X2HDR: HDR Image Generation in a Perceptually Uniform Space

Ronghuan Wu[1], **Wanchao Su**[2], **Kede Ma**[1], **Jing Liao**[1], **and Rafał K. Mantiuk**[3]
[1]City University of Hong Kong
[2]Monash University
[3]University of Cambridge
rh.wu@my.cityu.edu.hk, wanchao.su@monash.edu, kede.ma@cityu.edu.hk,
jingliao@cityu.edu.hk, rafal.mantiuk@cl.cam.ac.uk

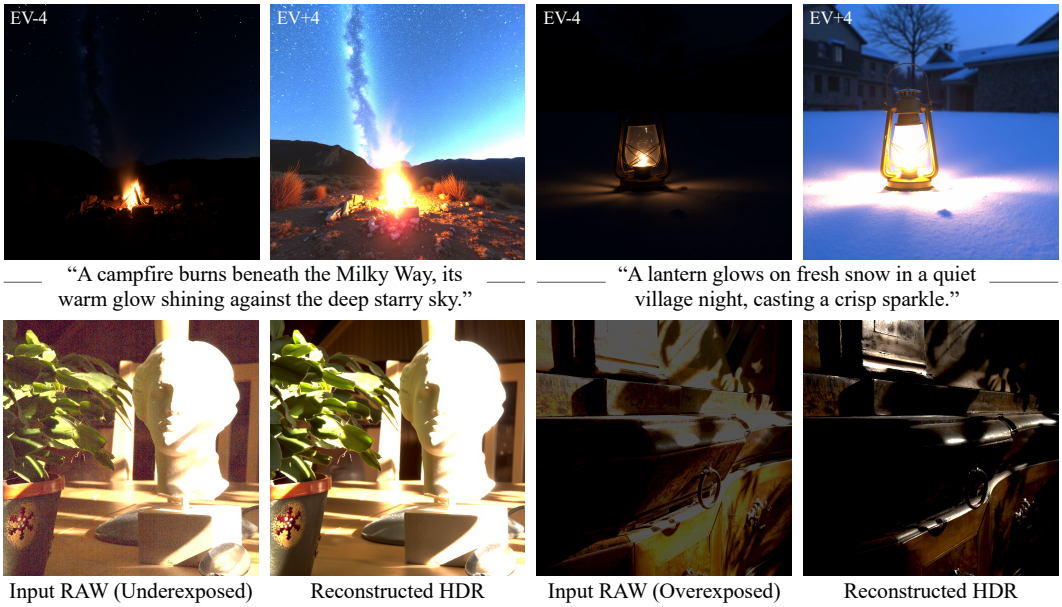

Figure 1: Qualitative results demonstrating the two modes supported by the proposed X2HDR. Top: text-to-HDR generation, visualized with exposure-adjusted views at EV -4 and EV $+4$, highlighting synthesized content across highlights and shadows. Bottom: single-image RAW-to-HDR reconstruction. From underexposed and overexposed RAW inputs, X2HDR reconstructs HDR structures by inpainting saturated regions and denoising in low-light areas.

Abstract

High-dynamic-range (HDR) formats and displays are becoming increasingly prevalent, yet state-of-the-art image generators (*e.g.*, Stable Diffusion and FLUX) typically remain limited to low-dynamic-range (LDR) output due to the lack of large-scale HDR training data. In this work, we show that existing pretrained diffusion models can be easily adapted to HDR generation without retraining from scratch. A key challenge is that HDR images are natively represented in linear RGB, whose intensity and color statistics differ substantially from those of sRGB-encoded LDR images. This gap, however, can be effectively bridged by converting HDR inputs into perceptually uniform encodings (*e.g.*, using PU21 or PQ). Empirically, we find that LDR-pretrained variational autoencoders (VAEs) reconstruct PU21-encoded HDR inputs with fidelity comparable to LDR data, whereas linear RGB inputs cause severe degradations. Motivated by this finding, we describe an efficient adap-

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

tation strategy that freezes the VAE and finetunes only the denoiser via low-rank adaptation in a perceptually uniform space. This results in a unified computational method that supports both text-to-HDR synthesis and single-image RAW-to-HDR reconstruction. Experiments demonstrate that our perceptually encoded adaptation consistently improves perceptual fidelity, text-image alignment, and effective dynamic range, relative to previous techniques. Complete HDR results and code are available at https://X2HDR.github.io/.

1 Introduction

High-dynamic-range (HDR) imagery provides a more faithful representation of natural scene luminance and color than conventional low-dynamic-range (LDR) formats, offering a more realistic and appealing visual appearance. Yet, creating HDR content in practice is still inconvenient. The standard solution—multi-exposure bracketing—often suffers from motion-induced misalignment and ghosting, and it increases both capture and processing cost. Specialized HDR sensors can mitigate these issues, but remain expensive and are not widely available.

In parallel, recent advances in text-to-image (T2I) diffusion models have made visual content creation accessible from simple inputs such as text prompts and coarse visual guidance (*e.g.*, scribbles, drafts, or low-resolution images). Extending this convenience to the HDR domain is highly desirable: users should be able to create HDR content directly from text or from a low-quality camera input, without relying on bracketing software or specialized hardware.

Several recent methods attempt to adapt pretrained T2I models to support HDR by emulating the classic HDR pipeline: generate multi-exposure "brackets" and then merge them into an HDR output [12]. While effective in principle, this strategy introduces substantial architectural and algorithmic complexity. For example, LEDiff [61] trains separate denoisers for highlight hallucination and shadow recovery, and additionally finetunes the variational autoencoder (VAE) decoder [29] to output linear HDR values. Bracket Diffusion [6] jointly denoises multi-exposure latents, incurring inference time and memory costs that scale with the number of brackets. These design choices complicate deployment and hinder adoption of newer, more memory-intensive backbone architectures.

We argue that much of this complexity stems from a simpler root cause: a representational mismatch between LDR and HDR imagery. Most contemporary diffusion models are trained on billions of display-encoded, nonlinearly compressed LDR images, whereas HDR and RAW data are natively expressed in a linear-light space prior to the image signal processor (ISP), resulting in dramatically different pixel-intensity statistics. In particular, human visual sensitivity to luminance changes is much higher in shadows than in highlights [40], making the linear HDR distribution poorly aligned with the perceptually shaped LDR data.

Building on this insight, we present X2HDR, a unified computational method for HDR image synthesis and reconstruction with minimal changes to existing T2I models. The key to X2HDR is to operate in a perceptually uniform space (*e.g.*, induced by PU21 [37] or PQ [45]). Specifically, we first map HDR data into this perceptual space, which compresses extreme highlights and reallocates precision toward shadows, in a way that reshapes HDR luminance statistics to better match LDR pretraining data. Empirically, this simple form of preprocessing empowers LDR-pretrained VAEs to recover PU21-encoded HDR inputs with fidelity comparable to LDR reconstructions, whereas linear HDR leads to distorted latents and severe artifacts. Leveraging this, X2HDR freezes the VAE and acquires HDR capability by finetuning only the denoiser through parameter-efficient low-rank adaptation (LoRA) [26].

We instantiate X2HDR on two representative tasks: 1) text-to-HDR generation, synthesizing HDR images directly from text, and 2) single-image RAW-to-HDR reconstruction, in which the strong generative prior of the pretrained T2I model facilitates plausible hallucination of missing content in overexposed regions while suppressing noise in underexposed areas. Qualitative and quantitative evaluations, together with a formal perceptual study (see Supplemental Section F.8) on a calibrated HDR display, show that the proposed X2HDR consistently improves perceptual fidelity, text-image alignment, and effective dynamic range over prior approaches (see Figure 1).

In summary, our main contributions are

- A simple yet effective practice for HDR adaptation—encoding HDR/RAW inputs in a perceptually uniform space—which avoids complex bracket-and-merge pipelines and enables pretrained T2I models to support HDR with minimal modification;
- A unified computational method for text-to-HDR generation and RAW-to-HDR reconstruction with substantial perceptual improvements over previous methods.

2 Perceptually Uniform HDR Representation

In this section, we present X2HDR's PU21-based HDR representation [37], supporting faithful reconstruction with an LDR-pretrained VAE.

2.1 HDR Image Representation

HDR (and RAW) images are typically stored in a linear color space, where pixel values are proportional to the underlying scene/sensor light signal up to a scale factor. However, linear representation is not perceptually uniform: an equal luminance increment is far more noticeable at low levels than at high levels. Perceptually uniform encodings address this mismatch by redistributing encoded values according to human visual sensitivity, including SMPTE PQ [45] and PU21 [37]. For example, under PQ, an increase of $1\,\text{cd/m}^2$ near darkness ($0.005\,\text{cd/m}^2$) is over $150\times$ more salient than the same increase at $100\,\text{cd/m}^2$. Beyond perceptual motivation, prior work [28] also shows that networks train more effectively on PQ/PU21-encoded HDR/RAW data, rather than represented in linear space. These observations suggest adopting a perceptually uniform representation for HDR synthesis.

In X2HDR, we encode linear HDR values using PU21, approximated by a log-quadratic function [28]:

$$V = f_{\mathrm{PU21}}(L) \;=\; a\left(\log_2(L) - L_{\min}\right)^2 \;+\; b\left(\log_2(L) - L_{\min}\right), \tag{1}$$

where L denotes an absolute linear RGB value within $[0.005, 10, 000]$. The fitted parameters are $a = 0.001908$ and $b = 0.0078$, and $L_{\min}$ is set to $\log_2(0.005)$. PU21 compresses extreme highlights while allocating more representational resolution to low-luminance values (see Supplemental Section E). The inverse transform is

$$L = f^{-1}_{\mathrm{PU21}}(V) \;=\; 2^{\frac{2aL_{\min} - b + \sqrt{b^2 + 4aV}}{2a}}. \tag{2}$$

For all HDR images, we globally rescale linear RGB values so that the maximum luminance corresponds to $L_{\mathrm{peak}} = 4,000\,\text{cd/m}^2$ (matching the peak luminance of commercially available HDR displays). We then apply $f_{\mathrm{PU21}}(\cdot)$ channel-wise to map the rescaled values to $[0, 1]$. For generated outputs, we apply $f^{-1}_{\mathrm{PU21}}(\cdot)$ to recover linear RGB values.

2.2 Pretrained VAEs for HDR Reconstruction

Modern diffusion models typically employ a VAE to map LDR images into a compact latent space, while reconstructing them with high fidelity. This raises a practical question central to X2HDR: must the VAE be finetuned for HDR reconstruction, or can an LDR-pretrained one already reconstruct HDR content accurately, provided an appropriate encoding is used?

Setups We curate a Blu-ray movie (namely, "BBC Blue Planet II Episode 2 - The Deep") that provides both LDR and HDR versions of the same content. Frames are time-synchronized by selecting the temporal offset that maximizes cross-correlation between corresponding pixel values. After temporal alignment and resolution matching, we randomly crop and resize paired frames to 768×768, producing a set of 512 (LDR, HDR) pairs with identical scene content. LDR frames are rescaled to $[0, 1]$, while HDR frames are processed in two ways: 1) converted into a perceptually uniform space using Equation (1), or 2) kept in linear space but normalized to $[0, 1]$. We then encode and decode all images using the pretrained `FLUX.1-dev` VAE. Reconstruction fidelity is evaluated by ColorVideoVDP [38], which reports perceptual just-objectionable-difference (JOD) scores (one JOD corresponds to roughly a 75% observer preference), together with standard LDR metrics [70, 62] applied via exposure-optimized inverse display models [8].

Table 1: VAE reconstruction fidelity on 512 aligned (LDR, HDR) frame pairs. The JOD score is given by ColorVideoVDP [38]. $Q^\star$ represents a family of LDR quality measures, equipped with exposure-optimized inverse display models [8].

Setting	JOD ↑	$Q^\star_{\mathrm{PSNR}}$ ↑	$Q^\star_{\mathrm{SSIM}}$ ↑	$Q^\star_{\mathrm{LPIPS}}$ ↓	$Q^\star_{\mathrm{DISTS}}$ ↓
LDR	**9.86**	**38.8**	0.961	0.026	**0.008**
HDR (Linear)	8.54	27.9	0.867	0.129	0.067
HDR (PU21)	9.44	35.6	**0.965**	**0.020**	0.010

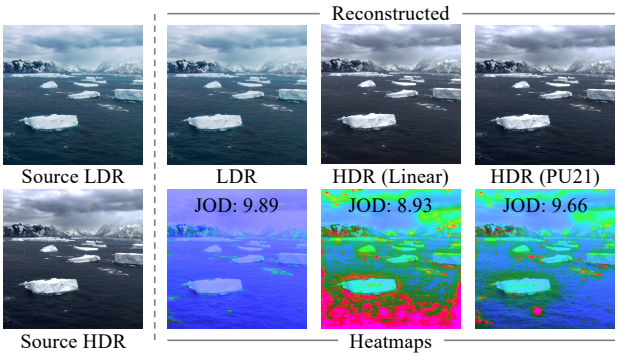

Figure 2: VAE reconstruction of an aligned (LDR, HDR) pair under different input encodings. Left: source LDR/HDR images fed to the VAE. Right: reconstructions (top row) for LDR, linear HDR, and PU21-encoded HDR representation, together with ColorVideoVDP perceptual error heatmaps (bottom row), where stronger colors indicate more noticeable differences. Insets report the corresponding JOD scores.

Results As summarized in Table 1, PU21-encoded HDR inputs are reconstructed nearly as well as LDR inputs: the JOD score drops only slightly ($9.86 \rightarrow 9.44$), and the other metrics show similarly small gaps. In contrast, linear HDR inputs exhibit substantial degradation across all metrics. These quantitative trends are consistent with Figure 2: LDR and PU21 reconstructions preserve scene structure and texture with only minor (and imperceptible) artifacts, whereas linear HDR reconstructions present widespread distortions.

Analysis PU21 improves VAE reconstruction by mapping HDR values to a perceptually meaningful representation that more closely matches the statistics of the LDR pretraining data. To illustrate this effect, we construct the inset histograms from 512 aligned LDR-HDR image pairs. We average the three RGB channels at each pixel, pool the resulting values across all images, and plot their probability densities on a logarithmic y-axis. Under linear HDR encoding, values are concentrated near zero, leaving much of the normalized input range underutilized. This highly skewed distribution differs substantially from that of the LDR training data. The mismatch also extends to the VAE latent space: linear HDR inputs produce a latent distribution that deviates from the LDR latent prior, with a Kullback-Leibler (KL) divergence of 0.231 across the 512 pairs. By compressing extreme highlights and allocating greater representational resolution to shadows and midtones, PU21 distributes HDR values more broadly across the normalized range and reduces their concentration near zero. This transformation brings the inputs into closer alignment with the statistics expected by the LDR-pretrained VAE, reducing the latent-space KL divergence to 0.033. Additional analysis and a t-SNE visualization of the VAE latent space are provided in Supplemental Sections D and F.4.

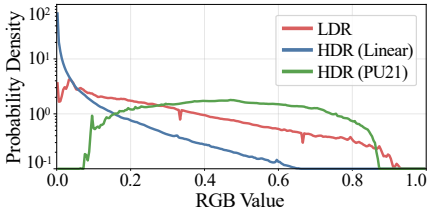

3 Text-to-HDR Image Generation

As demonstrated in Section 2, the pretrained VAE can faithfully encode and decode HDR images after PU21 mapping, achieving reconstruction fidelity close to that obtained on standard LDR images. This observation simplifies text-to-HDR image generation: we keep the text encoder and VAE fixed, and finetune only the denoiser (through LoRA) while operating entirely in a perceptually uniform space. Our system diagram is shown in Figure 3.

Training Data Preprocessing For each training HDR image I_{HDR}, we first apply a global rescaling so that its maximum luminance corresponds to $L_{\mathrm{peak}} = 4,000\,\mathrm{cd/m^2}$, and then map the rescaled

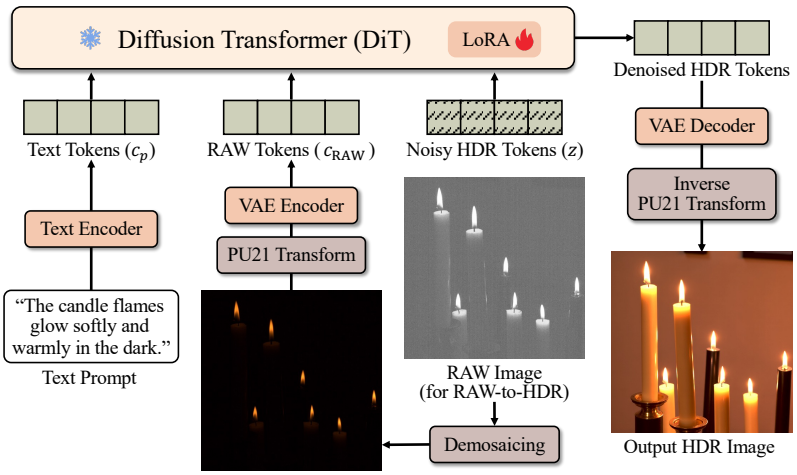

Figure 3: System diagram of X2HDR. For *Text-to-HDR*, the prompt is encoded into text tokens c_p; a noisy HDR latent z is sampled and denoised by the LoRA-adapted DiT; the final result is decoded by the frozen VAE and mapped back to linear HDR values via the inverse PU21 transform (Equation (2)). For *RAW-to-HDR*, a RAW capture is demosaicked, PU21-encoded (Equation (1)), and converted into image tokens c_{RAW}; conditioning on c_{RAW}, the same denoising and decoding pipeline reconstructs the HDR output. Both branches are shown in a single diagram for simplicity; in practice, training is performed separately with a task-specific LoRA for each branch.

values into the PU21 space. This conversion is the only departure from standard LDR finetuning; all subsequent steps follow the standard latent generative training procedure.

Latent Formulation and Objective We map the PU21-encoded HDR image into VAE latents x_0 and tokenize the paired text prompt into c_p. We then adopt the `FLUX.1-dev` T2I backbone and finetune it with flow matching. Specifically, for a timestep $t \in [0, 1]$, we sample $\epsilon \sim \mathcal{N}(0, I)$ and construct an interpolated noisy latent $z = (1 - t)x_0 + t\epsilon$. We then finetune the model by optimizing the flow-matching objective [2, 33, 20]:

$$\ell_{\text{flow}}\left(\Theta; c_p, x_0\right) = \mathbb{E}_{t,\epsilon}\ \left\| v_\Theta\left(z, t, c_p\right) - (\epsilon - x_0) \right\|_2^2, \tag{3}$$

where $v_\Theta(\cdot)$ denotes the learned velocity field with trainable parameters Θ, implemented via LoRA injection into the backbone.

LoRA Finetuning We employ LoRA [26], a parameter-efficient finetuning method to adapt the pretrained denoiser to the PU21-encoded HDR representation. LoRA represents the weight update of a linear layer as a low-rank decomposition: for a pretrained weight matrix $W \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$, instead of directly optimizing W, we freeze it and learn an additive update $\Delta W = BA$ with rank r, where $A \in \mathbb{R}^{r \times d_{\text{in}}}$ and $B \in \mathbb{R}^{d_{\text{out}} \times r}$. The adapted layer becomes

$$W' = W + \frac{\alpha}{r} BA, \tag{4}$$

where α is a scaling factor. We insert LoRA modules into linear layers of attention blocks (*e.g.*, query, key, value, and output projections) and optimize only the LoRA parameters while keeping the text encoder and VAE frozen. This yields an HDR-capable model with minimal trainable parameters and unchanged inference architecture.

Inference At test time, given a text prompt, we sample an initial noise latent and integrate the learned flow (with merged low-rank updates) to generate a PU21-encoded HDR latent. The final HDR image in linear space is produced by the frozen VAE decoding and the inverse PU21 transform.

4 RAW-to-HDR Image Reconstruction

We next show that the same paradigm extends naturally to HDR reconstruction from a single RAW capture. Unlike sRGB-encoded LDR images, RAW offers a physically meaningful (albeit noisy)

linear measurement that is approximately proportional to the incident light at the sensor before the ISP's nonlinear stages of processing.

We cast RAW-to-HDR as a direct mapping from a single RAW image to an HDR output, simultaneously denoising the sensor measurement and plausibly inpainting saturated or clipped regions. Although our method could be applied to enhance sRGB-encoded LDR images (see Supplemental Section F.7 for experiments in this setting), we focus on RAW inputs because 1) they typically retain more information in both underexposed and overexposed areas, and 2) they preserve approximately linear sensor values, avoiding the ill-posed problem of inverting in-camera processing. Accordingly, our goal is integration with the camera ISP rather than post-processing legacy LDR content.

Given a RAW input, we first demosaic it following Ramanath *et al.* [51]. We then rescale the peak luminance of both the RAW input and the ground-truth (GT) HDR target to $L_{\text{peak}} = 4,000\,\text{cd/m}^2$, and convert both to the PU21 space. We next encode the text prompt, RAW input, and target HDR to c_p, c_{RAW}, and x_0, respectively. Training is guided by the same flow-matching objective ℓ_{flow} with LoRA finetuning as in the text-to-HDR setting, except that the learned velocity field additionally conditions on the RAW input:

$$\ell_{\text{flow}}\left(\Theta; c_p, c_{\text{RAW}}, x_0\right) = \mathbb{E}_{t,\epsilon}\ \left\| v_\Theta\left(z, t, c_p, c_{\text{RAW}}\right) - (\epsilon - x_0)\right\|_2^2 . \tag{5}$$

Because our `FLUX.1-dev` backbone is DiT-based [49], supporting variable-length token sequences, we inject RAW conditioning by simply concatenating the RAW image tokens with the text and HDR latent tokens [71, 57]. Empirically, this simple design converges fast and reliably: the model learns to recover coarse scene structure and texture early in training, while dark-region denoising and bright-region inpainting improve steadily with continued optimization.

5 Experiments

In this section, we benchmark X2HDR against previous HDR adaptation techniques that follow a bracket-and-merge paradigm (*i.e.*, LEDiff [61] and Bracket Diffusion [6]) and a dedicated RAW-to-HDR reconstruction method (*i.e.*, RawHDR [73]) using metrics that capture 1) perceptual image quality and text-image alignment, 2) effective dynamic range, and 3) HDR reconstruction fidelity.

5.1 Evaluation Metrics

- *Perceptual image quality and text-image alignment*. We employ Q-Eval-100K [72], which provides two finetuned vision and language models (VLMs) (based on `Qwen2-VL-7B-Instruct`) that produce scalar scores in $[0, 1]$ for perceptual image quality and text-image alignment assessment for LDR images. To better match the VLMs' expected input statistics, HDR images are first PU21-encoded prior to evaluation. A dedicated sanity check is reported in Supplemental Section F.12.
- *Effective dynamic range*. For each HDR image generated in text-to-HDR, we compute an effective dynamic-range estimate in exposure "stops." We first apply Gaussian smoothing to suppress prediction noise, $\tilde{L} = G_\sigma * L$, using $\sigma = 3$ pixels. We then remove extreme outliers by retaining only the $[0.5, 99.5]$ percentile range, and define the dynamic range (in stops) as

$$\text{DR}_{\text{stops}} = \log_2(\tilde{L}_{99.5}/\tilde{L}_{0.5}), \tag{6}$$

  where $\tilde{L}_{0.5}$ and $\tilde{L}_{99.5}$ are the 0.5^{th} and 99.5^{th} percentiles of the filtered luminance, respectively.
- *HDR reconstruction fidelity*. For reconstruction evaluation, we use two complementary measures. We report JOD scores using ColorVideoVDP [38]. Meanwhile, we compute exposure-optimized variants of standard LDR metrics—PSNR, SSIM [62], LPIPS [70], and DISTS [16]—via an inverse display model [8] that decomposes both the reconstruction and reference into aligned exposure stacks and evaluates each LDR metric at the best-matching exposures.

5.2 Text-to-HDR Comparison

We compare X2HDR instantiated with two backbones (`FLUX.1-dev` and `SD-1.5`) against LEDiff [61] and Bracket Diffusion [6]. We evaluate on 100 text prompts generated by `ChatGPT-5`, spanning diverse lighting conditions (*e.g.*, sunlight, candles, aurora, and fireworks). To match baseline

Table 2: Quantitative results for text-to-HDR. We report image quality and text-image alignment scores by Q-Eval-100K [72] (higher is better), effective dynamic range by Equation (6) in stops (higher is better), and inference cost measured by runtime and peak memory (lower is better).

Method	Image Quality ↑	Text-Image Alignment ↑	Dynamic Range ↑	Time (sec) ↓	Memory (GB) ↓
LEDiff	0.488	0.623	4.3	7.4	8.6
Bracket Diffusion	0.524	0.671	11.1	557.5	38.0
Ours (SD)	0.568	0.672	**14.4**	**2.2**	**2.9**
Ours (FLUX)	**0.668**	**0.831**	14.2	11.2	33.0

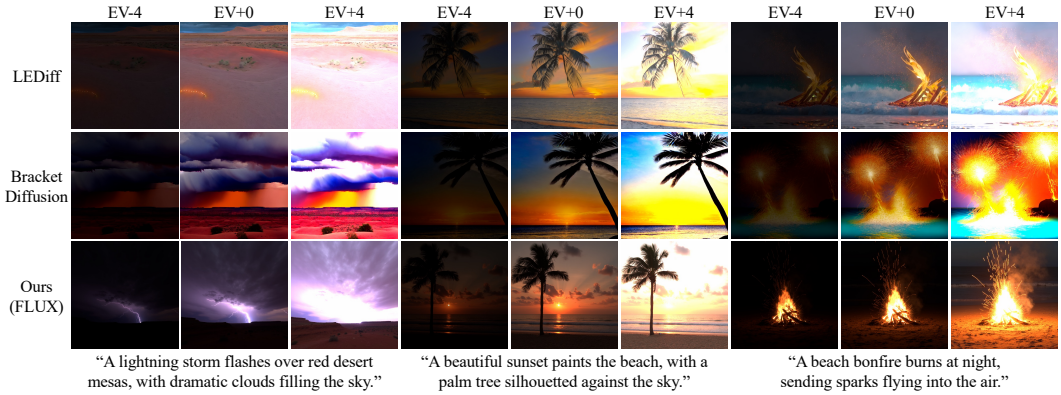

Figure 4: Visual comparison on text-to-HDR generation shown at three exposure settings (EV -4, EV $+0$, and EV $+4$). X2HDR better preserves light sources at low exposure (EV -4) and reveals informative shadow details at high exposure (EV $+4$), demonstrating a wider effective dynamic range. In contrast, LEDiff and Bracket Diffusion often underestimate peak luminance (dimming highlights at EV -4) and/or lose structure in underexposed areas (*e.g.*, shadow collapse in the palm tree example for Bracket Diffusion).

constraints, all methods generate 512×512 outputs for this comparison, though X2HDR supports higher resolutions. The full set of prompts, additional details for the `SD-1.5` variant, and results at higher resolutions are provided in Supplemental Sections C, F.5, and F.9.

Quantitative Comparison Table 2 shows the quantitative results. X2HDR achieves the best perceptual image quality and text-image alignment, with the `FLUX`-based variant performing strongest. Even when controlling for backbone (*i.e.*, `SD-1.5`), X2HDR remains competitive or better than competing methods while being much lighter at inference, owing to its parameter-efficient finetuning.

Beyond Q-Eval-100K scores, X2HDR produces HDR outputs with a wide effective dynamic range ($\approx$14 stops). In contrast, LEDiff yields significantly lower dynamic range (see also Figure 5). Bracket Diffusion reaches higher $\mathrm{DR}_{\mathrm{stops}}$ than LEDiff, but does so at extreme computational cost (*i.e.*, hundreds of seconds per image) and can exhibit luminance pathologies despite high $\mathrm{DR}_{\mathrm{stops}}$ statistics. For example, Bracket Diffusion may collapse large regions toward near-zero values, *i.e.*, a "shadow clipping" failure mode (see the palm tree in the second row of Figure 4), whereas X2HDR tends to maintain a more natural luminance distribution for each generation.

Qualitative Comparison Figure 4 visualizes results at multi-exposure settings (in the sRGB color space). At EV -4, X2HDR better preserves distinct highlight emitters (*e.g.*, lightning, bonfires, and the sun), indicating its capability of predicting sufficiently high peak luminance; LEDiff and Bracket Diffusion often appear dimmer with reduced highlight contrast. At EV $+4$, X2HDR reveals additional shadow structure, whereas LEDiff often washes out toward white and Bracket Diffusion tends to leave large regions nearly black due to shadow clipping. Overall, X2HDR better retains usable details across both highlights and shadows.

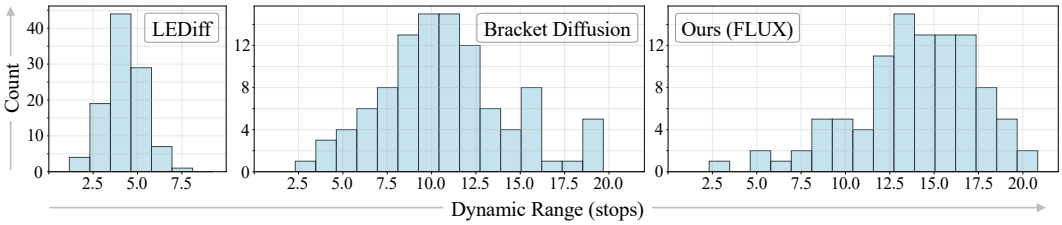

Figure 5: Distributions of effective dynamic range (in stops) over 100 generated HDR images.

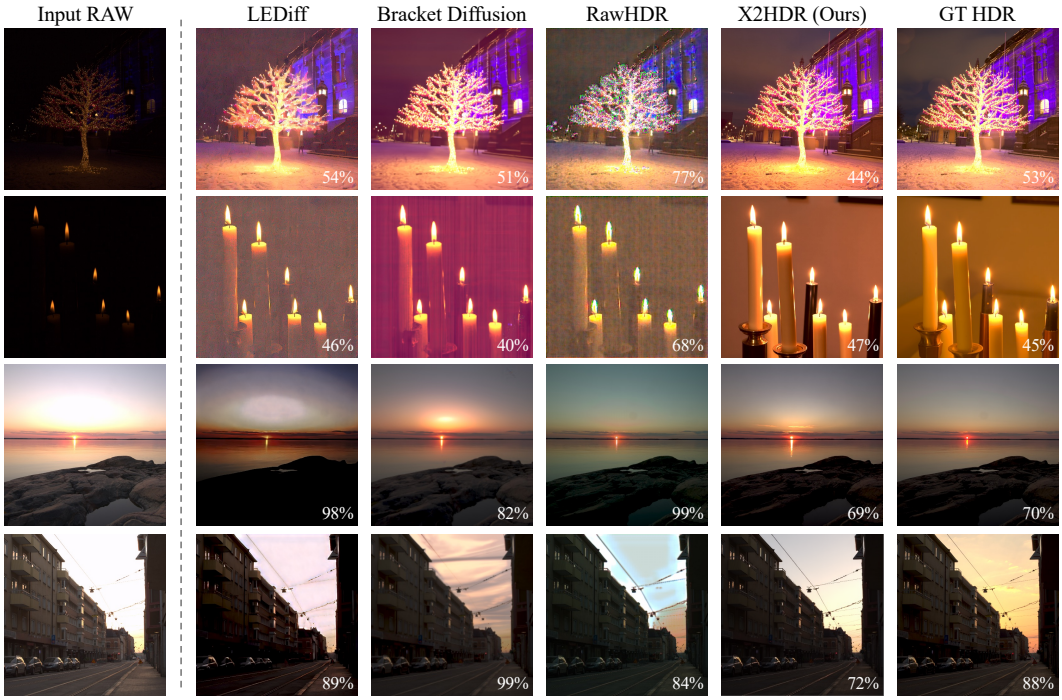

Figure 6: Visual comparison on RAW-to-HDR reconstruction. The competing methods exhibit similar failure modes: RawHDR is prone to noise amplification and color instability at low exposure, LEDiff often trades fidelity for excessive smoothing, and Bracket Diffusion's low conditioning resolution tends to soften edges and attenuate fine structure after upsampling. X2HDR, on the other hand, better balances denoising and detail preservation, and more consistently restores plausible content in clipped regions (*e.g.*, saturated sky) while maintaining stable color appearance. For better visual comparison, we map HDR outputs to LDR using the inverse display model of Mantiuk and Heidrich [39], optimizing the respective exposures (parameterized by luminance percentiles) to best align perceived brightness relative to the selected GT HDR exposure.

5.3 RAW-to-HDR Comparison

For RAW-to-HDR, we compare X2HDR to RawHDR [73] and additionally include LEDiff and Bracket Diffusion, adapting them to the RAW setting by a RAW-to-sRGB conversion before applying their original pipelines. Evaluation uses 96 RAW images from the SI-HDR dataset [24] at 512×512 (see Supplemental Section F.2). For methods supporting text conditioning, we use an empty prompt (*i.e.*, setting $c_p = \varnothing$) to ensure fairness. Text-guided hallucination results are deferred to Supplemental Section F.10.

Quantitative Comparison Table 3 reports the JOD scores by ColorVideoVDP [38] and exposure-optimized $Q^\star$ metrics [8]. X2HDR achieves the best performance across all reported measures, outperforming RawHDR as well as the LDR-to-HDR counterparts in this RAW setting. These trends are consistent with the methodological differences: RawHDR trains a small U-Net from scratch,

Table 3: Quantitative results for RAW-to-HDR.

Method	JOD ↑	$Q^\star_{\mathrm{PSNR}}$ ↑	$Q^\star_{\mathrm{SSIM}}$ ↑	$Q^\star_{\mathrm{LPIPS}}$ ↓	$Q^\star_{\mathrm{DISTS}}$ ↓
RawHDR	6.58	19.3	0.591	0.282	0.133
LEDiff	6.74	19.7	0.570	0.319	0.151
Bracket Diffusion	7.08	21.6	0.727	0.282	0.148
X2HDR (Ours)	**7.57**	**22.5**	**0.741**	**0.174**	**0.089**

Table 4: Ablation results for text-to-HDR and RAW-to-HDR.

(a) Text-to-HDR

Configuration	Image Quality ↑	Text-Image Alignment ↑	Dynamic Range ↑
HDR (Linear)	0.578	0.744	5.5
HDR (Gamma)	0.629	0.813	10.4
HDR (Log)	0.627	0.799	10.5
HDR (PQ)	0.661	0.822	**14.5**
HDR (PU21, default)	**0.668**	**0.831**	14.2

(b) RAW-to-HDR

Configuration	JOD ↑	$Q^\star_{\mathrm{PSNR}}$ ↑	$Q^\star_{\mathrm{SSIM}}$ ↑	$Q^\star_{\mathrm{LPIPS}}$ ↓	$Q^\star_{\mathrm{DISTS}}$ ↓
HDR (Linear)	5.60	12.9	0.306	0.551	0.302
HDR (Gamma)	5.99	14.9	0.435	0.436	0.222
HDR (Log)	6.12	15.8	0.476	0.416	0.208
HDR (PQ)	7.48	22.3	0.733	0.178	0.092
HDR (PU21, default)	**7.57**	**22.5**	**0.741**	**0.174**	**0.089**

limiting its ability to plausibly inpaint missing content, while LEDiff and Bracket Diffusion operate with weaker priors and face a larger domain gap when driven by sRGB-converted RAW inputs.

Qualitative Comparison Figure 6 shows representative results on four test RAW images. For underexposed inputs, RawHDR is inclined to introduce noticeable chromatic noise (*e.g.*, in the tree of the first row). For overexposed scenes, RawHDR often suffers from color shifts (*e.g.*, green cast) and struggles to inpaint coherent sky texture. LEDiff and Bracket Diffusion likewise fail to hallucinate plausible bright-region content. When conditioned on image inputs, Bracket Diffusion is further constrained by a lower operating resolution (256×256), leading to blur after upsampling. In contrast, X2HDR suppresses noise in dark regions and performs more spatially consistent inpainting in saturated areas, yielding cleaner and more realistic HDR reconstructions overall.

5.4 Ablation Studies

Necessity for Perceptually Uniform Representation To isolate the role of HDR representation, we replace the default PU21 encoding with linear encoding (normalized to $[0, 1]$), gamma encoding ($\gamma = 2.2$), log encoding ($\mu = 500$), and PQ encoding [45]. As shown in Table 4 and Figures 7 and 8, all non-perceptually-uniform alternatives substantially degrade HDR behavior: in text-to-HDR, linear encoding yields severely limited dynamic range (*i.e.*, $\mathrm{DR}_{\mathrm{stops}} = 5.5$), and in RAW-to-HDR, it introduces strong artifacts (*e.g.*, quantization), causing large drops across reconstruction metrics. Gamma and log encodings mitigate these issues to some extent, but both still exhibit reduced dynamic range and noticeable quantization artifacts. This arises because these representations assign an excessive portion of the available output range to highlights, resulting in a significant mismatch with sRGB statistics and lower image quality and text-image alignment scores under Q-Eval-100K. PQ, as expected, performs comparably to PU21 on both tasks, supporting the central claim that perceptually uniform HDR representation is critical for adapting LDR-pretrained T2I models to HDR synthesis and reconstruction.

Necessity for Finetuning We also test whether HDR outputs can be obtained without any finetuning by directly decoding and rescaling pretrained T2I latents, followed by the inverse PU21 transform. This text-to-HDR configuration frequently over-stretches the tonal and chromatic range, leading to two characteristic failure modes: 1) exaggerated global contrast and saturation caused by over-expanding LDR latents (see the third row of Figure 7) and 2) widespread clipping that removes recoverable details (*e.g.*, the grassland remains dark even at EV $+4$). We direct the readers to more visual examples in the HTML supplementary and the discussion in Supplemental Section F.12.

6 Conclusion

We have presented X2HDR, a simple and effective adaptation method that enables pretrained T2I diffusion models to operate in HDR for both text-to-HDR generation and RAW-to-HDR recon-

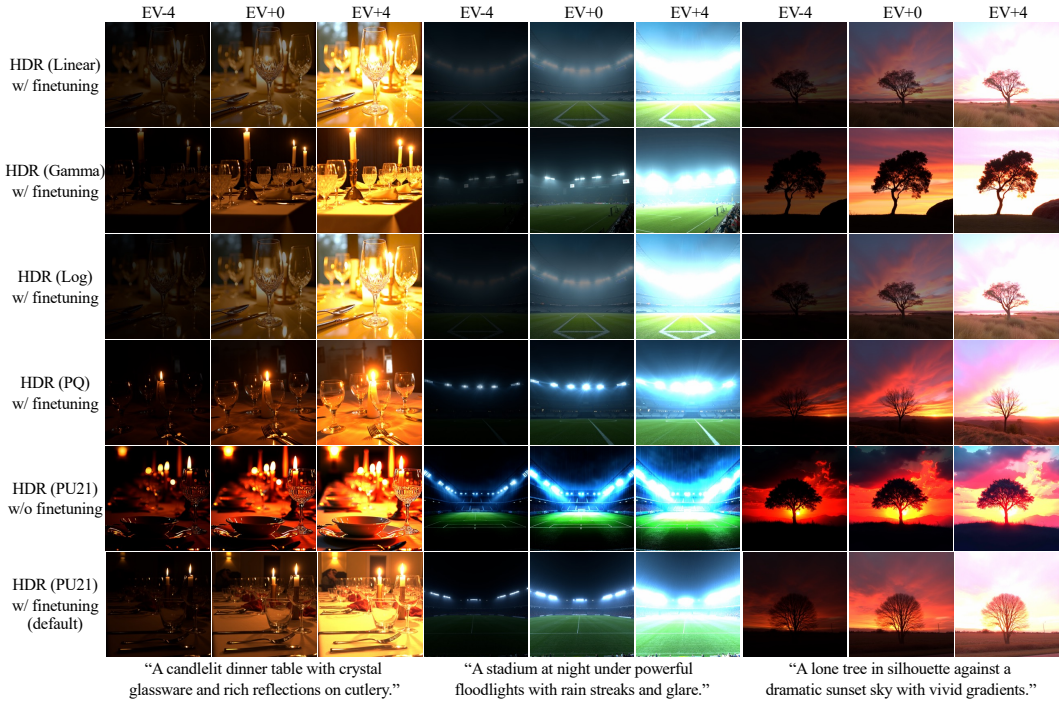

Figure 7: Text-to-HDR ablation on HDR representation and finetuning. Linear/gamma/log encodings severely compress bright emitters, most evident at EV -4, leading to a reduced effective dynamic range and poorer highlight/shadow recoverability. PQ closely matches PU21, preserving plausible details, whereas removing finetuning yields a "pseudo-HDR" appearance with exaggerated saturation/contrast and visible clipping artifacts.

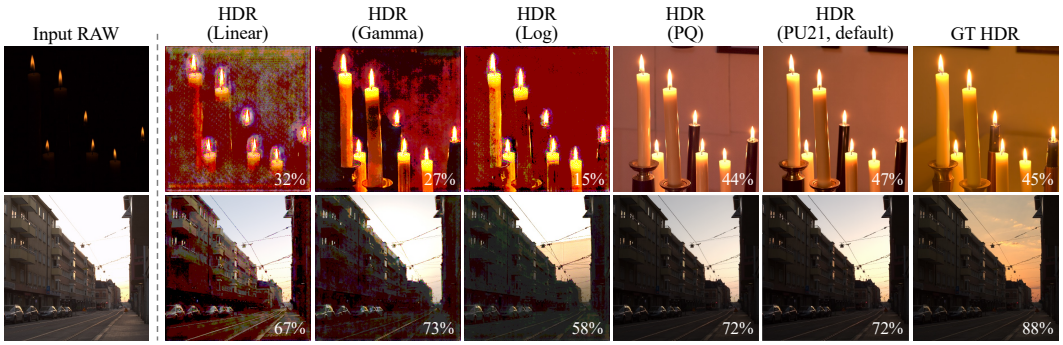

Figure 8: RAW-to-HDR ablation on HDR representation. Linear/gamma/log encodings introduce pronounced quantization and color distortions, whereas PQ and PU21 achieve comparable perceptual quality, with more faithful highlight recovery and reduced artifacts. In the candle scene, the background content varies between runs despite similar encoding behavior, explained by the stochasticity of the inpainting process.

struction. The key idea is to mitigate the distribution gap between linear-light HDR/RAW data and LDR-pretrained models by mapping HDR values into a perceptually uniform space. This simple preprocessing allows an off-the-shelf LDR-pretrained VAE to represent HDR content faithfully, so HDR capability can be acquired by parameter-efficient LoRA finetuning of the denoiser while keeping the rest of the computational structure unchanged. The resulting X2HDR is unified and deployment-friendly, and it yields reliable HDR generation and reconstruction results with improved perceptual fidelity and usable dynamic range.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China and the Research Grants Council of Hong Kong under the Collaborative Research Scheme (Project No. CRS_HKUST605/25), and by the Innovation and Technology Fund of the Hong Kong Innovation and Technology Commission (Project Nos. 9440379 and 9440390).

References

[1] Ahmet O. Akyüz, Roland Fleming, Bernhard E. Riecke, Erik Reinhard, and Heinrich H. Bülthoff. Do HDR displays support LDR content? A psychophysical evaluation. *ACM Transactions on Graphics*, 26(3):1–8, 2007.

[2] Michael S. Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *arXiv preprint arXiv:2209.15571*, 2022.

[3] Tunç O. Aydın, Rafał K. Mantiuk, and Hans-Peter Seidel. Extending quality metrics to full luminance range images. In *Human Vision and Electronic Imaging*, pages 109–118, 2008.

[4] Francesco Banterle, Patrick Ledda, Kurt Debattista, Marina Bloj, Alessandro Artusi, and Alan Chalmers. A psychophysical evaluation of inverse tone mapping techniques. *Computer Graphics Forum*, 28(1):13–25, 2009.

[5] Francesco Banterle, Patrick Ledda, Kurt Debattista, and Alan Chalmers. Expanding low dynamic range videos for high dynamic range applications. In *Spring Conference on Computer Graphics*, pages 33–41, 2008.

[6] Mojtaba Bemana, Thomas Leimkühler, Karol Myszkowski, Hans-Peter Seidel, and Tobias Ritschel. Bracket Diffusion: HDR image generation by consistent LDR denoising. *Computer Graphics Forum*, 44(2):1–13, 2025.

[7] Christophe Bolduc, Justine Giroux, Marc Hébert, Claude Demers, and Jean-François Lalonde. Beyond the Pixel: A photometrically calibrated HDR dataset for luminance and color prediction. In *IEEE International Conference on Computer Vision*, pages 8071–8081, 2023.

[8] Peibei Cao, Rafał K. Mantiuk, and Kede Ma. Perceptual assessment and optimization of HDR image rendering. In *IEEE Computer Vision and Pattern Recognition*, pages 22433–22443, 2024.

[9] Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart-α: Fast training of diffusion Transformer for photorealistic text-to-image synthesis. In *arXiv preprint arXiv:2310.00426*, 2023.

[10] Rufeng Chen, Bolun Zheng, Hua Zhang, Quan Chen, Chenggang Yan, Gregory Slabaugh, and Shanxin Yuan. Improving dynamic HDR imaging with fusion Transformer. In *Association for the Advancement of Artificial Intelligence*, pages 340–349, 2023.

[11] Zixuan Chen, Yujin Wang, Xin Cai, Zhiyuan You, Zheming Lu, Fan Zhang, Shi Guo, and Tianfan Xue. UltraFusion: Ultra high dynamic imaging using exposure fusion. In *IEEE Computer Vision and Pattern Recognition*, pages 16111–16121, 2025.

[12] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *ACM Special Interest Group on Computer Graphics and Interactive Techniques*, pages 643–652, 1997.

[13] Piotr Didyk, Rafał K. Mantiuk, Matthias Hein, and Hans-Peter Seidel. Enhancement of bright video features for HDR displays. *Computer Graphics Forum*, 27(4):1265–1274, 2008.

[14] Sebastian Dille, Chris Careaga, and Yağız Aksoy. Intrinsic single-image HDR reconstruction. In *European Conference on Computer Vision*, pages 161–177, 2024.

[15] Keyan Ding, Yi Liu, Xueyi Zou, Shiqi Wang, and Kede Ma. Locally adaptive structure and texture similarity for image quality assessment. In *ACM Multimedia*, pages 2483–2491, 2021.

[16] Keyan Ding, Kede Ma, Shiqi Wang, and Eero P. Simoncelli. Image quality assessment: Unifying structure and texture similarity. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(5):2567–2581, 2020.

[17] Frédéric Drago, Karol Myszkowski, Thomas Annen, and Norishige Chiba. Adaptive logarithmic mapping for displaying high contrast scenes. *Computer Graphics Forum*, 22(3):419–426, 2003.

[18] 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*, 36(6):1–15, 2017.

[19] Yuki Endo, Yoshihiro Kanamori, and Jun Mitani. Deep reverse tone mapping. *ACM Transactions on Graphics*, 36(6):1–10, 2017.

[20] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, Kyle Lacey, Alex Goodwin, Yannik Marek, and Robin Rombach. Scaling rectified flow Transformers for high-resolution image synthesis. In *International Conference on Machine Learning*, pages 1–28, 2024.

[21] Mark D. Fairchild. The HDR photographic survey. In *Color and Imaging Conference*, pages 233–238, 2007.

[22] Marc-André Gardner, Kalyan Sunkavalli, Ersin Yumer, Xiaohui Shen, Emiliano Gambaretto, Christian Gagné, and Jean-François Lalonde. Learning to predict indoor illumination from a single image. *ACM Transactions on Graphics*, 36(6):1–14, 2017.

[23] Param Hanji and Rafał K. Mantiuk. Robust estimation of exposure ratios in multi-exposure image stacks. *IEEE Transactions on Computational Imaging*, 9(1):721–731, 2023.

[24] Param Hanji, Rafał K. Mantiuk, Gabriel Eilertsen, Saghi Hajisharif, and Jonas Unger. Comparison of single image HDR reconstruction methods — the caveats of quality assessment. In *ACM Special Interest Group on Computer Graphics and Interactive Techniques*, pages 1–8, 2022.

[25] Param Hanji, Fangcheng Zhong, and Rafał K. Mantiuk. Noise-aware merging of high dynamic range image stacks without camera calibration. In *European Conference on Computer Vision Workshop on Advances in Image Manipulation*, pages 376–391, 2020.

[26] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *arXiv preprint arXiv:2106.09685*, 2021.

[27] Nima K. Kalantari and Ravi Ramamoorthi. Deep high dynamic range imaging of dynamic scenes. *ACM Transactions on Graphics*, 36(4):1–12, 2017.

[28] Andrew Yanzhe Ke, Lei Luo, Xiaoyu Xiang, Yuchen Fan, Rakesh Ranjan, Alexandre Chapiro, and Rafał K. Mantiuk. Training neural networks on RAW and HDR images for restoration tasks. In *arXiv preprint arXiv:2312.03640*, 2023.

[29] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *arXiv preprint arXiv:1312.6114*, 2013.

[30] Lingtong Kong, Bo Li, Yike Xiong, Hao Zhang, Hong Gu, and Jinwei Chen. SAFNet: Selective alignment fusion network for efficient HDR imaging. In *European Conference on Computer Vision*, pages 256–273, 2024.

[31] Black Forest Labs. FLUX: Official inference repository for FLUX.1 models. `https://github.com/black-forest-labs/flux`, 2024.

[32] Siyeong Lee, Gwon-Hwan An, and Suk-Ju Kang. Deep Recursive HDRI: Inverse tone mapping using generative adversarial networks. In *European Conference on Computer Vision*, pages 596–611, 2018.

[33] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *arXiv preprint arXiv:2210.02747*, 2022.

[34] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow Straight and Fast: Learning to generate and transfer data with rectified flow. In *arXiv preprint arXiv:2209.03003*, 2022.

[35] 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 *IEEE Computer Vision and Pattern Recognition*, pages 1651–1660, 2020.

[36] Zhen Liu, Yinglong Wang, Bing Zeng, and Shuaicheng Liu. Ghost-free high dynamic range imaging with context-aware Transformer. In *European Conference on Computer Vision*, pages 344–360, 2022.

[37] Rafał K. Mantiuk and Maryam Azimi. PU21: A novel perceptually uniform encoding for adapting existing quality metrics for HDR. In *Picture Coding Symposium*, pages 1–5, 2021.

[38] Rafał K. Mantiuk, Param Hanji, Maliha Ashraf, Yuta Asano, and Alexandre Chapiro. ColorVideoVDP: A visual difference predictor for image, video and display distortions. *ACM Transactions on Graphics*, 43(4):1–20, 2024.

[39] Rafał K. Mantiuk and Wolfgang Heidrich. Visualizing high dynamic range images in a web browser. *Journal of Graphics, GPU, and Game Tools*, 14(1):43–53, 2009.

[40] Rafał K. Mantiuk, Grzegorz Krawczyk, Karol Myszkowski, and Hans-Peter Seidel. Perception-motivated high dynamic range video encoding. *ACM Transactions on Graphics*, 23(3):733–741, 2004.

[41] 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. *Computer Graphics Forum*, 37(2):37–49, 2018.

[42] Belen Masia, Sandra Agustin, Roland W. Fleming, Olga Sorkine, and Diego Gutierrez. Evaluation of reverse tone mapping through varying exposure conditions. In *ACM Special Interest Group on Computer Graphics and Interactive Techniques Asia*, pages 1–8, 2009.

[43] Belen Masia, Ana Serrano, and Diego Gutierrez. Dynamic range expansion based on image statistics. *Multimedia Tools and Applications*, 76(1):631–648, 2017.

[44] Aliaksei Mikhailiuk, Clifford Wilmot, Maria Perez-Ortiz, Dingcheng Yue, and Rafał K. Mantiuk. Active sampling for pairwise comparisons via approximate message passing and information gain maximization. In *IEEE International Conference on Pattern Recognition*, pages 2559–2566, 2021.

[45] Scott Miller, Mahdi Nezamabadi, and Scott Daly. Perceptual signal coding for more efficient usage of bit codes. *Motion Imaging Journal*, 122(4):52–59, 2013.

[46] Aamir Mustafa, Hongjie You, and Rafał K. Mantiuk. A comparative study on the loss functions for image enhancement networks. In *London Imaging Meeting*, pages 11–15, 2022.

[47] 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. In *arXiv preprint arXiv:2112.10741*, 2021.

[48] Karen Panetta, Landry Kezebou, Victor Oludare, Sos Agaian, and Zehua Xia. TMO-Net: A parameter-free tone mapping operator using generative adversarial network, and performance benchmarking on large scale HDR dataset. *IEEE Access*, 9(1):39500–39517, 2021.

[49] William Peebles and Saining Xie. Scalable diffusion models with Transformers. In *IEEE International Conference on Computer Vision*, pages 4195–4205, 2023.

[50] Poly Haven. HDRIs. `https://polyhaven.com/hdris`, 2025.

[51] Rajeev Ramanath, Wesley E. Snyder, Griff L. Bilbro, and William A. Sander. Demosaicking methods for bayer color arrays. *Journal of Electronic Imaging*, 11(3):306–315, 2002.

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

[53] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed K. S. Ghasemipour, Burcu-Karagol Ayan, S. Sara Mahdavi, Rapha G. Lopes, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, pages 36479–36494, 2022.

[54] Marcel S. Santos, Tsang-Ing Ren, and Nima K. Kalantari. Single image HDR reconstruction using a cnn with masked features and perceptual loss. In *arXiv preprint arXiv:2005.07335*, 2020.

[55] Pradeep Sen, Nima K. Kalantari, Maziar Yaesoubi, Soheil Darabi, Dan B. Goldman, and Eli Shechtman. Robust patch-based HDR reconstruction of dynamic scenes. *ACM Transactions on Graphics*, 31(6):1–11, 2012.

[56] Jou-Won Song, Ye-In Park, Kyeongbo Kong, Jaeho Kwak, and Suk-Ju Kang. Selective TransHDR: Transformer-based selective HDR imaging using ghost region mask. In *European Conference on Computer Vision*, pages 288–304, 2022.

[57] Zhenxiong Tan, Songhua Liu, Xingyi Yang, Qiaochu Xue, and Xinchao Wang. OminiControl: Minimal and universal control for diffusion Transformer. In *IEEE International Conference on Computer Vision*, pages 14940–14950, 2025.

[58] Steven Tel, Zongwei Wu, Yulun Zhang, Barthélémy Heyrman, Cédric Demonceaux, Radu Timofte, and Dominique Ginhac. Alignment-free HDR deghosting with semantics consistent Transformer. In *IEEE International Conference on Computer Vision*, pages 12790–12799, 2023.

[59] Louis L. Thurstone. A law of comparative judgment. *Psychological Review*, 101(2):273–286, 1927.

[60] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 6000–6010, 2017.

[61] Chao Wang, Zhihao Xia, Thomas Leimkuhler, Karol Myszkowski, and Xuaner Zhang. LEDiff: Latent exposure diffusion for HDR generation. In *IEEE Computer Vision and Pattern Recognition*, pages 453–464, 2025.

[62] 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.

[63] Shangzhe Wu, Jiarui Xu, Yu-Wing Tai, and Chi-Keung Tang. Deep high dynamic range imaging with large foreground motions. In *European Conference on Computer Vision*, pages 117–132, 2018.

[64] Qingsen Yan, Dong Gong, Qinfeng Shi, Anton V. D. Hengel, Chunhua Shen, Ian Reid, and Yanning Zhang. Attention-guided network for ghost-free high dynamic range imaging. In *IEEE Computer Vision and Pattern Recognition*, pages 1751–1760, 2019.

[65] Qingsen Yan, Song Zhang, Weiye Chen, Hao Tang, Yu Zhu, Jinqiu Sun, Luc Van Gool, and Yanning Zhang. SMAE: Few-Shot learning for HDR deghosting with saturation-aware masked autoencoders. In *IEEE Computer Vision and Pattern Recognition*, pages 5775–5784, 2023.

[66] Qian Ye, Jun Xiao, Kin-man Lam, and Takayuki Okatani. Progressive and selective fusion network for high dynamic range imaging. In *ACM Multimedia*, pages 5290–5297, 2021.

[67] Hanning Yu, Wentao Liu, Chengjiang Long, Bo Dong, Qin Zou, and Chunxia Xiao. Luminance attentive networks for HDR image and panorama reconstruction. *Computer Graphics Forum*, 40(7):181–192, 2021.

[68] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *IEEE International Conference on Computer Vision*, pages 3836–3847, 2023.

[69] Ning Zhang, Yuyao Ye, Yang Zhao, and Ronggang Wang. Revisiting the stack-based inverse tone mapping. In *IEEE Computer Vision and Pattern Recognition*, pages 9162–9171, 2023.

[70] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Computer Vision and Pattern Recognition*, pages 586–595, 2018.

[71] Yuxuan Zhang, Yirui Yuan, Yiren Song, Haofan Wang, and Jiaming Liu. EasyControl: Adding efficient and flexible control for diffusion Transformer. In *arXiv preprint arXiv:2503.07027*, 2025.

[72] Zicheng Zhang, Tengchuan Kou, Shushi Wang, Chunyi Li, Wei Sun, Wei Wang, Xiaoyu Li, Zongyu Wang, Xuezhi Cao, Xiongkuo Min, Xiaohong Liu, and Guangtao Zhai. Q-Eval-100K: Evaluating visual quality and alignment level for text-to-vision content. In *IEEE Computer Vision and Pattern Recognition*, pages 10621–10631, 2025.

[73] Yunhao Zou, Chenggang Yan, and Ying Fu. RawHDR: High dynamic range image reconstruction from a single RAW image. In *IEEE International Conference on Computer Vision*, pages 12334–12344, 2023.

A Overview

This appendix provides supplementary results and implementation details that support and extend the findings of the main paper. Specifically, it includes:

- Prior HDR generation and reconstruction methods (Section B);
- Complete HDR results for text-to-HDR generation and RAW-to-HDR reconstruction, including comparisons, ablations, and perceptual experiments, as well as HDR outputs at multiple resolutions and for synthetic stimuli (Sections C, F.5, F.6, and F.9);
- Detailed analysis of PU21 perceptually uniform encoding (Section D);

- HDR encoding function visualization (Section E);
- Dataset and reproducibility details for training and evaluation (Sections F.1 and F.3);
- Extended VAE analysis, including failure characterization, VAE finetuning, and latent distribution visualization (Section F.4);
- LDR-to-HDR reconstruction experiments (Section F.7);
- Perceptual study details (Section F.8);
- Applications to text-guided hallucination (Section F.10);
- Limitations (Section G) and future research directions (Section H).

B Related Work

Our work connects three lines of research: T2I diffusion models in LDR, efforts to extend such models to HDR, and HDR reconstruction from a single image or multi-exposure brackets.

B.1 T2I Diffusion Models in LDR

Diffusion models have become a dominant paradigm for text-guided image synthesis, producing high-quality results across diverse scene types, semantic compositions, and visual styles [9, 52, 53, 47]. Latent diffusion models [52] improve scalability by learning the generative process in a compressed latent space: a VAE maps high-dimensional images into low-dimensional latents, and a denoiser learns to invert progressive corruption in this space. More recent work integrates Transformer backbones [60] into diffusion processes, giving rise to diffusion Transformers (DiT) [49]. Building on these developments, `FLUX` [31] combines Transformer-based architectures with flow matching objectives [33, 2, 34], achieving state-of-the-art T2I generation performance. Here we mainly adopt `FLUX.1-dev` as the backbone and investigate how to extend it to HDR tasks with minimal changes.

B.2 HDR Image Generation

Compared with LDR generation, text-guided HDR image synthesis remains relatively underexplored. Two closely related methods are LEDiff [61] and Bracket Diffusion [6]. Both emulate the classical HDR workflow by generating multi-exposure "brackets" and subsequently merging them into a single HDR output [12]. LEDiff finetunes specialized denoisers for highlight hallucination and shadow recovery, and adapts the VAE decoder to produce linear RGB values. Bracket Diffusion instead jointly denoises over multiple exposures, incorporating constraints that preserve exposure diversity while enforcing semantic consistency across the exposure stack. While effective, these bracket-and-merge designs add architectural and algorithmic complexity, increase inference-time and memory requirements, and can introduce unwanted fusion artifacts—factors that hinder portability to newer, more memory-intensive T2I backbones. In contrast, X2HDR adapts pretrained T2I models by operating directly in a perceptually uniform space, avoiding explicit bracket generation and fusion while remaining compatible with modern architectures.

B.3 HDR Image Reconstruction

Most HDR imaging techniques rely on multi-exposure bracketing: capturing several LDR exposures and merging them into HDR radiance maps [12]. For dynamic scenes, however, camera/scene motion causes misalignment and ghosting, motivating an extensive body of work on motion-aware HDR reconstruction [55]. Recent methods integrate stronger alignment and deghosting modules [27, 64, 65, 11, 30, 66] with powerful backbones such as U-Nets [63] and Transformers [10, 36, 56, 58]. Nevertheless, robust reconstruction under large motion and severe saturation remains challenging.

A complementary line of research considers single-image LDR-to-HDR reconstruction (also referred to as inverse tone-mapping). Classical inverse methods expand dynamic range via heuristic linearization and optional detail hallucination [1, 5, 4, 13, 42, 43]. Learning-based methods either approximately invert the ISP end-to-end [18, 35, 41, 54, 67, 14], or synthesize multi-exposures from a single image and then apply conventional HDR merging [19, 32, 69, 61, 6]. However, limited training data can reduce generalization, and bracket-synthesis approaches tend to introduce inter-exposure inconsistencies, especially under extreme dark and bright regions.

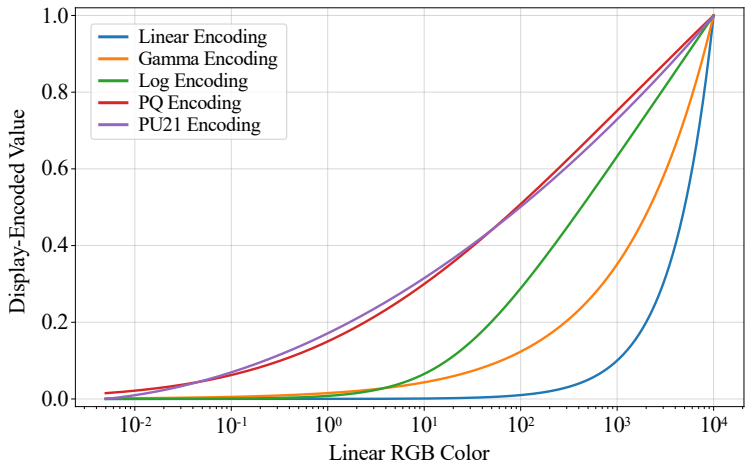

Figure 9: Encoding functions (linear, gamma, log, PQ, and PU21) mapping BT.2100-range linear luminance to display-encoded values (with log-scaled x-axis).

In this paper, we depart from the prevailing LDR-to-HDR paradigm and study HDR reconstruction from a single RAW capture, a setting that has received comparatively little attention. Zou *et al.* [73] explored a related setting with a rather elaborate workflow. In contrast, we propose a simple approach that leverages powerful generative priors to denoise and inpaint effectively during HDR reconstruction, yielding superior performance.

C Complete HDR Results

Because HDR content is difficult to reproduce faithfully in a static PDF, we provide extended results on the project website https://X2HDR.github.io/, including 1) 100 text-to-HDR images, 2) 96 RAW-to-HDR reconstructions, 3) ablation results under linear encoding, PQ encoding, and without finetuning for both text-to-HDR and RAW-to-HDR tasks, and 4) visual stimuli selected in the perceptual study with their corresponding JOD scores.

To ensure comparable perceived brightness, we normalize all HDR images by mapping the median luminance to $0.5\ \mathrm{cd/m^2}$. For efficient browser-based viewing, EXR files are converted to a compact JPEG gain-map representation using `gainmap-js`[1]. Note that a correct HDR viewing experience requires an HDR-capable display; on LDR monitors, appearance may deviate from true HDR content, and perceived quality depends on the specific HDR hardware.

D PU21 Perceptually Uniform Encoding

PU21 [37] is a scalar encoding that maps absolute HDR luminance values from 0.005 to $10{,}000\ \mathrm{cd/m^2}$ into an approximately perceptually uniform domain. In this domain, Euclidean distances correspond more closely to visible differences, allowing metrics such as PSNR and SSIM to be applied meaningfully to HDR images. The encoding is derived by numerically integrating contrast sensitivity, defined as the inverse of the detection threshold, over log-luminance. An analytic, invertible curve is then fitted and scaled to map the luminance range to $[0, 1]$.

Whereas PU08 [3] and PQ [45] are derived from peak contrast sensitivity, PU21 models the detection of banding artifacts caused by fixed-bit-depth quantization. It also accounts for the reduced visibility of differences in dark image regions due to disability glare—the scattering of light within the optics of the eye.

[1] `https://github.com/MONOGRID/gainmap-js`

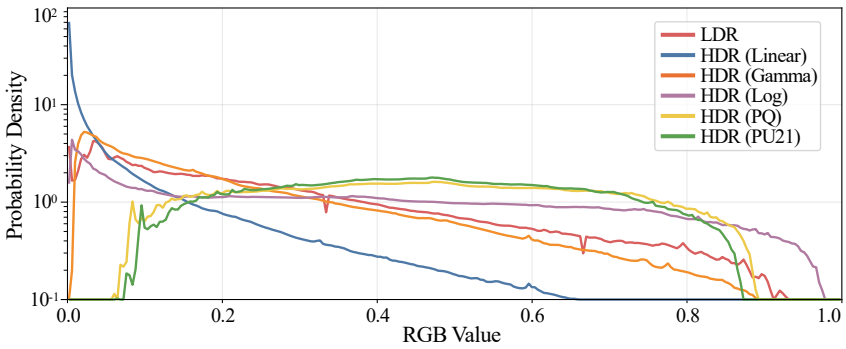

Figure 10: RGB value distributions under different encodings.

E Encoding Function Visualization

Figure 9 compares five encodings used to map linear RGB luminance to display-encoded values: linear encoding, gamma encoding ($\gamma = 2.2$), log encoding ($\mu = 500$), PQ [45], and PU21 [37]. Linear mapping assigns a disproportionately large range to highlights, heavily compressing low-luminance values. Gamma and log encodings partially rebalance this allocation, but they still do not achieve the same degree of perceptual uniformity as PQ and PU21. In contrast, PQ and PU21 nonlinearly compress extreme highlights while allocating more resolution to shadows and midtones, spreading low-luminance differences in the encoded domain. Although PQ and PU21 differ in functional form, both reshape HDR statistics toward distributions that better match LDR-pretrained models. We further visualize the RGB value distributions in Figure 10 using the histogram procedure described in Section 2.2. Gamma and log encodings produce distributions closer to that of LDR data than PQ and PU21 do. However, the ablation study in Section 5.4 shows that they underperform perceptually uniform encodings after finetuning, suggesting that LDR-like statistics alone are insufficient for HDR adaptation. One possible explanation lies in the inverse transformation applied after generation: the decoded image must be mapped back to linear RGB values, and this mapping can amplify errors introduced in the encoded domain. Although generic concave mappings such as gamma and log improve histogram alignment, they do not allocate encoded values according to perceptual sensitivity across the HDR luminance range. Consequently, small errors in the encoded domain can produce visible errors after inversion.

F More Experimental Details and Results

F.1 Training Dataset

Text-to-HDR We build the training set (3,278 HDR images) primarily from publicly available HDR datasets used in prior work [50, 7, 21, 22, 27, 35, 48, 58]. From each HDR image, we randomly extract 512×512 and $1,024\times1,024$ crops, retaining only those with an effective dynamic range of $\mathrm{DR}_{\mathrm{stops}} \geq 5$. For HDR panoramas, random perspective projections are applied before cropping. To better exploit high-resolution sources, we use an adaptive sampling scheme that draws proportionally more crops from larger images, producing $32,200$ patches at 512×512 and $6,714$ patches at $1,024 \times 1,024$. Text prompts are obtained by tone-mapping HDR crops [17] and captioning them with a VLM (*i.e.*, `Gemini-2.0-Flash`) using the following instruction:

```
You are a professional image-captioning assistant. Generate objective,
    accurate, and detailed captions based on the provided image. Produce two
    outputs: (i) a short caption (1--3 sentences) that summarizes the main
    content, and (ii) a long caption (one paragraph) that describes all
    salient details. Use precise, descriptive language; remain factual and
    avoid subjective interpretation. Output only the captions, formatted
    exactly as: '###Short:' and '###Long:'.
```

During training, we randomly select either the short or long caption at each iteration.

RAW-to-HDR The RAW-to-HDR training data are drawn from prior HDR datasets [21, 73], comprising 517 scenes with multiple RAW exposures. For each scene, bracketed RAW images are merged into a reference HDR image using HDRUtils[2], with exposure estimation and alignment [25, 23]. Paired (RAW, HDR) samples are formed by cropping spatially co-located regions. We extract 10 crops per RAW exposure, resulting in a total of 22,160 pairs at 512×512. HDR targets are tone-mapped and captioned with the same VLM.

F.2 RAW-to-HDR Test Dataset

Instead of splitting the collected training data, evaluation is performed on an independent SI-HDR dataset [24], consisting of 183 scenes with up to 7 RAW exposures and merged HDR references. It provides a clear domain gap and is well-suited for assessing the generalizability of RAW-to-HDR reconstruction methods. We randomly sample 96 scenes; for each scene, we discard the darkest and brightest exposures and then randomly choose one RAW image from the remaining. Each selected RAW image is cropped and resized to ensure pixel-alignment with the HDR reference at $1,888 \times 1,280$, and then further cropped and downsampled to 512×512 for evaluation.

F.3 Implementation Details

We use `FLUX.1-dev` as the default T2I backbone. For text-to-HDR, we set the LoRA rank to 32, with a learning rate of 1×10^{-4} and a dedicated LoRA trigger token, `[PU21]`. The per-GPU batch size is 8 for 512×512 and 2 for $1,024 \times 1,024$, with 4 gradient accumulation steps. Training converges after approximately 3,000 steps on 4 GPUs. For RAW-to-HDR, we adopt EasyControl [71] for conditional generation, which enhances the DiT backbone by introducing attention masking, position-encoding cloning, and conditional LoRA injection. We set the LoRA rank to 128 with learning rate 1×10^{-4}. During training, we randomly drop text prompts (*i.e.*, setting $c_p = \varnothing$) with 50% probability to implement RAW-only reconstruction. Training runs for 8 epochs on 512×512 images with batch size 2 on 4 GPUs. For both tasks, the LoRA scaling factor α is set equal to the LoRA rank.

F.4 Additional VAE Analysis

By default, X2HDR freezes the pretrained VAE and adapts only the denoiser of the T2I model. We further analyze when the frozen VAE fails on PU21-encoded HDR data and whether HDR-specific VAE finetuning provides additional benefits.

Failure characterization While PU21 generally induces high-fidelity reconstruction, the pretrained VAE occasionally under-reconstructs highlight regions (see Figure 11).

VAE finetuning To improve highlight reconstruction, we finetune the pretrained `FLUX.1-dev` VAE on PU21-encoded HDR data. Dataset preparation largely follows the text-to-HDR training pipeline (see Section F.1), but omits VLM captioning and fixes crop resolution to 768×768. This yields a total of 60,399 HDR images. Training uses learning rate 1×10^{-5}, batch size 1, 8 gradient accumulation steps, 2 epochs, and 4 GPUs. Several reconstruction losses are tested, including the mean absolute error (MAE), locally adaptive DISTS (ADISTS) metric [15], a differentiable histogram matching loss [46], and their variants combined with an inverse display model [39]. MAE is found sufficient.

After finetuning, average JOD on 512 test HDR images rises from 9.44 to 9.72, with visible improvements in some highlight cases (see the last column of Figure 11). However, gains are not consistent across content, and the finetuned VAE remains below the LDR reconstruction counterpart. The remaining gap may be due to limited HDR data scale and diversity (with many crops originating from relatively few scenes).

Latent distribution visualization Using 512 (LDR, HDR) pairs, we compare the VAE latent distributions for 1) LDR inputs, 2) linear HDR inputs, 3) PU21-encoded HDR inputs with the pretrained VAE, and 4) PU21-encoded HDR inputs with the finetuned VAE. A t-SNE projection with 2σ confidence ellipses in Figure 12 shows that PU21 significantly reduces the LDR-HDR mismatch relative to linear HDR, while VAE finetuning produces only marginal latent changes.

[2] `https://github.com/gfxdisp/HDRUtils`

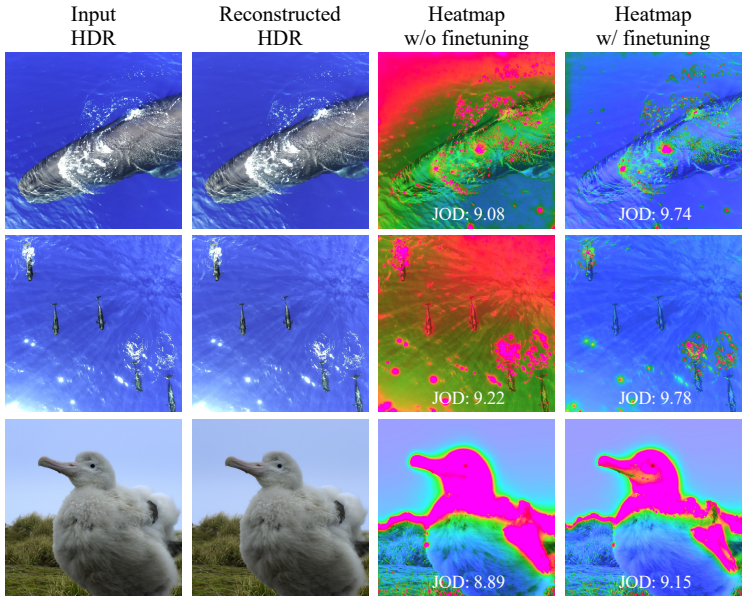

Figure 11: Representative VAE reconstruction failures on PU21-encoded HDR images, highlighted by ColorVideoVDP error maps.

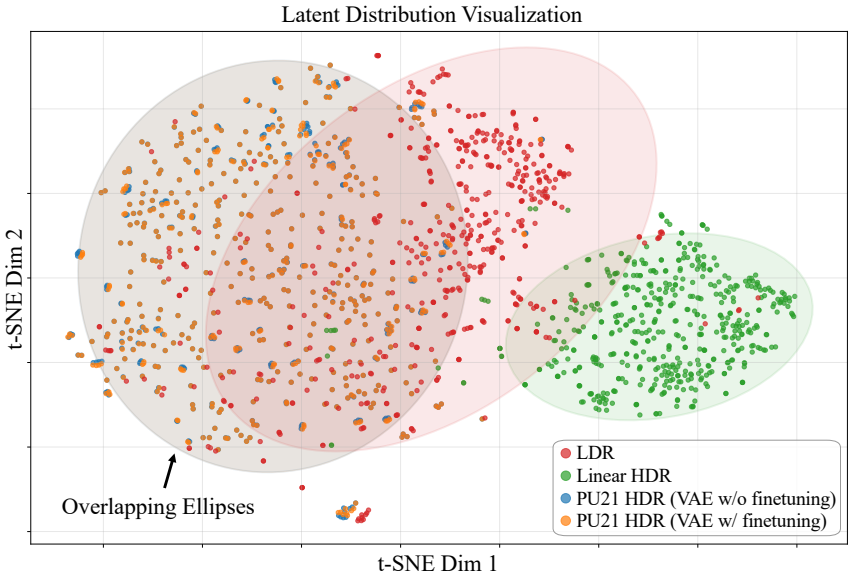

Figure 12: t-SNE visualization of VAE latent distributions under different encodings. PU21 substantially reduces the LDR-HDR mismatch, while VAE finetuning produces only marginal changes.

Effect of VAE finetuning on RAW-to-HDR Training a RAW-to-HDR model with the finetuned VAE (all else unchanged) does not yield consistent improvements. As shown in Table 5, aside from a modest gain in $Q^{\star}_{\mathrm{SSIM}}$, other measures slightly degrade, plausibly due to mild latent distribution shift that affects denoising and inpainting during reverse diffusion.

Summary Perceptually uniform encoding (*e.g.*, PU21) is the dominant factor facilitating alignment between LDR and HDR latent representations. Under the current data scale, VAE finetuning adds limited benefits, though it may be more valuable with larger and more diverse HDR datasets.

Table 5: Effect of VAE finetuning on RAW-to-HDR reconstruction.

Method	JOD ↑	$Q^{\star}_{\mathrm{PSNR}}$ ↑	$Q^{\star}_{\mathrm{SSIM}}$ ↑	$Q^{\star}_{\mathrm{LPIPS}}$ ↓	$Q^{\star}_{\mathrm{DISTS}}$ ↓
VAE w/ finetuning	7.34	22.3	**0.759**	0.197	0.111
VAE w/o finetuning (default)	**7.57**	**22.5**	0.741	**0.174**	**0.089**

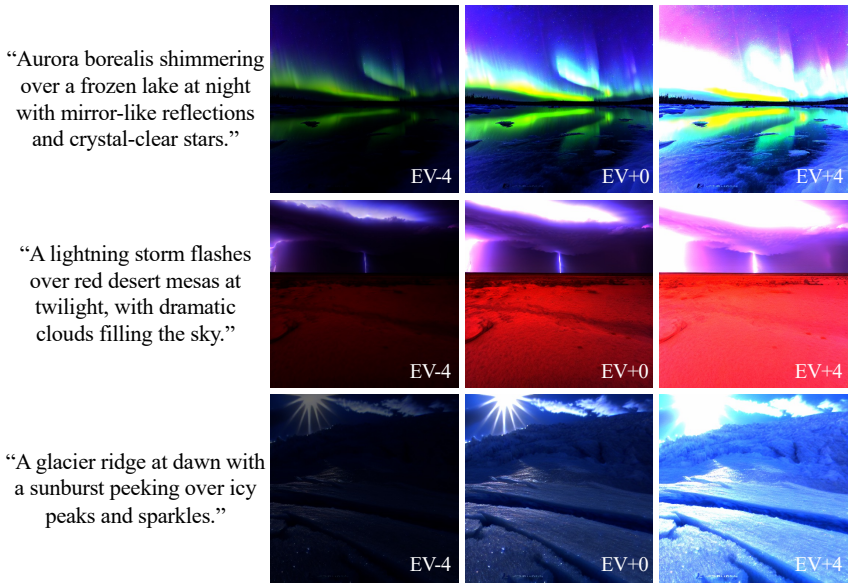

Figure 13: Sample HDR images generated by X2HDR with `SD-1.5` at different exposure levels.

F.5 Text-to-HDR with `SD-1.5`

To verify backbone generality (and support fair comparison with LEDiff and Bracket Diffusion), X2HDR is also instantiated with `SD-1.5`, using the same PU21 encoding and LoRA strategy. The `SD-1.5` setup uses LoRA rank 8, learning rate 1×10^{-5}, and batch size 32. Training is limited to 1,000 steps to avoid overfitting, and the trigger token `[PU21]` is prepended.

Additional visual results in Figure 13 show that highlights remain well separated at EV -4 and shadow details emerge at EV $+4$. Figure 14 further indicates that the `SD`-based variant achieves dynamic range comparable to the `FLUX`-based model, supporting the reliability of PU21-based adaptation across backbones.

F.6 RAW-to-HDR on Synthetic Data

To probe RAW-to-HDR behavior under controlled degradation, we construct a synthetic dataset using 512×512 HDR images with radial luminance gradients that decrease logarithmically from $4{,}000\,\mathrm{cd/m^2}$ at the center to $0.005\,\mathrm{cd/m^2}$ at the corners, instantiated across multiple chromatic channels (*e.g.*, yellow and red). Synthetic RAW inputs are simulated by applying a virtual exposure that clips values above $100\,\mathrm{cd/m^2}$ and adding Poisson-Gaussian noise calibrated to Sony A7R III.

Reconstructions in Figure 15 indicate that X2HDR effectively suppresses sensor noise, while competing techniques introduce structured artifacts. However, all methods struggle to recover ideal smooth gradients in heavily clipped regions, reflecting limited generalization to synthetic patterns absent from training. RawHDR [73] is excluded because it requires camera-specific metadata (*e.g.*, Bayer pattern and camera-to-RGB transformations) not defined in this synthetic setup.

F.7 LDR-to-HDR Image Reconstruction

We adapt X2HDR to single-image LDR-to-HDR reconstruction by conditioning the model directly on LDR inputs. We train it on 9,786 paired 512×512 LDR and HDR images from SingleHDR [35],

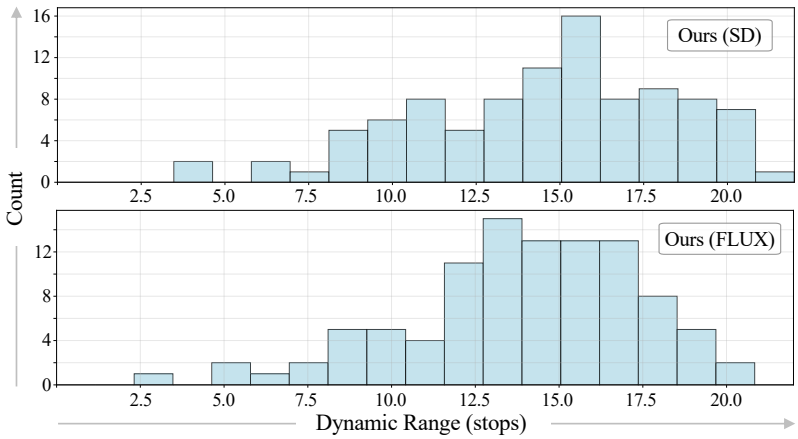

Figure 14: Distributions of effective dynamic range (in stops) over 100 generated HDR images. Both the `SD`-based and `FLUX`-based variants produce images with wide dynamic range.

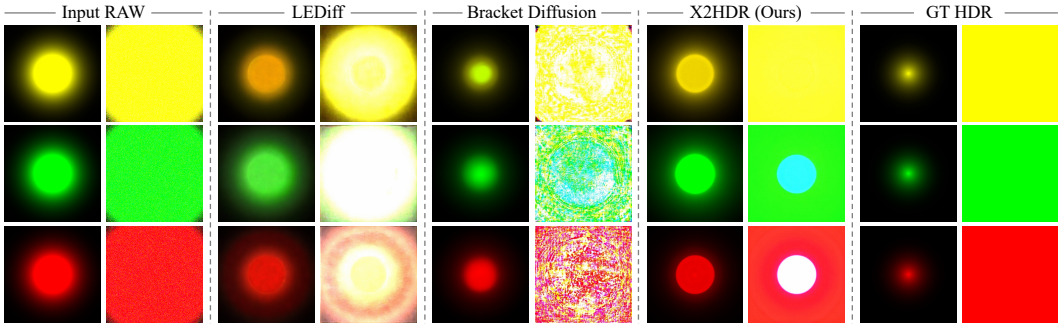

Figure 15: RAW-to-HDR reconstruction on synthetic data with controlled logarithmic luminance gradients. Rows correspond to different chromatic gradients (yellow, green, and red), and each method is visualized at two exposure levels. X2HDR strongly suppresses sensor noise, while LEDiff and Bracket Diffusion introduce structured artifacts. All methods remain challenged in recovering perfectly smooth gradients within severely clipped regions, indicating limited generalization to these synthetic patterns.

with VLM-generated captions for the HDR targets. The training setup follows that of RAW-to-HDR reconstruction, except that we reduce the number of training steps to 6,000 for the smaller dataset.

For evaluation, we randomly sample 50 images from the SingleHDR test set and compare X2HDR with LEDiff and Bracket Diffusion. The sampled images present challenging overexposed or underexposed regions. Visual results for all 50 test images are provided on the supplementary website. We focus on generative baselines rather than conducting an exhaustive comparison with earlier LDR-to-HDR methods, as prior work has established the advantages of generative approaches over conventional methods trained from scratch. As shown in Table 6, X2HDR achieves the best performance across all metrics among the evaluated generative methods. The visual comparisons in Figure 16 further show that X2HDR more faithfully reconstructs highlight structures, recovers low-light details, and produces fewer artifacts.

F.8 Perceptual Study

To complement the objective evaluations in Section 5, we conducted a controlled perceptual study on an HDR display to verify the perceptual gains obtained by X2HDR for both text-to-HDR generation and RAW-to-HDR reconstruction.

Stimuli For text-to-HDR, each trial compared a pair of HDR images generated from the same text prompt by two of four methods: LEDiff, Bracket Diffusion, and X2HDR instantiated with

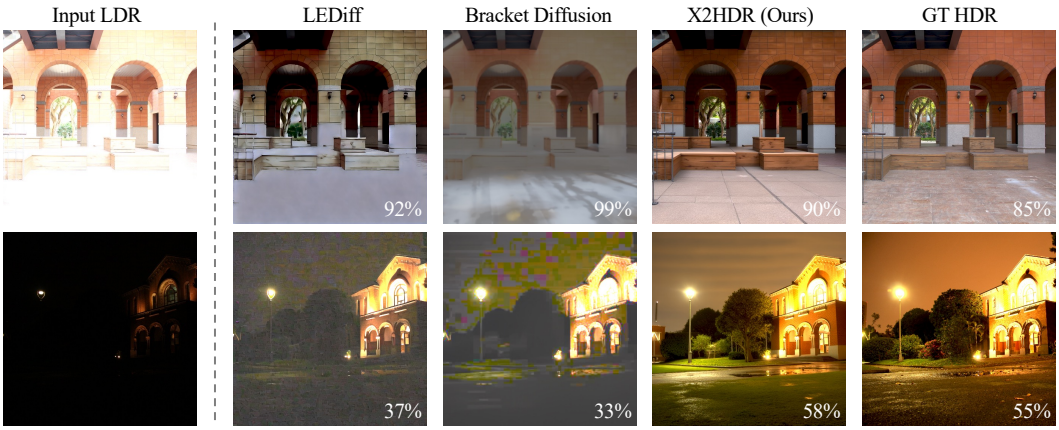

Figure 16: Visual comparison of LDR-to-HDR reconstruction. For the overexposed input (top), X2HDR recovers highlight structure more plausibly, while LEDiff and Bracket Diffusion produce visible distortions. For the underexposed input (bottom), X2HDR reveals more faithful shadow details with fewer artifacts.

Table 6: Quantitative results for LDR-to-HDR.

Method	JOD ↑	$Q^{\star}_{\mathrm{PSNR}}$ ↑	$Q^{\star}_{\mathrm{SSIM}}$ ↑	$Q^{\star}_{\mathrm{LPIPS}}$ ↓	$Q^{\star}_{\mathrm{DISTS}}$ ↓
LEDiff	5.74	17.7	0.569	0.252	0.130
Bracket Diffusion	5.86	18.4	0.641	0.272	0.160
X2HDR (Ours)	**6.47**	**19.3**	**0.680**	**0.187**	**0.101**

either `FLUX.1-dev` or `SD-1.5`. We randomly selected 20 prompts from the full prompt set. For RAW-to-HDR, each trial compared a pair of images drawn from two of six conditions: the input RAW capture, the GT HDR reference, and outputs from LEDiff, Bracket Diffusion, RawHDR, and X2HDR. We randomly sampled 20 test images for the study. All stimuli were resized to $1,280\times1,280$ via bilinear upsampling. To reduce bias from luminance shifts, we normalized each stimulus by mapping its median luminance to $8\,\mathrm{cd/m^2}$.

Apparatus and viewing conditions Stimuli were displayed on an ASUS ProArt Display PA32UCXR (32-inch, $3,840\times2,160$), a mini-LED HDR monitor supporting the Rec. 2020 color gamut. The measured peak luminance was $1,987\,\mathrm{cd/m^2}$ and the minimum luminance was $< 0.01\,\mathrm{cd/m^2}$. Participants viewed the display from approximately 80 cm ($\approx$76 pixels per degree) in a darkened room.

Procedure We recruited 26 participants (15 males and 11 females), aged between 20 and 30 years, all with normal or corrected-to-normal color vision. We employed a pairwise comparison protocol. On each trial, observers viewed two HDR images side by side and selected the one that appeared more natural, defined in the instructions as "closer to a high-quality photograph of the real world." For text-to-HDR evaluation, prompts were not shown since our adaptation is not designed to change text-image alignment of the base T2I model. Spatial left/right placement and temporal image order were randomized, and participants took as long as needed before responding via left/right keyboard. To improve rating efficiency, we used an active sampling strategy [44] to prioritize informative pairs; each participant completed 60 trials for text-to-HDR and 100 trials for RAW-to-HDR, requiring approximately 8 and 15 minutes, respectively.

Analysis Our pairwise comparison protocol used an active sampling strategy based on approximate message passing [44] to reduce comparisons while maintaining accurate perceptual scaling. Pairwise outcomes were converted to JOD scores using publicly available software[3] that performs maximum likelihood estimation under the Thurstone Case V model [59]. Because JOD is identifiable only up to an additive constant, we selected the origin of the JOD scale separately for each task. For

[3] `https://github.com/mantiuk/pwcmp`

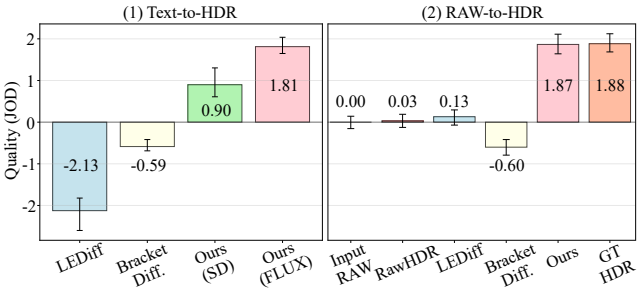

Figure 17: Quantitative results of the perceptual study, reported in the unit of JOD (higher is better). Bars show mean JOD across stimuli, and error bars denote 95% confidence intervals.

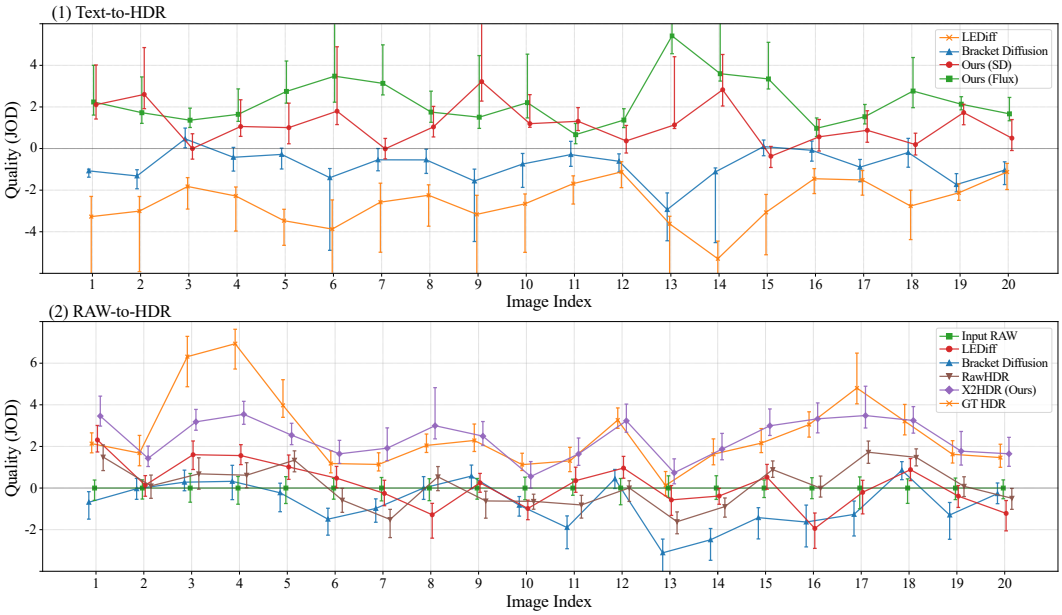

Figure 18: Per-image JOD scores from our perceptual study. The top and bottom panels report text-to-HDR and RAW-to-HDR results, respectively. Error bars denote 95% confidence intervals via bootstrapping. A small horizontal jitter is applied to separate methods for improved visibility.

RAW-to-HDR, we anchored the scale by assigning the input RAW image a JOD score of zero. As for text-to-HDR, where no natural reference condition exists, we centered the fitted scores by subtracting a constant so that the grand mean across methods is zero (*i.e.*, a purely conventional choice of origin). With the scale fixed, we summarize each method by reporting in Figure 17 its mean JOD across stimuli, together with 95% confidence intervals estimated via bootstrapping.

For text-to-HDR, X2HDR achieves the highest perceived quality, with the `FLUX`-based variant scoring 1.81 JOD, followed by the `SD`-based variant at 0.90 JOD. These results indicate a consistent preference for X2HDR over previous bracket-and-merge methods, and further suggest that backbone capacity (`FLUX.1-dev` vs. `SD-1.5`) notably impacts perceptual quality under our adaptation. For RAW-to-HDR, X2HDR achieves 1.87 JOD, nearly matching the GT HDR reference. In contrast, LEDiff and RawHDR yield only statistically insignificant improvements over the RAW input, while Bracket Diffusion performs worse with -0.60 JOD, which we believe is attributed to the resolution limitation under the current setting.

Per-image JOD scores are visualized in Figure 18 with 95% confidence intervals. In RAW-to-HDR, for extremely low-exposure cases (*e.g.*, the 3rd and 4th test scenes), the HDR references contain finer textural details than our reconstructions, resulting in a relatively larger perceptual gap. Otherwise, X2HDR is comparable to (and sometimes perceived as better than) GT.

Overall, the perceptual results corroborate the trends in Section 5, showing that X2HDR consistently delivers superior visual quality for both HDR generation and reconstruction tasks.

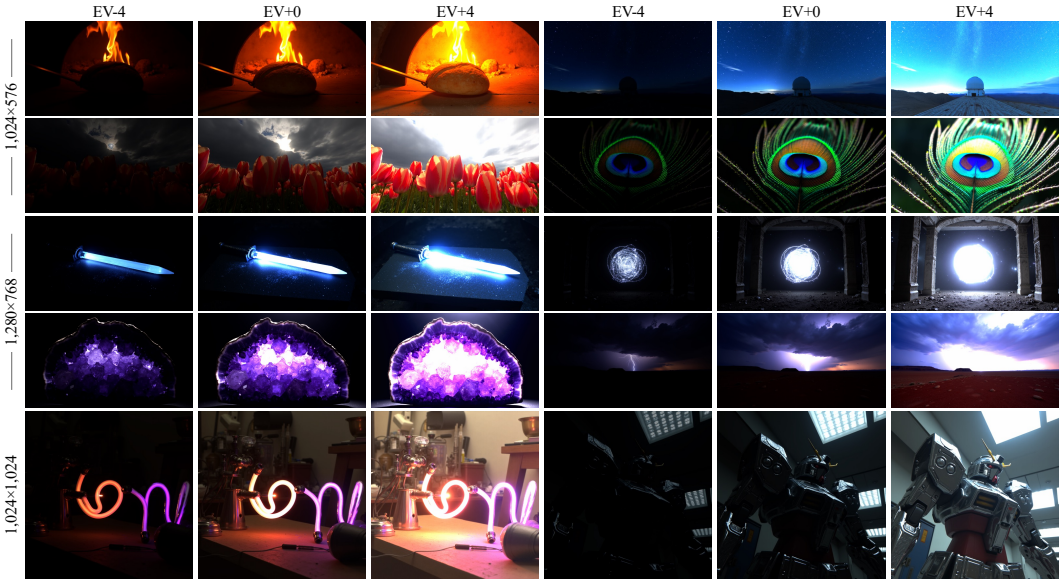

Figure 19: Multi-resolution HDR generation results. Examples are shown at three output resolutions, covering diverse lighting conditions and luminous phenomena (*e.g.*, artificial lights, fire, sunlight, and night skies). Each HDR image is visualized at EV -4/0/$+4$ to illustrate the effective dynamic range.

F.9 Higher-Resolution Results

To ensure a fair comparison, earlier experiments fix the output resolution to 512×512. In fact, X2HDR is built upon modern T2I backbones and has no inherent resolution constraint. Examples in Figure 19 demonstrate text-to-HDR synthesis at $1,024\times576$, $1,280\times768$, and $1,024\times1,024$ across diverse lighting conditions (*e.g.*, strong point emitters, outdoor daylight, and low-light astronomical scenes), with exposure-adjusted views showing preserved details and wide effective dynamic range. Similar generalizability extends to RAW-to-HDR reconstruction.

F.10 Single-RAW Multi-look HDR Authoring

Beyond "pure" reconstruction where text prompts are dropped, our RAW-to-HDR model can also be used as a one-to-many HDR authoring tool: conditioning the same RAW input on different prompts produces multiple HDR renderings with distinct global appearance and atmosphere (see Figure 20). Conventional workflows do so by applying tone-mapping and hand-crafted color grading in display space, which offer limited control over physically meaningful highlight behavior. In stark contrast, X2HDR performs prompt-driven tone manipulations in a physically grounded, perceptually uniform space, while preserving scene structure.

F.11 Discussion on Banding Effect

In text-to-HDR, we occasionally observe banding artifacts in a small fraction of generated HDR images (approximately 2 out of 100), predominantly in dark regions with smooth luminance gradients (see Figure 21). Under the default BF16 inference (16-bit representation with 8 bits for the exponent and only 7 bits for the mantissa), these smooth, low-luminance gradients are likely quantized into discrete bins, yielding visible banding. Using FP32 (23-bit mantissa) provides substantially finer resolution and effectively eliminates the artifacts. We therefore recommend FP32 inference when banding is visible, and retain BF16 as the default.

F.12 On the Use of Q-Eval-100K in Text-to-HDR

To validate the use of Q-Eval-100K [72] for evaluating HDR image quality and text-image alignment in text-to-HDR, we conduct a sanity check on 512 aligned (LDR, HDR) pairs. Captions are generated by `Gemini-2.0-Flash`. All HDR images are rescaled to $L_{\text{peak}} = 4{,}000\,\text{cd/m}^2$ and PU21-encoded

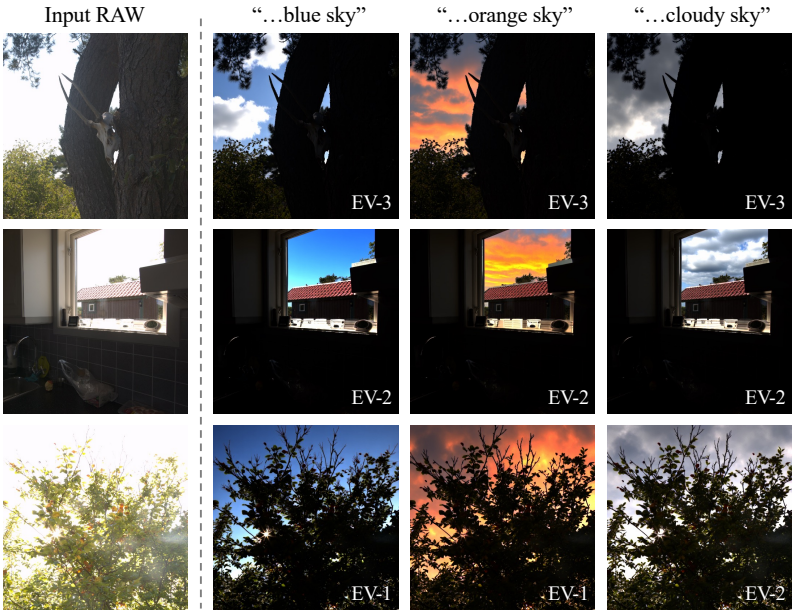

Figure 20: Single-RAW multi-look HDR authoring with text guidance. Given the same RAW input (left), our RAW-to-HDR model generates multiple HDR renderings conditioned on different prompts (*e.g.*, "blue sky," "orange sky," and "cloudy sky"), producing diverse global appearance and atmosphere while preserving scene structure.

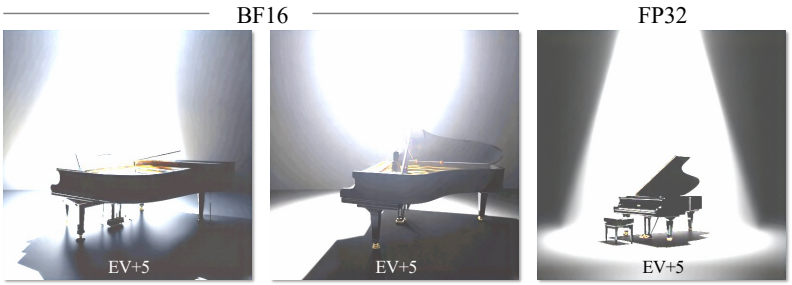

Figure 21: Banding artifacts induced by BF16 inference in text-to-HDR generation. When rendered at EV$+5$, quantization in smooth, low-luminance gradients produces visible banding (left/middle), whereas FP32 inference preserves subtle gradients and eliminates the artifacts (right).

before evaluation. Scores for LDR and PU21-encoded HDR images show strong cross-domain consistency: Pearson correlations are 0.933 for image quality and 0.830 for text-image alignment. Mean scores over the 512 pairs are 0.611 vs. 0.587 (image quality) and 0.736 vs. 0.695 (text-image alignment) for LDR vs. PU21-encoded HDR images, respectively, suggesting reasonable transfer of Q-Eval-100K (see also the score distributions in Figure 22).

Nonetheless, we observe a consistent bias in favor of LDR images, even when matched (LDR, HDR) pairs are perceptually equivalent. Therefore, Q-Eval-100K is not reported in settings that mix LDR and HDR outputs, where it could be misleading.

G Limitations

For text-to-HDR, our current training data are dominated by natural photographs, so the resulting model may generalize less reliably to out-of-domain styles (see the cartoons in Figure 23). For RAW-to-HDR, X2HDR can still fail in extremely underexposed or overexposed regions, occasionally introducing implausible hallucinations or local detail inconsistencies (see also Figure 23). In addition, the current X2HDR is display-agnostic, *i.e.*, it does not condition on target peak luminance or other

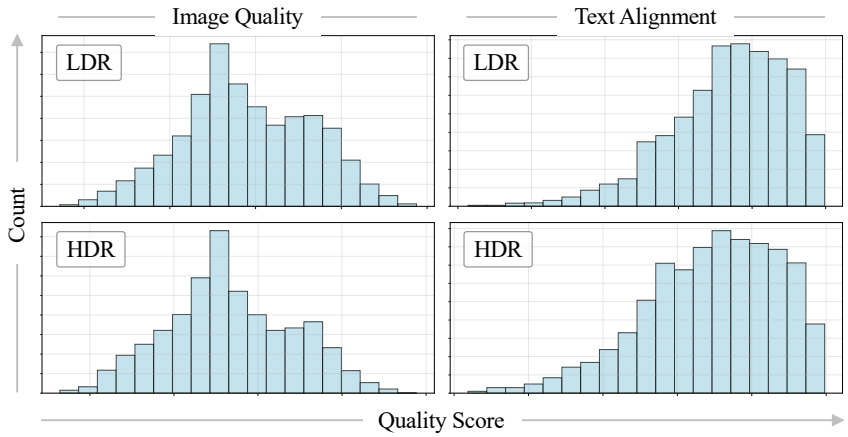

Figure 22: Q-Eval-100K score distributions for 512 aligned (LDR, HDR) pairs.



Figure 23: Limitations. Text-to-HDR: performance may degrade for out-of-domain styles, leading to less realistic appearance. RAW-to-HDR: implausible or inconsistent hallucinations may occur for large regions that are severely underexposed or overexposed.

device characteristics, and provides only limited support for controllable HDR image generation and interactive HDR editing.

H Future Work

First, for text-to-HDR, our training data are dominated by natural scenes, which can reduce reliability to out-of-domain styles (*e.g.*, cartoons in Figure 23). Style-robust finetuning and data augmentation, such as stylized HDR assets, synthetic relighting, and domain-mixing curricula, may improve generalization while preserving HDR plausibility. For RAW-to-HDR, failures in extremely underexposed or overexposed regions sometimes manifest as implausible or inconsistent local structure. This could be mitigated by incorporating perceptual objectives (*e.g.*, DISTS), locality-aware constraints, and uncertainty- or saturation-aware conditioning that explicitly prevents implausible extrapolation.

Second, because X2HDR adapts pretrained T2I models to HDR synthesis by LoRA finetuning, it is inherently modular and can be combined with complementary advances in the T2I ecosystem. A natural next step is to integrate X2HDR with controllable generation techniques (*e.g.*, ControlNet [68]) to unlock controllable HDR synthesis, and with emerging instruction- and context-based image editing models (*e.g.*, `FLUX.1-Kontext`) to support HDR-aware editing. Beyond the RAW-to-HDR setting studied here, our idea can be extended to single-image LDR-to-HDR reconstruction by conditioning on LDR inputs, and to multi-exposure inputs for HDR fusion and reconstruction.

Third, extending X2HDR from still images to videos is an important direction. Applying the same perceptually uniform encoding and adaptation strategy to video diffusion models could spur text-to-HDR video generation and RAW-to-HDR video reconstruction, with explicit mechanisms for temporal consistency (*e.g.*, motion-aware conditioning, recurrent features, or temporal regularizers) to prevent flickering and exposure instability. Meanwhile, conditioning generation on target display capabilities (*e.g.*, peak luminance, local-dimming behavior, and color volume) and incorporating power/thermal constraints during generation is of practical importance, which encourages content that "spends" dynamic range where it is most perceptually effective while respecting energy budgets.

Fourth, progress is currently constrained by the limited availability of large-scale, publicly accessible HDR data. Although our experiments suggest that freezing the pretrained VAE is sufficient for strong performance, future work could investigate VAE finetuning or distillation using larger HDR corpora to further improve reconstruction fidelity, particularly in extreme luminance ranges, where quantization and saturation effects are most salient.

Finally, HDR generation would benefit from more HDR-native evaluation and optimization. We currently use Q-Eval-100K [72] as a practical proxy for no-reference HDR assessment, but such evaluators are not designed for HDR-specific perceptual factors (*e.g.*, highlight naturalness, extreme luminance behavior, and wide-range local contrast). Developing dedicated no-reference HDR quality models is therefore a high-impact direction, unlocking scalable benchmarking, training-time filtering, and reinforcement-learning-style optimization.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction accurately summarize the paper's main contributions, including the proposed perceptually uniform HDR representation, the adaptation strategy for pretrained diffusion models, and the supported text-to-HDR and RAW-to-HDR tasks.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper discusses limitations, including reduced generalization to out-of-domain styles for text-to-HDR, possible failures in extremely underexposed or overexposed RAW inputs, the lack of display-specific conditioning, and limited support for controllable HDR generation and interactive editing.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: The paper does not present theoretical results in the form of theorems, lemmas, or formal proofs.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The paper provides the key information needed to reproduce the main experiments, including the HDR/RAW preprocessing pipeline, PU21 encoding and inverse transform, LoRA-based finetuning setup, training objectives for both text-to-HDR and RAW-to-HDR, evaluation datasets, baseline methods, metrics, and perceptual study protocol.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: The current submission does not provide open access to the code or complete reproduction scripts. The paper nevertheless describes the main methodological and experimental details, including preprocessing, training objectives, evaluation datasets, metrics, baselines, and perceptual study protocol.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: The paper specifies the experimental settings needed to understand the reported results, including the preprocessing of HDR/RAW inputs, the FLUX.1-dev and SD-1.5 backbones, LoRA-based finetuning objectives, evaluation datasets and test splits, baseline methods, image resolutions, evaluation metrics, and perceptual study protocol. Additional implementation details can be provided in the supplemental material.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: The paper reports statistical uncertainty for the controlled perceptual study, which directly supports the main claims on perceived HDR quality.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [No]

   Justification: The paper reports inference-time runtime and peak memory for the text-to-HDR comparison and provides implementation details for several training/evaluation settings in the appendix. However, it does not fully specify the compute workers, training time, memory usage, or total compute budget for each experiment.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research conforms to the NeurIPS Code of Ethics. The work focuses on HDR image generation and reconstruction, uses established public datasets and evaluation protocols, and reports a controlled perceptual study without collecting sensitive personal information. The submission also preserves author anonymity.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.

- If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [No]

    Justification: The paper discusses the technical benefits and limitations of X2HDR, including improved HDR generation and reconstruction quality, reduced architectural complexity, and possible failures under out-of-domain styles or extreme exposure conditions. However, it does not explicitly discuss broader societal impacts, such as potential positive applications in HDR content creation and photography, or potential negative uses of improved generative image quality.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [No]

    Justification: The paper studies HDR image generation and reconstruction, including a text-to-HDR image generation model, but it does not explicitly describe safeguards for responsible release, such as access restrictions, usage guidelines, safety filtering, or misuse monitoring. The paper also does not claim to release a high-risk dataset or model.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.

- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [No]

    Justification: The paper properly credits the main existing assets used in the experiments, including pretrained diffusion backbones, datasets, baselines, and evaluation tools. However, it does not explicitly state the licenses or terms of use for all existing assets, such as pretrained models, datasets, and software packages.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not release new assets such as a new dataset, codebase, or model checkpoint. It introduces a method for adapting pretrained diffusion models to HDR generation and reconstruction, and evaluates the method using existing datasets, baselines, and evaluation tools.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [No]

Justification: The paper includes a controlled perceptual study with human participants and reports the apparatus, viewing conditions, participant demographics, task procedure, number of trials, and analysis protocol. However, it does not include the full participant instructions, screenshots of the interface, or details about compensation.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [Yes]

Justification: The paper reports the controlled perceptual study protocol and states that participant risks were disclosed and that IRB approval or an equivalent institutional review was obtained, while preserving submission anonymity.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: The core method development in this research does not involve LLMs as an important, original, or non-standard component. The paper uses text prompts for text-to-HDR generation, but the proposed method is based on adapting diffusion models to HDR generation and reconstruction through perceptually uniform representations and LoRA finetuning.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.