DiffuseRAW: Generative ISP

In typical image processing pipelines, RAW data obtained from an imaging sensor undergoes several se-

quential module applications, including white balance, demosaicing, denoising, sharpening, color space conversion, gamma correction, etc., where these modules are frequently calibrated specifically for individual camera models. [62] suggested using numerous localized, linear, and learned filters collectively termed “L3” filters to mimic the intricate nonlinear pathways witnessed in current-day commercial imaging solutions. Previous deep-learned ISP methods are often trained using some kind of convolutional architecture [8, 62]. However, processing low-light images is an ill-posed problem since many possible reconstructed images could be consistent with the measurements. Using pre-trained generative models allows us to exploit learned priors and generate processed images that obey the learned natural image statistics. Our approach needs to learn these multiple stages of a typical image processing pipeline for it to be an end-to-end generative ISP. This pipeline is summarized in Figure 3.

Demosaicing. Demosaicing is a crucial step in image processing pipelines, particularly in digital imaging systems that use color filter arrays to capture images. In such systems, each pixel on the image sensor is assigned to detect only one primary color resulting in a spatially undersampled color image. Demosaicing reconstructs a full-color image from this mosaic by estimating the missing color values at each pixel location.

The sensor first gives us a Bayer RAW image and then our approach involves compressing the given information naively into four distinct channels. Consequently, there is a reduction in the original spatial resolution by half in both dimensions, $(\frac{H}{2}, \frac{W}{2})$ -sized images. Additionally, we then subtract the black level and scale the brightness of the input images by some amplification factor, α that is manually set. We then employ simple bicubic upsampling to upsample the input image back to (H, W) -sized images

Fine-tuning Approach. Incorporating a conditioning vector in classifier-free diffusion guidance [26] relies on training a conditional and an unconditional model for some conditioning \mathbf{c} . This score estimate is modified towards the conditional model and away from the unconditional model. So we can calculate the modified score estimate as follows:

$$\tilde{e}_\theta(\mathbf{z}_t, \mathbf{c}) = \underbrace{e_\theta(\mathbf{z}_t, \phi)}_{\text{unconditional model}} + s \left(\underbrace{e_\theta(\mathbf{z}_t, \mathbf{c})}_{\text{conditional model}} - \underbrace{e_\theta(\mathbf{z}_t, \phi)}_{\text{unconditional model}} \right) \quad (2)$$

move towards conditional model
and away from unconditional model

where s represents the level of guidance, \mathbf{c} represents the conditioning vector, ϕ represents a conditioning vector representing some fixed null value, and \mathbf{z}_t represents the latent representation.

Following this, in our case, we have two conditioning vectors: the input dark images \mathbf{c}_I and the text conditioning vector for high-level editing tasks \mathbf{c}_T . We use classifier-free diffusion guidance for both of these conditioning vectors and compose score estimates from these two conditioning vectors [41]. During fine-tuning, we randomly set only $\mathbf{c}_T = \phi_T$ for 5% of examples, and both $\mathbf{c}_I = \phi_I$ and $\mathbf{c}_T = \phi_T$ for 5% of examples which allows the model to be capable of conditional or unconditional denoising with respect to both or any of conditional inputs: on RAW images or text conditioning. As in InstructPix2Pix [41], we also have two separate guidance scales s_I representing the image guidance scale, and s_T representing the text guidance scale, which can be adjusted to trade off how strongly the generated samples correspond with the input RAW image and how strongly they correspond with the text conditioning. Thus, our score estimate is as follows:

$$\tilde{e}_\theta(\mathbf{z}_t, \mathbf{c}_I, \mathbf{c}_T) = \underbrace{e_\theta(\mathbf{z}_t, \phi_I, \phi_T)}_{\text{unconditional model}} + s_I \left(\underbrace{e_\theta(\mathbf{z}_t, \mathbf{c}_I, \phi_T)}_{\text{image conditioned}} - \underbrace{e_\theta(\mathbf{z}_t, \phi_I, \phi_T)}_{\text{unconditional model}} \right) + s_T \left(\underbrace{e_\theta(\mathbf{z}_t, \mathbf{c}_I, \mathbf{c}_T)}_{\text{image and text conditioned}} - \underbrace{e_\theta(\mathbf{z}_t, \mathbf{c}_I, \phi_T)}_{\text{image conditioned}} \right) \quad (3)$$

move towards \mathbf{c}_I conditioned model
and away from unconditional model

move towards conditioned model
and away from \mathbf{c}_I conditioned model

Our approach is based on fine-tuning standard latent diffusion models like Stable Diffusion [59] which

Model	$\times 100$	$\times 250$	$\times 300$
	LPIPS ↓	LPIPS ↓	LPIPS ↓
LRD [82]	0.2822	0.3013	0.3394
LSID [8]	0.3982	0.4342	0.4648
Ours	0.2793	0.3032	0.3184

Table 1. Quantitative results on the Sony subset of the SID dataset [8] in terms of LPIPS ↓ [84] for the best open-spurce models in the split. The results are conducted on different amplification ratios ($\times 100$, $\times 250$, $\times 300$).

Model	$\times 100$	$\times 200$
	LPIPS ↓	LPIPS ↓
PMN [17]	0.3197	-
LRD [82]	-	0.3184
Ours	0.3034	0.3101

Table 2. Quantitative results on the SonyA7S2 subset of the ELD dataset [71] in terms of LPIPS ↓ [84] for the best open-spurce models in the split. The results are conducted on different amplification ratios ($\times 100$, $\times 200$).

are often trained for the task of text-to-image generation. To support our suite of tasks we add additional channels to the first layer of the underlying UNet, concatenating the latent codes for variably noisy images \mathbf{z}_t and latent codes for the image conditioning $\mathcal{E}(\mathbf{c}_I)$ [41]. We use the same text encoder for the text conditioning vectors \mathbf{c}_T as in these text-to-image generation models which were trained with text conditioning contrary to the popular approach of replacing the text encoder with an independent image encoder [54]. To do so, our approach also relies on generating a set of text prompts for this task and uses these text prompts in addition to the RAW and long-exposure processed image pairs. These text prompts could also potentially be used to guide the model for a variety of other downstream high-level tasks as well.

3.3. Training

We fine-tune our models from Stable Diffusion v1.5 [54, 59] and fine-tune these models on the See in the Dark (SID) [8] and Extreme Low-light Denoising (ELD) [71] dataset. We fine-tune these models independently for each camera type. For both of these datasets, while training we use the short-exposure Bayer RAW images which are then packed into four channels (R-Gr-B-Gb) from the dataset, and the corresponding processed long-exposure images (us-

ing `rawpy` [57]) in the sRGB space. Following [8,71] we set the amplification ratio to be the exposure difference between the input and reference images (we use $\times 100$, $\times 250$, $\times 300$ for the SID dataset and $\times 100$, $\times 200$ for the ELD dataset) for both training and testing. We train the model on 512^2 patches which are made as suggested in [64].

4. Experiments

4.1. Implementation Details

Dataset. We use noisy-clean pairs from the Sony subset of the SID [8] and ELD [71] dataset. The images in the Sony subset of the SID dataset which are collected using Sony $\alpha 7SII$ in 231 static scenes across three low-light factors. There are 1865 image pairs for training, 234 for validation, and 598 for testing. The Sony subset of the ELD dataset consists of 60 low-light image pairs across two low-light factors, which are also captured using the same camera as the SID dataset.

We furthermore, augment all training samples with random crop, random horizontal flip, and use AutoAugment [12]. We use the Adam [34] optimizer, with a learning rate of $5e - 5$, 500 warmup steps, $1e - 2$ weight decay, cosine annealing, and use a batch-size of 64. We train the model for 300 epochs. Our code is in PyTorch 2.0 [53]. We use a number of open-source packages to develop our training workflows, par-

Model	$\times 100$		$\times 200$	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
BM3D [13]	37.69	0.803	34.06	0.696
N2N [38]	41.63	0.856	37.98	0.775
P+G [18, 71]	41.76	0.930	39.33	0.872
NoiseFlow [1]	41.05	0.925	39.23	0.889
Exposure Diffusion [69]	43.29	0.929	40.39	0.873
ELLE [67]	43.11	0.940	40.30	0.884
Starlight [48]	43.80	0.936	40.86	0.884
ELD [71]	45.45	0.975	43.43	0.954
LRD [82]	44.95	0.979	43.32	0.966
LSID [8]	44.47	0.968	41.97	0.928
PMN [17]	46.50	0.985	44.51	0.973
LLD* [7]	46.74	0.986	44.95	0.977
Ours	47.56	0.985	45.12	0.981

Table 3. Quantitative results on the SonyA7S2 subset of the ELD dataset [71] in terms of PSNR \uparrow and SSIM \uparrow . The results are conducted on different amplification ratios ($\times 100$, $\times 200$).

ticularly we use Huggingface Diffusers [66], Accelerate [21], Transformers [74], and `rawpy` [57]. We utilized mixed-precision training with PyTorch’s native AMP (through `torch.cuda.amp`) for mixed-precision training which allowed us to obtain significant boosts in the overall model training time. Our experiments are conducted on TPUv3-8 and we run test-scripts on Tesla A100.

4.2. Comparison with state-of-the-art methods

In order to demonstrate the reliability of our proposed generative camera ISP: DiffuseRAW, we conduct experiments on the SID [8] and ELD [71] datasets for learning the entire image processing pipeline from the RAW data.

We quantitatively compare our results on the Sony subset of the SID [8] dataset with non-learned methods like BM3D [13] and A-BM3D [47], models based on different noise models [18,38,71], and neural-net based learning methods [1, 8, 48, 69], which include state-of-the-art models. We summarize these results in Table 4 and Table 1 and demonstrate the performance of our approach over different low-light factors. We compare these models based on PSNR \uparrow , SSIM \uparrow , and LPIPS \downarrow [84]. As we can see Table 4 and Table 1, DiffuseRAW achieves better performance than other state-of-the-art



Figure 4. The scaled images and long-exposure ground-truth images for the same scene across three different low-light factors $\times 100$ (first row), $\times 250$ (second row), and $\times 300$ (third row) captured on Sony $\alpha 7SII$ with an ISO of 10000.

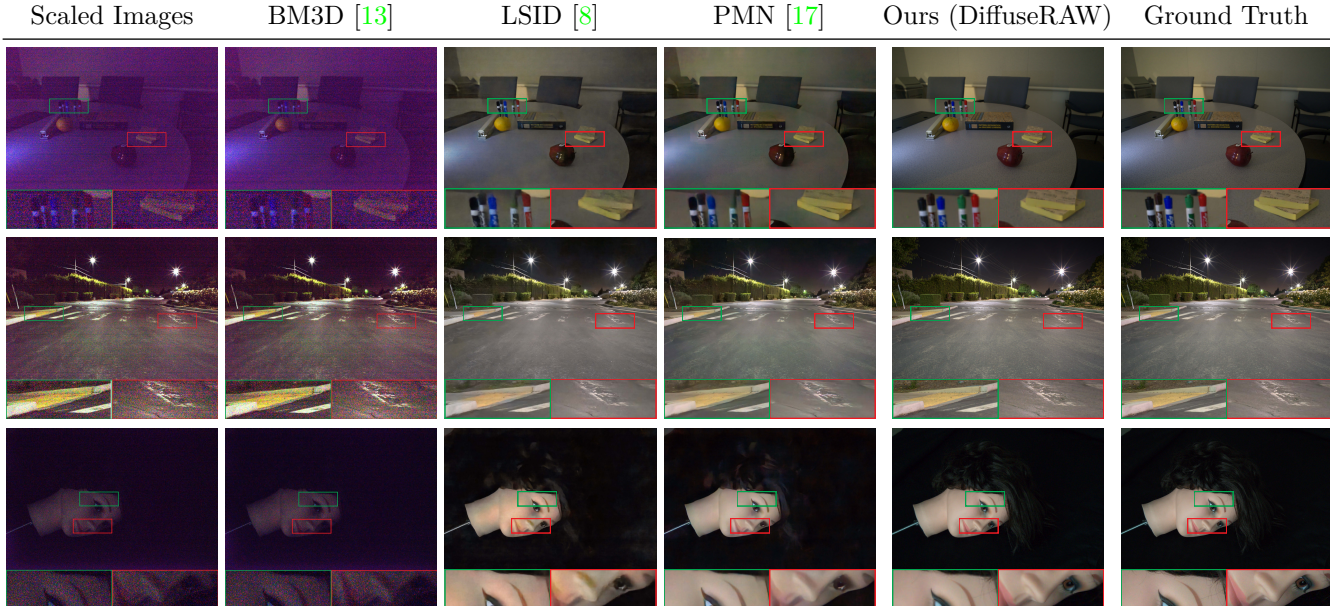


Figure 5. We show results on learning the entire image processing pipeline for extremely low-light enhancement and we compare our approach qualitatively to the best-performing open-source model. These images are taken with Sony α 7SII and the illuminance at the camera is < 0.1 lux. We show both indoor and outdoor scenes. **(Best viewed in color and with zoom.)**

learned methods across most metrics and most low-light factors. We particularly observe that DiffuseRAW performs significantly better than current state-of-the-art models for $\times 300$ low-light factor which is a more challenging task.

We also show visual results for our approach and compare it with other approaches in Figure 5. Though we notice that our model particularly performs very well for $\times 250$ and $\times 300$, we show the difference between different low-light factors in Figure 4, we show that our

model performs well for the $\times 100$ task as well which we show qualitatively in Figure 6. We also compare our model on the same scene across low-light factors which we show in Figure 7.

We quantitatively compare our results on the Sony subset of the ELD [71] dataset with non-learned methods, models based on different noise models [18, 38, 71], and neural-net based learning methods [1, 8, 48, 69], which include state-of-the-art models. We summarize these results in Table 3 and Table 2 and demonstrate the performance of our approach over different low-light factors. We compare these models based on PSNR \uparrow , SSIM \uparrow , and LPIPS \downarrow [84]. As we can see Table 3 and Table 2, DiffuseRAW achieves better performance than other state-of-the-art learned methods for $\times 200$ low-light factor which is a more challenging task and performs marginally better across the $\times 100$ low-light factor in terms of PSNR \uparrow .

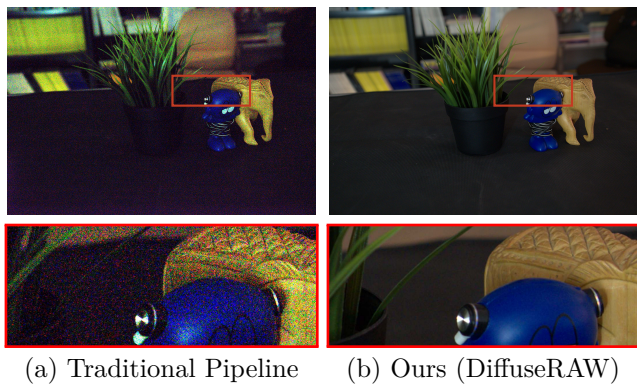


Figure 6. We qualitatively compare our model with a traditional pipeline with $\times 100$ low-light factor on the SID [8] dataset.

5. Conclusion

Our method creates a new generative camera ISP for extreme low-light image enhancement that learns the entire image processing pipeline as a diffusion model. Our approach showcases how we can make use of ample image priors in pre-trained text-to-image latent diffusion models for learning tasks in the RAW space. We not only demonstrate the effectiveness of working di-

Model	$\times 100$		$\times 250$		$\times 300$	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
BM3D [13]	32.92	0.758	29.56	0.686	28.88	0.674
A-BM3D [47]	33.79	0.743	27.24	0.518	26.52	0.558
N2N [38]	37.42	0.853	33.48	0.725	32.37	0.686
P+G [18, 71]	38.31	0.884	34.39	0.765	33.37	0.730
NoiseFlow [1]	38.89	0.929	35.80	0.867	32.29	0.801
Exposure	38.89	0.902	36.02	0.832	35.00	0.808
Diffusion [69]						
ELLE [67]	40.09	0.931	36.13	0.863	32.54	0.782
Starlight [48]	40.47	0.926	36.25	0.858	32.99	0.780
ELD [71]	41.95	0.953	39.44	0.931	36.36	0.911
LRD [82]	41.95	0.956	39.25	0.931	36.03	0.909
LSID [8]	42.06	0.955	39.60	0.938	36.85	0.923
SFRN [85]	42.29	0.951	40.22	0.938	36.87	0.917
PMN [17]	43.16	0.960	40.92	0.947	37.77	0.934
LLD* [7]	43.36	0.961	41.02	0.948	37.80	0.935
Ours	44.93	0.986	43.46	0.956	39.98	0.951

Table 4. Quantitative results on the Sony subset of the SID dataset [8] in terms of PSNR \uparrow and SSIM \uparrow . The results are conducted on different amplification ratios ($\times 100$, $\times 250$, $\times 300$).

rectly on RAW sensor data for such downstream tasks rather than operating on images processed through some traditional image-processing pipeline but also

show the effectiveness of diffusion models to learn in the RAW sensor data space. Parallel to how diffusion models tackle these problems today: learning on processed images, we hope that with this work, we foster further research on using diffusion models and other generative models for tackling the image processing problem and other problems in the RAW space.




Figure 7. We compare our model on the same scene on different low-light factors for which we observe that the model produces consistent results across different low-light factors, we demonstrate $\times 100$ for the first row and $\times 300$ for the second row.

Limitations. Although our approach, DiffuseRAW generates plausible results, there are several limitations of this work and avenues for future work. An important aspect of performing general low-light camera processing is being able to handle the changing exposure of an image. Secondly, our approach currently does not learn exposure settings well for general scenes or cameras. Our model in its current form is also only able to handle the ISP pipeline for images, an interesting avenue this work could be expanded to would be to explore how we could use similar approaches for video ISP tasks. Finally, we notice, that our approach heavily relies on the use of a proper augmentation regime which could make it challenging to train a model that works for multiple kinds of cameras and learns a universal ISP.

Acknowledgements

We would like to thank David Lindell of the University of Toronto for multiple insightful discussions. The authors would like to thank Google TPU Research Cloud (TRC) program ¹ for providing free access to TPUs.

References

- [1] Abdelrahman Abdelhamed, Marcus A. Brubaker, and Michael S. Brown. Noise flow: Noise modeling with conditional normalizing flows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 7, 8, 9
- [2] M. Abdullah-Al-Wadud, Md. Hasanul Kabir, M. Ali Akber Dewan, and Oksam Chae. A dynamic histogram equalization for image contrast enhancement. *IEEE Transactions on Consumer Electronics*, 53(2):593–600, 2007. 4
- [3] Miika Aittala and Fredo Durand. Burst image deblurring using permutation invariant convolutional neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 4
- [4] Marcelo Bertalmio, Vicent Caselles, Edoardo Provenzi, and Alessandro Rizzi. Perceptual color correction through variational techniques. *IEEE Transactions on Image Processing*, 16(4):1058–1072, 2007. 4
- [5] Siavash Arjomand Bigdeli and Matthias Zwicker. Image restoration using autoencoding priors, 2017. 4
- [6] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. 2
- [7] Yue Cao, Ming Liu, Shuai Liu, Xiaotao Wang, Lei Lei, and Wangmeng Zuo. Physics-guided iso-dependent sensor noise modeling for extreme low-light photography. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5744–5753, June 2023. 7, 9
- [8] 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 (CVPR)*, June 2018. 2, 3, 4, 5, 6, 7, 8, 9
- [9] Yunjin Chen and Thomas Pock. Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1256–1272, 2016. 3
- [10] H.D. Cheng and X.J. Shi. A simple and effective histogram equalization approach to image enhancement. *Digital Signal Processing*, 14(2):158–170, 2004. 4
- [11] Doo Hyun Choi, Ick Hoon Jang, Mi Hye Kim, and Nam Chul Kim. Color image enhancement using single-scale retinex based on an improved image formation model. In *2008 16th European Signal Processing Conference*, pages 1–5, 2008. 4
- [12] Ekin D. Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V. Le. Autoaugment: Learning augmentation policies from data, 2019. 7
- [13] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8):2080–2095, 2007. 2, 7, 8, 9
- [14] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007. 3
- [15] Mauricio Delbracio and Guillermo Sapiro. Burst deblurring: Removing camera shake through fourier burst accumulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 4
- [16] Michael Elad and Michal Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing*, 15(12):3736–3745, 2006. 3
- [17] Hansen Feng, Lizhi Wang, Yuzhi Wang, and Hua Huang. Learnability enhancement for low-light raw denoising: Where paired real data meets noise modeling. In *Proceedings of the 30th ACM International Conference on Multimedia*, MM ’22, page 1436–1444, New York, NY, USA, 2022. Association for Computing Machinery. 4, 6, 7, 8, 9
- [18] Alessandro Foi, Mejdí Trimeche, Vladimir Katkovnik, and Karen Egiazarian. Practical poissonian-gaussian noise modeling and fitting for single-image raw-data. *IEEE Transactions on Image Processing*, 17(10):1737–1754, 2008. 7, 8, 9
- [19] Xueyang Fu, Yinghao Liao, Delu Zeng, Yue Huang, Xiao-Ping Zhang, and Xinghao Ding. A probabilistic method for image enhancement with simultaneous illumination and reflectance estimation. *IEEE Transactions on Image Processing*, 24(12):4965–4977, 2015. 4
- [20] Shuhang Gu, Lei Zhang, Wangmeng Zuo, and Xianguo Feng. Weighted nuclear norm minimization with application to image denoising. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2862–2869, 2014. 3
- [21] Sylvain Gugger, Lysandre Debut, Thomas Wolf, Philipp Schmid, Zachary Mueller, Sourab Mangrulkar, Marc Sun, and Benjamin Bossan. Accelerate: Training and inference at scale made simple, efficient and adaptable. <https://github.com/huggingface/accelerate>, 2022. 7
- [22] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE/CVF*

¹<https://sites.research.google/trc/about/>

- Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [23] Xiaojie Guo, Yu Li, and Haibin Ling. Lime: Low-light image enhancement via illumination map estimation. *IEEE Transactions on Image Processing*, 26(2):982–993, 2017. 4
- [24] 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 Trans. Graph.*, 35(6), dec 2016. 2
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. 3
- [26] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. 6
- [27] Haidi Ibrahim and Nicholas Sia Pik Kong. Brightness preserving dynamic histogram equalization for image contrast enhancement. *IEEE Transactions on Consumer Electronics*, 53(4):1752–1758, 2007. 4
- [28] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Wespe: Weakly supervised photo enhancer for digital cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2018. 4
- [29] Andrey Ignatov and Radu Timofte. Ntire 2019 challenge on image enhancement: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019. 4
- [30] Anil K Jain. *Fundamentals of digital image processing*. Prentice-Hall, Inc., 1989. 2
- [31] Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou, and Zhangyang Wang. Enlightengan: Deep light enhancement without paired supervision. *IEEE Transactions on Image Processing*, 30:2340–2349, 2021. 2
- [32] D.J. Jobson, Z. Rahman, and G.A. Woodell. A multi-scale retinex for bridging the gap between color images and the human observation of scenes. *IEEE Transactions on Image Processing*, 6(7):965–976, 1997. 4
- [33] Manpreet Kaur, Jasdeep Kaur, and Jappreet Kaur. Survey of contrast enhancement techniques based on histogram equalization. *International Journal of Advanced Computer Science and Applications*, 2(7), 2011. 4
- [34] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. 7
- [35] Ashutosh Kulkarni, Shruti S. Phutke, and Subrahmanyam Murala. Unified transformer network for multi-weather image restoration. In Leonid Karlin-sky, Tomer Michaeli, and Ko Nishino, editors, *Computer Vision – ECCV 2022 Workshops*, pages 344–360, Cham, 2023. Springer Nature Switzerland. 4
- [36] Edwin H Land. The retinex theory of color vision. *Scientific american*, 237(6):108–129, 1977. 4
- [37] Chang-Hsing Lee, Jau-Ling Shih, Cheng-Chang Lien, and Chin-Chuan Han. Adaptive multiscale retinex for image contrast enhancement. In *2013 International Conference on Signal-Image Technology & Internet-Based Systems*, pages 43–50. IEEE, 2013. 4
- [38] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2Noise: Learning image restoration without clean data. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2965–2974. PMLR, 10–15 Jul 2018. 7, 8, 9
- [39] 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 (ICCV) Workshops*, pages 1833–1844, October 2021. 4
- [40] Haoning Lin and Zhenwei Shi. Multi-scale retinex improvement for nighttime image enhancement. *optik*, 125(24):7143–7148, 2014. 4
- [41] Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B. Tenenbaum. Compositional visual generation with composable diffusion models. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 423–439, Cham, 2022. Springer Nature Switzerland. 6
- [42] Zhen Liu, Wenjie Lin, Xinpeng Li, Qing Rao, Ting Jiang, Mingyan Han, Haoqiang Fan, Jian Sun, and Shuaicheng Liu. Adnet: Attention-guided deformable convolutional network for high dynamic range imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 463–470, 2021. 3
- [43] Kin Gwn Lore, Adedotun Akintayo, and Soumik Sarkar. Llnet: A deep autoencoder approach to natural low-light image enhancement. *Pattern Recognition*, 61:650–662, 2017. 2
- [44] Kin Gwn Lore, Adedotun Akintayo, and Soumik Sarkar. Llnet: A deep autoencoder approach to natural low-light image enhancement. *Pattern Recognition*, 61:650–662, 2017. 3
- [45] Calvin Luo. Understanding diffusion models: A unified perspective, 2022. 3
- [46] Paras Maharjan, Li Li, Zhu Li, Ning Xu, Chongyang Ma, and Yue Li. Improving extreme low-light image denoising via residual learning. In *2019 IEEE International Conference on Multimedia and Expo (ICME)*, pages 916–921, 2019. 2
- [47] Markku Makitalo and Alessandro Foi. Optimal inversion of the anscombe transformation in low-count

- poisson image denoising. *IEEE Transactions on Image Processing*, 20(1):99–109, 2011. 7, 9
- [48] Kristina Monakhova, Stephan R. Richter, Laura Waller, and Vladlen Koltun. Dancing under the stars: Video denoising in starlight. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16241–16251, June 2022. 2, 3, 4, 7, 8, 9
- [49] Igor Morawski, Yu-An Chen, Yu-Sheng Lin, Shushil Dangi, Kai He, and Winston H. Hsu. Genisp: Neural isp for low-light machine cognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 630–639, June 2022. 4
- [50] Cindy M. Nguyen, Eric R. Chan, Alexander W. Bergman, and Gordon Wetzstein. Diffusion in the dark: A diffusion model for low-light text recognition, 2023. 2, 3, 4
- [51] Rodrigo Palma-Amestoy, Edoardo Provenzi, Marcelo Bertalmío, and Vincent Caselles. A perceptually inspired variational framework for color enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(3):458–474, 2009. 4
- [52] Savvas Panagiotou and Anna S. Bosman. Denoising diffusion post-processing for low-light image enhancement, 2023. 2, 4
- [53] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019. 7
- [54] Justin Pinkney. Experiments with stable diffusion. <https://github.com/justinpinkney/stable-diffusion>, 2023. 6
- [55] Javier Portilla, Vasily Strela, Martin J Wainwright, and Eero P Simoncelli. Image denoising using scale mixtures of gaussians in the wavelet domain. *IEEE Transactions on Image processing*, 12(11):1338–1351, 2003. 3
- [56] Zia-ur Rahman, Daniel J Jobson, and Glenn A Woodell. Multi-scale retinex for color image enhancement. In *Proceedings of 3rd IEEE international conference on image processing*, volume 3, pages 1003–1006. IEEE, 1996. 4
- [57] Maik Riechert. Rawpy. <https://github.com/Letmaik/rawpy>. Accessed: 2023-12-01. 7
- [58] Alessandro Rizzi, Carlo Gatta, and Daniele Marini. From Retinex to Automatic Color Equalization: issues in developing a new algorithm for unsupervised color equalization. *Journal of Electronic Imaging*, 13(1):75–84, 2004. 4
- [59] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3, 5, 6
- [60] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992. 3
- [61] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023. 2
- [62] 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. 5
- [63] Tomoyoshi Shimobaba, Yutaka Endo, Ryuji Hirayama, Yuki Nagahama, Takayuki Takahashi, Takashi Nishitsuji, Takashi Kakue, Atsushi Shiraki, Naoki Takada, Nobuyuki Masuda, and Tomoyoshi Ito. Autoencoder-based holographic image restoration. *Appl. Opt.*, 56(13):F27–F30, May 2017. 4
- [64] Ivan Skorokhodov, Sergey Tulyakov, Yiqun Wang, and Peter Wonka. Epigraf: Rethinking training of 3d gans. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 24487–24501. Curran Associates, Inc., 2022. 7
- [65] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*, pages 839–846. IEEE, 1998. 2
- [66] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. 7
- [67] Jing Wang, Yitong Yu, Songtao Wu, Chang Lei, and Kuanhong Xu. Rethinking noise modeling in extreme low-light environments. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6, 2021. 7, 9
- [68] Yufei Wang, Renjie Wan, Wenhan Yang, Haoliang Li, Lap-Pui Chau, and Alex Kot. Low-light image enhancement with normalizing flow. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(3):2604–2612, Jun. 2022. 2
- [69] Yufei Wang, Yi Yu, Wenhan Yang, Lanqing Guo, Lap-Pui Chau, Alex C. Kot, and Bihan Wen. Exposediffusion: Learning to expose for low-light image enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12438–12448, October 2023. 4, 7, 8, 9
- [70] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 3
- [71] Kaixuan Wei, Ying Fu, Jiaolong Yang, and Hua Huang. A physics-based noise formation model for ex-

- treme low-light raw denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 4, 6, 7, 8, 9
- [72] Zebin Wei, Huiwen Zhang, Min Zhou, Yisong Zou, and Shanfeng Hou. Real-time low-light-level image-processing system. In *Photoelectronic Detection and Imaging: Technology and Applications' 93*, volume 1982, pages 427–436. SPIE, 1993. 4
- [73] Patrick Wieschollek, Bernhard Schölkopf, Hendrik P. A. Lensch, and Michael Hirsch. End-to-end learning for image burst deblurring. In Shang-Hong Lai, Vincent Lepetit, Ko Nishino, and Yoichi Sato, editors, *Computer Vision – ACCV 2016*, pages 35–51, Cham, 2017. Springer International Publishing. 4
- [74] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-Art Natural Language Processing. pages 38–45. Association for Computational Linguistics, Oct. 2020. 7
- [75] Jie Xiao, Xueyang Fu, Feng Wu, and Zheng-Jun Zha. Stochastic window transformer for image restoration. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 9315–9329. Curran Associates, Inc., 2022. 4
- [76] Rui kang Yang, Lin Yin, Moncef Gabbouj, Jaakko Astola, and Yrjö Neuvo. Optimal weighted median filtering under structural constraints. *IEEE transactions on signal processing*, 43(3):591–604, 1995. 2
- [77] Hui Yin, Yuanhao Gong, and Guoping Qiu. Fast and efficient implementation of image filtering using a side window convolutional neural network. *Signal Processing*, 176:107717, 2020. 3
- [78] 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 (CVPR)*, pages 5728–5739, June 2022. 4
- [79] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for real image restoration and enhancement. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV 16*, pages 492–511. Springer, 2020. 2, 4
- [80] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for fast image restoration and enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1934–1948, 2023. 2
- [81] Syed Waqas Zamir, Aditya Arora, Salman Khan, Fahad Shahbaz Khan, and Ling Shao. Learning digital camera pipeline for extreme low-light imaging. *Neurocomputing*, 452:37–47, 2021. 4
- [82] Feng Zhang, Bin Xu, Zhiqiang Li, Xinran Liu, Qingbo Lu, Changxin Gao, and Nong Sang. Towards general low-light raw noise synthesis and modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10820–10830, October 2023. 4, 6, 7, 9
- [83] 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. 2
- [84] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 6, 7, 8
- [85] 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 (ICCV)*, pages 4593–4601, October 2021. 9
- [86] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 2