T^{3}LIE: RETINEX-INNER TEST-TIME TRAINING FOR LOW-LIGHT IMAGE ENHANCEMENT

Anonymous authors
Paper under double-blind review

ABSTRACT

Low-light image enhancement remains challenging due to diverse illumination conditions and complex degradation patterns. Existing methods typically rely on offline-trained models with fixed parameters, which limits their ability to adapt to image-specific illumination during inference. To address this issue, we propose T^{3}LIE, a Retinex-inner Test-Time Training framework that performs illumination-aware adaptation by updating fast weights with self-supervised guidance. Specifically, T^{3}LIE introduces a Retinex-based additive enhancement formulation to decouple reflectance and illumination representations, whose features are progressively injected into the restoration process to guide test-time optimization. By exploiting image-specific illumination priors, T^{3}LIE enables adaptive enhancement while preserving scene details. Extensive experiments on multiple low-light benchmarks demonstrate that T^{3}LIE achieves state-of-the-art enhancement performance and improved generalization across diverse illumination.

1 INTRODUCTION

Low-light image enhancement (LLIE) aims to restore visibility and improve the perceptual quality of images captured under insufficient illumination. Its benefits extend to downstream vision tasks, including 3D reconstruction Cui et al. (2025); Liu et al. (2026), depth estimation Wang et al. (2021); Zheng et al. (2023), and object detection Hashmi et al. (2023); Sun et al. (2026), where illumination degradation obscures structural cues and compromises visual perception.

Recent deep learning-based LLIE methods Shi et al. (2026); Brateanu et al. (2026) have achieved substantial progress through convolutional neural networks Lore et al. (2017), Transformers Cai et al. (2023); Yan et al. (2025), and state-space models Weng et al. (2024); Bai et al. (2024). Retinex theory Chen Wei (2018) and other illumination-related priors Chen Wei (2018); Wang et al. (2024) have further informed these architectures, enabling illumination-conditioned feature processing through attention mechanisms and diffusion-based restoration. Nevertheless, most approaches Xiaogang Xu & Jia (2022); Wang et al. (2023b); Kong et al. (2026a) retain fixed parameters after offline training. Although their predictions depend on the input, their enhancement mappings are not explicitly optimized for the illumination of individual test images.

Test-Time Training (TTT) Sun et al. (2020) offers a mechanism for such adaptation by updating model parameters through self-supervision derived from the test input. Recent advances in visual TTT Han et al. (2026); Kong et al. (2026b) have demonstrated the potential of efficient inner-loop optimization for learning input-specific feature mappings. For LLIE, the effectiveness of adaptation depends on whether the inner objective captures the relationship between illumination degradation and the required enhancement. A generic feature reconstruction objective does not explicitly encode this relationship. Beyond advances in TTT mechanisms Zheng et al. (2026); Yang et al. (2026), a central question is therefore how to incorporate physical priors into the optimization objective itself, so that test-time updates are aligned with the enhancement task.

To this end, we introduce T^{3}LIE, a Retinex-inner Test-Time Training framework for low-light image enhancement. We first reformulate Retinex-based enhancement as an additive correction comprising illumination compensation and illumination-dependent reflectance refinement, in which illumination features condition the self-supervised reconstruction target used to update reflectance fast weights. Illumination guidance thus enters the inner optimization directly, enabling image-specific adapta-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

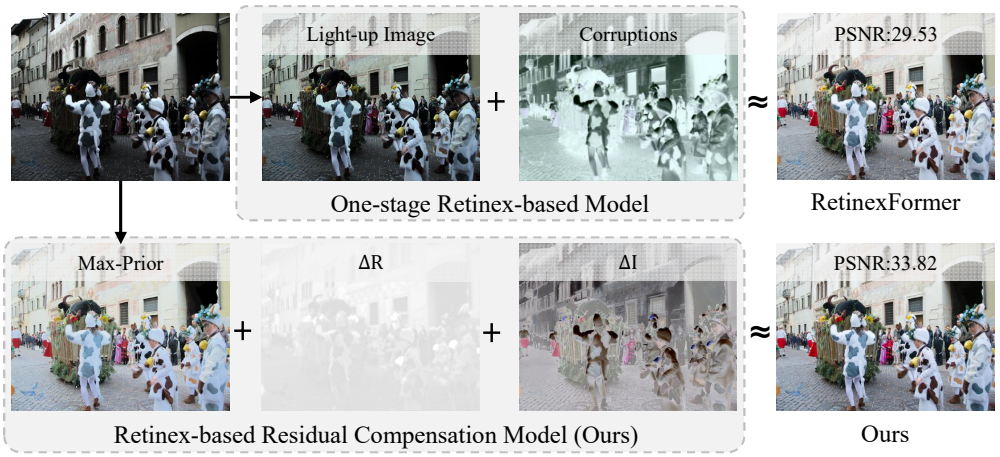


Figure 1: Comparison of previous Retinex-based decomposition methods (RetinexFormer Cai et al. (2023), RetinexMamba Bai et al. (2024)) and ours Retinex-based residual compensation model.

tion while the offline-trained backbone parameters remain fixed. The adapted representations are subsequently combined to reconstruct the enhancement residual. Our contributions are summarized as follows:

- We propose T^{3}LIE, a Retinex-inner TTT framework built on a residual compensation model that decomposes low-light enhancement into illumination-aware compensation and dynamic illumination-dependent reflectance refinement.
- We design an illumination-aware test-time training mechanism that couples illumination adaptation with reflectance refinement through a Retinex-guided inner self-supervised objective. The adapted illumination features condition the reflectance reconstruction targets, enabling image-specific fast-weight updates that account for spatially varying illumination.
- Extensive experiments on benchmark datasets and real-world low-light images demonstrate state-of-the-art enhancement performance and improved generalization across diverse illumination conditions.

2 Related Works

Low Light Enhancement. Deep learning-based methods have significantly advanced low-light image enhancement by leveraging powerful network architectures and illumination-related priors. Early approaches Lore et al. (2017); Guo et al. (2020) explored convolutional architectures and zero-reference learning strategies, while subsequent methods incorporated Retinex theory Ma et al. (2022), handcrafted enhancement priors Xiaogang Xu & Jia (2022), semantic information Zeng et al. (2025), and statistical modeling Kong et al. (2026a) to improve illumination estimation and detail restoration. Recent approaches further employ Transformers Cai et al. (2023); Zhang et al. (2025b), state-space models Weng et al. (2024); Zhang et al. (2025a), diffusion models Yi et al. (2023); Lan et al. (2025), and RWKV-based architectures Xu et al. (2025) for enhanced feature representation and restoration. However, most existing LLIE methods rely on fixed parameters during inference, limiting their ability to adapt to diverse illumination conditions.

Test-Time Training. Test-Time Training (TTT) Sun et al. (2020) is an emerging paradigm that enables models to adapt to individual test samples by optimizing model parameters with self-supervised objectives during inference. Recent studies have explored efficient fast-weight adaptation mechanisms and demonstrated the effectiveness of TTT in various vision tasks, including video generation Dalal et al. (2025), 3D reconstruction Chen et al. (2026), and image restoration Li et al. (2025); Zheng et al. (2026). However, existing TTT approaches mainly focus on improving adaptation efficiency, while lacking task-specific guidance for low-level vision problems. For LLIE, illumination variation is the key factor affecting enhancement behavior, yet how to incorporate illumination priors into TTT adaptation remains unexplored.

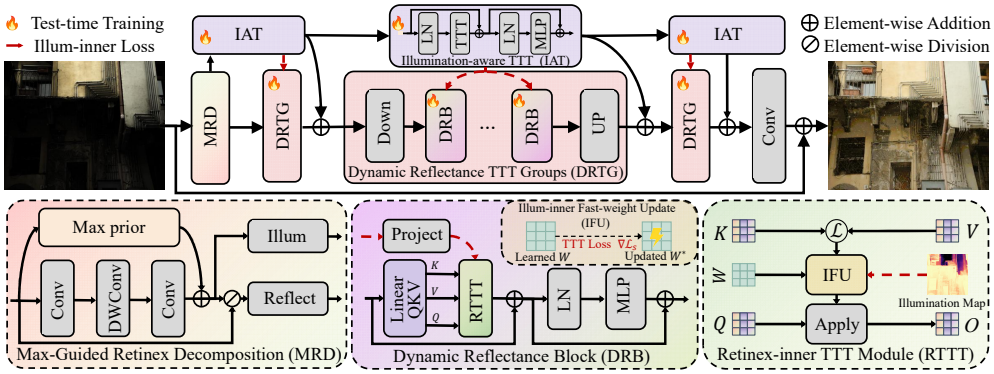


Figure 2: **The pipeline of our T^{3}LIE.** The input image is decomposed into illumination and reflectance via Max-guided Retinex decomposition, then refined by dynamic reflectance TTT groups (DRTG) and illumination-aware TTT (IAT) for Retinex-inner test-time adaptation.

3 METHODOLOGY

In this section, we first derive a Retinex-based residual compensation model that separates illumination compensation from illumination-dependent reflectance refinement. We then present our Retinex-inner test-time training (TTT) framework, comprising max-guided Retinex decomposition, dynamic illumination TTT, and illumination-aware reflectance TTT.

3.1 RESIDUAL RETINEX COMPENSATION MODEL (RRCM)

Retinex Model. Consider a low-light image $I \in \mathbb{R}^{H\times W\times 3}$. Following Retinex theory Chen Wei (2018), the observed image is modeled as the element-wise product of a reflectance component and an illumination component:

$$I = R \odot L, \tag{1}$$

where $R \in \mathbb{R}^{H\times W\times 3}$ denotes the reflectance component containing intrinsic scene properties, $L \in \mathbb{R}^{H\times W\times 1}$ represents the spatially varying illumination map, and $\odot$ denotes element-wise multiplication. Here, H and W represent the image height and width, respectively. In low-light enhancement, the reflectance component is commonly regarded as the illumination-independent representation to be recovered. However, the decomposition of a single image is inherently ill-posed, since multiple combinations of reflectance and illumination can produce the same observation.

RRCM. Based on the Retinex formulation, we derive a Residual Retinex Compensation Model, which interprets low-light enhancement as reflectance recovery through illumination-aware residual compensation. By rearranging the Retinex model in equation 1, the reflectance can be expressed as

$$R = I + \Delta I, \qquad \Delta I = R \odot (1 - L), \tag{2}$$

where $\Delta I \in \mathbb{R}^{H\times W\times 3}$ denotes the illumination-dependent correction required to compensate for attenuation under insufficient illumination.

Since the intrinsic reflectance R and illumination L cannot be directly observed from a single low-light image, we introduce their initial estimates $R_0 \in \mathbb{R}^{H\times W\times 3}$ and $L_0 \in \mathbb{R}^{H\times W\times 1}$, with the corresponding estimation residuals defined as $\delta R = R - R_0$ and $\delta L = L - L_0$. Substituting these estimates into equation 2 gives $\Delta I = (R_0 + \delta R) \odot [1 - (L_0 + \delta L)]$.

Accordingly, the enhancement residual can be decomposed into two complementary components:

$$\begin{aligned} \Delta L &= R_0 \odot [1 - (L_0 + \delta L)], \\ \Delta R &= \delta R \odot [1 - (L_0 + \delta L)]. \end{aligned} \tag{3}$$

Here, ΔL and ΔR denote illumination compensation and reflectance refinement, respectively. Both are formulated in the three-channel image domain, and together recover the complete enhancement residual:

$$R = I + \Delta L + \Delta R. \tag{4}$$

This decomposition reveals that illumination variation governs both compensation terms. While ΔL directly restores the attenuation caused by insufficient illumination, ΔR remains explicitly coupled with illumination estimation through $\Delta R = \delta R \odot (1 - L_0) - \delta R \odot \delta L$. Therefore, illumination errors affect not only intensity compensation but also the refinement of reflectance details. This coupled dependency motivates our illumination-aware adaptation strategy, where illumination information further guides reflectance refinement during test-time training.

Discussion. As shown in Figure 1, Conventional Retinex-based methods Ma et al. (2022); Cai et al. (2023); Bai et al. (2024) typically prioritize illumination enhancement and subsequently restore the accompanying degradations. Multinex Brateanu et al. (2026) adopts a residual form, but directly factorizes the enhancement delta into learned luminance and reflectance adjustment fields. In contrast, our RRCM derives the residual from the Retinex relation itself and decomposes it according to latent estimation errors. More importantly, the resulting reflectance refinement is explicitly conditioned on illumination, which further motivates our illumination-aware test-time adaptation.

3.2 Retinex-inner Test-Time Training Framework

Following the RRCM in equation 4, we implement illumination compensation and reflectance refinement through two coupled pathways, as illustrated in Fig. 2. Given a low-light image I, the Max-Guided Retinex Decomposition (MRD) first produces an initial illumination representation and a corresponding reflectance representation. The illumination pathway is progressively adapted by Illumination-aware TTT (IAT), while the reflectance pathway is refined by Dynamic Reflectance TTT Groups (DRTGs). At each stage, the adapted illumination features are injected into the reflectance stream and further provide spatial guidance for the fast-weight update inside each Dynamic Reflectance Block (DRB). Finally, the fused representation is projected to an enhancement residual and added to the input image to obtain the enhanced result.

Max-Guided Retinex Decomposition (MRD). Given a low-light image $I \in \mathbb{R}^{H\times W\times 3}$, MRD initializes the Retinex decomposition by combining a maximum-response illumination estimate with learnable residual refinement. Following the observation that the strongest RGB response provides an effective approximation of local illumination Guo et al. (2016), we first compute

$$P_{\max}(I) = \max_{c\in\{r,g,b\}} I_c. \tag{5}$$

We use $P_{\max}(I)$ as the initial illumination estimate and further refine it with a lightweight Conv–DWConv–Conv branch $\mathcal{E}_\theta$ that predicts an image-dependent residual correction:

$$L_0 = P_{\max}(I) + \mathcal{E}_\theta(I), \tag{6}$$

where the learned residual compensates for local deviations that cannot be captured by the maximum-response estimate alone. The corresponding reflectance is then recovered according to the Retinex relation:

$$R_0 = I \oslash (L_0 + \epsilon), \tag{7}$$

where $\oslash$ denotes element-wise division and $\epsilon = 10^{-6}$ stabilizes the division. The resulting illumination and reflectance representations are further encoded as F_L^0 and F_R^0, respectively, to initialize the subsequent adaptation pathways.

Illumination-aware TTT (IAT). Given an illumination feature $F_L \in \mathbb{R}^{h\times w\times C}$, IAT performs image-specific adaptation through a TTT sublayer followed by an MLP refinement, both with residual connections. Specifically, F_L is first normalized and linearly projected into key, value, and query features K_L, V_L, and Q_L. The key–value pair constructs a self-supervised reconstruction task, where the fast weights W_L are adapted to the current illumination condition and subsequently applied to the query:

$$W_L^* = W_L - \eta_L \nabla_{W_L} \ell_L(f_{W_L}(K_L), V_L), \qquad O_L = f_{W_L^*}(Q_L), \tag{8}$$

where $\ell_L(f_{W_L}(K_L), V_L) = -(s/N)\langle f_{W_L}(K_L), V_L\rangle_F$ is the unmodulated dot-product loss with scale s, and η_L is the inner learning rate. Unlike IFU, this illumination step follows the raw gradient. In this way, the TTT sublayer converts the current illumination statistics into an input-specific feature response rather than relying solely on the offline-learned mapping. The adapted response is fused with the input feature through a residual connection, followed by a pre-normalized MLP for further feature refinement. The complete IAT transformation is written as

$$\widetilde{F}_L = F_L + O_L, \qquad F_L^{\text{out}} = \widetilde{F}_L + \mathrm{MLP}\Big(\mathrm{LN}(\widetilde{F}_L)\Big). \tag{9}$$

The resulting F_L^{out} is propagated to the next illumination stage and simultaneously projected as illumination guidance $G_L = \mathcal{P}(F_L^{\text{out}})$ for the corresponding reflectance TTT modules. Therefore, IAT not only progressively adapts the illumination representation to each test image, but also provides the illumination state required to condition subsequent reflectance refinement.

Dynamic Reflectance TTT Group (DRTG). The reflectance pathway is organized in a U-Net-style hierarchy composed of Dynamic Reflectance TTT Groups (DRTGs). Each DRTG contains multiple Dynamic Reflectance Blocks (DRBs) for progressively refining reflectance features at the corresponding scale. Given the incoming reflectance feature F_R and the stage-aligned illumination feature F_L, the group can be written compactly $F_R^j = \mathrm{DRB}_j(F_R^{j-1}, F_L), j = 1, \ldots, M$, where M denotes the number of DRBs in the group, while the proposed DRB performs illumination-aware test-time adaptation within each stage.

Dynamic Reflectance Block (DRB). Each DRB dynamically refines the reflectance representation conditioned on the corresponding illumination state. Given F_R and F_L, the reflectance feature is first projected into key, value, and query representations K_R, V_R, and Q_R, while the illumination feature is projected as $G = \mathcal{P}(F_L)$ to match their spatial and channel dimensions. The resulting (K_R, V_R, Q_R, G) are fed into the Retinex-inner TTT Module (RTTT):

$$O_R = \mathrm{RTTT}(K_R, V_R, Q_R; G). \tag{10}$$

The adapted response is then incorporated through a residual connection followed by MLP:

$$\widetilde{F}_R = F_R + O_R, \qquad F_R^{\text{out}} = \widetilde{F}_R + \mathrm{MLP}\big(\mathrm{LN}(\widetilde{F}_R)\big). \tag{11}$$

In this way, each DRB preserves the incoming reflectance representation while introducing illumination-conditioned, image-specific adaptation.

Retinex-inner TTT Module (RTTT). RTTT realizes the illumination-dependent reflectance refinement derived in RRCM through fast-weight adaptation. For each head, let $K_R, V_R, Q_R, G \in \mathbb{R}^{N\times d}$, where N denotes the number of spatial locations. We employ a lightweight SwiGLU inner learner $f_W(X) = (XW_1) \odot \mathrm{SiLU}(XW_2)$, with $W = \{W_1, W_2\}$ denoting the fast weights.

The key features are first processed as $\widehat{V}_R = f_W(K_R)$. Instead of treating all reflectance responses equally, we modulate the reconstruction target using illumination guidance and define the Retinex-inner objective as

$$\mathcal{L}_{\text{inner}}^R = -\frac{s}{N}\left\langle \widehat{V}_R,\, G \odot V_R \right\rangle_F, \tag{12}$$

where $s = 1/3$ is the scaling factor. Its gradient with respect to the predicted feature is $\partial\mathcal{L}_{\text{inner}}^R/\partial\widehat{V}_R = -(s/N)(G\odot V_R)$, showing that G directly controls the feature-space signal used to update the reflectance learner. Consequently, regions and channels with different illumination states contribute differently to the inner optimization.

Illum-inner Fast-weight Update (IFU). Based on the illumination-modulated gradient, IFU updates the reflectance fast weights for the current input. To stabilize the inner optimization, the gradient of each fast-weight group is normalized before the update:

$$\bar{g}_m = \frac{g_m}{1 + \|g_m\|_2}, \qquad W_m^* = W_m - \eta_R \bar{g}_m, \tag{13}$$

where $g_m = \nabla_{W_m}\mathcal{L}_{\text{inner}}^R$ and η_R denotes the inner learning rate. The adapted weights are then applied to the query features:

$$O_R = f_{W^*}(Q_R), \tag{14}$$

where f_W denotes the lightweight inner learner composed of a simplified SwiGLU branch and a spatial depth-wise convolution branch. Thus, IFU adapts the reflectance mapping to the illumination characteristics of each input while retaining the offline-learned weights as initialization.

At the final stage, the two pathways are jointly aggregated to realize the complementary compensation terms characterized by RRCM. Accordingly, the enhanced result follows

$$\widehat{I} = I + \Delta\widehat{L} + \Delta\widehat{R}, \tag{15}$$

where $\Delta\widehat{L}$ and $\Delta\widehat{R}$ denote the illumination and reflectance compensations implicitly modeled by the two coupled pathways, respectively. Thus, the network architecture is directly instantiated from

Table 1: Quantitative comparisons on LOL-V2-real, LOL-V2-syn, SDSD-indoor, and SDSD-outdoor datasets. The best and second-best results are highlighted in **bold** and underlined.

Methods	Reference	LOL-V2-real		LOL-V2-syn		SDSD-indoor		SDSD-outdoor		Complexity	
		PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	FLOPs (G)	Params (M)
RetinexNet	BMVC'18	15.47	0.567	17.13	0.798	20.84	0.617	20.96	0.629	587.47	0.84
KinD	MM'19	14.74	0.641	13.29	0.578	21.95	0.672	21.97	0.654	34.99	8.02
EnGAN	TIP'21	18.93	0.617	16.57	0.734	23.29	0.703	24.90	0.693	61.01	14.35
Restormer	CVPR'22	19.94	0.827	21.41	0.830	25.67	0.827	24.79	0.802	144.25	26.13
SNR-Net	CVPR'22	21.48	0.849	24.14	0.928	29.44	0.894	28.66	0.866	26.35	4.01
RetinexFormer	ICCV'23	22.80	0.840	25.67	0.930	29.77	0.896	29.84	**0.877**	15.57	1.61
MambaLLIE	NeurIPS'24	22.95	0.847	25.87	0.940	30.12	0.900	30.00	0.869	20.85	2.28
CIDNet	CVPR'25	24.11	0.871	25.70	0.942	29.81	0.899	28.56	0.855	1.88	7.57
CWNet	ICCV'25	21.65	0.860	25.50	0.936	29.10	0.900	28.45	0.856	11.30	1.23
MultiNex	CVPR'26	23.04	0.860	25.04	0.930	29.77	0.897	28.58	0.864	2.50	0.05
TPCNet	ECCV'26	24.98	0.882	26.03	**0.943**	29.82	0.889	29.10	0.859	8.75	2.62
TTTIR	Arxiv'26	23.79	0.869	25.60	0.940	30.12	0.894	29.90	0.866	18.71	0.79
Ours	-	**25.02**	**0.885**	**26.05**	**0.943**	**30.28**	**0.903**	**30.12**	0.870	7.30	2.88

Table 2: Quantitative comparison on unpaired real-world low-light datasets.

	Zero-DCE	SCI	RetinexFormer	MambaLLIE	CIDNet	CWNet	TPCNet	TTTIR	Ours
MUSIQ↑	59.37	54.17	57.48	55.60	57.38	55.72	58.42	58.46	**59.42**
TOPIQ-NR↑	0.498	0.454	0.480	0.474	0.494	0.446	0.500	0.497	**0.508**
CLIP-IQA↑	0.432	0.406	0.442	0.473	0.429	0.315	0.466	0.474	**0.534**
LIQE↑	1.993	2.685	2.884	2.865	2.927	2.465	3.014	3.007	**3.179**

the RRCM decomposition, with the illumination and reflectance pathways respectively realizing the two compensation terms in equation 4.

During inference, the offline-learned parameters are frozen, while the fast weights are independently adapted to each test image using the corresponding inner objectives. Since these objectives are constructed solely from input-derived features, no well-exposed reference is required at test time. Further details of the loss functions are provided in Appendix A.2.

4 Experiments and Results

4.1 Benchmark Datasets and Implementation Details

Datasets. We perform comprehensive evaluations on four paired low-light datasets, which can be categorized into two groups: LOL-v2 and SDSD. The LOL-v2 dataset Yang et al. (2021) consists of two subsets, including LOL-v2-real and LOL-v2-syn. LOL-v2-real contains 689 training pairs and 100 testing pairs, while LOL-v2-syn provides 900 training pairs and 100 testing pairs. The SDSD dataset Ruixing Wang & Jia (2021) also contains two static subsets, i.e., SDSD-indoor and SDSD-outdoor, which include 62 and 116 training pairs, as well as 6 and 10 testing pairs, respectively. Beyond paired evaluation, we further examine the generalization capability of our model on five widely used real-world low-light datasets, including DICM Lee et al. (2013), LIME Guo et al. (2016), MEF Ma et al. (2015), NPE Wang et al. (2013), and VV Vonikakis et al. (2018).

Implementation Details. We implement T^{3}LIE in PyTorch Imambi et al. (2021) and train it on NVIDIA 4090 GPUs. Training images are randomly cropped to 128×128 patches with rotation, flipping, and illumination perturbation. We use Adam with batch size 4, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and an initial learning rate of 2×10^{-4} with cosine annealing for 300 epochs. For test-time training, we use one inner update step with $\eta_L = \eta_R = 1$ and $\epsilon = 10^{-6}$. We report PSNR and SSIM on paired datasets, and MUSIQ Ke et al. (2021), TOPIQ-NR Chen et al. (2024), CLIP-IQA Wang et al. (2023a), and LIQE Zhang et al. (2023) on unpaired real-world datasets.

Quantitative Results. As shown in Table 1, T^{3}LIE achieves strong performance on both synthetic and real paired benchmarks. It obtains the highest PSNR on all four datasets and reaches either the best or second-best SSIM. Compared with RetinexFormer, T^{3}LIE improves the average PSNR by 0.85 dB across the four benchmarks. The gain over TTTIR is 0.52 dB on average, showing that Retinex-guided inner-loop supervision leads to more effective image-specific adaptation. T^{3}LIE also

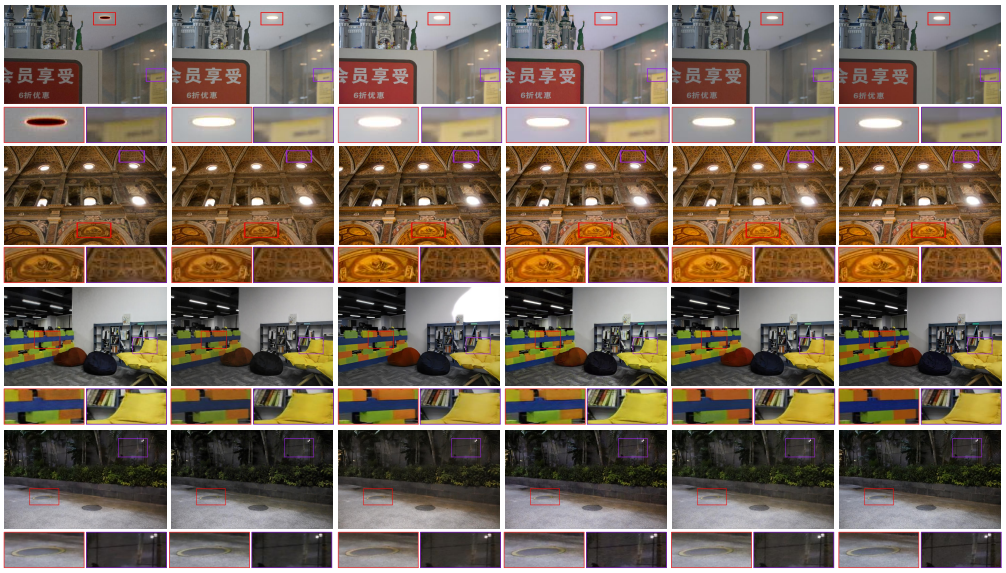


Figure 3: Qualitative comparisons on LOL-V2 and SDSD datasets. T^{3}LIE provides adaptive enhancement under diverse illumination conditions while preserving faithful colors and fine details.

Table 3: Quantitative comparisons on LOM for 3D reconstruction.

Methods	Bike		Buu		Chair		Shrub		Sofa		Average	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑
SCI	13.61	0.688	7.71	0.676	**19.52**	**0.829**	**17.23**	**0.551**	9.86	0.626	13.58	0.670
Zero-DCE	10.49	0.546	17.95	0.870	12.50	0.742	12.41	0.403	14.17	**0.810**	13.50	0.674
CIDNet	16.03	0.733	18.07	0.834	12.28	0.723	10.70	0.219	12.79	0.727	13.97	0.647
MultiNex	10.22	0.518	18.44	0.849	13.17	0.756	10.14	0.121	14.54	0.808	13.30	0.610
TTTIR	15.49	0.736	15.29	0.733	14.20	0.771	10.66	0.244	13.39	0.750	13.80	0.647
Ours	**16.61**	**0.746**	**18.91**	**0.882**	15.02	0.812	14.29	0.476	**15.01**	0.769	**15.97**	**0.737**

reduces FLOPs by 53.1% compared with RetinexFormer and by 61.0% compared with TTTIR. Its computational cost remains close to TPCNet while providing stronger enhancement performance.

We further evaluate generalization on unpaired real-world low-light datasets using MUSIQ, TOPIQ-NR, CLIP-IQA, and LIQE, as reported in Table 2. T^{3}LIE achieves the best results on all four metrics, showing that the proposed adaptation strategy remains effective beyond paired benchmarks and handles diverse real-world illumination conditions well.

Qualitative Results. Figures 3 and 4 present visual comparisons on paired benchmarks and unpaired real-world low-light datasets, respectively. For paired datasets, existing methods often suffer from insufficient illumination correction, color distortion, or detail loss when restoring degraded regions. In contrast, T^{3}LIE produces visually consistent enhanced results with the reference images by effectively recovering illumination while preserving structural details and color fidelity. For unpaired real-world images, where ground-truth references are unavailable, previous methods may introduce over-enhancement artifacts, unnatural brightness, or inaccurate color appearance under diverse illumination conditions. Benefiting from Retinex-inner test-time adaptation, T^{3}LIE dynamically adjusts the enhancement process according to each input image, achieving more natural illumination, better contrast, and improved visual quality in challenging real-world scenarios.

3D Reconstruction. We further evaluate T^{3}LIE on the LOM dataset for downstream 3D reconstruction. Images enhanced by different methods are used as reconstruction inputs, and the synthesized novel views are evaluated using PSNR and SSIM. T^{3}LIE achieves the highest average performance with 15.97 dB PSNR and 0.737 SSIM, improving the average PSNR by 2.00 dB over the second-best method. It performs particularly well on the Bike, Buu, and Sofa scenes, while remaining competitive on Chair and Shrub. The results show that Retinex-inner adaptation can alleviate low-light degradation while preserving scene structures that are important for consistent 3D reconstruction.

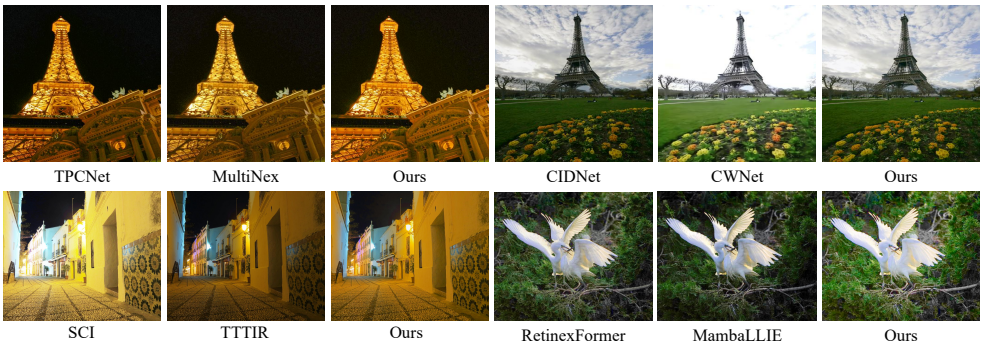


Figure 4: Qualitative comparisons on unpaired real-world low-light datasets. T^{3}LIE produces more natural illumination and improved details under real-world lighting conditions.

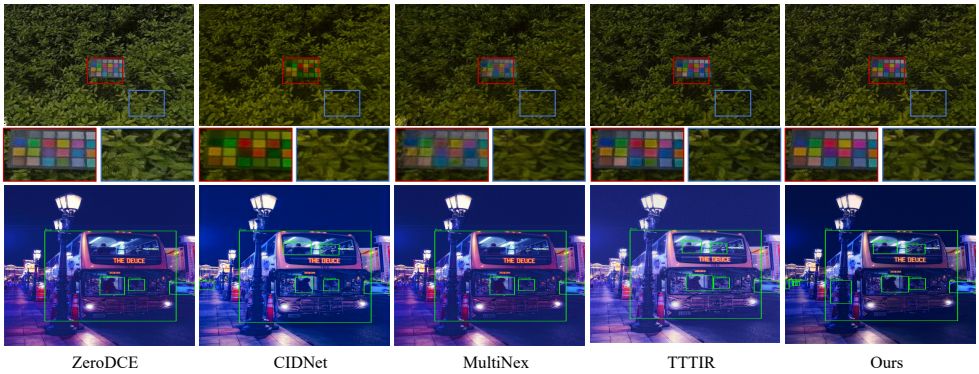


Figure 5: Qualitative results on downstream vision tasks, including 3D reconstruction and low-light object detection. T^{3}LIE preserves task-relevant structures while improving visibility under challenging illumination conditions.

Table 4: Object detection results on the ExDark dataset using different low-light enhancement methods. The best and second-best results are highlighted in **bold** and underlined, respectively.

Methods	Bicycle	Boat	Bottle	Bus	Car	Cat	Chair	Cup	Dog	Motorbike	People	Table	mAP↑
CIDNet	0.635	0.292	**0.324**	0.756	0.394	0.791	0.323	0.256	0.507	0.496	0.257	0.348	0.448
MultiNex	0.663	0.288	0.229	0.825	**0.405**	0.706	0.290	0.238	**0.550**	0.525	**0.396**	**0.369**	0.457
TPCNet	0.653	0.285	0.284	0.814	0.402	0.779	**0.355**	0.331	0.529	0.517	0.362	0.359	0.473
TTTIR	0.626	**0.310**	0.292	0.858	0.378	0.777	**0.355**	0.258	0.522	**0.534**	0.301	0.341	0.463
T^{3}LIE	**0.669**	0.291	0.295	**0.862**	0.396	**0.799**	0.338	**0.343**	0.544	0.526	0.324	0.325	**0.476**

Low-Light Object Detection. To evaluate the effect of low-light enhancement on downstream vision tasks, we conduct object detection experiments on the ExDark dataset . We randomly select 1,200 challenging low-light images from 12 object categories, with 100 images per category. Each image is first enhanced by different methods and then fed into LocateAnything, a vision foundation model for open-world object localization, using the same detector setting for all methods. As shown in Table 4, T^{3}LIE achieves the highest overall mAP of 0.476. It also obtains the best results on Bicycle, Bus, Cat, and Cup, together with strong performance on several other categories. These results show that the proposed Retinex-inner adaptation improves visibility while preserving task-relevant structures for downstream object recognition under challenging low-light conditions.

4.2 Ablation Study

We conduct ablation studies to examine the key design choices of T^{3}LIE from two aspects: the contribution of different components and the design of Retinex-inner test-time training. Detailed configurations of the ablation variants are provided in Appendix A.3.

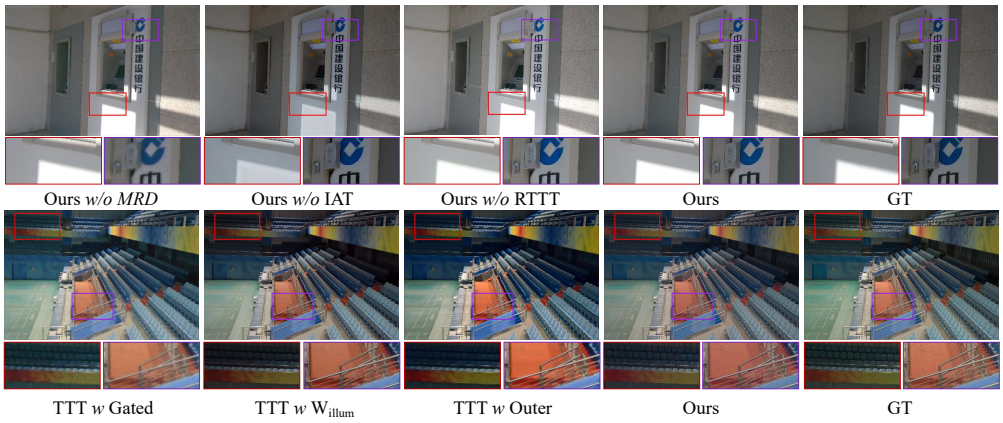


Figure 6: Visual comparison for ablation studies of the proposed T^{3}LIE.

Table 5: Ablation study on different components.

Base	MRD	IAT	DRB	PSNR↑	SSIM↑	Params	FLOPs
✓				23.07	0.860	1.10	2.00
✓		✓		23.40	0.864	2.39	5.16
✓			✓	23.88	0.866	1.59	3.65
✓		✓	✓	24.55	0.875	2.88	7.14
✓	✓	✓	✓	**25.02**	**0.885**	2.89	7.30

Table 6: Ablation study on the RTTT design.

TTT	Gated	W_{illum}	Outer	Retinex-inner Loss	PSNR↑	SSIM↑
✓					23.56	0.867
✓	✓				24.83	0.861
✓		✓			24.44	0.845
✓			✓		24.07	0.859
✓				✓	**25.02**	**0.885**

Effect of Different Components. Table 5 evaluates the contribution of different components in T^{3}LIE. The illumination and reflectance pathways exhibit complementary effects, and their combination achieves 24.55 dB PSNR and 0.875 SSIM. Incorporating the Max-Prior further improves the performance to 25.02 dB PSNR and 0.885 SSIM with only marginal additional complexity, demonstrating the effectiveness of the proposed component design and illumination guidance.

Effect of Test-Time Training Design. Table 6 compares different strategies for incorporating illumination information into test-time adaptation. Gating, illumination-conditioned weight initialization, and outer-level modulation provide partial improvements, whereas the proposed Retinex-inner objective achieves the best performance of 25.02 dB PSNR and 0.885 SSIM. These results demonstrate the benefit of introducing illumination guidance directly into the inner-loop optimization. Detailed configurations are provided in Appendix A.3.

5 Conclusion

We present T^{3}LIE, a Retinex-inner Test-Time Training framework that reformulates low-light image enhancement from fixed offline inference to adaptive inference. By leveraging Retinex priors to guide test-time parameter updates, T^{3}LIE dynamically adjusts enhancement strategies according to the illumination characteristics of individual images. Extensive experiments on diverse low-light benchmarks validate the effectiveness of the proposed framework and demonstrate the potential of task-aware test-time adaptation for low-level vision problems.

6 AI use statement

In this work, generative AI tools were used only for manuscript editing, including readability, grammar, and limited code debugging. They were not used for generating research ideas, developing the methodology, designing experiments, implementing the core method, or producing the main claims and conclusions. All AI-assisted outputs were manually reviewed and verified by the authors.

REFERENCES

Jiesong Bai, Yuhao Yin, Qiyuan He, Yuanxian Li, and Xiaofeng Zhang. Retinexmamba: Retinex-based mamba for low-light image enhancement. In *ICONIP*, pp. 427–442, 2024.

Alexandru Brateanu, Tingting Mu, Codruta O. Ancuti, and Cosmin Ancuti. Multinex: Lightweight low-light image enhancement via multi-prior retinex. In *CVPR*, 2026.

Yuanhao Cai, Hao Bian, Jing Lin, Haoqian Wang, Radu Timofte, and Yulun Zhang. Retinexformer: One-stage retinex-based transformer for low-light image enhancement. In *ICCV*, pp. 12504–12513, 2023.

Chaofeng Chen, Jiadi Mo, Jingwen Hou, Haoning Wu, Liang Liao, Wenxiu Sun, Qiong Yan, and Weisi Lin. Topiq: A top-down approach from semantics to distortions for image quality assessment. *IEEE Trans. Image Process.*, 33:2404–2418, 2024.

Xingyu Chen, Yue Chen, Yuliang Xiu, Andreas Geiger, and Anpei Chen. Ttt3r: 3d reconstruction as test-time training. In *ICLR*, volume 2026, pp. 50694–50718, 2026.

Wenhan Yang Jiaying Liu Chen Wei, Wenjing Wang. Deep retinex decomposition for low-light enhancement. In *BMVC*, 2018.

Ziteng Cui, Xuangeng Chu, and Tatsuya Harada. Luminance-gs: Adapting 3d gaussian splatting to challenging lighting conditions with view-adaptive curve adjustment. In *CVPR*, 2025.

Karan Dalal, Daniel Koceja, Jiarui Xu, Yue Zhao, Shihao Han, Ka Chun Cheung, Jan Kautz, Yejin Choi, Yu Sun, and Xiaolong Wang. One-minute video generation with test-time training. In *CVPR*, pp. 17702–17711. IEEE, 2025.

Chunle Guo 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 *CVPR*, pp. 1780–1789, 2020.

Xiaojie Guo, Yu Li, and Haibin Ling. Lime: Low-light image enhancement via illumination map estimation. *IEEE Trans. Image Process.*, 26(2):982–993, 2016.

Dongchen Han, Yining Li, Tianyu Li, Zixuan Cao, Ziming Wang, Jun Song, Yu Cheng, Bo Zheng, and Gao Huang. Vit3: Unlocking test-time training in vision. In *CVPR*, pp. 51–61, 2026.

Khurram Azeem Hashmi, Goutham Kallempudi, Didier Stricker, and Muhammad Zeshan Afzal. Featenhancer: Enhancing hierarchical features for object detection and beyond under low-light vision. In *ICCV*, pp. 6725–6735, 2023.

Sagar Imambi, Kolla Bhanu Prakash, and GR Kanagachidambaresan. Pytorch. In *Programming with TensorFlow: solution for edge computing applications*, pp. 87–104. 2021.

Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *ICCV*, pp. 5128–5137, 2021.

Derong Kong, Zhixiong Yang, Shengxi Li, Shuaifeng Zhi, Li Liu, Zhen Liu, and Jingyuan Xia. Luminance-aware statistical quantization: Unsupervised hierarchical learning for illumination enhancement. *NeurIPS*, 38:116286–116313, 2026a.

Quan Kong, Yanru Xiao, Yuhao Shen, and Cong Wang. Vision-ttt: Efficient and expressive visual representation learning with test-time training. In *ECCV*, pp. 401–420, 2026b.

Guanzhou Lan, Qianli Ma, Yuqi Yang, Zhigang Wang, Dong Wang, Xuelong Li, and Bin Zhao. Efficient diffusion as low light enhancer. In *CVPR*, pp. 21277–21286, June 2025.

Chulwoo Lee, Chul Lee, and Chang-Su Kim. Contrast enhancement based on layered difference representation of 2d histograms. *IEEE Trans. Image Process.*, 22(12):5372–5384, 2013.

Liangyan Li, Yimo Ning, Kevin Le, Wei Dong, Yunzhe Li, Jun Chen, and Xiaohong Liu. Moir\'exnet: Adaptive multi-scale demoir\'eing with linear attention test-time training and truncated flow matching prior. *arXiv preprint arXiv:2506.15929*, 2025.

Shuhong Liu, Chenyu Bao, Ziteng Cui, Yun Liu, Xuangeng Chu, Lin Gu, Marcos V Conde, Ryo Umagami, Tomohiro Hashimoto, Zijian Hu, et al. Realx3d: A physically-degraded 3d benchmark for multi-view visual restoration and reconstruction. *Int. J. Comput. Vis.*, 134(8):382, 2026.

Kin Gwn Lore, Adedotun Akintayo, and Soumik Sarkar. Llnet: A deep autoencoder approach to natural low-light image enhancement. *Pattern Recogn.*, 61:650–662, 2017.

Kede Ma, Kai Zeng, and Zhou Wang. Perceptual quality assessment for multi-exposure image fusion. *IEEE Trans. Image Process.*, 24(11):3345–3356, 2015.

Long Ma, Tengyu Ma, Risheng Liu, Xin Fan, and Zhongxuan Luo. Toward fast, flexible, and robust low-light image enhancement. In *CVPR*, pp. 5627–5636, 2022.

Chi-Wing Fu Jiangbo Lu Bei Yu Ruixing Wang, Xiaogang Xu and Jiaya Jia. Seeing dynamic scene in the dark: High-quality video dataset with mechatronic alignment. In *ICCV*, 2021.

Jing-Yi Shi, Ming-Fei Li, and Ling-An Wu. Tpcnet: A low-light image enhancement network inspired by triple physical constraints. In *ECCV*, pp. 627–645, 2026.

Fangtong Sun, Congyu Li, Ke Yang, Yuchen Pan, Hanwen Yu, Xichuan Zhang, and Yiying Li. Frbnet: Revisiting low-light vision through frequency-domain radial basis network. *NeurIPS*, 38: 99085–99113, 2026.

Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *ICML*, pp. 9229–9248, 2020.

Vassilios Vonikakis, Rigas Kouskouridas, and Antonios Gasteratos. On the evaluation of illumination compensation algorithms. *Multimed. Tools Appl.*, 77(8):9211–9231, 2018.

Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *AAAI*, 2023a.

Kun Wang, Zhenyu Zhang, Zhiqiang Yan, Xiang Li, Baobei Xu, Jun Li, and Jian Yang. Regularizing nighttime weirdness: Efficient self-supervised monocular depth estimation in the dark. In *ICCV*, pp. 16055–16064, 2021.

Shuhang Wang, Jin Zheng, Hai-Miao Hu, and Bo Li. Naturalness preserved enhancement algorithm for non-uniform illumination images. *IEEE Trans. Image Process.*, 22(9):3538–3548, 2013.

Tao Wang, Kaihao Zhang, Tianrun Shen, Wenhan Luo, Bjorn Stenger, and Tong Lu. Ultra-high-definition low-light image enhancement: A benchmark and transformer-based method. In *AAAI*, volume 37, pp. 2654–2662, 2023b.

Wenjing Wang, Huan Yang, Jianlong Fu, and Jiaying Liu. Zero-reference low-light enhancement via physical quadruple priors, 2024.

Jiangwei Weng, Zhiqiang Yan, Ying Tai, Jianjun Qian, Jian Yang, and Jun Li. Mamballie: Implicit retinex-aware low light enhancement with global-then-local state space. *NeurIPS*, 37:27440–27462, 2024.

Chi-Wing Fu Xiaogang Xu, Ruixing Wang and Jiaya Jia. Snr-aware low-light image enhancement. In *CVPR*, 2022.

Rui Xu, Yuzhen Niu, Yuezhou Li, Huangbiao Xu, Wenxi Liu, and Yuzhong Chen. Urwkv: Unified rwkv model with multi-state perspective for low-light image restoration. In *CVPR*, pp. 21267–21276, 2025.

Qingsen Yan, Yixu Feng, Cheng Zhang, Guansong Pang, Kangbiao Shi, Peng Wu, Wei Dong, Jinqiu Sun, and Yanning Zhang. Hvi: A new color space for low-light image enhancement. In *CVPR*, pp. 5678–5687. IEEE, 2025.

Wenhan Yang, Wenjing Wang, Haofeng Huang, Shiqi Wang, and Jiaying Liu. Sparse gradient regularized deep retinex network for robust low-light image enhancement. *IEEE Trans. Image Process.*, 30:2072–2086, 2021.

Zhiwen Yang, Jiayin Li, Hao Lu, Hui Zhang, Zihua Wang, and Yan Xu. U-ttt: Towards generalizable pet image denoising via test-time training. *arXiv preprint*, 2026.

Xunpeng Yi, Han Xu, Hao Zhang, Linfeng Tang, and Jiayi Ma. Diff-retinex: Rethinking low-light image enhancement with a generative diffusion model. In *ICCV*, pp. 12302–12311, 2023.

Xiangrui Zeng, Lingyu Zhu, Wenhan Yang, Howard Leung, Shiqi Wang, and Sam Kwong. Low-light image enhancement via diffusion models with semantic priors of any region. *IEEE Trans. Circuits Syst. Video Technol.*, 2025.

Tongshun Zhang, Pingping Liu, Mengen Cai, Zijian Zhang, Yubing Lu, and Qiuzhan Zhou. Bs-mamba: Brightness and semantic modeling for long-range interaction in low-light image enhancement. *arXiv preprint arXiv:2506.18346*, 2025a.

Tongshun Zhang, Pingping Liu, Yubing Lu, Mengen Cai, Zijian Zhang, Zhe Zhang, and Qiuzhan Zhou. Cwnet: Causal wavelet network for low-light image enhancement. In *ICCV*, pp. 8789–8799, 2025b.

Weixia Zhang, Guangtao Zhai, Ying Wei, Xiaokang Yang, and Kede Ma. Blind image quality assessment via vision-language correspondence: A multitask learning perspective. In *CVPR*, pp. 14071–14081, 2023.

Kaihang Zheng, Jun Li, Hang Guo, Hongyu Chi, Zimo Liu, Tao Dai, Jinpeng Wang, and Yaowei Wang. Tttir: Unlocking instance-specific state evolution via test-time training for image restoration. *arXiv preprint*, 2026.

Yupeng Zheng, Chengliang Zhong, Pengfei Li, Huan-ang Gao, Yuhang Zheng, Bu Jin, Ling Wang, Hao Zhao, Guyue Zhou, Qichao Zhang, et al. Steps: Joint self-supervised nighttime image enhancement and depth estimation. In *ICRA*, pp. 4916–4923, 2023.

A Appendix

A.1 Limitations

T^{3}LIE relies on paired low-light supervision to learn the offline parameters and the initial fast weights. The inner-loop updates at test time also add computation beyond a single forward pass, so the inference cost grows with the number of adaptation steps. In addition, the fast-weight updates are driven only by input-derived Retinex objectives. Illumination patterns or noise levels far from the training distribution can therefore remain difficult, even though the model adapts to each test image.

A.2 Loss Details

End-to-End Training and Inference. All modules are trained end-to-end under the final enhancement objective. Let θ denote the offline-learned parameters of the overall network. For each training image, the illumination and reflectance fast weights are updated through their corresponding inner objectives, yielding W_L^* and W_R^*, respectively. The outer optimization is formulated as

$$\min_{\theta} \mathbb{E}_{(I,I^{\mathrm{gt}})} \left[\mathcal{L}\left(\widehat{I}_\theta(I; W_L^*, W_R^*), I^{\mathrm{gt}} \right) \right], \tag{16}$$

where I^{gt} denotes the corresponding well-exposed reference. Gradients are propagated through the inner updates, allowing the offline parameters and fast-weight initializations to be jointly optimized toward the final enhancement objective.

Given the enhanced image $\widehat{I}$ and the ground-truth image I^{gt}, the enhancement loss is defined as

$$\mathcal{L}_{\mathrm{enh}} = \lambda_1 \mathcal{L}_1 + \lambda_s \mathcal{L}_{\mathrm{SSIM}} + \lambda_e \mathcal{L}_{\mathrm{edge}} + \lambda_p \mathcal{L}_{\mathrm{perc}}, \tag{17}$$

where $\mathcal{L}_1$ denotes the pixel-wise ℓ_1 loss, $\mathcal{L}_{\mathrm{SSIM}} = 1 - \mathrm{SSIM}(\widehat{I}, I^{\mathrm{gt}})$ measures structural consistency, $\mathcal{L}_{\mathrm{edge}}$ penalizes edge discrepancy, and $\mathcal{L}_{\mathrm{perc}}$ is the LPIPS perceptual loss. We set $\lambda_1 = 1$, $\lambda_s = 0.5$, $\lambda_e = 50$, and $\lambda_p = 0.1$.

Table 7: User study results on low-light image enhancement.

Metrics	MultiNex	HVI-CIDNet	TPCNet	RetinexFormer	TTTIR	T^{3}LIE
Naturalness↑	3.14	3.67	3.65	3.86	3.98	4.26
Color Fidelity↑	3.11	3.47	3.49	3.75	3.84	4.10
Texture↑	2.94	3.48	3.41	3.67	3.77	4.04
Noise Suppression↑	2.96	3.29	3.33	3.51	3.78	4.03

To regularize the Max-Guided Retinex Decomposition, we further supervise the initial reflectance estimate R_0 using the well-exposed reference:

$$\mathcal{L}_{\max} = \left\| R_0 - I^{\mathrm{gt}} \right\|_1 . \tag{18}$$

This auxiliary term encourages MRD to provide a reliable reflectance initialization for subsequent illumination and reflectance adaptation.

The overall training objective is

$$\mathcal{L} = \mathcal{L}_{\mathrm{enh}} + \lambda_{\max}\mathcal{L}_{\max}, \tag{19}$$

where $\lambda_{\max} = 0.1$.

A.3 Ablation Details

Table 6 compares how illumination is coupled to the Retinex-inner TTT module. All variants share the same backbone, $T = 1$ inner step, and inner learning rates $\eta_L = \eta_R = 1$. Each column turns on one coupling while the others stay off. The first row runs TTT alone, without Gated, W_{illum}, Outer, or the Retinex-inner loss.

Gated. The illumination feature is taken as an additional input of RTTT. It is passed through a gate and then detached, so the gated illumination signal conditions RTTT, while the inner-loop gradient does not flow back into the illumination branch along this path. The gate is not inside the inner objective; it only controls the input that RTTT receives.

W_{illum}**.** A separate weight is predicted from the illumination features and added to the original reflectance fast weight before the inner update,

$$W \leftarrow W + W_{\mathrm{illum}}. \tag{20}$$

Illumination therefore shifts the fast-weight initialization, rather than gating the features that enter RTTT or changing the Retinex-inner target.

Outer. The inner TTT update is run first, without illumination gating. After the fast weights have been adapted, a gate selects the adapted response. Illumination thus acts as a post-TTT selector outside the inner optimization.

Retinex-inner Loss. Gated, W_{illum}, and Outer are all turned off, and illumination enters the inner objective through $G \odot V_R$ in equation 12. This is the full model, and it obtains the best score in Table 6.

A.4 Gradient Visualization

To further examine the behavior of Retinex-inner test-time training, we visualize the gradient magnitude generated during inner-loop adaptation in Figure. 7. The gradients exhibit clear spatial variations across regions with different illumination conditions, with stronger responses generally concentrated in areas requiring larger enhancement corrections. Moreover, the response patterns vary across test images, indicating that the inner optimization adapts to image-specific illumination characteristics rather than applying a fixed update pattern. These observations demonstrate that T^{3}LIE performs spatially adaptive and image-specific refinement through Retinex-inner TTT.

A.5 Gradient Analysis of Retinex-inner TTT

We further analyze Retinex-inner TTT with the same normalized Illum-inner Fast-weight Update as equation 13. Let $K, V, Q \in \mathbb{R}^{N\times d}$ denote the key, value, and query features, and let the inner

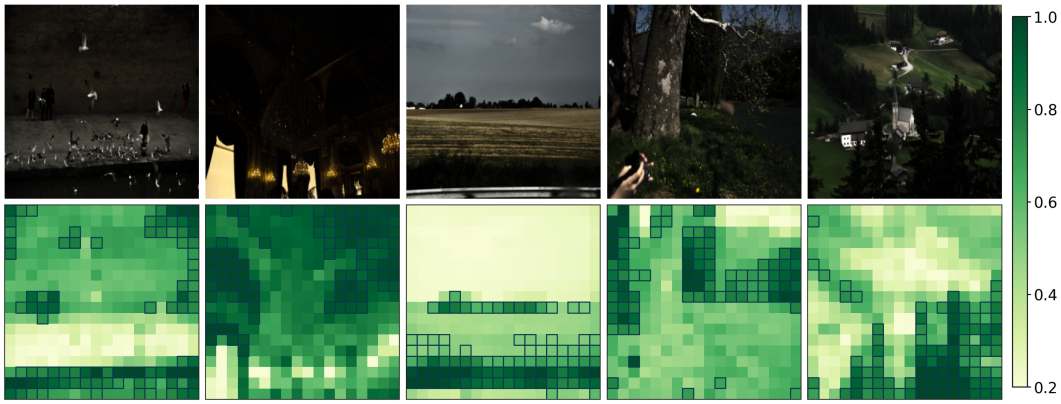


Figure 7: Gradient magnitude visualization during Retinex-inner test-time training. T^{3}LIE produces spatially varying gradient responses that adapt to the illumination characteristics of each input image.

learner predict $\widehat{V} = f_W(K)$. Denote the raw loss gradient by $g = \nabla_W \mathcal{L}$, and write $\mathcal{G}_W = g$ in the chain rules below. IFU does not step along g. Each column is normalized before the fast-weight update,

$$\bar{g}_{:,j} = \frac{g_{:,j}}{1 + \|g_{:,j}\|_2}, \qquad W^* = W - \eta_R \bar{g}, \tag{21}$$

with $\eta_R = 1$, $s = 1$, and $T = 1$ inner step, matching equation 13. Vanilla TTT and RTTT share this normalization. Illumination changes only the target that produces g, while $g \mapsto \bar{g}$ rescales column j by the positive factor $1/(1 + \|g_{:,j}\|_2)$ and does not remove the element-wise weights inside g. The adapted weights W^* are applied to the query Q.

Vanilla TTT. Without illumination conditioning, the inner optimization uses the value feature V directly as the reconstruction target. The dot-product objective can be written as

$$\mathcal{L}_{\mathrm{TTT}} = -\frac{s}{N} \left\langle \widehat{V}, V \right\rangle_F = -\frac{s}{N} \sum_{i=1}^{N} \sum_{j=1}^{d} \widehat{V}_{ij} V_{ij}, \tag{22}$$

where $s = 1$ is the scaling factor.

The gradient with respect to the predicted representation is

$$\frac{\partial \mathcal{L}_{\mathrm{TTT}}}{\partial \widehat{V}_{ij}} = -\frac{s}{N} V_{ij}. \tag{23}$$

Taking the derivative again with respect to V_{kl} gives

$$\frac{\partial^2 \mathcal{L}_{\mathrm{TTT}}}{\partial \widehat{V}_{ij} \partial V_{kl}} = -\frac{s}{N} \delta_{ik} \delta_{jl}, \tag{24}$$

where δ denotes the Kronecker delta. Therefore, after defining $\hat{v} = \mathrm{vec}(\widehat{V})$ and $v = \mathrm{vec}(V)$, the mixed derivative becomes

$$\frac{\partial^2 \mathcal{L}_{\mathrm{TTT}}}{\partial \hat{v}\, \partial v} = -\frac{s}{N} \mathbf{I}, \tag{25}$$

where $\mathbf{I} \in \mathbb{R}^{Nd \times Nd}$ is the identity matrix. Hence, vanilla TTT establishes a uniform element-wise coupling between $\widehat{V}$ and V.

To see how this affects the inner optimization, let θ_V denote the parameters of the value projection, and define the Jacobians $J_W = \partial \hat{v} / \partial W$ and $J_V = \partial v / \partial \theta_V$. The Jacobian of the raw gradient $g = \mathcal{G}_W$, before the column normalization in equation 21, is

$$\frac{\partial \mathcal{G}_W}{\partial \theta_V} = J_W^\top \frac{\partial^2 \mathcal{L}_{\mathrm{TTT}}}{\partial \hat{v}\, \partial v} J_V = -\frac{s}{N} J_W^\top J_V. \tag{26}$$

Thus, the value-dependent signal is propagated uniformly through the inner optimization.

Retinex-inner TTT. The RRCM analysis shows that reflectance correction is dependent on the illumination state. Accordingly, RTTT introduces the illumination guidance $G \in \mathbb{R}^{N\times d}$ into the inner target:

$$V^G = G \odot V. \tag{27}$$

The corresponding inner objective is

$$\mathcal{L}^R_{\text{inner}} = -\frac{s}{N}\left\langle \widehat{V}, G \odot V \right\rangle_F = -\frac{s}{N}\sum_{i=1}^{N}\sum_{j=1}^{d} \widehat{V}_{ij} G_{ij} V_{ij}. \tag{28}$$

Its first-order gradient with respect to $\widehat{V}$ is

$$\frac{\partial \mathcal{L}^R_{\text{inner}}}{\partial \widehat{V}_{ij}} = -\frac{s}{N} G_{ij} V_{ij}. \tag{29}$$

Further differentiating with respect to V_{kl} yields

$$\frac{\partial^2 \mathcal{L}^R_{\text{inner}}}{\partial \widehat{V}_{ij} \partial V_{kl}} = -\frac{s}{N} G_{ij} \delta_{ik} \delta_{jl}. \tag{30}$$

After vectorization, we obtain

$$\frac{\partial^2 \mathcal{L}^R_{\text{inner}}}{\partial \hat{v}\, \partial v} = -\frac{s}{N}\mathbf{D}_G, \qquad \mathbf{D}_G = \operatorname{Diag}\left(\operatorname{vec}(G)\right). \tag{31}$$

Compared with equation 25, the fixed identity coupling $\mathbf{I}$ in vanilla TTT is replaced by the illumination-dependent diagonal modulation $\mathbf{D}_G$. Consequently, the raw gradient that enters the normalized IFU satisfies

$$\frac{\partial \mathcal{G}_W}{\partial \theta_V} = -\frac{s}{N} J_W^\top \mathbf{D}_G J_V. \tag{32}$$

Therefore, different spatial locations and feature dimensions contribute to fast-weight adaptation according to the illumination state of the current input. IFU then steps with $W^* = W - \eta_R \bar{g}$ as in equation 13, rather than along g.

The difference between the two gradient couplings can be written explicitly as

$$\Delta\mathbf{H} = \frac{\partial^2 \mathcal{L}^R_{\text{inner}}}{\partial \hat{v}\, \partial v} - \frac{\partial^2 \mathcal{L}_{\text{TTT}}}{\partial \hat{v}\, \partial v} = -\frac{s}{N}\left(\mathbf{D}_G - \mathbf{I}\right). \tag{33}$$

When $G = \mathbf{1}$, we have $\mathbf{D}_G = \mathbf{I}$, and RTTT reduces exactly to the vanilla TTT objective. For spatially varying G, however, the inner optimization becomes input-dependent and assigns different optimization strengths to different spatial and feature responses.

Gradient Coupling with the Illumination Pathway. In the inner update, G is treated as fixed with respect to the reflectance fast weights W. However, it remains differentiable during outer-loop training. Therefore, the reflectance inner objective also establishes a learning path toward the illumination pathway.

Specifically, differentiating equation 28 with respect to G gives

$$\frac{\partial \mathcal{L}^R_{\text{inner}}}{\partial G_{ij}} = -\frac{s}{N} \widehat{V}_{ij} V_{ij}, \tag{34}$$

and the corresponding mixed derivative is

$$\frac{\partial^2 \mathcal{L}^R_{\text{inner}}}{\partial \widehat{V}_{ij} \partial G_{kl}} = -\frac{s}{N} V_{ij} \delta_{ik} \delta_{jl}. \tag{35}$$

Hence, after vectorization,

$$\frac{\partial^2 \mathcal{L}^R_{\text{inner}}}{\partial \hat{v}\, \partial g} = -\frac{s}{N}\mathbf{D}_V, \qquad \mathbf{D}_V = \operatorname{Diag}\left(\operatorname{vec}(V)\right), \tag{36}$$

where $g = \mathrm{vec}(G)$. This shows that the reflectance inner optimization is differentiably coupled with the illumination guidance during end-to-end training.

Overall, vanilla TTT and RTTT share the normalized IFU and differ only in the target that forms g:

$$\begin{aligned} \text{Vanilla TTT:} \quad & V \to g = \nabla_W \mathcal{L}_{\mathrm{TTT}} \to \bar{g} \to W^* = W - \eta_R \bar{g} \to f_{W^*}(Q), \\ \text{RTTT:} \quad & G \odot V \to g = \nabla_W \mathcal{L}^R_{\mathrm{inner}} \to \bar{g} \to W^* = W - \eta_R \bar{g} \to f_{W^*}(Q). \end{aligned} \tag{37}$$

RTTT therefore preserves the original TTT optimization mechanism while conditioning its inner target and the resulting fast-weight update on image-specific illumination. This provides a direct optimization-level realization of the illumination-dependent reflectance refinement derived by RRCM.

A.6 Implementation Details of Fast Weight Update

Fast-weight Architecture. The TTT module employs two lightweight fast-weight branches. Given the reflectance feature, a linear projection produces two sets of query, key, and value representations, (Q_1, K_1, V_1) and (Q_2, K_2, V_2). The first branch adopts a simplified SwiGLU inner learner:

$$\widehat{V}_1 = (K_1 W_1) \odot \mathrm{SiLU}(K_1 W_2), \tag{38}$$

where W_1 and W_2 are the corresponding fast weights. The second branch employs a 3×3 depth-wise convolution:

$$\widehat{V}_2 = \mathrm{DWConv}_{W_3}(K_2), \tag{39}$$

where W_3 denotes the depth-wise convolutional fast weights.

The illumination guidance modulates the value-side signal used for fast-weight adaptation in both branches. For the SwiGLU branch, each gradient column is normalized before updating the fast weights:

$$\bar{g}_{m,:,j} = \frac{g_{m,:,j}}{1 + \|g_{m,:,j}\|_2}, \qquad W_m^* = W_m - \eta_R \bar{g}_m, \quad m \in \{1, 2\}. \tag{40}$$

For the depth-wise convolution branch, the gradient of each spatial kernel is normalized as

$$\bar{g}_{3,c} = \frac{g_{3,c}}{1 + \|g_{3,c}\|_2}, \qquad W_3^* = W_3 - \eta_R \bar{g}_3, \tag{41}$$

where $g_{3,c} \in \mathbb{R}^{3\times3}$ denotes the gradient of the c-th depth-wise kernel.

The adapted fast weights are subsequently applied to the query representations:

$$O_1 = (Q_1 W_1^*) \odot \mathrm{SiLU}(Q_1 W_2^*), \qquad O_2 = \mathrm{DWConv}_{W_3^*}(Q_2). \tag{42}$$

Finally, the two adapted representations are concatenated and projected:

$$O_R = \mathrm{Proj}\left([O_1, O_2]\right). \tag{43}$$

A.7 Additional Qualitative Results

Additional qualitative comparisons are shown in Figs. 8–15.

A.8 User Study

We conduct a user study to evaluate the perceptual quality of enhanced images under diverse low-light conditions. We randomly select 64 low-light images from paired and unpaired datasets, covering different illumination levels and scene contents. For each participant, eight subjective evaluation questions are randomly sampled from this image pool, and the enhanced results are displayed in randomized order without revealing the method identities. The sampling strategy ensures that each participant evaluates at least four results generated by the same method. A total of 45 participants rate the enhanced images on a five-point scale from 1 (worst) to 5 (best).

Following the perceptual characteristics of low-light enhancement, we evaluate four aspects: *Naturalness*, which measures whether the enhanced image maintains a realistic visual appearance; *Color Fidelity*, which evaluates color consistency and the absence of color distortion; *Texture*, which assesses the preservation and recovery of local details; and *Noise Suppression*, which measures the effectiveness of reducing amplified noise and artifacts. The average scores across all participants are reported in Table 7.



Figure 8: Additional qualitative results.

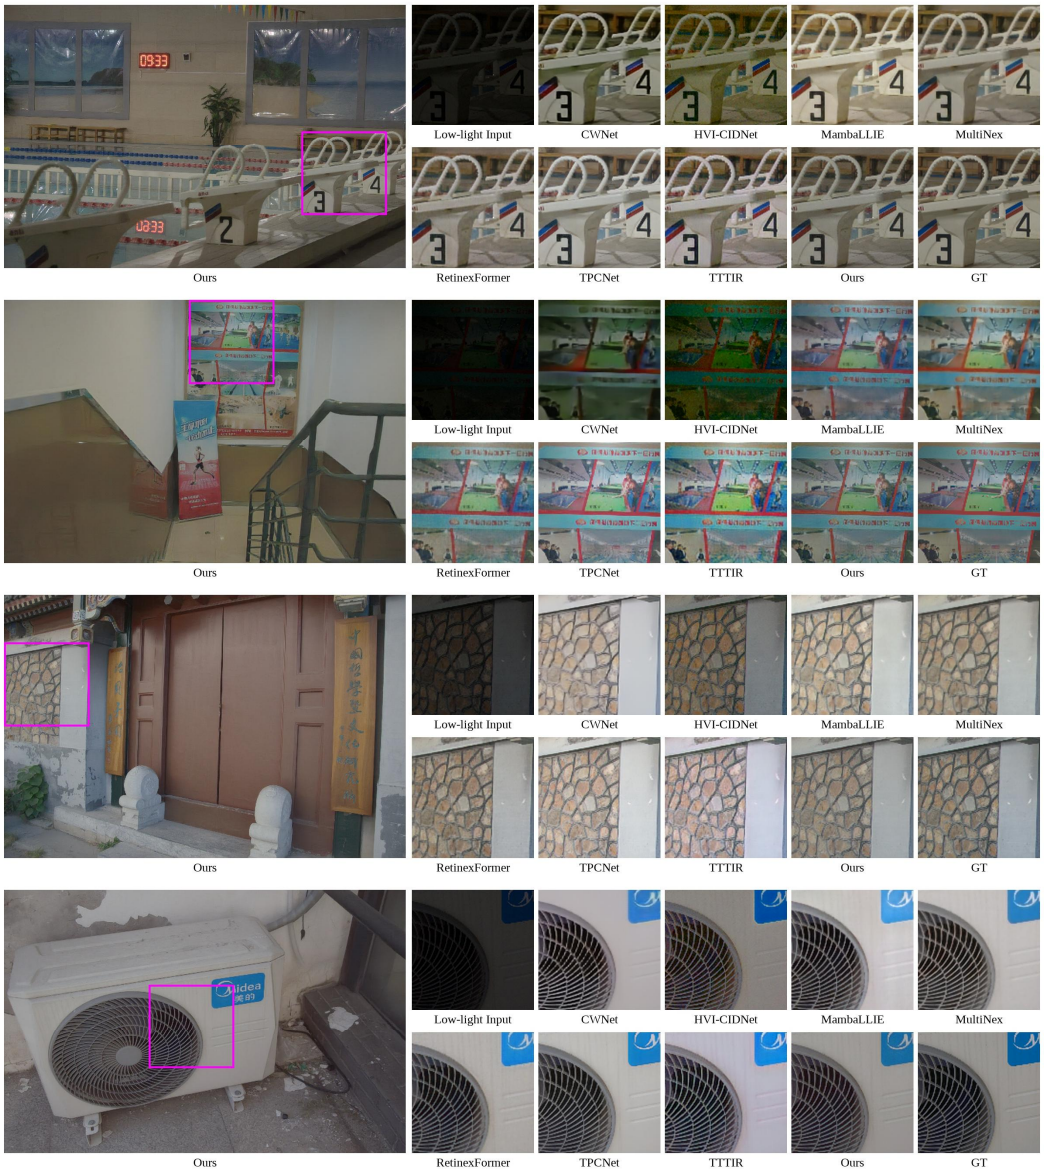


Figure 9: Additional qualitative results.

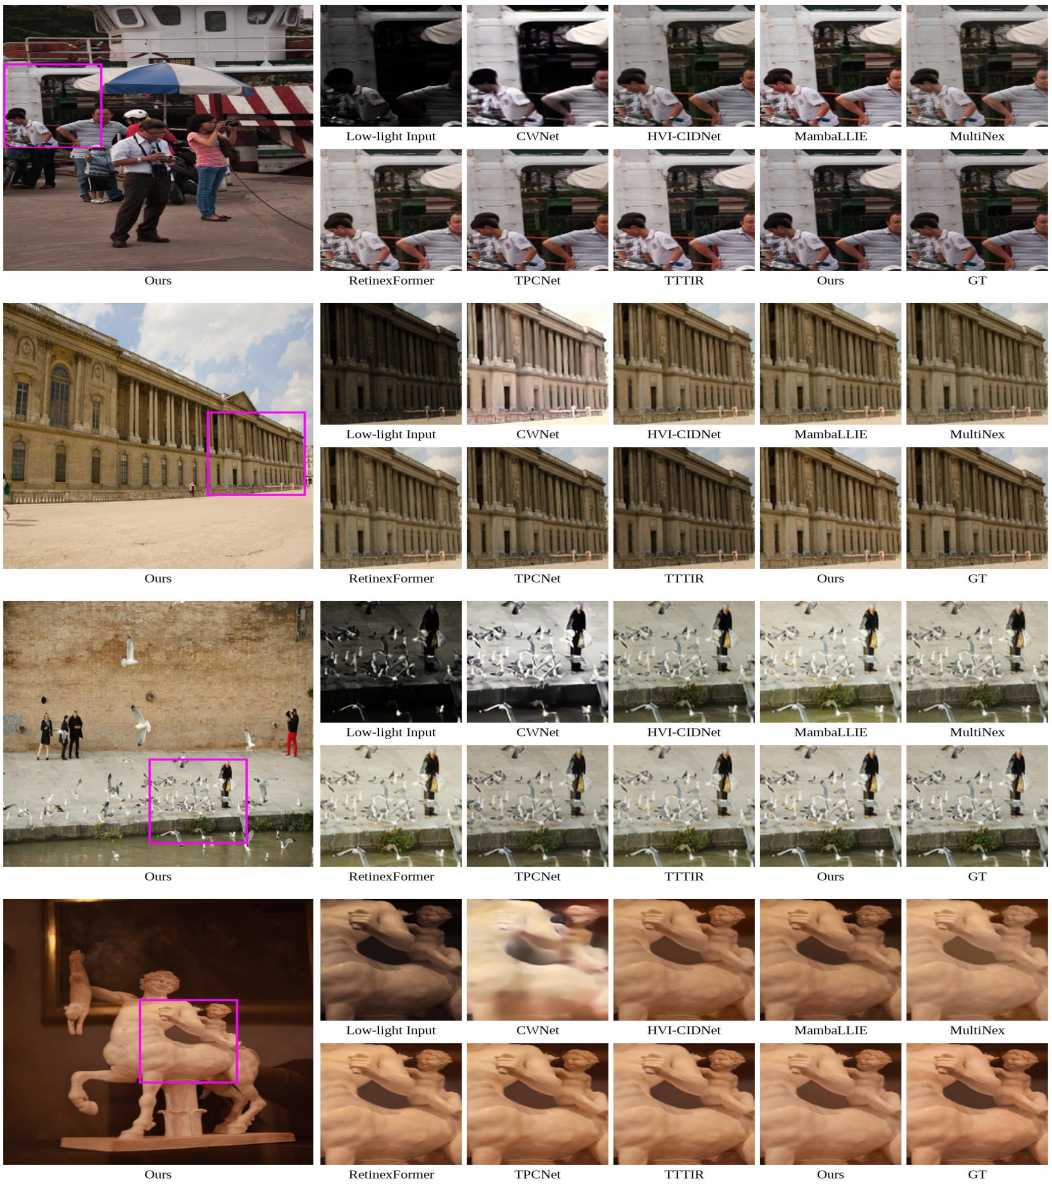


Figure 10: Additional qualitative results.

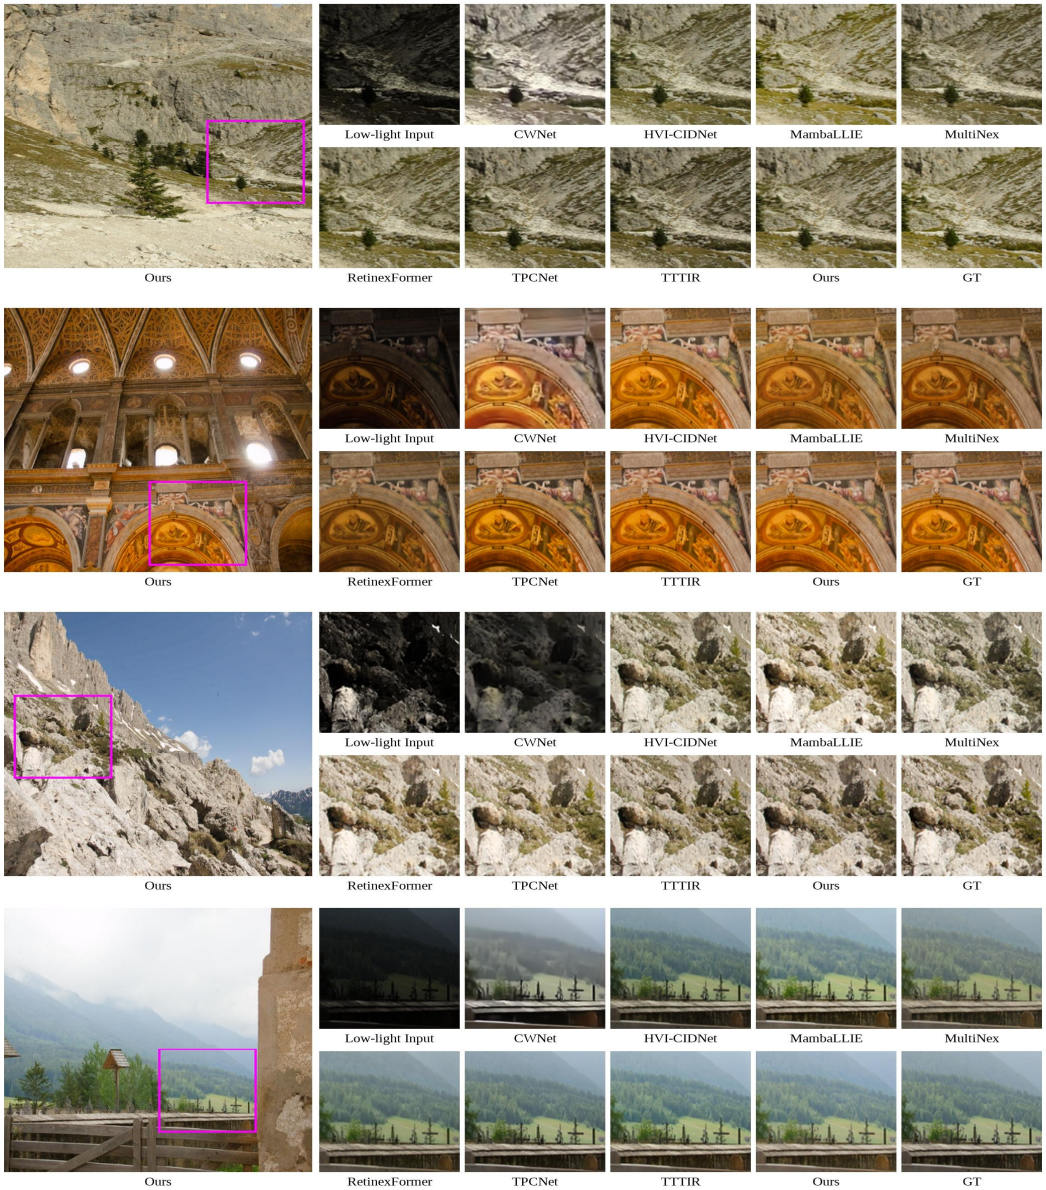


Figure 11: Additional qualitative results.

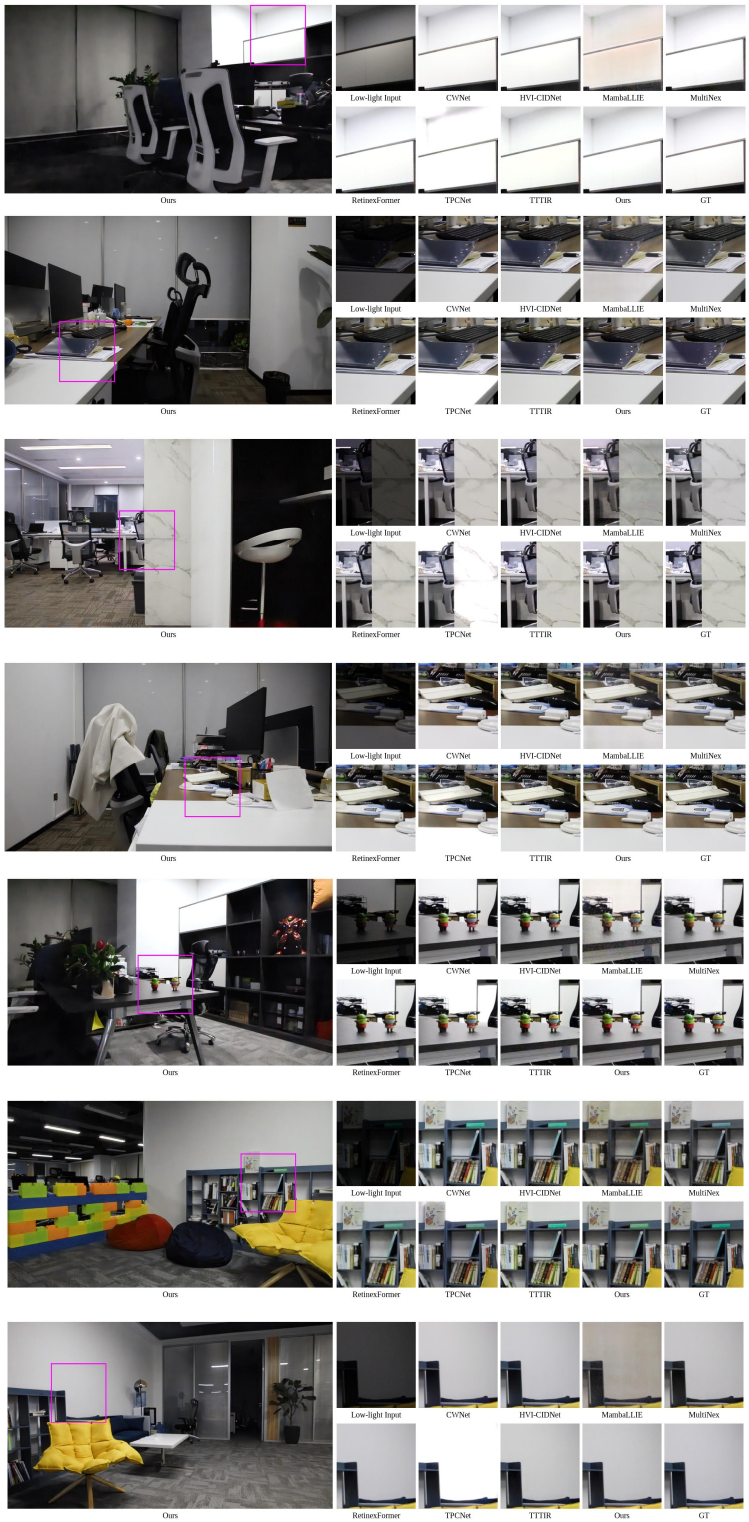


Figure 12: Additional qualitative results.



Figure 13: Additional qualitative results.

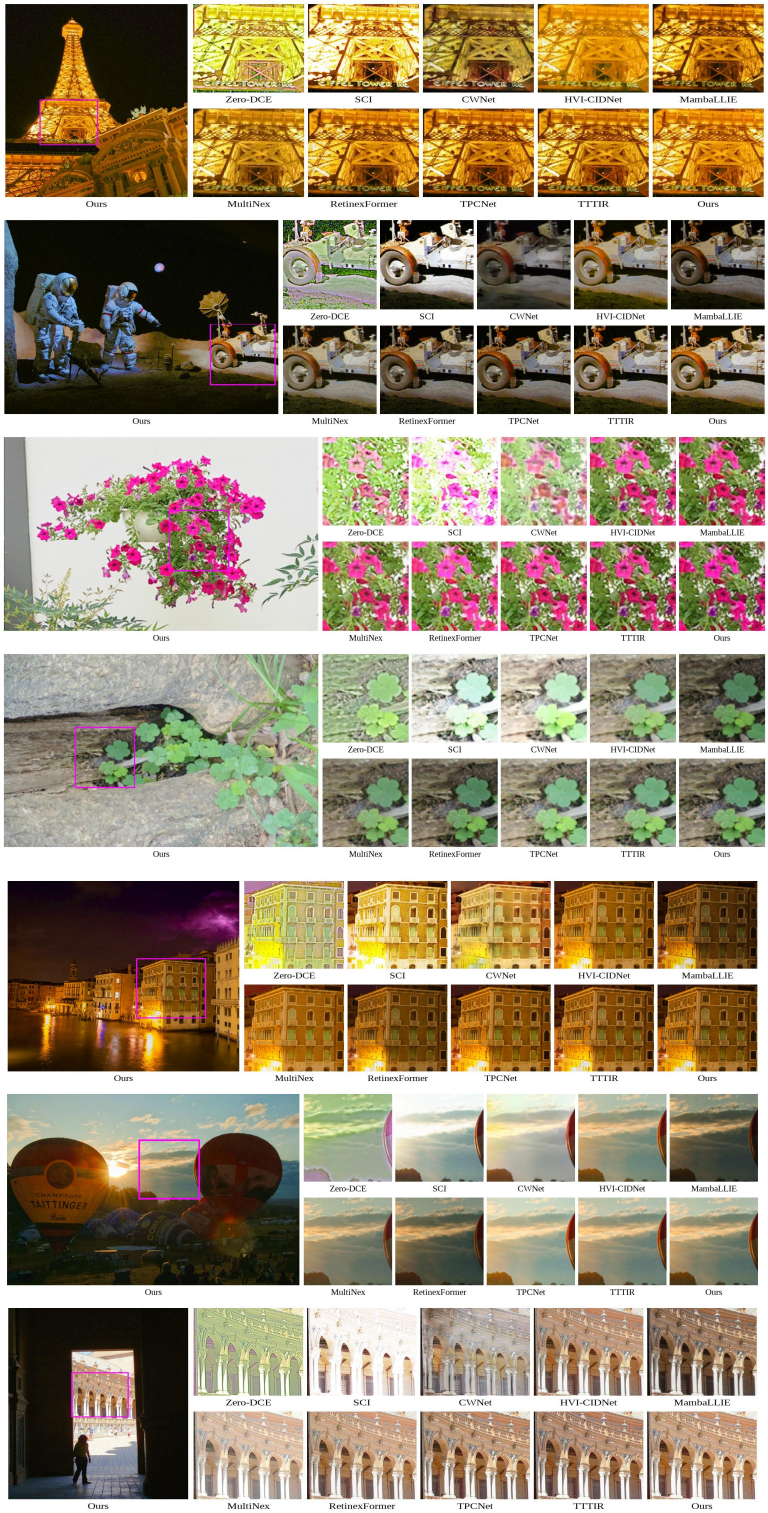


Figure 14: Additional qualitative results.

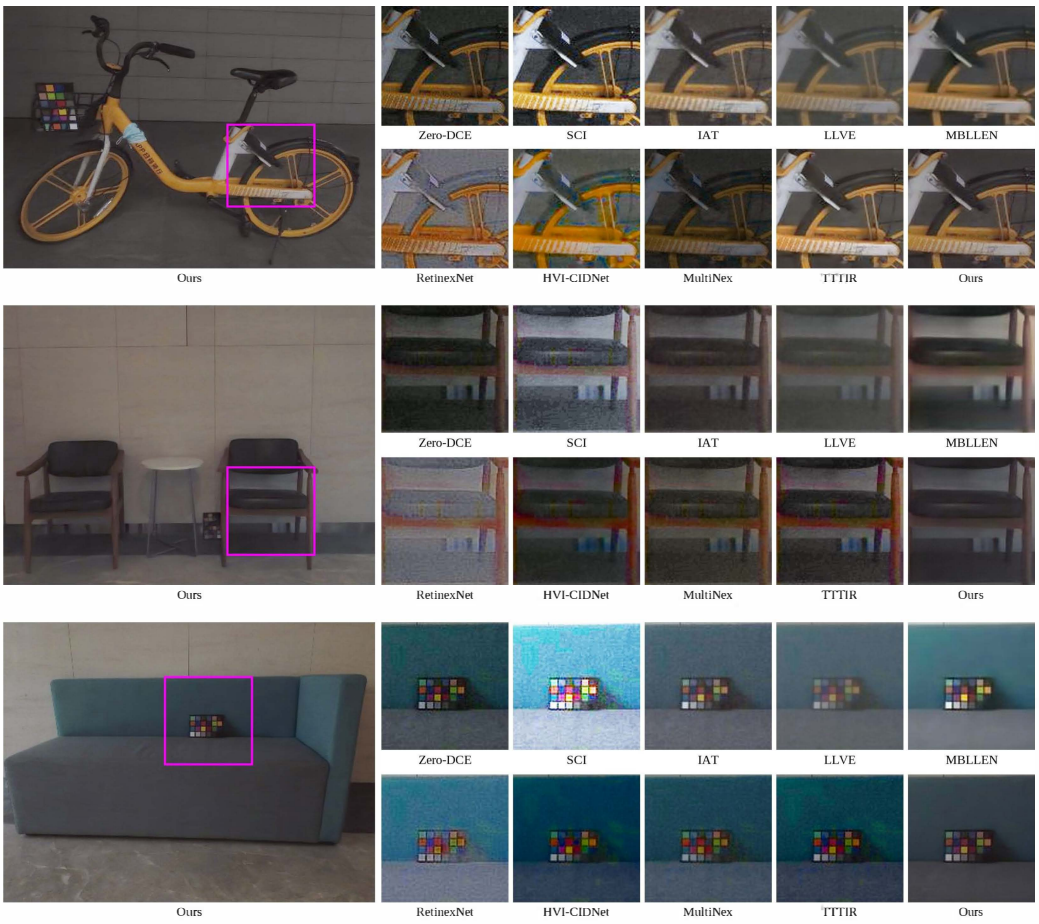


Figure 15: Additional qualitative results.

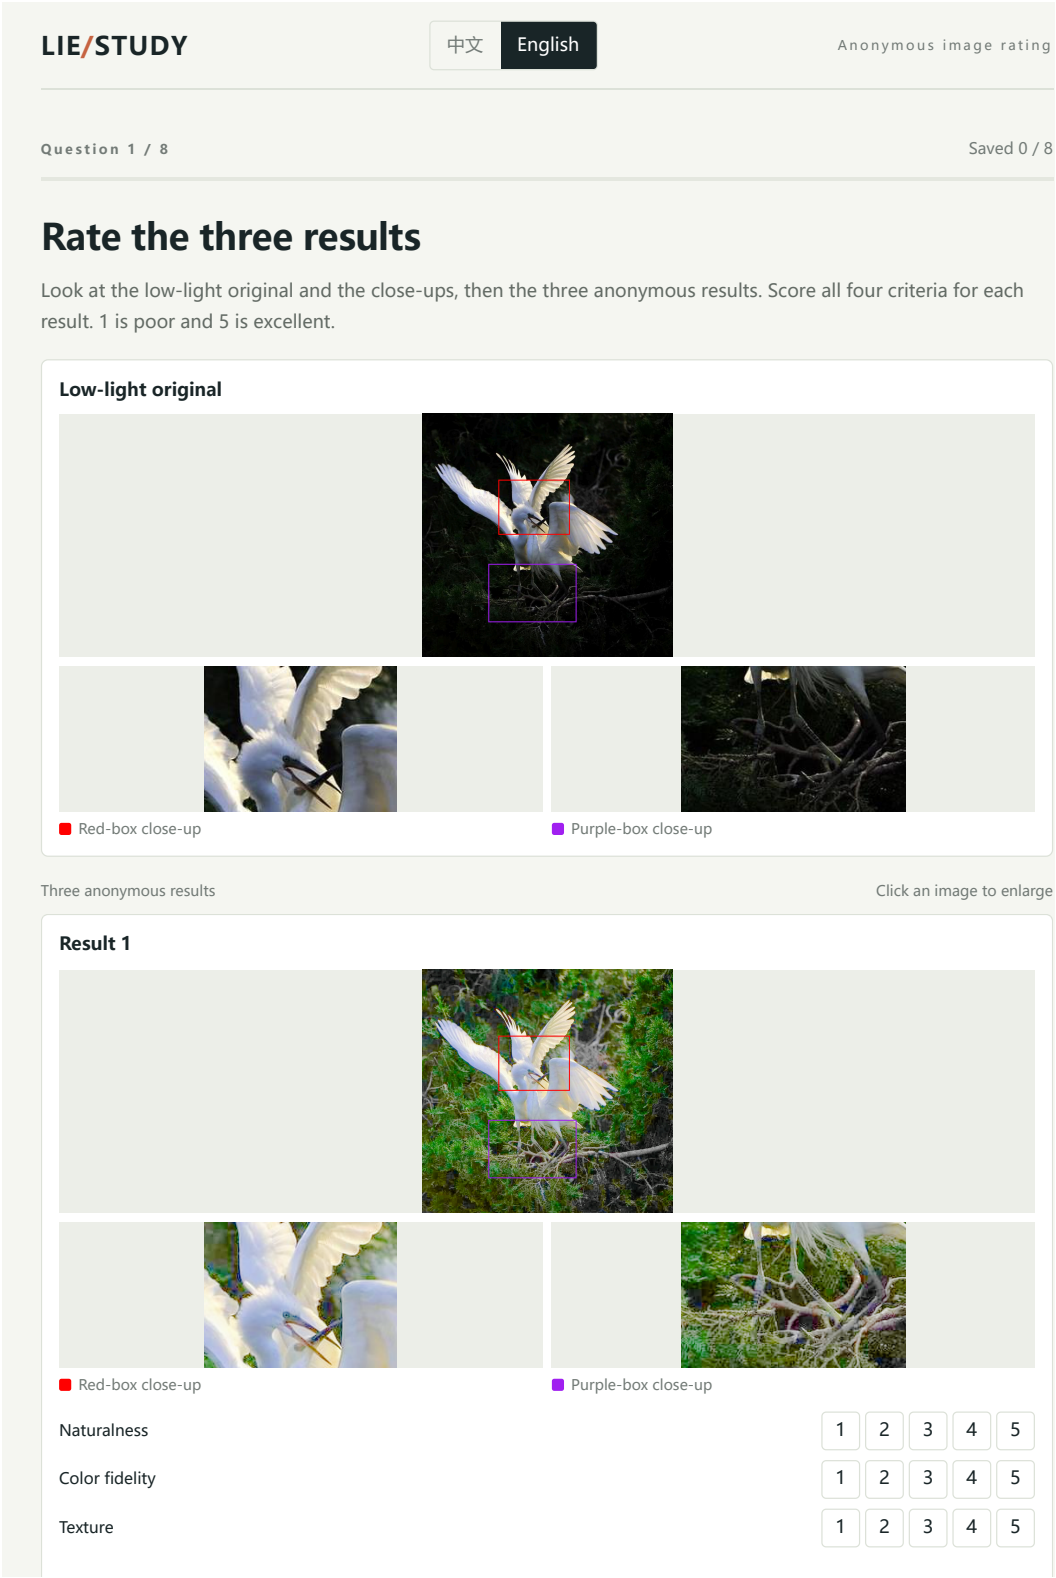


Figure 16: User study interface.