





Rawformer: Unpaired Raw-to-Raw Translation for Learnable Camera ISPs

Georgy Perevozchikov¹, Nancy Mehta¹*, Mahmoud Afifi²**, and Radu Timofte¹

¹ Computer Vision Lab, CAIDAS & IFI, University of Würzburg,
John Skilton Str. 4a, 97074 Würzburg, Germany
{georgii.perevozchikov,nancy.mehta,radu.timofte}@uni-wuerzburg.de
² York University, 4700 Keele St, Toronto, Ontario, Canada, M3J 1P3
m.3afifi@gmail.com

Abstract. Modern smartphone camera quality heavily relies on the image signal processor (ISP) to enhance captured raw images, utilizing carefully designed modules to produce final output images encoded in a standard color space (e.g., sRGB). Neural-based end-to-end learnable ISPs offer promising advancements, potentially replacing traditional ISPs with their ability to adapt without requiring extensive tuning for each new camera model, as is often the case for nearly every module in traditional ISPs. However, the key challenge with the recent learning-based ISPs is the urge to collect large paired datasets for each distinct camera model due to the influence of intrinsic camera characteristics on the formation of input raw images. This paper tackles this challenge by introducing a novel method for unpaired learning of raw-to-raw translation across diverse cameras. Specifically, we propose Rawformer, an unsupervised Transformer-based encoder-decoder method for raw-to-raw translation. It accurately maps raw images captured by a certain camera to the target camera, facilitating the generalization of learnable ISPs to new unseen cameras. Our method demonstrates superior performance on real camera datasets, achieving higher accuracy compared to previous state-of-the-art techniques, and preserving a more robust correlation between the original and translated raw images. The codes and the pretrained models are available at <https://github.com/gosha20777/rawformer>.

1 Introduction

Unlike the high-quality digital single-lens reflex (DSLR) cameras, mobile phone cameras possess inherent limitations owing to their smaller sensor size and fixed lenses with limited optical capabilities [55]. The mobile camera image signal processor (ISP) aims to overcome these limitations by applying a series of carefully designed modules (e.g., [1, 7, 11, 18, 25, 47]) to enhance the quality of the final image rendered in a standard color space, such as sRGB. These modules

* Corresponding Author

** Now at Google

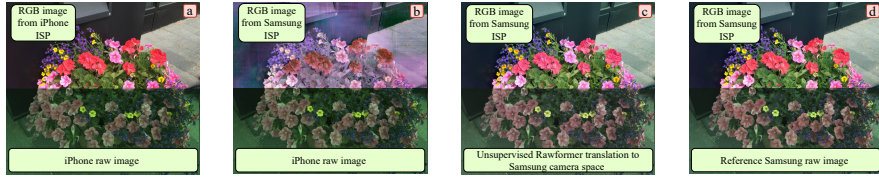


Fig. 1: We introduce Rawformer, an unsupervised method for raw-to-raw translation that allows the utilization of pre-trained neural-based ISPs to process raw images captured by previously unseen cameras. Shown are raw and sRGB images processed by a neural-based ISP [54]. a) Raw image from an iPhone X rendered by the neural-based ISP trained on iPhone X’s raw images. b) Raw image from an iPhone X rendered by the neural-based ISP trained on Samsung S9’s raw images. c) iPhone X raw image translated to Samsung S9’s raw space using our method, then processed by the Samsung S9 neural-based ISP. d) Raw image captured by Samsung S9’s camera rendered by its native camera ISP, provided as a reference for visual comparison.

work in tandem to optimize key parameters, such as sharpness, color accuracy, and noise reduction, resulting in a visually pleasing sRGB image from the raw image captured by the camera [13, 16].

Despite yielding promising results, achieving compatibility among camera ISP modules for each new camera device requires extensive tuning and adjustment efforts to attain the desired image quality. This motivates the use of deep neural networks to replace individual camera ISP modules with a single neural-based unit capable of accomplishing the task [12, 19, 44]. Nonetheless, the training of such neural-based ISPs demands an extensive dataset consisting of raw images taken with a specific camera, accompanied by corresponding “ground-truth” sRGB images (typically generated by a high-quality DSLR camera) [19, 22].

An issue emerges with the introduction of a new camera model, as pre-trained neural-based ISPs might encounter difficulty in accurately rendering raw images captured by the new camera (refer to Fig. 1(b)). This difficulty arises from the potential distinct characteristics of the new camera, such as sensor sensitivities, which influence the formation of its raw images [4, 5]. Consequently, inconsistencies arise in the interpretation of raw RGB colors between cameras [4], hindering the generalization of neural-based ISPs to new cameras not encountered during training. To address this issue, re-training or fine-tuning on a new paired dataset becomes necessary, in which the raw images are sourced from the target new camera. This process mirrors similar challenges encountered in traditional ISP development for new cameras.

An alternative strategy to address this challenge is domain adaptation, where the raw images from the new target camera are mapped to the raw space of the training camera to emulate those captured by the original camera used for training. In light of this, [4] proposed a semi-supervised learning framework for raw image translation, but the approach outlined in [4] inherits the drawbacks associated with the CNN-based architecture, notably the limited receptive field and inefficient encoding of global information. Thus, we present a fully unsupervised

Transformer-based raw-to-raw method, dubbed Rawformer, that maps raw images between camera models by efficiently encoding their global and semantic correlations. Our method eliminates the need for a paired raw-sRGB dataset from each new target camera, enabling the utilization of a neural-based ISP trained on a specific camera’s raw images to process raw images taken by new cameras with different characteristics without necessitating any re-training; see Fig. 1(c). Our contributions are summarized as follows:

- We propose Rawformer, a *fully unsupervised* encoder-decoder Transformer-based method designed for raw-to-raw translation facilitating the reuse of neural-based ISPs. Rawformer achieves state-of-the-art results in raw-to-raw translation, leading to better generalization for cross-camera learnable ISP rendering.
- We introduce contextual-scale aware downsampler and upsampler blocks that efficiently summarize the local-global contextual details in mixed scale representations via its condensed query attention block and scale perceptive feed-forward network.
- A novel cross-domain attention-driven discriminator is proposed along with a specialized discriminator head for stabilizing the network training.

2 Related Work

2.1 Neural-Based ISP

Several neural-based ISPs methods have been introduced [8, 17, 19, 24, 56, 59, 60], with the majority relying on CNN-based U-Net-like architecture [42]. These include earlier approaches, such as DeepISP [44] and PyNET [22], as well as more recent methods, such as MW-ISPNet [21], AW-Net [12], and LAN [54], which achieve improved image restoration quality by integrating discrete wavelet transform (DWT) and double attention modules (DAM) techniques. Such learnable ISPs are typically trained to map raw images captured by a *specific camera* model to target high-quality sRGB “ground-truth” images produced by some target camera ISP (usually with higher quality than the mobile native ISP; e.g., a DSLR camera [19, 22]). A significant challenge for neural-based ISPs is the need for paired datasets of raw-sRGB images. Collecting such datasets is tedious and typically involves using specific equipment. Additionally, it may require post-processing techniques to align the raw images with the target sRGB images [22]. The challenge intensifies when a mobile phone manufacturer introduces a new camera model with different characteristics, which happens frequently [49], requiring the recollection of paired raw-sRGB datasets for each new camera model.

One solution is to render images to the sRGB space from the device-independent CIE XYZ space [3]. However, this requires calibration for each camera model, and accuracy depends on specific camera conditions (e.g., the Luther condition) [15]. In contrast, our work introduces an unsupervised raw-to-raw mapping technique that eliminates the need for paired raw-sRGB datasets for re-training or camera calibration. Our method maps raw images from new cameras to the camera raw

space utilized during the training of neural-based ISPs, enabling the reuse of pre-trained ISPs.

2.2 Image-to-Image Domain Adaptation

Numerous generative adversarial network (GAN)-based methods address unpaired image-to-image (I2I) translation challenges [38]. Notably, CycleGAN [62], DualGAN [57], UNIT [32], STARGAN [10], SEAN [63], and U-GAT-IT [28] focus on cycle consistency, employing dual generator networks guided by a cycle-consistency constraint. ACLGAN [61] and CUT [39] represent advancements in this field. ACLGAN relaxes the cycle consistency constraint, favoring an adversarial constraint, potentially enhancing translation quality. CUT utilizes a contrastive loss, eliminating the need for multiple generators and enabling faster training. ITTR further refines CUT’s performance by altering the generator architecture. CycleGAN-based model, such as UVCGAN [50] incorporate U-Net-like generators with vision Transformers (ViTs).

In computational photography, attempts have been made for domain adaptation, including synthesizing raw night images from raw daylight images to enhance neural-ISP rendering quality [41], few-shot domain adaptation for low-light image enhancement [40], and mapping graphics images to a target camera’s raw space [45]. Raw-to-raw translation has received even limited attention. Classical approaches require a calibration object to learn global mapping [37], while recent work [4] proposes a semi-supervised CNN learning-based method using a small set of perfectly aligned images. Despite being effective at exploiting the local semantics, this CNN-based method lacks the potential to exploit the long-range/global semantics between different domains. In contrast, our method aims at solving the raw mapping by exploiting both local and global contextual information through *fully* unsupervised training, achieving state-of-the-art results.

3 Proposed Method

Existing raw-to-raw methodologies either rely on supervised approaches requiring paired datasets e.g., [37]) or semi-supervised approaches (e.g., [4]) with limited accuracy. Additionally, domain adaptation techniques (e.g., [39, 62]), which aim to extract local semantics between source and target domains, often fail to capture overall semantics and context-dependent relationships. Our proposed method, Rawformer, addresses these limitations by processing raw images in a *fully* unsupervised manner. Utilizing an encoder-decoder architecture within a CycleGAN framework, Rawformer effectively captures global semantics and maintains intrinsic consistency between translated images. The generator (Fig. 2(a)) in Rawformer employs contextual-scale aware downsampler and upsampler blocks to enhance content rendering and structural integrity. The discriminator (Fig. 3(a)) stabilizes training by incorporating key discriminator features, thereby preventing model collapse and reducing redundant computations. In the following, we provide a comprehensive elucidation of the proposed components.

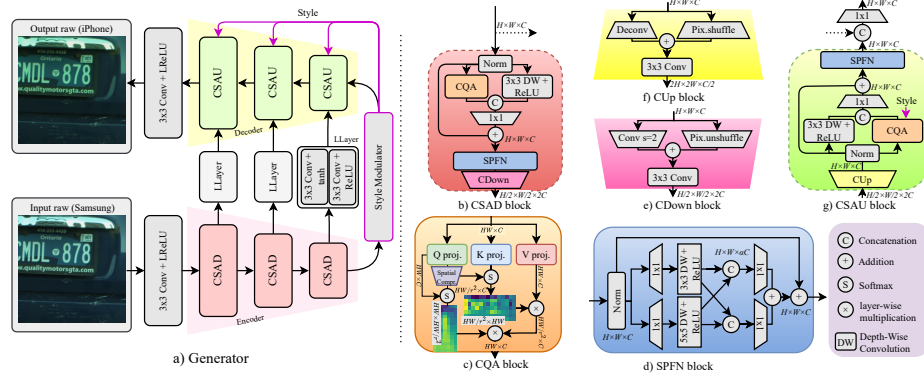


Fig. 2: a) Overview of the generator architecture (a) of Rawformer. The primary components of the generator: b) contextual-scale aware downsampler block (CSAD), c) condensed query attention block (CQA), d) scale perceptive feed-forward network (SPFN), e) composite downsampling block (CDown), f) composite upsampler (CUp) block, and g) contextual-scale aware upsampler block (CSAU).

3.1 Generator’s Encoder Network

As shown in Fig. 2(a), the encoder of the proposed generator in Rawformer comprises stacked contextual-scale aware downsampler (CSAD) blocks to bridge the semantic gap between the source and target domains. Each CSAD block, depicted in Fig. 2(b), offers a holistic approach of understanding the contextual dependencies efficiently between the domains by capturing both the global and local information via its condensed query attention (CQA) block, in parallel with a 3×3 depth-wise convolution layer. It further handles the intrinsic correlations among the extracted features in a multi-scale manner via maintaining a synergistic collaboration between the scale perceptive feed-forward network (SPFN), and the composite downsampler (CDown) block.

Condensed Query Attention Block: Self-attention is highly effective in modeling long-range contextual dependencies by comparing each token with all others, facilitating seamless information propagation within the network [58]. However, its computational demands pose challenges for tasks like unpaired raw-to-raw translation, where maintaining both visual quality and semantic consistency is crucial. To address these concerns and overcome the constraints of computation, we introduce the condensed query attention (CQA) block; see Fig. 2(c). Here, we make two key modifications to self-attention to enhance learning efficiency. Firstly, we spatially compress the query vector using average pooling followed by linear projection, and downsample the resulting feature map by a factor of r to confine the crucial information. This results in condensed queries, (\mathbf{Q}_c) used for similarity comparison, as shown in Eq. 1. Secondly, to balance global and regional modeling with computational efficiency, attention is applied hor-

horizontally and vertically to eliminate redundant computations. The CQA block operation is defined as:

$$\begin{aligned} CQA(\mathbf{Q}_c, \mathbf{K}, \mathbf{V}) &= \mathbf{A}_U \cdot (\mathbf{A}_H \cdot \mathbf{V}); \quad \mathbf{Q}_c = W_p(AvgPool(\mathbf{Q})) \\ \mathbf{A}_H &= Softmax(\frac{\mathbf{Q}_c \mathbf{K}^T}{\sqrt{d}}), \quad \mathbf{A}_U = Softmax(\frac{\mathbf{Q} \mathbf{Q}_c^T}{\sqrt{d}}) \end{aligned} \quad (1)$$

where, the duet of key (\mathbf{K}), and value (\mathbf{V}) vectors are obtained through plain linear convolutions, and the dimensions of each of these features are shown in the Fig. 2(c). $W_p(\cdot)$ denotes a 1×1 pointwise convolution, and \mathbf{A}_H , and \mathbf{A}_U represents the compressed attention map between the key-condensed query pair horizontally and query-condensed query pair (for focusing on fine grained contextual information) vertically, respectively. *The choice of operations and the reason for condensing the queries is investigated in the supplementary of the paper.*

Scale Perceptive Feed-forward Network: To proficiently capture the intricacies of the local image structure across various scales by encoding the information from neighboring pixel locations, we propose a scale perceptive feed-forward network (SPFN); see Fig. 2(d). Unlike the conventional feed-forward networks that integrate single-scale depth-wise convolutions [58], we emphasize on enhancing the locality and inter-correlations among the incoming raw features of different domains by inserting multi-scale depth-wise convolutions. Given an input tensor at layer s and scale i , such that $\mathbf{X}_s^i \in \mathbb{R}^{H^i \times W^i \times C}$, SPFN is formulated as:

$$\begin{aligned} \mathbf{X}_s^1 &= \phi(W_d^3(W_p^1(LN(\mathbf{X}_{s-1})))); \mathbf{X}_s^2 = \phi(W_d^5(W_p^2(LN(\mathbf{X}_{s-1})))), \\ \mathbf{Y}_s &= W_p^3(\mathbf{X}_s^1 \textcircled{\text{C}} \mathbf{X}_s^2) + \mathbf{X}_{s-1}. \end{aligned} \quad (2)$$

Here, ϕ represents a LeakyReLU activation function [35], W_d^k represents a depth-wise convolution with filter size k , LN denotes the layer normalization [6], and $\textcircled{\text{C}}$ represents the concatenation operation. This enriches a more nuanced understanding of the underlying image features at different hierarchical levels.

Composite Downsampler: In general, the essence of downsampling is the retrieval of the multi-scale information in the image. Unlike the common encoder architectures for various image signal processing pipelines [31, 48] that deploy strided convolutions for downsampling, we design a hybrid strategy of leveraging the benefits of information exchange between different resolution features. To adaptively recover the underlying content information from the input domain at multiple scales, we design a composite down (CDown) block that intelligently combines pixel-unshuffle (to prevent the information loss) and strided convolution (to enhance the model’s expressiveness) as shown in Fig. 2(e).

Style Modulator: Different cameras have their own distinctive characteristics to be handled or restored. To augment the expressiveness of Rawformer for

handling camera perturbations, we design a style modulator in the generator using extended pixel-wise ViT [51]. It aids in calibrating the decoder features and subsequently encourages the recovery of crucial details while translation. As demonstrated in Fig. 2(a), the style modulator applies style tokens in each contextual aware upsampler (CSAU) block, where every style token models the weight of the query vector of the CQA block to enrich the style specific context. This addition of style token at each stage in the decoder helps in the flexible adjustment of the feature maps. Unlike StyleGANv2 [27], where the style vectors are generated from the random prior, our approach resorts to learnable style token for conjecturing the important target style directly from the source domain, thus enhancing the consistency and control in raw image translation. *More details on style modulation are provided in the supplementary.*

3.2 Generator’s Decoder Network

Unlike the popular additive or multiplicative skip connections between the encoder-decoder, we also present a learnable linking layer (LLayer) as shown in the Fig. 2(a) to facilitate the feature preservation in the contextual scale aware upsampler (CSAU) block. The overall design of the CSAU block closely resembles that of the CSAD block, ensuring consistency and coherence in the network architecture in a bottom-up fashion. The primary distinction lies within the composite upsampling block, where for inducing the relevant spatial information into the target domain for the incoming features from the style modulator, we design a composite upsampler block (CUp). It embodies a parallel combination of pixel-shuffle (for preventing the checkerboard artifacts in the target domain) and deconvolution (for combating with aliasing) as shown in Fig. 2(f) and (g).

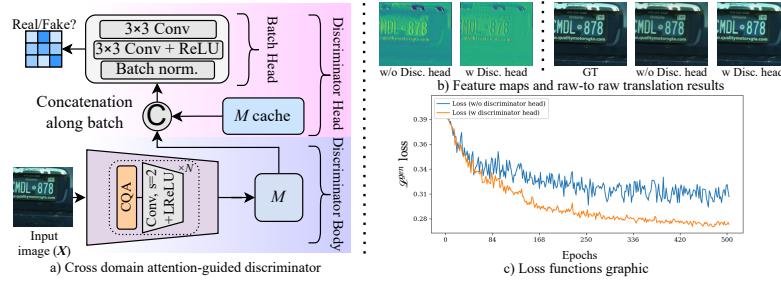


Fig. 3: a) The proposed cross-domain attention-guided discriminator. b) Feature map visualizations without (w/o) and with (w) the discriminator head. The inclusion of the discriminator head aids in refining the overall results, and the discriminator training (as shown in c) benefits from the inclusion of the discriminator head, particularly with the cache (M) component.

3.3 Cross-Domain Attention-Guided Discriminator

Leveraging batch statistics is an effective strategy to mitigate mode collapse and enhance diversity in GAN models [43]. Traditional minibatch discrimination techniques require large batch sizes [26], posing challenges due to GPU limitations. Inspired by [50], we address this by decoupling batch size from the minibatch discrimination approach. Our core innovation involves using a cache of previous discriminator features as an alternative to larger batch sizes. During training, batch features are stored in four separate caches (M); real images from the source and target cameras, and fake images from both domains. This memory bank-style cache utilizes a composite discriminator with a head and body. The head captures batch statistics via batch normalization and two 3×3 convolutional layers, processing a concatenated input of current and historical discriminator outputs from the cache, as shown in Fig. 3(a). The discriminator body includes a CQA block followed by strided convolution to handle detailed global features across domains. This integration of batch statistics, caching, and attention mechanisms significantly enhances the discriminator’s ability to manage multi-resolution image data and stabilize the overall training (Figs. 3 b and c).

3.4 Self-Supervised Pre-Training

Self-supervised pre-training optimizes the initial weights of a network by engaging it in auxiliary tasks formulated from the inherent structure of the original data. In general, despite modifications in the cross-domain attention-guided discriminator, GAN training still suffers from mode collapse, where the generator repeatedly tends to produce a limited set of images. Thus, we opt to pre-train the generator of our Rawformer using self-supervised inpainting tasks to enhance initial weight optimization and mitigate this mode collapse problem. *More details about the mode of pretraining is provided in the supplementary (Section 1) of the paper.*

4 Experiments

We evaluated our method using two types of datasets : 1) raw-to-raw datasets, which provide unpaired raw images for training and aligned raw images from source and target cameras for assessing unsupervised raw-to-raw translation accuracy (Sec. 4.3), and 2) raw-to-sRGB datasets, offering aligned raw and corresponding sRGB images for evaluating transformed raw image quality followed by learnable neural-based ISP rendering (Sec. 4.4). This evaluation mirrors real-world usage scenarios where raw images from a new camera are translated to the raw space of the camera used for ISP training. In all the experiments, raw images were demosaiced using the Menon algorithm [36] and normalized after black-level subtraction. During our method’s training, each raw image was segmented into non-overlapping crops of 224×224 pixels. Similarly, we used paired raw-sRGB non-overlapping crops of 224×224 pixels to train the camera-specific LAN ISP [54] on each raw-to-sRGB dataset.

4.1 Datasets

We utilized two raw-to-raw datasets: 1) the Raw-to-Raw dataset [4] and 2) the NUS dataset [9], and three raw-to-sRGB datasets: 1) the Zurich raw-to-RGB dataset [20], 2) the Samsung S7 dataset [44], and 3) the Mobile AIM21 dataset [19]. The Raw-to-Raw dataset [4] includes raw images from the Samsung Galaxy S9 and the iPhone X, featuring both unpaired and paired sets. The NUS dataset [9] comprises scenes captured by eight DSLR cameras. Following the evaluation in [4], we utilized the Nikon D5200 and Canon EOS 600D DSLR cameras from the NUS dataset [9]. Aligned paired raw images were generated using the same approach as in the Raw-to-Raw dataset [4]. In both datasets, paired raw images are dedicated for testing, while unpaired images are utilized for unsupervised training.

The Zurich raw-to-RGB dataset [20] includes pairs of raw and sRGB images captured synchronously using a Huawei P20 smartphone camera and a Canon 5D Mark IV DSLR camera. The Samsung S7 dataset [44] comprises paired raw and sRGB images captured by the Samsung S7 smartphone’s main camera. The Mobile AIM21 dataset [19] includes images captured using a Sony IMX586 mobile sensor and a Fujifilm GFX100 DSLR, employing an advanced dense correspondence algorithm for matching due to imperfect alignment [52].

4.2 Training details

Inspired by the effective fine-tuning capabilities demonstrated by the pre-trained BERT model for the text completion task [14], we initially trained each generator network in a self-supervised fashion for image completion. Initially, the generator was trained on 32×32 pixel patches extracted from training images, with 40% of these patches being randomly masked out. The objective of the generator during this stage was to reconstruct the original unmasked image from its partially obscured version. Training was performed using the AdamW optimizer [34] with betas set to (0.9, 0.99) and a weight decay of 0.05 for 500 epochs to minimize the following pixel-wise loss function:

$$\mathcal{L}_{pixel-wise} = L_1 + (1 - SSIM) + VGG \quad (3)$$

where L_1 , $SSIM$, VGG refer to L1 loss, Structural Similarity Index (SSIM) [53], and VGG perceptual loss [30] between original unmasked image and the reconstructed image. The learning rate was initialized at 0.005, and the learning schedule was managed using the cosine annealing warm restarts strategy [33]. This self-supervised pre-training was shown to improve the raw mapping results (see Table 1).

Thereafter, the weights of both generators were fine-tuned for our unpaired raw-to-raw translation task for an additional 500 epochs, while simultaneously training the discriminator networks. We use the Adam optimizer [29] with betas set to (0.5, 0.99) and a learning rate of 0.0001 for the discriminator networks and 0.00005 for the generator networks. The discriminator networks were optimized to minimize the following loss functions:

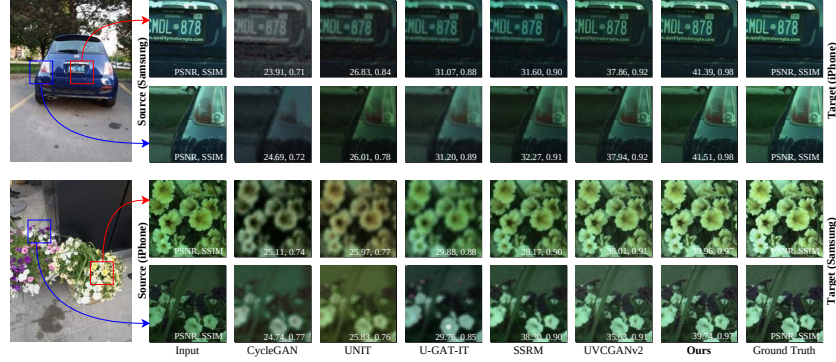


Fig. 4: Shown are two images captured by Samsung S9 and iPhone X in sRGB (left) and two cropped patches from each image in raw (right). On the right, we show the input raw patch from the corresponding camera and the corresponding ground-truth raw patch from the other camera, along with the results by other methods. Our proposed Rawformer is better at preserving the domain consistent features.

$$\mathcal{L}_A^{dis} = \ell_{gan}(D_A(G_{B \rightarrow A}(b)), 0) + \ell_{gan}(D_A(a), 1) \quad (4)$$

$$\mathcal{L}_B^{dis} = \ell_{gan}(D_B(G_{A \rightarrow B}(a)), 0) + \ell_{gan}(D_B(b), 1) \quad (5)$$

where a and b represent images from camera A and B , respectively, while $G_{B \rightarrow A}(b)$ and $G_{A \rightarrow B}(a)$ are the translated images produced by each generator network, respectively. The labels 0 and 1 represent “fake” and “real” raw images, respectively, and $\ell_{gan}(\cdot)$ computes the cross-entropy loss.

During this stage, the weights of both generators were optimized to minimize the following loss function:

$$\mathcal{L}^{gen} = \beta_1(\mathcal{L}_A^{gan} + \mathcal{L}_B^{gan}) + \beta_2(\mathcal{L}_A^{idt} + \mathcal{L}_B^{idt}) + \beta_3(\mathcal{L}_A^{cyc} + \mathcal{L}_B^{cyc}) \quad (6)$$

where $\beta_1 = 1$, $\beta_2 = 10$, $\beta_3 = 0.5$, $\mathcal{L}_A^{gan} = \ell_{gan}(D_B(G_{A \rightarrow B}(a)), 1)$, $\mathcal{L}_A^{idt} = \mathcal{L}_{pixel-wise}(G_{B \rightarrow A}(a), a)$, and $\mathcal{L}_A^{cyc} = \mathcal{L}_{pixel-wise}(G_{B \rightarrow A}(G_{A \rightarrow B}(a)), a)$. Note that at inference time, only a single generator network (e.g., $G_{A \rightarrow B}$) is required to translate raw images captured by the source camera (e.g., A) to the target camera (e.g., B).

4.3 Raw-to-Raw Translation Results

We compare our raw-to-raw translation results with those of the recent semi-supervised raw-to-raw mapping (SSRM) method [4]. Additionally, we compare our method against several established unsupervised generic I2I methodologies, including UVCANv2 [51], U-GAT-IT [28], UNIT [32], and the classic CycleGAN [62]. All other methods were trained on the same training data used for

Table 1: Raw-to-raw translation results on NUS dataset [9]. The mapping results of Nikon D5200’s raw images to Canon EOS 600D DSLR camera and vice versa using our method and other methods are shown. Here, $\ell_a = L_1$, $\ell_b = L_1 + VGG$, $\ell_c = \mathcal{L}_{pixel-wise}$ (Eq. 3), w, and w/o denote the results with and without training on a corresponding configuration, respectively. Additionally, we show the GFLOPs of each method for the 256×256 input image. *Best results are in bold.*

| Methods | Canon-to-Nikon | | | | Nikon-to-Canon | | | | FLOPs (G) |
|--|----------------|-------------|------------------------|--------------|----------------|-------------|------------------------|--------------|-----------|
| | PSNR↑ | SSIM↑ | MAE↓ ($\times 1e-2$) | ΔE ↓ | PSNR↑ | SSIM↑ | MAE↓ ($\times 1e-2$) | ΔE ↓ | |
| CycleGAN [62] | 27.32 | 0.83 | 3.21 | 15.40 | 26.81 | 0.82 | 3.33 | 11.49 | 36.8 |
| UNIT [32] | 25.55 | 0.81 | 3.17 | 18.22 | 24.73 | 0.79 | 3.23 | 18.64 | 47.3 |
| U-GAT-IT [28] | 29.41 | 0.87 | 3.02 | 7.12 | 27.86 | 0.85 | 3.29 | 8.01 | 62.9 |
| SSRM [4] | 32.36 | 0.93 | 2.41 | 6.21 | 30.81 | 0.93 | 2.47 | 5.95 | 37.1 |
| UVCGANv2 [51] | 37.11 | 0.96 | 2.32 | 4.34 | 37.29 | 0.96 | 2.28 | 4.28 | 69.6 |
| Ours (w/o self-supervised pre-training; ℓ_a loss) | 37.13 | 0.96 | 2.29 | 4.65 | 37.27 | 0.96 | 2.31 | 4.46 | 52.7 |
| Ours (w/o self-supervised pre-training; ℓ_b loss) | 37.94 | 0.97 | 2.21 | 3.96 | 37.99 | 0.96 | 2.23 | 3.93 | 52.7 |
| Ours (w/o self-supervised pre-training; ℓ_c loss) | 38.73 | 0.97 | 2.19 | 3.64 | 38.26 | 0.97 | 2.22 | 3.72 | 52.7 |
| Ours (w self-supervised pre-training; ℓ_c loss) | 41.89 | 0.98 | 1.46 | 2.04 | 41.37 | 0.98 | 1.41 | 2.53 | 52.7 |

Table 2: Raw-to-raw translation results on the Raw-to-Raw dataset [4]. Shown are the mapping results of Samsung Galaxy S9’s raw images to iPhone X and vice versa using our method and other methods. Additionally, the total number of parameters of each method is shown in millions (M), and inference time in milliseconds (ms) without optimization on NVIDIA GeForce RTX 4090 GPU. *Results on inference time with optimization are in supplementary.*

| Methods | Samsung-to-iPhone | | | | iPhone-to-Samsung | | | | Params (M) | Time (ms) |
|---------------|-------------------|-------------|------------------------|--------------|-------------------|-------------|------------------------|--------------|-------------|-----------|
| | PSNR↑ | SSIM↑ | MAE↓ ($\times 1e-2$) | ΔE ↓ | PSNR↑ | SSIM↑ | MAE↓ ($\times 1e-2$) | ΔE ↓ | | |
| CycleGAN [62] | 24.63 | 0.71 | 5.40 | 14.71 | 24.37 | 0.73 | 4.34 | 12.88 | 21.7 | 19 |
| UNIT [32] | 23.91 | 0.70 | 5.15 | 13.23 | 26.42 | 0.78 | 4.27 | 13.16 | 32.2 | 29 |
| U-GAT-IT [28] | 28.22 | 0.89 | 3.11 | 5.98 | 31.04 | 0.89 | 2.39 | 5.04 | 34.4 | 35 |
| SSRM [4] | 29.65 | 0.89 | 2.22 | 6.32 | 28.58 | 0.90 | 2.43 | 6.53 | 20.2 | 18 |
| UVCGANv2 [51] | 36.32 | 0.94 | 2.09 | 4.21 | 36.46 | 0.92 | 2.02 | 4.73 | 32.6 | 31 |
| Ours | 40.98 | 0.97 | 1.33 | 2.09 | 41.48 | 0.98 | 1.21 | 1.99 | 26.1 | 26 |

our method, for 500 epochs. Since SSRM is not entirely unsupervised, we incorporated 22 image pairs from the paired subset for its training, as discussed in [4]. Fig. 4 shows qualitative comparisons on the Raw-to-Raw dataset [4], while Tables 1 and 2 present the quantitative results on the NUS dataset [9] and the Raw-to-Raw dataset, respectively. In Tables 1 and 2, we adopted the same quantitative metrics used in [4]: peak signal-to-noise ratio (PSNR), SSIM, mean absolute error (MAE), and ΔE 2000 [46]. As can be seen in both the tables, our Rawformer significantly outperforms the alternative methods across all the quantitative metrics by a large margin (around +12 dB against SSRM [4]).

Evaluation on Low-Light Images: Translating raw-to-raw for low-light images presents challenges due to various variations among cameras, including differences in raw normalization and noise characteristics [2,4]. To further prove the efficacy of our method against other methods on low-light images, we selected a subset of 15 paired low-light images from the Raw-to-Raw dataset [4]. The results on this subset are presented in Table 3, where our proposed method significantly outperforms SSRM [4], and UVCGANv2 [51] for translating low-light raw images captured by an iPhone X to the Samsung S9’s camera.

Table 3: Raw-to-Raw dataset [4] results on low-light raw images (a subset).

| | | | | |
|---|---|------|-------------|----|
| iPhone | PSNR, SSIM | SSRM | 23.86, 0.85 | S9 |
|  |  | | | |
| |  | | | |
| |  | | | |
| Input | Prediction | | | GT |

| Method | PSNR↑ | SSIM↑ |
|---------------|--------------|-------------|
| SSRM [4] | 23.87 | 0.85 |
| UVCGANv2 [51] | 35.22 | 0.94 |
| Ours | 40.21 | 0.97 |

Table 4: Ablation results on the impact of various configurations of the generator model on the NUS dataset [9]. ‘CQA’, ‘Style’, ‘SPFN’ and ‘CUp-CDown’ refers to the condensed query attention, style modulator, scale perceptive feed-forward network, and the composite up and downsampler, respectively.

| Configuration | Modules | | | | | Canon-to-Nikon | | Nikon-to-Canon | |
|---------------|----------|-----|-------|------|-----------|----------------|-------------|----------------|-------------|
| | Baseline | CQA | Style | SPFN | CUp-Cdown | PSNR↑ | SSIM↑ | PSNR↑ | SSIM↑ |
| G_1 | ✓ | | | | | 37.23 | 0.96 | 37.31 | 0.96 |
| G_2 | ✓ | ✓ | | | | 38.21 | 0.96 | 37.97 | 0.96 |
| G_3 | ✓ | ✓ | ✓ | | | 39.08 | 0.97 | 38.99 | 0.96 |
| G_4 | ✓ | ✓ | ✓ | ✓ | | 39.72 | 0.97 | 39.68 | 0.97 |
| G_5 | ✓ | ✓ | ✓ | ✓ | ✓ | 39.93 | 0.97 | 39.86 | 0.97 |

Table 5: Ablation results on the impact of various discriminator model configurations on the NUS dataset [9]. ‘Baseline’ refers to the G_5 config. from Tab. 4.

| Configuration | Modules | | | | Canon-to-Nikon | | Nikon-to-Canon | |
|---------------|----------|-----|-----------|------------|----------------|-------------|----------------|-------------|
| | Baseline | CQA | Batchnorm | Disc. head | PSNR↑ | SSIM↑ | PSNR↑ | SSIM↑ |
| D_1 | ✓ | | | | 39.93 | 0.97 | 39.86 | 0.97 |
| D_2 | ✓ | ✓ | | | 40.12 | 0.97 | 40.09 | 0.97 |
| D_3 | ✓ | ✓ | ✓ | | 40.24 | 0.97 | 40.21 | 0.97 |
| D_4 | ✓ | ✓ | ✓ | ✓ | 41.89 | 0.98 | 41.37 | 0.98 |

Ablation Studies: We conducted a set of ablation studies to validate the impact of the proposed components in the generator and discriminator networks. In this series of experiments, we trained our Rawformer on the NUS dataset [9], in addition to a baseline model trained using the same training settings as our Rawformer. The baseline model comprises a UNet architecture with channel and spatial attention [50] for the generator, and a PatchGAN discriminator [23]. As demonstrated in Table 4, the integration of the proposed modules into the baseline consistently improves the network’s capabilities, showcasing an enhancement in overall accuracy. We also conducted further ablation studies that explore various configurations of the discriminator, as shown in Table 5. The results clearly underscore the improvement by the inclusion of each proposed component in the discriminator network (attention, batch normalization, and discriminator head (Disc. head)) in improving the accuracy of our raw-to-raw translation.

4.4 Raw Translation for Learnable ISP

To comprehend the practical application of our model in real-world scenarios, we utilized three datasets: Zurich raw-to-RGB, Samsung S7 ISP, and Mobile AIM21, encompassing various source and target domain combinations. Specifically, for each dataset, we trained the LAN ISP [54] on a single paired raw-to-sRGB

Table 6: Cross-camera raw-to-sRGB results (PSNR/SSIM) comparing: a) UVCGANv2 [51] with b) our Rawformer. Each entry compares ground-truth sRGB images generated by the function: $ISP_B(F^B(raw_A))$, where F^B denotes translation of raw images, raw_A , captured by camera A to camera B . $ISP_B(\cdot)$ produces sRGB images using the LAN ISP [54] trained on raw images from camera B . Diagonal results (in **red**) represent the *camera-specific* ISP case and are the sRGB images produced by the ISP model tested on the same camera used for training. Best cross-camera results are in bold.

| Training camera A | a) UVCGANv2 [51]+ LAN ISP [54] | | | b) Our Rawformer+ LAN ISP [54] | | |
|-------------------|--------------------------------|-------------------|-------------------|--------------------------------|-------------------|-------------------|
| | Testing camera B | | | Testing camera B | | |
| | ZRR [20] | AIM21 [19] | S7 [44] | ZRR [20] | AIM21 [19] | S7 [44] |
| ZRR [20] | 19.46/0.73 | 17.72/0.70 | 17.81/0.70 | 19.46/0.73 | 18.97/0.72 | 19.02/0.72 |
| AIM21 [19] | 22.14/0.79 | 23.48/0.87 | 22.39/0.79 | 23.25/0.86 | 23.48/0.87 | 23.08/0.86 |
| S7 [44] | 21.07/0.80 | 20.93/0.79 | 22.16/0.81 | 22.08/0.80 | 22.03/0.80 | 22.16/0.81 |

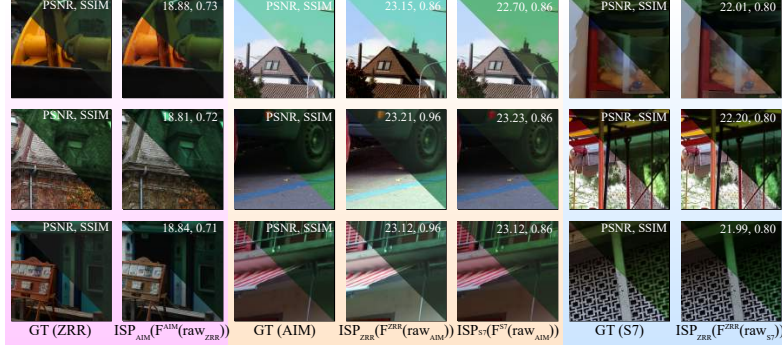


Fig. 5: ISP rendering results with our raw translation on various datasets. Each set includes ground-truth (GT) raw-sRGB paired images and LAN ISP [54] results on mapped raw images from different cameras. S7, ZRR, and AIM stand for the Samsung S7 ISP dataset [44], Zurich raw-to-RGB dataset [20], and Mobile AIM21 dataset [19], respectively. The shown results are consistent with the ground-truth, demonstrating the proficiency of our model.

dataset, where raw images originated from a single camera. Then, we transferred all raw images from other cameras in the other datasets using our Rawformer before rendering with the trained LAN model. We report the PSNR and SSIM of the rendered sRGB images compared to the ground-truth sRGB images for each set as shown in Table 6. This includes the results of *camera-specific* LAN ISP models trained and tested on the same datasets. The camera-specific results represent the optimal results achievable by the neural-based ISP, which requires a paired raw-sRGB dataset for training on each camera model. As can be seen, our unsupervised raw-to-raw translation method demonstrates robust improvement, exhibiting only a marginal 3% reduction in accuracy compared to the camera-specific LAN model trained and tested on the same camera. This underscores the effectiveness of our method in addressing real-world cross-domain broadcasting challenges. See Fig. 5 for qualitative examples.

4.5 Comparison with Few-Shot Domain Adaptation

We further compared our Rawformer against Prabhakar’s Few-Shot Domain Adaptation (PDA) [40]. Specifically, we conducted a 10-shot domain adaptation from the Samsung S7 ISP dataset’s camera [44] to the Zurich raw-to-RGB dataset’s camera [20], and vice-versa. The PDA model underwent initial training for 10 epochs using the Samsung S7 ISP dataset. Subsequently, domain adaptation was performed on 10 images from the Zurich raw-to-RGB dataset’s camera over 4 epochs. This process resulted in performance metrics of 21.07 dB, and 17.81 dB in terms of PSNR and 0.80 and 0.70 for SSIM for our Rawformer, representing approximately 11% superior performance compared to PDA (see Table 7).

Table 7: Comparison with PDA [40]. S7 and ZRR denote the Samsung S7 ISP [44] and the Zurich Raw-to-RGB [20] datasets, respectively. Results are produced after translating each raw image to the target camera using PDA and our Rawformer, subsequently rendering with a LAN ISP [54] trained on the target camera’s raw images.

| Method | S7-to-ZRR | | | | ZRR-to-S7 | | | |
|----------|-----------------|-----------------|------------------------------------|----------------------|-----------------|-----------------|------------------------------------|----------------------|
| | PSNR \uparrow | SSIM \uparrow | MAE \downarrow ($\times 1e-2$) | $\Delta E\downarrow$ | PSNR \uparrow | SSIM \uparrow | MAE \downarrow ($\times 1e-2$) | $\Delta E\downarrow$ |
| PDA [40] | 20.33 | 0.79 | 9.41 | 13.99 | 17.12 | 0.69 | 18.03 | 15.86 |
| Ours | 22.08 | 0.80 | 8.18 | 10.27 | 19.02 | 0.72 | 15.27 | 11.74 |

5 Conclusion and Future Work

In this work, we presented a new approach, Rawformer, for raw-to-raw image translation using unpaired data. Rawformer sets a new standard for performance on real camera datasets, showcasing its superior capabilities in raw-to-raw translation when measured against other alternatives. The proposed method reduces the need for the extensive collection of raw-to-sRGB paired datasets to re-train learnable ISPs for new cameras, thereby decreasing associated costs. Our results indicate that existing raw-to-sRGB paired datasets can reliably train neural-based learnable ISPs, which can then be utilized to render raw images taken by unseen cameras, lacking such unpaired data. We believe that our approach provides valuable insights for learnable ISP, raw image translation research, and applications. Our method (without optimization) runs at 26 milliseconds per frame on GPU (~ 1 second on CPU), which may be impractical for real-time rendering in camera previews on devices with limited computational power. Future work will focus on introducing lighter models capable of real-time performance on CPU.

Acknowledgments

This work was partly supported by The Alexander von Humboldt Foundation.

References

1. A Sharif, S., Naqvi, R.A., Biswas, M.: Beyond joint demosaicking and denoising: An image processing pipeline for a pixel-bin image sensor. In: CVPR (2021)
2. Abdelhamed, A., Lin, S., Brown, M.S.: A high-quality denoising dataset for smart-phone cameras. In: CVPR (2018)
3. Affi, M., Abdelhamed, A., Abuolaim, A., Punnappurath, A., Brown, M.S.: CIE XYZ net: Unprocessing images for low-level computer vision tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(9), 4688–4700 (2021)
4. Affi, M., Abuolaim, A.: Semi-supervised raw-to-raw mapping. In: BMVC (2021)
5. Affi, M., Brown, M.S.: Sensor-independent illumination estimation for dnn models. In: BMVC (2019)
6. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. *arXiv preprint arXiv:1607.06450* (2016)
7. Barron, J.T., Tsai, Y.T.: Fast Fourier color constancy. In: CVPR (2017)
8. Chen, C., Chen, Q., Xu, J., Koltun, V.: Learning to see in the dark. In: CVPR (2018)
9. Cheng, D., Prasad, D.K., Brown, M.S.: Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution. *JOSA A* **31**(5), 1049–1058 (2014)
10. Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In: CVPR (2018)
11. Conde, M.V., Vasluianu, F., Vazquez-Corral, J., Timofte, R.: Perceptual image enhancement for smartphone real-time applications. In: CVPR (2023)
12. Dai, L., Liu, X., Li, C., Chen, J.: AWNNet: Attentive wavelet network for image ISP. In: ECCV (2020)
13. Delbracio, M., Kelly, D., Brown, M.S., Milanfar, P.: Mobile computational photography: A tour. *Annual Review of Vision Science* **7**, 571–604 (2021)
14. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
15. Finlayson, G.D., Zhu, Y.: Designing color filters that make cameras more colorimetric. *IEEE Transactions on Image Processing* **30**, 853–867 (2020)
16. Hasinoff, S.W., Sharlet, D., Geiss, R., Adams, A., Barron, J.T., Kainz, F., Chen, J., Levoy, M.: Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics* **35**(6), 1–12 (2016)
17. He, X., Hu, T., Wang, G., Wang, Z., Wang, R., Zhang, Q., Yan, K., Chen, Z., Li, R., Xie, C., et al.: Enhancing RAW-to-sRGB with decoupled style structure in Fourier domain. In: AAAI (2024)
18. Herrmann, C., Bowen, R.S., Wadhwa, N., Garg, R., He, Q., Barron, J.T., Zabih, R.: Learning to autofocus. In: CVPR (2020)
19. Ignatov, A., Chiang, C.M., Kuo, H.K., Sycheva, A., Timofte, R.: Learned smartphone ISP on mobile NPUs with deep learning, mobile AI 2021 challenge: Report. In: CVPRW (2021)
20. Ignatov, A., Timofte, R., Ko, S.J., Kim, S.W., Uhm, K.H., Ji, S.W., Cho, S.J., Hong, J.P., Mei, K., Li, J., et al.: AIM 2019 challenge on raw to rgb mapping: Methods and results. In: ICCVW (2019)
21. Ignatov, A., Timofte, R., Zhang, Z., Liu, M., Wang, H., Zuo, W., Zhang, J., Zhang, R., Peng, Z., Ren, S., et al.: AIM 2020 challenge on learned image signal processing pipeline. In: ECCV (2020)

22. Ignatov, A., Van Gool, L., Timofte, R.: Replacing mobile camera ISP with a single deep learning model. In: CVPRW (2020)
23. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017)
24. Jeong, W., Jung, S.W.: RAWtoBit: A fully end-to-end camera isp network. In: ECCV (2022)
25. Jiang, Y., Wronski, B., Mildenhall, B., Barron, J.T., Wang, Z., Xue, T.: Fast and high quality image denoising via malleable convolution. In: ECCV (2022)
26. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196 (2017)
27. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: CVPR (2020)
28. Kim, J., Kim, M., Kang, H., Lee, K.: U-GAT-IT: Unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. arXiv preprint arXiv:1907.10830 (2019)
29. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
30. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: CVPR (2017)
31. Liang, Z., Cai, J., Cao, Z., Zhang, L.: Cameranet: A two-stage framework for effective camera isp learning. *IEEE Transactions on Image Processing* **30**, 2248–2262 (2021)
32. Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. In: NeurIPS (2017)
33. Loshchilov, I., Hutter, F.: SGDR: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
34. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
35. Maas, A.L., Hannun, A.Y., Ng, A.Y., et al.: Rectifier nonlinearities improve neural network acoustic models. In: ICML (2013)
36. Menon, D., Andriani, S., Calvagno, G.: Demosaicing with directional filtering and a posteriori decision. *IEEE Transactions on Image Processing* **16**(1), 132–141 (2006)
37. Nguyen, R., Prasad, D.K., Brown, M.S.: Raw-to-raw: Mapping between image sensor color responses. In: CVPR (2014)
38. Pang, Y., Lin, J., Qin, T., Chen, Z.: Image-to-image translation: Methods and applications. *IEEE Transactions on Multimedia* **24**, 3859–3881 (2021)
39. Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: ECCV (2020)
40. Prabhakar, K.R., Vinod, V., Sahoo, N.R., Babu, R.V.: Few-shot domain adaptation for low light raw image enhancement. arXiv preprint arXiv:2303.15528 (2023)
41. Punnapurath, A., Abuolaim, A., Abdelhamed, A., Levinshtein, A., Brown, M.S.: Day-to-night image synthesis for training nighttime neural isps. In: CVPR (2022)
42. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention (2015)
43. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. *Advances in neural information processing systems* **29** (2016)

44. Schwartz, E., Giryes, R., Bronstein, A.M.: DeepISP: Toward learning an end-to-end image processing pipeline. *IEEE Transactions on Image Processing* **28**(2), 912–923 (2018)
45. Seo, D., Punnappurath, A., Zhao, L., Abdelhamed, A., Tedla, S.K., Park, S., Choe, J., Brown, M.S.: Graphics2RAW: Mapping computer graphics images to sensor raw images. In: *ICCV* (2023)
46. Sharma, G., Wu, W., Dalal, E.N.: The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application* **30**(1), 21–30 (2005)
47. Šindelář, O., Šroubek, F.: Image deblurring in smartphone devices using built-in inertial measurement sensors. *Journal of Electronic Imaging* **22**(1), 011003–011003 (2013)
48. Souza, M., Heidrich, W.: Crispnet: Color rendition isp net. *arXiv preprint arXiv:2203.10562* (2022)
49. Tominaga, S., Nishi, S., Ohtera, R.: Measurement and estimation of spectral sensitivity functions for mobile phone cameras. *Sensors* **21**(15), 4985 (2021)
50. Torbunov, D., Huang, Y., Yu, H., Huang, J., Yoo, S., Lin, M., Viren, B., Ren, Y.: UVCGAN: UNet vision transformer cycle-consistent GAN for unpaired image-to-image translation. In: *WACV* (2023)
51. Torbunov, D., Huang, Y., Yu, H., Huang, J., Yoo, S., Lin, M., Viren, B., Ren, Y.: UVCGAN v2: An improved cycle-consistent GAN for unpaired image-to-image translation. *arXiv preprint arXiv:2303.16280* (2023)
52. Truong, P., Danelljan, M., Van Gool, L., Timofte, R.: Learning accurate dense correspondences and when to trust them. In: *CVPR* (2021)
53. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004)
54. Wirzberger Raimundo, D., Ignatov, A., Timofte, R.: LAN: Lightweight attention-based network for raw-to-rgb smartphone image processing. In: *CVPRW*. pp. 807–815 (2022)
55. Wronski, B., Garcia-Dorado, I., Ernst, M., Kelly, D., Krainin, M., Liang, C.K., Levoy, M., Milanfar, P.: Handheld multi-frame super-resolution. *ACM Transactions on Graphics (ToG)* **38**(4), 1–18 (2019)
56. Xing, Y., Qian, Z., Chen, Q.: Invertible image signal processing. In: *CVPR* (2021)
57. Yi, Z., Zhang, H., Tan, P., Gong, M.: DualGAN: Unsupervised dual learning for image-to-image translation. In: *ICCV*. pp. 2849–2857 (2017)
58. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 5728–5739 (2022)
59. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: CycleISP: Real image restoration via improved data synthesis. In: *CVPR* (2020)
60. Zhang, Z., Wang, H., Liu, M., Wang, R., Zhang, J., Zuo, W.: Learning raw-to-srgb mappings with inaccurately aligned supervision. In: *ICCV* (2021)
61. Zhao, Y., Wu, R., Dong, H.: Unpaired image-to-image translation using adversarial consistency loss. In: *ECCV* (2020)
62. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *ICCV* (2017)
63. Zhu, P., Abdal, R., Qin, Y., Wonka, P.: SEAN: Image synthesis with semantic region-adaptive normalization. In: *CVPR* (2020)