

Invertible Image Signal Processing

Yazhou Xing* Zian Qian* Qifeng Chen
The Hong Kong University of Science and Technology

Abstract

Unprocessed RAW data is a highly valuable image format for image editing and computer vision. However, since the file size of RAW data is huge, most users can only get access to processed and compressed sRGB images. To bridge this gap, we design an Invertible Image Signal Processing (InvISP) pipeline, which not only enables rendering visually appealing sRGB images but also allows recovering nearly perfect RAW data. Due to our framework’s inherent reversibility, we can reconstruct realistic RAW data instead of synthesizing RAW data from sRGB images without any memory overhead. We also integrate a differentiable JPEG compression simulator that empowers our framework to reconstruct RAW data from JPEG images. Extensive quantitative and qualitative experiments on two DSLR demonstrate that our method obtains much higher quality in both rendered sRGB images and reconstructed RAW data than alternative methods.

1. Introduction

Professional photographers can choose to process RAW images by themselves instead of RGB images to produce images with better visual effects as the RAW data captures unprocessed scene irradiance at each in 12-14 bits by a camera. Due to its linear relationship with scene irradiance, raw sensor data is also a better choice than RGB images for many image editing and computer vision tasks, such as photometric stereo, intrinsic image decomposition, image denoising, reflection removal, and image super resolution [4, 7, 27, 16, 30, 39, 40, 45]. However, accessing RAW images can be quite hard due to their memory-demanding property: RAW images may be discarded during the process of data storing, transferring, and sharing. In this paper, we are interested in the question: can users get access to the real RAW data without explicitly storing it?

Due to the great advantages of RAW images, there have been many approaches to provide the mapping from sRGB images to their RAW counterparts [2, 7, 29, 31, 36, 44]. Nguyen et al. [31] suggest explicitly storing the parameters

of sRGB-RAW mapping functions into the JPEG metadata for the prospective RAW reconstruction. Brooks et al. [7] use the prior information of the cameras (e.g., color correction matrices and digital gains) to reverse the ISP step-by-step. Another line of work [29, 36, 44] follows the inverse order of ISP and proposes learning-based methods to synthesize RAW data from sRGB images. However, these methods still rely on the underlying lossy in-camera ISP pipeline, and the recovered RAW images are inaccurate and may be different from the original ones.

In this work, we propose a novel and effective learned solution by redesigning the camera image signal processing pipeline as an invertible one, which can be aptly called *Invertible ISP (InvISP)*. Our learning-based InvISP enables rendering visually appealing RGB images in the forward process, and recovering nearly perfect quality raw sensor data from compressed RGB images through the inverse process. Our reconstructed RAW data is nearly identical with real RAW data and enables computer vision applications, such as image retouching and HDR reconstruction, as shown in Figure 1.

Designing an invertible ISP is not a trivial task for at least three reasons. First, some steps in the traditional ISP, such as denoising, tone mapping, and quantization, can lead to inevitable information lost from wide-range (12-bit or 14-bit) raw sensor data to 8-bit RGB images. Second, the invertible ISP should not produce visual artifacts such as halo and ghosting artifacts [18]. To render visually appealing sRGB images, denoising, demosaicing, color correction, white balance gain, tone mapping, and color enhancement must be designed carefully in ISP. Third, modern digital cameras store RGB images in the JPEG format, where the lossy compression process makes reconstructing high-quality RAW data highly challenging.

To overcome these challenges, we take advantage of the inherent reversibility of normalizing-flow-based models [12, 25] and design both the RAW-to-RGB and RGB-to-RAW mapping in our invertible ISP with one single invertible neural network. We deeply analyze the properties of traditional ISP and design specific modules that can not only well approximate the camera ISP but also reconstruct almost identical RAW data with the camera RAW data.

*Joint first authors

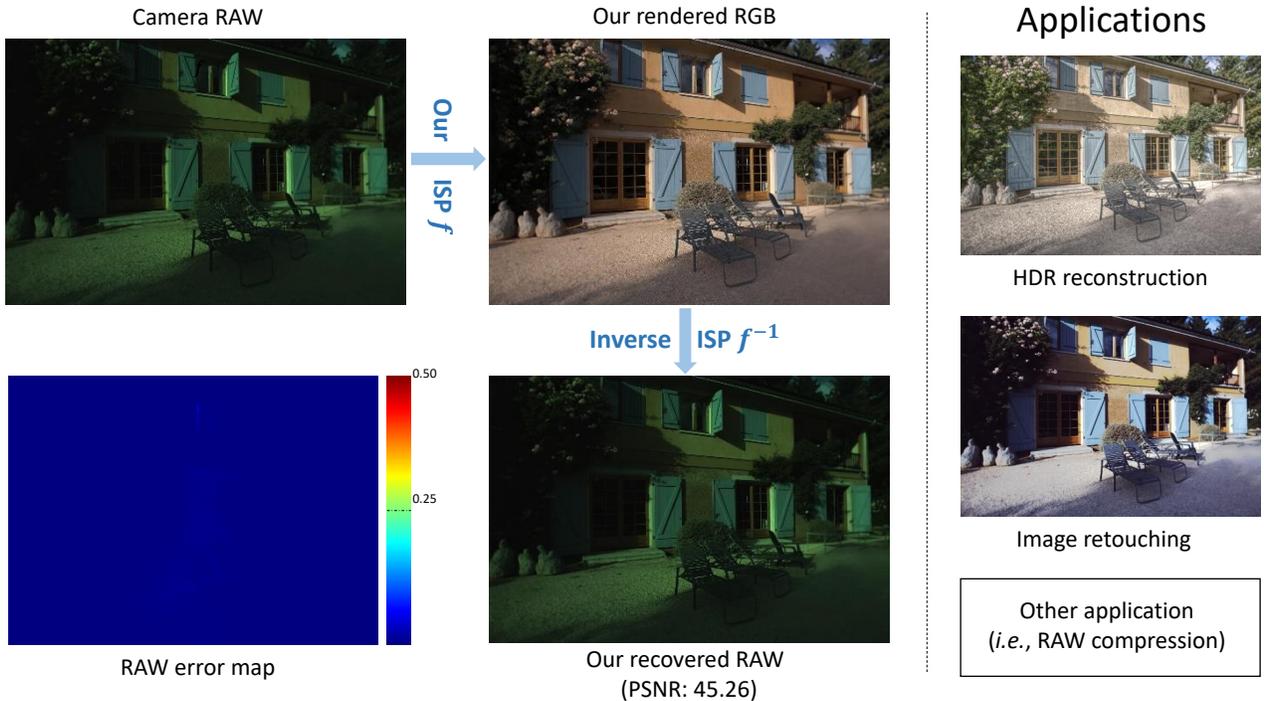


Figure 1. Our ISP model can not only render visually pleasing RGB images but also recover RAW images that are nearly the same as the original RAW data. The recovered RAW data are valuable for photographers and benefit a number of computer vision tasks such as HDR reconstruction [32], image retouching [22], and RAW compression. Here, the RAW images are visualized with bilinear demosaicing.

Specifically, we design our model with the composition of a stack of affine coupling layers and utilize the invertible 1×1 convolution as the learnable permutation function between the coupling layers. Besides, to empower our model to recover realistic RAW data from JPEG images, we integrate a differentiable JPEG simulator into our invertible neural network. We leverage the idea from Fourier transformation to replace the non-differentiable quantization step in JPEG compression. Thus, our end-to-end InvISP framework bypasses traditional ISP modules and minimizes the information loss for the RAW data and RGB image conversion. We bidirectionally train our network to optimize the RGB and RAW reconstruction process jointly. We experimentally prove that our framework can recover much better RAW data than state-of-the-art baselines without sacrificing the RGB reconstruction performance.

To the best of our knowledge, our framework is the first attempt for RAW data reconstruction from the perspective of redesigning the camera ISP as an invertible one. Our method can address the information loss issue in ISP modules and is robust to the JPEG compression step. We demonstrate the effectiveness of our method on two DSLR cameras and show that our method outperforms state-of-the-art baselines to a large extent. Moreover, we also exhibit potential applications through RAW data compression, image retouching, and HDR reconstruction.

2. Related Work

RAW Image Reconstruction. Recovering RAW from sRGB images has been well-studied [36, 31, 2, 7, 44, 30, 29]. Nguyen et al. [31] encode the parameters in ISP into JPEG metadata with 64KB overhead and use them to reconstruct RAW from JPEG images. Brooks et al. [7] propose to inverse the ISP pipeline step by step with camera priors. CIE-XYZ Net [2] proposes to recover RAW from sRGB images to the camera independent CIE-XYZ space. CycleISP [44] proposes to model the RGB-RAW-RGB data conversion cycle for synthesizing RAW from sRGB images. Unlike previous methods, we aim to fundamentally solve the RAW reconstruction problem by re-designing the camera ISP into an invertible one.

Image Signal Processing (ISP). Image signal processing pipeline (ISP) aims at converting raw sensor data to human-readable RGB images [20, 10, 9, 39, 45, 43, 28, 26, 16]. Heide et al. [20] merge the steps in the traditional ISP pipeline to avoid the accumulative error. Gharbi et al. [16] propose a method with end-to-end networks to learn RAW demosaicing and denoising jointly. Hasinoff et al. [19] propose a low-light imaging system for mobile devices. Other works [10, 39] focus on learning low-light enhancement ISP pipelines with CNNs. Zhang et al. [45] process RAW

for super-resolution task with U-net [38] to preserve high-frequency information. CameraNET [28] splits the ISP into two learning stages for CNN. Unlike the encoder-decoder style network adopted in previous work, we demonstrate that invertible neural networks own great potential for ISP pipeline and enable accurate RAW reconstruction.

Invertible Neural Networks. Normalizing flow-based invertible neural networks [21, 25, 11, 12] have become a popular choice in image generation tasks. Normalizing flow transforms a simple posterior distribution to a complex real-world distribution through a series of invertible transformations. NICE [11] is the first learning-based normalizing flow framework with the proposed additive coupling layers. RealNVP [12] modifies the additive coupling layer to both multiplication and addition, and composes the coupling layer in an alternating pattern such that all the inputs can be altered with equal chance. Kingma et al. [25] propose ActNorm layer and generalize channel-shuffle operations with invertible 1×1 convolution. Flow++ [21] modifies the affine coupling layer to logistics mixture CDF coupling flows and applies self-attention module.

3. Traditional ISP analysis

Modern digital cameras apply a series of operations, which form the image signal processing pipeline (ISP), to render RAW data to human-readable RGB images. These operations include white balance, demosaicing, denoising, color space transformation, tone mapping, and others [24]. Traditionally, every step of an ISP needs labor-intensive tuning for specific cameras, and inverting the traditional ISP steps is quite challenging. In this section, we analyze the existing modules with information loss in the traditional ISP. We show that the lossy steps in traditional ISP restrict the RAW reconstruction performance of a series of works [31, 7, 44] that aim at synthesizing RAW from sRGB images. Different from previous works, we re-design the ISP into an end-to-end invertible one that can bypass the traditional modules to minimize information loss during the RAW data and JPEG image conversion, which further enables recovering high-quality RAW data.

Quantization and tone mapping. Some ISP steps like demosaicing and gamma compression may involve float-point operations, and thus quantization is inevitable to transform the data into the integer type. For instance, the rounding function can bring $(-0.5, 0.5)$ intensity error to a pixel in theory. In the context of ISP, however, the tone mapping step can enlarge the intensity error much greater than ± 0.5 . The tone mapping curve is usually designed as S-curve that compresses the high-intensity value and low-intensity value more than mid-intensity values [37, 5]. As illustrated in Figure 2, for a 14-bit raw image, gamma compression makes pixel intensity at [16313, 16383] all be rounded to the max

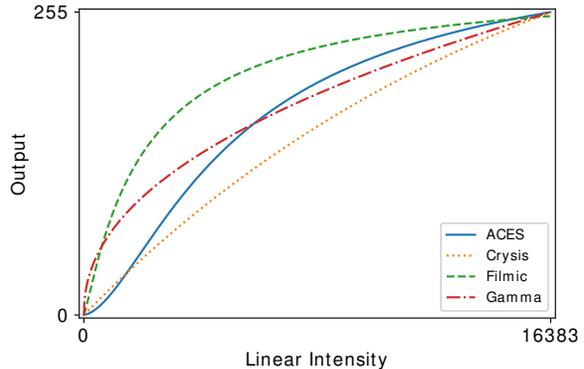


Figure 2. Some popular tone mapping curves used in games and industries [37, 5]. Although the tone mapping function itself is lossless, the following quantization causes a great loss of information in over-exposed and under-exposed pixels. For instance, in a 14-bit linear RAW image, the pixel intensity lies in [16313, 16383] will all be quantized to the maximum pixel intensity 255 of an 8-bit RGB image.

intensity 255 after normalized to $(0, 255)$. This step may cause a 0.004 RMSE error at this single pixel. Thus, it is challenging for existing works [31, 7, 44] to directly synthesize the 14-bit RAW data from its 8-bit sRGB counterparts, especially at the over-exposed regions. We show the comparison of our recovered RAW with previous works in Figure 5. Our method can preserve much more detail of RAW data, even at high-intensity pixels.

Out-of-range value clipping. Value clipping is a common step to normalize the raw value within a reasonable range, which may happen after color space transformation, demosaicing, denoising, and tone mapping [1, 15, 34, 14]. Most commonly used value clipping operation is like $\min(\max(x, 0), 1)$, which will discard the out-of-range pixels at over- and under-exposed regions. Note that this restricts the image capacity for further adjustment. Moreover, traditional ISPs are manually tuned in isolation by experts, which accumulates the clip error among ISP steps to bring further information lost. Our end-to-end pipeline jointly optimizes all the ISP steps and alleviates the clip error accumulation problem to recover more realistic RAW images.

JPEG compression. Modern digital cameras store RGB images in JPEG format, whose information loss further brings challenges to RAW image reconstruction. JPEG encoding pipeline consists of four main steps: color space transformation, discrete cosine transformation (DCT), quantization, and entropy encoding [33]. In reality, quantization is the only lossy and non-differentiable step in JPEG compression. Note that the JPEG information loss is quite hard to reverse. Thus we take a compromised step by integrating the JPEG compression procedure into our network optimization process to alleviate the information loss.

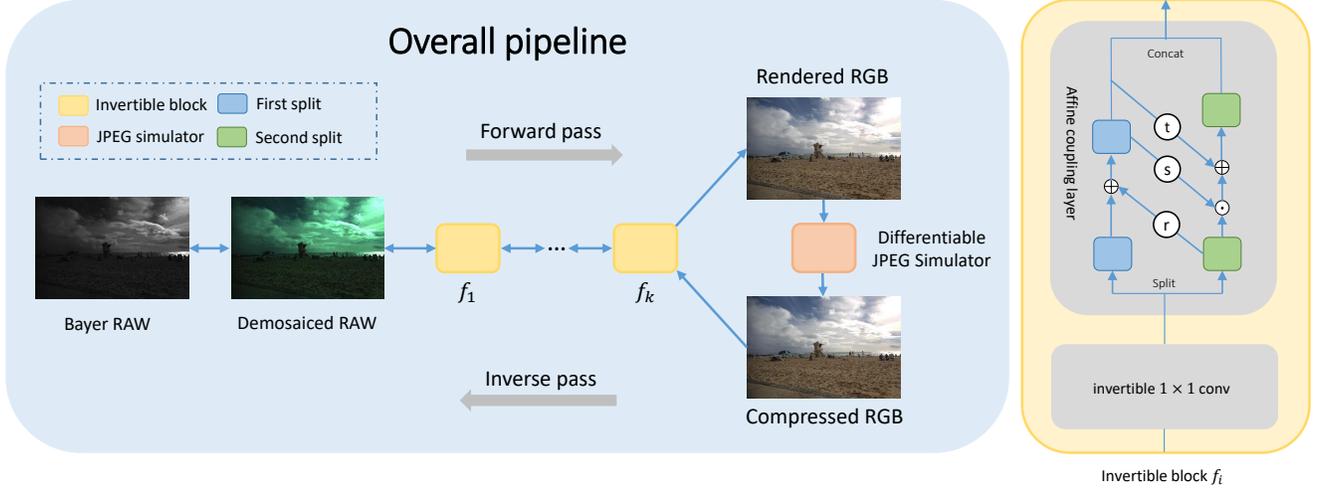


Figure 3. Our Invertible ISP (InvISP) framework. InvISP is composed of both forward and inverse passes. In the forward pass, the Bayer RAW is first bilinearly demosaiced and then transformed to an RGB image by a stack of bijective functions $\{f_i\}_{i=0}^k$. Our model integrates a differentiable JPEG simulator to account for compression information lost. During the training time, to invert the ISP, the backward pass takes a compressed RGB image as input and reverses all the bijective functions and the bilinear demosaicing to obtain the original RAW image. Note that the backward pass takes real JPEG images as input at test time. We illustrate the details of the invertible block on the right. r , s , and t are transformations defined in the bijective functions $\{f_i\}_{i=0}^k$.

To achieve this, we design a differentiable JPEG simulator by carefully simulating the JPEG compression procedure and replacing the quantization step with differentiable Fourier transformations.

4. Method

4.1. Invertible Image Signal Processing (InvISP)

We denote the RAW data space as \mathcal{X} and sRGB data space as \mathcal{Y} . Our goal is to find the invertible and bijective function which can map the data point from RAW data space to sRGB data space, denoted as $f : \mathcal{X} \rightarrow \mathcal{Y}$. To achieve this, classical neural networks need two separate networks to approximate $\mathcal{X} \rightarrow \mathcal{Y}$ and $\mathcal{Y} \rightarrow \mathcal{X}$ mappings respectively, which leads to inaccurate bijective mapping and may accumulate the error of one mapping into the other. We take an alternative method and use the affine coupling layers in [12, 25] to enable invertibility of one single network. We design our invertible ISP with the composition of a stack of invertible and tractable bijective functions $\{f_i\}_{i=0}^k$, i.e. $f = f_0 \circ f_1 \circ f_2 \circ \dots \circ f_k$. For a given observed data sample \mathbf{x} , we can derive the transformation to target data sample \mathbf{y} through

$$\mathbf{y} = f_0 \circ f_1 \circ f_2 \circ \dots \circ f_k(\mathbf{x}), \quad (1)$$

$$\mathbf{x} = f_k^{-1} \circ f_{k-1}^{-1} \circ \dots \circ f_0^{-1}(\mathbf{y}). \quad (2)$$

The bijective model f_i is implemented through affine coupling layers. In each affine coupling layer, given a D dimensional input \mathbf{m} and $d < D$, the output \mathbf{n} is calculated

as

$$\mathbf{n}_{1:d} = \mathbf{m}_{1:d}, \quad (3)$$

$$\mathbf{n}_{d+1:D} = \mathbf{m}_{d+1:D} \odot \exp(s(\mathbf{m}_{1:d})) + t(\mathbf{m}_{1:d}), \quad (4)$$

where s and t represent scale and translation functions from $R^d \mapsto R^{D-d}$, and \odot is the Hadamard product. Note that the scale and translation functions are not necessarily invertible, and thus we realize them by neural networks.

As stated in [12], the coupling layer leaves some input channels unchanged, which greatly restricts the representation learning power of this architecture. To alleviate this problem, we firstly enhance [42] the coupling layer (3) by

$$\mathbf{n}_{1:d} = \mathbf{m}_{1:d} + r(\mathbf{m}_{d+1:D}), \quad (5)$$

where r can be arbitrary function from $R^{D-d} \mapsto R^d$. The inverse step is easily obtained by

$$\mathbf{m}_{d+1:D} = (\mathbf{n}_{d+1:D} - t(\mathbf{n}_{1:d})) \odot \exp(-s(\mathbf{n}_{1:d})), \quad (6)$$

$$\mathbf{m}_{1:d} = \mathbf{n}_{1:d} - r(\mathbf{m}_{d+1:D}). \quad (7)$$

Next, we utilize the invertible 1×1 convolution proposed in [25] as the learnable permutation function to reverse the order of channels for the next affine coupling layer.

We remove the spatial checkerboard mask as it brings no evident performance improvement [25]. We follow the implementation of [10] and disable batch normalization [23] and weight normalization used in [12]. For our image-to-image translation task, we directly learn the RAW-to-RGB mapping without explicitly modeling the latent distribution to stabilize the training process.

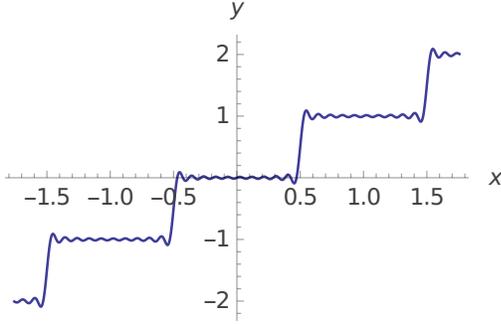


Figure 4. The curve of our approximation rounding function for quantization in our differentiable JPEG simulator.

Note that the input size of invertible neural networks must be identical to the output size. Thus, we take the bilinear demosaiced RAW data as input, which will not destroy the RAW data quality, and reversing the bilinear demosaicing is trivial [7]. For the affine coupling layer, we split the input into two parts. We note that although three-channel input cannot be split evenly, the invertible 1×1 convolution ensures that unchanged components are updated in the next invertible block. Thus R, G, and B channels are still treated equally. We also do an online gamma correction (*i.e.* without storing on disk) to RAW data to compress the dynamic range for faster convergence speed.

The forward pass of our InvISP produces the sRGB images, and the reverse pass aims at recovering realistic RAW data. We conduct bi-directional training with L_1 loss to optimize our framework.

$$L = \|f(\mathbf{x}) - \mathbf{y}\|_1 + \lambda \|f^{-1}(\mathbf{y}) - \mathbf{x}\|_1, \quad (8)$$

where λ is the hyper-parameter used to balance the accuracy between RGB and RAW reconstruction. We set λ to 1 in our main experiments.

4.2. Differentiable JPEG Simulator

Our goal is to train a robust invertible ISP that can tolerate the distortion by JPEG compression to recover accurate RAW. However, the JPEG compression algorithm is not differentiable, which can not be directly integrated into our end-to-end framework. Thus, we propose a differentiable JPEG simulator to enable our network robust to the JPEG compression through the optimization process. Since entropy encoding is lossless and goes after quantization, we skip this step and only simulate color space transformation, DCT, and quantization steps.

To simulate the DCT process, we compute the DCT coefficients and split the input into 8×8 blocks. Then each block is multiplied by DCT coefficients to get the DCT map. In JPEG compression, the DCT map is divided by quantization tables and rounding to the integer type. Since the

rounding function is not differentiable, we design a differentiable rounding function base on the Fourier series, which can be defined as:

$$Q(I) = I - \frac{1}{\pi} \sum_{k=1}^K \frac{(-1)^{k+1}}{k} \sin(2\pi kI), \quad (9)$$

where I is the input map after divided by quantization tables in JPEG compression, and K is used for the tradeoff between approximation accuracy and computation efficiency. As K increases, the simulation function is closer to the real round function, but the running time will also increase. We empirically set K to 10. The rounding process is illustrated in Figure 4.

In the decoding phase of JPEG compression, I is multiplied by the quantization table. The inverse DCT and color space transformation are then applied to reconstruct the simulated JPEG images.

Discussion. Differentiable rounding function is widely used in network quantization research. To fairly prove the effectiveness of our proposed rounding function, we also compare with the rounding function in [17], as shown in Table 1. Our method can achieve a better balance between RGB rendering and RAW reconstruction.

5. Experiments

5.1. Experimental setup

Datasets. We collect the Canon EOS 5D subset (777 image pairs) and the Nikon D700 subset (590 image paris) from the MIT-Adobe FiveK dataset [8] as the training and test data. We train our model for each camera separately. We randomly split each of the two sets (Canon, Nikon) into training and test sets with a ratio of 85:15. We use the LibRaw library to process the RAW images to render ground-truth sRGB images. In general, LibRaw conducts most representative ISP steps in modern digital cameras to render sRGB images, including color space conversion, demosaicing, denoising, white balancing, exposure compensation, gamma compression, and global tone mapping.

Implementation details. We utilize random crop, random rotation, and random flip as data augmentation to train our model. We preprocess the raw data using the white balance parameters provided by camera metadata since estimating white balance directly from raw images is a research topic in itself [3]. To test the effectiveness of our JPEG simulator, we store the ground truth RGB into JPEG format, whose quality is set to 90 (Q=90, most representative JPEG quality in modern digital cameras). We also conduct experiments without preprocessing white balance and with another JPEG quality, whose quantitative results are accessible in the supplement. At test time, we first conduct the forward pass of our network to render RGB images and save

Method	NIKON D700			Canon EOS 5D		
	RGB		RAW	RGB		RAW
	PSNR	SSIM	PSNR	PSNR	SSIM	PSNR
UPI [7]	-	-	30.12	-	-	26.31
CycleISP [44]	-	-	30.19	-	-	34.48
InvGrayscale [41]	24.13	0.8258	33.28	28.22	0.8714	38.00
U-net [10]	36.48	0.9342	41.17	33.44	0.8893	41.14
Ours (w/o JPEG simulation)	37.44	0.9309	44.19	33.45	0.8923	45.73
Ours (JPEG with DSQ [17])	37.44	0.9467	45.25	33.15	0.8946	48.22
Ours (JPEG with Fourier)	37.47	0.9473	45.23	33.61	0.9007	48.57

Table 1. Quantitative evaluation among our model and baselines. Various perceptual metrics show that our proposed ISP model outperforms all the baselines. Our method with JPEG simulation using proposed Fourier quantization outperforms the other two alternative models.

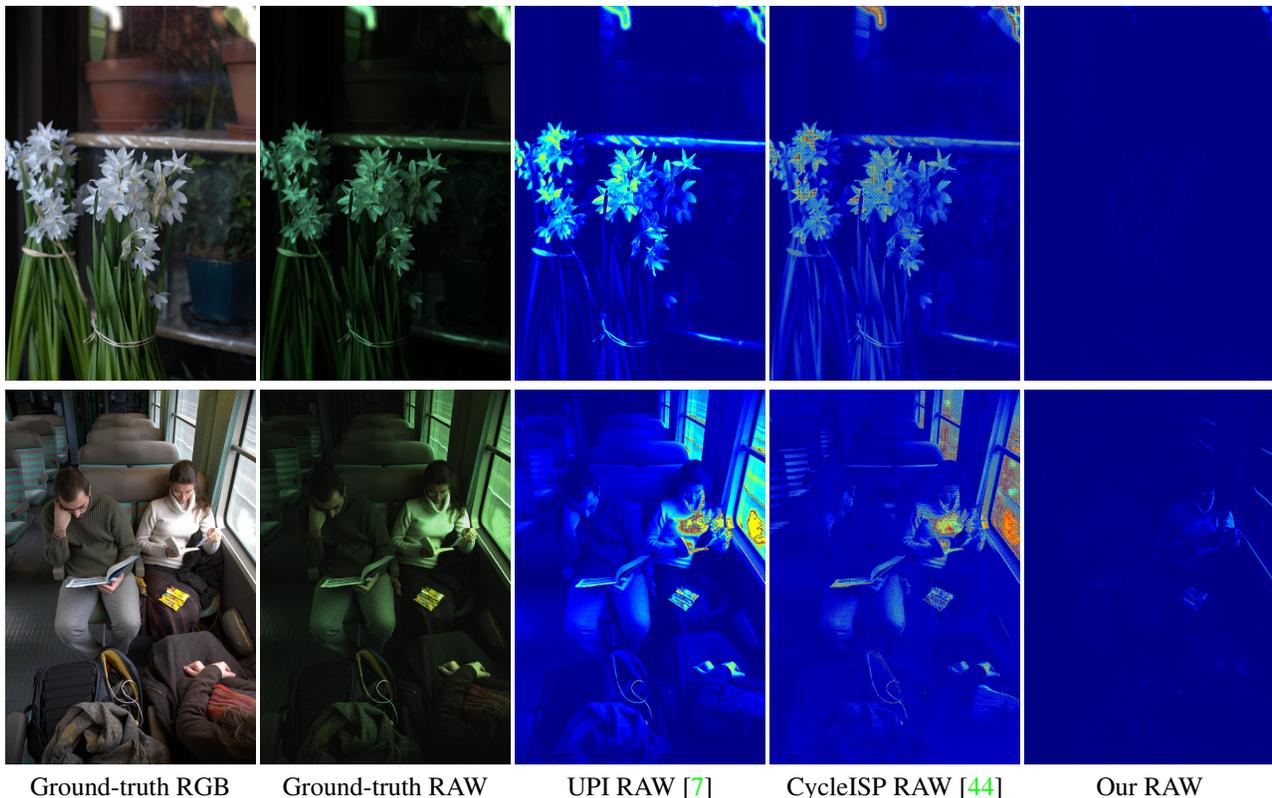


Figure 5. The qualitative comparison among UPI [7], CycleISP [44] and our method. UPI and CycleISP synthesize RAW data from 8-bit compressed RGB, which is inevitable to suffer from the information loss of traditional ISP. Unlike theirs, our model forms a RAW-RGB-RAW cycle and is inherently reversible to recover the realistic RAW image. The GT RAW image is visualized through bilinear demosaicing, and other RAW images are visualized through error maps. This figure is best viewed in the electronic version.

them into JPEG images. Then we load the saved JPEG images and conduct the inverse step to recover RAW images.

5.2. Baselines

UPI. Brooks et al. [7] unprocess the sRGB images to synthesize high-quality RAW images for learned RAW de-

noising. They adopt camera priors to inverse ISP step-by-step, such as digital gain, tone mapping curves, white balance, and color correction matrices [7]. Since the metadata such as color correction matrix, white balance, and digital gain are camera dependent, we modified these parameters in their method to fit our dataset. We use their described

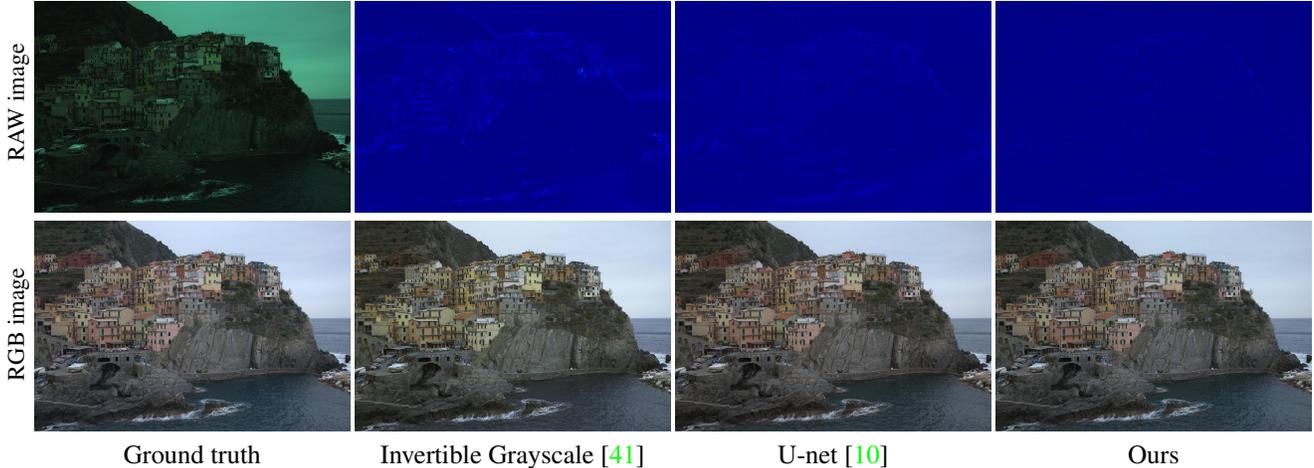


Figure 6. Comparison with baselines. Invertible Grayscale [41] fails at learning a good balance between RGB rendering and RAW recovering, which results in relatively poor performance in both RGB and RAW images. The U-net [10] can render comparable RGB performance with ours but perform worse at RAW recovering. Our invertible ISP can both render visually pleasing RGB images and reconstruct realistic RAW data. The GT RAW is visualized through bilinear demosaicing, and other RAW images are visualized through error maps. This figure is best viewed in the electronic version.

method to estimate metadata for our datasets.

CycleISP. We select the state-of-the-art learning-based RAW synthesizing method, CycleISP [44], as another baseline for synthetic RAW direction. Note that their model has access to RGB images at test time, and thus we only need to compare with their synthesized RAW images. We directly utilize their pretrained model since their framework is trained on the MIT-Adobe FiveK dataset, and their proposed color attention unit can be generalized to different cameras.

U-net. U-net is a representative architecture for ISP in recent year publications [10, 45], thus we implement an encoder-decoder baseline using U-net [38]. Both the encoder and decoder consist of an independent U-net. Same as [10, 45], we pack the Bayer pattern RAW into R-G-G-B channels for encoder input and utilize the depth-to-space operation to restore the RGB resolution. We utilize the same data augmentation strategies as our InvISP. We jointly train the encoder and decoder of our U-net baseline using L1 loss from scratch on all our datasets.

Invertible Grayscale. Invertible Grayscale [41] is a general framework to learn the forward and inverse mapping between two space, such as color-image space and grayscale-image space. The encoder of Invertible Grayscale takes a 3-channel RGB image as input and processes it to a single-channel grayscale image. The decoder recovers the original sRGB image with the same color from the grayscale image. Similar to their settings, we change the input from sRGB image to RAW data after bilinear demosaicing and set the output of the encoder to the 3-channel RGB image. Since the lightness loss function is not suitable for our tasks, we remove it for our experiments.

5.3. Results

Quantitative results. To quantitatively evaluate our method, we use PSNR and SSIM for rendered RGB images, and PSNR for recovered RAW images. The comparison with baselines is reported in Table 1. Compared with the RAW synthesizing method UPI and CycleISP, our model can recover more accurate RAW data, which is proved by more than 13 dB improvement of PSNR. The results are not surprising because lots of information lost in the ISP is quite hard to invert, which results in poor performance for synthetic RAW reconstruction methods. However, our InvISP can jointly optimize RGB rendering and RAW recovering process and thus is better to handle the information lost in quantization, JPEG compression, and saturated value clipping problem in ISP. For the Invertible Grayscale and the U-net baselines, the results indicate that our method contributes a better ISP as well as a stronger model for recovering RAW data. This is because using two separate networks for ISP and inverse ISP will cause the error accumulation problem, which further degrades the RAW reconstruction performance. Our methods take the inherent reversibility of invertible neural networks thus can recover higher-quality RAW images than baselines.

Qualitative results. We show qualitative comparisons against baseline methods in Figure 5 and Figure 6. In Figure 5, the synthetic RAW by CycleISP and UPI differs a lot from ground truth RAW images, especially at over-exposed regions, which indicates that their model performs poorly to handle the information loss of ISP. Our model, however, can recover the RAW information much better than synthetic RAW methods, even at challenging highlight pixels, which

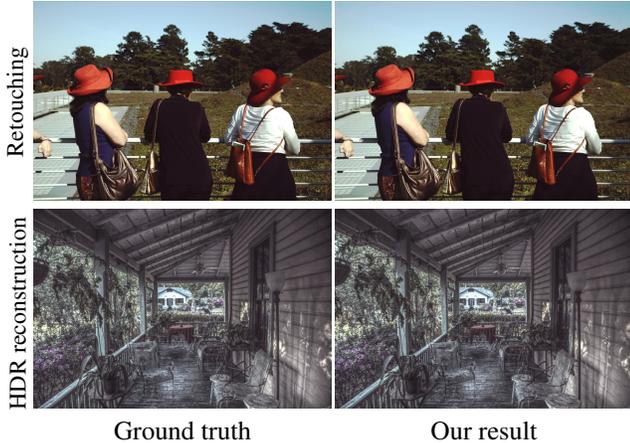


Figure 7. Our model can enable image retouching [22] and HDR reconstruction [32] applications. Note that we use the pretrained model of [22] for retouching.

raises the potential for prospective photo editing tasks. In Figure 6, Invertible Grayscale fails to pursue a good balance between RGB rendering and RAW reconstruction. Our naive U-net baseline can achieve comparable performance in terms of RGB rendering but not perform well at RAW recovering. Our method reconstructs higher-quality RAW images on edges and over-exposed areas without sacrificing the RGB rendering performance.

6. Applications

6.1. RAW data compression

One important application of our framework is RAW data compression for cameras. Traditionally, users need to explicitly store RAW data for further applications. Using our technique, however, only JPEG images need to be stored, and users can reconstruct the corresponding RAW data from JPEG images. To evaluate the reduced file size, we calculate the compression ratio and the bit per pixel (BPP). The compression ratio C_{ratio} [35] is calculated by

$$C_{ratio} = \frac{\text{Uncompressed size}}{\text{Compressed size}} = \frac{B_{BMP}}{B_{JPEG}}, \quad (10)$$

where B_{BMP} is the file size of RAW data in BMP format and B_{JPEG} is the file size of rendered sRGB image in JPEG format. Note that B_{BMP} is calculated by [6]

$$B_{BMP} = 54 + \frac{H \times W \times b}{8}, \quad (11)$$

where H , W and b are the height, width and the bit depth of the RAW data. We further compare our compression effectiveness with Adobe lossy DNG. As shown in Table 2, the file size is highly reduced, even compared with lossy DNG.

Dataset	Compression ratio \uparrow		BPP \downarrow	
	Lossy DNG	Ours	Lossy DNG	Ours
NIKON D700	1.61	34.98	8.73	0.4655
Canon EOS 5D	1.52	27.37	6.56	0.5237

Table 2. Comparison of the compression ratio of the file size and bit per pixel (BPP) between our method and lossy DNG. The file size is significantly reduced by our framework.

6.2. Image retouching

Professional photographers choose to retouch images from RAW data for better visual quality. We demonstrate that our recovered RAW can be directly taken as input for high-quality image retouching. We use an automatic deep learning based image retouching method Exposure [22] as an example. We preprocess the recovered RAW and ground truth RAW through demosaicing and white balancing, following the setting of the paper [22]. We directly utilize their pretrained model that is also trained on the MIT-Adobe FiveK dataset. As illustrated in Figure 7, our reconstructed RAW data has an indistinguishable visual quality to the RAW data captured by the camera.

6.3. HDR reconstruction and tone manipulation

Inferring a high dynamic range image from a single low dynamic range input is challenging [13, 14] since the information lost in saturated and under-exposed regions are hard to invert accurately. Our invertible ISP framework fundamentally alleviates these difficulties and thus enables single image HDR reconstruction. Further, the recovered HDR image can be tone mapped to display much more details than the original RGB input. In Figure 7, we use [32] as tone mapper to demonstrate the potential of our method.

7. Conclusion

We have proposed an end-to-end invertible image signal processing (InvISP) framework to generate visually pleasing RGB images and recover nearly perfect quality RAW data. We leverage the idea from invertible neural networks to design our invertible structure and integrate a differentiable JPEG simulator to enhance the network stability to JPEG compression. We use LibRaw to simulate ground-truth ISP on the MIT-Adobe FiveK dataset. We evaluate our method through comparisons with other frameworks and RAW data synthesis methods. We also demonstrate that our framework enables RAW data compression, image retouching, and HDR reconstruction tasks. We hope our method can inspire further research on RAW image reconstruction.

References

- [1] Abdelrahman Abdelhamed, Stephen Lin, and Michael S. Brown. A high-quality denoising dataset for smartphone cameras. In *CVPR*, 2018. 3
- [2] Mahmoud Afifi, Abdelrahman Abdelhamed, Abdullah Abuolaim, Abhijith Punnappurath, and Michael S Brown. Cie xyz net: Unprocessing images for low-level computer vision tasks. *arXiv:2006.12709*, 2020. 1, 2
- [3] Jonathan T Barron. Convolutional color constancy. In *ICCV*, 2015. 5
- [4] Sean Bell, Kavita Bala, and Noah Snavely. Intrinsic images in the wild. *ACM Transactions on Graphics (TOG)*, 33(4):1–12, 2014. 1
- [5] R. Boitard, D. Thoreau, R. Cozot, and K. Bouatouch. Impact of temporal coherence-based tone mapping on video compression. In *EUSIPCO*, pages 1–5, 2013. 3
- [6] Paul Bourke. Bmp image format. *BMP Files*. July, 1998. 8
- [7] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T Barron. Unprocessing images for learned raw denoising. In *CVPR*, 2019. 1, 2, 3, 5, 6
- [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. 5
- [9] Chen Chen, Qifeng Chen, Minh N. Do, and Vladlen Koltun. Seeing motion in the dark. In *ICCV*, 2019. 2
- [10] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *CVPR*, 2018. 2, 4, 6, 7
- [11] Laurent Dinh, David Krueger, and Yoshua Bengio. NICE: non-linear independent components estimation. In *ICLR*, 2015. 3
- [12] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *ICLR*, 2017. 1, 3, 4
- [13] Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafal K Mantiuk, and Jonas Unger. Hdr image reconstruction from a single exposure using deep cnns. *ACM transactions on graphics (TOG)*, 36(6):1–15, 2017. 8
- [14] Yuki Endo, Yoshihiro Kanamori, and Jun Mitani. Deep reverse tone mapping. *ACM Transactions on Graphics (TOG)*, 36(6):177–1, 2017. 3, 8
- [15] Alessandro Foi. Clipped noisy images: Heteroskedastic modeling and practical denoising. *Signal Processing*, 89(12):2609–2629, 2009. 3
- [16] Michaël Gharbi, Gaurav Chaurasia, Sylvain Paris, and Frédo Durand. Deep joint demosaicking and denoising. *ACM Transactions on Graphics (TOG)*, 35(6):1–12, 2016. 1, 2
- [17] Ruihao Gong, Xianglong Liu, Shenghu Jiang, Tianxiang Li, Peng Hu, Jiazhen Lin, Fengwei Yu, and Junjie Yan. Differentiable soft quantization: Bridging full-precision and low-bit neural networks. In *ICCV*, 2019. 5, 6
- [18] Gabriele Guarnieri, Stefano Marsi, and Giovanni Ramponi. High dynamic range image display with halo and clipping prevention. *IEEE Transactions on Image Processing*, 20(5):1351–1362, 2010. 1
- [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 (TOG)*, 35(6):1–12, 2016. 2
- [20] Felix Heide, Markus Steinberger, Yun-Ta Tsai, Mushfiqur 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. 2
- [21] Jonathan Ho, Xi Chen, Aravind Srinivas, Yan Duan, and Pieter Abbeel. Flow++: Improving flow-based generative models with variational dequantization and architecture design. In *ICML*, 2019. 3
- [22] Yuanming Hu, Hao He, Chenxi Xu, Baoyuan Wang, and Stephen Lin. Exposure: A white-box photo post-processing framework. *ACM Transactions on Graphics (TOG)*, 37(2):26, 2018. 2, 8
- [23] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 4
- [24] Hakki Can Karaimer and Michael S Brown. A software platform for manipulating the camera imaging pipeline. In *ECCV*, 2016. 3
- [25] Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *NeurIPS*, 2018. 1, 3, 4
- [26] Teresa Klatzer, Kerstin Hammernik, Patrick Knobelreiter, and Thomas Pock. Learning joint demosaicing and denoising based on sequential energy minimization. In *ICCP*, 2016. 2
- [27] Chenyang Lei, Xuhua Huang, Mengdi Zhang, Qiong Yan, Wenxiu Sun, and Qifeng Chen. Polarized reflection removal with perfect alignment in the wild. In *CVPR*, 2020. 1
- [28] Zhetong Liang, Jianrui Cai, Zisheng Cao, and Lei Zhang. Cameranet: A two-stage framework for effective camera ISP learning. *ACM Transactions on Graphics (TOG)*, 30:2248–2262, 2021. 2, 3
- [29] Yu-Lun Liu, Wei-Sheng Lai, Yu-Sheng Chen, Yi-Lung Kao, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Single-image HDR reconstruction by learning to reverse the camera pipeline. In *CVPR*, 2020. 1, 2
- [30] Ben Mildenhall, Jonathan T Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. Burst denoising with kernel prediction networks. In *CVPR*, 2019. 1, 2
- [31] Rang MH Nguyen and Michael S Brown. Raw image reconstruction using a self-contained srgb-jpeg image with small memory overhead. *International journal of computer vision*, 126(6):637–650, 2018. 1, 2, 3
- [32] Sylvain Paris, Samuel W Hasinoff, and Jan Kautz. Local laplacian filters: Edge-aware image processing with a laplacian pyramid. *ACM Transactions on Graphics (TOG)*, 30(4):68, 2011. 2, 8
- [33] William B Pennebaker and Joan L Mitchell. *JPEG: Still image data compression standard*. Springer Science & Business Media, 1992. 3
- [34] Tobias Plotz and Stefan Roth. Benchmarking denoising algorithms with real photographs. In *CVPR*, 2017. 3

- [35] Charles Poynton. *Digital video and HD: Algorithms and Interfaces*. Elsevier, 2012. 8
- [36] A. Punnappurath and M. S. Brown. Learning raw image reconstruction-aware deep image compressors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(4):1013–1019, 2020. 1, 2
- [37] Erik Reinhard, Michael Stark, Peter Shirley, and James Ferwerda. Photographic tone reproduction for digital images. In *SIGGRAPH*, 2002. 3
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 3, 7
- [39] 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
- [40] Boxin Shi, Zhe Wu, Zhipeng Mo, Dinglong Duan, Sai-Kit Yeung, and Ping Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. In *CVPR*, 2016. 1
- [41] Menghan Xia, Xueting Liu, and Tien-Tsin Wong. Invertible grayscale. *ACM Transactions on Graphics (TOG)*, 37(6):1–10, 2018. 6, 7
- [42] Mingqing Xiao, Shuxin Zheng, Chang Liu, Yaolong Wang, Di He, Guolin Ke, Jiang Bian, Zhouchen Lin, and Tie-Yan Liu. Invertible image rescaling. In *ECCV*, 2020. 4
- [43] Xiangyu Xu, Yongrui Ma, and Wenxiu Sun. Towards real scene super-resolution with raw images. In *CVPR*, 2019. 2
- [44] 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 *CVPR*, 2020. 1, 2, 3, 6, 7
- [45] Xuaner Zhang, Qifeng Chen, Ren Ng, and Vladlen Koltun. Zoom to learn, learn to zoom. In *CVPR*, 2019. 1, 2, 7