

HDR Reconstruction Boosting with Training-Free and Exposure-Consistent Diffusion

Yo-Tin Lin¹ Su-Kai Chen² Hou-Ning Hu² Yen-Yu Lin¹ Yu-Lun Liu¹

¹National Yang Ming Chiao Tung University ²MediaTek Inc.

<https://github.com/EusdenLin/HDR-Reconstruction-Boosting>

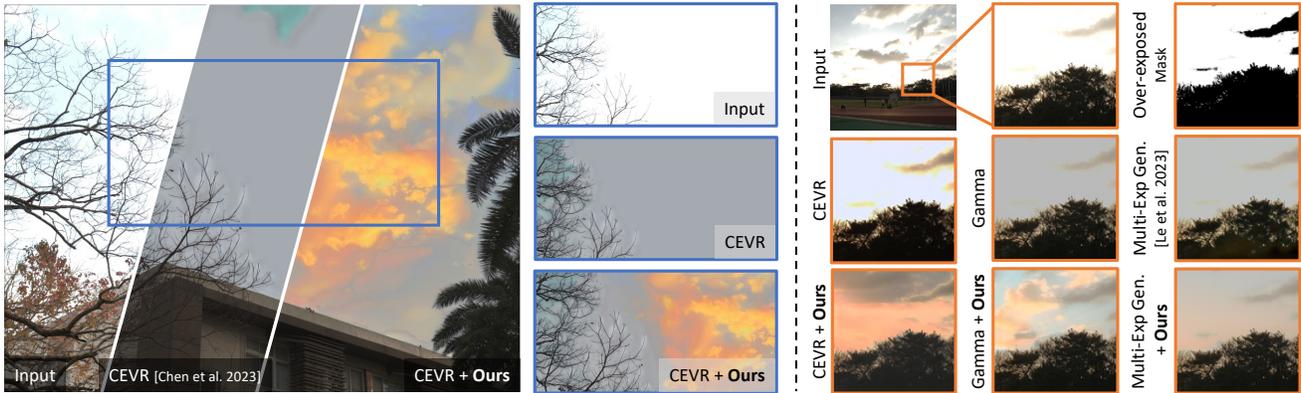


Figure 1. **Our training-free, diffusion-based approach enhances existing HDR reconstruction methods for over-exposed regions.** (Left) Our method improves CEVR [7] by generating natural sky details in the over-exposed area. (Right) Comparisons with different baseline methods (CEVR, SingleHDR [25], and Multi-Exposure Generation [19]) demonstrate how our approach consistently improves reconstruction quality across different methods without requiring additional training. Here are the results of the *real-world* cases captured by a Fujifilm X-T30 DSLR camera.

Abstract

Single LDR to HDR reconstruction remains challenging for over-exposed regions where traditional methods often fail due to complete information loss. We present a training-free approach that enhances existing indirect and direct HDR reconstruction methods through diffusion-based inpainting. Our method combines text-guided diffusion models with SDEdit refinement to generate plausible content in over-exposed areas while maintaining consistency across multi-exposure LDR images. Unlike previous approaches requiring extensive training, our method seamlessly integrates with existing HDR reconstruction techniques through an iterative compensation mechanism that ensures luminance coherence across multiple exposures. We demonstrate significant improvements in both perceptual quality and quantitative metrics on standard HDR datasets and in-the-wild captures. Results show that our method effectively recovers natural details in challenging scenarios while preserving the advantages of existing HDR reconstruction pipelines.

Project page: [here](#).

1. Introduction

High Dynamic Range (HDR) imaging aims to capture and reproduce the wide spectrum of luminance levels present in real-world scenes. This encompasses both extremely bright and very dark regions, which typically exceed the dynamic range of conventional imaging sensors and standard displays. To address this, HDR reconstruction methods have been developed to recover information lost during the signal processing pipeline of Low Dynamic Range (LDR) images. However, as illustrated in Fig. 1, while existing methods such as CEVR [7] excel at reconstructing correctly exposed regions, they often struggle with over-exposed areas due to the limited information in these regions.

Over-exposed regions pose a significant challenge in enhancing image quality and realism. These regions, characterized by the complete loss of texture and detail, undermine the visual coherence and fidelity of reconstructed images. Conventional methods, including those based on tone

mapping curves or learning-based approaches, are often inadequate as evidenced by the gamma correction results in Fig. 1. These regions typically appear as uniform patches devoid of texture, making accurate reconstruction an ill-posed problem. Even advanced GAN-based methods, such as [21, 50], despite their promise, suffer from generalization issues and inconsistent reconstruction quality, particularly in severe over-exposed scenarios. Moreover, as illustrated in Fig. 3, direct application of existing diffusion inpainting models to independently inpaint different exposures often leads to artifacts such as ghosting and color inconsistencies due to misalignment across exposure levels and inadequate control of inter-texture consistency.

In this work, we propose a generative method that leverages diffusion priors to address the limitations of existing HDR reconstruction techniques. Our method employs an image inpainting paradigm to hallucinate plausible content in over-exposed regions, prioritizing perceptual quality and natural transitions over strict adherence to ground truth. As demonstrated in Fig. 1 (Right), our approach can significantly enhance existing methods like CEVR and Multi-Exposure Generation by generating natural and consistent details in over-exposed areas.

To achieve this, we incorporate HDR reconstruction through multi-exposure image synthesis. By aligning and compensating for discrepancies across low dynamic range (LDR) images with different exposure values (EVs), we ensure consistency and luminance accuracy in the generated HDR image. Our pipeline includes three key components: (1) *a diffusion-based inpainting backbone* to ensure high-quality generation tailored to over-exposed regions, (2) *iterative refinement with SDEdit* [29] to maintain luminance consistency and structural coherence and to gradually enhance the generated content, and (3) *compensation after every inpainting step* to make the generated content exceed the expected luminance lower bounds, avoiding unrealistic brightness or structural artifacts.

In summary, the key contributions of this work include:

- Introducing a diffusion-based inpainting pipeline that effectively hallucinates over-exposed regions, significantly enhancing visual quality and aesthetic appeal.
- Proposing a robust iterative compensation and inpainting strategy using SDEdit to align luminance and maintain texture consistency across LDR in the brackets.
- Showing the versatility of our approach, which seamlessly integrates with existing indirect and direct HDR reconstruction methods without additional training.

2. Related Work

HDR reconstruction. Single-image HDR methods are either direct or indirect. Direct methods, including HDR-CNN [11], ExpandNet [28], CNN models [16, 20, 54], masked features [43], and pipeline reversal [25], predict

HDR directly but struggle with ill-posed over-exposed regions. Indirect methods merge generated multi-exposure LDR images via inverse response estimation [8, 30, 41]. While methods like CEVR [7], Deep Chain [12, 21], and Recursive HDRI [22] excel in proper exposures, they often fail in saturated areas. Recent approaches explore cycle-consistency [23], exposure control [14], differentiable learning [17], generative models [58], GANs [50], and decomposition [10]. Our method improves the indirect approach using diffusion priors [9, 13] to reconstruct over-exposed content while retaining exposure bracketing benefits.

Over-exposed region handling for HDR. Reconstructing over-exposed regions remains challenging due to information loss [11, 25, 43, 58]. While GAN-based methods like GlowGAN [50] (using StyleGAN-XL [44]) exist, they often lack generalization or require scene-specific training [22, 36, 58]. In contrast, diffusion approaches like Exposure Diffusion [5], DiffusionLight [37], and LEDiff [51] generalize well across scenes [9, 34, 42]. Similarly, LightsOut [48] employs diffusion to reconstruct lost light sources. Leveraging diffusion’s synthesis [1, 26] and control capabilities [13, 45], our method achieves diverse generation without scene-specific training.

Diffusion & image inpainting. Recent diffusion-based inpainting methods, such as Blended Latent Diffusion [2], use CLIP guidance [39] and noise blending for smooth transitions [35, 42]. Several studies explore controlled generation [9, 13], such as RePaint [26], ControlNet [57], SDEdit [29], and IP-Adapter [56]. Paint-by-Example [55] shows effective CLIP-guided inpainting, while others explore personalization [47] and specialized architectures [38, 49] to enhance inpainting capabilities. To improve generation fidelity, CorrFill [24] introduces correspondence guidance for reference-based inpainting. Furthermore, recent works have expanded diffusion priors to traditional regression tasks, such as matting [53]. Unlike methods requiring scene-specific training [36, 50, 58] or fine-tuning [37], our training-free approach achieves high-quality results by controlling pre-trained diffusion models to maintain cross-exposure consistency.

3. Method

Given an LDR image containing over-exposed regions, our objective is to reconstruct an HDR image, with particular emphasis on restoring details within the over-exposed areas. Using existing indirect HDR reconstruction methods (Section 3.2), we can generate predicted LDR images with varying exposure value (EV) offsets (e.g., EV: -1, -2, -3) and obtain the estimated inverse camera response function (CRF) using Debevec’s method [8]. After applying inpainting techniques to these images, we can create low-EV im-

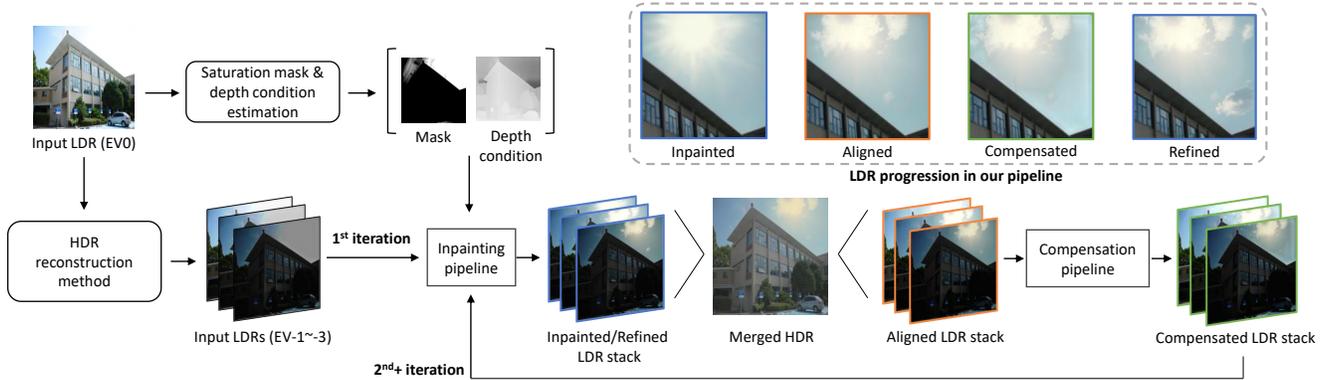


Figure 2. **Overview of our training-free HDR reconstruction pipeline.** Given an input LDR image (EV0), we generate bracketed LDR images using an existing HDR reconstruction method. Our iterative pipeline then enhances these results through (1) an inpainting stage guided by exposure and condition maps, (2) HDR merging and alignment of the generated content, and (3) a compensation stage to ensure physical consistency. The top-right inset shows the progressive refinement of the over-exposed regions through our pipeline stages. The upper sky region in the aligned case shows the clear alignment across EV images. The lower sky region in the compensation and refined cases demonstrates our algorithm’s ability to hallucinate realistic over-expose regions with plausible intensity and texture. The upper part of the sky region in the aligned case shows the clear alignment across EV images. The lower parts of the sky region in the compensation and refined cases show the ability of our algorithm to hallucinate realistic over-expose regions with reasonable intensity and texture.

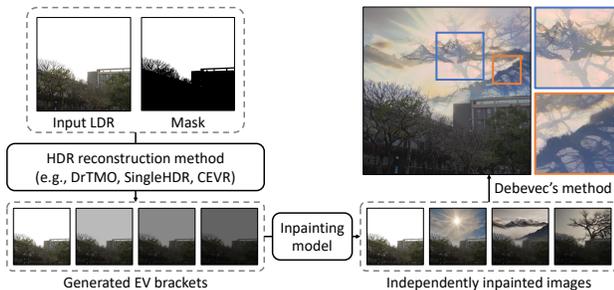


Figure 3. **Limitations arising from naively combining indirect HDR reconstruction methods and over-exposed regions inpainting.** Independent inpainting of each EV bracket, without cross-EV alignment, can introduce ghosting artifacts in the merged HDR result. These artifacts stem from inconsistencies in the generated content, becoming apparent after merging the independently inpainted exposures using Debevec’s method.

ages containing reconstructed over-exposed details, which are utilized to synthesize the final HDR image.

However, as shown in Fig. 4, if the generated details violate the lower-bound constraint, *i.e.*, the minimum luminance level in over-exposed regions, the inverse CRF predicted by reconstruction methods can be disrupted. This disruption leads to color shifts and other artifacts, causing the reconstructed HDR image to deviate from the original LDR input. Furthermore, while diffusion models can produce consistent textures with varying brightness levels under a fixed random seed, the iterative nature of inpainting may result in a loss of cross-EV consistency. This leads to chaotic results in the final HDR image, as shown in Fig. 3.

To address the challenges of lower-bound adherence and cross-EV alignment, we propose a novel inpainting

framework comprising two key components: the Inpainting Pipeline (Section 3.3) and the Compensation Pipeline (Section 3.4). An overview of our algorithm is shown in Fig. 2. Our method employs an iterative inpainting algorithm with a single LDR image as input. At each iteration, we refine the saturated areas using a corresponding text prompt, ensuring that the inpainted pixels satisfy the lower-bound constraint and remain consistent across different EV levels.

3.1. Preliminaries

ControlNet [57]. It provides a mechanism to condition the generation process on additional structural guidance, such as edges or depth maps. A standard latent diffusion model (LDM) operates through a denoising process: $z_{t-1} = f_{\theta}(z_t, t, y)$, where z_t is the noisy latent variable at timestep t , and y is the text prompt. The function f_{θ} represents the U-Net denoiser parameterized by θ . ControlNet introduces an external conditioning signal c , modifying the generation process to $z_{t-1} = f_{\theta}(z_t, t, y, c)$, where c ensures that the generated image preserves structural consistency. Using ControlNet prevents the model from generating unrealistic textures in over-exposed regions, making the generated content align with the physical scene structure.

Inpainting Models. To combine ControlNet with the inpainting diffusion process [1], we use a conditioning mechanism to ensure that the inpainted region adheres to both structural constraints and textual guidelines. The modified denoising function is

$$z_{t-1} = f_{\theta}(z_t, t, y, c) \odot (1 - M) + \hat{X}_t \odot M, \quad (1)$$

where M is the saturation mask, \hat{X}_t is the target image at timestep t in the forward process, c is the ControlNet guidance (depth map), and \odot is the element-wise multiplication.

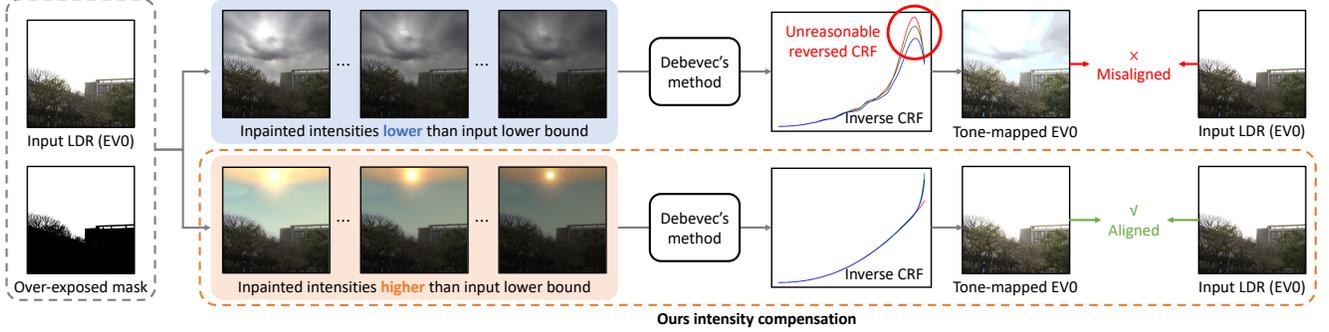


Figure 4. **Motivation for luminance compensation in HDR reconstruction.** The top row shows the existence of inpainted intensities below the lower bound of over-exposed regions may lead to unreasonable inverse CRF estimation and misaligned results. The bottom row demonstrates our intensity compensation approach, which ensures bounded intensities, resulting in proper inverse CRF estimation and exposure-aligned reconstruction that matches the input LDR.

3.2. LDR Preprocess

Given an LDR image, our objective at this stage is to estimate its appearance under different exposure values, which is essential to CRF estimation. To this end, we utilize an existing indirect or direct HDR reconstruction method to generate LDR images across varying exposure values (EVs).

For cases where a direct HDR reconstruction method, such as [25], is preferred, we can generate the baseline LDR stacks using the following gamma function:

$$X_{EV} = (H \times 2^{EV})^\gamma, \quad (2)$$

where X_{EV} represents the EV-adjusted image of the normalized HDR input H at the exposure level 0 (EV0). We set γ to 2.2 in our implementation and apply this gamma function to all direct methods in the experiments.

Using the generated LDR stacks, we estimate the inverse CRF via Debevec’s method [8]. This estimated function serves as the reference for both tone-mapping and HDR merging, and is also used within our compensation pipeline (Section 3.4).

Additionally, we identify saturated regions in the image by using the soft mask formula [11]:

$$m_i = \frac{\max(0, \max_c(X_{i,c} - \tau))}{255 - \tau}, \quad (3)$$

where X is the input image with pixel index i and channel index c , and $X_{i,c} \in [0, 255]$. In our implementation, we set the threshold τ to 245. To mitigate edge side effects caused by iterative inpainting, we apply this soft mask, marking saturated areas for the first iteration, to reduce edge artifacts.

3.3. Inpainting Pipeline

A primary challenge during inpainting is that blended diffusion inpainting models [2] often disregard the pixel values generated by baseline methods in over-exposed regions. They tend to produce visually appealing but brightness-inconsistent results.

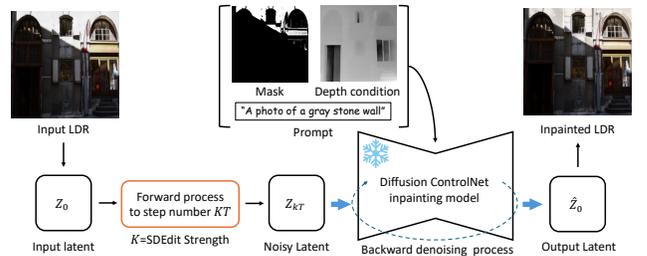


Figure 5. **Inpainting pipeline in our method.** The depth-conditioned inpainting pipeline with scheduled strength is able to generate reasonable and consistent content across different EVs and iterations. The scheduled SDEdit strength enables us to balance detail preservation with creative generation. In the early iterations, we allow the model to explore a broader range of plausible scene details, while in later iterations, we retain the refined details from previous steps and only update pixels that do not meet the physical constraints.

To address this issue, we employ SDEdit [29]. Unlike the standard forward process, which adds noise up to the full timestep T , we stop at a partial timestep t , where $t < T$. Inspired by the observation that brightness information is primarily determined in early timesteps, we set t to the range $[0.85T, 0.95T]$ in our implementation. Furthermore, we find that smaller values of t typically result in closer alignment with the results in the previous iteration. Therefore, we develop a scheduler that starts at $0.95T$ and then gradually reduces the value of t throughout the iterative inpainting process. This mechanism helps balance the diversity of generated results while maintaining texture consistency across iterations.

As depicted in Fig. 5, we use the standard diffusion inpainting model as the backbone model. The proposed inpainting pipeline takes a prompt, the soft mask, and the depth map as the extra guidance, and can generate more plausible LDR images for HDR reconstruction.

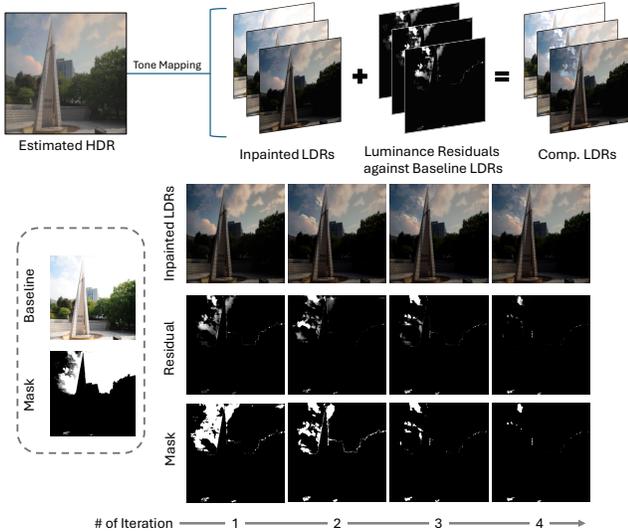


Figure 6. **Compensation pipeline in our method.** Our pipeline ensures proper exposure relationships through iterative refinement. (Top) Overview of how inpainted LDRs are combined with luminance residuals to produce compensated results after tone-mapping. (Bottom) Visualization over four iterations shows inpainted results, decreasing residuals, and shrinking mask regions. The process maintains proper luminance lower bounds while focusing refinement on problematic areas through selective masking, preventing CRF estimation issues.

3.4. Compensation Pipeline

As shown in Fig. 4, inpainting with intensities below the input’s lower bound leads to inaccurate inverse CRF estimation and misaligned results. When the generated content fails to maintain consistent luminance relationships across different exposures, it disrupts the inverse CRF predicted by reconstruction methods and causes artifacts in the final output. To address this issue, we introduce a compensation pipeline that ensures cross-EV consistency and proper intensity lower bounds.

To maintain cross-EV consistency among LDR images throughout the iterative process, we utilize the inverse CRF estimated by baseline methods: The LDR stacks are first merged into an HDR image [4], which is then converted back into LDR stacks with matching EV levels, using the same CRF.

After obtaining the aligned LDR stacks, we ensure that the generated regions meet the desired lower-bound constraints. To achieve this, we transform both the generated image and the baseline image into the YUV space to extract luminance. By calculating the residual between the baseline and the generated image and multiplying it by a compensation scale, we derive the values to adjust the generated image. These adjusted values are passed to the next iteration for refinement.

To prevent abrupt brightness increases that could cause

Algorithm 1 Inpainting Pipeline

Input: Baseline LDR images $X = \{X_0, X_{-1}, X_{-2}, X_{-3}\}$, text prompt p , total step number T , condition image c , SDEdit strength K , compensation scale S , saturation mask m , total inpainting iterations N , estimated inverse CRF $inv\text{-}CRF$, exposure times $times$

Output: HDR image H

$\hat{X} \leftarrow X$

for $i = 1, 2, \dots, N$ **do**

for $EV = -1, -2, -3$ **do**

$Z_{EV,t} = \text{noise}(\hat{X}_{EV,t})$

for all t from $T \times K$ to l **do**

$Z_{EV,t,fg} = \text{noise}(\hat{X}_{EV,t-1})$

$Z_{EV,t,bg} = \text{denoise}(Z_{EV,t}, c, p)$

$Z_{EV,t-1} = Z_{EV,t,fg} \odot m + Z_{EV,t,bg} \odot (1 - m)$

end for

$\hat{X}_{EV,\text{inpaint}} = Z_{EV,0}$

end for

$H = \text{merge}(\hat{X}_{\text{inpaint}}, inv\text{-}CRF, times)$

$\hat{X}_{\text{align}} = \text{tonemap}(H, inv\text{-}CRF, times)$

$\text{residual} = X - \hat{X}_{\text{align}}$

$\hat{X}_{\text{comp}} = \hat{X}_{\text{align}} + \text{residual} \times m \times S_i$

$m \leftarrow m \times (\text{residual} > 0)$

$\hat{X} \leftarrow \hat{X}_{\text{comp}}$

end for

return H

artifacts and disrupt smooth transitions between adjacent regions, the compensation scale is initially set to a low value and gradually increases with each iteration.

To ensure that valid generated regions remain unchanged in subsequent iterations, as shown in Fig. 6, we create a mask during the residual calculation. Valid pixels are assigned a mask value of 0, while invalid pixels are set to 1. This shrinks the inpainting region, allowing the model to focus specifically on areas that do not meet the standard. As illustrated in the figure, the compensation pipeline effectively reduces the residual, ultimately ensuring that all pixels meet the expected criteria. Alg. 1 summarizes the process from the LDR stacks to the final HDR results.

4. Experiments

4.1. Experimental Setup

Implementation Details. We use SDXL-base1.0 with an off-the-shelf depth-conditioned ControlNet [57] as the diffusion inpainting backbone model. In this work, we use a fixed and generic set of parameters and input prompts for evaluating all the cases. During inference, we use 50 sampling steps, a 5.0 guide scale, and a 0.5 ControlNet conditioning scale. Additionally, we set the saturation threshold to 245. For SDEdit strength scheduling, we use a linear decay schedule, starting at 0.95 and decreasing by 0.05 per iteration, *i.e.*, [0.95, 0.90, 0.85, 0.80, ...]. The compensation

Table 1. **Quantitative comparisons.** Our algorithm consistently improves state-of-the-art methods on both VDS [21] and HDR-Eye [33] datasets without additional training.

Dataset	Method	Non reference											Ref.-based (vs. GT)		
		BRISQUE ↓		NIQE ↓		NIMA ↑		MUSIQ ↑		CLIP-IQA ↑		PU21 ↓	KID ↓		HDR- ↑
		TMO	RH's	KK's	PIQE	RH's	KK's								
VDS	SingleHDR [25]	75.73	73.08	14.49	14.25	4.021	4.352	2.032	0.295	2.063	0.318	90.16	183.9	205.1	7.5610
	SingleHDR + Ours	64.82	60.34	12.17	11.17	4.056	4.419	2.151	0.301	2.159	0.324	78.35	186.2	202.8	8.8604
	GlowGAN [50]	79.02	77.49	14.54	14.25	4.008	4.209	1.890	0.342	1.900	0.346	93.79	256.7	275.4	6.0815
	GlowGAN + Ours	69.19	66.79	11.82	11.47	4.021	4.248	2.020	0.363	2.027	0.394	78.41	231.7	258.0	6.8441
	Deep Recursive HDRI [22]	68.29	66.45	14.30	13.40	4.207	4.594	2.161	0.307	2.199	0.305	88.27	111.7	119.0	9.3884
	Deep Recursive HDRI + Ours	64.24	64.24	13.96	12.75	4.289	4.673	2.250	0.320	2.309	0.393	85.51	101.1	110.3	9.0358
	Multi-Exp. Gen. [19]	71.51	65.40	13.89	12.71	3.900	4.410	2.013	0.209	1.964	0.262	66.12	228.3	275.8	7.8906
	Multi-Exp. Gen. + Ours	62.47	56.12	11.08	10.01	3.984	4.445	2.154	0.216	2.117	0.421	63.40	300.2	319.0	8.0233
	CEVR [7]	72.42	70.21	14.33	13.48	4.044	4.421	2.043	0.299	2.093	0.299	90.44	123.4	135.3	9.2704
	CEVR + Ours	71.17	68.90	13.71	12.97	4.100	4.456	2.072	0.322	2.123	0.369	86.12	126.8	138.0	9.0358
HDR-Eye	SingleHDR [25]	73.27	71.19	14.29	14.18	4.033	4.375	1.736	1.820	0.250	0.235	91.06	221.6	227.6	7.1277
	SingleHDR + Ours	62.50	57.72	12.02	11.12	4.087	4.446	1.875	1.955	0.251	0.235	78.39	213.6	227.0	6.8385
	GlowGAN [50]	70.07	70.51	14.14	13.92	4.211	4.346	1.873	1.886	0.259	0.259	89.48	226.5	245.2	6.4705
	GlowGAN + Ours	61.00	61.67	11.98	11.75	4.277	4.460	1.972	2.011	0.260	0.260	79.29	181.3	225.1	6.4286
	Deep Recursive HDRI [22]	65.61	64.98	13.21	11.70	4.225	4.620	1.900	1.933	0.185	0.176	87.68	199.1	200.5	7.2791
	Deep Recursive HDRI + Ours	65.04	62.95	12.21	10.94	4.260	4.674	1.950	2.019	0.190	0.180	84.55	204.8	192.6	7.0809
	Multi-Exp. Gen. [19]	66.71	61.07	13.38	11.53	4.000	4.539	1.791	1.859	0.158	0.158	67.98	229.0	228.9	6.1145
	Multi-Exp. Gen. + Ours	62.74	56.14	11.06	9.80	4.113	4.559	2.075	2.106	0.186	0.199	67.44	338.8	302.0	5.8216
	CEVR [7]	69.31	68.04	14.69	13.57	4.082	4.447	1.792	1.823	0.197	0.174	89.64	202.8	201.2	7.2253
	CEVR + Ours	68.25	66.16	12.46	11.29	4.100	4.478	1.792	1.851	0.205	0.183	83.80	219.9	208.6	7.0871

scale has a linear increasing schedule, starting at 0.2 and increasing by 0.1 per iteration, *i.e.*, [0.2, 0.3, 0.4, 0.5, ...].

To avoid a single specially chosen prompt from disproportionately boosting performance, we use a fixed, generic text prompt “A beautiful photo with bright background. High-resolution image with a lot of details and sharpness. 4K, Ultra Quality. Good photo.” to minimize the sensitivity of generative models to prompt variations.

Baseline Methods. We use four indirect HDR reconstruction methods: One is the simple gamma 2.2 tone mapping curve, and the others are existing indirect HDR reconstruction methods, including CEVR [7], Multi-Exposure Generation [19], and Deep Recursive HDRI [21]. We also use two direct methods, including GlowGAN [50] and single-HDR [25], for comprehensive evaluations.

Datasets. We conduct the experiments on two public HDR datasets, including the VDS [21] and HDR-Eye [33] datasets. Both of them cover outdoor and indoor scenes.

Evaluation Metrics. To comprehensively evaluate the performance of our method, we use the full-reference metrics HDR-VDP-3 [27] and KID score [6], which assess the fidelity of generation. In addition, we use the non-reference metrics BRISQUE [31], NIQE [32], NIMA [46], MUSIQ [15], CLIP-IQA [52], and PU21-PIQE [3] for the overall image quality evaluation.

Tone Mapping Operators. The majority of image quality metrics are not compatible with HDR file inputs. Therefore, we perform tone mapping to convert the HDR results

to LDR prior to comparison, ensuring that they share the same dynamic range.

For this process, we employ two different tone-mapping operators (TMOs): Reinhard’s (RH’s) method [40] and Kim and Kautz’s (KK’s) method [18]. These operators compress the full HDR dynamic range into a more limited range while preserving details without clipping.

4.2. Results

Quantitative Comparisons. Our method demonstrates notably superior performance in non-reference metrics and consistently outperforms the baseline methods across both VDS and HDR-Eye datasets, as reported in Tab. 1. While we observe slightly lower scores in reference-based metrics like KID and HDR-VDP-3, this is expected since our method focuses on plausible reconstruction of over-exposed regions rather than strict adherence to the ground truth. Standard reference-based metrics inherently penalize creative reconstruction of lost information, even when the results are perceptually compelling. These results substantiate the effectiveness of our inpainting pipeline in reconstructing challenging areas where traditional methods struggle.

Qualitative Comparisons. With our diffusion inpainting pipeline, we can generate natural and diverse over-exposed scenes for the HDR images. Qualitative results, as displayed in Fig. 7, illustrate the differences between the baseline methods and the boosted baseline results one by one. In addition, we compare our results to those by GlowGAN. Our method leverages the diffusion prior and can produce more feasible HDR scenes than GlowGAN. We also explore

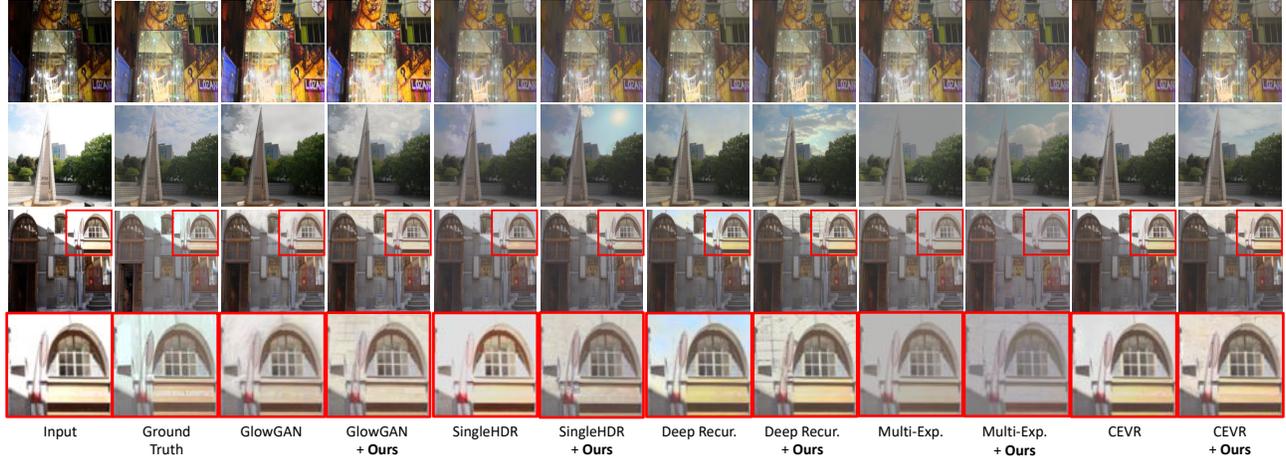


Figure 7. **Qualitative comparisons of HDR reconstruction methods.** Our approach enhances various baseline methods (indicated by “+Ours”) across diverse scenes, including outdoor architecture, indoor lighting, and urban landscapes. While existing methods often produce unrealistic content or artifacts in over-exposed regions, our method generates more natural details and lighting patterns while maintaining structural consistency without training. Results show improved reconstruction of sky details, cloud patterns, and lighting effects across different baseline methods. The top row shows the main difference between the baseline results and hallucinated results. Though we can generate a similar texture and color tone, the generated content and ground truth might not be considered the same by the KID score and HDR-VDP-3.



Figure 8. **Inpainting diversity results.** By leveraging the diffusion prior, our inpainting pipeline gains the ability to generate diverse scenes with different lighting conditions and sky appearances from the same input LDR image.

Table 2. **Ablation study on diffusion backbones.** Our algorithm improves state-of-the-art methods on the VDS dataset [21] with different diffusion backbones without additional training.

Method	NIQE ↓		NIMA ↑		PU21-PIQE ↓
	TMO	RH's	KK's	RH's	
CEVR [7]	14.33	13.48	4.044	4.421	90.44
CEVR + Ours (SDXL)	13.71	12.97	4.100	4.456	86.12
CEVR + Ours (SD turbo)	12.36	11.60	4.150	4.482	78.46

the impact of different text prompts. As visualized in Fig. 8, our method exhibits adaptability to all the text prompts and hallucinates plausible results. More examples of in-the-wild or extreme cases can be found in supplementary materials.

4.3. Ablation Studies

SDEdit Strength Scheduling. In the inpainting pipeline, we use varying SDEdit strengths to control the balance

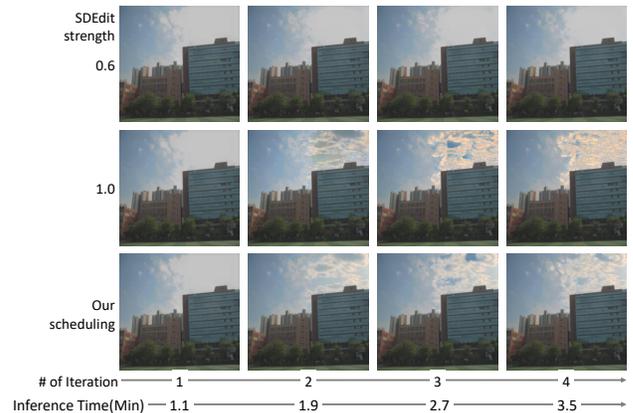


Figure 9. **SDEdit strength scheduling ablation study.** Low strength, *e.g.*, 0.6, limits content generation in over-exposed regions, while high strength, *e.g.*, 1.0, causes inconsistency between iterations. Our scheduling gradually decreases noise strength to balance content generation with detail refinement and consistency across all the four iterations.

of creativity and consistency of texture between iterations. Fig. 9 shows the results by using different SDEdit strengths. When the strength is fixed to a lower value, *e.g.*, 0.6, the results of the next iteration would lose the generative diversity and be the same as those in the previous iteration. In contrast, a fixed high strength of 1.0 prevents the inpainting model from taking the results of the previous iteration into consideration, resulting in the production of infeasible images.

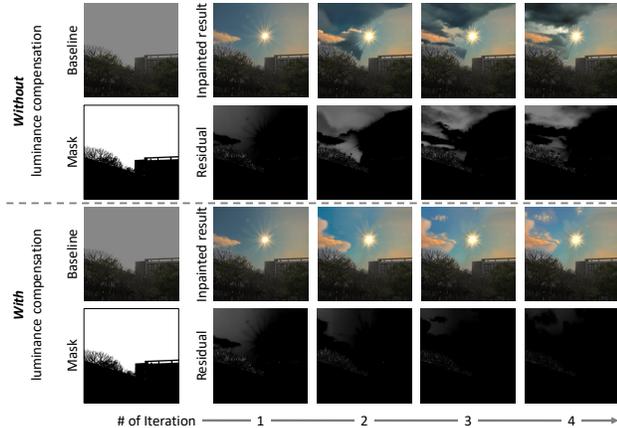


Figure 10. **Luminance compensation ablation study.** (Top) Without compensation, the inpainted content may be unrealistically dark. (Bottom) Our compensation adjusts pixel values to maintain physically plausible brightness levels, as shown by decreasing residuals across iterations.

Table 3. **Quantitative ablation studies.** We conducted a quantitative evaluation on the VDS dataset and demonstrated that our approach yields improvements in non-reference image quality metrics.

Components (RH’s TMO)	NIQE ↓	NIMA ↑	PU21-PIQE ↓
Baseline (CEVR)	14.33	4.044	90.44
w/o Merge & Comp.	13.57	4.045	85.15
w/o Merge	13.60	4.021	84.68
w/o Comp.	13.39	4.028	85.07
Ours	13.71	4.100	86.12

Compensation Block. When generating details for over-exposed regions, we aim to prevent the generation process from modifying the original input LDR in a way that leads to an unrealistic inverse CRF. To ensure this, we constrain the baseline methods by setting their pixel values to the lower-bound of the overexposed areas. If inpainted pixels fall below this lower-bound, we correct them by adding the absolute difference between the inpainted value and the lower-bound, scaled by a compensation factor. Additionally, we refine the mask to better align with the adjusted pixel values. Fig. 10 shows the results with and without the compensation block. As demonstrated in the second row, we find that the residual fails to diminish without the compensation block as the iteration progresses. In the fourth row, it is evident that the residual is significantly reduced with the aid of this compensation block. The quantitative evaluation of each component of our approach is reported in Tab. 3.

Diffusion Backbone. The training-free nature of our method ensures that the inpainting pipeline exhibits consistent performance across various diffusion backbone models. To demonstrate this, we conduct experiments using two distinct diffusion models with depth-conditioned Con-

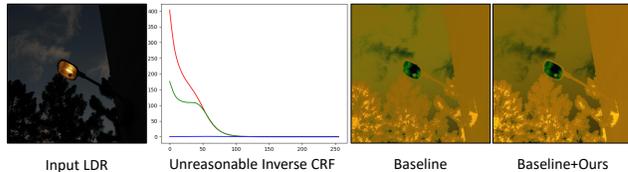


Figure 11. **Limitation.** Our method is restricted by the estimated inverse CRF. When the estimated inverse CRF is unreasonable, our results suffer from color shifting as well.

trolNet, including SDXL, the original model employed in our method, and SDXL Turbo, a distilled variant of SDXL characterized by an accelerated inference speed but diminished image quality and prompt alignment. As presented in Tab. 2, our method with SDXL turbo still beats the baseline method in all non-reference-based metrics. Notably, it even surpasses the performance with SDXL in certain metrics.

Inference Time. We conduct the evaluations with different numbers of iterations on the VDS dataset, using the output image size of 1024×1024 . Fig. 9 shows the computation time with different numbers of iterations. We conducted this analysis with single RTX 3090.

4.4. Limitation

Although our method effectively boosts the baseline methods by iteratively inpainting the LDR stacks, its performance is still contingent upon the baseline method producing a reasonable inverse CRF. As shown in Fig. 11, the baseline method yields a non-monotonic inverse CRF. The non-monotonic inverse CRF disrupts the assumed one-to-one relationship between pixel intensity and scene radiance. Thus, both the baseline and the inpainted results suffer from color-shifting side effects.

5. Conclusion

We present a training-free, diffusion-based pipeline for HDR restoration of over-exposed regions. Our method leverages diffusion priors with SDEdit refinement to generate plausible content while preserving exposure consistency. Through iterative refinement and compensation mechanisms, our method consistently and substantially enhances existing HDR reconstruction methods, particularly in over-exposed areas. Comprehensive results demonstrate significant improvements across metrics and datasets, highlighting the effectiveness of our approach across diverse scenes without the need for additional training. As a promising direction for future work, we plan to extend our framework to handle under-exposed regions, aiming for a unified solution that addresses both extremes of the dynamic range.

Acknowledgments. This work was supported in part by National Science and Technology Council (NSTC) under grants 111-2628-E-A49-025-MY3, 112-2221-E-A49-090-MY3, 111-2634-F-002-023, 111-2634-F-006-012, 110-2221-E-A49-065-MY3, 111-2634-F-A49-010, 112-2222-E-A49-004-MY2, and 113-2628-E-A49-023-. This work was funded in part by MediaTek. Yu-Lun Liu acknowledges the Yushan Young Fellow Program by the MOE in Taiwan.

References

- [1] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18208–18218, 2022. 2, 3
- [2] Omri Avrahami, Ohad Fried, and Dani Lischinski. Blended latent diffusion. *ACM Transactions on Graphics (TOG)*, 42(4):1–11, 2023. 2, 4
- [3] Maryam Azimi et al. Pu21: A novel perceptually uniform encoding for adapting existing quality metrics for hdr. In *2021 Picture Coding Symposium (PCS)*, pages 1–5. IEEE, 2021. 6
- [4] Francesco Banterle, Alessandro Artusi, Kurt Debattista, and Alan Chalmers. *Advanced High Dynamic Range Imaging (2nd Edition)*. AK Peters (CRC Press), Natick, MA, USA, 2017. 5
- [5] Mojtaba Bemana, Thomas Leimkühler, Karol Myszkowski, Hans-Peter Seidel, and Tobias Ritschel. Exposure diffusion: Hdr image generation by consistent ldr denoising. *arXiv preprint arXiv:2405.14304*, 2024. 2
- [6] Mikołaj Bińkowski, Dougal J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD GANs. In *International Conference on Learning Representations*, 2018. 6
- [7] Su-Kai Chen, Hung-Lin Yen, Yu-Lun Liu, Min-Hung Chen, Hou-Ning Hu, Wen-Hsiao Peng, and Yen-Yu Lin. Learning continuous exposure value representations for single-image hdr reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12990–13000, 2023. 1, 2, 6, 7
- [8] Paul E Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 643–652. 2023. 2, 4
- [9] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. 2
- [10] Sebastian Dille, Chris Careaga, and Yağız Aksoy. Intrinsic single-image hdr reconstruction. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part XI*, page 161–177, Berlin, Heidelberg, 2024. Springer-Verlag. 2
- [11] Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafał K Mantiuk, and Jonas Unger. Hdr image reconstruction from a single exposure using deep cnns. *ACM Transactions on Graphics (TOG)*, 36(6):1–15, 2017. 2, 4
- [12] Yuki Endo, Yoshihiro Kanamori, and Jun Mitani. Deep reverse tone mapping. *ACM Trans. Graph.*, 36(6):1–10, 2017. 2
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2
- [14] So Yeon Jo, Siyeong Lee, Namhyun Ahn, and Suk-Ju Kang. Deep arbitrary hdri: Inverse tone mapping with controllable exposure changes. *IEEE Transactions on Multimedia*, 24:2713–2726, 2021. 2
- [15] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5128–5137, 2021. 6
- [16] Zeeshan Khan, Mukul Khanna, and Shanmuganathan Ramam. Fhdr: Hdr image reconstruction from a single ldr image using feedback network. In *2019 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 1–5, 2019. 2
- [17] Junghee Kim, Siyeong Lee, and Suk-Ju Kang. End-to-end differentiable learning to hdr image synthesis for multi-exposure images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1780–1788, 2021. 2
- [18] Min H. Kim and Jan Kautz. Consistent tone reproduction. In *Proceedings of the Tenth IASTED International Conference on Computer Graphics and Imaging*, page 152–159, USA, 2008. ACTA Press. 6
- [19] Phuoc-Hieu Le, Quynh Le, Rang Nguyen, and Binh-Son Hua. Single-image hdr reconstruction by multi-exposure generation. In *Proceedings of the IEEE/CVF winter Conference on Applications of Computer Vision*, pages 4063–4072, 2023. 1, 6
- [20] Min Jung Lee, Chi-hyoung Rhee, and Chang Ha Lee. Hsvnet: Reconstructing hdr image from a single exposure ldr image with cnn. *Applied Sciences*, 12(5), 2022. 2
- [21] Siyeong Lee, Gwon Hwan An, and Suk-Ju Kang. Deep recursive hdri: Inverse tone mapping using generative adversarial networks. In *proceedings of the European Conference on Computer Vision (ECCV)*, pages 596–611, 2018. 2, 6, 7
- [22] Siyeong Lee, Gwon Hwan An, and Suk-Ju Kang. Deep recursive hdri: Inverse tone mapping using generative adversarial networks. In *proceedings of the European Conference on Computer Vision (ECCV)*, pages 596–611, 2018. 2, 6
- [23] Siyeong Lee, So Yeon Jo, Gwon Hwan An, and Suk-Ju Kang. Learning to generate multi-exposure stacks with cycle consistency for high dynamic range imaging. *IEEE Transactions on Multimedia*, 23:2561–2574, 2020. 2
- [24] Kuan-Hung Liu, Cheng-Kun Yang, Min-Hung Chen, Yu-Lun Liu, and Yen-Yu Lin. Corrfill: Enhancing faithfulness in reference-based inpainting with correspondence guidance in diffusion models. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1618–1627. IEEE, 2025. 2
- [25] Yu-Lun Liu, Wei-Sheng Lai, Yu-Sheng Chen, Yi-Lung Kao, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Single-image hdr reconstruction by learning to reverse the camera pipeline. In *Proceedings of the IEEE/CVF Confer-*

- ence on *Computer Vision and Pattern Recognition*, pages 1651–1660, 2020. 1, 2, 4, 6
- [26] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022. 2
- [27] Rafal K. Mantiuk, Dounia Hammou, and Param Hanji. Hdrvdp-3: A multi-metric for predicting image differences, quality and contrast distortions in high dynamic range and regular content. 2023. 6
- [28] Demetris Marnerides, Thomas Bashford-Rogers, Jonathan Hatchett, and Kurt Debattista. Expandnet: A deep convolutional neural network for high dynamic range expansion from low dynamic range content. In *Computer Graphics Forum*, pages 37–49. Wiley Online Library, 2018. 2
- [29] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 2, 4
- [30] Tomoo Mitsunaga and Shree K Nayar. Radiometric self calibration. In *Proceedings. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)*, pages 374–380. IEEE, 1999. 2
- [31] 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. 6
- [32] Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2013. 6
- [33] Hiromi Nemoto, Pavel Korshunov, Philippe Hanhart, and Touradj Ebrahimi. Visual attention in ldr and hdr images. In *9th International Workshop on Video Processing and Quality Metrics for Consumer Electronics (VPQM)*, 2015. 6
- [34] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 2
- [35] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. 2
- [36] Yuzhen Niu, Jianbin Wu, Wenxi Liu, Wenzhong Guo, and Rynson WH Lau. Hdr-gan: Hdr image reconstruction from multi-exposed ldr images with large motions. *IEEE Transactions on Image Processing*, 30:3885–3896, 2021. 2
- [37] Pakkapon Phongthawee, Worameth Chinchuthakun, Nontaphat Sinsunthithet, Varun Jampani, Amit Raj, Pramook Khungurn, and Supasorn Suwajanakorn. Diffusionlight: Light probes for free by painting a chrome ball. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 98–108, 2024. 2
- [38] Konpat Preechakul, Nattanat Chatthee, Suttisak Wizadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10619–10629, 2022. 2
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2
- [40] Erik Reinhard, Michael Stark, Peter Shirley, and James Ferwerda. Photographic tone reproduction for digital images. *ACM Trans. Graph.*, 21(3):267–276, 2002. 6
- [41] Mark A Robertson, Sean Borman, and Robert L Stevenson. Estimation-theoretic approach to dynamic range enhancement using multiple exposures. *Journal of Electronic Imaging*, 12(2):219–228, 2003. 2
- [42] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 2
- [43] Marcel Santana Santos, Tsang Ing Ren, and Nima Khademi Kalantari. Single image hdr reconstruction using a cnn with masked features and perceptual loss. *arXiv Preprint arXiv:2005.07335*, 2020. 2
- [44] Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022. 2
- [45] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2
- [46] Hossein Talebi and Peyman Milanfar. Nima: Neural image assessment. *IEEE Transactions on Image Processing*, 27(8):3998–4011, 2018. 6
- [47] Luming Tang, Nataniel Ruiz, Qinghao Chu, Yuanzhen Li, Aleksander Holynski, David E Jacobs, Bharath Hariharan, Yael Pritch, Neal Wadhwa, Kfir Aberman, et al. Realfill: Reference-driven generation for authentic image completion. *ACM Transactions on Graphics (TOG)*, 43(4):1–12, 2024. 2
- [48] Shr-Ruei Tsai, Wei-Cheng Chang, Jie-Ying Lee, Chih-Hai Su, and Yu-Lun Liu. Lightsout: Diffusion-based outpainting for enhanced lens flare removal. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6353–6363, 2025. 2
- [49] Bram Wallace, Akash Gokul, and Nikhil Naik. Edict: Exact diffusion inversion via coupled transformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22532–22541, 2023. 2
- [50] Chao Wang, Ana Serrano, Xingang Pan, Bin Chen, Karol Myszkowski, Hans-Peter Seidel, Christian Theobalt, and Thomas Leimkühler. Glowgan: Unsupervised learning of hdr images from ldr images in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10509–10519, 2023. 2, 6
- [51] Chao Wang, Zhihao Xia, Thomas Leimkuehler, Karol Myszkowski, and Xuaner Zhang. Lediff: Latent ex-

- posure diffusion for hdr generation. *arXiv preprint arXiv:2412.14456*, 2024. [2](#)
- [52] Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *AAAI*, 2023. [6](#)
- [53] Zhixiang Wang, Baiang Li, Jian Wang, Yu-Lun Liu, Jinwei Gu, Yung-Yu Chuang, and Shin’Ichi Satoh. Matting by generation. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. [2](#)
- [54] Guotao Wu, Ran Song, Mingxin Zhang, Xiaolei Li, and Paul L. Rosin. Litmnet: A deep cnn for efficient hdr image reconstruction from a single ldr image. *Pattern Recogn.*, 127(C), 2022. [2](#)
- [55] Binxin Yang, Shuyang Gu, Bo Zhang, Ting Zhang, Xuejin Chen, Xiaoyan Sun, Dong Chen, and Fang Wen. Paint by example: Exemplar-based image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18381–18391, 2023. [2](#)
- [56] Hu Ye, Jun Zhang, Sibao Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023. [2](#)
- [57] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. [2](#), [3](#), [5](#)
- [58] Yang Zhang and TO Aydın. Deep hdr estimation with generative detail reconstruction. In *Computer Graphics Forum*, pages 179–190. Wiley Online Library, 2021. [2](#)