

ParamISP: Learned Forward and Inverse ISPs using Camera Parameters

Woohyeok Kim^{1*} Geonu Kim^{1*} Junyong Lee^{2†}
 Seungyong Lee¹ Seung-Hwan Baek¹ Sunghyun Cho¹
¹POSTECH ²Samsung AI Center Toronto

Abstract

RAW images are rarely shared mainly due to its excessive data size compared to their sRGB counterparts obtained by camera ISPs. Learning the forward and inverse processes of camera ISPs has been recently demonstrated, enabling physically-meaningful RAW-level image processing on input sRGB images. However, existing learning-based ISP methods fail to handle the large variations in the ISP processes with respect to camera parameters such as ISO and exposure time, and have limitations when used for various applications. In this paper, we propose ParamISP, a learning-based method for forward and inverse conversion between sRGB and RAW images, that adopts a novel neural-network module to utilize camera parameters, which is dubbed as ParamNet. Given the camera parameters provided in the EXIF data, ParamNet converts them into a feature vector to control the ISP networks. Extensive experiments demonstrate that ParamISP achieve superior RAW and sRGB reconstruction results compared to previous methods and it can be effectively used for a variety of applications such as deblurring dataset synthesis, raw deblurring, HDR reconstruction, and camera-to-camera transfer.

1. Introduction

A camera ISP converts a RAW image into a visually pleasing sRGB image, which is typically saved as a JPEG image. A camera ISP performs a series of operations including defective pixel correction, denoising, lens shading correction, white balance, color filter array interpolation, color space conversion, tone reproduction, and non-linear contrast enhancement. Detailed operations of camera ISPs are typically sealed to the public and vary from camera to camera.

Unlike sRGB images, RAW images provide physically-meaningful and interpretable information such as noise distributions as they have the linear relationship between image intensity and radiant energy incident on a camera

sensor. Such properties of RAW images have been exploited for denoising [1, 17, 23], HDR [5, 14], and super-resolution [24, 32], leading to superior quality than using only sRGB images. However, a RAW image demands large memory to store due to the use of high-precision bits without any compression. Therefore, only sRGB images are often shared without their RAW counterparts, making it difficult to utilize the useful properties of RAW images.

Recently, several approaches have been proposed to reconstruct RAW images from sRGB images and vice versa by modeling forward and inverse ISPs [3, 4, 7, 29, 30]. However, despite notable improvements, they still suffer from limitations. First, real-world ISPs adjust their operations based on the camera parameters, e.g., exposure time and sensor sensitivity, as shown in Fig. 1. However, previous methods overlook this adaptive nature of real-world ISPs, and learn average ISP operations, which leads to low reconstruction performance. Second, previous methods adopt rather simple network architectures disregarding the complexity of the ISP operations, which leads to low RAW reconstruction quality. They construct a single network by stacking invertible or residual blocks [29, 30], or organize modules by simply arranging convolution layers [3].

In this paper, we present a novel forward and inverse ISP framework, *ParamISP*. ParamISP is designed to faithfully reflect the real-world ISP operations that change based on camera parameters. To this end, ParamISP leverages camera parameters provided in the EXIF metadata of a JPEG image. To effectively incorporate the camera parameters, ParamISP consists of a pair of forward and inverse ISP networks that include a novel neural network module *ParamNet*. ParamNet extracts a feature vector from the camera parameters including exposure time, sensitivity, aperture size, and focal length, and feeds it to the forward and inverse ISP subnetworks to control their behaviors.

To learn ISP operations for varying camera parameters, we need to address the following problems. First, the camera parameters have significantly different scales, such as an exposure time of 0.01 sec. and a sensitivity of 800. Moreover, some parameters such as exposure time and sensor sensitivity have non-linearly increasing parameter val-

* Equal contribution.

† Work done prior to joining Samsung.

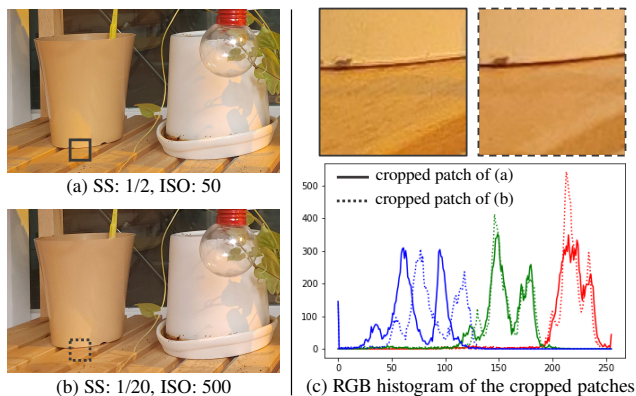


Figure 1. Impact of camera parameters on a camera ISP. Images (a) and (b) are taken by a Samsung Galaxy S22 with different camera parameters. SS and ISO indicate the shutter speed and sensor sensitivity, respectively. In (c), we visualize RGB histograms of each cropped patch of (a) and (b). Despite the same target scene, the captured images exhibit distinct histograms, implying complex ISP operations dependent on the camera parameters.

ues, e.g., the pre-set exposure time of commodity cameras roughly increases exponentially (1/250 sec., 1/125 sec., 1/60 sec., ...). Second, existing datasets lack of diversity in camera parameters [8, 21, 23], while a significant effort is required to collect a sufficient number of images for possible combinations of different camera parameters. Such scale difference and insufficient amount of training data make it difficult to reliably learn ISP operations for different camera parameters. Therefore, we propose a non-linear equalization scheme that adjusts the scales of the camera parameters and a random-dropout-based learning strategy to effectively learn the effects of all the camera parameters.

Finally, we present novel network architectures that reflect the real-world ISP operations. Specifically, our ISP networks consist of four subnetworks: *CanoNet*, *LocalNet* and *GlobalNet*, along with *ParamNet*. *CanoNet* performs common ISP operations such as demosaicing, white balance, and color space conversion using fixed operations without learnable weights. *GlobalNet* performs global tone manipulation with respect to the camera parameters. *LocalNet* performs other residual operations that are not captured by *CanoNet* and *GlobalNet*.

By incorporating camera parameters, novel network architectures, and training schemes, *ParamISP* achieves accurate reconstruction performance with a smaller network size. Our extensive evaluation shows that *ParamISP* surpasses previous learning-based ISP models by an average of 1.93 dB and 1.84 dB in RAW and sRGB reconstruction, respectively (Sec. 4). Furthermore, *ParamISP* is robustly applicable to various applications (e.g., deblurring dataset synthesis, RAW deblurring, HDR reconstruction, camera-to-camera transfer) (Sec. 5).

Our contributions are summarized as follows:

- We propose *ParamISP*, a novel learning-based forward and inverse ISP framework that leverages camera parameters (e.g., exposure time, sensitivity, aperture size, and focal length) for high-quality sRGB/RAW reconstruction.
- To effectively incorporate camera parameters, we develop *ParamNet*, a novel neural network module that controls the forward and inverse ISP networks according to camera parameters. We also introduce a non-linear equalization scheme and a random-dropout-based learning strategy for stable and effective learning of ISP operations with respect to camera parameters.
- We present novel network architectures for the forward and inverse ISP networks that better reflect real-world ISPs. (*CanoNet*, *LocalNet*, *GlobalNet*, *ParamNet*)
- We demonstrate the performance of *ParamISP* and its internal modules on many cameras, and show its applicability to various applications.

2. Related Work

Parametric ISPs To approximate the mapping between RAW and sRGB images, ISP approaches based on non-learnable parametric operations [4, 7] have been proposed. They employ simple invertible ISPs that are composed of non-learnable operations using either camera parameters [4] (e.g., white balance and color correction matrices) or parameters learned from DNNs [7]. However, they do not consider complex nonlinear ISP operations (e.g., local tone mapping and denoising), resulting in inaccurate reconstruction of sRGB and RAW images.

Learnable ISPs For more accurate approximation to forward and inverse ISP operations, DNN-based approaches [3, 29, 30] have been proposed, for which symmetrical forward and inverse ISP networks [3, 30] and an invertible ISP network [29] are employed to learn mapping between sRGB and RAW images. However, they are designed without considering camera parameters (e.g., exposure time, sensor sensitivity, etc), which limit their RAW and sRGB reconstruction quality.

Moreover, some of the previous methods primarily focus on cyclic reconstruction (sRGB-to-RAW-to-sRGB) and try to minimize the difference between the input sRGB image and the sRGB image restored back by the inverse and forward ISPs. Specifically, *CycleISP* [30] uses the input sRGB image for restoring the tone when reconstructing an sRGB image back from a RAW image, which makes it unsuitable for applications that manipulate the tone in the RAW space. *InvISP* [29] employs a single normalizing flow-based invertible network for both inverse and forward processes. While this design choice leads to near perfect reconstruction quality for cyclic reconstruction, its quality significantly degrades when the intermediate RAW images are altered, making it unsuitable for applications that manipulate im-

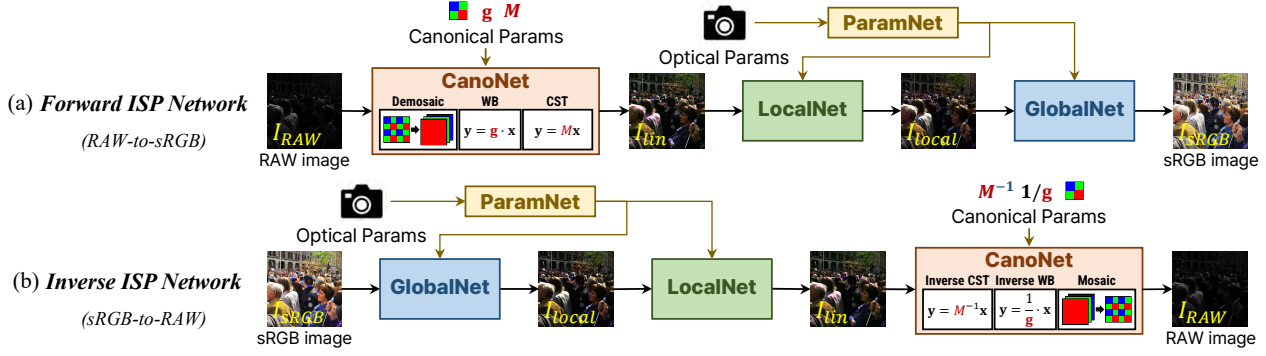


Figure 2. Overview of the proposed ParamISP framework. The full pipeline is constructed by combining learnable networks (ParamNet, LocalNet, GlobalNet) with invertible canonical camera operations (CanoNet). CanoNet consists of differentiable operations without learnable weights, where WB and CST denote white balance and color space transform, respectively.

ages in the RAW space as demonstrated in Sec. 5. In contrast, we design our forward and inverse ISP networks to operate independently during inference time to cater to a broader range of applications.

RAW Reconstruction using Additional Information To achieve precise RAW reconstruction from an sRGB image, another line of research that utilizes additional information has been introduced [11, 18, 20, 26]. This approach encodes necessary metadata such as a small portion of a RAW image within an sRGB image at capture time to reconstruct the original RAW image with high accuracy. However, it necessitates a modification to the existing ISP process to store the metadata. In contrast, our approach utilizes the EXIF metadata that commodity cameras already provide.

3. ParamISP

Given a target camera, our goal is to learn its forward and inverse ISP processes that change with respect to camera parameters. To accomplish this, ParamISP is designed to have a pair of forward (RAW-to-sRGB) and inverse (sRGB-to-RAW) ISP networks (Fig. 2). Both networks are equipped with ParamNet so that they adaptively operate based on camera parameters.

In ParamISP, we classify camera parameters into two distinct categories: optical parameters (including exposure time, sensitivity, aperture size, and focal length) and canonical parameters (Bayer pattern, white balance coefficients, and a color correction matrix). The canonical parameters directly influence fundamental ISP operations like demosaicing, white balance adjustments, and color space conversion. These operations are relatively straightforward, and their relationship with the canonical parameters is well-defined.

While previous approaches have leveraged the canonical parameters [4, 7], the optical parameters have remained untapped. In contrast, ParamISP exploits both sets of parameters to achieve highly accurate sRGB and RAW reconstruction. To harness the canonical parameters, our ISP networks

incorporate CanoNet, a subnetwork that performs canonical ISP operations without learnable weights. For the optical parameters, we introduce ParamNet, which is the key component to dynamically control the behavior of the ISP networks based on the optical parameters.

In the following, we describe ParamNet and our training strategy for stable and effective training. We then explain the forward and inverse ISP subnetworks in detail.

3.1. ParamNet

As described in Fig. 3, ParamNet takes optical parameters as input and converts them into an optical parameter feature vector z , which is then fed to both LocalNet and GlobalNet. ParamNet consists of a non-linear equalization layer, and fully-connected layers. The non-linear equalization layer computes a normalized feature vector for each optical parameter to compensate for the scale difference between the optical parameters. The equalized feature vectors are then arithmetically summed together, and fed to the fully-connected layers to obtain an optical parameter feature vector z .

Non-linear Equalization As discussed in Sec. 1, the optical parameters exhibit significantly different scales, and their values are non-linearly distributed. While some optical parameters may substantially alter the behavior of the camera ISP, others may have minimal effect. Moreover, the influence of optical parameters may vary across different camera models. As a result, naively incorporating the optical parameters leads to unstable training of ParamNet.

To resolve this issue, the non-linear equalization layer of ParamNet applies various non-linear mapping to each optical parameter and learns the best composition of them that enables stable training. Specifically, we first apply non-linear mapping functions such as x , $1/x$, \sqrt{x} , $x^{-1/2}$, $x^{1/4}$, $x^{-1/4}$, $\log(x)$, $\sin(\log(x))$, $\cos(\log(x))$, $\sin(c \cdot x)$, and $\cos(c \cdot x)$ to each optical parameter, where x is the value of an optical parameter, and c controls the frequency of the sinusoidal functions, for which we use three different val-

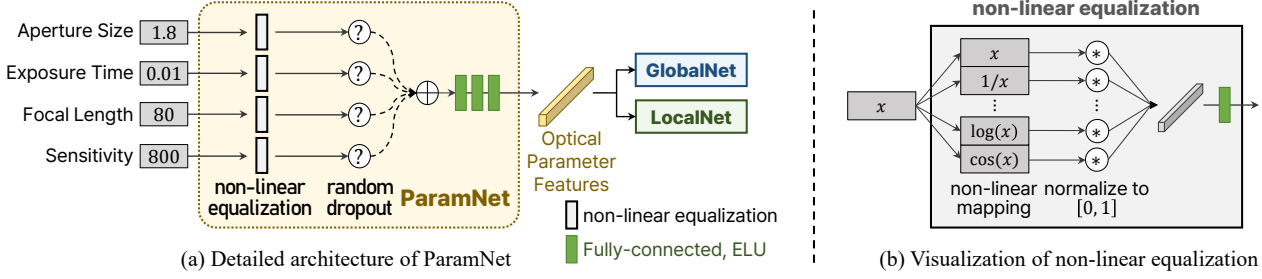


Figure 3. Architecture of ParamNet. (a) Given camera optical parameters, ParamNet estimates optical parameter features used for modulating the LocalNet and GlobalNet. (b) In order to deal with different scales and non-linearly distributed values of optical parameters, we propose to use non-linear equalization that exploits multiple non-linear mapping functions.

ues empirically chosen. We then normalize each of the non-linear mapping results into $[0, 1]$. As a result, we obtain a 15-dim. vector for each optical parameter. Then, each vector is processed through a subsequent fully-connected layer to extract useful information from each parameter vector.

Our non-linear equalization layer incorporates a range of diverse non-linear mapping functions, rather than relying on a predefined set of carefully-chosen functions, in order to cover the wide spectrum of potential relationships between each optical parameter and the camera ISP. Thus, some mapping functions may look redundant or unnecessary. Nevertheless, the subsequent fully-connected layer can successfully learn to extract only useful information from them by combining them with different weights.

Random Dropout Even with the non-linear equalization, ParamNet can still be trained to overfit to a subset of the optical parameters due to the scale difference and the insufficient amount of training data. To mitigate the problem, during training, we randomly drop out the equalized feature vector of each optical parameter. Specifically, each optical parameter is randomly dropped out by the probability of 20% at each training iteration.

3.2. Forward ISP Network

The forward ISP network consists of four subnetworks: CanoNet, LocalNet, GlobalNet, and ParamNet (Fig. 2(a)). CanoNet performs canonical ISP operations: demosaicing [9], white balance, and color space conversion, exploiting the canonical parameters. LocalNet performs local operations of an ISP such as denoising, sharpening and local tone mapping in addition to compensating for the residual errors of the canonical ISP operations of CanoNet. GlobalNet performs global tone manipulation with respect to the camera parameters. ParamNet is connected to LocalNet and GlobalNet and controls their behavior. In the following, we explain each of CanoNet, LocalNet and GlobalNet in detail.

CanoNet CanoNet takes a RAW image $I_{RAW} \in \mathbb{R}^{1 \times H \times W}$ and first performs demosaicing [9] to produce $I_{Dem} \in \mathbb{R}^{3 \times H \times W}$ according to the Bayer pattern of I_{RAW} . The demosaicing algorithm of CanoNet may differ from that

of the target camera ISP. Nonetheless, such discrepancy is compensated by the subsequent LocalNet. It then performs white balance using the coefficients $g_{WB} \in \mathbb{R}^3$. Lastly, it transforms the color space from the RAW space to the linear sRGB space using the color correction matrix $M_{Cam} \in \mathbb{R}^{3 \times 3}$. We denote the resulting image as $I_{lin} \in \mathbb{R}^{3 \times H \times W}$.

LocalNet Fig. 4 shows an overview of LocalNet. LocalNet takes I_{lin} and an optical parameter feature vector z from ParamNet as input, and performs local ISP operations. As well as I_{lin} , LocalNet also takes additional handcrafted features: a gradient map, a soft histogram map, and an over-exposure mask computed from I_{lin} , as they help improve reconstruction quality, as shown in [12–14]. For more details, refer to the supplementary material. Finally, LocalNet predicts a residual image, which is then added to I_{lin} to obtain the output image I_{local} .

As the goal of LocalNet is to learn local ISP operations, LocalNet adopts a network architecture based on the U-Net [22], which has been proven to be highly effective for various image restoration and enhancement tasks [2, 6, 10, 27, 31]. Specifically, the input image and the additional input features computed from the input image are concatenated, and fed to LocalNet. Besides, the optical parameter feature vector z is duplicated horizontally and vertically to build a feature map of the same spatial size as the input image. The feature map is then added to an intermediate feature map of LocalNet. Then, the resulting feature map is processed through multi-scale residual blocks (ResBlocks) with convolutional block attention modules (CBAMs) [28], and converted to a residual image.

GlobalNet GlobalNet globally adjusts tone and color according to the content of its input image and the optical parameter feature vector z . To this end, GlobalNet adopts parametric global operations and predicts their parameters. Specifically, GlobalNet adopts two types of parametric global operations: quadratic transformation f_q and gamma correction f_g . Let us denote the input image to global adjustment operation as I , and its pixel as $p = (p_r, p_g, p_b)$. Then, the quadratic transformation f_q is defined as:

$$f_q(p) = Wp' \quad (1)$$

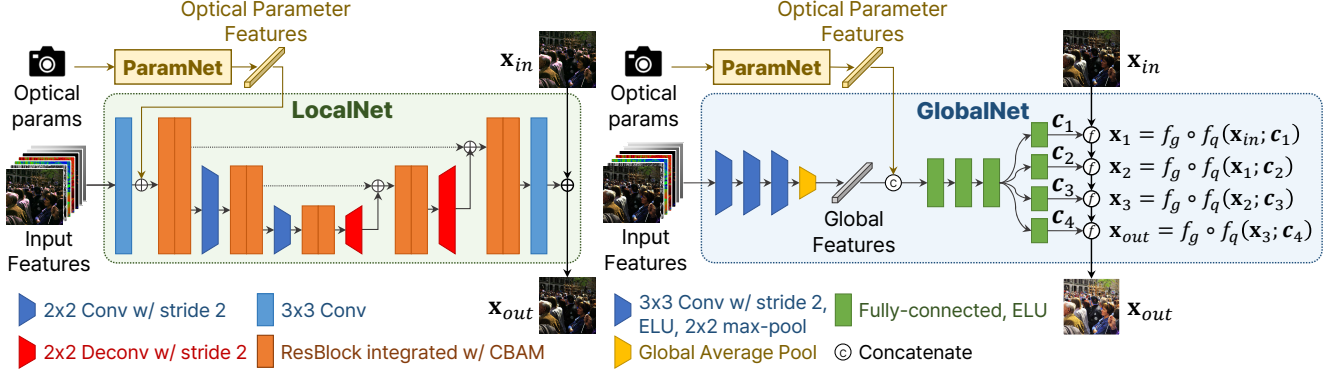


Figure 4. Detailed architecture of LocalNet and GlobalNet. In GlobalNet, f_g and f_q represent the gamma correction and quadratic transformation, respectively, while c_n represents the n -th coefficients G^n and W^n , as explained in Sec. 3.2.

where p' is a 10-dim. vector defined as $p' = [p_r^2, p_g^2, p_b^2, p_r p_g, p_g p_b, p_b p_r, p_r, p_g, p_b, 1]^T$. W is a 3×10 matrix with the coefficients of the quadratic transformation, which is uniformly applied to all the pixels in I . The gamma correction operator f_g is defined as:

$$f_{g,c}(p_c) = \frac{(\alpha_c p_c + \beta_c)^{\gamma_c} - \beta_c^{\gamma_c}}{(\alpha_c + \beta_c)^{\gamma_c} - \beta_c^{\gamma_c}} \quad (2)$$

where the subscript c is an index to each color channel, i.e., $c \in \{r, g, b\}$. γ_c is a gamma parameter, while α_c and β_c are a scale and an offset, respectively. The set of the gamma correction coefficients is denoted by G , i.e., $G = \{\alpha_r, \beta_r, \gamma_r, \alpha_g, \beta_g, \gamma_g, \alpha_b, \beta_b, \gamma_b\}$. Similar to W , G is also uniformly applied to all the pixels in I . For the sake of notational simplicity, we represent the quadratic transformation and gamma correction of an entire image I as $f_q(I)$ and $f_g(I)$, respectively, in the rest of the paper.

To support a wide range of potential non-linear tone adjustments of commodity camera ISPs, GlobalNet models the global tone adjustment as a series of gamma correction and quadratic transformation. Specifically, given an input image I , GlobalNet performs global tone adjustment as:

$$\hat{I} = f_{gq}^N \circ \dots \circ f_{gq}^1(I) \quad (3)$$

where $f_{gq}^n(I)$ is defined as $f_{gq}^n = f_g^n \circ f_q^n(I)$. f_g^n and f_q^n are the gamma correction and quadratic transformation with the n -th coefficients G^n and W^n , respectively. In our experiments, we use $N = 4$ as it leads to the best quality (Refer to Tab. S3 in the supplementary material).

The quadratic transformation is also adopted by previous modular learnable ISPs [3, 23]. However, a single quadratic transformation cannot accurately model diverse non-linear tone adjustments performed by commodity camera ISPs including gamma correction. Therefore, in our approach, we extend the global tone adjustment by adopting the gamma correction and iteratively applying both gamma correction and quadratic transformation.

Fig. 4 shows an overview of GlobalNet. To predict the coefficients $\{\dots, (G^n, W^n), \dots\}$, GlobalNet takes a concatenation of I_{local} and the handcrafted features computed from I_{local} as done in LocalNet as input, and extracts a global feature vector through a series of convolution layers with max pooling, and a global average pooling layer. The global feature vector is then added to the optical parameter feature vector, and processed through fully-connected layers to obtain the coefficients.

3.3. Inverse ISP Network

The inverse ISP network also consists of CanoNet, LocalNet, GlobalNet and ParamNet with an inverse order of the forward ISP network (Fig. 2(b)). Specifically, GlobalNet takes an sRGB image as input and performs inverse tone manipulation. Next, LocalNet executes inverse local operations, and finally, CanoNet handles inverse color space conversion, inverse white balance, and mosaicing. GlobalNet and LocalNet use the same network architectures as those in the forward ISP network. On the other hand, CanoNet consists of the inverse operations of those of the forward ISP network, but in reverse order.

4. Experiments

Implementation We used PyTorch [19] to implement our models. We train our model in two stages: pre-training and fine-tuning. In the pre-training stage, we train our models with multiple datasets captured from multiple cameras. In the fine-tuning stage, we train our models on a specific target camera. Although our models aim at learning the ISP of a single target camera model, we found that pre-training with multiple cameras substantially improves the reconstruction quality. In the pre-training stage, we use the RAISE dataset [8] from Nikon D7000, D90, and D40, the RealBlur dataset [21] from Sony A7R3, and the S7 ISP dataset [23] from Samsung Galaxy S7. More details such as the effect of the pre-training, the statistics of the datasets, and how we split the datasets to training, validation, and test

Baseline (Ours w/o ParamNet)			PSNR↑	Param↓
ParamNet (w/ Opt. Params)	Non-linear Equalization	Random Dropout		
			34.77	0.68M
✓			-	-
✓	✓		35.64	0.71M
✓	✓	✓	36.21	0.71M

Table 1. Ablation study on the components of ParamNet.

Optical Params	w/o <i>A</i>	w/o <i>B</i>	w/o <i>C</i>	w/o <i>D</i>	Full
PSNR↑	35.13	35.25	35.87	35.46	36.21
SSIM↑	0.9678	0.9718	0.9724	0.9705	0.9724

Table 2. Ablation study on the impact of each optical parameter. *A, B, C*, and *D* represent sensor sensitivity, exposure time, aperture size, and focal length, respectively.

sets can be found in the supplementary material.

We train our network in the pre-training stage for 520 epochs with an initial learning rate of 2.0×10^{-4} , and in the fine-tuning stage for 2,140 epochs with an initial learning rate of 2.0×10^{-5} . For both pre-training and fine-tuning, we employ AdamW [15] and the learning rate is step-decayed with a rate of 0.8 every 10 epochs. In each epoch in the pre-training stage, we randomly sample 1024 cropped patches of size 448×448 from the dataset of each camera, resulting in a total of 5120 patches. In the fine-tuning stage, we sample 1024 patches from the dataset of a target camera in each epoch. We evaluated our models and other models on a PC equipped with an NVIDIA RTX A6000. For a fair comparison, we also apply both pre-training and fine-tuning to other models in our experiments.

We train our forward and inverse ISP networks in an end-to-end fashion using losses \mathcal{L}_{for} and \mathcal{L}_{inv} :

$$\mathcal{L}_{for} = \|\hat{I}_{sRGB} - I_{sRGB}\|_1 \quad (4)$$

$$\mathcal{L}_{inv} = \|\hat{I}_{RAW} - I_{RAW}\|_1 \quad (5)$$

respectively, where \hat{I}_{sRGB} , I_{sRGB} , \hat{I}_{RAW} , and I_{RAW} are a reconstructed sRGB image, its ground-truth sRGB image, a reconstructed RAW image, and its ground-truth RAW image, respectively.

4.1. Ablation Study

For all the ablation studies in this section, we train inverse (sRGB-to-RAW) ISP network using the D7000 training images of the RAISE dataset [8]. Additional ablation studies using the forward ISP network and different camera datasets, and on the effect of the input features and training approach can be found in the supplementary material.

ParamNet We first evaluate the effect of the non-linear equalization and random optical parameter dropout of ParamNet. To this end, we compare variants of our model in Tab. 1. The first row in the table corresponds to a baseline model without ParamNet. Introducing ParamNet without non-linear equalization (2nd row) results in the failure of

Baseline (CIE XYZ Net [3])		PSNR↑	SSIM↑	Param↓
LocalNet	GlobalNet			
		30.04	0.9461	1.3M
✓		33.12	0.9644	1.7M
✓	✓	33.66	0.9646	0.6M

Table 3. Ablation study on the effects of LocalNet and GlobalNet.

training with diverging losses. On the other hand, the non-linear equalization (3rd row) not only stabilizes the training but also improves the reconstruction quality compared to the baseline model in the first row, which proves that the non-linear equalization is essential for training, and that exploiting the optical parameters is necessary to improve the reconstruction quality. Finally, random dropout of the optical parameters (4th row) further enhances the reconstruction quality as it prevents overfitting and enhances the generalization ability of ParamNet.

Optical Parameters We now analyze the impact of the optical parameters. To this end, we compare four variants of ParamNet, for each of which, we exclude each one of the four optical parameters. Tab. 2 shows the quantitative ablation results. Compared to the model that uses all the optical parameters, all the other variants exhibit performance degradation, indicating that all the optical parameters need to be considered for high-quality reconstruction. We can also observe that excluding the sensitivity and exposure time shows the largest performance degradation. While how commodity camera ISPs utilize the optical parameters is unknown, these optical parameters may affect the image enhancement operation of camera ISPs, such as denoising, as they are closely related to noise. The impact of the aperture size and focal length is not as significant as the others, but the reconstruction quality still decreases without them. This is because aperture size and focal length are related to defocus blur and lens aberrations, so the camera ISP operations may alter based on them.

ISP Network Architecture We verify the effectiveness of the network architecture of our ISP subnetworks. As our modular network structure is inspired by CIE XYZ Net [3] that consists of a pair of local and global mapping modules, we consider its local and global mapping networks as our baseline for LocalNet and GlobalNet. Specifically, we build a baseline model for the inverse ISP network by replacing the LocalNet and GlobalNet with the local and global mapping networks of CIE XYZ Net, then compare the performance of the baseline and its variants to ours. The local mapping network of CIE XYZ Net simply stacks convolution layers, while the global mapping network models global tone manipulation as a single quadratic function and estimate its parameters using fully-connected layers. Tab. 3 shows the comparison result. The table shows that LocalNet significantly improves the PSNR by more than 3 dB over the local mapping of CIE XYZ Net. Moreover, GlobalNet further improves the PSNR by around 0.5 dB and leads to




Figure 5. sRGB-to-RAW reconstruction. We show error maps between reconstructed and GT RAW images.




Figure 6. RAW-to-sRGB reconstruction. We show error maps between reconstructed and GT sRGB images.

Method	sRGB→RAW measured in PSNR \uparrow					Param. \downarrow
	D7000	D90	D40	S7	A7R3	
UPI [4]	20.67	26.57	22.05	29.98	30.48	-
CIE XYZ Net [3]	30.04	32.62	38.57	33.24	36.42	1.3M
CycleISP [30]	35.52	35.85	42.83	34.55	45.35	3.1M
InvISP [29]	33.48	35.39	45.08	34.29	47.14	1.4M
ParamISP (Ours)	38.49	37.06	45.97	35.20	48.33	0.7M

Table 4. Quantitative comparison on RAW reconstruction.

a significantly reduced model size thanks to our global tone adjustment model and efficient network architecture.

4.2. RAW & sRGB Reconstruction

In this section, we compare the sRGB-to-RAW and RAW-to-sRGB reconstruction ability of ParamISP with existing state-of-the-art methods: UPI [4], CIE XYZ Net [3], CycleISP [30], and InvISP [29]. UPI is a non-learnable parametric method that approximates the general structure of the camera ISP. CIE XYZ Net, CycleISP, and InvISP are learning-based approaches that employ symmetrical forward and inverse ISP networks. For CIE XYZ Net whose output is an image in the CIE XYZ color space, we convert its output to a RAW image using canonical ISP operations as in CanoNet for a fair comparison. For InvISP whose output is white-balanced RAW image, we apply inverse white balancing and mosaicing operations to obtain a RAW image. These learning-based methods train the forward and inverse networks together, but we found that separately training them for each yields better performance. Therefore, we train them separately in the same manner as ours.

Tab. 4 shows quantitative results of the inverse sRGB-to-RAW reconstruction in terms of PSNR evaluated on our test set. UPI [4] results in high reconstruction errors due to its parametric model-based pipeline. CIE XYZ Net [3], CycleISP [30], and InvISP [29] show better reconstruction quality than UPI, but their quality is still limited compared to our method. ParamISP achieves the best reconstruction quality despite its much smaller model size, thanks to leveraging the optical parameters and our carefully-designed net-

Method	RAW→sRGB measured in PSNR \uparrow					Param. \downarrow
	D7000	D90	D40	S7	A7R3	
UPI [4]	18.81	20.30	16.01	20.05	19.37	-
CIE XYZ Net [3]	26.76	27.61	34.84	27.63	37.19	1.3M
InvISP [29]	30.20	28.89	37.86	28.96	43.93	1.4M
ParamISP (Ours)	34.14	30.83	39.54	29.02	45.51	0.7M

Table 5. Quantitative comparison on sRGB reconstruction.

work architecture. Fig. 5 shows the qualitative results of error maps between reconstructed RAW images and ground-truth RAW images. As the figure shows, our method produces much less error than the others.

Tab. 5 reports quantitative results of the forward reconstruction (RAW-to-sRGB). CycleISP is not included in this comparison because it needs an input sRGB image for sRGB reconstruction in the sRGB-to-RAW-to-sRGB cyclic reconstruction. Similar to the inverse sRGB-to-RAW reconstruction, ParamISP clearly outperforms the other methods, showing the effectiveness of our approach.

5. Applications

Forward and inverse ISP models can benefit various applications as shown in [3, 29, 30]. In this section, we demonstrate a couple of applications of ParamISP: RAW deblurring and HDR reconstruction. Other additional applications including deblurring dataset synthesis and camera-to-camera transfer are also provided in the supplementary material. Before applying ParamISP to applications, we further perform joint fine-tuning on separately trained forward and inverse ISP networks. For details on the joint fine-tuning, refer to the supplementary material.

Deblurring in RAW Space While the irradiance of a blurred image has a linear relationship with its latent sharp image, such a linear relationship no longer holds in sRGB images after non-linear operations have been applied by camera ISPs. Such non-linearity varies across different cameras. Consequently, learning sRGB image deblurring requires camera-specific training datasets, which leads to



Figure 7. Qualitative results of RAW deblurring. In (c) and (d), red artifacts are visible in white text, while in (b) and (e), the results are blurrier than ours.

Method	UPI [4]	CIE XYZ Net [3]	InvISP [29]	CycleISP [30]	ParamISP (Ours)
PSNR \uparrow	24.63	28.30	28.82	29.94	30.70
SSIM \uparrow	0.8011	0.8444	0.8721	0.8855	0.8984
Param \downarrow	-	2.7M	1.4M	7.4M	1.4M

Table 6. Quantitative results of RAW deblurring.

an excessive amount of training time and memory space to support diverse camera models. Instead, we can effectively reconstruct a RAW image from an sRGB image using ParamISP, apply a camera-independent deblurring model to the RAW image, and obtain a deblurred sRGB image. While ParamISP also requires camera-wise training, it requires a much smaller number of parameters ($2 \times 0.7M$) compared to deblurring models (Stripformer: 20M [25] and Uformer-B: 51M [27]), and supports a wide range of applications.

Tab. 6 and Fig. 7 show quantitative and qualitative comparisons of deblurring results in the RAW space using different ISP networks. For deblurring in the RAW space, we use Stripformer [25] trained on the RealBlur-R dataset, which is a real-world RAW blurry image dataset [21]. The evaluation is done on the RealBlur-J test set [21]. As the comparisons show, ParamISP clearly outperforms all the other ISP networks in deblurring performance thanks to its high-quality sRGB and RAW reconstruction.

HDR Reconstruction As shown in CIE XYZ Net [3], RAW image reconstruction can also be used for high-dynamic-range (HDR) image reconstruction from a single low-dynamic-range (LDR) image as RAW images provide a wider tonal value range. Specifically, we first convert an LDR sRGB image to an LDR RAW image, and multiply the reconstructed RAW image by synthetic digital gains (0.1, 1.4, 2.7, 4.0) to create four RAW images. Each of these is then passed through the forward ISP to reconstruct sRGB images. Finally, we apply an off-the-shelf exposure-fusion




Figure 8. Qualitative results of HDR reconstruction.

algorithm [16] to obtain a single HDR image.

Fig. 8 shows qualitative results using different models. Previous methods primarily focus on cyclic reconstruction. Specifically, CycleISP uses an input LDR sRGB image for the forward ISP to adjust the tone, while InvISP, composed of a single flow-based network, is not robust when intermediate RAW images are altered. On the other hand, ParamISP successfully produces an HDR reconstruction result compared to other methods. For more qualitative results, refer to the supplementary material.

6. Conclusion

In this paper, we present ParamISP, a novel learning-based forward and inverse ISP framework that leverages camera parameters. To incorporate camera parameters effectively, we introduce ParamNet to control the forward and inverse ISP networks, proposing a stable and effective training strategy with respect to camera parameters. We also present novel network architectures for ISP networks that better reflect real-world ISP operations. Through our extensive experiments, ParamISP shows the state-of-the-art performance in RAW and sRGB reconstruction and the robust applicability to various applications.

Limitations and Future Work While we have quantitatively confirmed that incorporating the aperture size and focal length improves the reconstruction quality in Tab. 2, it is still unclear how such parameters affect the behaviors of the ISP operations. While we use datasets captured by various cameras in the pre-training stage, different ISPs may behave differently with respect to different optical parameters, and such inconsistency may potentially have negative impact on a pre-trained ISP model. A more sophisticated training strategy may resolve such issue.

Acknowledgments This work was supported by the NRF grants (RS-2023-00211658, RS-2023-00280400, 2022R1A6A1A03052954, 2023R1A2C200494611) and IITP grant (2019-0-01906, Artificial Intelligence Graduate School Program (POSTECH)) funded by the Korea government (MSIT). This work was also supported by Samsung Research Funding Center (SRFCIT1801-52) and Samsung Electronics Co.

References

- [1] Abdelrahman Abdelhamed, Stephen Lin, and Michael S Brown. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [2] Abdullah Abuolaim and Michael S Brown. Defocus deblurring using dual-pixel data. In *Proceedings of the European conference on computer vision (ECCV)*, 2020. 4
- [3] Mahmoud Afifi, Abdelrahman Abdelhamed, Abdullah Abuolaim, Abhijith Punnappurath, and Michael S Brown. CIE XYZ Net: Unprocessing images for low-level computer vision tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 44(9):4688–4700, 2021. 1, 2, 5, 6, 7, 8
- [4] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T Barron. Unprocessing images for learned raw denoising. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 3, 7, 8
- [5] 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)*, 2018. 1
- [6] Yu-Sheng Chen, Yu-Ching Wang, Man-Hsin Kao, and Yung-Yu Chuang. Deep photo enhancer: Unpaired learning for image enhancement from photographs with GANs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 4
- [7] Marcos V. Conde, Steven McDonagh, Matteo Maggioni, Ales Leonardis, and Eduardo Pérez-Pellitero. Model-based image signal processors via learnable dictionaries. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 36(1):481–489, 2022. 1, 2, 3
- [8] Duc-Tien Dang-Nguyen, Cecilia Pasquini, Valentina Conotter, and Giulia Boato. RAISE: A raw images dataset for digital image forensics. In *Proceedings of the 6th ACM multimedia systems conference (MMSys)*, 2015. 2, 5, 6
- [9] Pascal Getreuer. Malvar-he-cutler linear image demosaicking. *Image Processing on Line*, 1:83–89, 2011. 4
- [10] Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, and Jiří Matas. Deblurgan: Blind motion deblurring using conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8183–8192, 2018. 4
- [11] Leyi Li, Huijie Qiao, Qi Ye, and Qinmin Yang. Metadata-based raw reconstruction via implicit neural functions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [12] Steve Lin and L. Zhang. Determining the radiometric response function from a single grayscale image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005. 4
- [13] Steve Lin, Jinwei Gu, Shuntaro Yamazaki, and Harry Shum. Radiometric calibration from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004.
- [14] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 4
- [15] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. 6
- [16] Tom Mertens, Jan Kautz, and Frank Van Reeth. Exposure fusion. In *15th Pacific Conference on Computer Graphics and Applications (PG)*, 2007. 8
- [17] Ben Mildenhall, Jonathan T Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. Burst denoising with kernel prediction networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [18] Seonghyeon Nam, Abhijith Punnappurath, Marcus A Brubaker, and Michael S Brown. Learning srgb-to-raw-rgb de-rendering with content-aware metadata. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [19] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch. In *Proceedings of the Neural Information Processing Systems Workshops (NeurIPSW)*, 2017. 5
- [20] Abhijith Punnappurath and Michael S Brown. Spatially aware metadata for raw reconstruction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2021. 3
- [21] Jaesung Rim, Haeyun Lee, Jucheol Won, and Sunghyun Cho. Real-world blur dataset for learning and benchmarking deblurring algorithms. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2, 5, 8
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention (MICCAI)*, 2015. 4
- [23] Eli Schwartz, Raja Giryes, and Alex M Bronstein. DeepISP: Toward learning an end-to-end image processing pipeline. *IEEE Transactions on Image Processing (TIP)*, 28(2):912–923, 2018. 1, 2, 5
- [24] Chengzhou Tang, Yuqiang Yang, Bing Zeng, Ping Tan, and Shuaicheng Liu. Learning to zoom inside camera imaging pipeline. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1
- [25] Fu-Jen Tsai, Yan-Tsung Peng, Yen-Yu Lin, Chung-Chi Tsai, and Chia-Wen Lin. Stripformer: Strip transformer for fast image deblurring. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 8
- [26] Yufei Wang, Yi Yu, Wenhan Yang, Lanqing Guo, Lap-Pui Chau, Alex C Kot, and Bihan Wen. Raw image reconstruction with learned compact metadata. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3

- [27] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 4, 8
- [28] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. CBAM: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, 2018. 4
- [29] Yazhou Xing, Zian Qian, and Qifeng Chen. Invertible image signal processing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 7, 8
- [30] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Cycleisp: Real image restoration via improved data synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2, 7, 8
- [31] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 4
- [32] Xuaner Zhang, Qifeng Chen, Ren Ng, and Vladlen Koltun. Zoom to learn, learn to zoom. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1