R2RFLOW: REFLECTANCE-TO-RESTORATION FLOW WITH COMPLEMENTARY EVIDENCE FOR LOW LIGHT IMAGE ENHANCEMENT

Anonymous authors
Paper under double-blind review

ABSTRACT

Low-light image enhancement aims to recover visually pleasing normal-light images while preserving the scene content concealed by complex illumination degradation. However, existing generative approaches often start from uninformative noise without directly using the scene content preserved in the low-light observation to initialize the transport. Existing flow objectives primarily supervise the underlying flow field, without directly exploiting the known restoration endpoint available in paired restoration. To address these limitations, we propose R2RFlow, a reflectance-to-restoration flow framework integrating complementary evidence for one-step low-light image enhancement. R2RFlow initializes the transport from an observation-derived reflectance state and constructs finite-interval transport toward the paired normal-light image, explicitly leveraging scene content available in the degraded observation. To complement restoration-relevant information not fully retained by the reflectance source, we further incorporate structural, illumination, and frequency cues through role-specific conditioning pathways. In addition, a restoration-aligned multi-path learning strategy supervises the shared transport operator from the deployment source, reference intermediate states, and model-induced rollout states through their final restoration outcomes. Experiments on the LOL-v1 and LOL-v2-Real benchmarks show that R2RFlow achieves competitive performance against existing low-light enhancement methods. When trained exclusively on LOL-v1, R2RFlow also achieves favorable performance on LSRW-Nikon and LSRW-Huawei without target-domain fine-tuning, demonstrating favorable cross-dataset transfer.

1 INTRODUCTION

Low-light image enhancement (LLIE) aims to recover normal-light images from observations degraded by insufficient and spatially non-uniform illumination while preserving scene structure and fine details. Recent learning-based methods have made substantial progress with Retinex priors, illumination modeling, and generative approaches. In particular, diffusion- and flow-based methods provide flexible restoration models, but often retain the noise-to-data paradigm of generic generation despite the degraded observation already containing substantial scene information.

This distinction is particularly important for paired LLIE, where the low-light observation and its normal-light reference depict the same scene with strong spatial correspondence. Rather than starting from uninformative noise, an observation-derived representation may therefore provide a more informative source for restoration. Meanwhile, paired restoration provides another task-specific advantage: the desired normal-light endpoint is directly available during training. This enables finite-interval transport to be supervised through the restoration outcomes it produces, beyond conventional supervision of the underlying flow field.

Motivated by these observations, we propose **R2RFlow**, a reflectance-to-restoration flow framework with complementary evidence for one-step low-light image enhancement. R2RFlow initializes finite-interval transport from an observation-derived reflectance representation, which exhibits a smaller discrepancy to the restoration target than alternative source representations. Since source proximity does not imply that all restoration-relevant information is preserved, we further incor-

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.

porate structural, illumination, and frequency evidence through role-specific conditioning pathways. Finally, restoration-aligned multi-path learning supervises the shared transport operator from the deployment source, reference intermediate states, and model-induced rollout states through their final restoration outcomes, while retaining one-step inference.

Experiments on LOL-v1 (**?**) and LOL-v2-Real (Yang et al., 2021) demonstrate competitive performance against representative low-light enhancement methods. The LOL-v1-trained model also retains favorable performance on LSRW-Nikon and LSRW-Huawei without target-domain fine-tuning. Ablation studies further validate the three components of R2RFlow.

Our main contributions are summarized as follows:

- We formulate LLIE as a **reflectance-to-restoration finite-interval flow**, replacing noise initialization with an observation-derived reflectance source.
- We develop **complementary evidence conditioning** that incorporates structural, illumination, and frequency information through role-specific pathways to complement the reflectance source.
- We introduce **restoration-aligned multi-path learning**, supervising the shared transport operator from deployment, reference, and model-induced states through their restoration outcomes while retaining one-step inference.

2 Related Work

2.1 Low-Light Image Enhancement

Learning-based low-light image enhancement has increasingly incorporated decomposition and degradation-aware priors to recover normal-light appearance. RetinexNet (Wei et al., 2018) learns reflectance–illumination decomposition, while Retinexformer (Cai et al., 2023) uses illumination representations to guide feature interactions. SNR-Net (Xu et al., 2022) and HVI-CIDNet (Yan et al., 2025) further exploit signal-to-noise estimates and decoupled color representations, respectively. Generative approaches model restoration distributions through conditional normalizing flows, as in LLFlow (Wang et al., 2022), or diffusion. Diff-Retinex (Yi et al., 2023) restores decomposed Retinex components, while GSAD (Hou et al., 2023) incorporates global structure into the diffusion process. R2RFlow instead separates source construction from complementary conditioning: reflectance defines the transport source, while structural, illumination, and frequency cues provide role-specific restoration evidence.

2.2 Flow-Based and One-Step Generative Modeling

Flow-based generative models construct transport between source and target distributions through learned transformations or continuous dynamics. Flow Matching (Lipman et al., 2023) learns instantaneous velocity fields along probability paths and typically relies on numerical integration at inference. Recent work has increasingly explored reducing this iterative sampling process. Consistency Models (Song et al., 2023) learn to map states along a trajectory toward a common endpoint, while MeanFlow (Geng et al., 2025) models average velocities over finite time intervals to support one-step generation.

R2RFlow considers this problem in paired image restoration, where the degraded observation and its desired normal-light endpoint are both available. Rather than transporting from an uninformed noise prior, R2RFlow constructs the source from observation-derived reflectance and adopts a finite-interval transport parameterization. The shared operator is further optimized through restoration outcomes from the deployment source, reference intermediate states, and model-induced rollout states.

3 Method

The proposed R2RFlow framework formulates low-light image enhancement as a one-step transport from an observation-derived reflectance state to its normal-light counterpart. We first construct

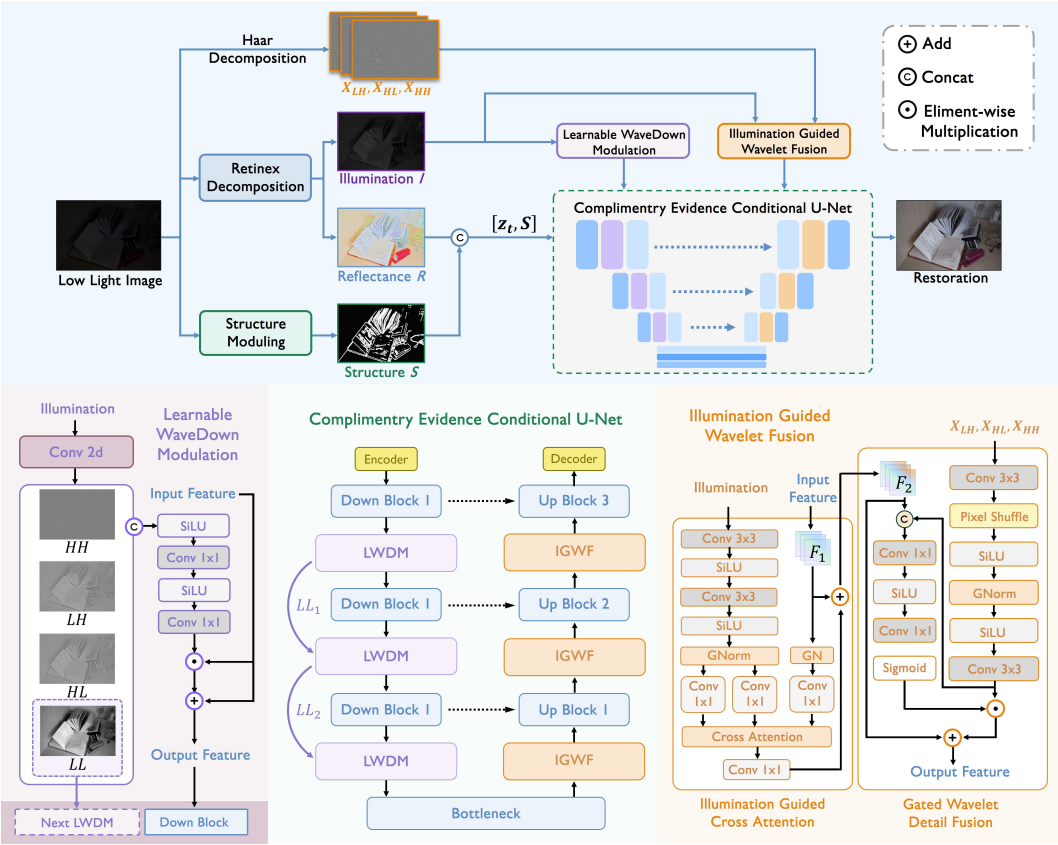


Figure 1: Overview of R2RFlow. Reflectance serves as the source state for one-step restoration, while complementary structural, illumination, and frequency evidence is incorporated through input enrichment, Learnable WaveDown Modulation (LWDM), and Illumination-Guided Wavelet Fusion (IGWF), respectively. The lower panels detail the multi-scale architecture and the designs of LWDM and IGWF.

a reflectance-informed source and define a finite-interval transport operator for restoration (Section 3.1). We further incorporate complementary structural, illumination, and frequency evidence through role-specific conditioning pathways (Section 3.2). Finally, we introduce restoration-aligned multi-path learning, which supervises the shared operator through its final restoration outcomes from the deployment source, reference intermediate states, and model-induced rollout states (Section 3.3).

3.1 Reflectance-to-Restoration Flow

In paired low-light image enhancement, the low-light observation and its normal-light reference depict the same scene with strong spatial correspondence. This inherent correspondence motivates initializing the transport from an observation-derived state rather than an uninformative prior. We therefore seek an informative source representation with reduced discrepancy to the restoration target.

Following the Retinex formulation, we decompose the low-light observation into reflectance and illumination components to obtain an observation-derived representation less dominated by illumination degradation. We employ a pre-trained Retinex decomposition network $\mathcal{D}$ to obtain a re-

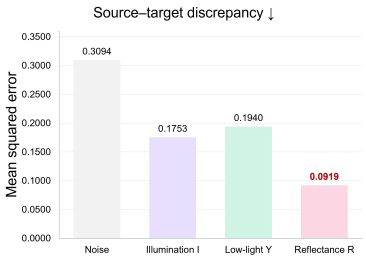


Figure 2: Source–target MSE for different representations. Reflectance achieves the lowest discrepancy. Lower is better.

flectance component R and an illumination map I:

$$(R, I) = \mathcal{D}(Y). \tag{1}$$

We treat R as a reflectance-informed representation of the observed scene, without requiring it to be a physically exact reflectance estimate. As shown in Fig. 2, R exhibits the smallest source–target MSE among the candidate representations, reducing the discrepancy from Y by 52.6%. This motivates our choice of R as the source state of the restoration flow, a design further validated by the initialization ablation in Table 3.

With this source choice, we formulate enhancement as transport from R to X. For each paired sample, we define a linear reference path and its conditional velocity:

$$z_t = (1-t)X + tR, \qquad v_c = \frac{\mathrm{d}z_t}{\mathrm{d}t} = R - X, \qquad t \in [0, 1]. \tag{2}$$

Here, $z_1 = R$ is the source state and $z_0 = X$ is the restoration endpoint. The interpolation also provides reference intermediate states used by the multi-path learning strategy in Section 3.3.

To represent transport over finite intervals, we adopt the finite-interval parameterization of MeanFlow (Geng et al., 2025). An interval-conditioned predictor $u_\theta(z_t, t, r; \mathcal{C})$ parameterizes the transport over the requested interval $t \to r$, where $0 \le r < t \le 1$. We define the corresponding learned transport operator as

$$T_\theta(z_t, t, r; \mathcal{C}) = z_t - (t-r)u_\theta(z_t, t, r; \mathcal{C}), \tag{3}$$

where $\mathcal{C}$ denotes the complementary restoration evidence introduced in Section 3.2. This parameterization allows the same predictor to represent both complete source-to-target restoration and partial transport between intermediate states.

At inference, R2RFlow initializes the flow from R and applies the operator over the full interval from $t = 1$ to $r = 0$:

$$\widehat{X} = T_\theta(R, 1, 0; \mathcal{C}) = R - u_\theta(R, 1, 0; \mathcal{C}). \tag{4}$$

Once the source and conditioning evidence have been extracted, restoration requires a single evaluation of the flow predictor, without iterative numerical integration. However, a source representation with reduced target discrepancy does not necessarily retain all restoration-relevant information. We therefore complement the reflectance initialization with evidence preserved in the illumination component and the original observation, as described next.

3.2 Complementary Evidence Conditioning

The reflectance source exhibits the smallest pixel-space discrepancy to the restoration target, but source proximity alone does not imply that all restoration-relevant information is preserved. As shown in Fig. 3, both the illumination component I and the original observation Y exhibit higher local structural correlation with the target than R. The representation best suited for initializing transport therefore need not contain all evidence useful for restoration. We retain I and Y as complementary evidence alongside the reflectance initialization.

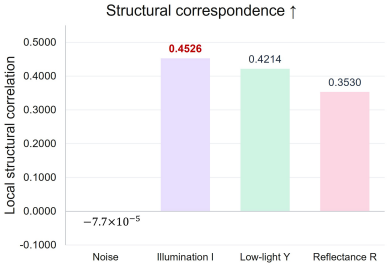


Figure 3: Local structural correlation with the normal-light target. Illumination and low-light observations show higher correlations than reflectance. Higher is better.

Rather than treating these cues as interchangeable conditions, we incorporate them according to their distinct roles in restoration. A structural prior extracted from Y enriches the flow-state representation, illumination-derived features modulate hierarchical encoding through Learnable WaveDown Modulation (LWDM), and Illumination-Guided Wavelet Fusion (IGWF) combines illumination guidance with observation-derived frequency details during decoding.

Evidence-enriched input. Given the low-light observation Y, we obtain a structural prior $S = \mathcal{G}_{\mathrm{SMG}}(Y)$ using a pretrained SMG model (Xu et al., 2023). We concatenate S with the current flow state z_t to form the input representation:

$$F^0 = \phi_{\mathrm{in}}\left([z_t \| S]\right), \tag{5}$$

where $\|$ denotes channel-wise concatenation and ϕ_{in} is the input projection. This directly supplements the state to be transported with observation-derived structural evidence.

Learnable WaveDown Modulation. The illumination component I retains spatial variations that provide useful cues for restoration. We introduce LWDM to extract this information hierarchically and modulate encoder representations at multiple scales. Starting from an illumination embedding $F_I^0 = \phi_I(I)$, we construct a hierarchy of illumination-derived features and condition the encoder as

$$F_I^\ell = \mathcal{W}_\ell(F_I^{\ell-1}), \qquad \widetilde{E}^\ell = \mathcal{M}_\ell(E^\ell, F_I^\ell), \tag{6}$$

where E^ℓ denotes the encoder feature at scale ℓ. The WaveDown operator $\mathcal{W}_\ell$ employs Haar-initialized learnable wavelet filters to extract multi-scale frequency responses, while $\mathcal{M}_\ell$ adaptively modulates the corresponding encoder representation.

Illumination-Guided Wavelet Fusion. During decoding, IGWF complements the incoming representation with illumination guidance from I and local frequency details retained in the original observation Y. It consists of two successive operations: Illumination-Guided Cross Attention (IGCA) and Gated Wavelet Detail Fusion (GWDF).

We obtain the detail evidence by applying a multi-level Haar decomposition to Y and projecting the three high-frequency subbands at each scale:

$$H_Y^\ell = \phi_H^\ell \left(\left[Y_{\text{LH}}^\ell \| Y_{\text{HL}}^\ell \| Y_{\text{HH}}^\ell \right] \right). \tag{7}$$

Here, Y_{LH}^ℓ, Y_{HL}^ℓ, and Y_{HH}^ℓ denote the high-frequency subbands at level ℓ. These observation-derived subbands retain local variations that may be attenuated during Retinex decomposition.

Let D^ℓ denote the incoming decoder feature. At each decoder scale, an illumination-conditioning feature C_I^ℓ is independently constructed from I at the corresponding resolution. IGWF updates the decoder representation as

$$\begin{aligned} \overline{D}^\ell &= D^\ell + \text{IGCA}_\ell(D^\ell, C_I^\ell), \\ \widetilde{D}^\ell &= \text{GWDF}_\ell(\overline{D}^\ell, H_Y^\ell). \end{aligned} \tag{8}$$

IGCA introduces spatially varying illumination guidance, while GWDF adaptively controls the contribution of observation-derived high-frequency details, which may contain both useful structures and noise.

Together, these role-specific pathways provide complementary evidence to the shared transport predictor $u_\theta(z_t, t, r; \mathcal{C})$. We next optimize this predictor through restoration outcomes from multiple transport states.

3.3 Restoration-Aligned Multi-Path Learning

Conventional flow objectives primarily supervise a velocity field or its temporal relations. In paired image restoration, however, the desired normal-light endpoint X is directly available for each training pair. This allows us to supervise finite-interval transport through the restoration outcome it induces, directly aligning the learned operator with the task objective. We therefore optimize the shared transport operator through three complementary paths covering the deployment source, reference intermediate states, and model-induced rollout states.

Deployment path. We first supervise the complete transport used during one-step inference. Starting from the reflectance source R, the model predicts the normal-light endpoint as

$$\widehat{X}_{\text{dep}} = T_\theta(R, 1, 0; \mathcal{C}), \qquad \mathcal{L}_{\text{dep}} = \rho(\widehat{X}_{\text{dep}}, X), \tag{9}$$

where ρ denotes the restoration criterion defined below. This path directly aligns training with the full-interval transport used at inference.

Reference intermediate path. To extend supervision beyond the complete source-to-target interval, we additionally restore from intermediate states along the paired reference path. For a sampled

intermediate time τ, we construct z_τ by interpolating between R and X and transport it to the restoration endpoint:

$$\begin{aligned} z_\tau &= (1-\tau)X + \tau R, \\ \widehat{X}_{\text{int}} &= T_\theta(z_\tau, \tau, 0; \mathcal{C}), \\ \mathcal{L}_{\text{int}} &= \rho(\widehat{X}_{\text{int}}, X). \end{aligned} \tag{10}$$

Varying τ exposes the shared operator to different states and interval lengths along the reference path, rather than only the complete $R \to X$ transport.

Predicted rollout path. Reference intermediate states are constructed from paired images and do not depend on the learned transport. We therefore further train with intermediate states produced by the model itself through a two-segment rollout. Given a reference state z_t and times $0 < r < t \leq 1$, the model first transports z_t to time r and then completes the restoration:

$$\begin{aligned} \widehat{z}_r &= T_\theta(z_t, t, r; \mathcal{C}), \\ \widehat{X}_{\text{roll}} &= T_\theta(\widehat{z}_r, r, 0; \mathcal{C}), \\ \mathcal{L}_{\text{roll}} &= \rho(\widehat{X}_{\text{roll}}, X). \end{aligned} \tag{11}$$

The two segments share the same predictor, and gradients propagate through $\widehat{z}_r$ so that the first transport segment is optimized through its effect on the final restoration. This encourages compatible transport across interval decompositions without imposing a separate reconstruction target on the predicted intermediate state.

Overall objective. We use a common restoration criterion for all three paths:

$$\rho(\widehat{X}, X) = \lambda_1 \mathcal{L}_1(\widehat{X}, X) + \lambda_2 \mathcal{L}_{\text{SSIM}}(\widehat{X}, X) + \lambda_3 \mathcal{L}_{\text{perc}}(\widehat{X}, X), \tag{12}$$

where the three terms measure pixel-wise, structural, and perceptual discrepancies, respectively. The total training objective is

$$\mathcal{L}_{\text{total}} = \mathbb{E}\left[\mathcal{L}_{\text{dep}} + \beta \mathcal{L}_{\text{int}} + \gamma \mathcal{L}_{\text{roll}}\right], \tag{13}$$

where β and γ control the contributions of the two auxiliary paths. The auxiliary paths are used only during training; inference remains the single full-interval transport in Eq. 4, requiring one evaluation of the predictor.

4 Experiments

4.1 Experimental Settings

Datasets. We evaluate R2RFlow on LOL-v1 and LOL-v2-Real using their standard training and test splits. LOL-v1 contains 485 training and 15 test pairs, while LOL-v2-Real contains 689 training and 100 test pairs. For the main comparison, we train and evaluate a separate model on each dataset. For cross-dataset generalization, the model trained only on LOL-v1 is directly evaluated on LOL-v2-Real, LSRW-Nikon, and LSRW-Huawei (Hai et al., 2023), containing 100, 20, and 30 test pairs, respectively, without fine-tuning.

Evaluation metrics. We report PSNR, SSIM (Wang et al., 2004), and LPIPS, with LPIPS (Zhang et al., 2018) computed using AlexNet features. All metrics are computed from saved PNG outputs under the No GT-mean protocol, without adjusting enhanced-image brightness using the ground-truth mean intensity.

Methods under comparison. We compare with five regression-based methods: RetinexNet (Wei et al., 2018), SNR-Net (Xu et al., 2022), Retinexformer (Cai et al., 2023), HVI-CIDNet (Yan et al., 2025), and Multinex (Brateanu et al., 2026); and three generative methods: LLFlow (Wang et al., 2022), Diff-Retinex (Yi et al., 2023), and GSAD (Hou et al., 2023). For cross-dataset evaluation, we additionally include the unsupervised methods PairLIE (Fu et al., 2023), LightenDiffusion (Jiang et al., 2024), and CLODE (Jung et al., 2025). All supervised methods use their LOL-v1 checkpoints in this setting.

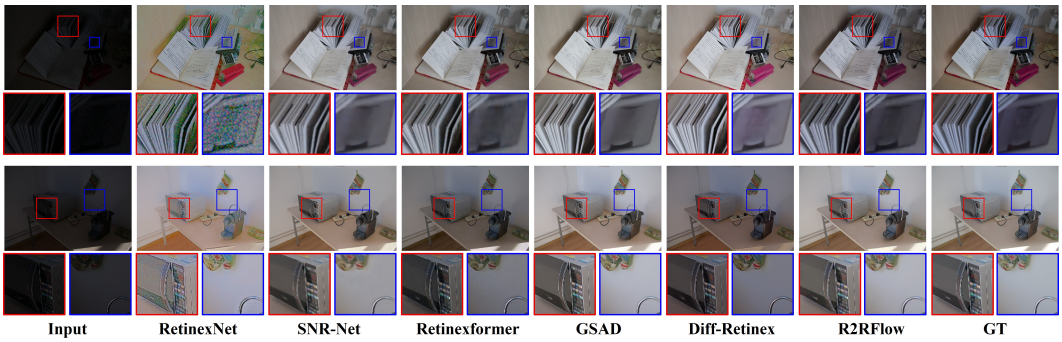


Figure 4: Qualitative comparison on the LOL-v1 and LOL-v2-Real benchmarks (top to bottom). The red and blue boxes indicate representative regions, with their enlarged patches shown below each image. Compared with existing methods, R2RFlow recovers more natural illumination while preserving fine structures and local details, producing results closer to the ground truth.

Table 1: Quantitative comparison on LOL-v1 and LOL-v2-Real. Higher PSNR and SSIM and lower LPIPS indicate better performance. The best and second-best results are highlighted in red and blue, respectively.

Type	Method	LOL-v1			LOL-v2-Real		
		PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
Regression-based	RetinexNet (2018 BMVC)	16.77	0.4191	0.4739	18.37	0.723	0.365
	SNR-Net (2022 CVPR)	24.61	0.8419	0.1510	21.48	0.8489	0.1568
	Retinexformer (2023 ICCV)	25.15	0.8454	0.1314	22.79	0.8397	0.1707
	HVI-CIDNet (2025 CVPR)	23.81	0.8574	0.0856	23.90	0.8656	0.1219
	Multinex (2026 CVPR)	23.19	0.8420	0.1247	22.95	0.8497	0.1788
Generative	LLFlow (2022 AAAI)	21.13	0.8534	0.1185	17.45	0.8316	0.1757
	Diff-Retinex (2023 ICCV)	22.69	0.8550	0.0981	21.04	0.8325	0.1627
	GSAD (2023 NeurIPS)	23.15	0.8532	0.1022	20.18	0.8468	0.1125
	R2RFlow (Ours)	25.26	0.8578	0.1018	24.04	0.8566	0.1214

Implementation details. We optimize R2RFlow using AdamW (Loshchilov & Hutter, 2019) with a batch size of 16, zero weight decay, and a learning rate annealed from 1×10^{-4} to 1×10^{-7} using cosine annealing. We set $(\lambda_1, \lambda_2, \lambda_3) = (1, 0.5, 0.01)$ and $(\beta, \gamma) = (0.5, 0.5)$. On LOL-v1, we train on full images for 800 epochs. On LOL-v2-Real, we train on 400×400 patches for 200 epochs, followed by 1,300 epochs on full images. At inference, we set $(t, r) = (1, 0)$ and restore from the reflectance initialization with one predictor evaluation (1 NFE).

4.2 Quantitative Comparison

Table 1 compares R2RFlow with regression-based and generative methods on LOL-v1 and LOL-v2-Real. On LOL-v1, R2RFlow achieves the highest PSNR and SSIM, with 25.26 dB and 0.8578, respectively. Its PSNR exceeds Retinexformer by 0.11 dB, while its LPIPS remains competitive with the best-performing methods. These results show strong reconstruction and structural fidelity under one-step inference.

On LOL-v2-Real, R2RFlow achieves the highest PSNR of 24.04 dB and ranks second in both SSIM and LPIPS. In particular, it improves PSNR by 0.14 dB over HVI-CIDNet while remaining among the top two methods across all three metrics, demonstrating a favorable balance of reconstruction, structural, and perceptual quality.

4.3 Cross-Dataset Generalization

We assess cross-dataset generalization by directly applying the model trained on LOL-v1 to LSRW-Nikon and LSRW-Huawei without target-domain fine-tuning. Table 2 reports the results, with all

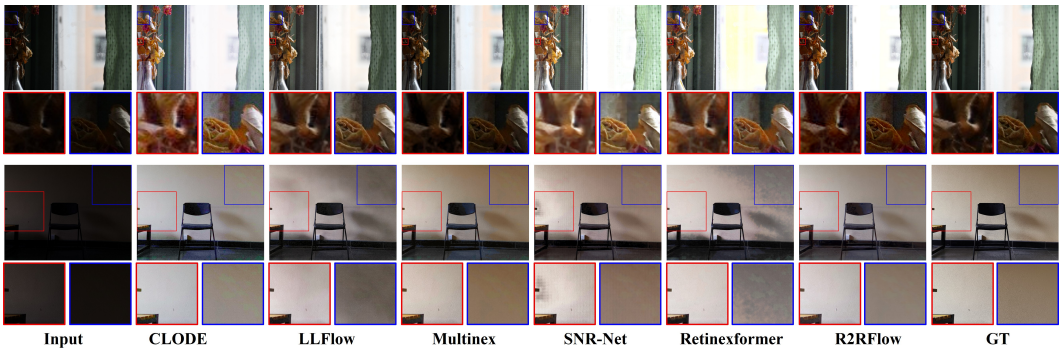


Figure 5: Qualitative comparison of cross-dataset generalization on LSRW-Nikon and LSRW-Huawei. The red and blue boxes indicate representative regions, with their enlarged patches shown below. R2RFlow maintains consistent illumination and preserves structural details across unseen datasets.

Table 2: Cross-dataset generalization on LSRW-Nikon (20 images) and LSRW-Huawei (30 images). Supervised methods use LOL-v1 checkpoints without target-domain fine-tuning. LPIPS is computed using AlexNet. Within supervised methods, the best and second-best results are highlighted in red and blue, respectively. Within unsupervised methods, the best results are highlighted with a light blue background.

Type	Method	LSRW-Nikon			LSRW-Huawei		
		PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
Unsupervised	PairLIE	15.52	0.4347	0.2569	18.99	0.5632	0.3769
	LightenDiffusion	16.43	0.4692	0.2480	19.89	0.5787	0.3673
	CLODE	15.54	0.4871	0.3883	18.43	0.5668	0.4826
Supervised	RetinexNet	13.49	0.2937	0.3754	16.82	0.3953	0.4702
	LLFlow	15.38	0.4519	0.2153	18.97	0.5446	0.3701
	SNR-Net	15.47	0.4603	0.2555	18.56	0.5882	0.4169
	Retinexformer	15.81	0.4493	0.2240	19.09	0.5644	0.3833
	Diff-Retinex	14.89	0.4216	0.2678	16.65	0.5346	0.3899
	GSAD	15.96	0.4414	0.2219	18.14	0.5574	0.3355
	HVI-CIDNet	15.43	0.4536	0.2289	18.89	0.5710	0.3571
	Multinex	15.96	0.4595	0.2142	19.82	0.5895	0.3518
	R2RFlow (Ours)	16.15	0.4728	0.2312	20.08	0.5998	0.3790

supervised methods evaluated using their LOL-v1 checkpoints and unsupervised methods included as additional references.

R2RFlow achieves the highest PSNR and SSIM among supervised methods on both target datasets, reaching 16.15 dB and 0.4728 on LSRW-Nikon and 20.08 dB and 0.5998 on LSRW-Huawei. On LSRW-Huawei, its PSNR and SSIM also exceed those of the evaluated unsupervised methods, while LightenDiffusion and CLODE obtain higher PSNR and SSIM, respectively, on LSRW-Nikon. Perceptual performance varies across methods, with Multinex and GSAD achieving the lowest LPIPS on LSRW-Nikon and LSRW-Huawei, respectively. Overall, R2RFlow shows favorable cross-dataset reconstruction and structural fidelity without target-domain adaptation.

4.4 Ablation Studies

We conduct ablation studies on LOL-v1 to examine the effects of flow initialization, evidence conditioning, and training objectives under one-step inference.

Flow initialization. Table 3 compares different source states while keeping the conditioning pathways and training objective fixed. Observation-derived initializations consistently outperform noise, supporting informed initialization for paired restoration. Among them, reflectance achieves the best performance across all three metrics, validating its use as the source state of R2RFlow.

Table 3: Ablation study on flow initialization. All variants share the same conditioning pathways and training objective.

Initialization	SSIM ↑	LPIPS ↓	
Noise	18.57	0.7102	0.1711
Low-light image Y	23.58	0.8206	0.1560
Illumination I	23.59	0.8356	0.1278
Reflectance R (Ours)	**25.26**	**0.8578**	**0.1018**

Evidence conditioning. Table 4 progressively introduces the role-specific conditioning pathways. Structural conditioning and IGCA provide substantial gains in reconstruction and structural fidelity, while GWDF further incorporates observation-derived frequency details. Adding LWDM yields the complete configuration, which achieves the best performance across all three metrics. Overall, the full model improves PSNR by 2.17 dB and SSIM by 0.0240 while reducing LPIPS by 0.0227 over the variant without complementary conditioning.

Table 4: Ablation study on evidence conditioning modules.

Structure	IGWF		LWDM	PSNR ↑	SSIM ↑	LPIPS ↓
	IGCA	GWDF				
				23.09	0.8338	0.1245
✓				23.67	0.8345	0.1335
✓	✓			24.35	0.8411	0.1139
✓	✓	✓		24.54	0.8433	0.1145
✓	✓	✓	✓	**25.26**	**0.8578**	**0.1018**

Training objectives. Table 5 compares different objectives under the same reflectance initialization and conditioning architecture. Replacing the MeanFlow objective with deployment endpoint supervision leaves PSNR nearly unchanged while substantially improving SSIM and LPIPS, indicating that restoration-level supervision better captures structural and perceptual quality in this setting. Extending the same supervision to reference intermediate and model-induced rollout states further improves all three metrics. These results support restoration-aligned multi-path learning beyond supervising the deployment transport alone.

Table 5: Ablation study on training objectives with fixed reflectance initialization and evidence conditioning.

Training objective	PSNR ↑	SSIM ↑	LPIPS ↓
MeanFlow objective	25.09	0.8238	0.1305
Deployment objective	25.10	0.8512	0.1071
Multi-path objective (Ours)	**25.26**	**0.8578**	**0.1018**

5 Conclusion

We presented R2RFlow, a reflectance-to-restoration flow framework for one-step low-light image enhancement. R2RFlow constructs finite-interval transport from observation-derived reflectance and incorporates complementary structural, illumination, and frequency evidence through role-specific conditioning pathways. A restoration-aligned multi-path learning strategy further optimizes the shared transport operator through restoration outcomes from the deployment source, reference intermediate states, and model-induced rollout states. Experiments on LOL-v1 and LOL-v2-Real demonstrate competitive enhancement performance under one-step inference. The model trained only on LOL-v1 also retains favorable performance on LOL-v2-Real, LSRW-Nikon, and LSRW-Huawei without fine-tuning. Ablation studies further validate the contributions of reflectance initialization, complementary evidence conditioning, and multi-path restoration supervision.

AI use statement

In this work, generative AI tools were used only to assist with language polishing and presentation of the manuscript, including improving grammar, readability, clarity of expression, and organization of author-written text.

References

Alexandru Brateanu, Tingting Mu, Codruta O. Ancuti, and Cosmin Ancuti. Multinex: Lightweight low-light image enhancement via multi-prior retinex. In *Proceedings of the IEEE/CVF Confer-*

ence on Computer Vision and Pattern Recognition, pp. 29887–29896, 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12504–12513, 2023.

Zhenqi Fu, Yan Yang, Xiaotong Tu, Yue Huang, Xinghao Ding, and Kai-Kuang Ma. Learning a simple low-light image enhancer from paired low-light instances. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22252–22261, 2023.

Zhengyang Geng, Mingyang Deng, Xingjian Bai, J. Zico Kolter, and Kaiming He. Mean flows for one-step generative modeling. *arXiv preprint arXiv:2505.13447*, 2025.

Jiang Hai, Zhu Xuan, Ren Yang, Yutong Hao, Feng 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.

Jinhui Hou, Zhiyu Zhu, Junhui Hou, Hui Liu, Huanqiang Zeng, and Hui Yuan. Global structure-aware diffusion process for low-light image enhancement. In *Advances in Neural Information Processing Systems*, 2023.

Hai Jiang, Ao Luo, Xiaohong Liu, Songchen Han, and Shuaicheng Liu. Lightendiffusion: Unsupervised low-light image enhancement with latent-retinex diffusion models. In *European Conference on Computer Vision*, 2024.

Donggoo Jung, Daehyun Kim, and Tae Hyun Kim. Continuous exposure learning for low-light image enhancement using neural odes. In *International Conference on Learning Representations*, 2025.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *International Conference on Machine Learning*, 2023.

Yufei Wang, Renjie Wan, Wenhan Yang, Haoliang Li, Lap-Pui Chau, and Alex C. Kot. Low-light image enhancement with normalizing flow. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004.

Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. In *British Machine Vision Conference*, 2018.

Xiaogang Xu, Ruixing Wang, Chi-Wing Fu, and Jiaya Jia. Snr-aware low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17714–17724, 2022.

Xiaogang Xu, Ruixing Wang, and Jiangbo Lu. Low-light image enhancement via structure modeling and guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5678–5687, 2025.

Wenhan Yang, Shiqi Wang, Yuming Fang, Yue Wang, and Jiaying Liu. Sparse gradient regularized deep retinex network for robust low-light image enhancement. *IEEE Transactions on Image Processing*, 30:2072–2086, 2021.

Xunpeng Yi, Han Xu, Hao Zhang, Linfeng Tang, and Jiayi Ma. Diff-retinex: Rethinking low-light image enhancement with a generative diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12302–12311, 2023.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 586–595, 2018.