

Implicit Neural Representation for Cooperative Low-light Image Enhancement

Shuzhou Yang^{1,2}, Moxuan Ding[†], Yanmin Wu¹, Zihan Li³, Jian Zhang^{*1}

¹Peking University Shenzhen Graduate School, China

²Peng Cheng Laboratory, China ³University of Washington, USA

szyang@stu.pku.edu.cn, zhangjian.sz@pku.edu.cn

Abstract

The following three factors restrict the application of existing low-light image enhancement methods: unpredictable brightness degradation and noise, inherent gap between metric-favorable and visual-friendly versions, and the limited paired training data. To address these limitations, we propose an implicit **Neural Representation** method for **Cooperative** low-light image enhancement, dubbed **NeRCo**. It robustly recovers perceptual-friendly results in an unsupervised manner. Concretely, NeRCo unifies the diverse degradation factors of real-world scenes with a controllable fitting function, leading to better robustness. In addition, for the output results, we introduce semantic-oriented supervision with priors from the pre-trained vision-language model. Instead of merely following reference images, it encourages results to meet subjective expectations, finding more visual-friendly solutions. Further, to ease the reliance on paired data and reduce solution space, we develop a dual-closed-loop constrained enhancement module. It is trained cooperatively with other affiliated modules in a self-supervised manner. Finally, extensive experiments demonstrate the robustness and superior effectiveness of our proposed NeRCo. Our code is available at <https://github.com/Ysz2022/NeRCo>.

1. Introduction

Due to the degraded brightness in low-light images covering objects and reducing contrast, low-light images have severely impacted the subsequent high-level computer visual tasks (e.g., semantic segmentation [15], and object detection [26], etc.). Hence, it is of practical importance to remedy the brightness degradation for assisting the exploration of sophisticated dark environment. Low-light image enhancement, which aims to recover the desired content in degraded regions, has drawn wide attention in recent years [10, 11, 12, 16, 27, 34].

Corresponding author*. Independent researcher[†].

This work was supported in part by Shenzhen General Research Project under Grant JCYJ20220531093215035.




Figure 1: Comparison with two state-of-the-art methods on the LIME [12] dataset. One can see that we recovered more authentic color and visual-friendly contents.

Over the past few years, prolific algorithms have been proposed to address this classic ill-posed problem, which can be roughly categorized into two groups: conventional model-based methods (e.g., gamma correction [32], Retinex-based model [33], and histogram equalization [34]) and recent deep learning-based methods [16, 25, 47]. The former formulates the degradation as a physical model and treats enhancement as the problem of estimating model parameters, but is limited in characterizing diverse low-light factors and requires massive hand-crafted priors. The latter elaborates various models to adjust tone and contrast, which is able to learn from massive data automatically. Essentially, they are trained to learn a mapping from input to output domain. In real-world scenarios, however, many samples are far away from the feature space of input domain, causing a trained model to lack stable effect. We propose to normalize the degradation before enhancement to bring these samples closer to the input domain. Besides, existing supervised methods highly rely on paired training data and mainly attempt to produce metric-favorable results,

i.e., similar to the ground truth. But the limited supervised datasets and the inherent gap between metric-oriented and visual-friendly versions inevitably impact their effectiveness. We develop a cooperative training strategy to address it. As shown in Fig. 1, we test on the LIME [12] dataset, which only consists of low-light images without normal-light references. One can see that even the recently proposed top-performing algorithms perform severe color cast.

Specifically, our key insights are: **i) Normalizing the input with a controllable fitting function to reduce the unpredictable degradation features in real-world scenarios.** We adopt neural representation to reproduce the degraded scene before the enhancement operation. By manipulating the positional encoding, we selectively avoid regenerating extreme degradation, which objectively realizes normalization and thereby decreases enhancement difficulty. **ii) Supervising the output with different modalities to achieve both metric-favorable and perceptual-oriented enhancement.** We employ multi-modal learning to supervise from both textual and image perspectives. Compared with image supervision, which contains varying brightness across different samples, the feature space of the designed prompt is more stable and accurate in describing brightness. During training, our results are not only encouraged to be similar to references, but also forced to match their related prompts. In this way, we bridge the gap between the metric-favorable and the perceptual-friendly versions. **iii) Developing an unsupervised training strategy to ease the reliance on the paired data.** We propose to train the enhancement module with a dual-closed-loop cooperative adversarial constraint procedure, which learns in an unsupervised manner. More related loss functions are also proposed to further reduce the solution space. Benefiting from these, we recover more authentic tone and better contrast (see Fig. 1). Overall, our contributions are as follows:

- We are the first to utilize the controllable fitting capability of neural representation in low-light image enhancement. It normalizes lightness degradation and removes natural noise without any additional operations, providing new ideas for future work.
- For the first time, we introduce multi-modal learning to low-light image enhancement. Benefiting from its efficient vision-language priors, our method learns diverse features, resulting in perceptually better results.
- We develop an unsupervised cooperative adversarial learning strategy to ease the reliance on the paired training data. In which the appearance-based discrimination ensures authenticity from both color and detail levels, improving the quality of the restored results.
- Extensive experiments are conducted on representative benchmarks, manifesting the superiority of our NeRCo

against a rich set of state-of-the-art algorithms. Especially, it even outperforms some supervised methods.

2. Related Work

2.1. Low-light Image Enhancement

To improve the visibility of low-light images, model-based methods are first widely adopted. Retinex theory [43] decomposes the observation into illumination and reflectance (*i.e.*, clear prediction), but tends to over-expose the appearance. Various hand-crafted priors are further introduced into models as regularization terms. Fu *et al.* [10] developed a weighted variational model to simultaneously estimate reflectance and illumination layers. Cai *et al.* [2] proposed an edge-preserving smoothing algorithm to model brightness. Guo *et al.* [12] predicted the illumination by adopting the relative total variation [48]. However, these defined priors are labor-intensive and perform poor generalization towards real-world scenarios.

Due to these limitations, researchers took advantage of deep learning to recover in a data-driven manner [3, 11, 25, 47], which exploits priors from massive data automatically. For example, Guo *et al.* [11] formulated light enhancement as a task of image-specific curve estimation with a lightweight deep model. Jiang *et al.* [16] introduced adversarial training for learning from unpaired supervision. Wei *et al.* [46] designed an end-to-end trainable RetinexNet but still troubled by heavy noise. To ameliate it, Zhang *et al.* [53] proposed a decomposition-type architecture to impose constraint on reflectance. Liu *et al.* [25] employed architecture search and built an unrolling network. Although these well-designed models have realized impressive effectiveness, they are not stable in real-world applications. To improve robustness, we pre-modulate the degradation to a uniform level with neural representation before the enhancement procedure.

2.2. Neural Representation for Images

Recently, neural representation has been widely adopt to depict images. Chen *et al.* [5] firstly utilized implicit image representation for continuous image super-resolution. However, the MultiLayer Perceptron (MLP) tends to distort high-frequency components. To address this issue, Lee *et al.* [20] developed a dominant-frequency estimator to predict local texture for natural images. Lee *et al.* [19] further utilized implicit neural representation to warp images into continuous shape. Dupont *et al.* [9] tried to produce different objects with one MLP by manipulating the latent code from its hidden layers. Saragadam *et al.* [40] adopted multiple MLPs to represent a single image in a multi-scale manner. Sun *et al.* [41] predicted continuous information based on the captured tomographic features. Tancik *et al.* [42] introduced meta-learning to initialize the parameters of MLP

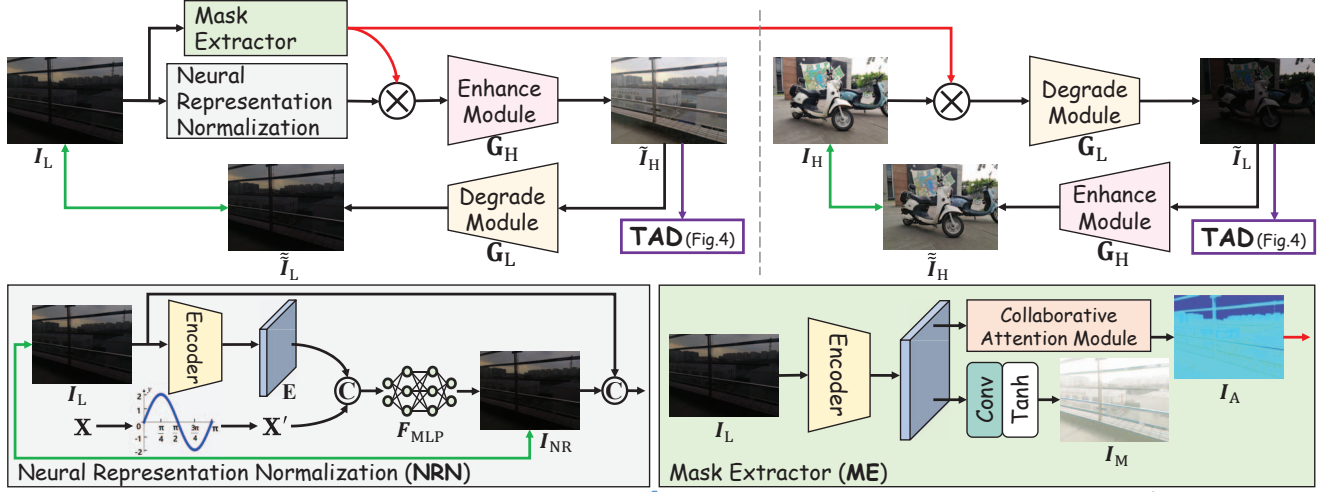


Figure 2: Workflow of our NeRCo. It presents a cooperative adversarial enhancement process containing dual-closed-loop branches, each of which contains an enhancement operation and a degradation operation. We embed a Mask Extractor (ME) to portrait the degradation distribution and a Neural Representation Normalization (NRN) module to normalize the degradation condition of the input low-light image. All of them are trained together to constrain each other, locking on to a more accurate target domain. The red means the transfer of the attention map.

to accelerate training. Reed *et al.* [38] adopted neural representation and parametric motion fields to predict the shape and location of organs. Further, some researchers adopted neural representation to compress videos [1, 4, 54].

However, existing neural representation is mainly applied on image compressing, denoising and depicting continuous information, *etc.* We are the first to apply its controllable fitting capability to low-light image enhancement.

2.3. Multi-modal Learning

In recent years, learning across modalities has attracted extensive attention [6, 7, 22, 24, 49]. Various vision-language models are developed. Radford *et al.* [35] proposed to learn visual model from language supervision, called CLIP. After training on 400 million image-text pairs, it can describe any visual concept with natural language and transfer to other tasks without any specific training. Furthermore, Zhou *et al.* [56] developed soft prompts to replace the hand-crafted ones, which uses learnable vectors to model context words and obtains task-relevant context. To further refine prompts to the instance-level, Rao *et al.* [37] designed context-aware prompting to combine prompts with visual features. Cho *et al.* [8] shared priors across different tasks by updating a uniform framework towards a common target of seven multi-modal tasks. Ju *et al.* [17] adopted the pre-trained CLIP model to video understanding.

Existing methods mainly focus on high-level computer vision tasks such as image classification. For the first time, we apply priors of the pre-trained vision-language model to low-light image enhancement, developing semantic-oriented guidance and realizing better performance.

3. Our Method

3.1. Framework Architecture

As shown in Fig. 2, given a low-light image I_L , we first normalize it by neural representation (NRN, Sec. 3.2) to improve the robustness of the model to different degradation conditions. Then the Mask Extractor (ME, Sec. 3.4) module extracts the attention mask from the image to guide the enhancement of different regions. After that, the Enhance Module G_H (represented by ResNet) generates a high-light image \tilde{I}_H . To ensure its quality, we design a Text-driven Appearance Discriminator (TAD, Sec. 3.3) to supervise image generation, where text-driven supervision guarantees semantic reliability, and appearance supervision guarantees visual reliability. \tilde{I}_H is then passed through the Degrade Module G_L to convert back to the low-light domain \tilde{I}_L and calculate the consistency loss (Sec. 3.4) with the original low-light image I_L . The upper right branch of Fig. 2 inputs the high-light image I_H , implemented in a similar way.

The network is realized in a dual-loop way (Sec. 3.4) to achieve stable constraints based on the unpaired data. It operates bidirectional mapping: enhance-degrade ($I_L \rightarrow \tilde{I}_H \rightarrow \tilde{I}_L$) and degrade-enhance ($I_H \rightarrow \tilde{I}_L \rightarrow \tilde{I}_H$). This dual loop constraint fully exploits the latent general distinction between the low-light and high-light domains. Besides, the cooperative loss (Sec. 3.4) encourages all components in the framework to supervise each other collaboratively, which further reduces the solution space.

During **training**, we run the whole process in Fig. 2. We input two images (*i.e.*, low-light I_L and high-light I_H). I_L is

enhanced to \tilde{I}_H , then translated back to low-light \tilde{I}_L . And vice versa for I_H . Noting that I_H is used for training purposes only, *i.e.*, training model in an unsupervised manner for better enhancement rather than degradation. Hence, we only use NRN to enhance I_L but not degrade I_H . For **inference**, I_L is directly enhanced to \tilde{I}_H without any other operations, as shown in the top left part of Fig. 2.

3.2. Neural Representation for Normalization

Motivation. Images captured in real-world typically exhibit varying degradation levels due to lighting conditions or camera parameters, as shown in Fig. 3 (a)(b). We report their pixel value distribution on the Y channel in Fig. 3 (c). The inconsistency between these samples is challenging for a well-trained model. We attempt to normalize degradation level (see Fig. 3 (d)(e)) with neural representations (NR) to obtain a more consistent degradation distribution (see Fig. 3 (f)) to reduce the difficulty of subsequent operations.

Neural Representation. Concretely, in NR, image I_L is transformed into a feature map $\mathbf{E} \in \mathbb{R}^{H \times W \times C}$, where H and W are image resolution. While the location of each pixel is recorded in a coordinate set $\mathbf{X} \in \mathbb{R}^{H \times W \times 2}$, where 2 means horizontal and vertical coordinates. I_L can thus be represented by its features and a set of coordinates. As shown in the Neural Representation Normalization (NRN) module of Fig. 2, we fuse \mathbf{X} and \mathbf{E} , and use a decoding function F_{MLP} to output image I_{NR} , which is parameterized as a MultiLayer Perceptron (MLP). The neural representation of the image is expressed as:

$$I_{NR}[i,j] = F_{MLP}(\mathbf{E}[i,j], \mathbf{X}[i,j]), \quad (1)$$

where $[i,j]$ is the location of a pixel and $I_{NR}[i,j]$ is the generated RGB value. By predicting RGB of each pixel, an image I_{NR} is reproduced. We encourage I_{NR} to be similar to I_L through l_1 -norm. This NR-related loss expressed as:

$$\mathcal{L}_{NR} = \|I_{NR} - I_L\|_1. \quad (2)$$

Why Neural Representation Works. With the trained F_{MLP} , each feature map \mathbf{E} can form a function $F_{MLP}(\mathbf{E}, \cdot) : \mathbf{X} \rightarrow I_{NR}$, which maps coordinates to its predicted RGB values. Without \mathbf{E} , it is impossible for F_{MLP} to depict various RGB values with the same coordinates \mathbf{X} . Without \mathbf{X} , we cannot normalize degradation by adjusting fitting capability, which is explained below.

According to [36], neural networks tend to portrait lower frequency information. Despite our decoding function can approximate RGB values, some high-frequency components may be discarded during rendering. For example, for adjacent pixels around the edge, their RGB values vary a lot but coordinates vary little. It means F_{MLP} should output different results based on similar inputs, which is difficult. Inspired by [30], to fit high-frequency variation better, we

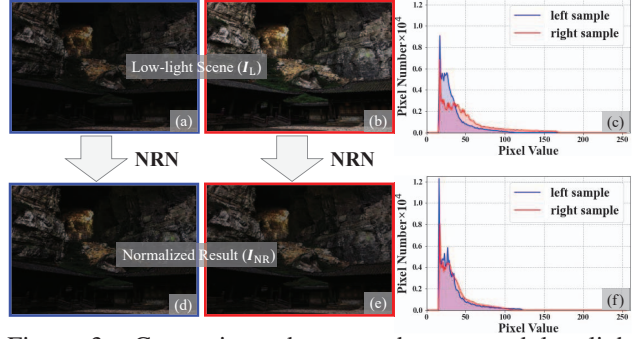


Figure 3: Comparisons between the captured low-light scenes (top row) and the results of NRN (bottom row). The low-light samples are from the SICE [3] dataset. It contains numerous image sets, each with common content and varying lighting conditions. The pixel value distribution of images on the Y channel is given on the right. One can see that NRN normalizes the brightness to be similar.

map the input coordinates to a higher dimensional space before passing them to F_{MLP} , which is called positional encoding. As shown in the gray region of Fig. 2, before fusing coordinates with image feature, we use a high-frequency function $\gamma(\cdot)$ to map the original coordinates \mathbf{x} from \mathbb{R} into a higher dimensional space \mathbb{R}^{2L} , expressed as:

$$\gamma(\mathbf{x}) = (\cdots, \sin(2^i \pi \mathbf{x}), \cos(2^i \pi \mathbf{x}), \cdots), \quad (3)$$

where i values from 0 to $L - 1$, L is a hyperparameter that determines the dimension value. The final coordinates are composed as: $\mathbf{x}' = \gamma(\mathbf{x})$. Noting that by manipulating the value of L , we can change the fitting capacity of our NRN module, *i.e.*, a bigger L results in a more precise fit.

However, stronger fitting capability is not always better. Since NRN aims to normalize various degradation, its output is not expected to be exactly the same as the input, especially the degradation components. We want to choose an L that does not overfit to remain all information while faithfully preserves desired content. Chen *et al.* [4] have demonstrated that NR is robust to perturbations and can denoise without any special design. We think this is because MLPs lack spatial correlation priors, some extreme information is thus hard to be reproduced faithfully. Hence, this underfitting property objectively limits the unpredictable degradation. We further found in experiments that MLPs tend to learn an average brightness range of training data rather than fitting the unique brightness of different images. As shown in Fig. 3, the output lightness of a trained MLP is similar. By visualizing their pixel value distribution in the Y channel, we further prove the normalization property of our NRN. More results are given in the [supplementary material \(SM\)](#). Thereby, we set $L = 8$ for the trade-off between degradation normalization and content fidelity. By reducing these fickle degradation signals, we decrease the difficulty of subsequent enhancement procedure.

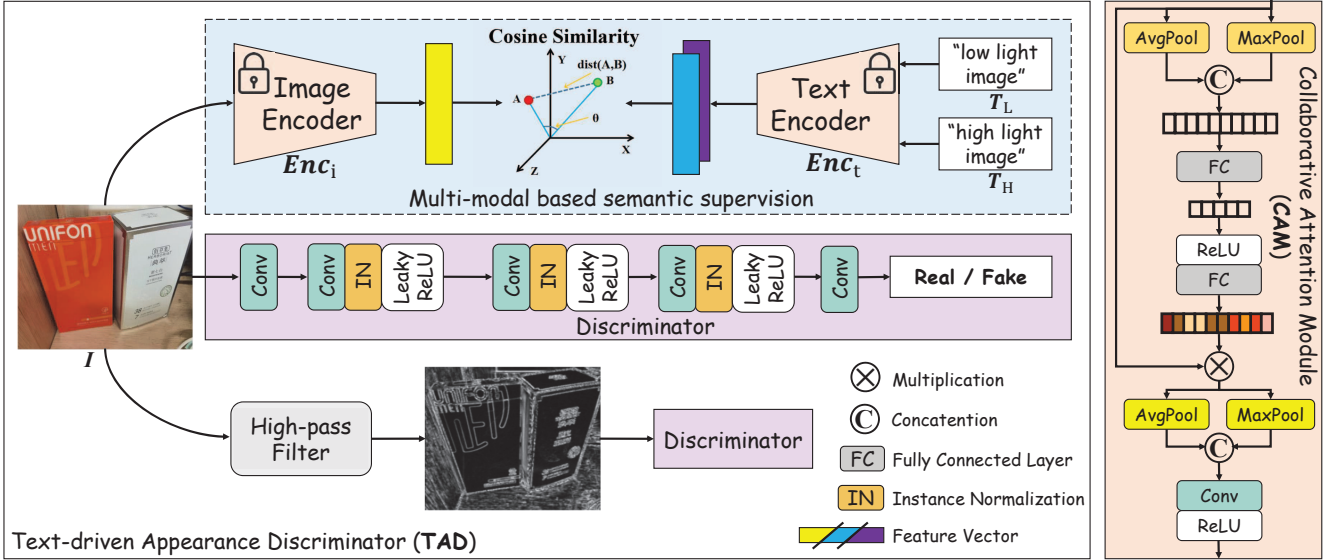


Figure 4: The details of our proposed text-driven appearance discriminator (in the left region) and our collaborative attention module (in the right region). The former supervises the input with text and image modalities, and focuses on high-frequency components. The latter adjusts attention to different channels and outputs attention map.

3.3. Text-driven Appearance Discriminator

Motivation. Existing methods mainly adopt image-level supervision, *i.e.*, forcing the output to be close to the target images. However, the brightness across different references varies greatly, confusing model training; some references are visually poor (*e.g.*, unnatural lightness), leading to visual-unfriendly results. To reduce training difficulty and bridge the gap between the metric-favorable and visual-friendly versions, we design a Text-driven Appearance Discriminator (TAD) to supervise image generation from the semantic level and appearance level, respectively.

Text-driven Discrimination. We denote the low-light domain as \mathcal{L} and the high-light domain as \mathcal{H} . As shown in Fig. 4, we introduce multi-modal learning to supervise images with both image and text modalities. Concretely, inspired by Radford *et al.* [35], we employ the recent well-known CLIP model to get efficient priors. It consists of two pre-trained encoders, *i.e.*, Enc_t for text and Enc_i for image. We first manually design two prompts, *i.e.*, *low-light image* and *high-light image*, to describe \mathcal{L} and \mathcal{H} respectively, denoted as T_L and T_H in Fig. 4. Experiments on more other texts are given in the SM. Enc_t extracts two feature vectors of size 1×512 from two prompts. Similarly, Enc_i extracts a vector of the same size from our intermediate result I . We compute the cosine similarity between the image vector and the text vector to measure their discrepancy, formulated as:

$$D_{cos}(I, T) = \frac{\langle Enc_i(I), Enc_t(T) \rangle}{\|Enc_i(I)\| \|Enc_t(T)\|}, \quad (4)$$

where I denotes the predicted image and T is the prompt. For the enhanced results (*e.g.*, \tilde{I}_H and $\tilde{\tilde{I}}_H$ in Fig. 2), we en-

courage their vectors to be similar to those of T_H and away from T_L , and vice versa for low-light predictions. In this way, we encourage semantically consistent outputs. This cosine objective function is formulated as:

$$\mathcal{L}_{cos}(I_H, T_L, T_H) = D_{cos}(I_H, T_L) - D_{cos}(I_H, T_H), \quad (5)$$

$$\mathcal{L}_{cos}(I_L, T_L, T_H) = D_{cos}(I_L, T_H) - D_{cos}(I_L, T_L). \quad (6)$$

Appearance-based Discrimination. Admittedly, text descriptions cannot provide low-level guidance like images. To generate faithful content, image supervision is necessary. As shown in the purple region of Fig. 4, we stack a discriminator to distinguish the predicted results from real images, encouraging image-level authenticity (*e.g.*, color, texture, and structure). Considering detail distortion in image processing, we embed a high-frequency path consisting of a high-pass filter and a discriminator of the same structure. The filter extracts high-frequency components and the discriminator supervises them at edge-level. Based on this double-path color-edge discriminative structure, we realize the trade-off between color and detail.

During training, TAD plays an adversarial role in learning bidirectional mapping relationship between \mathcal{L} and \mathcal{H} . We develop an adversarial loss on each generation loop to realize it. As shown in Fig. 2, for the enhancement operation $G_H(\cdot): I_L \rightarrow \tilde{I}_H$, we apply a TAD module and dub its appearance discrimination as D_H . Developing an adversarial objective function as:

$$\begin{aligned} \mathcal{L}_{adv}(I_L, I_H, G_H, D_H) &= \mathcal{L}_{cos}(G_H(I_L), T_L, T_H) \\ &+ \mathbb{E}_{I_H \sim \mathcal{H}}[\log D_H(I_H)] \\ &+ \mathbb{E}_{I_L \sim \mathcal{L}}[\log(1 - D_H(G_H(I_L)))], \end{aligned} \quad (7)$$

where D_H aims to determine whether an image is captured or generated, that is, distinguish the enhanced results $G_H(I_L)$ from the real high-light domain \mathcal{H} . While $G_H(\cdot)$ aims at deceiving D_H , *i.e.*, generating results close to \mathcal{H} . Simultaneously, the aforementioned cosine constraint is also adopted to supervise $G_H(\cdot)$. The reverse mapping $G_L(\cdot): I_H \rightarrow \tilde{I}_L$ adopts a similar objective function, which is supervised by another TAD.

In both functions, generators try to minimize the objective, whilst TAD modules maximize it. During this adversarial learning, our model realizes better semantic consistency, which is proved by experiments presented in SM.

3.4. Dual Loop Generation Procedure

Previous methods mainly map the low-light image I_L to its high-light version \tilde{I}_H directly. To provide stable constraint without the paired data, we stack a forward enhancement module with a backward degradation one for bidirectional mapping, which operates in an unsupervised manner.

Dual Loop. Specifically, the forward enhancement procedure aims to realize the mapping $G_H(\cdot): I_L \rightarrow \tilde{I}_H$. While the other does the opposite, *i.e.*, depicting a low-light scene from a clean image I_H with the mapping $G_L(\cdot): I_H \rightarrow \tilde{I}_L$. Our generation loop is composed of alternating these two operations. As shown in the left end of Fig. 2, our input is the observed low-light image. It is first extracted an attention-guidance I_A and normalized by a neural representation module. Then the subsequent procedure first translates it to the high-light domain, and maps the enhanced image \tilde{I}_H back to the low-light version $\tilde{\tilde{I}}_L$. This enhancement-degradation generation branch is formulated as:

$$\tilde{\tilde{I}}_L = G_L(\tilde{I}_H) = G_L(G_H(ME(I_L) \otimes NRN(I_L))), \quad (8)$$

where $G_H(\cdot)$ and $G_L(\cdot)$ denote enhancement and degeneration operations, respectively, and NRN is the neural representation normalization module discussed in Sec. 3.2. ME means mask extractor, as shown in the green region of Fig. 2, in which we develop a Collaborative Attention Module (CAM) to extract the attention map I_A . The details of CAM are displayed at the right end of Fig. 4.

The degeneration-enhancement generation branch, as shown in the right end of Fig. 2, is formulated similarly:

$$\tilde{\tilde{I}}_H = G_H(\tilde{I}_L) = G_H(G_L(ME(I_L) \otimes I_H)). \quad (9)$$

To constrain this bidirectional mapping, during training, we develop cycle consistency to directly impose supervision at the pixel-level. For example, for the left branch in Fig. 2, we ensure that: $I_L \approx \tilde{\tilde{I}}_L = G_L(G_H(I_L))$. Accordingly, the other cycle follows: $I_H \approx \tilde{\tilde{I}}_H = G_H(G_L(I_H))$. We adopt the l_1 -norm to measure discrepancy and develop the




Figure 5: Details of the proposed cooperative loss function.

consistency constraint as:

$$\begin{aligned} \mathcal{L}_{con} = & \mathbb{E}_{I_L \sim \mathcal{L}} [\|G_L(G_H(I_L)) - I_L\|_1] \\ & + \mathbb{E}_{I_H \sim \mathcal{H}} [\|G_H(G_L(I_H)) - I_H\|_1]. \end{aligned} \quad (10)$$

Cooperative Loss. To reduce solution space and enhance attention guidance, we elaborate Cooperative Loss (CL). Inspired by Lee *et al.* [21], this function trains different modules cooperatively by imposing mutual constraints. As shown in Fig. 5, we indirectly supervise attention guidance I_A and provide stronger constraints for all modules.

Concretely, as shown in Mask Extractor (ME) of Fig. 2, since attention map I_A is generated based on the extracted image features, the content of features heavily impacts the quality of I_A . To get better I_A , we generate a **lightness mask** I_M from the same features and co-supervise it with other information, including our enhanced image.

As shown in the left end of Fig. 5, for the low-light input I_L , on the one hand, we extract I_M with ME, and obtain a pseudo-high-light image \tilde{I}_H by adding I_M with I_L . On the other hand, we subtract I_L at the pixel level from the predicted high-light result \tilde{I}_H , obtaining a pseudo-lightness mask \tilde{I}_M . By using consistency loss, we encourage the estimated \tilde{I}_M to be consistent with the extracted I_M , the calculated \tilde{I}_H to be similar to the depicted \tilde{I}_H , expressed as:

$$\mathcal{L}_{rec1} = \|\tilde{I}_M - I_M\|_1 + \|\tilde{I}_H - \tilde{I}_H\|_1. \quad (11)$$

For another branch of given the high-light input I_H , as shown in the right part of Fig. 5, the loss is similar:

$$\mathcal{L}_{rec2} = \|\tilde{\tilde{I}}_H - I_H\|_1 + \|I_H - \tilde{I}_L - \tilde{I}_M\|_1. \quad (12)$$

For the calculated images based on I_M , *i.e.*, \tilde{I}_H and \tilde{I}_L , we further use the double-path discriminator from our TAD to inspect their authenticity, formulated as:

$$\begin{aligned} \mathcal{L}_{insp}(\tilde{I}_H, \tilde{I}_L, D_H, D_L) = & \mathbb{E}_{\tilde{I}_H \sim \tilde{\mathcal{H}}} [\log(1 - D_H(\tilde{I}_H))] \\ & + \mathbb{E}_{\tilde{I}_L \sim \tilde{\mathcal{L}}} [\log(1 - D_L(\tilde{I}_L))], \end{aligned} \quad (13)$$

where $\tilde{\mathcal{H}}$ and $\tilde{\mathcal{L}}$ mean the pseudo high-light and low-light image domain respectively. The final cooperative loss is:

$$\mathcal{L}_{CL} = \mathcal{L}_{rec1} + \mathcal{L}_{rec2} + \mathcal{L}_{insp}. \quad (14)$$

Combined with the NR-related loss in Sec. 3.2, the adversarial constraint in Sec. 3.3, and the consistency loss in Sec. 3.4, our final objective function is expressed as:

$$\mathcal{L} = \mathcal{L}_{NR} + \mathcal{L}_{adv} + \mathcal{L}_{con} + \mathcal{L}_{CL}. \quad (15)$$

Datasets	Metrics	Model-based Methods		Supervised Learning Methods			Unsupervised Learning Methods					
		LECARM	SDD	RetinexNet	KinD	URetinexNet	ZeroDCE	SSIENet	RUAS	EnGAN	SCI	Ours
LOL [46]	PSNR \uparrow	14.41	13.34	16.77	17.65	19.54	14.80	19.50	16.40	17.48	14.78	19.84
	SSIM \uparrow	0.5448	0.6342	0.4249	0.7614	0.7621	0.5607	0.7003	0.5034	0.6515	0.5254	0.7713
	NIQE \downarrow	12.34	13.77	12.51	14.81	11.39	12.62	15.89	11.19	12.53	11.72	11.26
	LOE \downarrow	187.9	263.8	486.2	350.8	158.0	216.6	224.1	125.6	366.2	102.3	117.7
LSRW [13]	PSNR \uparrow	15.34	14.71	15.48	16.41	18.10	15.80	16.14	14.11	17.06	15.24	19.00
	SSIM \uparrow	0.4212	0.4849	0.3468	0.4760	0.5149	0.4450	0.4627	0.4112	0.4601	0.4192	0.5360
	NIQE \downarrow	18.31	11.68	10.31	11.13	10.76	11.83	12.70	11.08	11.94	10.22	9.23
	LOE \downarrow	146.3	218.5	535.6	255.4	202.4	216.0	196.0	198.9	385.1	234.6	189.5
LIME [12]	NIQE \downarrow	12.80	15.21	11.88	14.72	14.48	12.85	16.16	12.44	14.59	12.38	11.01
	LOE \downarrow	261.7	217.5	589.6	249.6	166.7	192.1	216.6	288.7	421.1	212.6	187.2

Table 1: Comparison on three benchmarks. The best and the second best results are highlighted in red and blue respectively.




Figure 6: Subjective comparison on the LSRW dataset among state-of-the-art low-light image enhancement algorithms.

4. Experiments

In this section, we first present the implementation details of our approach. Then we compare it with the state-of-the-art methods through multiple benchmarks. To identify the contribution of each component, we further conduct ablation analyses. All experiments are implemented with PyTorch and conducted on a single NVIDIA Tesla V100 GPU.

4.1. Implementation Details

Parameter Settings. For training, we adopt Adam optimizer [18] with the hyper parameters $\beta_1 = 0.5$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. Our model is trained for 300 epochs with an initial learning rate of 2×10^{-4} and decaying linearly to 0 in the last 200 epochs. Batch size is set to be 1 and patch size is resized to 256×256 for training in the concern of efficiency. Heuristically, we adopt the MLP with 3 hidden layers to normalize the degradation level.

Benchmarks and Metrics. To validate the effectiveness of our method, we train and test the model on LSRW dataset [13], which contains 1000 low-light-normal-light image pairs for training and 50 pairs for evaluation. Each pair consists of a degraded image and a well-exposed reference, which are captured from real world with different exposure times. For a more convincing comparison, we further extend evaluation to other benchmarks such as LOL [46] and LIME [12]. As LOLv1 only contains 15 im-

ages for evaluation, we randomly sample 35 images from the test set of LOLv2 (not used for training) and evaluate on these 50 images. To demonstrate the generalization to real-world degradation scenarios, we test on LIME with the model trained on LSRW. Noting that to prove the superiority of our unsupervised learning manner, during training, we only adopt the low-light part of the paired training data, and replace the references with 300 images from BSD300 dataset [29]. We use two full-reference metrics, *i.e.*, Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) [45], and two no-reference metrics, namely NIQE [31] and LOE [44], to evaluate the effectiveness of different algorithms objectively. In general, a higher PSNR or SSIM means more authentic restored results, while a lower NIQE or LOE represents higher quality details, lightness, and tone.

4.2. Comparison with the State-of-the-Art

For a more comprehensive analysis, we compare our method with two recently-proposed model-based methods, namely LECARM [39] and SDD [14], three advanced supervised learning methods (*i.e.*, RetinexNet [46], KinD [55], and URetinex-Net [47]), and five unsupervised learning methods, including ZeroDCE [11], SSIENet [53], RUAS [25], EnGAN [16], and SCI [28].

Quantitative analysis. We obtain quantitative results of other methods by adopting official pre-trained models and




Figure 7: Visual results of ablation study. The full set performs best, especially in the regions boxed in green and red.

running their respective public codes. As shown in Tab. 1, our method achieves nearly SOTA performance in both full-reference and no-reference metrics across all benchmarks. It validates the superior effect of the proposed framework. Noting that our method even performs better than the supervised ones. Compared with some recent competitive unsupervised approaches such as EnGAN [16], and SCI [28]. We provide stronger constraints than EnGAN, including visual-oriented guidance from prompts. Besides, due to the normalization property, our method outperforms SCI especially on some challenging scenes. The visual results of all methods (including NeRCo and SCI) in a scenario with severe brightness degradation are given below.

Qualitative analysis. For a more intuitive comparison, we report the visual results of all approaches in Fig. 6. Our input is a severely degraded image. One can see that the recent traditional methods can not recover enough brightness. While the advanced deep learning-based methods over-smooth background or introduce unknown veils, resulting in miserable artifacts and unnatural tone. In particular, SCI fails to effectively enhance such a challenging scenario. By comparison, our model realizes the best visual quality with prominent contrast and vivid colors. Details are also remained well. More results are presented in SM.

4.3. Ablation Study

As shown in Tab. 2, we consider four ablation settings by adding the proposed components to the dual loop successively. All ablation studies are conducted on the LSRW dataset. **i)** “#1” is a naive dual loop without any other operations, which only adopts a vanilla color discriminator. This framework has achieved comparable scores, indicating its effectiveness. **ii)** “#2” adds the mask extractor (ME) and the related cooperative loss (CL) to “#1”. ME provides attention guidance and CL reduces solution space. Note that

index	NRN	TAD	CL&ME	PSNR \uparrow	SSIM \uparrow	NIQE \downarrow	LOE \downarrow
#1	\times	\times	\times	16.77	0.4565	12.34	272.4
#2	\times	\times	\checkmark	17.65	0.5023	10.60	247.9
#3	\times	\checkmark	\checkmark	18.32	0.5201	10.83	230.9
#4	\checkmark	\times	\checkmark	18.62	0.5239	9.63	218.8
NeRCo	\checkmark	\checkmark	\checkmark	19.00	0.5360	9.23	189.5

Table 2: Quantitative evaluation on the enhanced results obtained from different settings. The best and the second best results are highlighted in red and blue respectively.

“#2” further adopts edge-path discriminator but without text supervision. **iii)** “#3” adds TAD to “#2”, which employs visual-oriented guidance from textual modality. Further ablation study on TAD will be given in SM. **iv)** “#4” adds NRN to “#2” but without TAD, which is designed to compare the performance gain of TAD and NRN. **v)** Finally, we adopt a complete NeRCo. Due to the improved robustness to different degradation conditions and visual-friendly guidance, this setting achieves apparent performance gain.

We report qualitative results of all settings in Fig. 7. One can see that in the first sample, “#1” increases contrast and roughly recovers objects in dark region, but the enhanced tone is still inauthentic. “#2” relieves color cast phenomenon to a certain degree but still remains undesired veils. Both “#3” and “#4” generate cleaner results with realistic lightness. Further, the result of NeRCo faithfully preserves the most details and performs the best perceptual effectiveness, especially in the regions boxed in green and red. We attribute this to the neural representation normalization and text-driven appearance discriminator, the former unifies degradation level and reduces the difficulty of enhancement task, while the latter guides visual-friendly optimization.

5. Conclusion

We proposed an implicit Neural Representation method for Cooperative low-light image enhancement, dubbed NeRCo, to recover visual-friendly results in an unsupervised manner. Firstly, for the input degraded image, we employed neural representation to normalize degradation levels (*e.g.*, dark lightness and natural noise). Besides, for the output enhanced image, we equipped the discriminator with a high-frequency path and utilized priors from the pre-trained vision-language model to impart perceptual-oriented guidance. Finally, to ease the reliance on paired data and enhance low-light scenarios in a self-supervised manner, a dual-closed-loop cooperative constraint was developed to train the enhancement module. It encourages all components to constrain each other, further reducing solution space. Experiments proved the superiority of our method compared with other top-performing ones. The proposed components provide valuable inspiration for other low-level tasks, such as image dehazing [50], compressive sensing [23], and hyperspectral imaging [51, 52].

References

- [1] Yunpeng Bai, Chao Dong, and Cairong Wang. Ps-nerv: Patch-wise stylized neural representations for videos. *arXiv preprint arXiv:2208.03742*, 2022.
- [2] Bolun Cai, Xianming Xu, Kailing Guo, Kui Jia, Bin Hu, and Dacheng Tao. A joint intrinsic-extrinsic prior model for retinex. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
- [3] Jianrui Cai, Shuhang Gu, and Lei Zhang. Learning a deep single image contrast enhancer from multi-exposure images. *IEEE Transactions on Image Processing*, 2018.
- [4] Hao Chen, Bo He, Hanyu Wang, Yixuan Ren, Ser Nam Lim, and Abhinav Shrivastava. Nerv: Neural representations for videos. *Advances in Neural Information Processing Systems*, 2021.
- [5] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [6] Xuxin Cheng, Qianqian Dong, Fengpeng Yue, Tom Ko, Mingxuan Wang, and Yuexian Zou. M3st: Mix at three levels for speech translation. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [7] Xuxin Cheng, Zhihong Zhu, Hongxiang Li, Yaowei Li, and Yuexian Zou. Ssvmr: Saliency-based self-training for video-music retrieval. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [8] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *Proceedings of International Conference on Machine Learning*, 2021.
- [9] Emilien Dupont, Yee Whye Teh, and Arnaud Doucet. Generative models as distributions of functions. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, 2022.
- [10] Xueyang Fu, Delu Zeng, Yue Huang, Xiao-Ping Zhang, and Xinghao Ding. A weighted variational model for simultaneous reflectance and illumination estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [11] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [12] Xiaojie Guo, Yu Li, and Haibin Ling. Lime: Low-light image enhancement via illumination map estimation. *IEEE Transactions on Image Processing*, 2017.
- [13] Jiang Hai, Zhu Xuan, Ren Yang, Yutong Hao, Fengzhu Zou, Fang Lin, and Songchen Han. R2rnet: Low-light image enhancement via real-low to real-normal network. *Journal of Visual Communication and Image Representation*, 2023.
- [14] Shijie Hao, Xu Han, Yanrong Guo, Xin Xu, and Meng Wang. Low-light image enhancement with semi-decoupled decomposition. *IEEE Transactions on Multimedia*, 2020.
- [15] Md Jahidul Islam, Chelsey Edge, Yuyang Xiao, Peigen Luo, Muntaqim Mehtaz, Christopher Morse, Sadman Sakib Enan, and Junaed Sattar. Semantic segmentation of underwater imagery: Dataset and benchmark. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.
- [16] Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou, and Zhangyang Wang. Enlightengan: Deep light enhancement without paired supervision. *IEEE Transactions on Image Processing*, 2021.
- [17] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *European Conference on Computer Vision*, 2022.
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [19] Jaewon Lee, Kwang Pyo Choi, and Kyong Hwan Jin. Learning local implicit fourier representation for image warping. In *European Conference on Computer Vision*, 2022.
- [20] Jaewon Lee and Kyong Hwan Jin. Local texture estimator for implicit representation function. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [21] Kanggeun Lee and Won-Ki Jeong. Iscl: Interdependent self-cooperative learning for unpaired image denoising. *IEEE Transactions on Medical Imaging*, 2021.
- [22] Hongxiang Li, Meng Cao, Xuxin Cheng, Yaowei Li, Zhihong Zhu, and Yuexian Zou. G2l: Semantically aligned and uniform video grounding via geodesic and game theory. In *Proceedings of the IEEE International Conference on Computer Vision*, 2023.
- [23] Weiqi Li, Bin Chen, and Jian Zhang. D3c2-net: Dual-domain deep convolutional coding network for compressive sensing. *arXiv preprint arXiv:2207.13560*, 2022.
- [24] Yaowei Li, Bang Yang, Xuxin Cheng, Zhihong Zhu, Hongxiang Li, and Yuexian Zou. Unify, align and refine: Multi-level semantic alignment for radiology report generation. *arXiv preprint arXiv:2303.15932*, 2023.
- [25] Risheng Liu, Long Ma, Jiaao Zhang, Xin Fan, and Zhongxuan Luo. Retinex-inspired unrolling with cooperative prior architecture search for low-light image enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [26] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European Conference on Computer Vision*, 2016.
- [27] Yun Liu, Zhongsheng Yan, Sixiang Chen, Tian Ye, Wenqi Ren, and Erkang Chen. Nighthazeformer: Single nighttime haze removal using prior query transformer. *arXiv preprint arXiv:2305.09533*, 2023.
- [28] Long Ma, Tengyu Ma, Risheng Liu, Xin Fan, and Zhongxuan Luo. Toward fast, flexible, and robust low-light image enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [29] 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 of the IEEE International Conference on Computer Vision*, 2001.
- [30] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, 2020.
- [31] Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters*, 2013.
- [32] Nathan Moroney. Local color correction using non-linear masking. In *Color and Imaging Conference*, 2000.
- [33] Michael K. Ng and Wei Wang. A total variation model for retinex. *SIAM Journal on Imaging Sciences*, 2011.
- [34] Etta D Pisano, Shuquan Zong, Bradley M Hemminger, Marla DeLuca, R Eugene Johnston, Keith Muller, M Patricia Braeuning, and Stephen M Pizer. Contrast limited adaptive histogram equalization image processing to improve the detection of simulated spiculations in dense mammograms. *Journal of Digital imaging*, 1998.
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of International Conference on Machine Learning*, 2021.
- [36] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *Proceedings of International Conference on Machine Learning*, 2019.
- [37] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [38] Albert W. Reed, Hyojin Kim, Rushil Anirudh, K. Aditya Mohan, Kyle Champley, Jingu Kang, and Suren Jayasuriya. Dynamic ct reconstruction from limited views with implicit neural representations and parametric motion fields. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021.
- [39] Yurui Ren, Zhenqiang Ying, Thomas H. Li, and Ge Li. Lecarm: Low-light image enhancement using the camera response model. *IEEE Transactions on Circuits and Systems for Video Technology*, 2019.
- [40] Vishwanath Saragadam, Jasper Tan, Guha Balakrishnan, Richard G Baraniuk, and Ashok Veeraraghavan. Miner: Multiscale implicit neural representation. In *European Conference on Computer Vision*, 2022.
- [41] Yu Sun, Jiaming Liu, Mingyang Xie, Brendt Egon Wohlberg, and Ulugbek S Kamilov. Coil: Coordinate-based internal learning for imaging inverse problems. *IEEE Transactions on Computational Imaging*, 2021.
- [42] Matthew Tancik, Ben Mildenhall, Terrance Wang, Divi Schmidt, Pratul P. Srinivasan, Jonathan T. Barron, and Ren Ng. Learned initializations for optimizing coordinate-based neural representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [43] Zia ur Rahman, Daniel J. Jobson, and Glenn A. Woodell. Retinex processing for automatic image enhancement. *Journal of Electronic Imaging*, 2004.
- [44] Shuhang Wang, Jin Zheng, Hai-Miao Hu, and Bo Li. Naturalness preserved enhancement algorithm for non-uniform illumination images. *IEEE Transactions on Image Processing*, 2013.
- [45] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 2004.
- [46] Chen Wei, Wenjing Wang, Wenhao Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. *arXiv preprint arXiv:1808.04560*, 2018.
- [47] Wenhui Wu, Jian Weng, Pingping Zhang, Xu Wang, Wenhao Yang, and Jianmin Jiang. Uretinex-net: Retinex-based deep unfolding network for low-light image enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [48] Li Xu, Qiong Yan, Yang Xia, and Jiaya Jia. Structure extraction from texture via relative total variation. *ACM Trans. Graph.*, 2012.
- [49] Yuhang Yang, Wei Zhai, Hongchen Luo, Yang Cao, Jiebo Luo, and Zheng-Jun Zha. Grounding 3d object affordance from 2d interactions in images. In *Proceedings of the IEEE International Conference on Computer Vision*, 2023.
- [50] Tian Ye, Yunchen Zhang, Mingchao Jiang, Liang Chen, Yun Liu, Sixiang Chen, and Erkang Chen. Perceiving and modeling density for image dehazing. In *European Conference on Computer Vision*, 2022.
- [51] Xuanyu Zhang, Bin Chen, Wenzhen Zou, Shuai Liu, Yongbing Zhang, Ruiqin Xiong, and Jian Zhang. Progressive content-aware coded hyperspectral compressive imaging. *arXiv preprint arXiv:2303.09773*, 2023.
- [52] Xuanyu Zhang, Yongbing Zhang, Ruiqin Xiong, Qilin Sun, and Jian Zhang. Herosnet: Hyperspectral explicable reconstruction and optimal sampling deep network for snapshot compressive imaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [53] Yu Zhang, Xiaoguang Di, Bin Zhang, and Chunhui Wang. Self-supervised image enhancement network: Training with low light images only. *arXiv preprint arXiv:2002.11300*, 2020.
- [54] Yunfan Zhang, Ties van Rozendaal, Johann Brehmer, Markus Nagel, and Taco Cohen. Implicit neural video compression. In *ICLR Workshop on Deep Generative Models for Highly Structured Data*, 2022.
- [55] Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. Kindling the darkness: A practical low-light image enhancer. In *Proceedings of the 27th ACM International Conference on Multimedia*, 2019.
- [56] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 2022.