

Experts-Guided Unbalanced Optimal Transport for ISP Learning from Unpaired and/or Paired Data

Georgy Perevozchikov¹ Nancy Mehta^{1*} Egor Ershov^{2,3,4} Radu Timofte¹

¹Computer Vision Lab, CAIDAS & IFI, University of Würzburg

²Institute for Information Transmission Problems, RAS

³Moscow Institute of Physics and Technologies

⁴Artificial Intelligence Research Institute

{georgii.perevozchikov,nancy.mehta,radu.timofte}@uni-wuerzburg.de ershov@iitp.ru

Abstract

Learned Image Signal Processing (ISP) pipelines offer powerful end-to-end performance but are critically dependent on large-scale paired raw-to-sRGB datasets. This reliance on costly-to-acquire paired data remains a significant bottleneck. To address this challenge, we introduce a novel, unsupervised training framework based on Optimal Transport capable of training arbitrary ISP architectures in both unpaired and paired modes. We are the first to successfully apply Unbalanced Optimal Transport (UOT) for this complex, cross-domain translation task. Our UOT-based framework provides robustness to outliers in the target sRGB data, allowing it to discount atypical samples that would be prohibitively costly to map. A key component of our framework is a novel “committee of expert discriminators,” a hybrid adversarial regularizer. This committee guides the optimal transport mapping by providing specialized, targeted gradients to correct specific ISP failure modes, including color fidelity, structural artifacts, and frequency-domain realism. To demonstrate the superiority of our approach, we retrained existing state-of-the-art ISP architectures using our paired and unpaired setups. Our experiments show that while our framework, when trained in paired mode, exceeds the performance of the original paired methods across all metrics, our unpaired mode concurrently achieves quantitative and qualitative performance that rivals, and in some cases surpasses, the original paired-trained counterparts. The code and pre-trained models are available at: <https://github.com/gosha20777/EGUOT-ISP.git>.

1. Introduction

Modern smartphone photography is a triumph of computational imaging, relying on a complex Image Signal Processor (ISP) [24, 33, 46] to overcome the physical limita-

*Corresponding Author




Figure 1. A conceptual overview of our framework’s robustness to outliers. *Top*: Standard GANs (RawFormer [63]), when trained on unpaired datasets, are highly sensitive to data outliers (e.g., overexposed photos). This leads to training instability and a collapsed output. *Bottom*: Our proposed EGUOT framework, leveraging Experts Guided Unbalanced Optimal Transport, is inherently robust to such data outliers and focuses on the high-quality core of the target distribution to produce perceptually better results.

tions of small sensors and compact optics [36, 77]. The ISP converts noisy, monochromatic, camera-specific raw sensor data into the perceptually-pleasing sRGB images users expect [24, 31]. While traditional hardware ISPs use fixed, hand-tuned modules [12], recent work has shifted to learned end-to-end pipelines [37, 78, 84]. These deep learning models offer superior performance and flexibility, bridging the quality gap between mobile phones and high-end DSLR cameras [37, 38, 41, 42, 66].

Despite their success, these state-of-the-art learned ISPs [37, 38, 40–42, 66, 76, 84] are critically dependent on large-scale, paired training datasets [38, 42, 71]. This requires a difficult and costly data acquisition process: capturing thousands of scenes simultaneously with a smartphone (raw)

and a high-end reference DSLR (sRGB), followed by meticulous alignment and manual photo finishing [27, 38, 42]. This dependency forms a significant bottleneck, as the entire data collection and alignment effort must be repeated for every new camera sensor [59, 63].

The obvious solution is unpaired training to break this reliance on paired data. Recent attempts [6, 60, 63] have explored this using standard Generative Adversarial Network (GAN) frameworks [61, 74, 87], but these approaches often struggle with training stability and produce sub-optimal results [60]. We assume that a core reason for this failure is the presence of distributional outliers (e.g., blurry, over/under-exposed images) within the raw-to-sRGB datasets [38, 42, 71]. Standard adversarial losses, including those based on classical Optimal Transport (OT) [30, 45, 51], are highly sensitive to these outliers [17], as model performance is degraded by attempting to learn from such suboptimal data.

To overcome these challenges, we introduce **EGUOT** (**Experts-Guided Unbalanced Optimal Transport**), a novel training framework robust to dataset outliers (Fig. 1) that operates in *paired* and *unpaired* modes. Our framework is the first to apply Unbalanced Optimal Transport (UOT) [64] for the raw-to-sRGB translation task. By relaxing classical OT constraints [45, 51], our method is inherently robust to outliers, discounting low-quality data to focus on the high-quality core of the target distribution. While UOT provides robustness [17, 64], it is insufficient for ISP high-fidelity demands. We guide the transport mapping with a central component: a Committee of Expert Discriminators. This novel hybrid regularizer provides targeted, multi-modal gradients to correct specific ISP failure modes, such as semantic color inaccuracies, lack of structural detail, and unrealistic frequency artifacts.

We demonstrate the power and architecture-agnostic nature of our framework. We use our unsupervised method to retrain existing paired state-of-the-art supervised ISP backbones (e.g., MW-ISPNet [37], Restormer [82]). Experiments confirm our EGUOT framework’s flexibility. When trained without paired data, EGUOT achieves a new state-of-the-art for unpaired translation, handily beating prior work [6, 63]. Its performance also proves competitive with, and at times better than, the original fully-paired models. Additionally, in a paired-data setting, our method establishes a new performance ceiling, uniformly improving upon the original models’ results across all metrics.

Contribution

Our contributions can be summarized as follows:

- We propose a novel framework for raw-to-sRGB translation, trainable in both *unpaired* and *paired* modes. To our knowledge, we are the first to successfully apply Unbalanced Optimal Transport (UOT) [17] to this task, leveraging its robustness to dataset outliers, a critical flaw in

prior adversarial approaches.

- We introduce a hybrid regularization method, the *Committee of Expert Discriminators*, which guides the UOT mapping. This committee provides targeted, multi-modal perceptual gradients (color, structure, frequency), working in synergy with the UOT loss to ensure a high-fidelity image reconstruction rich in perceptual detail.
- We demonstrate that our *architecture-agnostic framework* can train existing SOTA paired ISP backbones. In *unpaired* mode, our models achieve performance competitive with, or even surpassing, their original, fully-paired results. In *paired* mode, our framework surpasses these original results on standard benchmarks.

2. Related Work

Our work is positioned at the intersection of three research domains: (1) learned ISP pipelines, (2) unpaired image-to-image translation, and (3) generative modeling with Optimal Transport.

2.1. Learned Image Signal Processing

The camera Image Signal Processor (ISP) [15, 24, 32, 33, 39, 43, 46, 78, 81, 84] solves a complex image-to-image translation task, converting noisy, sensor-specific raw data into a plausible sRGB image [24, 31]. While traditional ISPs rely on a cascade of hand-tuned hardware modules [1, 9, 12, 20, 33, 34, 44, 46, 72], this paradigm has recently been challenged by recent end-to-end deep learning approaches [37, 38, 41, 42, 58, 66, 78, 84, 86].

Foundational works (DeepISP [69], PyNET [38]) demonstrated that a U-Net-like network [67] could learn the entire ISP mapping and subsequent research improved their performance. MW-ISPNet [37] and AW-Net [23] introduced Discrete Wavelet Transforms (DWT) to preserve high-frequency details. LAN-ISP [76] and [41, 42, 84] integrated attention mechanisms for improved spatial processing. More recently, Transformer-based [26] backbones like Restormer [82], ISPFormer [66], and [10, 32, 35, 86] were adapted for this task, leveraging powerful long-range dependency modeling.

Concurrently, specialized networks were developed for specific ISP sub-problems [3, 5, 13, 20, 21, 70], including raw-to-raw translation [2, 63] for generalizing ISPs to new camera sensors, and color mapping (e.g., cmKAN [60], Ni-LUT [22], SepLUT [79], and [4, 11, 14, 25, 47, 53, 54]) for high-fidelity color matching.

Despite architectural diversity and strong performance, these state-of-the-art methods are all trained in a fully paired manner, fundamentally depending on large-scale, paired raw-to-sRGB datasets [8, 27, 38, 42, 71]. This data is notoriously difficult and costly to acquire, requiring meticulous spatial and temporal alignment between a source (e.g., smartphone) and a target (e.g., DSLR) camera. This paired

data bottleneck [2, 6, 60, 63] is the field’s primary limitation, as the entire data collection process must be repeated for every new camera sensor. This limitation motivates our work in developing a robust unpaired framework.

2.2. Unpaired Image-to-Image Translation

The most common approach for unpaired image-to-image translation is the Generative Adversarial Network (GAN) [61]. CycleGAN [87] introduced cycle-consistency, using a pair of generators and discriminators to enforce that a translated image can map back to its original domain. This inspired a wide range of methods based on shared latent spaces (e.g., DualGAN [80], UNIT [55], STARGAN [18], SEAN [88]), adaptive normalization (e.g., U-GAT-IT [48], UVCGANv2 [74]), contrastive learning (e.g., CUT [62]), and recent [16, 19, 73, 85].

Last years, this unpaired approach was explored for the raw-to-sRGB task [6], proposing the usage of a standard GAN [61] with a multi-term loss, including pre-trained networks (VGG [52], LPIPS [83]) for content preservation. While pioneering, this approach relies on standard adversarial losses [61], not designed to handle distributional outliers in real-world datasets [17, 60]. We hypothesize that standard GAN losses are negatively impacted by these outliers, leading to sub-optimal results quality, thus necessitating a more robust adversarial framework.

2.3. Optimal Transport in Generative Modeling

Optimal Transport (OT) [45] offers a powerful theoretically-grounded alternative to traditional GAN losses. While many methods use the OT cost as a loss function (e.g., Wasserstein GANs [7, 49]), recent works have shown the OT plan itself can serve as a generative model [28–30, 50, 51, 57, 65, 75].

However, a well-known limitation of classical OT is its high sensitivity to outliers [17, 51, 64]. The standard OT problem requires transporting every point in the source distribution to the target distribution. This means a few outlier samples can unduly influence the data [17], drastically altering the transport map and degrading the final result.

To solve this, the theory of *Unbalanced Optimal Transport* (UOT) was proposed [64], relaxing the hard marginal constraints of classical OT [45, 51]. This relaxation permits partial transport, which, in practical terms, means the model can learn to *ignore or discard outliers* rather than being forced to match them.

While the UOT theory was powerful, a stable, practical algorithm was needed. The UOTM paper [17] introduced a novel, stable training framework based on the semi-dual formulation of UOT [64]. The authors demonstrated that UOTM is significantly more robust to outliers and converges faster than classical OT models [45, 51].

Our work is the first to connect these fields. We are the

first to apply the Experts Guided UOT framework to the paired and unpaired ISP problem. Thus, we leverage the robustness of UOT [17] to solve a task where prior methods struggle precisely because of the outlier sensitivity that UOT was designed to fix [17].

3. Method

Our goal is to train a raw-to-sRGB mapping network T_θ for *paired* and *unpaired* settings, which we achieve with EGUOT, a flexible training framework (Fig. 2). The framework consists of a generator T_θ trained by a novel hybrid loss system. In unpaired mode, this system uses a cost function $c(x, T_\theta(x))$ for robust content mapping and a Committee of Expert Discriminators D_ψ for perceptual fidelity. In paired mode, the system has a direct pixel loss $c(y, T_\theta(x))$, retaining the expert committee.

3.1. Preliminaries: UOT for Image Translation

Inspired by UOTM [17], which learns a transport T_θ and potential P_ω network via a minimax problem, we utilize its core components:

- T_θ : The transport network that learns the transport map.
- P_ω : The potential network that acts as the discriminator.
- $c(x, y)$: A function measuring the x to y mapping cost.
- Φ_1, Φ_2 : A non-decreasing, convex function pair (e.g., e^x) applying a soft penalty that relaxes transport constraints, enabling outlier robustness.

The objective for the potential network P_ω , augmented with a stabilizing R_1 gradient penalty [68], is:

$$\begin{aligned} \mathcal{L}_P = & \mathbb{E}_{x \sim \mathbb{X}}[\Phi_1(-c(x, T_\theta(x)) + P_\omega(T_\theta(x)))] \\ & + \mathbb{E}_{y \sim \mathbb{Y}}[\Phi_2(-P_\omega(y))] + \frac{\gamma}{2} \mathbb{E}_{y \sim \mathbb{Y}}[\|\nabla_y P_\omega(y)\|_2^2] \end{aligned} \quad (1)$$

The objective for the transport network T_θ is:

$$\mathcal{L}_T = \mathbb{E}_{x \sim \mathbb{X}}[c(x, T_\theta(x)) - P_\omega(T_\theta(x))] \quad (2)$$

UOTM [17] application to the raw-to-sRGB task requires two considerations: (1) a valid cost function $c(\cdot, \cdot)$ must be defined for the paired/unpaired case, as raw and sRGB domains are incompatible; and (2) the basic transport loss (Eq. 2) is insufficient alone for the high-fidelity translation quality. Our EGUOT framework addresses both.

3.2. The EGUOT Framework

We propose the EGUOT framework to solve these two challenges. First, we define the cost functions to utilize Eq. (1) and (2). Second, we introduce a hybrid loss function that regularizes the objective (1) with a committee of expert discriminators.




Figure 2. **An overview of our EGUOT training framework.** (a) *The main pipeline*, where the generator (T_θ) is trained by three parallel objectives: (1) a content-preserving cost function ($c(\cdot, \cdot)$) for paired and unpaired mode, (2) a robust Potential network (P_ω), and (3) an Experts committee (D_ψ). (b) *The detailed breakdown of the Experts committee*, showing its three specialized discriminators for *Color*, *Structure*, and *Frequency*, each analyzing different features of the real (Y) and predicted ($\hat{Y}$) sRGB images.

3.2.1. Cost Functions

We define the transport cost $c(\cdot, \cdot)$ for paired and unpaired training.

Unpaired Setting: The cost $c(x, T_\theta(x))$ is computed between the input x (4-channel raw) and output $T_\theta(x)$ (3-channel sRGB). To resolve this domain incompatibility, we use a fixed pre-processing function P (e.g., basic demosaicing) mapping the raw image $x \in \mathbb{X}$ to a simple RGB image $P(x) \in \mathcal{X}$. The unpaired cost is the L_2 distance:

$$c(x, T_\theta(x)) = \tau \|T_\theta(x) - P(x)\|_2^2$$

Paired Setting: With access to the ground-truth sRGB image y , we use L_1 norm as a stronger paired cost.

3.2.2. Experts-Guided Hybrid Loss

The UOTM [17] objective, while robust, is insufficient to enforce the high-quality raw-to-sRGB mapping.

To address this, we introduce a committee of expert discriminators $D_\psi = \{D_i\}_{i=1}^M$, which runs in parallel to the Potential Network P_ω providing targeted adversarial gradients. We use a Hinge loss $h(t) = \max(0, 1 + t)$ for this committee.

The expert discriminators D_ψ are trained to distinguish real sRGB images from fake ones by minimizing $\mathcal{L}_D$:

$$\mathcal{L}_D = \sum_{D_i \in D_\psi} (\mathbb{E}_{y \sim \mathbb{Y}}[h(-D_i(y))] + \mathbb{E}_{x \sim \mathbb{X}}[h(D_i(T_\theta(x)))])$$

The generator T_θ is then simultaneously trained to fool this committee by minimizing the expert generator loss $\mathcal{L}_{T_exp}$:

$$\mathcal{L}_{T_exp} = \sum_{D_i \in D_\psi} \mathbb{E}_{x \sim \mathbb{X}}[-D_i(T_\theta(x))]$$

The final, total loss for our generator is:

$$\mathcal{L}_{T_total} = \mathcal{L}_T + \lambda \cdot \mathcal{L}_{T_exp}$$

3.3. Network Architectures

We selected SOTA components for our approach.

Transport Network (T_θ): A learnable ISP pipeline for the complete raw-to-sRGB mapping. As our framework is architecture-agnostic, we adopt multiple SOTA ISP backbones (e.g., [42, 60]) to serve as T_θ in our experiments.

Potential Network (P_ω): This network serves as the potential function [51]. We use a ConvNeXt [56] backbone, whose features are passed through an MLP head to output a single scalar.

Experts Committee (D_ψ): The committee is comprised of three specialist networks:

- **Color Expert:** A frozen ConvNeXt [56] encoder from a pre-trained colorization U-Net [67] that judges semantic color plausibility in the L^* channel.
- **Structure Expert:** A multi-scale PatchGAN discriminator [63] operating on the sRGB image to enforce local texture fidelity.
- **Frequency Expert:** A lightweight CNN operating on the 2D log-magnitude FFT spectrum of the grayscale image to penalize high-frequency artifacts.

3.4. Training Algorithm

We train our framework with an alternating optimization scheme. We use a discriminator-heavy $N : 1$ training ratio: both the Potential Network P_ω and Experts Committee D_ψ are updated N steps for every single Transport Network T_θ update. The complete EGUOT training procedure is in Algorithm 1.

4. Experiments

To validate our EGUOT framework, we conduct a comprehensive set of experiments. We demonstrate that (1) our framework can train existing supervised ISP backbones in paired and unpaired modes to outperform (in the paired

case) or achieve performance competitive with (in the unpaired case) their original settings; (2) our framework sets a new state-of-the-art for unpaired raw-to-sRGB translation; (3) the core components of our method—the UOT objective and the expert committee—are both essential for its success.

Algorithm 1 ISP with EGUOT

Input: The source $\mathbb{X}$ and the target $\mathbb{Y}$ distributions; transport $T_\theta : \mathbb{R}^X \rightarrow \mathbb{R}^Y$; potential $P_\omega : \mathbb{R}^Y \rightarrow \mathbb{R}$; experts $D_\psi = \{D_i : \mathbb{R}^Y \rightarrow \mathbb{R}\}_{i=1}^M$; experts loss $h : \mathbb{R} \rightarrow \mathbb{R}$; cost $c : \mathbb{Y} \times \mathbb{Y} \rightarrow \mathbb{R}$ for *paired* or $c : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$ for *unpaired* mode; non-decreasing function pair (Φ_1, Φ_2) ; number of inner iterations N ; hyperparameters γ, λ .

Output: learned transport network T_θ .

```

1: repeat                                ▷ Discriminator/Potential training loop
2:   for  $k = 1, \dots, N$  do
3:     Sample batch  $X \sim \mathbb{X}, Y \sim \mathbb{Y}$ ;
4:      $\mathcal{L}_P \leftarrow \frac{1}{|X|} \sum_{x \in X} \Phi_1(-c(\cdot, T_\theta(x)) + P_\omega(T_\theta(x)))$ 
         $+ \frac{1}{|Y|} \sum_{y \in Y} \Phi_2(-P_\omega(y))$ 
         $+ \frac{\gamma}{2} \cdot \frac{1}{|Y|} \sum_{y \in Y} \|\nabla_y P_\omega(y)\|_2^2$ ;
5:     Update  $\omega$  by using  $\nabla_\omega \mathcal{L}_P$ ;
6:      $\mathcal{L}_D \leftarrow \sum_{D_i \in D_\psi} \left( \frac{1}{|Y|} \sum_{y \in Y} h(-D_i(y)) \right.$ 
         $\left. + \frac{1}{|X|} \sum_{x \in X} h(D_i(T_\theta(x))) \right)$ ;
7:     Update  $\psi$  by using  $\nabla_\psi \mathcal{L}_D$ ;
8:   end for
9:   Sample batch  $X \sim \mathbb{X}$ ;                ▷ Transport loss
10:   $\mathcal{L}_T \leftarrow \frac{1}{|X|} \sum_{x \in X} \left( \underbrace{c(\cdot, T_\theta(x))}_{\mathcal{L}_{T\_cost}} - \underbrace{P_\omega(T_\theta(x))}_{\mathcal{L}_{T\_pot}} \right) +$ 
     $\lambda \cdot \underbrace{\sum_{D_i \in D_\psi} \frac{1}{|X|} \sum_{x \in X} [-D_i(T_\theta(x))]}_{\mathcal{L}_{T\_exp}}$ ;
11:  Update  $\theta$  by using  $\nabla_\theta \mathcal{L}_T$ ;
12: until not converged;

```

4.1. Datasets

Zurich raw-to-sRGB (ZRR) dataset [38] This dataset consists of 48,043 paired raw (Huawei P20) and sRGB (Canon 5D Mark IV) image patches (448×448). We follow the official dataset splits for training and testing.

ISPIW dataset [71] This dataset comprises raw images from a Huawei Mate 30 Pro and corresponding sRGB references from a Canon 5D Mark IV. We use all 192 full-resolution images, which are split into 5,724 patches, using a 70%/20%/10% split for training, validation, and testing.

Mobile AI dataset [37] This dataset was acquired using a Sony IMX586 Quad Bayer sensor and a Fujifilm GFX100 DSLR. The dataset contains 95K cropped 256×256 image pairs. We again follow the standard 70%/20%/10% split.

4.2. Training Details

All experiments were conducted on a single NVIDIA GeForce RTX 4090 GPU. All networks (Transport T_θ , Potential P_ω , and Expert Committee D_ψ) were trained using the Adam optimizer with $\beta_1 = 0.5$ and $\beta_2 = 0.999$. We used a Cosine Annealing scheduler for all learning rates. The learning rates for the transport, potential, and expert networks were set to $lr_{\text{transport}} = 1e-4$, $lr_{\text{potential}} = 2e-4$, and $lr_{\text{experts}} = 2e-4$, respectively. The R_1 gradient penalty coefficient was $\gamma = 1.5$. The UOT content cost weight was $\tau = 1e-3$, and the expert committee weight was $\lambda = 1.0$. All models were trained for 300 epochs. For the final model weights, we apply Stochastic Weight Averaging (SWA) for the last 10 epochs. For all unpaired training tasks, the source and target datasets were shuffled independently.

4.3. Results

Our primary results demonstrate EGUOT’s performance and flexibility. First, EGUOT can train diverse SOTA ISP architectures in an unpaired manner, achieving competitive performance, while our paired mode achieves superior performance. Finally, our framework sets a new state-of-the-art for unpaired raw-to-sRGB translation. *See supp. materials for additional results.*

ZRR Results

We begin our analysis with the ZRR dataset. Table 1 and Figure 3 show the results of applying our framework (paired and unpaired) to nine SOTA ISP backbones. The results are compelling: (1) *our paired* mode consistently surpasses the original paired-trained models across all metrics and architectures, showing the advantage of our expert-guided regularization; (2) *our unpaired* mode consistently rivals its paired counterparts. In particular, our EGUOT-trained models even *surpass* the original results in several key metrics, such as SSIM for Restormer [82] and MW-ISP [37], and ΔE for LAN-ISP [76] and AW-Net [23]. This demonstrates our framework is truly architecture-agnostic and can effectively replace the need for paired data.

Next, we establish a new state-of-the-art for unpaired translation on this dataset. Table 2 compares unpaired training frameworks using a fixed Restormer backbone. Our EGUOT framework significantly outperforms all other methods. It surpasses standard GANs (CycleGAN [87], WGAN [7]) by a massive margin and clearly outperforms recent specialized methods like RawFormer [63] and LLUIISP [6]. Crucially, EGUOT also outperforms other OT-based methods (NOT [51], UOTM [17]), demonstrating




Figure 3. Our framework’s unpaired mode achieves the paired performance across diverse architectures on the Zurich raw-to-sRGB dataset [38]. The top row (a) shows the results of SOTA ISP backbones (e.g., Restormer [63], cmKAN [60]) trained with their *original paired* settings. The bottom row (b) — results of the same backbones trained in *unpaired mode*, highlighting the effectiveness of our method. Best viewed in the electronic version.

Table 1. Results on the Zurich raw-to-sRGB dataset [38], comparing SOTA backbones retrained with our EGUOT framework against their *original paired-data results*. Our *paired* mode consistently outperforms the original, while our *unpaired* mode achieves competitive performance. Better results are highlighted in yellow.

Backbone	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$	#Params	FLOPs	Architecture Type
	Paired Mode (original)				Paired Mode (ours)				Unpaired Mode (ours)						
Restormer [82]	21.51	0.83	9.71	0.19	21.73	0.84	9.42	0.18	21.69	0.84	9.93	0.20	5.03M	583G	U-Net-like
MW-ISP [37]	21.88	0.82	10.33	0.21	22.15	0.84	9.61	0.20	21.62	0.83	9.97	0.20	29.22M	3.6T	U-Net-like
AW-Net [23]	21.58	0.75	11.02	0.19	22.02	0.77	9.98	0.18	21.71	0.76	10.17	0.19	54.83M	3.1T	U-Net-like
LAN-ISP [76]	20.76	0.81	10.41	0.22	21.19	0.81	9.36	0.21	20.50	0.80	9.88	0.22	49.2K	56.9G	U-Net-like
LiteISP [84]	22.18	0.83	10.28	0.19	22.29	0.84	9.79	0.19	21.43	0.83	10.04	0.21	9.04M	174G	U-Net-like
DH-ISP [40]	19.68	0.72	11.46	0.26	19.88	0.74	10.99	0.24	19.69	0.74	11.25	0.25	3.1K	31.7G	Forward CNN
MicroISP [41]	20.30	0.78	11.14	0.24	20.61	0.79	11.08	0.24	19.81	0.79	11.65	0.24	105K	37G	Forward CNN
NVNet [42]	20.95	0.80	10.96	0.22	21.13	0.80	10.05	0.21	20.15	0.78	10.31	0.23	3.5K	31.8G	ResNet-like
cmKAN [60]	24.41	0.85	7.27	0.17	24.63	0.86	7.10	0.17	24.46	0.84	8.31	0.18	76.4K	40G	Hyper-network




Figure 4. Quantitative results of *unpaired frameworks* on the Zurich raw-to-sRGB dataset [38]. We compare against modern GAN-based (e.g., RawFormer [63], UVCANv2 [73]), other OT-based (e.g., NOT [51], UOTM [17]), and recent unpaired ISP methods (LLUIISP [6]). Our framework sets a *new state-of-the-art performance across all metrics*. Best viewed in the electronic version.

our novel experts-guided hybrid loss is critical for achiev-

Table 2. Results of various unpaired image-to-image translation approaches on the Zurich raw-to-sRGB dataset [38]. We retrain all methods on a fixed Restormer backbone for a fair comparison. Our EGUOT framework achieves the best performance.

Method	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$	Type
CycleGAN [87]	7.73	0.32	27.8	0.82	GAN-based
UVCANv2 [74]	16.25	0.61	12.2	0.49	GAN-based
RawFormer [63]	17.09	0.68	11.9	0.45	GAN-based
WGAN [7]	14.13	0.52	14.3	0.67	GAN-based
W2GAN [49]	15.62	0.58	13.9	0.60	GAN-based
LLUIISP [6]	19.07	0.72	11.7	0.32	GAN-based
NOT [51]	16.03	0.60	12.6	0.54	OT-based
UOTM [17]	18.96	0.71	11.0	0.31	OT-based
LOT [65]	17.77	0.68	12.2	0.39	OT-based
Ours	21.69	0.84	9.93	0.20	OT-based

ing high-fidelity results. Figure 4 qualitatively shows our method’s superior color and texture fidelity.




Figure 5. Quantitative results on the on the ISPIW raw-to-sRGB dataset [71], demonstrating the architecture-agnostic nature of our EGUOT framework. We apply (a) the original, *paired* training and (b) our proposed *unpaired* training to a wide range of SOTA ISP backbones. Our unpaired mode achieves performance that is highly competitive with the original settings. Best viewed in the electronic version.

Table 3. Results on the ISPIW raw-to-sRGB dataset [71], comparing SOTA backbones retrained with our framework against their *original paired-data results*. Our *unpaired* mode achieves competitive performance. Better results are highlighted in yellow.

Backbone	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$
	Paired Mode (original)				Unpaired Mode (ours)			
Restormer [82]	20.93	0.79	6.85	0.14	20.81	0.79	6.99	0.14
MW-ISP [37]	21.90	0.81	7.03	0.11	21.81	0.81	6.99	0.11
AW-Net [23]	21.75	0.81	6.99	0.12	21.93	0.81	7.30	0.11
LAN-ISP [76]	22.09	0.81	6.97	0.11	21.95	0.81	6.42	0.11
LiteISP [84]	22.14	0.81	6.31	0.11	22.01	0.80	6.88	0.12
DH-ISP [40]	19.92	0.76	7.89	0.17	19.59	0.77	7.40	0.17
MicroISP [41]	20.70	0.77	6.92	0.15	20.79	0.77	6.33	0.15
NVNet [42]	23.80	0.82	5.81	0.10	23.39	0.82	6.08	0.10
cmKAN [60]	24.22	0.83	5.29	0.09	24.21	0.83	5.73	0.09

ISPIW Results

We confirm these findings with the ISPIW dataset. As shown in Table 3 and Figure 5, our unpaired training again achieves performance that is statistically on-par with the original paired training for nearly all architectures. For several models (e.g., AW-Net [23], MicroISP [41]), our unpaired method even achieves a higher PSNR or lower ΔE than the paired-data baseline, further highlighting the framework’s effectiveness.

Mobile AI Results

Finally, we validate our approach on the Mobile AI dataset [42] in Table 4 and Figure 6. The results are consistent with our findings on the other two datasets. Our unpaired mode remains highly competitive with paired settings, and in the case of MW-ISP [37], even outperforms the original model in both PSNR (24.44 vs 24.31) and ΔE (6.09 vs 6.27).




Figure 6. Quantitative results on the on the Mobile AI raw-to-sRGB dataset [42], demonstrating the architecture-agnostic nature of our EGUOT framework. We apply (a) the original, *paired* training and (b) our proposed *unpaired* training to a wide range of SOTA ISP backbones. Our unpaired mode achieves performance that is highly competitive with the original settings. Best viewed in the electronic version.

Table 4. Results on the Mobile AI raw-to-sRGB dataset [42], comparing SOTA backbones retrained with our framework against their *original paired-data results*. Our *unpaired* mode achieves competitive performance. Better results are highlighted in yellow.

Backbone	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$
	Paired Mode (original)				Unpaired Mode (ours)			
Restormer [82]	23.99	0.88	6.74	0.13	22.51	0.86	7.27	0.14
MW-ISP [37]	24.31	0.88	6.27	0.11	24.44	0.87	6.09	0.11
AW-Net [23]	24.17	0.87	6.35	0.12	24.06	0.87	6.33	0.12
LAN-ISP [76]	23.48	0.87	7.18	0.15	23.15	0.86	7.22	0.14
LiteISP [84]	24.05	0.86	6.43	0.13	24.01	0.86	6.58	0.13
DH-ISP [40]	23.20	0.84	8.52	0.17	22.98	0.84	8.49	0.17
MicroISP [41]	23.87	0.85	7.04	0.16	23.79	0.85	7.19	0.16
NVNet [42]	24.09	0.85	6.19	0.13	23.87	0.84	7.03	0.13
cmKAN [60]	24.51	0.88	5.31	0.10	24.25	0.88	6.07	0.10



Figure 7. Visual ablation study on the ZRR dataset [38]. Our full method (C_4) produces the best results. The standard GAN baseline (C_1) suffers from artifacts and color issues. The UOT-Only baseline (C_2) fails to reconstruct fine textures and accurate color. The $\tau = 0$ baseline (C_3) is not shown as it suffered from total mode collapse.

4.4. Ablation Studies

We conduct a series of ablation studies on the ZRR dataset [38] using the Restormer [82] backbone to validate our core design choices.

Table 5. Ablation study of our framework’s core components on the ZRR [38] dataset, using a fixed Restormer [82] backbone. Our full method (C_4) significantly outperforms a standard GAN baseline (C_1), while C_2 and C_3 demonstrate that both the expert committee and the UOT content cost ($\tau > 0$) are essential for high-fidelity results.

Config	Adversarial Obj.	Expert Committee	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$
C_1	Hinge Loss	✓ (All)	19.23	0.73	10.8
C_2	UOT-Loss	—	19.15	0.72	12.2
C_3	UOT-Loss ($\tau=0$)	✓ (All)	Fails (Mode Collapse)		
C_4 (Ours)	UOT-Loss	✓ (All)	21.69	0.84	9.93




Figure 8. Activation heatmaps (red=fake, blue=real) showing our potential and the discriminator committee analyzing a fake sRGB image ($\hat{Y}$). Each network learns to spot specific errors: *Potential Network* identifies global implausibility; *Color Expert* highlights semantically incorrect colors; *Structure Expert* detects bad local textures; *Frequency Expert* uses the FFT spectrum to find unnatural high-frequency artifacts (spikes).

Table 6. Ablation study of the expert committee. All models are trained with our UOT objective on the ZRR dataset [38] with the Restormer [82] backbone. The baseline (C_2) uses no expert discriminators. Our full method (C_4) utilizes all experts, and removing any of them, especially the Color Expert (E_1), leads to a significant drop in performance.

Config	Color Exp.	Structure Exp.	Frequency Exp.	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$
C_2 (UOT-Only)	—	—	—	19.15	0.72	12.2
E_1 (w/o Color)	—	✓	✓	19.88	0.75	14.8
E_2 (w/o Structure)	✓	—	✓	20.16	0.79	10.4
E_3 (w/o Frequency)	✓	✓	—	20.95	0.82	9.91
C_3 (Ours)	✓	✓	✓	21.69	0.84	9.93

Framework Ablation We first validate our hybrid EGUOT framework in Table 5. We test four configurations: (C_1) a standard GAN baseline using Hinge Loss and our expert committee; (C_2) a "UOT-Only" baseline, using only the UOT objective without the expert committee; (C_3) our full model, but with the content cost weight τ set to 0; and (C_4) our full EGUOT model. The results clearly demonstrate that all components are essential. The GAN baseline (C_1) and the UOT-Only baseline (C_2) produce poor results, with low PSNR/SSIM and high color error. C_3 fails due to mode collapse, proving that the content cost $c(\cdot, \cdot)$ is essential to anchor the UOT mapping. Our full model (C_4) significantly outperforms all other configurations, showing that the UOT objective and the expert committee work synergistically.

Expert Committee Ablation In Table 6, we analyze the contribution of each expert discriminator. We start from the C_2 (UOT-Only) baseline and also show the results of removing one expert at a time from our full model (C_4). The results show that all three experts contribute to the final performance. The UOT-Only model (C_2) is the weakest. Removing any single expert degrades performance, with the Color Expert being the most critical (E_1); its removal causes the ΔE to spike from 9.93 to 14.8, confirming its vital role in maintaining color fidelity. Figure 8 provides a qualitative visualization of this specialization, showing how each component learns to identify distinct error types.

5. Conclusion

In this work, we addressed the critical bottleneck in modern ISP development: the reliance on large-scale, paired raw-sRGB datasets. We identified that prior unpaired methods often fail due to their sensitivity to the distributional outliers found in real-world target sRGB datasets.

To solve this, we proposed **EGUOT**, a novel and robust framework for raw-to-sRGB translation, designed to operate in both paired and unpaired modes. Our framework is the first to apply Unbalanced Optimal Transport (UOT) for the ISP task, leveraging its theoretical robustness to dataset outliers. To achieve the high fidelity required of a modern ISP, we introduced a Committee of Expert Discriminators. This hybrid regularizer guides the mapping by providing targeted perceptual gradients for color, structure, and frequency, and is a key component in both training modes.

Our comprehensive experiments demonstrated the power and flexibility of our approach. We showed that EGUOT is architecture-agnostic, capable of training numerous state-of-the-art ISPs in a fully unpaired manner. Our paired training mode consistently exceeds the performance of the original settings across all metrics. Concurrently, our unpaired results are highly competitive with, and in some cases surpass, their original fully-paired counterparts. Furthermore, our unpaired mode sets a new state-of-the-art for unpaired raw-to-sRGB translation, significantly outperforming all prior GAN and OT-based methods. Our ablation studies empirically confirmed our core hypotheses: the UOT objective is essential for outlier robustness, while the expert committee is critical for perceptual fidelity.

We believe this approach will make ISP development more efficient by allowing developers to focus on the desired distribution properties instead of spending extensive time on data collection.

Limitations and Future Work

While our expert committee is effective, it adds complexity compared to a single-discriminator framework. Future work could explore methods to prune this committee. Moreover, the demonstrated effectiveness of EGUOT

opens promising avenues for applying this framework to other challenging unpaired image restoration tasks, such as low-light enhancement, denoising, and HDR reconstruction, where data outliers are common.

Acknowledgments

This work was partly supported by The Alexander von Humboldt Foundation. The authors also express sincere gratitude to Alexander Korotin for his valuable advice on the theoretical background of Optimal Transport.

References

- [1] SM A Sharif, Rizwan Ali Naqvi, and Mithun Biswas. Beyond joint demosaicking and denoising: An image processing pipeline for a pixel-bin image sensor. In *CVPR*, 2021. 2
- [2] Mahmoud Afifi and Abdullah Abuolaim. Semi-supervised raw-to-raw mapping. In *BMVC*, 2021. 2, 3
- [3] Mahmoud Afifi and Michael S Brown. Deep white-balance editing. In *CVPR*, 2020. 2
- [4] Mahmoud Afifi, Marcus A Brubaker, and Michael S Brown. Histogan: Controlling colors of gan-generated and real images via color histograms. In *CVPR*, 2021. 2
- [5] Mahmoud Afifi, Konstantinos G Derpanis, Bjorn Ommer, and Michael S Brown. Learning multi-scale photo exposure correction. In *CVPR*, 2021. 2
- [6] Andrei Arhire and Radu Timofte. Learned lightweight smartphone isp with unpaired data. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1878–1887, 2025. 2, 3, 5, 6
- [7] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017. 3, 5, 6
- [8] Nikola Banić, Egor Ershov, Artyom Panshin, Oleg Karasev, Sergey Korchagin, Shepelev Lev, Alexandr Startsev, Daniil Vladimirov, Ekaterina Zaychenkova, Dmitrii R Iarchuk, et al. Ntire 2024 challenge on night photography rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. 2
- [9] Jonathan T Barron and Yun-Ta Tsai. Fast Fourier color constancy. In *CVPR*, 2017. 2
- [10] Alexandru Brateanu, Raul Balmez, Adrian Avram, Ciprian Orhei, and Cosmin Ancuti. Lyt-net: Lightweight yuv transformer-based network for low-light image enhancement. *IEEE Signal Processing Letters*, 2025. 2
- [11] Alexandru Brateanu, Raul Balmez, Ciprian Orhei, Cosmin Ancuti, and Codruta Ancuti. Enhancing low-light images with kolmogorov–arnold networks in transformer attention. *Sensors*, 25(2):327, 2025. 2
- [12] Michael S Brown. Color processing for digital cameras. *Fundamentals and Applications of Colour Engineering*, pages 81–98, 2023. 1, 2
- [13] Yuanhao Cai, Hao Bian, Jing Lin, Haoqian Wang, Radu Timofte, and Yulun Zhang. Retinexformer: One-stage Retinex-based transformer for low-light image enhancement. In *ICCV*, 2023. 2
- [14] Huiwen Chang, Ohad Fried, Yiming Liu, Stephen DiVerdi, and Adam Finkelstein. Palette-based photo recoloring. *ACM Transactions on Graphics*, 34(4):139–1, 2015. 2
- [15] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *CVPR*, 2018. 2
- [16] Yu Cheng, Zhe Gan, Yitong Li, Jingjing Liu, and Jianfeng Gao. Sequential attention gan for interactive image editing. In *Proceedings of the 28th ACM international conference on multimedia*, pages 4383–4391, 2020. 3
- [17] Jaemoo Choi, Jaewoong Choi, and Myungjoo Kang. Generative modeling through the semi-dual formulation of unbalanced optimal transport. *Advances in Neural Information Processing Systems*, 36:42433–42455, 2023. 2, 3, 4, 5, 6
- [18] Yunje Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, 2018. 3
- [19] Yunje Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8188–8197, 2020. 3
- [20] Marcos V Conde, Florin Vasluianu, Javier Vazquez-Corral, and Radu Timofte. Perceptual image enhancement for smartphone real-time applications. In *CVPR*, 2023. 2
- [21] Marcos V Conde, Gregor Geigle, and Radu Timofte. Instructir: High-quality image restoration following human instructions. In *European Conference on Computer Vision*, pages 1–21. Springer, 2024. 2
- [22] 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. 2
- [23] Linhui Dai, Xiaohong Liu, Chengqi Li, and Jun Chen. AWNet: Attentive wavelet network for image ISP. In *ECCV*, 2020. 2, 5, 6, 7, 1
- [24] 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
- [25] Zhicheng Ding, Panfeng Li, Qikai Yang, Siyang Li, and Qingtian Gong. Regional style and color transfer. In *CVIDL*, 2024. 2
- [26] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [27] Egor Ershov, Alexey Savchik, Illya Semenov, Nikola Banić, Alexander Belokopytov, Daria Senshina, Karlo Košćević, Marko Subašić, and Sven Lončarić. The cube++ illumination estimation dataset. *IEEE access*, 8:227511–227527, 2020. 2
- [28] Milena Gazdieva, Petr Mokrov, Litu Rout, Alexander Korotin, Andrey Kravchenko, Alexander Filippov, and Evgeny Burnaev. An optimal transport perspective on unpaired image super-resolution. *Journal of Optimization Theory and Applications*, 207(2):40, 2025. 3

- [29] Jonathan Geuter, Gregor Kornhardt, Ingimar Tomasson, and Vaios Laschos. Universal neural optimal transport. In *Forty-second International Conference on Machine Learning*, 2025.
- [30] Nikita Gushchin, Alexander Kolesov, Alexander Korotin, Dmitry P Vetrov, and Evgeny Burnaev. Entropic neural optimal transport via diffusion processes. *Advances in Neural Information Processing Systems*, 36:75517–75544, 2023. 2, 3
- [31] Samuel W Hasinoff, Dillon Sharlet, Ryan Geiss, Andrew Adams, Jonathan T Barron, Florian Kainz, Jiawen Chen, and Marc Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics*, 35(6):1–12, 2016. 1, 2
- [32] Xuanhua He, Tao Hu, Guoli Wang, Zejin Wang, Run Wang, Qian Zhang, Keyu Yan, Ziyi Chen, Rui Li, Chengjun Xie, et al. Enhancing RAW-to-sRGB with decoupled style structure in Fourier domain. In *AAAI*, 2024. 2
- [33] Felix Heide, Markus Steinberger, Yun-Ta Tsai, Mushfiqu Rouf, Dawid Pająk, Dikpal Reddy, Orazio Gallo, Jing Liu, Wolfgang Heidrich, Karen Egiazarian, et al. Flexisp: A flexible camera image processing framework. *ACM Transactions on Graphics*, 33(6):1–13, 2014. 1, 2
- [34] Charles Herrmann, Richard Strong Bowen, Neal Wadhwa, Rahul Garg, Qiurui He, Jonathan T Barron, and Ramin Zabih. Learning to autofocus. In *CVPR*, 2020. 2
- [35] Zhiyong Hong, Dexin Zhen, Liping Xiong, Xuechen Li, and Yuhua Lin. srsrc-net: A new low-light image enhancement network via raw image reconstruction. *Applied Sciences*, 15(1):361, 2025. 2
- [36] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Dslr-quality photos on mobile devices with deep convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 3277–3285, 2017. 1
- [37] Andrey Ignatov, Radu Timofte, Zhilu Zhang, Ming Liu, Haolin Wang, Wangmeng Zuo, Jiawei Zhang, Ruimao Zhang, Zhanglin Peng, Sijie Ren, et al. AIM 2020 challenge on learned image signal processing pipeline. In *ECCVW*, 2020. 1, 2, 5, 6, 7
- [38] Andrey Ignatov, Luc Van Gool, and Radu Timofte. Replacing mobile camera ISP with a single deep learning model. In *CVPRW*, 2020. 1, 2, 5, 6, 7, 8
- [39] Andrey Ignatov, Cheng-Ming Chiang, Hsien-Kai Kuo, Anastasia Sycheva, and Radu Timofte. Learned smartphone ISP on mobile NPUs with deep learning, mobile AI 2021 challenge: Report. In *CVPRW*, 2021. 2
- [40] Andrey Ignatov, Cheng-Ming Chiang, Hsien-Kai Kuo, Anastasia Sycheva, and Radu Timofte. Learned smartphone ISP on mobile NPUs with deep learning, mobile ai 2021 challenge: Report. In *CVPR*, 2021. 1, 6, 7, 2
- [41] Andrey Ignatov, Anastasia Sycheva, Radu Timofte, Yu Tseng, Yu-Syuan Xu, Po-Hsiang Yu, Cheng-Ming Chiang, Hsien-Kai Kuo, Min-Hung Chen, Chia-Ming Cheng, et al. MicroISP: processing 32mp photos on mobile devices with deep learning. In *ECCV*, 2022. 1, 2, 6, 7
- [42] Andrey Ignatov, Georgii Perevozchikov, Radu Timofte, Cheng Li, Lian Liu, Jun Cao, Heng Sun, Wu Pan, Song Wang, KeQiang Yu, et al. Learned smartphone isp on mobile gpus, mobile ai 2025 challenge: Report. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1934–1946, 2025. 1, 2, 4, 6, 7
- [43] Woosik Jeong and Seung-Won Jung. RAWtoBit: A fully end-to-end camera isp network. In *ECCV*, 2022. 2
- [44] Yifan Jiang, Bartłomiej Wronski, Ben Mildenhall, Jonathan T Barron, Zhangyang Wang, and Tianfan Xue. Fast and high quality image denoising via malleable convolution. In *ECCV*, 2022. 2
- [45] Leonid Kantorovitch. On the translocation of masses. *Management science*, 5(1):1–4, 1958. 2, 3
- [46] Hakki Can Karaim and Michael S Brown. A software platform for manipulating the camera imaging pipeline. In *ECCV*, 2016. 1, 2
- [47] Zhanghan Ke, Yuhao Liu, Lei Zhu, Nanxuan Zhao, and Rynson WH Lau. Neural preset for color style transfer. In *CVPR*, 2023. 2
- [48] Junho Kim, Minjae Kim, Hyeonwoo Kang, and Kwanghee Lee. U-GAT-IT: Unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. *arXiv preprint arXiv:1907.10830*, 2019. 3
- [49] Alexander Korotin, Vage Egiazarian, Arip Asadulaev, Alexander Safin, and Evgeny Burnaev. Wasserstein-2 generative networks. *arXiv preprint arXiv:1909.13082*, 2019. 3, 6
- [50] Alexander Korotin, Alexander Kolesov, and Evgeny Burnaev. Kantorovich strikes back! wasserstein gans are not optimal transport? *Advances in Neural Information Processing Systems*, 35:13933–13946, 2022. 3
- [51] Alexander Korotin, Daniil Selikhanovych, and Evgeny Burnaev. Neural optimal transport. *arXiv preprint arXiv:2201.12220*, 2022. 2, 3, 4, 5, 6
- [52] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, 2017. 3
- [53] Kaijiang Li, Hao Li, Haining Li, Peisen Wang, Chunyi Guo, and Wenfeng Jiang. SIRLUT: Simulated infrared fusion guided image-adaptive 3D lookup tables for lightweight image enhancement. In *ACM MM*, 2024. 2
- [54] Chengxu Liu, Huan Yang, Jianlong Fu, and Xueming Qian. 4D LUT: learnable context-aware 4D lookup table for image enhancement. *IEEE Transactions on Image Processing*, 32:4742–4756, 2023. 2
- [55] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, 2017. 3
- [56] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022. 4
- [57] Ashok Makkuva, Amirhossein Taghvaei, Sewoong Oh, and Jason Lee. Optimal transport mapping via input convex neural networks. In *International Conference on Machine Learning*, pages 6672–6681. PMLR, 2020. 3

- [58] Zhipeng Mo, Wenbo Li, and Sihao Ding. Diffuse and refine latent prior with transformers for neural isp. In *2025 IEEE International Conference on Image Processing (ICIP)*, pages 1456–1461. IEEE, 2025. 2
- [59] Rang Nguyen, Dilip K Prasad, and Michael S Brown. Raw-to-raw: Mapping between image sensor color responses. In *CVPR*, 2014. 2
- [60] Artem Nikonorov, Georgy Perevozchikov, Andrei Korepanov, Nancy Mehta, Mahmoud Afifi, Egor Ershov, and Radu Timofte. Color matching using hypernetwork-based kolmogorov-arnold networks. *arXiv preprint arXiv:2503.11781*, 2025. 2, 3, 4, 6, 7, 1
- [61] Yingxue Pang, Jianxin Lin, Tao Qin, and Zhibo Chen. Image-to-image translation: Methods and applications. *IEEE Transactions on Multimedia*, 24:3859–3881, 2021. 2, 3
- [62] Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *ECCV*, 2020. 3
- [63] Georgy Perevozchikov, Nancy Mehta, Mahmoud Afifi, and Radu Timofte. Rawformer: Unpaired raw-to-raw translation for learnable camera ISPs. In *ECCV*, 2024. 1, 2, 3, 4, 5, 6
- [64] Khiem Pham, Khang Le, Nhat Ho, Tung Pham, and Hung Bui. On unbalanced optimal transport: An analysis of sinkhorn algorithm. In *International Conference on Machine Learning*, pages 7673–7682. PMLR, 2020. 2, 3
- [65] Aram-Alexandre Pooladian, Carles Domingo-Enrich, Ricky TQ Chen, and Brandon Amos. Neural optimal transport with lagrangian costs. *arXiv preprint arXiv:2406.00288*, 2024. 3, 6
- [66] Yang Ren, Xinhan Niu, Hai Jiang, Zhen Liu, Ting Jiang, Guanghui Liu, and Shuaicheng Liu. Isformer: Learning raw-to-srgb mappings with wavelet-based self-attention. *Neurocomputing*, page 131084, 2025. 1, 2
- [67] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 2015. 2, 4
- [68] Kevin Roth, Aurelien Lucchi, Sebastian Nowozin, and Thomas Hofmann. Stabilizing training of generative adversarial networks through regularization. *Advances in neural information processing systems*, 30, 2017. 3
- [69] Eli Schwartz, Raja Giryes, and Alex M Bronstein. DeepISP: Toward learning an end-to-end image processing pipeline. *IEEE Transactions on Image Processing*, 28(2):912–923, 2018. 2
- [70] Tim Seizinger, Marcos V Conde, Manuel Kolmet, Tom E Bishop, and Radu Timofte. Efficient multi-lens bokeh effect rendering and transformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1633–1642, 2023. 2
- [71] Ardhendu Shekhar Tripathi, Martin Danelljan, Samarth Shukla, Radu Timofte, and Luc Van Gool. Transform your smartphone into a dslr camera: Learning the isp in the wild. In *European Conference on Computer Vision*, pages 625–641. Springer, 2022. 1, 2, 5, 7
- [72] Ondřej Šindelář and Filip Šroubek. Image deblurring in smartphone devices using built-in inertial measurement sensors. *Journal of Electronic Imaging*, 22(1):011003–011003, 2013. 2
- [73] Dmitrii Torbunov, Yi Huang, Huan-Hsin Tseng, Haiwang Yu, Jin Huang, Shinjae Yoo, Meifeng Lin, Brett Viren, and Yihui Ren. Rethinking cycleGAN: Improving quality of GANs for unpaired image-to-image translation. *arXiv preprint arXiv:2303.16280*, 2023. 3, 6
- [74] Dmitrii Torbunov, Yi Huang, Haiwang Yu, Jin Huang, Shinjae Yoo, Meifeng Lin, Brett Viren, and Yihui Ren. UVCGAN v2: An improved cycle-consistent GAN for unpaired image-to-image translation. *arXiv preprint arXiv:2303.16280*, 2023. 2, 3, 6
- [75] Vamsi Krishna Vasa, Peijie Qiu, Wenhui Zhu, Yujian Xiong, Oana Dumitrascu, and Yalin Wang. Context-aware optimal transport learning for retinal fundus image enhancement. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 4016–4025. IEEE, 2025. 3
- [76] Daniel Wirzberger Raimundo, Andrey Ignatov, and Radu Timofte. LAN: Lightweight attention-based network for raw-to-rgb smartphone image processing. In *CVPRW*, 2022. 1, 2, 5, 6, 7
- [77] Bartłomiej Wronski, Ignacio Garcia-Dorado, Manfred Ernst, Damien Kelly, Michael Krainin, Chia-Kai Liang, Marc Levoy, and Peyman Milanfar. Handheld multi-frame super-resolution. *ACM Transactions on Graphics (ToG)*, 38(4):1–18, 2019. 1
- [78] Yazhou Xing, Zian Qian, and Qifeng Chen. Invertible image signal processing. In *CVPR*, 2021. 1, 2
- [79] Canqian Yang, Meiguang Jin, Yi Xu, Rui Zhang, Ying Chen, and Huaidda Liu. SepLUT: Separable image-adaptive lookup tables for real-time image enhancement. In *ECCV*, 2022. 2
- [80] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. DualGAN: Unsupervised dual learning for image-to-image translation. In *ICCV*, pages 2849–2857, 2017. 3
- [81] 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
- [82] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5728–5739, 2022. 2, 5, 6, 7, 8, 1
- [83] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 3
- [84] Zhilu Zhang, Haolin Wang, Ming Liu, Ruohao Wang, Jiawei Zhang, and Wangmeng Zuo. Learning raw-to-sRGB mappings with inaccurately aligned supervision. In *ICCV*, 2021. 1, 2, 6, 7
- [85] Yihao Zhao, Ruihai Wu, and Hao Dong. Unpaired image-to-image translation using adversarial consistency loss. In *ECCV*, 2020. 3

- [86] Wenli Zheng, Huiyuan Fu, Xicong Wang, Hao Kang, Chuanming Wang, Jin Liu, Zekai Xu, Heng Zhang, and Huadong Ma. Evraw: Event-guided structural and color modeling for raw-to-srgb image reconstruction. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 209–218, 2025. [2](#)
- [87] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. [2](#), [3](#), [5](#), [6](#)
- [88] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. SEAN: Image synthesis with semantic region-adaptive normalization. In *CVPR*, 2020. [3](#)

Experts-Guided Unbalanced Optimal Transport for ISP Learning from Unpaired and/or Paired Data

Supplementary Material

In this supplementary material, we first present additional quantitative and qualitative results in Sec. 6 to further validate the architecture-agnostic nature of our framework. Next, we provide a detailed robustness study in Sec. 7, empirically demonstrating the superiority of Unbalanced Optimal Transport over standard GANs in the presence of dataset outliers.

6. Additional results

In this section, we verify the performance of the proposed EGUOT framework across three diverse benchmarks: the Zurich Raw-to-RGB (ZRR) dataset [38], the Mobile AI (MAI) dataset [42], and the ISPIW dataset [71].

Visual Quality. We provide extensive visual comparisons in Fig. 9, Fig. 10, and Fig. 11. These qualitative results demonstrate that our Experts-Guided Unbalanced Optimal Transport (EGUOT) framework successfully reconstructs high-fidelity sRGB images across different sensors and scenes. Crucially, our *unpaired* training mode produces results that are perceptually indistinguishable from, and in some cases sharper than, the original *paired* baselines. The Expert Committee—specifically, the frequency and structure experts—ensures that the generated images remain free of the hallucinatory artifacts often observed in standard GAN-based unpaired translations.

Quantitative Evaluation. We report detailed quantitative metrics for our paired training mode in Table 7 (Mobile AI) and Table 8 (ISPIW). Consistent with the results reported in the main paper, our framework establishes a new state-of-the-art for supervised training. By utilizing the Expert Committee as a regularizer alongside the pixel-wise loss, we consistently improve upon the original paired baselines across all metrics (PSNR, SSIM, ΔE , and LPIPS).

For instance, on the Mobile AI dataset (Table 7), our method improves the MW-ISP [37] backbone by nearly 0.5 dB in PSNR while reducing the color error (ΔE) from 6.27 to 5.91. Similarly, on the ISPIW dataset (Table 8), our training scheme boosts the performance of the recent Transformer-based Restormer [82] and LiteISP [84] architectures. These results confirm that EGUOT is a robust, architecture-agnostic training paradigm that enhances ISP learning regardless of the underlying network capacity.

7. Robustness Study

A core motivation of our work is the fragility of standard GAN losses when facing noisy datasets—a common sce-




Figure 9. Quantitative results on the on the ZRR raw-to-sRGB dataset [38], demonstrating the architecture-agnostic nature of our EGUOT framework. We apply (a) the original, *paired* training and our proposed (b) *paired* / (c) *unpaired* training to a wide range of SOTA ISP backbones. Our *paired* mode consistently outperforms the original, while our *unpaired* mode achieves competitive performance. Best viewed in the electronic version.

Table 7. Results on the Mobile AI raw-to-sRGB dataset [42], comparing SOTA backbones retrained with our framework against their *original paired-data results*. Our *paired* mode achieves the best performance. Better results are highlighted in yellow.

Backbone	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$	Paired Mode (original)				Paired Mode (ours)			
					PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$
Restormer [82]	23.99	0.88	6.74	0.13	24.09	0.88	6.67	0.12	24.09	0.88	6.67	0.12
MW-ISP [37]	24.31	0.88	6.27	0.11	24.79	0.89	5.91	0.10	24.79	0.89	5.91	0.10
AW-Net [23]	24.17	0.87	6.35	0.12	24.33	0.88	6.12	0.11	24.33	0.88	6.12	0.11
LAN-ISP [76]	23.48	0.87	7.18	0.15	23.59	0.88	7.07	0.14	23.59	0.88	7.07	0.14
LiteISP [84]	24.05	0.86	6.43	0.13	24.22	0.87	6.08	0.12	24.22	0.87	6.08	0.12
DH-ISP [40]	23.20	0.84	8.52	0.17	23.29	0.84	8.21	0.16	23.29	0.84	8.21	0.16
MicroISP [41]	23.87	0.85	7.04	0.16	24.00	0.86	6.99	0.15	24.00	0.86	6.99	0.15
NVNet [42]	24.09	0.85	6.19	0.13	24.12	0.86	6.11	0.13	24.12	0.86	6.11	0.13
cmKAN [60]	24.51	0.88	5.31	0.10	24.55	0.88	5.28	0.10	24.55	0.88	5.28	0.10

nario in real-world unpaired learning where the source and target domains contain different outliers. To validate our hypothesis that the Unbalanced Optimal Transport (UOT) framework provides superior robustness, we conduct a controlled corruption experiment.

We create a “Dirty” version of the ZRR [38] target




Figure 10. Quantitative results on the on the MAI raw-to-sRGB dataset [42], demonstrating the architecture-agnostic nature of our EGUOT framework. We apply (a) the original, *paired* training and our proposed (b) *paired* / (c) *unpaired* training to a wide range of SOTA ISP backbones. Our *paired* mode consistently outperforms the original, while our *unpaired* mode achieves competitive performance. Best viewed in the electronic version.

Table 8. Results on the ISPIW raw-to-sRGB dataset [71], comparing SOTA backbones retrained with our framework against their *original paired-data results*. Our *paired* mode achieves the best performance. Better results are highlighted in yellow.

Backbone	Paired Mode (original)				Paired Mode (ours)			
	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$	LPIPS $\downarrow$
Restormer [82]	20.93	0.79	6.85	0.14	21.04	0.80	6.78	0.13
MW-ISP [37]	21.90	0.81	7.03	0.11	21.96	0.81	6.85	0.10
AW-Net [23]	21.75	0.81	6.99	0.12	22.04	0.82	6.87	0.11
LAN-ISP [76]	22.09	0.81	6.97	0.11	22.11	0.82	6.40	0.10
LiteISP [84]	22.14	0.81	6.31	0.11	22.33	0.82	6.09	0.10
DH-ISP [40]	19.92	0.76	7.89	0.17	21.15	0.77	7.27	0.17
MicroISP [41]	20.70	0.77	6.92	0.15	20.88	0.78	6.29	0.14
NVNet [42]	23.80	0.82	5.81	0.10	23.96	0.83	5.80	0.10
cmKAN [60]	24.22	0.83	5.29	0.09	24.30	0.84	5.15	0.09

Table 9. Robustness to Outliers. We compare our UOT-based framework against a standard GAN when trained on a “clean” vs. “dirty” (15% outliers) ZRR dataset [38]. Our method’s performance remains stable, demonstrating its robustness.

Config	Training Dataset	PSNR $\uparrow$	SSIM $\uparrow$	ΔE $\downarrow$
C_1	Clean	19.23	0.73	10.8
	Dirty	18.65	0.72	11.17
C_4 (Ours)	Clean	21.69	0.84	9.93
	Dirty	21.41	0.84	10.09

dataset by polluting it with 15% outlier images. Specifically, we apply severe random augmentations (random color jitter, contrast shifts, and brightness modifications) to 15%




Figure 11. Quantitative results on the on the ISPIW raw-to-sRGB dataset [71], demonstrating the architecture-agnostic nature of our EGUOT framework. We apply (a) the original, *paired* training and our proposed (b) *paired* / (c) *unpaired* training to a wide range of SOTA ISP backbones. Our *paired* mode consistently outperforms the original, while our *unpaired* mode achieves competitive performance. Best viewed in the electronic version.

of the training data to simulate low-quality samples often found in scraped datasets. We then train our full model (C_4) and a standard GAN baseline (C_1 , utilizing a Hinge loss) on both the “Clean” and “Dirty” datasets.

As shown in Table 9, the standard GAN’s performance is highly sensitive to the data quality. When trained on the polluted data, the GAN fails to discount the outliers, leading to a significant performance degradation: its PSNR drops by 0.58 dB and ΔE worsens by 0.37 units. In stark contrast, our EGUOT framework demonstrates remarkable stability. Consequently, our method maintains high fidelity on the “Dirty” dataset, suffering only a negligible 0.28 dB drop in PSNR and 0.18 units in ΔE . This empirically proves that UOT-based framework is fundamentally more robust to the noisy data distributions encountered in unpaired environments.