

Edit-aware RAW Reconstruction

Abhijith Punnappurath Luxi Zhao* Ke Zhao*
Hue Nguyen Radek Grzeszczuk Michael S. Brown
AI Center–Toronto, Samsung Electronics

{abhijith.p, lucy.zhao, k.zhao, h.nguyen1, radek.g, michael.b1}@samsung.com

Abstract

Users frequently edit camera images post-capture to achieve their preferred photofinishing style. While editing in the RAW domain provides greater accuracy and flexibility, most edits are performed on the camera’s display-referred output (e.g., 8-bit sRGB JPEG) since RAW images are rarely stored. Existing RAW reconstruction methods can recover RAW data from sRGB images, but these approaches are typically optimized for pixel-wise RAW reconstruction fidelity and tend to degrade under diverse rendering styles and editing operations. We introduce a plug-and-play, edit-aware loss function that can be integrated into any existing RAW reconstruction framework to make the recovered RAWs more robust to different rendering styles and edits. Our loss formulation incorporates a modular, differentiable image signal processor (ISP) that simulates realistic photofinishing pipelines with tunable parameters. During training, parameters for each ISP module are randomly sampled from carefully designed distributions that model practical variations in real camera processing. The loss is then computed in sRGB space between ground-truth and reconstructed RAWs rendered through this differentiable ISP. Incorporating our loss improves sRGB reconstruction quality by up to 1.5–2 dB PSNR across various editing conditions. Moreover, when applied to metadata-assisted RAW reconstruction methods, our approach enables fine-tuning for target edits, yielding further gains. Since photographic editing is the primary motivation for RAW reconstruction in consumer imaging, our simple yet effective loss function provides a general mechanism for enhancing edit fidelity and rendering flexibility across existing methods.

1. Introduction

Users edit camera-captured images in their photo galleries to achieve a preferred photofinishing style. Gallery images are the output of cameras’ on-board image signal processor (ISP), which renders the RAW sensor image into a fi-

*Equal contribution.

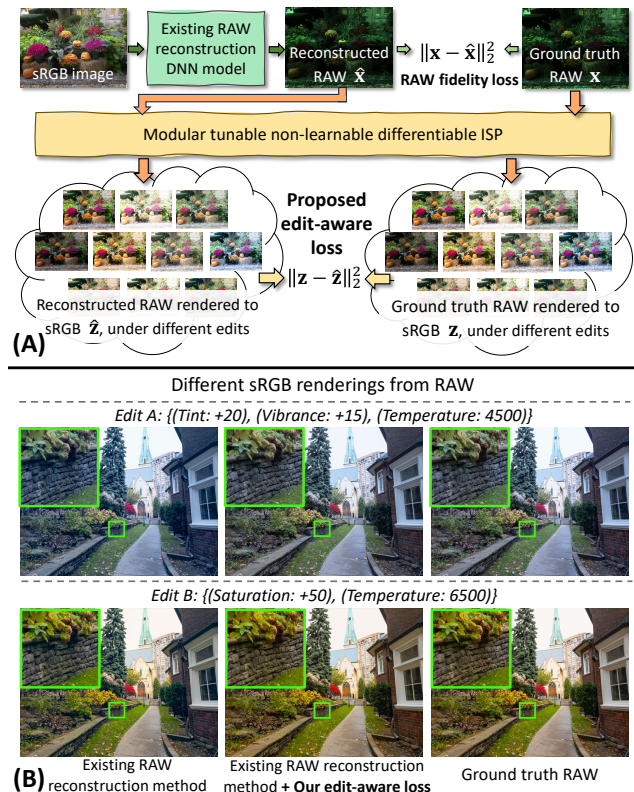


Figure 1. (A) Existing RAW reconstruction methods typically employ loss functions focused on accurate per-pixel RAW recovery (green path). However, the primary use of a reconstructed RAW image is to allow high-quality editing targeting sRGB outputs. This paper introduces an edit-aware loss for RAW reconstruction that yields higher sRGB fidelity across diverse edits (yellow path). (B) Two example edits applied to the RAW estimated by a conventional reconstruction method [27] (left), the RAW reconstructed by augmenting [27] with our edit-aware loss (middle), and the ground-truth RAW (right). Renderings are produced in Adobe Photoshop using the indicated settings. Our results more closely match the ground truth across edits.

nal photofinished output image [14]. The ISP follows a default rendering pipeline governed by fixed or discretely

selectable photofinishing styles determined at capture time. Users often perform post-capture editing because they prefer a different style than the ISP default, or find that the ISP’s automatically selected settings produced a photofinished image that did not match their intent.

Post-capture editing in the RAW domain offers significant advantages over manipulating the photofinished image [17]. The RAW image represents the minimally processed sensor measurement of the physical scene, maintaining a linear response to scene radiance and preserving higher bit depth and dynamic range. In contrast, the ISP applies a series of operations—such as denoising, demosaicing, exposure and white balance correction, and color and tone mapping—to produce the final rendered output [14]. Many of these operations are non-invertible and introduce information loss through tonal compression, clipping, and quantization. Consequently, the photofinished image—typically encoded in a perceptual color space (e.g., sRGB or Display P3) with 8–10 bits per channel and compressed using formats like JPEG or HEIC—contains reduced information content. In consumer photography, the most common output remains 8-bit sRGB JPEG. Although editing RAW images allows more accurate and flexible adjustments due to their higher bit depth and tonal range, RAW files are rarely retained because of their large size and limited compatibility with standard display and sharing workflows.

RAW reconstruction methods aim to recover RAW sensor measurements from their rendered sRGB counterparts. Existing approaches can be broadly categorized into metadata-assisted and blind methods. Metadata-assisted approaches leverage additional capture-time information such as RAW samples [19, 23, 26] or RAW latent vectors [30], and typically achieve higher reconstruction fidelity than blind methods [27, 33, 35], which rely solely on the sRGB input. The vast majority of RAW recovery methods optimize exclusively for pixel-level RAW reconstruction accuracy, without considering downstream use cases such as photographic rendering or editing. A few approaches [33, 35] include cyclic losses that reprocess the recovered RAW to verify that it matches the unedited sRGB input image. While the fidelity loss is in the sRGB domain, such supervision does not provide robustness to diverse photofinishing styles or post-capture edits. Some methods (e.g., [5]) have explored task-specific RAW reconstruction for computer vision applications (e.g., object detection), but these objectives are not aligned with the goals of photographic editing. In consumer photography, the primary motivation for RAW reconstruction is to enable high-quality, flexible image editing—a setting in which existing RAW-fidelity-driven approaches may produce suboptimal results. See the example in the first column of Fig. 1(B).

Contributions: The limitations discussed above highlight the need for a reconstruction objective that directly ac-

counts for the downstream goal of robust RAW-image editing supporting a wide variety of sRGB renderings. To address this gap, we introduce a plug-and-play edit-aware loss that can be integrated into any RAW reconstruction framework to improve robustness across a wide range of editing conditions. Our loss is formulated as a modular and tunable differentiable ISP that emulates realistic photofinishing pipelines by rendering RAW inputs to sRGB under a wide range of sampled edit settings. Importantly, our method does not assume access to the camera ISP that produced the input sRGB image. In practice, camera ISPs are proprietary black-box pipelines. Our approach instead uses a simplified differentiable ISP only as a training-time loss, without requiring knowledge of or access to the actual camera ISP. During training, parameters for each ISP module are randomly drawn from distributions designed to reflect common operating ranges, ensuring realistic variability. The loss is computed in sRGB space between the ground-truth and reconstructed RAW images, with both rendered through our differentiable ISP using the same random settings. See Fig. 1(A). By incorporating our edit-aware loss, the reconstructed RAW is optimized for editability and rendering robustness, trading a small amount of pixel-wise RAW-space fidelity for significantly improved sRGB rendering accuracy and consistency across edits.

We evaluate our approach across three representative RAW reconstruction models from both metadata-assisted and blind categories. We demonstrate that our loss generalizes to real-world editing workflows by applying diverse photographic edits in Adobe Photoshop. The introduction of our edit-aware loss, implemented through our lightweight differentiable ISP, consistently improves performance across all tested settings and we observe gains of up to 1.5–2 dB PSNR in sRGB space under various edits. An example is shown in the second column of Fig. 1(B). Overall, our approach provides a simple yet powerful mechanism for enhancing the practical utility of RAW reconstruction methods in real-world photographic editing scenarios.

2. Related work

We first review metadata-assisted approaches to RAW reconstruction, which leverage additional capture-time information to achieve higher accuracy than purely blind methods. One of the earliest works [34] stored a downsampled RAW image and performed guided upsampling to recover the full-resolution RAW. Nguyen and Brown [24, 25] computed and stored metadata in the form of estimated parameters modeling a global transformation from sRGB to RAW. Punnappurath and Brown [26] instead saved a sparse set of uniformly sampled RAW values and used radial basis function interpolation to capture non-global ISP effects. Li et al. [19] replaced this analytical interpolation with a neural implicit representation to better model nonlinear map-

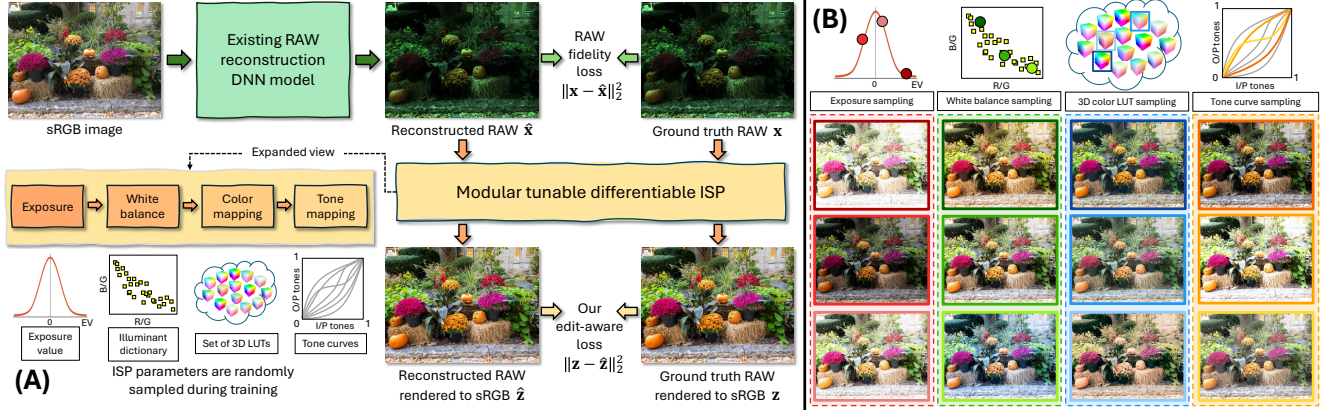


Figure 2. Overview of our edit-aware loss framework and differentiable ISP design. (A) Conventional RAW reconstruction models optimize for per-pixel RAW fidelity (green path), which often yields suboptimal results under diverse rendering or editing operations. We introduce an edit-aware loss (yellow path) that complements existing objectives by promoting robustness across different rendering styles. Our loss pipeline mimics a modular tunable differentiable ISP whose module parameters—exposure, white balance, color, and tone—are randomly sampled from distributions modeling realistic ISP settings. Both the ground-truth and reconstructed RAWs are rendered to sRGB using the same sampled parameters, and the loss is computed in sRGB space to encourage edit-consistent reconstruction. (B) Each column shows three samples illustrating the diversity of exposure, white balance, color, and tone adjustments encountered during training.

tings. CAM [23] improved upon fixed sampling by introducing content-aware sampling and a jointly trained decoder to enhance spatial adaptivity. Wang et al. [30, 31] further proposed storing metadata as compact RAW latent features instead of direct pixel samples. Although these methods differ in sampling strategy and model class, they uniformly optimize for per-pixel RAW fidelity (sometimes with auxiliary regularizers, such as super-pixel losses for content-aware sampling [23] or entropy coding for adaptive bit allocation [30]) and do not account for the need to re-render the reconstructed RAW to sRGB. In contrast, our edit-aware loss operates directly in sRGB space, optimizing reconstructed RAWs for rendering and editing robustness.

A complementary line of work explores blind RAW reconstruction methods that operate solely on the sRGB input [3, 6, 10, 11, 18, 22, 27, 33, 35]. While RAW recovery targeting machine vision tasks, such as object detection, have been explored [5], most methods focus exclusively on reconstruction accuracy without task-specific objectives. Early work on RAW reconstruction was calibration-based [7, 8, 13, 15, 17, 21]. Later methods, such as UPI [6], reversed the sRGB image back to RAW by inverting the ISP stage-by-stage using non-learnable parametric operations. Nam et al. [22] introduced a deep model trained on paired RAW-sRGB data for jointly learning forward and inverse ISP transformations, later extended by methods such as CIE-XYZ-Net [3], CycleISP [35], InvISP [33], ParamISP [18], and model-based ISP approaches [10]. Recently, diffusion-based methods [27] have also been explored for RAW reconstruction, but they continue to optimize losses in RAW space. While the forward-inverse ap-

proaches [3, 10, 18, 22, 33, 35] sometimes include cyclic consistency loss terms to enforce sRGB reconstruction from the recovered RAW, the supervision always corresponds to the same input sRGB without any edits. Consequently, they cannot explicitly account for diverse photofinishing styles. Our work differs by introducing a stochastic, edit-aware loss computed in sRGB space via a tunable differentiable ISP, thereby encouraging robustness across diverse rendering styles and enabling accurate post-capture RAW editing targeting higher-quality sRGB outputs—the practical end goal for most RAW recovery workflows.

3. Methodology

We first briefly review the loss objectives used in existing RAW reconstruction methods. We then describe the ISP modules that form our edit-aware loss pipeline and explain how their tunable parameters are sampled during training.

Let \mathbf{x} denotes the ground truth RAW image and \mathbf{y} is the ISP’s sRGB rendering of \mathbf{x} , then the objective of RAW reconstruction is to recover the RAW image $\hat{\mathbf{x}}$, given \mathbf{y} , as:

$$\hat{\mathbf{x}} = f_{\theta}(\mathbf{y}), \quad (1)$$

where $f_{\theta} : \mathcal{Y} \rightarrow \mathcal{X}$ denotes the RAW reconstruction DNN model with learnable parameters θ . In the case of metadata-assisted RAW reconstruction methods, the model f_{θ} may receive additional capture-time information as input, in addition to the sRGB image \mathbf{y} . Existing methods are trained to minimize the RAW reconstruction error between the ground truth RAW \mathbf{x} and the network’s prediction $\hat{\mathbf{x}}$:

$$\mathcal{L}_{\text{RAW}}(\mathbf{x}, \hat{\mathbf{x}}) = \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2, \quad (2)$$

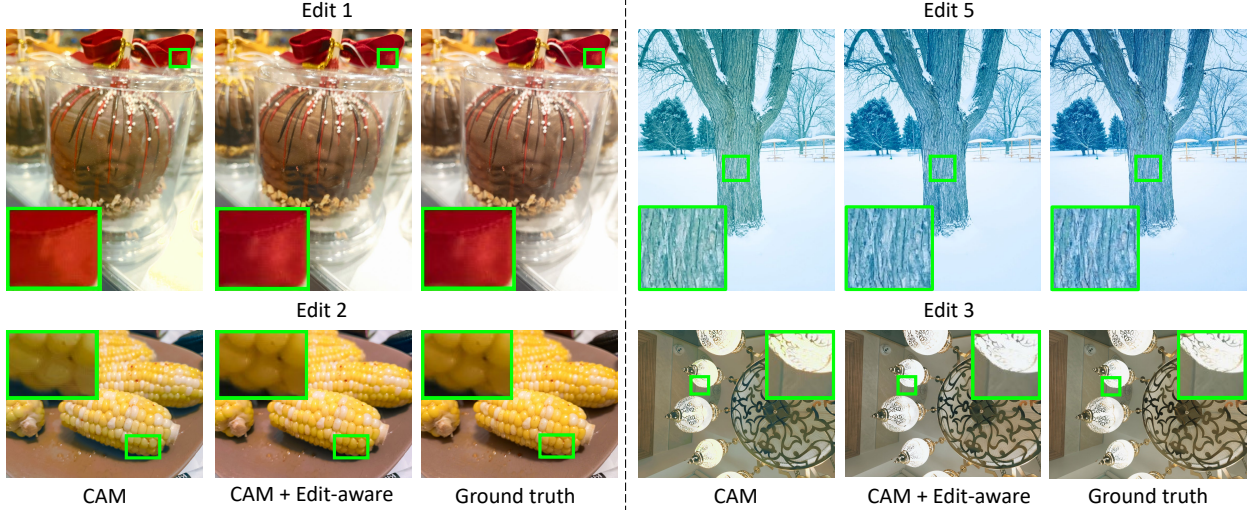


Figure 3. Results of adding our edit-aware loss to the RAW reconstruction method in CAM [23]. See Table 1 for details of the edits. While the baseline model suffers from banding and color distortions, adding our edit-aware loss improves color and tone reproduction accuracy.

typically, using ℓ_2 error.

As shown in Fig. 2(A), we propose adding a loss function $\mathcal{L}_{\text{sRGB}}$ computed between the ground truth RAW \mathbf{x} and the network’s prediction $\hat{\mathbf{x}}$, both rendered to sRGB using our differentiable ISP pipeline $g_\phi: \mathcal{X} \rightarrow \mathcal{Y}$.

$$\begin{aligned} \mathbf{z} &= g_\phi(\mathbf{x}), & \hat{\mathbf{z}} &= g_\phi(\hat{\mathbf{x}}), \\ \mathcal{L}_{\text{sRGB}}(\mathbf{z}, \hat{\mathbf{z}}) &= \|\mathbf{z} - \hat{\mathbf{z}}\|_2^2. \end{aligned} \quad (3)$$

Here, the parameters ϕ of the ISP g_ϕ are not learned but sampled during training from distributions designed to closely mimic realistic ISP renderings such that the output of g_ϕ lies in the same distribution as \mathbf{y} i.e., $\mathbf{z}, \hat{\mathbf{z}}, \mathbf{y} \in \mathcal{Y}$.

Our edit-aware loss pipeline g_ϕ models four ISP modules, namely (i) exposure e_ε , (ii) white balance w_ω , (iii) color manipulation c_ρ , and (iv) tone mapping t_τ , such that:

$$\phi = (\varepsilon, \omega, \rho, \tau), \quad g_\phi = t_\tau \circ c_\rho \circ w_\omega \circ e_\varepsilon. \quad (4)$$

Note that on a typical ISP, black-, white-level normalization, demosaicing and denoising are the first set of operations to be performed. Most RAW reconstruction methods recover a 3-channel denoised RAW image normalized to the range $[0, 1]$. As such, we assume normalization, demosaicing and denoising have already been applied, and do not include these stages in our loss pipeline. Without loss of generality, in the subsequent discussion, we adopt a vectorized notation for all images $\in \mathbb{R}^{3 \times N}$, where N denotes the total number of pixels in the image. The structure of our tunable and modular ISP framework and the details of these four modules are described next.

3.1. Exposure module

Exposure adjustments are typically applied directly to the RAW sensor image. We model exposure as:

$$\mathbf{p}_1 = e_\varepsilon(\mathbf{p}_0) = \mathbf{p}_0 \cdot 2^\varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2), \quad \mathbf{p}_0 \in \{\mathbf{x}, \hat{\mathbf{x}}\}, \quad (5)$$

where ε represents the exposure value. This parameter is randomly sampled from a normal distribution \mathcal{N} for each image in each mini-batch during training. Note that the same random sample is applied to \mathbf{x} and $\hat{\mathbf{x}}$. Examples of exposure sampling are shown in the first column of Fig. 2(B).

3.2. White-balance module

Cameras apply a white-balance correction to the RAW image to remove the color cast due to the scene illumination. The idea is to divide each color channel of the image by the scene illumination, producing an image that appears lit under neutral illumination. However, users may still wish to adjust the white-point post-capture for aesthetic reasons (e.g., if they prefer a warmer or cooler image).

The illuminant estimated by the ISP is stored in the ‘AsShotNeutral’ (ASN) tag in the RAW DNG file. To randomly sample an illuminant during training, we construct an illuminant dictionary \mathcal{D} by extracting the ASN tags from a set of DNGs captured under different illuminations. We fit a 2D multivariate Gaussian distribution of joint chromaticity values $\begin{bmatrix} r \\ g \end{bmatrix}$ to \mathcal{D} , and randomly sample an illuminant ω from this distribution as follows:

$$\omega \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (6)$$

$$\boldsymbol{\Sigma} = \frac{1}{M} \sum_{m=1}^M \left(\left(\begin{bmatrix} r \\ g \end{bmatrix}_m - \boldsymbol{\mu} \right) \left(\begin{bmatrix} r \\ g \end{bmatrix}_m - \boldsymbol{\mu} \right)^\top \right), \quad (7)$$

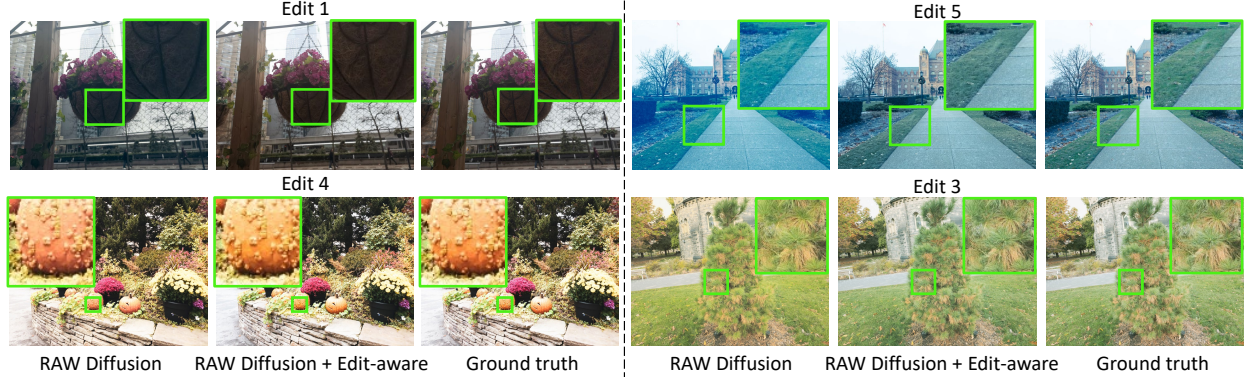


Figure 4. Results of adding our edit-aware loss to the RAW recovery method in RAW Diffusion [27]. See Table 1 for details of the edits. Compared to the baseline, which shows pronounced global color distortions, our edit-aware method more closely matches the ground truth.

where μ and Σ are the mean and covariance of the normalized chromaticity values in \mathcal{D} , M is the number of illuminants in \mathcal{D} , and $\omega, \mu \in \mathbb{R}^2$ and $\Sigma \in \mathbb{R}^{2 \times 2}$. We enforce that the sampled illuminant ω falls within the convex hull of \mathcal{D} . We further constrain it to lie within a Euclidean threshold in chromaticity space relative to the ASN value of the image, because users typically make only small adjustments to the white balance during editing. White balance is performed by multiplying the image by a 3×3 diagonal matrix W constructed directly from the sampled illuminant ω .

Next, the white-balanced RAW image is transformed to the canonical CIE XYZ space using a full 3×3 matrix C called the color space transform (CST) matrix. The CST matrix is dependent on the illumination ω and is interpolated between two factory-calibrated color matrices specified in the DNG metadata [2]. We apply the matrices as:

$$\mathbf{p}_2 = w_\omega(\mathbf{p}_1) = C_\omega W_\omega \mathbf{p}_1. \quad (8)$$

Examples of our white balance sampling are provided in the second column of Fig. 2(B).

3.3. Color-manipulation module

Color and tone manipulations are typically performed using lookup tables (LUTs) on conventional ISPs [16]. Specifically, color manipulations to impart a desired photofinishing style are implemented using 3D LUTs. A 3D LUT approximates a nonlinear 3D color transform by sparsely sampling the transformation on a discrete 3D lattice, and acts as a global color operator that maps a source RGB color to a target RGB color as $\tilde{c}([r, g, b]) = [r', g', b']$ [20, 36]. However, LUTs are not differentiable; therefore, we approximate the LUT using a multi-layer perceptron (MLP) c , similar to [12], such that c is continuous and differentiable, and $c([r, g, b]) \approx \tilde{c}([r, g, b])$. We select a set of K 3D LUTs and train an MLP corresponding to each to obtain $\{c_\rho\}_{\rho=1}^K$. Note that the MLPs are learned offline and their weights are frozen during the training of the RAW reconstruction

network. We pick a random LUT for each mini-batch and apply the color mapping as:

$$\mathbf{p}_3 = c_\rho(\mathbf{p}_2), \quad \rho \sim \mathcal{U}\{1, \dots, K\}, \quad (9)$$

where \mathcal{U} denotes a uniform distribution. Representative examples are shown in the third column of Fig. 2(B).

3.4. Tone-mapping module

A tone curve is typically specified as a 1D LUT that is applied uniformly to all three color channels. We train an MLP ψ that approximates the Adobe tone curve [16], which we select as our default tone operator. We then introduce random perturbations around this tone curve by applying low-degree polynomials. Specifically, we generate random polynomials S with degree τ that are monotonically non-decreasing. Finally, we apply a fixed 3×3 matrix T to map from CIE XYZ space to linear-sRGB and a gamma function to obtain the final sRGB image.

$$\mathbf{p}_4 = t_\tau(\mathbf{p}_3) = (TS_\tau\psi(\mathbf{p}_3))^{\frac{1}{2.2}}, \quad \tau \sim \mathcal{U}\{1, \dots, d\}. \quad (10)$$

Examples of tone sampling can be seen in the last column of Fig. 2(B). Although we do not explicitly model local tone mapping for the sake of simplicity and speed, most RAW reconstruction methods adopt patch-based training, and our random sampling of the tone curves for individual patches can offer a degree of robustness to local tonal adjustments.

3.5. Combined loss

The combined loss is computed as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{RAW}}(\mathbf{x}, \hat{\mathbf{x}}) + \mathcal{L}_{\text{misc}} + \lambda \mathcal{L}_{\text{sRGB}}(\mathbf{z}, \hat{\mathbf{z}}), \quad (11)$$

where λ is a scalar weighting factor, and $\mathcal{L}_{\text{misc}}$ denotes any other additional losses employed by the selected RAW recovery method that do not directly target per-pixel RAW image reconstruction. For example, CAM[23] uses a super-pixel loss for content-aware sampling.

Our four tunable ISP modules provide a lightweight yet expressive model of real-world ISP variability, enabling our loss $\mathcal{L}_{\text{sRGB}}$ to deliver edit-aware supervision that generalizes across diverse rendering styles and editing workflows.

4. Results

We assess the effectiveness of our edit-aware loss on two representative RAW reconstruction methods from prior work—CAM [23] and RAW Diffusion [27]—to demonstrate its generality and impact. CAM serves as a metadata-assisted approach, while RAW Diffusion represents a blind model. Additionally, we designed a vanilla UNet-based [28] model to further evaluate our loss in a metadata-assisted setting.

We use the smartphone RAW dataset from [4] for training and evaluation. The dataset contains 3,224 images captured with the main rear camera of a Samsung S24 Ultra, encompassing a diverse range of scenes and lighting conditions. Each image has a resolution of 3000×4000 pixels. For training pairs, we use the provided denoised and demosaiced linear RAW images and the corresponding camera-ISP-rendered sRGB JPEGs. The dataset is split into 2,619 training, 205 validation, and 400 test images.

We fix the parameters defining the sampling distributions of our modular ISP across all experiments. Specifically, we set $\sigma = 0.75$, $M = 2,619$ (corresponding to all training images), $K = 15$, and $d = 5$. For the approximator MLPs, we follow the architecture and training protocol of [12]: the LUT MLP c uses three input and output features, and the tone MLP ψ uses one input and output feature. The edit-aware loss $\mathcal{L}_{\text{sRGB}}$ is computed using the standard ℓ_2 error.

For evaluation, reconstructed RAW images are saved in DNG format [1] and edited using Adobe Photoshop. Photoshop represents a full-featured, complex software ISP that is independent of our training pipeline. Our method does not assume access to or differentiability of the ISP used at inference. We evaluate using Photoshop to demonstrate robustness to unknown processing. We select edits that are diverse, realistic, and representative of common user adjustments. Specifically, we apply the five edits listed in Table 1 using Adobe Camera RAW, Photoshop’s plugin for RAW image editing. Each edit specifies the adjusted settings or parameters, and their values; a single edit may modify multiple parameters simultaneously. For example, ‘Presets’ in Table 1 adjust several ISP parameters at once to achieve a particular style. For consistency and reproducibility, the same fixed set of edits are applied to all test images.

We first evaluate our edit-aware loss using the metadata-assisted CAM [23] method. CAM employs a learned sampler at capture time to select a subset of RAW pixels. These RAW samples, saved as metadata, guide the decoder at inference to reconstruct the full RAW image from the sRGB input. Following the training procedure and hyperparam-

Name	Adobe Camera RAW settings
Edit 1	Adobe Camera RAW default rendering
Edit 2	Preset → Color → Bright
Edit 3	Preset → Creative → Flat and green Light → Exposure → +0.5
Edit 4	Preset → Creative → Warm contrast Light → Exposure → -0.25 Curve → Lights → +50
Edit 5	Preset → Creative → Cool matte Light → Exposure → +1.5 Curve → Darks → +50 Color → Temperature → 3500K

Table 1. A description of the five different edits applied at inference to the RAW DNGs. The settings are applied in Adobe Camera RAW, which is Adobe’s plugin for RAW image processing.

eters recommended by the authors, we train two separate models: a baseline model using the original CAM implementation and an edit-aware variant under identical settings, augmented with our loss (we set $\lambda = 2$ in Eq. 11). Evaluation is performed on the 400-image test split of the smartphone RAW dataset [4]. In Table 2, we report RAW reconstruction quality using PSNR (dB) and evaluate the sRGB renderings after Photoshop edits using PSNR, SSIM [32], and ΔE [29] (lower is better). Our method consistently improves all metrics across edits, with the largest gains—up to 2 dB—observed under more substantial edits, such as Edit 5. Interestingly, RAW PSNR also shows a slight improvement. Qualitative results are shown in Fig. 3, where the baseline CAM model exhibits banding artifacts and color distortions due to collapsed tones, whereas our method produces smoother, more accurate color and tonal responses. See Sec. S4 of supplementary material for additional evaluations of the CAM model on the NUS dataset [9].

Next, we perform a similar experiment using the RAW Diffusion [27] method, a blind reconstruction model that takes only the sRGB image as input. RAW Diffusion employs a diffusion-based framework with RGB-guided residual blocks to reconstruct high-fidelity RAW images. As before, we follow the training procedure and hyperparameters recommended by the authors, using a $\lambda = 4$ weighting factor for our edit-aware loss. Results are summarized in Table 2. Our method trades off a small amount of RAW fidelity but significantly improves sRGB reconstruction across different edits. Note that because RAW Diffusion is blind, overall performance is generally lower than metadata-assisted CAM, as expected. Qualitative results in Fig. 4 show that our method achieves more accurate color reproduction and closely matches the ground truth, while the baseline exhibits noticeable global color deviations.

Method	RAW	Edit 1			Edit 2			Edit 3			Edit 4			Edit 5		
	PSNR	PSNR	SSIM	ΔE	PSNR	SSIM	ΔE	PSNR	SSIM	ΔE	PSNR	SSIM	ΔE	PSNR	SSIM	ΔE
CAM	37.17	27.27	0.8644	7.12	26.95	0.8632	7.48	28.91	0.8951	5.91	27.70	0.8999	6.12	25.43	0.8456	8.00
CAM + Edit-aware	37.57	29.24	0.8860	5.20	28.97	0.8819	5.54	30.72	0.9068	4.33	29.23	0.9121	4.54	27.43	0.8637	5.98
RAWDiff	34.18	24.27	0.8584	9.50	23.96	0.8608	9.99	26.09	0.8965	8.25	25.33	0.9037	8.21	23.29	0.8484	9.91
RAWDiff + Edit-aware	33.37	25.44	0.8728	8.52	25.46	0.8730	8.74	27.17	0.9014	7.26	25.66	0.9019	7.74	25.03	0.8590	8.60
UNet	38.82	28.52	0.8693	5.89	28.17	0.8664	6.38	30.09	0.8956	4.90	28.73	0.9016	4.93	26.44	0.8463	6.88
UNet + Edit-aware	35.62	29.26	0.8832	5.30	29.01	0.8793	5.68	30.83	0.9032	4.36	29.32	0.9097	4.48	28.02	0.8635	5.75

Table 2. Quantitative evaluation of the proposed edit-aware loss integrated into different RAW reconstruction frameworks. Results are reported for CAM [23], RAW Diffusion [27], and a UNet-based [28] model, evaluated on 400 test images from the RAW smartphone dataset of [4]. We report PSNR (dB) for reconstructed RAWs, and PSNR (dB), SSIM, and ΔE (lower is better) for the corresponding edited sRGB renderings. Editing operations follow the settings described in Table 1. The **best** results are highlighted.

Method	EV +2			CCT 3000K			EV +2 & CCT 3000K		
	PSNR	SSIM	ΔE	PSNR	SSIM	ΔE	PSNR	SSIM	ΔE
UNet: FT image	30.26	0.8777	3.39	30.84	0.8859	3.79	30.12	0.8756	3.44
UNet + Edit-aware: FT image	31.31	0.8889	3.17	31.11	0.8898	3.73	31.09	0.8861	3.23
UNet + Edit-aware: FT image & FT edit	31.51	0.8894	3.04	31.20	0.8898	3.71	31.26	0.8870	3.18

Table 3. Quantitative results of fine-tuning (FT) a UNet-based [28] metadata-assisted RAW reconstruction model. Results are averaged over the 400 test images from the RAW smartphone dataset of [4]. Metrics include PSNR (dB), SSIM, and ΔE for the edited sRGB renderings. EV denotes exposure value and CCT represents correlated color temperature. The **best** and **second-best** results are highlighted.

Finally, we evaluate our method in a metadata-assisted setting using a standard UNet [28] architecture. The extra capture-time metadata is a downsampled version of the RAW image \mathbf{x}_d , and reconstruction is performed given the sRGB image and downsampled RAW: $\hat{\mathbf{x}} = f_\theta(\mathbf{y}, \mathbf{x}_d)$. We use an $8 \times$ downsampling factor along both axes, corresponding to $\approx 1.5\%$ metadata overhead in RAW samples. The UNet receives a 6-channel input: the first three channels are the sRGB image, and the last three are the downsampled RAW image bilinearly upsampled back to the original image resolution. The network outputs the reconstructed RAW image. The baseline model is trained using the RAW reconstruction loss \mathcal{L}_{RAW} , while for our method, we use only the edit-aware sRGB loss $\mathcal{L}_{\text{sRGB}}$, omitting \mathcal{L}_{RAW} as a test of whether our loss in isolation can guide edit-robust RAW reconstruction. Results in Table 2 show a tradeoff in RAW fidelity, as expected without direct RAW supervision, but our method produces significantly better reconstructions across edits. See Sec. S1 of supplementary material for training details and qualitative results.

4.1. Target edit-aware fine-tuning

Metadata-assisted methods offer the additional capability of model fine-tuning at inference time, typically using RAW samples stored as metadata, to adapt the model weights to a specific image. With our edit-aware loss, fine-tuning can be directed not only to the image, but also to the target edit. We demonstrate this ability of our method using the UNet model, where the fine-tuning loss is computed as $\mathcal{L}_{\text{sRGB-FT}}(\mathbf{z}_d, \hat{\mathbf{z}}_d) = \|\mathbf{z}_d - \hat{\mathbf{z}}_d\|_2^2$, where $\mathbf{z}_d = g_\phi(\mathbf{x}_d)$, $\hat{\mathbf{z}}_d = g_\phi(\hat{\mathbf{x}}_d)$, and $\hat{\mathbf{x}}_d$ is obtained by downsampling the

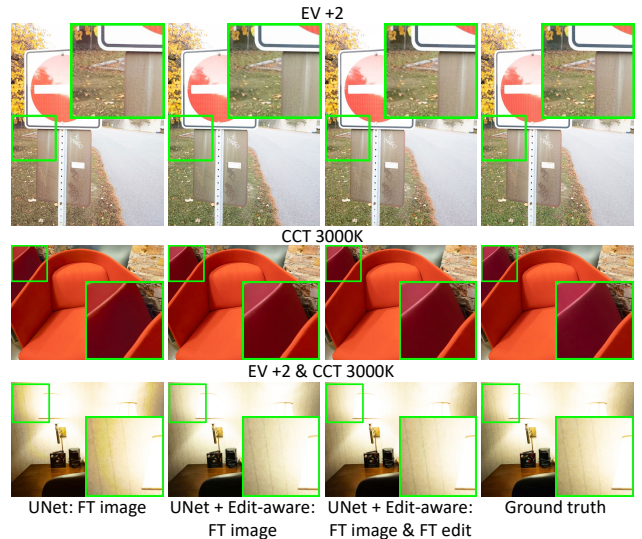


Figure 5. Results of fine-tuning (FT) a UNet-based [28] metadata-assisted RAW reconstruction model using our edit-aware loss. Our loss enables fine-tuning not only on the specific test image but also with respect to the target edit, yielding reconstructions that better align with the desired photofinishing outcome.

output of the network $\hat{\mathbf{x}}$. Importantly, during fine-tuning, the parameters ϕ can be fixed to match the target edit (e.g., $\varepsilon = 0.5$ for a +0.5 exposure edit) or randomly sampled as during training. We perform fine-tuning for 100 iterations using random 1024-pixel crops; \mathbf{z}_d is computed once.

We demonstrate this idea on exposure and white-balance edits, as the two most common and representative user adjustments. Table 3 shows three scenarios: the first row is

Model	Sampling ✓	No sampling ●			Disabled ✗	sRGB PSNR
	Exp.	WB	Color	Tone		
Exp. only	✓	✗	✗	✗	23.22	
WB only	✗	✓	✗	✗	22.35	
Color only	✗	✗	✓	✗	20.54	
Tone only	✗	✗	✗	✓	23.77	
Fixed pipeline	●	●	●	●	24.20	
Ours	✓	✓	✓	✓	25.15	

Table 4. Ablation study on the contribution of different loss modules and sampling configuration. Results are reported using the UNet reconstruction model on a subset of 50 challenging images from the test split of the RAW smartphone dataset of [4]. We report sRGB PSNR (dB) for Edit 5 (see Table 1 for edit details).

the baseline model fine-tuned for the test image, the second row is our model fine-tuned with ϕ randomly sampled, and the last row is our model fine-tuned with ϕ fixed to the target edit. Even with standard fine-tuning (first and second rows), our method maintains an advantage over the baseline model, and fixing ϕ to the target edit (last row) further improves performance. The qualitative results in Fig. 5 show that target edit-aware fine-tuning produces more accurate and visually consistent outputs under the selected edits.

4.2. Ablations

We perform ablations on the individual loss modules and the ISP parameter sampling to assess their contributions. Results are summarized in Table 4 using the UNet model, evaluated on a subset of 50 challenging images from the test split. We report sRGB PSNR for Edit 5 (see Table 1). The first four rows examine the contribution of each loss module. The fifth row tests a configuration using all modules but with parameter sampling disabled (i.e., a fixed ISP pipeline). This is equivalent to the cyclic consistency loss term used by forward-inverse methods [3, 33, 35] that enforces sRGB reconstruction from the recovered RAW *without* any edit sampling. As shown in the final row, our complete configuration, including all modules with randomized parameter sampling, yields the best performance.

4.3. User study

We conducted a blind, forced-choice perceptual study in which participants selected between the results with and without our edit loss that best matched the edited ground truth in terms of color and brightness while minimizing artifacts. The study used 20 images from the challenging 50-image set in Sec. 4.2. Each user rated all 20 images for the 3 methods in Table 2, with a random edit from Table 1 selected per image. We surveyed 25 users, yielding 1500 responses. Our method was preferred 83% of the time; preferences for CAM [23], RAW Diffusion [27], and UNet [28] were 77%, 88%, and 82%.

Method	CAM [23]		RAWDiff [27]		UNet [28]	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Model	22.11	0.6499	20.57	0.6927	23.30	0.6794
Model + EA	24.36	0.7087	21.59	0.7019	24.58	0.7199

Table 5. Results with manual image-specific edits applied to a subset of 50 challenging images from the RAW smartphone dataset of [4]. Our edit-aware method (EA) generalizes well to image-specific edits across various RAW reconstruction methods.

4.4. Image-specific edits

In Table 1, we used a fixed edit set for consistency and batch efficiency in Photoshop. Practical edits are typically image- and content-specific. We additionally conducted an experiment where manual image-specific edits were applied to the difficult 50-image subset in Sec. 4.2. PSNRs reported in Table 5 are lower due to more challenging images and edits but we observe the same improvement trend as with fixed edits, indicating good generalization to image-specific edits.

5. Discussion and conclusion

We presented an edit-aware loss framework that enhances the robustness of RAW reconstruction models under diverse rendering styles and post-capture edits. Unlike prior approaches that optimize solely for pixel-level RAW fidelity, our method explicitly aligns the training objective with downstream photographic editing by introducing a differentiable, modular, and tunable ISP that models realistic photofinishing variations. This plug-and-play loss formulation can be integrated into existing RAW reconstruction pipelines, improving edit fidelity and rendering consistency without any architectural changes.

While the inclusion of our edit-aware loss increases training time due to the added ISP sampling (by approximately 15% in our experiments), inference remains unchanged, allowing existing RAW reconstruction models to benefit from enhanced editability without impacting deployment efficiency. Although our framework introduces multiple ISP modules and diverse parameter sampling to cover a wide range of potential edits, the design remains highly flexible. For example, if specific target edit conditions are known *a priori*, the sampling can be restricted, enabling training for those specific types of edits while still leveraging the edit-aware supervision paradigm.

Extensive experiments demonstrate that our approach yields consistent gains across both metadata-assisted and blind reconstruction models, producing higher-quality sRGB renderings under a wide range of edits. Our method performs best for strong edits inducing large tonal, color, or contrast changes. Observed failures primarily occurred in extreme low-light scenes or uncommon lighting (e.g., green LEDs) not seen during training.

References

- [1] Adobe Digital Negative (DNG). <https://helpx.adobe.com/camera-raw/digital-negative.html>, . Accessed: 2025-11-04. 6
- [2] DNG Specification. https://helpx.adobe.com/content/dam/help/en/photoshop/pdf/dng_spec_1_6_0_0.pdf, . Accessed: 2025-11-04. 5
- [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*, 44(9):4688–4700, 2021. 3, 8
- [4] Mahmoud Afifi, Luxi Zhao, Abhijith Punnappurath, Mohamed A Abdelsalam, Ran Zhang, and Michael S Brown. Time-aware auto white balance in mobile photography. In *ICCV*, 2025. 6, 7, 8
- [5] Radu Berdan, Beril Besbinar, Christoph Reinders, Junji Otsuka, and Daisuke Iso. ReRAW: RGB-to-RAW image reconstruction via stratified sampling for efficient object detection on the edge. In *CVPR*, 2025. 2, 3
- [6] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T Barron. Unprocessing images for learned raw denoising. In *CVPR*, 2019. 3
- [7] Ayan Chakrabarti, Daniel Scharstein, and Todd E. Zickler. An empirical camera model for internet color vision. In *BMVC*, 2009. 3
- [8] Ayan Chakrabarti, Ying Xiong, Baochen Sun, Trevor Darrell, Daniel Scharstein, Todd Zickler, and Kate Saenko. Modeling radiometric uncertainty for vision with tone-mapped color images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(11):2185–2198, 2014. 3
- [9] Dongliang Cheng, Dilip K Prasad, and Michael S Brown. Illuminant estimation for color constancy: Why spatial-domain methods work and the role of the color distribution. *JOSA A*, 31(5):1049–1058, 2014. 6
- [10] Marcos V Conde, Steven McDonagh, Matteo Maggioni, Ales Leonardis, and Eduardo Pérez-Pellitero. Model-based image signal processors via learnable dictionaries. In *AAAI*, 2022. 3
- [11] Marcos V Conde, Radu Timofte, Yibin Huang, Jingyang Peng, Chang Chen, Cheng Li, Eduardo Pérez-Pellitero, Fenglong Song, Furui Bai, Shuai Liu, et al. Reversed image signal processing and RAW reconstruction. AIM 2022 challenge report. In *ECCV AIM Workshop*, 2022. 3
- [12] Marcos V Conde, Javier Vazquez-Corral, Michael S Brown, and Radu Timofte. NILUT: Conditional neural implicit 3D lookup tables for image enhancement. In *AAAI*, 2024. 5, 6
- [13] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *ACM SIGGRAPH*, 2008. 3
- [14] Mauricio Delbracio, Damien Kelly, Michael S Brown, and Peyman Milanfar. Mobile computational photography: A tour. *Annual review of vision science*, 7(1):571–604, 2021. 1, 2
- [15] Michael D. Grossberg and Shree K. Nayar. Determining the camera response from images: What is knowable? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(11):1455–1467, 2003. 3
- [16] Hakki Can Karaimer and Michael S Brown. A software platform for manipulating the camera imaging pipeline. In *ECCV*, 2016. 5
- [17] Seon Joo Kim, Hai Ting Lin, Zheng Lu, Sabine Süsstrunk, Stephen Lin, and Michael S. Brown. A new in-camera imaging model for color computer vision and its application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(12):2289–2302, 2012. 2, 3
- [18] Woohyeok Kim, Geonu Kim, Junyong Lee, Seungyong Lee, Seung-Hwan Baek, and Sunghyun Cho. ParamISP: Learned forward and inverse ISPs using camera parameters. In *CVPR*, 2024. 3
- [19] Leyi Li, Huijie Qiao, Qi Ye, and Qinmin Yang. Metadata-based RAW reconstruction via implicit neural functions. In *CVPR*, 2023. 2
- [20] Hai Ting Lin, Zheng Lu, Seon Joo Kim, and Michael S Brown. Nonuniform lattice regression for modeling the camera imaging pipeline. In *ECCV*, 2012. 5
- [21] Tomoo Mitsunaga and Shree K. Nayar. Radiometric self calibration. In *CVPR*, 1999. 3
- [22] Seonghyeon Nam and Seon Joo Kim. Modelling the scene dependent imaging in cameras with a deep neural network. In *ICCV*, 2017. 3
- [23] Seonghyeon Nam, Abhijith Punnappurath, Marcus A Brubaker, and Michael S Brown. Learning sRGB-to-raw-RGB de-rendering with content-aware metadata. In *CVPR*, 2022. 2, 3, 4, 5, 6, 7, 8
- [24] Rang MH Nguyen and Michael S Brown. RAW image reconstruction using a self-contained sRGB-JPEG image with only 64 KB overhead. In *CVPR*, 2016. 2
- [25] Rang MH Nguyen and Michael S Brown. RAW image reconstruction using a self-contained sRGB-JPEG image with small memory overhead. *International Journal of Computer Vision*, 126(6):637–650, 2018. 2
- [26] Abhijith Punnappurath and Michael S Brown. Spatially aware metadata for raw reconstruction. In *WACV*, 2021. 2
- [27] Christoph Reinders, Radu Berdan, Beril Besbinar, Junji Otsuka, and Daisuke Iso. RAW-diffusion: RGB-guided diffusion models for high-fidelity RAW image generation. In *WACV*, 2025. 1, 2, 3, 5, 6, 7, 8
- [28] O. Ronneberger, P.Fischer, and T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015. 6, 7, 8
- [29] Gaurav Sharma and Raja Bala. *Digital Color Imaging Handbook*. CRC Press, 2nd edition, 2013. 6
- [30] 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 *CVPR*, 2023. 2, 3
- [31] Yufei Wang, Yi Yu, Wenhan Yang, Lanqing Guo, Lap-Pui Chau, Alex C Kot, and Bihan Wen. Beyond learned metadata-based raw image reconstruction. *International Journal of Computer Vision*, 132(12):5514–5533, 2024. 3
- [32] Zhou Wang, Alan Bovik, Hamid Sheikh, and Eero Simoncelli. Image quality assessment: From error visibility to

- structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. [6](#)
- [33] Yazhou Xing, Zian Qian, and Qifeng Chen. Invertible image signal processing. In *CVPR*, 2021. [2](#), [3](#), [8](#)
- [34] Lu Yuan and Jian Sun. High quality image reconstruction from RAW and JPEG image pair. In *ICCV*, 2011. [2](#)
- [35] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. CycleISP: Real image restoration via improved data synthesis. In *CVPR*, 2020. [2](#), [3](#), [8](#)
- [36] Hui Zeng, Jianrui Cai, Lida Li, Zisheng Cao, and Lei Zhang. Learning image-adaptive 3D lookup tables for high performance photo enhancement in real-time. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(4):2058–2073, 2020. [5](#)