Bi-Bridge: Bidirectional Diffusion Bridges for Low-Light Image Enhancement

Supplementary Material

1. DDBM Framework for Low-Light Image Enhancement

This section provides a detailed mathematical formulation of DDBM [18] for LLIE. The concepts here form the technical foundation for our **Bi-Bridge** framework and expand upon the overview in Section 3.

1.1. Reverse Process and Probability Flow ODE

The core of the DDBM generation process is the time-reversal of a guided forward diffusion bridge. As shown in Eq. (2), this reverse process is described by a stochastic differential equation (SDE):

$$\begin{aligned} d\mathbf{x}_t = \big[\mathbf{f}(\mathbf{x}_t, t) - g(t)^2 \left(\mathbf{s}_\theta(\mathbf{x}_t, t, \mathbf{x}_T, T)\right. \\ \left. - \mathbf{h}(\mathbf{x}_t, t, \mathbf{x}_T, T)\right)\big]\, dt + g(t) d\bar{\mathbf{w}}_t, \end{aligned} \tag{1}$$

where $\bar{\mathbf{w}}_t$ is a standard Wiener process when time flows from T to 0.

The corresponding deterministic process for faster sampling is given by the probability flow Ordinary Differential Equation (ODE) [12, 18]:

$$\begin{aligned} d\mathbf{x}_t = [\mathbf{f}(\mathbf{x}_t, t) - g^2(t)(\frac{1}{2}\mathbf{s}_\theta(\mathbf{x}_t, t, \mathbf{x}_T, T) \\ - \mathbf{h}(\mathbf{x}_t, t, \mathbf{x}_T, T))]dt. \end{aligned} \tag{2}$$

Our hybrid sampler, introduced in Section 3.4, leverages both the SDE for introducing beneficial stochasticity and the ODE for stable, efficient integration.

1.2. Score Function and Network Parameterization

Our network D_θ is trained to predict the starting point $\mathbf{x}_0$ from a noised sample $\mathbf{x}_t$ via "pred-x" parameterization [8]. The network's prediction, $\hat{\mathbf{x}}_0 = D_\theta(\mathbf{x}_t, t, \mathbf{x}_T)$, is used to construct score function $\mathbf{s}_\theta$, which in turn approximates the true score of the reverse process:

$$\begin{aligned} &\nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t \mid \mathbf{x}_T) \approx \mathbf{s}_\theta(\mathbf{x}_t, t, \mathbf{x}_T, T) \\ &= -\frac{\mathbf{x}_t - \left(\alpha_t \left(1 - \frac{\mathrm{SNR}_T}{\mathrm{SNR}_t}\right) \hat{\mathbf{x}}_0 + \frac{\alpha_t}{\alpha_T} \frac{\mathrm{SNR}_T}{\mathrm{SNR}_t} \mathbf{x}_T\right)}{\sigma_t^2 \left(1 - \frac{\mathrm{SNR}_T}{\mathrm{SNR}_t}\right)}. \end{aligned} \tag{3}$$

To improve stability, the network D_θ is implemented with a preconditioning scheme [8]:

$$D_\theta(\mathbf{x}_t, t) = c_{\text{skip}}(t)\mathbf{x}_t + c_{\text{out}}(t) F_\theta(c_{\text{in}}(t)\mathbf{x}_t, c_{\text{noise}}(t)), \tag{4}$$

where F_θ is the neural network backbone that predicts x_0. The time-dependent scaling factors are:

$$c_{\text{in}}(t) = \frac{1}{\sqrt{a_t^2\sigma_T^2 + b_t^2\sigma_0^2 + 2a_t b_t \sigma_{0T} + c_t}}, \tag{5}$$

$$c_{\text{out}}(t) = \sqrt{a_t^2\left(\sigma_T^2\sigma_0^2 - \sigma_{0T}^2\right) + \sigma_0^2 c_t}\, c_{\text{in}}(t), \tag{6}$$

$$c_{\text{skip}}(t) = \left(b_t\sigma_0^2 + a_t\sigma_{0T}\right) c_{\text{in}}^2(t), \tag{7}$$

$$c_{\text{noise}}(t) = \frac{1}{4}\log(t), \tag{8}$$

Here, the coefficients a_t, b_t, and c_t are defined as:

$$a_t = \frac{\alpha_t}{\alpha_T}\frac{\mathrm{SNR}_T}{\mathrm{SNR}_t}, \tag{9}$$

$$b_t = \alpha_t\left(1 - \frac{\mathrm{SNR}_T}{\mathrm{SNR}_t}\right), \tag{10}$$

$$c_t = \sigma_t^2\left(1 - \frac{\mathrm{SNR}_T}{\mathrm{SNR}_t}\right), \tag{11}$$

The terms σ_0^2, σ_T^2, and σ_{0T} represent the variance of $\mathbf{x}_0$, the variance of $\mathbf{x}_T$, and their covariance. The loss weighting is $w(t) = 1/c_{\text{out}}(t)^2$.

The DDBM framework can be instantiated with different diffusion processes, most notably Variance Preserving (VP) and Variance Exploding (VE) [12], which define the specific forms of the drift and diffusion coefficients. These instantiations are summarized in Table 1.

2. Bidirectional Inference

Once the unified predictor D_θ is trained, it can be seamlessly deployed for both enhancement and degradation tasks. The inference process involves numerically solving the reverse SDE in Eq. (2) backward in time.

To ensure both high-fidelity results and sampling efficiency, we adopt a higher-order hybrid sampler [8, 18]. This predictor–corrector method combines a stochastic Euler–Maruyama step with a deterministic Heun step. The corrector step provides a more stable and accurate update by numerically integrating the corresponding probability flow ODE in Eq. (2). The core of this process is the dynamic computation of the SDE drift at each step, which combines the analytical h-transform $\mathbf{h}(\cdot)$ with our learned score function $\mathbf{s}_\theta(\cdot)$ derived from the network prediction $\hat{\mathbf{x}}_0 = D_\theta(\mathbf{x}_t, t, \mathbf{x}_T)$.

The specific application of this inference process depends on the choice of the conditional endpoint $\mathbf{x}_T$.

Enhancement. To perform enhancement (X_A to X_B), the process is conditioned on the source low-light image by

	$\mathbf{f}(\mathbf{x}_t, t)$	$g^2(t)$	$p(\mathbf{x}_t \mid \mathbf{x}_0)$	SNR_t	$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_T \mid \mathbf{x}_t)$
VP	$\frac{d \log \alpha_t}{dt} \mathbf{x}_t$	$\frac{d}{dt}\sigma_t^2 - 2\frac{d \log \alpha_t}{dt}\sigma_t^2$	$\mathcal{N}(\alpha_t \mathbf{x}_0, \sigma_t^2 \mathbf{I})$	α_t^2/σ_t^2	$\frac{(\alpha_t/\alpha_T)\mathbf{x}_T - \mathbf{x}_t}{\sigma_t^2(\mathrm{SNR}_t/\mathrm{SNR}_T - 1)}$
VE	$\mathbf{0}$	$\frac{d}{dt}\sigma_t^2$	$\mathcal{N}(\mathbf{x}_0, \sigma_t^2 \mathbf{I})$	$1/\sigma_t^2$	$\frac{\mathbf{x}_T - \mathbf{x}_t}{\sigma_T^2 - \sigma_t^2}$

Table 1. VP and VE bridge parameterizations.

setting the conditional endpoint $\mathbf{x}_T = \mathbf{x}_A$. The sampler is initialized at $\mathbf{x}_{t_N=T} = \mathbf{x}_T$ and solves the SDE backward in time to generate the enhanced normal-light image $\mathbf{x}_B$. This procedure is detailed in Algorithm 1.

Algorithm 1 Enhancement with Bi-Bridge Hybrid Sampler

Require: Model D_θ, low-light image $\mathbf{x}_A$, time steps $\{t_i\}_{i=0}^N$ with $t_N = T$, step ratio s.

Set endpoint: $\mathbf{x}_T \leftarrow \mathbf{x}_A$
Initialize state: $\mathbf{x}_{t_N} \leftarrow \mathbf{x}_A$
for $i = N, \dots, 1$ **do**
  $t \leftarrow t_i, \quad t_{\text{prev}} \leftarrow t_{i-1}, \quad \mathbf{x}_t \leftarrow \mathbf{x}_{t_i}$
// 1. score, h-transform, and drift at current state $(t, \mathbf{x}_t)$
  $s_i \leftarrow \mathbf{s}_\theta(\mathbf{x}_t, t, \mathbf{x}_T, T)$ ▷ Eq. (3)
  $h_i \leftarrow \mathbf{h}(\mathbf{x}_t, t, \mathbf{x}_T, T)$ ▷ Table 1
  $d_i \leftarrow \mathbf{f}(\mathbf{x}_t, t) - g^2(t)\,(s_i - h_i)$
// 2. predictor: stochastic Euler step to intermediate time
  $t_{\text{mid}} \leftarrow t + s\,(t_{\text{prev}} - t)$
  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $i > 1$ else $\mathbf{z} \leftarrow \mathbf{0}$
  $\Delta t_1 \leftarrow t_{\text{mid}} - t$
  $\mathbf{x}_{\text{mid}} \leftarrow \mathbf{x}_t + d_i\,\Delta t_1 + g(t)\sqrt{-\Delta t_1}\,\mathbf{z}$
// 3. corrector (part 1): drift d_{mid} *at* $(t_{mid}, \mathbf{x}_{mid})$
  $s_{\text{mid}} \leftarrow \mathbf{s}_\theta(\mathbf{x}_{\text{mid}}, t_{\text{mid}}, \mathbf{x}_T, T)$
  $h_{\text{mid}} \leftarrow \mathbf{h}(\mathbf{x}_{\text{mid}}, t_{\text{mid}}, \mathbf{x}_T, T)$
  $d_{\text{mid}} \leftarrow \mathbf{f}(\mathbf{x}_{\text{mid}}, t_{\text{mid}}) - g^2(t_{\text{mid}})\,(\frac{1}{2}s_{\text{mid}} - h_{\text{mid}})$
// 4. corrector (part 2): Heun update over $[t_{prev}, t]$
  $\Delta t \leftarrow t_{\text{prev}} - t$
  $\mathbf{x}_{t_{i-1}} \leftarrow \mathbf{x}_t + 0.5\,(d_i + d_{\text{mid}})\,\Delta t$
  if $i > 1$ **then**
// 5. re-correction at $(t_{prev}, \mathbf{x}_{t_{i-1}})$
    $s_{\text{prev}} \leftarrow \mathbf{s}_\theta(\mathbf{x}_{t_{i-1}}, t_{\text{prev}}, \mathbf{x}_T, T)$
    $h_{\text{prev}} \leftarrow \mathbf{h}(\mathbf{x}_{t_{i-1}}, t_{\text{prev}}, \mathbf{x}_T, T)$
    $d_{\text{prev}} \leftarrow \mathbf{f}(\mathbf{x}_{t_{i-1}}, t_{\text{prev}}) - g^2(t_{\text{prev}})\,(\frac{1}{2}s_{\text{prev}} - h_{\text{prev}})$
    $\mathbf{x}_{t_{i-1}} \leftarrow \mathbf{x}_t + 0.5\,(d_i + d_{\text{prev}})\,\Delta t$
  end if
end for
return $\mathbf{x}_{t_0}$ ▷ enhanced normal-light image $\approx \mathbf{x}_B$

Degradation. Conversely, for degradation (X_B to X_A), the process is conditioned on the source normal-light image by setting the conditional endpoint $\mathbf{x}_T = \mathbf{x}_B$. The sampler is initialized at $\mathbf{x}_{t_N=T} = \mathbf{x}_T$ and follows the same numerical procedure to generate the degraded low-light image $\mathbf{x}_A$. This is detailed in Algorithm 2.

Algorithm 2 Degradation with Bi-Bridge Hybrid Sampler

Require: Model D_θ, normal-light image $\mathbf{x}_B$, time steps $\{t_i\}_{i=0}^N$ with $t_N = T$, step ratio s.

Set endpoint: $\mathbf{x}_T \leftarrow \mathbf{x}_B$
Initialize state: $\mathbf{x}_{t_N} \leftarrow \mathbf{x}_B$
The remaining steps are identical to Algorithm 1.
return $\mathbf{x}_{t_0}$ ▷ degraded low-light image $\approx \mathbf{x}_A$

3. Implementation Details

We implement our **Bi-Bridge** framework based on DDBM [18] using the ADM [3] architecture. For 256×256 resolution, we employ a wide residual U-Net with a base channel dimension of 128, and multi-head self-attention at resolutions of 32, 16, and 8. We apply a dropout rate of 0.1 and use concatenation for conditioning at the input level. Regarding the bridge parameterization, we adopt the VP schedule to ensure stable bidirectional learning. Following [18], we utilize a time-invariant drift function $f(x_t, t) = -0.5\beta_0 x_t$ with a symmetric noise schedule ($\beta_1 = \beta_0$), and set the boundary noise variances as $\sigma_0 = \sigma_T = 0.5$ with covariance $\sigma_{0T} = \sigma_0^2/2$.

We train the model using AdamW with a learning rate of 1×10^{-4} and no weight decay. The training is conducted on 4 NVIDIA A800 (80GB) GPUs. To ensure strictly fair comparisons, all evaluated models are trained from scratch without using any pre-trained weights. For paired benchmarks, we adhere to the official training and testing splits; for unpaired benchmarks, we directly apply the model trained on LOL-v2-Syn without any fine-tuning. We set the per-GPU batch size to 8 and employ gradient accumulation every 4 steps, resulting in an effective batch size of 128. To enhance robustness, random horizontal flipping is applied jointly to both source and target images during training. For inference, we adopt the hybrid sampler described in Sec. 2, which combines deterministic ODE steps with stochastic noise injection. Unless otherwise specified, we use 80 NFEs for all reported results.

For quantitative evaluation, we follow the "GT Mean" protocol [9, 14, 16, 17], which aligns the mean luminance of the predicted image with its reference before computing metrics. As discussed in Sec. 1, LLIE is inherently ill-posed and one-to-many, where dataset targets serve as reference exposures rather than absolute ground truth. By factoring out global exposure differences, this protocol en-

courages metrics to better reflect restoration fidelity in terms of structural details and realistic textures, rather than brightness matching alone. All results for both our method and baselines are reported under this same evaluation protocol.

4. More Quantitative Results

4.1. More Results on Paired Datasets

We further evaluate our **Bi-Bridge** on the LSRW [6] and SICE [1] paired datasets. As shown in Table 2, our method consistently outperforms prior SOTA models, demonstrating strong generalization capabilities across diverse lighting conditions.

Table 2. Extended evaluation on additional paired datasets. We report PSNR (↑). Best results are in **bold**.

Dataset	RetinexNet [13]	KinD [16]	Zero-DCE [5]	DarkIR [4]	CIDNet [15]	Ours
LSRW	15.90	16.47	15.83	18.93	19.21	**19.49**
SICE	12.42	12.98	12.45	13.90	13.43	**15.37**

4.2. More Results on Unpaired Datasets

To supplement the NIQE results in the main, Table 3 reports the BRISQUE [10] scores on the five unpaired datasets. **Bi-Bridge** achieves highly competitive performance, further validating its robustness against diverse real-world degradations.

Table 3. Zero-shot generalization performance on unpaired datasets. We report BRISQUE scores (↓). Best results are in **bold**.

Method	DICM	LIME	MEF	NPE	VV
Retinexformer [2]	16.34	21.77	13.78	15.58	22.99
CIDNet [15]	21.47	16.25	13.77	18.92	30.63
ReDDiT [9]	18.13	**12.00**	9.93	**10.42**	18.36
DarkIR [4]	18.69	21.62	13.90	12.88	26.87
Ours	**16.28**	14.52	**9.58**	12.01	**17.72**

5. More Visual Comparisons

To provide a more comprehensive assessment of perceptual quality, we present additional visual comparisons on the benchmark datasets.

5.1. Additional Results on LOL-v1

Figure 1 presents a challenging low-light scenario from the LOL-v1 dataset, featuring a swimming pool environment with a digital scoreboard. As detailed in the zoomed-in regions (red boxes), accurately restoring the legibility of the timestamp while maintaining the correct color temperature is particularly challenging. It can be observed that competing methods such as CIDNet [15] and ReDDiT [9] introduce noticeable color shifts, rendering the scene with unnatural reddish or greenish casts. Meanwhile, Retinexformer [2] and SNR-Aware [14] struggle to reconstruct the sharp edges of the digital digits. In contrast, our **Bi-Bridge** effectively suppresses noise and restores high-fidelity textual details with natural colors, achieving the visual quality closest to the GT.

5.2. Additional Results on LOL-v2-Real

We further evaluate our method on the LOL-v2-Real dataset using representative scenes shown in Figure 2. In the top row, recovering details from the extremely dark region (highlighted in the red box) proves difficult for baselines like Retinexformer [2] and CIDNet [15], which suffer from severe noise artifacts and unnatural greenish casts, whereas our method successfully restores the distinct wall structures. Similarly, in the storefront scene (bottom row), while competitors tend to blur the strokes of the characters and produce washed-out colors, our **Bi-Bridge** preserves sharp font edges and vibrant red hues, demonstrating superior fidelity in both structural and semantic detail recovery.

5.3. Additional Results on LOL-v2-Synthetic

Figure 3 presents results on LOL-v2-Synthetic. The outdoor statue example (top row) corresponds to a challenging backlit scenario. Many SOTA methods (e.g., Retinexformer [2] and CIDNet [15]) under-enhance the illumination, leaving the sculpture significantly darker than the reference and with muted colors. Our **Bi-Bridge** produces a brighter, more natural rendering with richer local details.

In the indoor sculpture scene (bottom row), faithfully preserving the warm ambient lighting is crucial. GSAD [7] produces an almost desaturated gray appearance, while ReDDiT [9] tends to generate hazy, low-contrast outputs. **Bi-Bridge**, in comparison, retains both the depth cues of the relief and the warm color tone of the original scene, leading to a closer match to the reference.

6. Limitations and Future Work

Limitations. The main limitation of **Bi-Bridge** lies in its inference efficiency. As an iterative diffusion-based model, it remains substantially slower than regression-based LLIE methods, especially at higher resolutions. To address this, we plan to explore acceleration techniques, including knowledge distillation and advanced sampling algorithms, to bridge the speed gap with regression-based methods.

Future Work. Beyond acceleration, our model's ability to synthesize paired low-light/normal-light data opens up opportunities to leverage large-scale generative priors, such as Stable Diffusion [11], for LLIE. We plan to investigate how such synthetic datasets can facilitate the transfer of foundation models to complex low-light degradations and broader real-world scenarios. Furthermore, while our symmetry-consistent training has proven empirically effective, conducting a deeper probabilistic analysis to theoretically formalize its regularization properties remains a compelling avenue for future research.

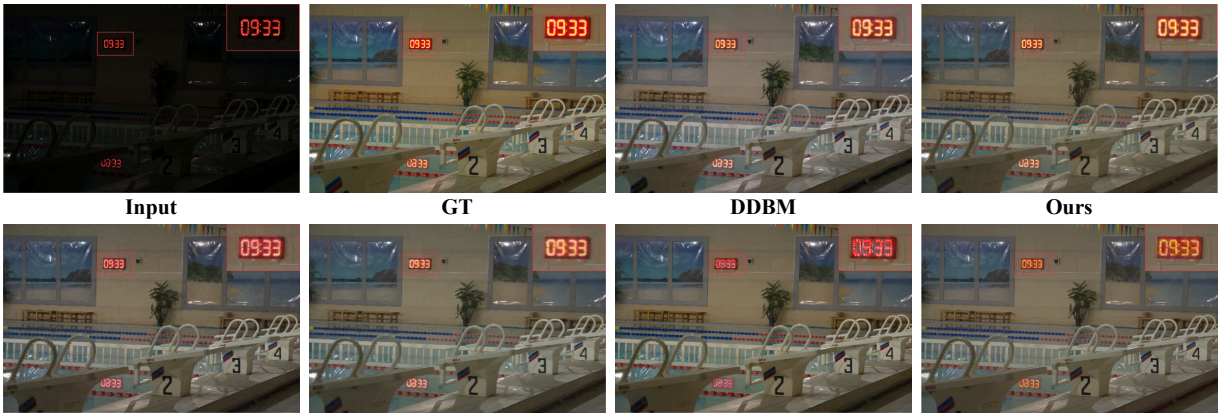

Figure 1. **Additional qualitative comparison on the LOL-v1 dataset.**

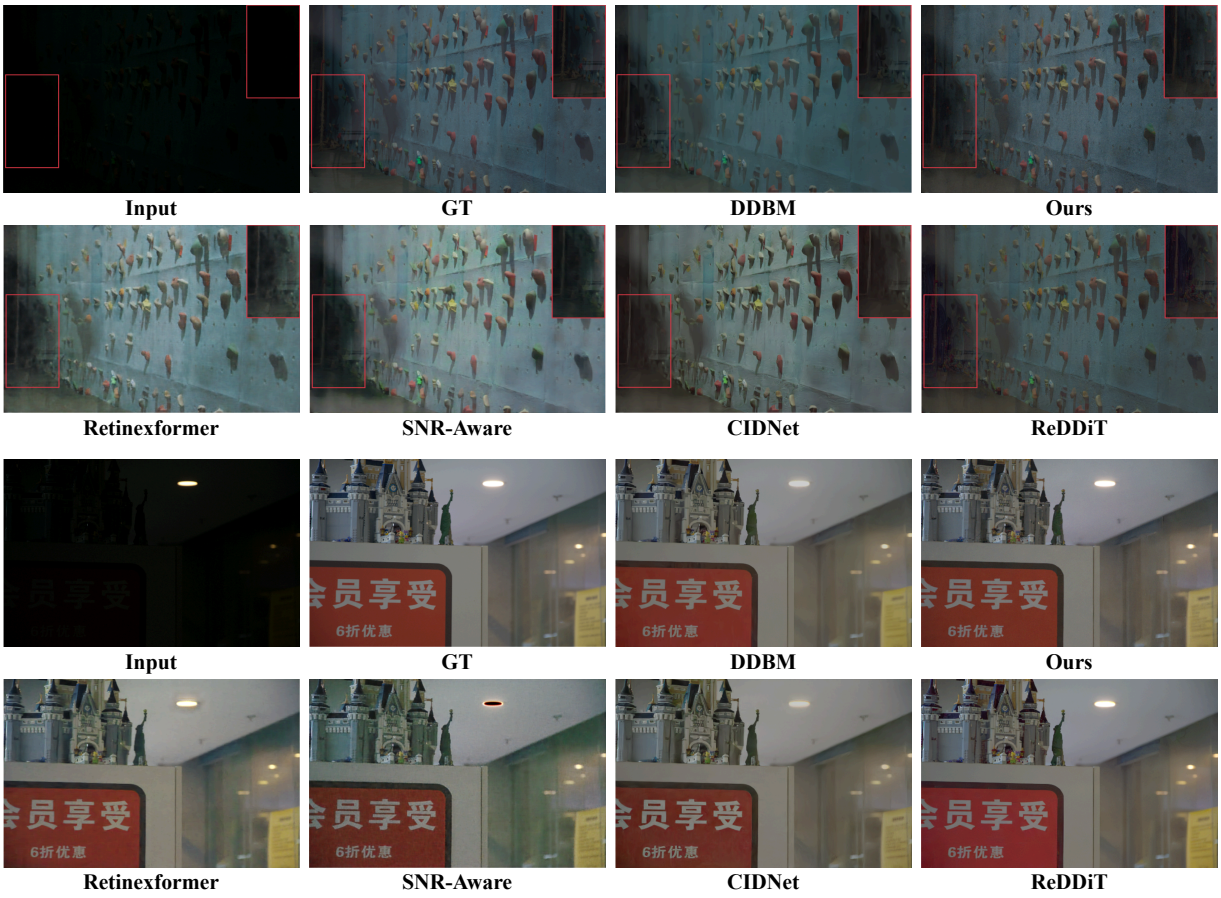

Figure 2. **Additional qualitative comparison on the LOL-v2-Real dataset.**

Figure 3. **Additional qualitative comparison on the LOL-v2-synthetic dataset.**

References

[1] Jianrui Cai, Shuhang Gu, and Lei Zhang. Learning a deep single image contrast enhancer from multi-exposure images. *IEEE transactions on image processing*, 27(4):2049–2062, 2018. 3

[2] 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 *International Conference on Computer Vision (ICCV)*, pages 12504–12513, 2023. 3

[3] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 8780–8794, 2021. 2

[4] Daniel Feijoo, Juan C Benito, Alvaro Garcia, and Marcos V Conde. Darkir: Robust low-light image restoration. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 10879–10889, 2025. 3

[5] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1780–1789, 2020. 3

[6] Jiang Hai, Zhu Xuan, Ren Yang, Yutong Hao, Fengzhu Zou, Fang Lin, and Songchen Han. R2rnet: Low-light image enhancement via real-low to real-normal network. *Journal of Visual Communication and Image Representation*, 90: 103712, 2023. 3

[7] 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 (NeurIPS)*, pages 79734–79747, 2023. 3

[8] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 26565–26577, 2022. 1

[9] Guanzhou Lan, Qianli Ma, Yuqi Yang, Zhigang Wang, Dong Wang, Xuelong Li, and Bin Zhao. Efficient diffusion as low light enhancer. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21277–21286, 2025. 2, 3

[10] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing*, 21(12): 4695–4708, 2012. 3

[11] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. 3

[12] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 1

[13] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. *arXiv preprint arXiv:1808.04560*, 2018. 3

[14] Xiaogang Xu, Ruixing Wang, Chi-Wing Fu, and Jiaya Jia. Snr-aware low-light image enhancement. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17714–17724, 2022. 2, 3

[15] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5678–5687, 2025. 3

[16] Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. Kindling the darkness: A practical low-light image enhancer. In *ACM International Conference on Multimedia (ACM MM)*, pages 1632–1640, 2019. 2, 3

[17] Dewei Zhou, Zongxin Yang, and Yi Yang. Pyramid diffusion models for low-light image enhancement. *arXiv preprint arXiv:2305.10028*, 2023. 2

[18] Linqi Zhou, Aaron Lou, Samar Khanna, and Stefano Ermon. Denoising diffusion bridge models. In *International Conference on Learning Representations (ICLR)*, 2024. 1, 2