

# Learning Degradation-Independent Representations for Camera ISP Pipelines

Yanhui Guo\*

McMaster University, Canada

guoy143@mcmaster.ca

Fangzhou Luo\*

McMaster University, Canada

luof1@mcmaster.ca

Xiaolin Wu†

McMaster University, Canada

xwu@mcmaster.ca

## Abstract

*Image signal processing (ISP) pipeline plays a fundamental role in digital cameras, which converts raw Bayer sensor data to RGB images. However, ISP-generated images usually suffer from imperfections due to the compounded degradations that stem from sensor noises, demosaicing noises, compression artifacts, and possibly adverse effects of erroneous ISP hyperparameter settings such as ISO and gamma values. In a general sense, these ISP imperfections can be considered as degradations. The highly complex mechanisms of ISP degradations, some of which are even unknown, pose great challenges to the generalization capability of deep neural networks (DNN) for image restoration and to their adaptability to downstream tasks. To tackle the issues, we propose a novel DNN approach to learn degradation-independent representations (DiR) through the refinement of a self-supervised learned baseline representation. The proposed DiR learning technique has remarkable domain generalization capability and consequently, it outperforms state-of-the-art methods across various downstream tasks, including blind image restoration, object detection, and instance segmentation, as verified in our experiments.*

## 1. Introduction

In our information era, digital cameras have become ubiquitous and indispensable in all walks of socioeconomic life. They are arguably the most common type of sensory device for both humans and their intelligent agents to sense the environment. The applications of digital imaging fall into two categories depending on their end purposes: human vision oriented and machine vision oriented. Examples of the first category include smartphones, social media, internet videos, etc., while robotics, automation, and medical imaging are examples of the second category.

For human vision oriented applications, the image quality metrics should correspond to human visual perception, because the images are primarily intended to please our

eyes. On the other hand, for machine vision oriented applications [18, 19, 40, 49], the image quality requirements may not completely align with those of human perception. In the latter case, what matters most is the preservation of semantically critical image features that are needed to solve computer vision problems of recognition and classification.

For either human viewing or downstream computer vision tasks, the image acquisition process of digital cameras is the same. All images are generated by a hardware image signal processing (ISP) pipeline that transforms raw Bayer sensor data into RGB images. The ISP pipeline consists of a set of cascaded steps, including denoising, color demosaicking, and compression. Each of the ISP steps can cause a level of degradation in image quality. At the beginning of the ISP pipeline, there are sensor noises, particularly in dark environments with a low signal-to-noise ratio. The noisy RAW data are then demosaiced to produce the corresponding RGB images. This demosaicking process introduces interpolation noises of its own. Finally, the RGB image is compressed for bandwidth efficiency, which adds extra compression noises on top of sensor noises and demosaicking noises. The resulting compound effects of these cascaded noises are very difficult to model precisely [11]. The underlying complex image quality degradation mechanism prevents the synthesis of clean and degraded image pairs for supervised learning.

Yet, there is another type of ISP degradation that exerts adverse effects on the applications of ISP-generated images in downstream tasks. The camera ISP involves a number of adjustable hyperparameters such as gamma correction, ISO, gain control, and white balance. Their optimal settings may vary from application to application, *i.e.*, there exists no one-fits-all solution. Most cameras have a default hyperparameter setting that is tailored to perceptual quality rather than optimized for a specific computer vision task. We consider the deviations of ISP hyperparameters from their optimal settings, in a general sense, to be degradations as well. This second type of ISP degradation can lead to underperformances or even outright failures of downstream computer vision algorithms [32, 38]. Like the first type of signal-level degradations, the second type of semantic-level degradations is extremely difficult, if not impossible,

\* Equal contribution.

† Corresponding author.

to model and synthesize, not only because the space of ISP hyperparameters is of very high dimensions but also because the two types of ISP degradations are intertwined.




Figure 1. Comparison of the canonical learning paradigm and our DiR learning technique.

As discussed above, it is either intractable or highly expensive to obtain the ISP-degraded and the corresponding latent degradation-free (ground truth) image pairs. Moreover, exhausting all ISP degradation variants in a brute force way is impossible, many of which are not even known or understood [27, 28]. To meet these challenges, we introduce a notion of degradation-independent latent representation (DiR) that is robust against camera-captured non-ideal real-world images and propose a new deep neural network approach to learn the DiR. Fig. 1 illustrates the difference between our DiR learning paradigm and the canonical learning paradigm. Within the framework of our DiR learning, the learned DiR representations contain intrinsic information of the original degradation-free image and hence it is more general and robust with respect to unseen degradation types and severity when used in downstream computer vision and image processing tasks. The DiR learning, carried out by a deep neural network called DiRNet, is self-supervised, and the DiRNet is trained with two learning objectives: any two degraded images are mapped to the same DiR, and all DiRs are sampled from a normal distribution, which is learning a variational auto-encoder (VAE) by the multi-view mutual information (MMI) maximization [7, 15, 43].

Moreover, we take advantage of known priors of the ISP workflow to synthesize original and degraded pairs, so that the above baseline DiR can be reinforced by an auxiliary self-supervised learning branch. This auxiliary branch learns the reference representation  $\mathbf{r}^*$  of the degradation-free images, using high-quality images (surrogate ground truth). The baseline DiR result  $\mathbf{r}^{(0)}$  is statistically aligned with the result  $\mathbf{r}^*$  of the auxiliary branch by an alignment network module. This alignment module is trained using the aforementioned synthesized pairs and learns to align the baseline DiR  $\mathbf{r}^{(0)}$  to the degradation-free representation reference (DfR)  $\mathbf{r}^*$ , yielding a refined DiR  $\mathbf{r}^+$ ; finally, the refined DiR  $\mathbf{r}^+$  is ready for being used in downstream tasks. Overall, the refined DiR benefits from the best of both self-supervised and supervised

learning, and strikes a good balance between the model’s robustness and precision for applications.

Our main contributions and key results can be summarized as follows:

- A novel method is proposed for self-supervised learning of the DiR representation by multivariate mutual information maximization, which requires only degraded training images.
- A joint learning approach is employed to refine the DiR representation for optimal performance in the targeted end-task network.
- The versatility of the proposed methods across a wide range of tasks, including image restoration, object detection, and instance segmentation, is supported by ample empirical evidence.

## 2. Related Work

**Degradation blind image restoration network** There is a long history of research on image restoration tasks that are blind to degradation parameters, such as blind image deconvolution algorithms [23]. Accompanying the proliferation of deep learning based image restoration methods, recent years have seen a number of papers [9, 10, 12, 24, 26, 28, 36, 47] on the topic of improving the generalization ability of image restoration networks. Wang *et al.* [47] used a high-order degradation model to mimic real-world degradations. Luo *et al.* [28] proposed a novel adversarial neural degradation model to simulate degraded training data. Li *et al.* [24] introduced a distortion invariant representation learning paradigm from the causality perspective, to improve the generalization ability of restoration networks for unknown degradations.

**End-to-end neural network solutions for ISP pipelines** A more holistic design of neural networks [34, 41, 51] for image restoration is to replace the cascaded stages of the hardware ISP pipeline with a single end-to-end mapping from raw Bayer sensor data to the intended ideal RGB images, instead of postprocessing the output images of the camera ISP. For instance, Chen *et al.* [2] proposed a CNN method, called SID (see in the dark), which directly converts a low lighting noisy Bayer image to a corresponding properly exposed RGB image, without generating the underexposed ineligible RGB image first and then restoring it. Tseng *et al.* [44] leveraged a convolutional neural network to parameterize a differentiable mapping between the ISP configuration space and evaluation metrics of downstream tasks and validated their method on the tasks of object detection, and extreme low-light imaging. Yu *et al.* [52] replaced the conventional ISP with a trainable ISP framework that improves the quality of ISP-generated results by using a differential neural architecture search algorithm to search for the optimal ISP architecture and hyperparameters.

**Image enhancement for high-level tasks** In general practice, image restoration networks tend to be optimized for

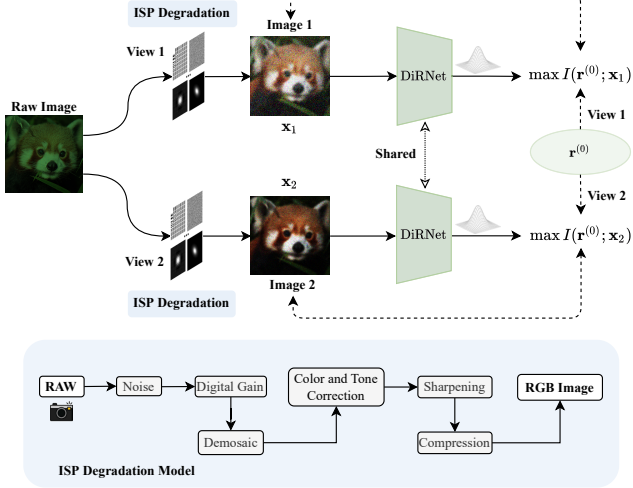


Figure 2. Self-supervised MMI maximization for DiRNet.

human visual perception rather than for computer vision tasks. However, the human vision-centric neural network may not provide optimal results on high-level computer vision tasks. This motivates the development of a family of image restoration networks that are tuned in line with the subsequent computer vision tasks, such as recognition, detection, and classification [4, 39, 52]. For example, Cui *et al.* [4] proposed a multitask auto-encoding transformation model to learn the intrinsic visual structure by considering the physical sensor noise model in ISP, demonstrating effectiveness for dark object detection. Mosleh *et al.* [32] investigated a so-called “hardware-in-the-loop” strategy to directly optimize the hyperparameters of hardware ISP pipelines, by solving a nonlinear multi-objective optimization problem with a 0th-order stochastic solver. Qin *et al.* [39] designed a sequential CNN model to recurrently adjust the ISP hyperparameters for downstream tasks and demonstrate the advantages of the model.

### 3. Approach

In this section, we first detail the DiRNet mentioned in the introduction to learn the baseline DiR  $\mathbf{r}^{(0)}$  of multiple degraded images. The DiRNet training is driven by the MMI maximization. Then, we discuss how to learn the degradation-free latent representation  $\mathbf{r}^*$  using a variant of DiRNet. The two latent representations  $\mathbf{r}^{(0)}$  and  $\mathbf{r}^*$  are prepared to train an alignment network to produce a refined DiR  $\mathbf{r}^+$ . Finally, we elaborate on the process to jointly optimize the alignment network and an auxiliary network for a downstream task.

#### 3.1. Optimizing DiRNet by MMI Maximization

**Mutual information maximization** To begin, we pave the way for computing the multi-view mutual information (MMI) maximization by introducing a mutual information estimator. For estimating mutual information

$I(\mathbf{x}; \mathbf{y})$  of two random variables  $\mathbf{x}$  and  $\mathbf{y}$ , it is equivalent to computing the Kullback-Leibler (KL) divergence  $\mathcal{D}_{\text{KL}}(p(\mathbf{x}, \mathbf{y}) || p(\mathbf{x})p(\mathbf{y}))$ . However, the KL divergence has no theoretical upper bound. Maximizing it cannot guarantee the convergence of the optimization process. Following previous works [35, 45], we exploit a non-KL divergence named Jensen-Shannon divergence (JSD) to devise the lower bound of mutual information, which satisfies the below proposition. **Proposition 1** Let  $\mathbf{x}$  and  $\mathbf{y}$  represent two random variables, the  $\mathbf{J} = p(\mathbf{x}, \mathbf{y})$  and  $\mathbf{M} = p(\mathbf{x})p(\mathbf{y})$  are the joint and the product of marginals of the two variables, respectively. The mutual information between the variables satisfies<sup>1</sup>

$$I(\mathbf{x}; \mathbf{y}) := \mathcal{D}_{\text{JSD}}(\mathbf{J}; \mathbf{M}) \geq \mathbb{E}_{z \sim \mathbf{J}} [-\sigma(-\mathcal{F}_\omega(z))] - \mathbb{E}_{z' \sim \mathbf{M}} [\sigma(\mathcal{F}_\omega(z'))], \quad (1)$$

where  $\sigma(t) = \log(1 + e^t)$ , and the discriminator function  $\mathcal{F}_\omega$  [35] is modeled by a neural network with parameters  $\omega$ . **Multi-view mutual information maximization** The baseline DiR  $\mathbf{r}^{(0)}$  is a degradation-independent representation shared among all degraded images that originate from an ideal degradation-free image. From the perspective of information theory, the DiR representation should satisfy the criterion of shared information maximization between it and multiple degraded observations. This shared information maximization is so-called multi-view mutual information (MMI) maximization[7] a.k.a interaction information maximization [43] for the three variables. We state its definition as follows: Given two degraded images  $\mathbf{x}_1$  and  $\mathbf{x}_2$ , their shared latent DiR  $\mathbf{r}^{(0)}$  can be found by maximizing their MMI  $I(\mathbf{r}^{(0)}; \mathbf{x}_1; \mathbf{x}_2)$  that is defined as

$$I(\mathbf{r}^{(0)}; \mathbf{x}_1; \mathbf{x}_2) = I(\mathbf{r}^{(0)}; \mathbf{x}_1) - \underbrace{I(\mathbf{r}^{(0)}; \mathbf{x}_1 | \mathbf{x}_2)}_{\text{Redundancy}} \quad (2)$$

$$= I(\mathbf{r}^{(0)}; \mathbf{x}_2) - I(\mathbf{r}^{(0)}; \mathbf{x}_2 | \mathbf{x}_1), \quad (3)$$

where Eq. (3) holds due to symmetry. The conditional mutual information  $I(\mathbf{r}^{(0)}; \mathbf{x}_1 | \mathbf{x}_2)$  represents the incremental mutual information between  $\mathbf{r}^{(0)}$  and the degraded observation  $\mathbf{x}_1$  when given another degraded observation  $\mathbf{x}_2$ . This incremental information denotes the information not shared by the observations  $\mathbf{x}_1$  and  $\mathbf{x}_2$ , due to the differences in various degradation sources, which should be minimized. Maximizing  $I(\mathbf{r}^{(0)}; \mathbf{x}_1; \mathbf{x}_2)$  encourages  $\mathbf{r}^{(0)}$  to encode the degradation-independent information shared by the given multiple degraded images.

The network DiRNet for learning the baseline DiR is of VAE type [14]. The DiRNet takes the ISP-generated images (*i.e.* camera captured) as inputs and maps them into a canonical representation space *i.e.* the DiR latent space, by removing the degradation-specific information. Fig. 2 shows the training process for the DiRNet. To continue the development, we introduce the following proposition [15]:

<sup>1</sup>The derivation is in Sec.B and Sec.C of the supplementary material.

**Proposition 2** Given two degraded images  $\mathbf{x}_1$  and  $\mathbf{x}_2$ , and the baseline DiR representation  $\mathbf{r}^{(0)}$ , the conditional mutual information  $I(\mathbf{r}^{(0)}; \mathbf{x}_1 | \mathbf{x}_2)$  satisfies:

$$I(\mathbf{r}^{(0)}; \mathbf{x}_1 | \mathbf{x}_2) \leq \mathbb{E}_{p(\mathbf{x}_1, \mathbf{x}_2)} [\mathcal{D}_{\text{KL}}(p(\mathbf{r}^{(0)} | \mathbf{x}_1, \mathbf{x}_2) || p_\varphi(\mathbf{r}^{(0)} | \mathbf{x}_2))]. \quad (4)$$

To maximize  $I(\mathbf{r}^{(0)}; \mathbf{x}_1; \mathbf{x}_2)$ , the overall training objective<sup>2</sup> of DiRNet is defined by the average of Eq. (2) and Eq. (3):

$$\begin{aligned} \mathcal{L}_0 = & - \frac{I(\mathbf{r}^{(0)}; \mathbf{x}_1) + I(\mathbf{r}^{(0)}; \mathbf{x}_2)}{2} \\ & + \lambda \mathbb{E}_{p(\mathbf{x}_1, \mathbf{x}_2)} [\mathcal{D}_{\text{AKL}}] \\ & + \beta \mathcal{D}_{\text{KL}} [p(\mathbf{r}^{(0)}) || \mathcal{N}(0, \mathbf{I})], \end{aligned} \quad (5)$$

in which we can maximize the upper bound of the conditional mutual information according to Proposition 2. In Eq.(5),  $\lambda$  and  $\beta$  are the trade-off hyperparameters. The DiR representation is regularized by the Gaussian prior  $\mathcal{N}(0, \mathbf{I})$  as in VAE [14].  $\mathcal{D}_{\text{AKL}}$  stands for the average of  $\mathcal{D}_{\text{KL}}(p(\mathbf{r}^{(0)} | \mathbf{x}_1, \mathbf{x}_2) || p_\varphi(\mathbf{r}^{(0)} | \mathbf{x}_1))$  and  $\mathcal{D}_{\text{KL}}(p(\mathbf{r}^{(0)} | \mathbf{x}_1, \mathbf{x}_2) || p_\varphi(\mathbf{r}^{(0)} | \mathbf{x}_2))$ . The trainable network is the DiRNet  $p_\varphi$  that produces the baseline DiR and an auxiliary network to estimate the upper bound [15].

### 3.2. Degradation-Free Representation Reference

Although the baseline DiR  $\mathbf{r}^{(0)}$  aims to distill the degradation-invariant intrinsic elements and structures, it may have biases toward the degraded data because the learning process of  $\mathbf{r}^{(0)}$  has not observed degradation-free samples. To compensate for this insufficiency, we collect a large set of high-quality RGB images and assume them to be free of degradation. As these images are captured by real non-idealistic cameras the degradation-free assumption is not strictly true, but it is acceptable for practical purposes. These degradation-free images are used to learn the degradation-free latent representation (DfR)  $\mathbf{r}^*$ ; the above two latent representations  $\mathbf{r}^{(0)}$  and  $\mathbf{r}^*$  will be aligned to produce a refined DiR  $\mathbf{r}^+$  by a so-called alignment network, whose construction and rationale will be discussed in Sec. 3.3.

Similar to training DiRNet for the baseline DiR  $\mathbf{r}^{(0)}$ , we propose the DfRNet to learn the DfR  $\mathbf{r}^*$  from a degradation-free image (surrogate ground truth). Since this DfR learning is with respect to a single image, the MMI training for  $\mathbf{r}^*$  is simplified to the mutual information maximization *i.e.* maximizing  $I(\mathbf{r}^*; \mathbf{y}^*)$ , which can be solved by maximizing the lower bound derived by Proposition 1. Moreover, following VAE [14], we train the decoder  $D_*$  that forces the learned DfR to be able to reconstruct the degradation-free image. The objective of the above DfR learning is to minimize  $\mathcal{L}_*$

<sup>2</sup>The full derivation of  $\mathcal{L}_0$  can be found in Sec.B of the supplementary material.

defined as follows

$$\begin{aligned} \mathcal{L}_* = & - I(\mathbf{r}^*; \mathbf{y}^*) + \mathbb{E} [\|D_*(\mathbf{r}^*) - \mathbf{y}^*\|_1] \\ & + \beta^* \mathcal{D}_{\text{KL}} [p(\mathbf{r}^*) || \mathcal{N}(0, \mathbf{I})], \end{aligned} \quad (6)$$

where  $\beta^*$  is the trade-off hyperparameter. The defined training objective in Eq. (6) is the same as the InfoVAE [54].

### 3.3. Alignment Network for Refined DiR

**Alignment network** As mentioned in the introduction, we can neutralize, to a certain extent, possible biases in the baseline DiR  $\mathbf{r}^{(0)}$  by aligning  $\mathbf{r}^{(0)}$  with the above DfR  $\mathbf{r}^*$  in distribution. This is accomplished by a so-called alignment network  $\mathcal{A}$ . To guide the statistical alignment of  $\mathbf{r}^{(0)}$  and  $\mathbf{r}^*$ , we also feed the degraded training image  $\mathbf{x}$  to network DfR-Net and use the resulting latent representation  $\vec{\mathbf{r}} = \text{DfR}(\mathbf{x})$ , called the pilot representation. This guided alignment process is a mapping function

$$\mathcal{A}(\mathbf{r}^{(0)}, \vec{\mathbf{r}}) : \mathbf{r}^{(0)} \xrightarrow[\vec{\mathbf{r}} = \text{DfR}(\mathbf{x})]{\text{Guided by } \vec{\mathbf{r}}} \mathbf{r}^+. \quad (7)$$

Next, we describe the architecture of the alignment network  $\mathcal{A}$  and discuss the role of the pilot representation  $\vec{\mathbf{r}}$ .




Figure 3. Illustration of the alignment network.

**Network structure** Fig. 3 shows the network structure of the alignment network  $\mathcal{A}$ . It consists of three cascaded modules  $\mathcal{A}_1$ ,  $\mathcal{A}_2$  and  $\mathcal{A}_3$ . The front module  $\mathcal{A}_1$  of  $m_1$  convolution layers is used to extract the features from the input image. The middle module  $\mathcal{A}_2$  of  $m_2$  layers is designed to modulate the latent features with the guidance of the pilot DfR  $\vec{\mathbf{r}}$ . The end module  $\mathcal{A}_3$  of  $m_3$  convolution layers is used to produce the final refined DiR  $\mathbf{r}^+$ .

In the guidance module  $\mathcal{A}_2$ ,  $\vec{\mathbf{r}}$  is fed to convolution layers with  $3 \times 3$  kernels, followed by the average pooling layer to obtain the modulation feature tensors. These modulation tensors are then reshaped to form  $k$  adaptive convolution kernels. We use a so-called G-Conv module formed by these adaptive convolution kernels to convolve with the input features to produce the guided features  $f_k$  which are embedded with the information of the pilot DfR  $\vec{\mathbf{r}}$ . In addition,



we exploit a spatial attention branch to assist the guidance information extraction. This attention branch predicts the attention maps to strengthen the feature units that have a high correlation to pilot DfR  $\vec{r}$ , yielding the attention-weighted intermediate features  $f_a$ . Finally, the guided features  $f_k$  and the attention-weighted features  $f_a$  are fused by an addition layer to produce the refined DiR  $\mathbf{r}^+$ , which is ready for use in downstream tasks. In our experiments, we *jointly* train this alignment network  $\mathcal{A}$  and a task-aligned auxiliary network  $\mathcal{T}$ , which will be introduced in Sec. 3.4.

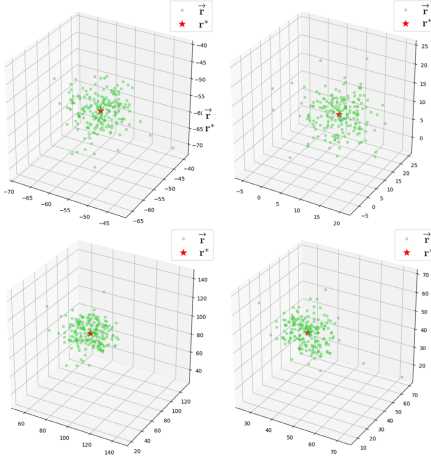


Figure 4. Visualization of the pilot DfR representations  $\vec{r}$  and the expected DfR representations  $\mathbf{r}^*$ . The information provided by the pilot DfR can guide the alignment to find the optimal DfR.

**Discussion on the pilot DfR  $\vec{r}$**  To explain the rationale behind the guidance design of the alignment network, we conduct experiments for visualization. More specifically, we choose four different degradation-free images to extract their corresponding DfRs  $\mathbf{r}^*$ . We then use the ISP pipeline as in [39] by adding noise to generate 200 various degraded copies of each degradation-free image. These degraded samples, which share the same image content but have different ISP degradations, are fed into the network DfR to produce the pilot DfRs  $\mathbf{r}^*$ . Fig. 4 shows the visualized results by using t-SNE [46], where the red stars denote the expected ground truth DfRs  $\mathbf{r}^*$  for the alignment network training, and the green stars stand for the pilot DfRs  $\vec{r}$ . We can observe that the green points are projected near the red points, suggesting that the pilot DfRs have a strong correlation with the ground truth DfRs. This is because the pilot DfRs and the ground truth DfRs share the same degradation-free content. Inspired by this phenomenon, the alignment network is designed to take into account the assistant information provided by  $\vec{r}$  to guide the refinement of the baseline DiR.

### 3.4. Joint Task-Aligned Refinement Learning

For training the alignment network, we simulate the Bayer RAW data from the degradation-free images  $\mathbf{y}^*$  (surrogate ground truth) by using the inverse ISP [21]. These RAW data

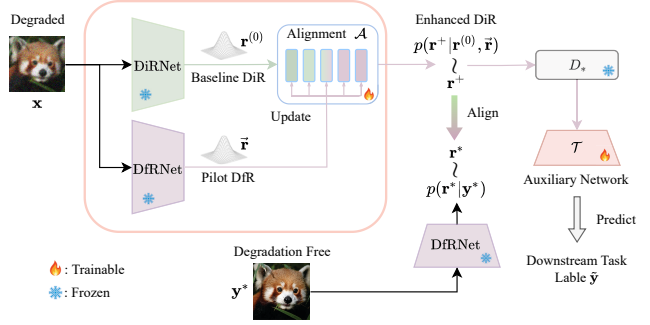


Figure 5. Illustration of the joint task-aligned refinement learning.

are used to generate the degraded RGB images  $\mathbf{x}$  with the ISP software tool [39]. Fig. 5 illustrates the joint task-aligned refinement learning process. The DiRNet and DfRNet are frozen to extract the paired representations  $(\mathbf{r}^{(0)}, \mathbf{r}^*)$  from the paired degraded and degradation-free images  $(\mathbf{x}, \mathbf{y}^*)$ . The alignment network is driven by the pairs  $(\mathbf{r}^{(0)}, \mathbf{r}^*)$  to learn the mapping function  $\mathcal{A} : \mathbf{r}^{(0)} \rightarrow \mathbf{r}^+$ , where the final refined DiR  $\mathbf{r}^+$  lies in the degradation-free latent representation space determined by the DfRNet.

The paired images  $(\mathbf{x}, \mathbf{y}^*)$  and the associated labels  $\tilde{\mathbf{y}}$  w.r.t. the downstream task are used to supervise the joint training of the alignment network  $\mathcal{A}$  and the task-related auxiliary network  $\mathcal{T}$ . As the auxiliary network for the downstream task usually takes an RGB image as input, we adopt the decoder  $D_*$  trained for the DfRNet to reconstruct the RGB images from the refined DiRs  $\mathbf{r}^+$ . The joint training loss in the alignment phase is to minimize  $\mathcal{L}_+$  defined as

$$\mathcal{L}_+ = \mathbb{E} [\|\mathbf{r}^+ - \mathbf{r}^*\|_1] + \gamma_1 \mathcal{L}_{\text{task}}(\mathcal{T}(D_*(\mathbf{r}^+)), \tilde{\mathbf{y}}) + \gamma_2 \mathbb{E} [\|D_*(\mathbf{r}^+) - \mathbf{y}^*\|_1], \quad (8)$$

where  $\gamma_1$  and  $\gamma_2$  are the trade-off parameters, the DfR  $\mathbf{r}^* = \text{DfR}(\mathbf{y}^*)$ , and the refined DiR representation  $\mathbf{r}^+ = \mathcal{A}(\mathbf{r}^{(0)}, \vec{r})$ , where  $\mathbf{r}^{(0)} = \text{DiR}(\mathbf{x})$  and  $\vec{r} = \text{DfR}(\mathbf{x})$ . The alignment network  $\mathcal{A}$  and the auxiliary network  $\mathcal{T}$  are jointly optimized, whereas the pretrained networks DiRNet and DfRNet as well as the decoder  $D_*$  are frozen. The second term  $\mathcal{L}_{\text{task}}$  is a task-related loss for a downstream task, such as the cross-entropy for image segmentation. The third term is an image reconstruction loss that forces the refined DiR to inherit the complete information from the degradation-free image through the pixel-wise reconstruction restriction. We summarize the DiR learning process in Algorithm 1.

## 4. Experiments

### 4.1. Implementation Details

**Datasets and evaluation** As previously introduced, DiR learning comprises two stages: Stage I, dedicated to training the DiRNet and DfRNet, and Stage II, where the alignment network  $\mathcal{A}$  and the task-related auxiliary network  $\mathcal{T}$  are jointly trained. For the training dataset, we adopt the

---

**Algorithm 1** DiR Learning and Task-Aligned Refinement

---

**Input:** Degradation-free RAW image  $\mathbf{y}$  and corresponding degradation-free RGB image  $\mathbf{y}^*$ , task-related label  $\tilde{\mathbf{y}}$

**Stage I:**

- 1: **repeat**
- 2: Generate two images  $\mathbf{x}_1$  and  $\mathbf{x}_2$  from a degradation-free image  $\mathbf{y}$  using two randomly sampled ISP degradations.
- 3: Train DiRNet using  $(\mathbf{x}_1, \mathbf{x}_2)$  ( $\min \mathcal{L}_0$  in Eq. (5))
- 4: Train DfRNet using  $\mathbf{y}^*$  ( $\min \mathcal{L}_*$  in Eq. (6))
- 5: **until** reaching the maximum number of epochs

**Stage II:**

- 6: **repeat**
  - 7: Generate an image  $\mathbf{x}$  from a degradation-free image  $\mathbf{y}$  using randomly sampled ISP degradations.
  - 8:  $\vec{\mathbf{r}} = \text{DfR}(\mathbf{x}), \mathbf{r}^* = \text{DfR}(\mathbf{y}^*)$
  - 9:  $\mathbf{r}^+ = \mathcal{A}(\mathbf{r}^{(0)}, \vec{\mathbf{r}})$
  - 10: Train the alignment network  $\mathcal{A}$  using the paired  $(\mathbf{x}, \mathbf{r}^*)$  with the guidance of the pilot DfR  $\vec{\mathbf{r}}$ , and jointly train the auxiliary  $\mathcal{T}$  using the pairs of  $(D_*(\mathbf{r}^+), \tilde{\mathbf{y}})$  ( $\min \mathcal{L}_+$  in Eq. (8))
  - 11: **until** reaching the maximum number of epochs
- 

DIV2K dataset [1] and degrade the images by synthetic ISP degradations. To simulate ISP degradations, we begin by converting RGB images into RAW data by applying the reversed ISP model [21]. Subsequently, we employ a synthetic ISP pipeline [32, 39] with random parameters to generate degraded ISP-generated images from the simulated RAW data<sup>3</sup>. In Stage II, we employ the identical simulation process used in Stage I to synthesize ISP-degraded images and use the MS COCO dataset [25] for experiments.

For evaluation, we test the proposed method on three downstream tasks, including 1) the image restoration experiments on real-world datasets including the PolyU-Real [50], NC12 [22], Nam [33], and DND<sup>4</sup> [37] datasets, where the test images involve real complex ISP degradations, 2) the object detection experiments on the MS COCO [25] dataset, and 3) the image segmentation experiments on the MS COCO dataset and the ADE20k dataset [55].

**Training settings** At Stage I, we randomly select the contrastive degraded pairs to enable the MMI maximization for the training of the DiRNet. For training the DfRNet, we follow the same setting as in [54], by setting the parameter  $\beta^* = 1$  in Eq. (6). We adopt the encoder network of [17] as the architectures of the DiRNet and the DfRNet. During training, the learning rate is initially set to  $10^{-4}$  and reduced to  $10^{-6}$  after 200 epochs. At Stage II, we randomly form pairs of degraded and degradation-free RGB images for training the alignment network. The learning rate is initially set to  $10^{-3}$  and reduced to  $10^{-6}$  after 300 epochs. For the image restoration task, we set  $\gamma_1 = 0, \gamma_2 = 1$ , because the decoder  $D_*$  can serve as the auxiliary network for the RGB image reconstruction using the refined DiR  $\mathbf{r}^+$ . For the object detection and image segmentation tasks, the trade-off parameters are set to  $\gamma_1 = 2, \gamma_2 = 1$  and  $\gamma_1 = 5, \gamma_2 = 1$ ,

<sup>3</sup>Additional details can be found in the supplementary material

<sup>4</sup>We randomly select 500 images for the experiments

respectively. The hyperparameters  $m_1, m_2$ , and  $m_3$  in the alignment network are separately set to 4, 2, and 4.

## 4.2. Results on ISP-Degraded Image Restoration

First, we present the experimental results of the proposed method and compare them with those of eight competing methods [5, 6, 8, 24, 36, 48, 53], on the restoration of ISP-degraded images. The task is to remove ISP artifacts caused by complex coupling of multiple degradation sources, including sensor noises, demosaicing noises, compression noises, and possibly adverse effects of incorrect settings of ISO and gamma correction. Here, like in all other image restoration methods for human viewers, our goal is to pursue the highest possible perceptual image quality.

In Fig. 6 the reader can compare the visual quality of the DiR learning based restoration method against those of SOTA image restoration methods [6, 8, 24, 26, 36, 48, 53], on restoring real-world ISP-degraded images taken by four different cameras, whose true degradation models are unknown. We can observe that the proposed method succeeds in achieving its design goal of blind restoration, as it removes most of the ISP artifacts caused by unknown degradation sources. On the contrary, the competing methods fail to remove complex artifacts while preserving the sharpness and clarity of edges and textures. Table 1 reports the quantitative performance evaluation results. Since there is no ideal degradation-free ground truth, we adopt the non-reference image quality assessment methods, including NIQE [31], NRQM [29], and BIQA [42]. In agreement with the visual comparisons in Fig. 6, the proposed DiR method also outperforms the competitors in terms of all three metrics. Both subjective and objective evaluations demonstrate superior robustness of the proposed method when being applied to data unseen in the training, indicating that the DiR learning approach can distill degradation-free representations from ISP-degraded images, regardless of whether observed or not.

Method	NIQE ↓ / NRQM ↑ / BIQA ↑			
	PolyU-Real	NC12	Nam	DND
BM3D [5]	9.72/2.60/40.5	12.78/3.75/30.8	12.58/2.56/41.2	12.89/1.84/40.1
CBDNet [8]	7.32/3.78/44.1	10.67/6.88/33.1	10.26/3.62/42.1	9.74/2.96/41.1
LIR [6]	8.37/4.52/48.1	11.41/7.19/33.7	10.69/4.25/49.8	11.86/3.19/44.3
MPRNet [53]	<b>6.12/4.23/38.7</b>	11.21/6.30/32.9	<b>7.25/3.94/33.9</b>	9.53/2.77/45.2
Uformer [48]	6.25/3.96/37.4	10.54/6.05/31.2	7.92/3.62/34.6	9.58/2.72/45.1
SCPGab [26]	8.71/3.25/41.1	11.75/3.93/31.9	10.60/3.25/40.7	9.80/2.53/42.0
DIL [24]	7.28/4.10/41.8	10.88/6.52/30.8	9.11/3.59/40.8	<b>7.86/3.72/44.7</b>
SDAP [36]	9.62/3.00/46.1	11.34/3.88/32.5	12.38/2.95/45.7	10.74/2.33/45.8
<b>Ours</b>	<b>6.57/4.91/51.9</b>	<b>9.06/7.41/34.4</b>	<b>8.96/4.64/53.6</b>	<b>8.62/3.59/46.9</b>

Table 1. Quantitative results on real-world ISP-degraded images.

## 4.3. Results on Object Detection

Following [39], we utilize the MS COCO dataset [25] to evaluate our proposed method’s performance on object detection. We compare our method with the SOTA neural ISP solvers [38, 39], which replace the conventional ISP pipeline

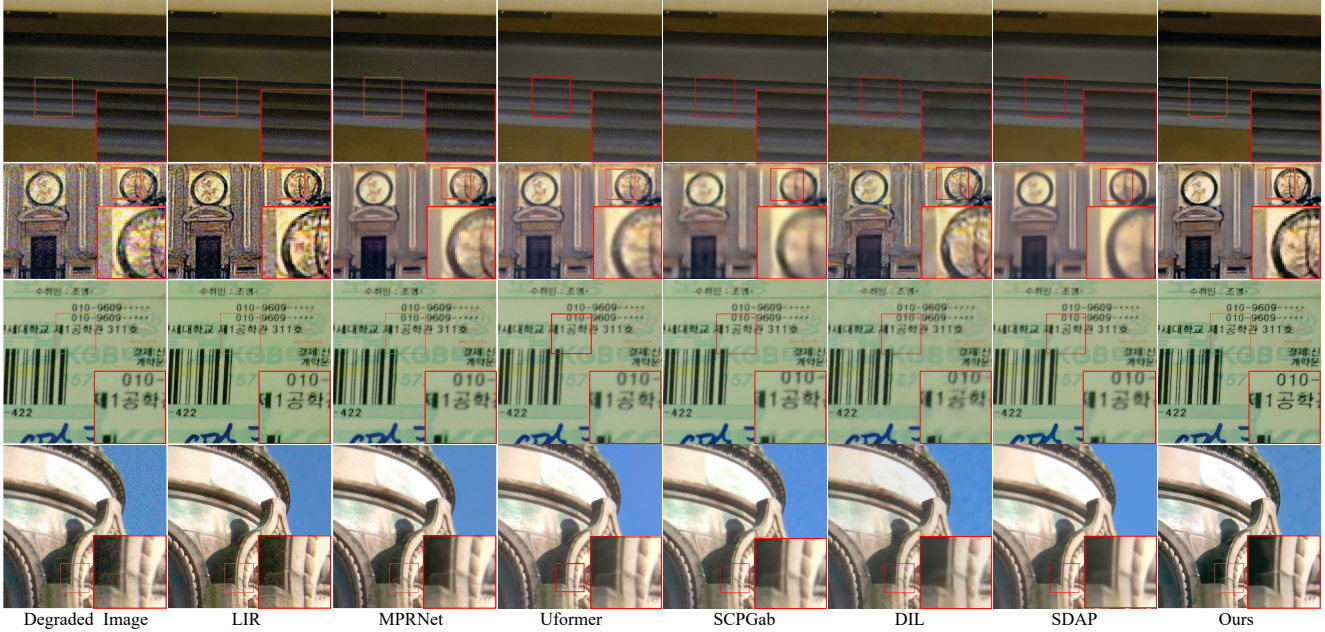


Figure 6. Qualitative comparison on real-world ISP-degraded images. The degraded images, from top to bottom, are from the real-world datasets PolyU-Real [50], NC12 [22], Nam [33], and DND [37]. Please zoom in for better visualization.

with end-to-end trained neural networks while considering the downstream task as their optimization objectives. Following the approach in [32, 39], we choose the ISP pipeline with default parameters as the baseline. Our method utilizes these ISP-generated results with default settings as inputs. For comparison, we adopt the pre-trained object detection model [40] for performance benchmarking, and we assess the results using the mean accuracy score (mAP) as the evaluation metric. The quantitative results are reported in Table 2. Our method demonstrates a significant improvement compared to the baseline and outperforms the SOTA solutions [38, 39]. The left side of Fig 7 illustrates the qualitative results. As a reference, we include the results produced by the expert-tuned ISP for visual comparison. We can observe that while the expert-tuned ISP can produce more visually appealing results aligned with human perception preferences, it is inferior to our method on the task of object detection.

#### 4.4. Results on Image Segmentation

Similar to the object detection experiments, we adopt the MS COCO dataset [25] to evaluate the performance of image segmentation. We benchmark the performance using a pre-trained segmentation model [13] and report the results in Table 3. As shown, our method outperforms the SOTA method Attention-aware [38] by about 3.9% on the mAP@0.5. Quantitative results are shown on the right side of Fig 7. These results are in line with the observations in the object detection experiments. Our method consistently outperforms the expert-tuned ISP in achieving better segmentation results.

Methods	ISP Model	mAP $\uparrow$		
		0.5	0.75	0.5:0.95
Default Params	ISP [32]	15.	-	-
Blockwise-tuned [34]		20.	-	-
Hardware-tuned [32]		39.	-	-
Default Params	ISP [39]	34.1	22.4	21.4
Attention-aware [38]		61.0	44.1	41.0
Sequential-tuned [39]		62.8	45.2	42.3
<b>Ours</b>		<b>63.9</b>	<b>46.4</b>	<b>43.3</b>

Table 2. Quantitative results on object detection.

Methods	ISP Model	mAP $\uparrow$		
		0.5	0.75	0.5:0.95
Default Params	ISP [32]	12.	-	-
Hardware-tuned [32]		32.	-	-
Default Params	ISP [39]	22.7	13.2	12.0
Attention-aware [38]		52.3	33.4	31.5
Sequential-tuned [39]		54.2	35.1	32.6
<b>Ours</b>		<b>56.3</b>	<b>37.1</b>	<b>34.8</b>

Table 3. Quantitative results on image segmentation.

Additionally, we conduct experiments to evaluate the performance of the proposed method in low signal-to-noise ratio situations, such as dark environments where the noise level is very high. To mimic the degradations, we increase the noise level in the aforementioned synthetic ISP pipeline by adding Gaussian noise with the standard deviation  $\sigma$  randomly sampled between the range of [0.15, 0.35] and the Poisson noise with the magnitude  $\lambda$  randomly sampled between the range of [0.02, 0.04]. We compare the proposed DiR learning based method with the SOTA methods [16, 20] for ISP degradation removal and present the results in Table 4. To evaluate performance, we use two pretrained segmentation models, including DeeplabV3 [3] and SegFormer [49]. We observe that our method outperforms all



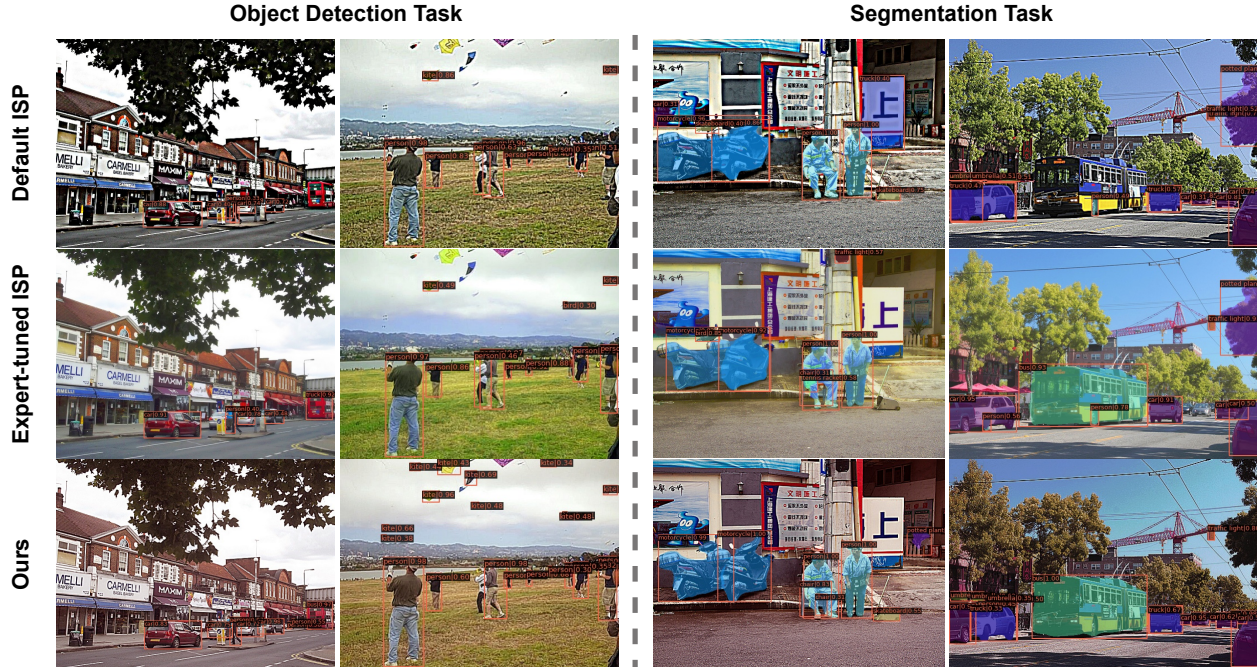


Figure 7. We evaluate different ISP-generated images for the tasks of object detection (left) and instance segmentation (right) on the COCO dataset. From top to bottom, we show the results from the images generated by default ISP hyperparameters, expert-tuned ISP hyperparameters, and the refined DiR with the proposed method. The proposed method achieves better performance for downstream tasks.

Methods	DeeplabV3 [3]			SegFormer [49]		
	aAcc $\uparrow$	mIoU $\uparrow$	mAcc $\uparrow$	aAcc $\uparrow$	mIoU $\uparrow$	mAcc $\uparrow$
Default ISP	64.8	22.6	28.9	74.8	36.9	45.1
PyNET [16]	67.7	25.1	31.5	74.3	37.4	45.8
DNF [20]	71.4	29.6	36.9	75.4	37.8	46.2
<b>Ours</b>	<b>75.5</b>	<b>34.6</b>	<b>44.5</b>	<b>79.7</b>	<b>42.4</b>	<b>52.9</b>

Table 4. Results of image segmentation on the ADE20k [55].

the competitors. For example, compared to the DNF [20], our method increases the mAcc by 7.6% and 6.7% when using the DeeplabV3 and SegFormer models, respectively.

#### 4.5. Ablation Study

Ablation studies are conducted to assess the effectiveness of the proposed components: the baseline DiR  $\mathbf{r}^{(0)}$ , the alignment network  $\mathcal{A}$ , and the pilot DfR  $\vec{\mathbf{r}}$ . For experiments, we apply the synthetic ISP process as described in [39] to generate ISP-degraded test images using the Set14 and BSD100 [30] datasets. To simulate degradation, we introduce noise into the synthetic ISP process. The added noise includes random Gaussian noise with a standard deviation ranging from 0.05 to 0.1 and JPEG compression noise with a quality factor randomly chosen in the range [10, 30]. As a reference, we train an auto-encoder model, labeled as “Auto-Encoder”. This model adopts the same encoder architecture as DiR-Net and the same decoder architecture as  $D_*$ . The training process of this auto-encoder model follows the canonical paradigm, *i.e.*, it learns a direct mapping from degraded images to their degradation-free counterparts. Table 5 shows comparisons: when excluding the pilot DfR (2nd row) from

the complete DiR learning technique (3rd row), we observe a decrease in PSNR by 0.2 dB and 0.4 dB for the Set14 and BSD100 datasets. In comparison to using only the baseline DiR, the inclusion of the alignment network  $\mathcal{A}$  leads to a PSNR improvement from 22.4 dB to 23.9 dB and 23.8 dB to 24.6 dB on the two datasets. This enhancement is attributed to the alignment network’s ability to refine the baseline DiR by aligning it with the degradation-free representation.

Method	Set14		BSD100	
	PSNR $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	LPIPS $\downarrow$
$\mathbf{r}^{(0)}$				
$\mathcal{A}$				
$\vec{\mathbf{r}}$				
$\mathcal{A}$ + $\vec{\mathbf{r}}$				
Auto-Encoder	21.6	0.378	23.1	0.408

Table 5. Ablation studies on each module in DiR learning.

## 5. Conclusion

This paper presents a novel method of learning degradation-independent representations for ISP-degraded images. In the proposed method, we devise a DiRNet to learn baseline DiR representations from diverse ISP-degraded observations in a self-supervised learning manner, ensuring adaptability to various degradation scenarios. Furthermore, we employ a guided alignment network to refine the baseline DiRs by aligning them with latent degradation-free representations. Our proposed approach demonstrates superior performance and scalability, as validated by extensive empirical results from synthetic and real-world datasets across various downstream applications.



## References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017. 6
- [2] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3291–3300, 2018. 2
- [3] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 7, 8
- [4] Ziteng Cui, Guo-Jun Qi, Lin Gu, Shaodi You, Zenghui Zhang, and Tatsuya Harada. Multitask aet with orthogonal tangent regularity for dark object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2553–2562, October 2021. 3
- [5] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007. 6
- [6] Wenchao Du, Hu Chen, and Hongyu Yang. Learning invariant representation for unsupervised image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14483–14492, 2020. 6
- [7] Marco Federici, Anjan Dutta, Patrick Forré, Nate Kushman, and Zeynep Akata. Learning robust representations via multi-view information bottleneck. In *International Conference on Learning Representations*, 2019. 2, 3
- [8] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward convolutional blind denoising of real photographs. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 6
- [9] Yanhui Guo, Xue Ke, Jie Ma, and Jun Zhang. A pipeline neural network for low-light image enhancement. *IEEE Access*, 7:13737–13744, 2019. 2
- [10] Yanhui Guo, Fangzhou Luo, and Shaoyuan Xu. Self-supervised face image restoration with a one-shot reference. *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2024. 2
- [11] Yanhui Guo, Xiaolin Wu, and Xiao Shu. Data acquisition and preparation for dual-reference deep learning of image super-resolution. *IEEE Transactions on Image Processing*, 31:4393–4404, 2022. 1
- [12] Yanhui Guo, Xi Zhang, and Xiaolin Wu. Deep multi-modality soft-decoding of very low bit-rate face videos. *Proceedings of the 28th ACM International Conference on Multimedia*, Oct 2020. 2
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2961–2969, 2017. 7
- [14] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016. 3, 4
- [15] HyeonJoo Hwang, Geon-Hyeong Kim, Seunghoon Hong, and Kee-Eung Kim. Variational interaction information maximization for cross-domain disentanglement. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 22479–22491. Curran Associates, Inc., 2020. 2, 3, 4
- [16] Andrey Ignatov, Luc Van Gool, and Radu Timofte. Replacing mobile camera isp with a single deep learning model. *arXiv preprint arXiv:2002.05509*, 2020. 7, 8
- [17] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1125–1134, 2017. 6
- [18] Yang Jiao, Yan Gao, Jingjing Meng, Jin Shang, and Yi Sun. Learning attribute and class-specific representation duet for fine-grained fashion analysis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11050–11059, 2023. 1
- [19] Yang Jiao, Ning Xie, Yan Gao, Chien-Chih Wang, and Yi Sun. Fine-grained fashion representation learning by online deep clustering. In *European Conference on Computer Vision*, pages 19–35. Springer, 2022. 1
- [20] Xin Jin, Ling-Hao Han, Zhen Li, Chun-Le Guo, Zhi Chai, and Chongyi Li. Dnf: Decouple and feedback network for seeing in the dark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18135–18144, 2023. 7, 8
- [21] 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. 5, 6
- [22] Marc Lebrun, Miguel Colom, and Jean-Michel Morel. The noise clinic: a blind image denoising algorithm. *Image Processing On Line*, 5:1–54, 2015. <https://doi.org/10.5201/ipol.2015.125>. 6, 7
- [23] Anat Levin, Yair Weiss, Fredo Durand, and William T. Freeman. Understanding and evaluating blind deconvolution algorithms. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1964–1971, 2009. 2
- [24] Xin Li, Bingchen Li, Xin Jin, Cuiling Lan, and Zhibo Chen. Learning distortion invariant representation for image restoration from a causality perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1714–1724, 2023. 2, 6
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014. 6, 7
- [26] Xin Lin, Chao Ren, Xiao Liu, Jie Huang, and Yinjie Lei. Unsupervised image denoising in real-world scenarios via self-collaboration parallel generative adversarial branches. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 2, 6
- [27] Fangzhou Luo, Xiaolin Wu, and Yanhui Guo. Functional neural networks for parametric image restoration problems. In *Thirty-fifth Annual Conference on Neural Information Processing Systems (NeurIPS)*, volume 34, pages 6762–6775, 2021. 2
- [28] Fangzhou Luo, Xiaolin Wu, and Yanhui Guo. And: Ad-

- versarial neural degradation for learning blind image super-resolution. In *Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS)*, 2023. 2
- [29] Chao Ma, Chih-Yuan Yang, Xiaokang Yang, and Ming-Hsuan Yang. Learning a no-reference quality metric for single-image super-resolution, 2016. 6
- [30] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision (ICCV)*, volume 2, pages 416–423, July 2001. 8
- [31] Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2013. 6
- [32] Ali Mosleh, Avinash Sharma, Emmanuel Onzon, Fahim Mannan, Nicolas Robidoux, and Felix Heide. Hardware-in-the-loop end-to-end optimization of camera image processing pipelines. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 3, 6, 7
- [33] Seonghyeon Nam, Youngbae Hwang, Yasuyuki Matsushita, and Seon Joo Kim. A holistic approach to cross-channel image noise modeling and its application to image denoising. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1683–1691, 2016. 6, 7
- [34] Jun Nishimura, Timo Gerasimow, Rao Sushma, Aleksandar Sutic, Chyuan-Tyng Wu, and Gilad Michael. Automatic isp image quality tuning using nonlinear optimization. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 2471–2475. IEEE, 2018. 2, 7
- [35] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 271–279, 2016. 3
- [36] Yizhong Pan, Xiao Liu, Xiangyu Liao, Yuanzhouhan Cao, and Chao Ren. Random sub-samples generation for self-supervised real image denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12150–12159, October 2023. 2, 6
- [37] Tobias Plötz and Stefan Roth. Benchmarking denoising algorithms with real photographs, 2017. 6, 7
- [38] Haina Qin, Longfei Han, Juan Wang, Congxuan Zhang, Yanwei Li, Bing Li, and Weiming Hu. Attention-aware learning for hyperparameter prediction in image processing pipelines. In *European Conference on Computer Vision*, pages 271–287. Springer, 2022. 1, 6, 7
- [39] Haina Qin, Longfei Han, Weihua Xiong, Juan Wang, Wentao Ma, Bing Li, and Weiming Hu. Learning to exploit the sequence-specific prior knowledge for image processing pipelines optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22314–22323, 2023. 3, 5, 6, 7, 8
- [40] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 1, 7
- [41] 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
- [42] Shaolin Su, Qingsen Yan, Yu Zhu, Cheng Zhang, Xin Ge, Jinqiu Sun, and Yanning Zhang. Blindly assess image quality in the wild guided by a self-adaptive hyper network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 6
- [43] Naftali Tishby, Fernando C. Pereira, and William Bialek. The information bottleneck method, 2000. 2, 3
- [44] Ethan Tseng, Felix Yu, Yuting Yang, Fahim Mannan, Karl St. Arnaud, Derek Nowrouzezahrai, Jean-Francois Lalonde, and Felix Heide. Hyperparameter optimization in black-box image processing using differentiable proxies. *ACM Transactions on Graphics (TOG)*, 38(4), 7 2019. 2
- [45] Aaron Van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv e-prints*, pages arXiv–1807, 2018. 3
- [46] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(11), 2008. 5
- [47] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, pages 1905–1914, 2021. 2
- [48] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17683–17693, June 2022. 6
- [49] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34, 2021. 1, 7, 8
- [50] Jun Xu, Hui Li, Zhetong Liang, David Zhang, and Lei Zhang. Real-world noisy image denoising: A new benchmark, 2018. 6, 7
- [51] Masakazu Yoshimura, Junji Otsuka, Atsushi Irie, and Takeshi Ohashi. Rawgment: noise-accounted raw augmentation enables recognition in a wide variety of environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14007–14017, 2023. 2
- [52] Ke Yu, Zexian Li, Yue Peng, Chen Change Loy, and Jinwei Gu. Reconfigisp: Reconfigurable camera image processing pipeline. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4248–4257, 2021. 2, 3
- [53] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *CVPR*, 2021. 6
- [54] Shengjia Zhao, Jiaming Song, and Stefano Ermon. Infovae: Information maximizing variational autoencoders. *arXiv preprint arXiv:1706.02262*, 2017. 4, 6
- [55] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5122–5130, 2017. 6, 8