FaceProbe: Recovering HDR Environment Map via Masked Diffusion and Physical Preference

Peng Zhao[1,2] William Jing[1] Juyan Ba[4] Kairui Feng[2,3,*] Xuanhong Chen[1,*]

[1]Shanghai Jiao Tong University, Shanghai 200241, China

[2]Shanghai Innovation Institute, Shanghai 200231, China

[3]Shanghai Research Institute for Intelligent Autonomous Systems & State Key Laboratory Autonomous Intelligent Unmanned Systems, Tongji University, Shanghai 201210, China

[4]Shenzhen University, Shenzhen 518060, China

[*]Corresponding authors: Kelvinfkr@tongji.edu.cn;chenxuanhong@sjtu.edu.cn

Abstract

Recovering 360° HDR environment maps from unconstrained portraits is a fundamental yet ill-posed challenge in inverse rendering. Existing approaches typically suffer from a critical trade-off: traditional facial reflection-based methods yield low-fidelity, blurry results that lack practical utility, while recent generative models often resort to unconstrained panoramic hallucination (i.e., synthesizing environments with no physical grounding), which leads to severe overfitting and inaccurate lighting. In this paper, we present FaceProbe, an inpainting-driven framework that treats the human face as an observational light probe carrying physically produced illumination evidence. Instead of synthesizing the environment from scratch, we reformulate the task as a masked completion and outpainting problem. By preserving the peripheral context of the input image and utilizing the portrait as a structural condition, our architecture leverages the generative priors of Diffusion Transformers (DiT) to extend sparse observations into high-fidelity, sharp 360° HDR panoramas. To further eliminate color shifts and highlight inaccuracies, we introduce OLAT-GRPO, a post-training alignment mechanism based on Group Relative Policy Optimization. By subjecting the generated maps to image-based relighting on One-Light-At-a-Time (OLAT) datasets, we evaluate their physical validity through relighting consistency. This allows us to explicitly prune inconsistent denoising paths and align the model with multi-objective rewards, ensuring energy conservation and spectral accuracy. Extensive evaluations demonstrate that FaceProbe achieves a paradigm-shifting improvement. Most notably, in the critical portrait relighting task, our method reduces estimation errors by 54.4% in Angular Error and 48.5% in si-RMSE over state-of-the-art methods, delivering a robust and highly usable solution for photorealistic rendering and world model applications.

1 Introduction

High-Dynamic-Range (HDR) environment map estimation aims to recover panoramic radiance from sparse visual cues and is crucial for inverse rendering and human-centric world models. While recent methods [1, 2, 3, 4, 5, 6, 7] achieve visual plausibility, they often degenerate into unconstrained panorama generation. By prioritizing aesthetics over physics, they suffer from severe energy distribution mismatches and fail to produce consistent relighting effects. To address this, we delve into the subtle reflection cues within human portraits to recover high-fidelity HDR maps that strictly conform to true energy distributions, ensuring physically accurate illumination.

To date, existing technical routes for environment map estimation primarily bifurcate into traditional reflection-based inverse rendering and recent direct generative paradigms. Despite their respective

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

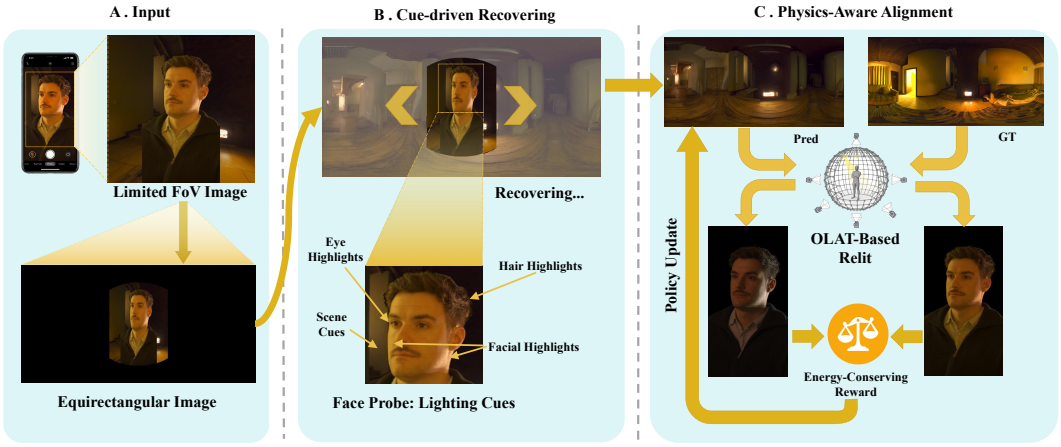

Figure 1: **Overview of our proposed lighting estimation framework.** (A) Input: A limited Field-of-View (FoV) portrait image is projected into an equirectangular space. (B) Cue-driven Recovering: Authentic facial illumination cues (e.g., corneal reflections, hair sheen, and skin specularities) are explicitly conditioned via a dedicated reference pathway to guide the recovery of the unobserved 360° environment map. (C) Physics-Aware Alignment: Both the predicted and ground truth environment maps are applied to an OLAT-based relighting operator. An energy-conserving reward evaluates the relit faces to formulate a policy update, strictly aligning the generative recovery process with real-world physical reflection laws.

merits, these trajectories are inherently hindered by three critical flaws: 1) **Information Bottleneck**: Traditional approaches [1, 2, 3, 4, 5, 6] rely exclusively on limited facial specularities (e.g., corneal reflections), discarding visible peripheral context and consequently yielding blurry, low-resolution estimations; 2) **Unconstrained Hallucination**: Recent generative models [7] directly synthesize full 360° panoramas from cropped portraits. Without anchoring to the observable background, they hallucinate unobservable regions with no physical basis, leading to severe structural mismatches. 3) **Illumination Misalignment**: Most data-driven frameworks [1, 8, 7] are optimized purely via statistical noise-matching. The absence of explicit physical feedback inevitably causes spectral shifts and inaccurate high-frequency highlights.

To bridge the gap between generative plausibility and physical rigor, we present **FaceProbe** (Fig. 1), a specialized framework designed to distill high-fidelity HDR environment maps from sparse facial cues. Unlike existing methods that attempt to hallucinate full panoramas from isolated portraits, our approach reformulates the task as a joint inpainting and outpainting problem. By preserving the visible peripheral context of the input image as a structural scaffold, we introduce explicit boundary conditions that effectively constrain the solution space. This masked learning paradigm significantly reduces the ill-posed complexity of inverse rendering, yielding environment maps that are not only visually sharp but also structurally grounded in the observed scene.

However, standard diffusion models are optimized purely via statistical noise matching and lack a closed-loop physical feedback mechanism. Intuitively, one might enforce relighting consistency using direct pixel-wise or feature-wise alignment losses (e.g., MSE or perceptual loss). Yet, imposing such rigid spatial constraints severely disrupts the generative prior, inevitably triggering severe visual degradation such as pronounced checkerboard artifacts. To fundamentally address this, we introduce OLAT-GRPO coupled with an adversarial loss. We subject the synthesized environment maps to image-based relighting using One-Light-At-a-Time (OLAT) data. Instead of rigid pixel matching, we formulate physical alignment as a preference optimization problem, utilizing GRPO to explicitly penalize non-physical denoising trajectories. Concurrently, the integrated adversarial loss acts as an indispensable regularizer, strictly anchoring the fine-tuned distribution to the natural HDR manifold. This joint strategy successfully bypasses pixel-level artifacts, ensuring that the synthesized environments are both strictly energy-consistent and visually flawless.

Our primary contributions are summarized as follows:

- We propose FaceProbe, reformulating HDR environment map estimation as a masked completion task. By anchoring generation to contextual boundaries and facial cues, it significantly reduces the ill-posed complexity of inverse rendering.
- We introduce OLAT-GRPO coupled with an adversarial regularizer. By framing physical consistency as a preference alignment problem, this joint strategy successfully bypasses the severe checkerboard artifacts caused by naive pixel-alignment losses, ensuring both energy conservation and high-fidelity photorealism.
- Extensive evaluations—including matched-input baseline comparisons, facial-cue ablations, and an independent third-party benchmark—demonstrate state-of-the-art performance: our method reduces angular error by 54.4% and si-RMSE by 48.5% in downstream relighting tasks, providing a physically accurate and highly usable solution for inverse rendering applications.

2 Related Work

Single-Image Environment Map Estimation. The trajectory of environment map estimation from sparse visual observations can be broadly summarized into three distinct evolutionary phases. *Unconstrained Direct Generation.* Initially, to move beyond cumbersome hardware captures [9, 10, 11], methods compressed radiance into parametric or volumetric models [12, 13, 14, 15, 16, 17, 18, 19, 20]. The field then shifted to direct generative synthesis [21, 22, 23, 24, 25, 26, 8], sometimes inserting virtual proxies or relying on intrinsic material decomposition [1, 2, 3, 4, 5, 6, 27, 28, 29]. Operating essentially by guessing unseen regions without reliable physical anchors, they inevitably fail to produce physically correct energy distributions. *Portrait-Guided Regression.* To introduce physical grounding, subsequent research began exploiting human faces as natural light probes. Alongside advances in fine-grained facial attribute and appearance modeling [30, 31], pioneering works like facelight [7] extract lighting cues directly from facial appearances. However, these attempts regress the illumination into highly compressed representations, such as an extremely blurry chrome ball or a low-resolution grid, completely discarding high-frequency reflections required for realistic rendering. *Masked-Based Learning (Ours).* Our framework fundamentally breaks from the previous paradigms. Rather than hallucinating panoramas from scratch or regressing to blurry proxies, we formulate the estimation as a masked completion task. The critical difference lies in our explicit utilization of native facial reflection cues alongside the visible background as rigorous boundary conditions. This effectively constrains the generation space, drastically reducing the ill-posed complexity to recover high-resolution, physically accurate panoramas.

Reinforcement Learning for Generative Alignment. Standard diffusion models optimize for noise-matching, struggling with non-differentiable physical laws. DDPO [32] pioneered the formulation of diffusion denoising as a Markov Decision Process. Recently, Group Relative Policy Optimization (GRPO) [33] eliminated the compute-heavy value model, and Flow-GRPO [34] successfully adapted this mechanism to continuous generative architectures. We uniquely adapt GRPO for *physical inverse rendering*. By defining multi-objective rewards based on relighting accuracy and energy conservation, our framework explicitly forces the diffusion model to satisfy rigorous physical constraints, bridging generative AI and physically grounded lighting estimation.

3 Method

We formulate single-image lighting estimation as a cue-driven masked recovering problem in the 360° omnidirectional domain.

Notation and System Definition. Given a limited Field-of-View (FoV) LDR portrait image denoted as $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$, our primary goal is to reconstruct the unobserved global illumination. To explicitly define the knowns and unknowns, we first project the visible background of $\mathbf{I}$ into a 360° equirectangular projection (ERP) space, yielding a partial environment map $\mathbf{E}_{masked} \in \mathbb{R}^{H' \times W' \times 3}$. This projection is accompanied by a binary visibility mask $\mathbf{M}$, where $\mathbf{M}_{i,j} = 0$ indicates the missing unobserved regions to be hallucinated, and 1 indicates the visible camera observations. Instead of relying solely on the generative prior to blindly fill in the mask $\mathbf{M}$, we treat the human portrait within $\mathbf{I}$ as an observational light probe: natural facial appearance inherently captures real-world illumination responses (e.g., ocular and corneal highlights [35, 36], skin specularities, shading gradients, and hair

responses). We explicitly condition the network on these facial cues, denoted as a conditional tensor $\mathbf{C}_{face}$, via a dedicated reference pathway. Concretely, $\mathbf{C}_{face}$ is the observed portrait in the same ERP coordinate system—the full portrait encoded by a separate conditioning stream (Sec. 3.1), not a hand-crafted highlight map or a thresholded reflection tensor.

What "physical" means in this paper. The term is operational, not analytic: we neither invert a facial BSSRDF nor assume a known BRDF. A portrait does contain physically produced illumination responses which, being linear in incident illumination, constrain HDR recovery and can be checked by light-transport superposition (Sec. 3.3) without separating diffuse, specular and subsurface components (Appendix B).

The Recovering Objective. Formally, our objective is to learn a generative mapping function $\mathcal{G}_\theta$ (parameterized by a Diffusion Transformer) that leverages the observed background and facial cues to reconstruct the complete 360° HDR environment map:

$$\mathbf{E}_{pred} = \mathcal{G}_\theta(\mathbf{E}_{masked}, \mathbf{M}, \mathbf{C}_{face}), \tag{1}$$

where $\mathbf{E}_{pred} \in \mathbb{R}^{H' \times W' \times 3}$ is the predicted full-resolution HDR environment map.

HDR Target Representation. Standard VAEs cannot faithfully encode raw HDR radiance due to their bounded $[-1, 1]$ input requirement. To bridge this numerical range mismatch without losing radiometric ordering, we compress the ground-truth HDR panorama $\mathbf{E}$ (clipped to a maximum threshold $T_{max} = 10,000$) into a numerically stable bounded space via a stabilized logarithmic transformation:

$$\mathbf{E}_{norm} = \frac{\log_{10}(\mathbf{E} + \epsilon) - \log_{10}(\epsilon)}{\log_{10}(T_{max} + \epsilon) - \log_{10}(\epsilon)}, \tag{2}$$

where the offset $\epsilon = 10^{-2}$ robustly handles near-zero regions. This tensor is then linearly mapped to the VAE input interval: $\mathbf{E}_{target} = 2\mathbf{E}_{norm} - 1 \in [-1, 1]$. Crucially, rather than predicting pixel-level representations directly, we map this bounded HDR panorama into the latent space via the frozen VAE encoder $\mathcal{E}$ to formulate a single, unified generative target: $\mathbf{z} = \mathcal{E}(\mathbf{E}_{target})$. While log-normalization resolves the unbounded numerical range, the remaining domain adaptation is efficiently handled through Low-Rank Adaptation (LoRA) on the generative prior.

3.1 Network Architecture

As shown in Figure 2, we build our backbone upon the pre-trained Wan2.1-Fun-14B-Control [37] model, operating entirely within the VAE latent space. The network processes three inputs: a masked image for ambient context, a binary mask, and a reference image providing explicit facial lighting cues. Specifically, the noisy latent, masked latent, and binary mask are concatenated along the channel dimension and fed into the native patch embedding layer. To explicitly decouple the conditional guidance role of the reference image from the spatial denoising targets, its latent representation is processed by a newly introduced Reference Patch Embedding layer. The patch tokens from both pathways are then concatenated along the sequence dimension and forwarded through the DiT blocks. To ensure parameter-efficient fine-tuning and mitigate catastrophic forgetting, we freeze the pre-trained DiT weights and apply Low-Rank Adaptation (LoRA) exclusively to the self-attention layers. Finally, the token sequence corresponding to the reference image is truncated, and the remaining sequence is decoded by the VAE to output the final HDR environment map.

3.2 Mask Learning and Face Cues

Predicting complex out-of-view radiance from a narrow FoV is mathematically underdetermined. Consequently, unconstrained direct generation [21, 22, 23, 24, 25, 26, 8] and methods [1, 2, 3, 4, 28, 29] relying on virtual object proxies inevitably degrade into "probabilistic hallucinations." To address this, we adopt a masked learning paradigm driven by explicit face cues. The partial background $\mathbf{E}_{masked}$ provides continuous ambient context, while the reference portrait $\mathbf{C}_{face}$ acts as an observational light probe indicating primary source directions and intensities. As detailed in Sec. 3.1, we instantiate $\mathcal{G}_\theta$ to fuse these conditions and optimize the base model via continuous-time Flow-Matching using the Mean Squared Error (MSE) loss (Eq. 3):

$$\mathcal{L}_{base} = \mathbb{E}_{\mathbf{z},\epsilon,t\sim\mathcal{U}(0,1)} \left[\| \mathcal{G}_\theta(\mathbf{z}_t, t, \mathbf{E}_{masked}, \mathbf{M}, \mathbf{C}_{face}) - (\epsilon - \mathbf{z}) \|_2^2 \right], \tag{3}$$

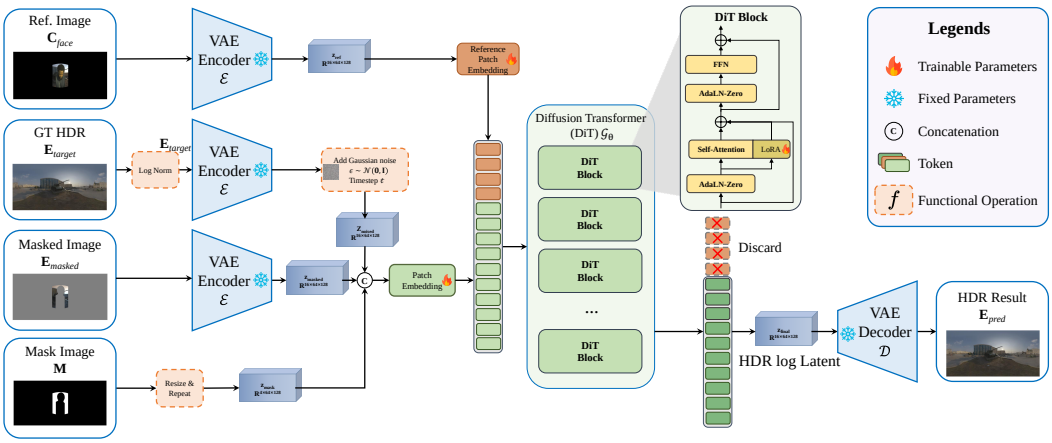

Figure 2: Architecture of FaceProbe. We formulate HDR environment map estimation as a masked completion task. A DiT leverages facial reflection cues and the visible peripheral background as strict boundary conditions to extrapolate high-fidelity, physically consistent panoramas.

where $\mathbf{z}$ represents the target latent of the ground-truth HDR environment map, $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ is the sampled Gaussian noise, $t \in [0, 1]$ denotes the continuous timestep, and $\mathbf{z}_t = (1-t)\mathbf{z} + t\epsilon$ defines the linear probability flow trajectory.

3.3 Physics-Aware Alignment via OLAT-GRPO

While the aforementioned cue-driven generation successfully learns structural panoramic distributions, it inherently lacks a physical feedback mechanism. The network merely mimics statistical pixel patterns without understanding the underlying light transport, inevitably leading to physical inconsistencies in downstream tasks, such as color shifts and inaccurate highlights.

The Virtual Light Stage Simulator. To bridge this gap, we introduce an innovative alignment strategy: **OLAT-GRPO**. We employ the One-Light-At-a-Time (OLAT) Image-Based Relighting process as a "Virtual Light Stage" simulator. For a predicted environment map $\mathbf{E}_{pred}$ (the log-normalized output of Eq. (2), inverted back to linear HDR before use), the relit portrait $\mathbf{I}_{relit}$ is synthesized via standard linear light transport superposition across $K = 331$ captured LightStage basis images:

$$\mathbf{I}_{relit} = \sum_{k=1}^{K} w_k \mathbf{B}_k, \quad \text{where} \quad w_k = \sum_{p} \mathbf{S}_{k,p} \mathbf{E}_{pred}(p), \tag{4}$$

where $\mathbf{B}_k$ denotes the real captured portrait basis under the k-th directional light, and $\mathbf{S}_{k,p}$ represents the solid-angle integration weight for angular support k over pixel p. Importantly, because each basis $\mathbf{B}_k$ is physically captured from real human subjects, it intrinsically contains all complex transport phenomena—including multi-layer skin reflectance, cast shadows, eye/hair reflections, and localized subsurface scattering. This formulation completely bypasses the need for simplified analytic BRDF models.

Reward Function Formulation. For a given input, the policy π_θ samples a group of G candidate HDR maps $\{\mathbf{E}_{pred}^{(1)}, \ldots, \mathbf{E}_{pred}^{(G)}\}$. Each candidate is fed into the Virtual Light Stage to render a relit face. Based on physical constraints, we formulate the total reward $r^{(g)}$ as a weighted sum of four components:

$$r^{(g)} = 0.2 r_{pixel}^{(g)} + 0.35 r_{hist}^{(g)} + 0.15 r_{color}^{(g)} + 0.3 r_{adv}^{(g)}, \tag{5}$$

where r_{pixel} (Log-MAE) penalizes highlight misplacements, r_{hist} (Histogram Matching) enforces global energy conservation, r_{color} (Cosine Similarity) prevents color shifts, and r_{adv} acts as an adversarial score to preserve high-frequency panoramic realism.

Advantage and Policy Update. To circumvent the computationally prohibitive value network, GRPO evaluates the relative quality of candidates within the sampled group. The advantage function $\hat{A}^{(g)}$ is

computed by standardizing the rewards:

$$\hat{A}^{(g)} = \frac{r^{(g)} - \mu_r}{\sigma_r + \epsilon}, \quad \text{where} \quad \mu_r = \frac{1}{G}\sum_{i=1}^{G} r^{(i)}, \quad \sigma_r = \sqrt{\frac{1}{G}\sum_{i=1}^{G}(r^{(i)} - \mu_r)^2}. \tag{6}$$

The policy is then optimized by maximizing the clipped surrogate objective, augmented with a Kullback-Leibler (KL) divergence penalty to prevent excessive deviation from the pre-trained reference model π_{ref}:

$$\mathcal{L}_{GRPO} = -\frac{1}{G}\sum_{g=1}^{G}\left[\min\left(\rho_t^{(g)}\hat{A}^{(g)}, \text{clip}\left(\rho_t^{(g)}, 1-c, 1+c\right)\hat{A}^{(g)}\right)\right] + \beta_{KL}\mathbb{D}_{KL}(\pi_\theta \| \pi_{ref}), \tag{7}$$

where $\rho_t^{(g)}$ is the importance sampling ratio, c is the clipping threshold, and β_{KL} controls the KL penalty magnitude.

4 Experiments

Dataset. We construct a relighting dataset using OLAT bases of 120 subjects and 3,527 HDR maps (private captures and PolyHaven [38]). To prevent distribution memorization, we enforce a strictly disjoint split: 100 subjects and 3,000 maps yield 300,000 training pairs, while 20 unseen subjects and 527 unseen maps provide 10,540 pairs exclusively for testing. To definitively eliminate training-distribution familiarity, we quantitatively evaluate on an independent third-party benchmark of 105 commercially sourced HDR scenes. Additionally, we conduct qualitative assessments on arbitrary in-the-wild portraits. These inherently lack ground truth—as simultaneous in-the-wild 360° HDR capture is practically infeasible—serving as a rigorous real-world generalization test (Appendix E). Our primary scope remains controlled image-based relighting (Appendix B).

Baselines and Evaluation Protocol. For fairness, we compare against official open-source implementations of DiffusionLight [1], StyleLight [24], and LuxDiT [8] (excluding works with unavailable code). This is a zero-shot comparison under official inference settings. We do not retrain baselines on our data, as modifying their architectures for our dual-condition input would invalidate the evaluation of their officially released methods. All metrics (si-RMSE, Angular Error, N-RMSE, PSNR, SSIM) are computed using the official `DiffusionLight-evaluation` codebase. To ensure input fairness, we include matched-input comparisons providing baselines with the identical person-containing images (Appendix C). Under this protocol, we comprehensively evaluate across rendered spheres, 360° HDR panoramas, downstream portrait relighting, an independent third-party benchmark dataset, qualitative in-the-wild lighting estimation, and extensive ablations.

4.1 Evaluation on Rendered Spheres

Quantitative Results. Our framework outperforms all baselines across all metrics (Table 1). Reductions in Angular Error and si-RMSE reveal that unconstrained priors (e.g., DiffusionLight, LuxDiT) suffer from severe spatial misalignment and structural hallucination. Furthermore, under our fixed-exposure linear protocol, Normalized RMSE explicitly captures absolute radiometric differences. This confirms that leveraging face probes and OLAT-GRPO energy conservation transforms an underdetermined guessing game into a constrained physical recovery, guaranteeing precise directional mapping and global exposure consistency.

Table 1: Quantitative evaluation of relighting quality on virtual spheres.

Method	si-RMSE			Angular Error			Normalized RMSE		
	Diffuse	Matte	Mirror	Diffuse	Matte	Mirror	Diffuse	Matte	Mirror
DiffusionLight [1]	0.7106	2.4104	17.7523	6.9903	7.4993	10.6625	0.4394	0.4410	7.6355
StyleLight [24]	0.6846	2.4267	17.7387	8.0669	9.7648	14.5966	0.4129	0.4676	7.6135
LuxDiT [8]	0.6748	2.1152	17.2268	6.6138	7.1764	10.2357	0.3455	0.3670	11.2195
Ours(w/ cross-attn)	0.4416	1.5227	12.5808	3.9799	3.8031	6.4758	0.2185	0.2367	7.1766
Ours(w/o GRPO)	**0.3294**	**1.1880**	**11.0929**	3.2781	3.1401	**5.7669**	0.1682	0.1890	7.0114
Ours(Full Pipeline)	0.3342	1.2425	11.1726	**3.1685**	**3.0819**	5.7771	**0.1642**	**0.1888**	**6.9728**

Table 2: Quantitative Evaluation on 360° HDR panoramas

Method	si-RMSE	Ang. Err.	N-RMSE	PSNR	SSIM
DiffusionLight [1]	19.1128	8.7134	2.7444	9.2396	0.4945
StyleLight [24]	19.0760	11.3186	2.7463	9.0099	0.3907
LuxDiT [8]	18.7746	8.2269	2.9873	8.7136	0.5765
Ours (w/ CA)	14.0573	4.2870	2.5023	12.3162	0.7107
Ours (w/o GRPO)	**12.9084**	3.8918	2.4056	12.8669	0.7495
Ours (Full)	13.0174	**3.8420**	**2.3274**	**13.3969**	**0.7523**

Table 3: Quantitative Evaluation on Portrait Relighting

Method	si-RMSE	Ang. Err.	N-RMSE	PSNR	SSIM
DiffusionLight [1]	0.1239	5.4047	0.1609	21.3372	0.7989
StyleLight [24]	0.1102	6.2116	0.1443	22.4622	0.8293
LuxDiT [8]	0.0806	4.5583	0.1163	25.0298	0.8520
Ours (w/ CA)	0.0663	2.8826	0.0778	29.7832	0.9175
Ours (w/o GRPO)	0.0429	2.0900	0.0588	31.8476	0.9482
Ours (Full)	**0.0415**	**2.0785**	**0.0580**	**32.1927**	**0.9489**

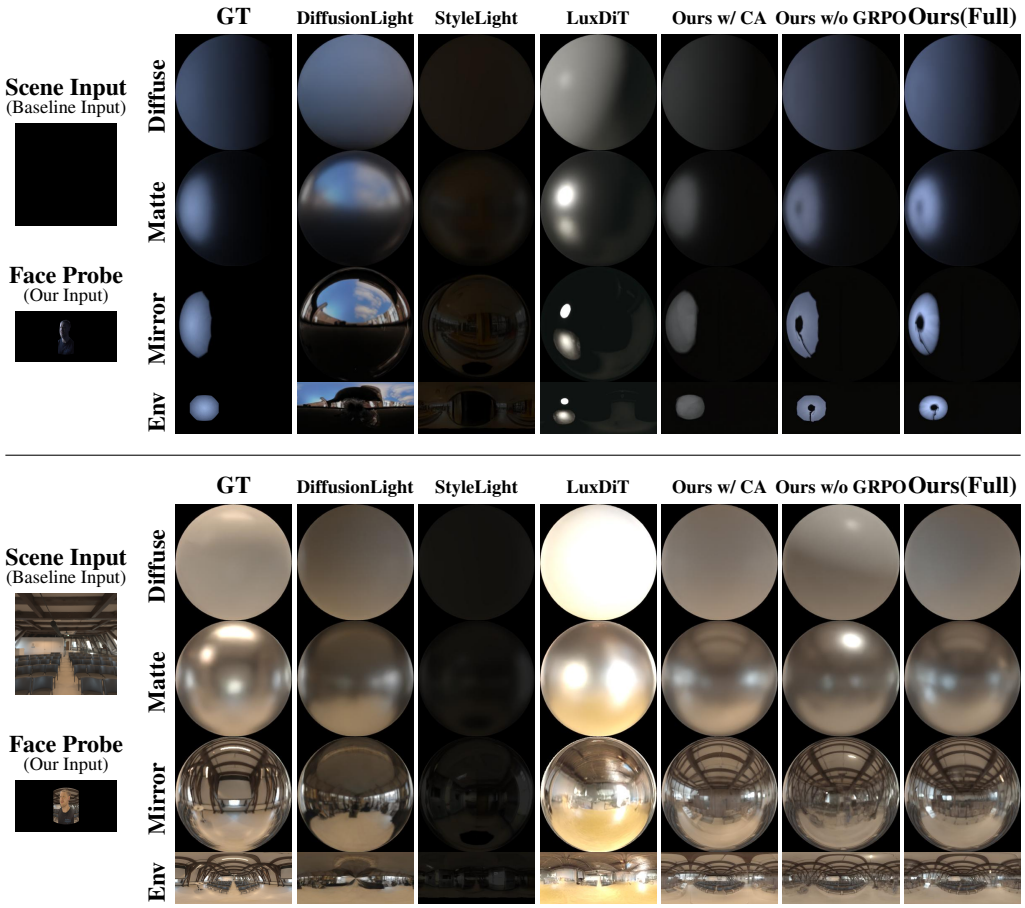

Figure 3: Qualitative comparison of relighting on virtual spheres. All panels are rendered from linear HDR under one fixed exposure shared by all methods and the GT (Sec. 4.1). Our full model is closest to the GT in structure and energy; other methods show a clear gap. Extended comparisons and enlarged high-resolution views are provided in Appendix F.3 and the supplementary material.

Rendering Protocol. Probe renderings and their derived metrics (Table 1) are computed directly from linear HDR under a single fixed exposure shared by all methods and the GT within each scene. We intentionally bypass per-image percentile tone mapping, which is mathematically invariant to global radiance scale and artificially hides energy differences (Appendix B).

Qualitative Comparison. Visual results (Fig. 3) corroborate this. Given zero-context inputs (top), baselines degenerate into probabilistic hallucinations biased by dataset priors, whereas our explicit face probes robustly recover specific out-of-view sources. In complex environments (bottom), baselines yield severe color shifts. Constrained by OLAT-GRPO, our pipeline accurately reconstructs sharp high-frequency lighting morphologies and absolute exposure consistency. Finally, the predicted environment maps maintain high input consistency because masked observations provide pixel-aligned spatial constraints, significantly reducing learning complexity.

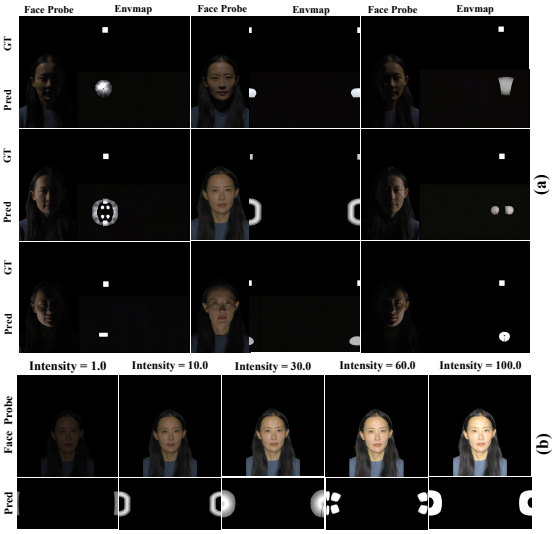

Figure 4: Analysis of Light Source Direction (a) and Intensity (b). Although the appearance of the light sources is not perfectly reconstructed, our method accurately estimates their positions and intensities. This is crucial for practical applications.

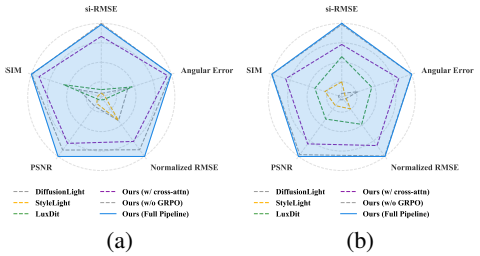

(a) (b)

Figure 5: Radar Comparison. (a) Metrics on environment map recovery. (b) Metrics on portrait relighting.

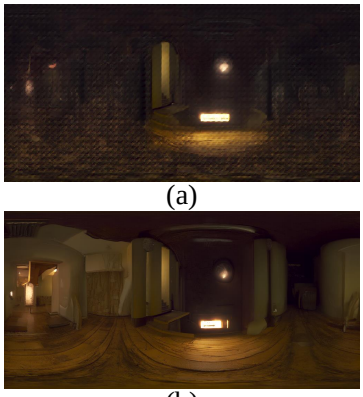

(a)

(b)

Figure 6: Checkerboard artifacts of GRPO.

4.2 Evaluation on 360° HDR Panoramas

Quantitative Evaluation on 360° Panoramas. Evaluating directly within the linear HDR radiance space (Table 2) further exposes the limitations of existing methods. Halving the Angular Error compared to the strongest baseline alongside sweeping improvements across all metrics highlights a fundamental distinction: unconstrained baselines blindly hallucinate light source locations and intensities, resulting in severe spatial and energetic deviations. By contrast, our framework leverages the face probe to deterministically anchor light directions (drastically reducing Angular Error) and employs OLAT-GRPO to calibrate absolute energy, indicating that it reconstructs the true illumination distribution rather than merely synthesizing visually plausible pixels.

Observed-Region Contamination. Restricting evaluation to unobserved pixels (Appendix C) leaves the margin intact and slightly *larger*, so the gains come from outpainting and lighting recovery, not from background copying; for reference, the observed region is only 5.05% of ERP pixels and is conditioned rather than pixel-pasted.

Isolating the Facial Probe. To strictly separate the effect of facial cues from the background outpainting prior, we quantitatively evaluate the model without background context (portrait only on a pure black background). Conditioned solely on the isolated portrait, the dominant-source direction error reduces significantly from $104.58°$ (when neither background nor portrait is available) to $55.14°$ (Appendix D). This directly demonstrates that the portrait provides explicit illumination evidence independent of contextual outpainting, while the full input performs best ($37.82°$), confirming that facial cues and ambient context are complementary.

4.3 Evaluation on Portrait Relighting

Quantitative Results. The ultimate litmus test for an estimated environment map is its physical validity in downstream rendering. As shown in Table 3, the dramatic reductions in si-RMSE (by 48.5%) and Angular Error (by 54.4%) compared to LuxDiT highlight a crucial advantage of our

pipeline. Because portrait rendering is extraordinarily sensitive to illumination variations, baseline models based on unconstrained inference inevitably suffer from erroneous radiance distributions and energy non-conservation, yielding incorrect skin exposures. Conversely, by grounding generation on deterministic face cues and explicitly penalizing facial rendering discrepancies via OLAT-GRPO, our framework accurately recovers both directional lighting and absolute energy.

Qualitative Analysis & Radiometric Nature of OLAT-GRPO. While OLAT-GRPO consistently improves all metrics in Table 3 and four of five in Table 2, its numerical delta in scale-invariant metrics appears modest. This is theoretically expected—si-RMSE analytically factors out global scale, angular error is invariant to positive scaling—so metrics that normalize away absolute radiance suppress the primary quantity targeted by GRPO; on linear scale-aware quantities (Section 4.4) the reduction is substantial. Figure 7 visually corroborates this mechanism: lowered exposure visualizations (red boxes) reveal that OLAT-GRPO effectively rectifies intensity deviations, yielding facial shading and highlights (red arrows) that more closely align with physical light transport laws.

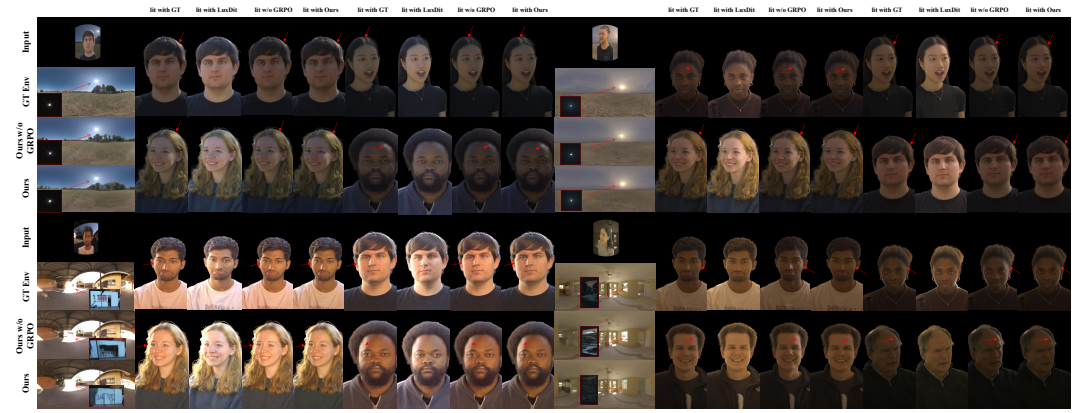

Figure 7: Comparison for Portrait Relighting. Enlarged, full-page high-resolution views of each panel and further results are provided in Appendix F.2, Appendix F.4, and the supplementary material.

4.4 Evaluation on Independent Third-Party Benchmark

To strictly verify true out-of-distribution generalization and eliminate the confound of training-distribution familiarity, we evaluate on an independent benchmark of 105 strictly unseen scenes. Standard metrics (e.g., si-RMSE) are scale-invariant and normalize away absolute radiance, masking the radiometric calibration targeted by OLAT-GRPO. Therefore, we introduce strictly scale-aware physical metrics (Table 4) to explicitly evaluate absolute energy.

FaceProbe retains an unambiguous advantage over LuxDiT with strictly non-overlapping 95% confidence intervals on key metrics, confirming the margin stems from fundamental methodological superiority rather than dataset memorization. Finally, comparing Ours (Full) against Ours (w/o GRPO) shows that physical alignment consistently shrinks errors across all energy-domain quantities, validating GRPO's ability to calibrate absolute lighting energy.

Table 4: **Evaluation on Independent Third-Party Benchmark (105 Scenes).** Mean and [95% bootstrap confidence intervals of the mean, 10,000 resamples].

Metric ↓	Ours (w/o GRPO)	Ours (Full)	LuxDiT [8]
Linear-HDR RMSE	2.2945 [1.6662, 3.0720]	**1.9624** [1.4537, 2.5753]	4.2052 [3.1335, 5.5011]
Integrated Energy Error (%)	161.52 [71.68, 314.65]	**133.94** [63.30, 252.25]	811.45 [181.02, 2010.83]
Envmap Exposure Error (EV)	0.9654 [0.7801, 1.1663]	**0.9440** [0.7631, 1.1371]	1.6478 [1.4078, 1.9096]
OLAT-Relit Face Exposure Error (EV)	0.9582 [0.7826, 1.1451]	**0.9445** [0.7736, 1.1234]	2.6998 [2.4900, 2.9216]
Per-Channel Energy Error (%)	171.9104 [85.8706, 314.1391]	**143.1726** [72.8178, 257.2359]	697.6978 [161.3508, 1708.8512]
Log-Radiance Hist Distance	1.1937 [0.9849, 1.4194]	**1.1328** [0.9298, 1.3492]	1.6984 [1.4523, 1.9595]

4.5 Ablation Studies

Foreground Cue Injection. Tables 1-3 demonstrate that sequence concatenation (applying LoRA exclusively to Self-Attention) consistently outperforms conventional Cross-Attention injection. Be-

cause all inputs are uniformly mapped to an equirectangular projection, their patch embeddings form a unified sequence with stable relative positional relationships. Self-Attention naturally leverages positional encodings, easily capturing the critical geometric correlations between facial specular cues and light source locations. Conversely, the Cross-Attention mechanism in the Wan architecture lacks relative positional encoding; for a task fundamentally driven by mapping directional face cues to spatial light positions, this missing geometric inductive bias places Cross-Attention at a distinct disadvantage, aside from introducing redundant trainable parameters.

Facial Cue Perturbations. To verify that the network actively exploits facial cues rather than defaulting to background outpainting when context is present, we evaluate systematic perturbations under pixel-identical backgrounds (Appendix D). Systematically degrading only the reference portrait (blurring, blanking, and identity-swapping [39, 40]) triggers a monotonic, statistically significant performance drop across all metrics ($p < 10^{-7}$). This monotonic degradation proves that the model actively attends to fine-grained facial reflections rather than treating the portrait as a passive occlusion to be bypassed.

Effectiveness of OLAT-GRPO. As established in Sec. 4.4, RL fine-tuning calibrates absolute radiometry rather than displacing spatial structures: it achieves a 14.5% reduction in linear-HDR RMSE and a 17.1% drop in integrated energy error, with gains strictly concentrated in the energy-domain quantities the reward acts on.

Role of the Discriminator. Optimizing solely for physical relighting rewards induces severe "reward hacking," generating heavy checkerboard artifacts to artificially satisfy the relighting metric (Fig. 6(a)). Integrating an adversarial discriminator explicitly penalizes unrealistic panoramas, ensuring that the model explores physically grounded illuminations without compromising structural fidelity (Fig. 6(b)).

Holistic Performance. Radar charts (Fig. 5) synthesize all five metrics, confirming that our full pipeline achieves the most balanced performance across panorama recovery and downstream relighting.

5 Conclusion and Limitations

This paper presents FaceProbe, a novel generative framework for reconstructing high-fidelity 360° HDR environment maps from a narrow-FoV portrait. By reformulating the task as a cue-driven masked recovery process and introducing OLAT-GRPO physical alignment, we bridge the gap between unconstrained generation and physics-grounded inverse rendering, achieving state-of-the-art performance in panoramic recovery and downstream relighting. Matched-input comparisons, facial perturbation ablations and an independent third-party benchmark attribute these gains to the portrait pathway and physical alignment, rather than to the outpainting prior or to training-distribution familiarity.

Limitations and Future Work. While OLAT-GRPO accurately recovers light direction and absolute intensity, it occasionally struggles with the exact morphological reconstruction of infinitely sharp synthetic light shapes. Because the reward prioritizes facial photometric consistency, the model favors global energy conservation over precise unobserved high-frequency geometric boundaries. Expanding the training data scale and exploring directional reflectance priors remain promising avenues for future research.

Acknowledgements and Disclosure of Funding

We thank DreamX (USA) for providing financial support and computational resources for this work, as well as for its key contributions to data preparation, annotation, and cleaning. We also thank Wen Li and Chenxi Fan from Artirizon AI for their critical contributions to the data curation system. Kairui Feng was supported by the National Natural Science Foundation of China (Grant No. T2600685), and the Xiaomi Foundation. We also thank the Shanghai Innovation Institute.

References

[1] Pakkapon Phongthawee, Worameth Chinchuthakun, Nontaphat Sinsunthithet, Amit Raj, Varun Jampani, Pramook Khungurn, and Supasorn Suwajanakorn. Diffusionlight: Light probes for free by painting a chrome ball. *IEEE*, 2023.

[2] Worameth Chinchuthakun, Pakkapon Phongthawee, Amit Raj, Varun Jampani, Pramook Khungurn, and Supasorn Suwajanakorn. Diffusionlight-turbo: Accelerated light probes for free via single-pass chrome ball inpainting, 2026.

[3] Ruofan Liang, Zan Gojcic, Merlin Nimier-David, David Acuna, Nandita Vijaykumar, Sanja Fidler, and Zian Wang. Photorealistic object insertion with diffusion-guided inverse rendering. In *European Conference on Computer Vision*, pages 446–465. Springer, 2024.

[4] Zian Wang, Wenzheng Chen, David Acuna, Jan Kautz, and Sanja Fidler. Neural light field estimation for street scenes with differentiable virtual object insertion. In *European Conference on Computer Vision*, pages 380–397. Springer, 2022.

[5] Peter Kocsis, Vincent Sitzmann, and Matthias Nießner. Intrinsic image diffusion for indoor single-view material estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5198–5208, 2024.

[6] Ruofan Liang, Zan Gojcic, Huan Ling, Jacob Munkberg, Jon Hasselgren, Chih-Hao Lin, Jun Gao, Alexander Keller, Nandita Vijaykumar, Sanja Fidler, et al. Diffusion renderer: Neural inverse and forward rendering with video diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26069–26080, 2025.

[7] Chloe Legendre, Wan Chun Ma, Rohit Pandey, Sean Fanello, Christoph Rhemann, Jason Dourgarian, Jay Busch, and Paul Debevec. Learning illumination from diverse portraits, 2020.

[8] Ruofan Liang, Kai He, Zan Gojcic, Igor Gilitschenski, Sanja Fidler, Nandita Vijaykumar, and Zian Wang. Luxdit: Lighting estimation with video diffusion transformer, 2025.

[9] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 145–156, 2000.

[10] Pramod Rao, Abhimitra Meka, Xilong Zhou, Gereon Fox, Mallikarjun B R, Fangneng Zhan, Tim Weyrich, Bernd Bickel, Hanspeter Pfister, Wojciech Matusik, Thabo Beeler, Mohamed Elgharib, Marc Habermann, and Christian Theobalt. 3dpr: Single image 3d portrait relight using generative priors, 2025.

[11] Timo Teufel, Pulkit Gera, Xilong Zhou, Umar Iqbal, Pramod Rao, Jan Kautz, Vladislav Golyanik, and Christian Theobalt. Humanolat: A large-scale dataset for full-body human relighting and novel-view synthesis, 2025.

[12] Dachuan Cheng, Jian Shi, Yanyun Chen, Xiaoming Deng, and Xiaopeng Zhang. Learning scene illumination by pairwise photos from rear and front mobile cameras. In *Computer Graphics Forum*, volume 37, pages 213–221. Wiley Online Library, 2018.

[13] Mathieu Garon, Kalyan Sunkavalli, Sunil Hadap, Nathan Carr, and Jean Franois Lalonde. Fast spatially-varying indoor lighting estimation. *IEEE*, 2020.

[14] Marc-André Gardner, Yannick Hold-Geoffroy, Kalyan Sunkavalli, Christian Gagné, and Jean-François Lalonde. Deep parametric indoor lighting estimation. *IEEE*, 2020.

[15] Fangneng. Zhan, Yingchen. Yu, Changgong. Zhang, Rongliang. Wu, Wenbo. Hu, Shijian. Lu, Feiying. Ma, Xuansong. Xie, and Ling. Shao. Gmlight: Lighting estimation via geometric distribution approximation. *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, 31:2268–2278, 2022.

[16] Pratul P. Srinivasan, Ben Mildenhall, Matthew Tancik, Jonathan T. Barron, Richard Tucker, and Noah Snavely. Lighthouse: Predicting lighting volumes for spatially-coherent illumination, 2020.

[17] Xuecan Wang, Shibang Xiao, and Xiaohui Liang. Lightoctree: Lightweight 3d spatially-coherent indoor lighting estimation, 2024.

[18] Shiyuan Shen, Zhongyun Bao, Hong Ding, Wenju Xu, Tenghui Lai, and Chunxia Xiao. Stglight: Online indoor lighting estimation via spatio-temporal gaussian fusion. *ACM Transactions on Graphics (TOG)*, 44(6):1–14, 2025.

[19] JunYong Choi, SeokYeong Lee, Haesol Park, Seung-Won Jung, Ig-Jae Kim, and Junghyun Cho. Mair: multi-view attention inverse rendering with 3d spatially-varying lighting estimation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8392–8401. IEEE, 2023.

[20] Zhengqin Li, Li Yu, Mikhail Okunev, Manmohan Chandraker, and Zhao Dong. Spatiotemporally consistent hdr indoor lighting estimation. *ACM Transactions on Graphics*, 42(3):1–15, 2023.

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

[22] Shuran Song and Thomas Funkhouser. Neural illumination: Lighting prediction for indoor environments. 2019.

[23] Mohammad Reza Karimi Dastjerdi, Yannick Hold-Geoffroy, Jonathan Eisenmann, and Jean Franois Lalonde. Everlight: Indoor-outdoor editable hdr lighting estimation. *ArXiv*, abs/2304.13207, 2023.

[24] Guangcong Wang, Yinuo Yang, Chen Change Loy, and Ziwei Liu. Stylelight: Hdr panorama generation for lighting estimation and editing. *Springer, Cham*, 2022.

[25] Chloe LeGendre, Wan-Chun Ma, Graham Fyffe, John Flynn, Laurent Charbonnel, Jay Busch, and Paul Debevec. Deeplight: Learning illumination for unconstrained mobile mixed reality. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5918–5928, 2019.

[26] Shiyuan Shen, Zhongyun Bao, Wenju Xu, and Chunxia Xiao. Illumidiff: Indoor illumination estimation from a single image with diffusion model. *IEEE transactions on visualization and computer graphics*, 2025.

[27] Jacob Munkberg, Zian Wang, Ruofan Liang, Tianchang Shen, and Jon Hasselgren. Videomat: Extracting pbr materials from video diffusion models. In *Computer Graphics Forum*, volume 44, page e70180. Wiley Online Library, 2025.

[28] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021.

[29] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. Mitsuba 3 renderer, 2022. https://mitsuba-renderer.org.

[30] Xuanhong Chen, Bingbing Ni, Naiyuan Liu, Ziang Liu, Yiliu Jiang, Loc Truong, and Qi Tian. Coogan: A memory-efficient framework for high-resolution facial attribute editing. In *European Conference on Computer Vision*, pages 670–686. Springer, 2020.

[31] Qiaoqiao Jin, Xuanhong Chen, Meiguang Jin, Ying Chen, Rui Shi, Yucheng Zheng, Yupeng Zhu, and Bingbing Ni. Toward tiny and high-quality facial makeup with data amplify learning. In *European Conference on Computer Vision*, pages 340–356. Springer, 2024.

[32] Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models with reinforcement learning, 2024.

[33] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024.

[34] Jie Liu, Gongye Liu, Jiajun Liang, Yangguang Li, Jiaheng Liu, Xintao Wang, Pengfei Wan, Di Zhang, and Wanli Ouyang. Flow-grpo: Training flow matching models via online rl, 2025.

[35] Xuanhong Chen, Muchun Chen, Yugang Chen, Yinxin Lin, Bilian Ke, and Bingbing Ni. Large generative model impulsed lightweight gaze estimator via deformable approximate large kernel pursuit. *IEEE Transactions on Image Processing*, 34:1149–1162, 2025.

[36] Muchun Chen, Yinxin Lin, Yuhan Li, Yugang Chen, Xuanhong Chen, Bilian Ke, and Bingbing Ni. Learning gaze synthesizer via 3d-eye controlled diffusion and cross-domain feature alignment. *IEEE Transactions on Image Processing*, 2026.

[37] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wente Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.

[38] Poly Haven Contributors. Poly haven: The public 3d asset library. `https://polyhaven.com/`, 2025. Accessed: 2025-10-08.

[39] Renwang Chen, Xuanhong Chen, Bingbing Ni, and Yanhao Ge. Simswap: An efficient framework for high fidelity face swapping. In *Proceedings of the 28th ACM international conference on multimedia*, pages 2003–2011, 2020.

[40] Xuanhong Chen, Bingbing Ni, Yutian Liu, Naiyuan Liu, Zhilin Zeng, and Hang Wang. Simswap++: Towards faster and high-quality identity swapping. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(1):576–592, 2023.

A Implementation Details

We initialize our backbone with `Wan2.1-Fun-14B-Control` and employ Low-Rank Adaptation (LoRA, $r = 256$) across transformer blocks to avoid prohibitive full-parameter tuning. The training comprises two phases: *1) Task-Specific Adaptation:* We optimize the network via the MSE objective to learn foundational visual mappings. This phase is trained for 30K steps using the AdamW optimizer (lr=1×10^{-5}) with `bfloat16` precision on two NVIDIA A100 GPUs. *2) Physics-Aware Alignment (OLAT-GRPO):* We enforce physical constraints via RL fine-tuning for 2K steps (lr $= 1 \times 10^{-5}$). To optimize sample efficiency and gradient quality, we truncate trajectories to 10 denoising steps, generate $G = 16$ candidates per condition for group-relative advantages, and reuse rollouts for 2 consecutive updates.

B Scope, Protocol and Rendering Notes

This appendix collects the statements that the main text refers to but that are too long for it. None of them changes a result; they fix the meaning of our terms and the protocol under which the numbers and figures were produced.

Operational meaning of "physical grounding". FaceProbe does not assume that a face is a calibrated mirror, a Lambertian surface, or an object with a known BRDF. Nor does it attempt to explicitly recover a high-dimensional facial BSSRDF from one portrait—that decomposition would itself be severely ill-posed and is not required by our formulation. Our claim is that a portrait contains physically produced illumination observations (including directional specularities, shading, eye reflections, and colour responses) that constrain HDR recovery when combined with visible context and a learned prior. "Face probe" describes this observational role; it does not mean that the face directly images the environment like a mirror ball.

Complex transport, including subsurface scattering, also does not invalidate this formulation. Under fixed geometry and material, the OLAT response remains linear with respect to incident illumination. Each captured basis $\mathbf{B}_k$ already contains the subject's actual combined transport response—including diffuse and specular reflection, cast shadows, eye and hair appearance, and localized subsurface effects. We therefore do not need to analytically separate or approximate these components. In typical frontal/side portrait conditions, the strongest subsurface transmission is localized, becoming most visible under intense back/side lighting in thin regions such as the ears or nose. It does not systematically erase the directional surface specularities used by the model. More importantly, even in these cases, the effect remains present in the captured OLAT basis and is evaluated through the same image-based relighting operator. The physical preference is also evaluated at the level where illumination becomes observable: candidate HDR maps are evaluated by whether they reproduce the GT portrait through a standard physical light-transport superposition.

Evaluation scope. Our benchmark is a controlled image-based relighting benchmark constructed from real captured LightStage OLAT responses and real HDR panoramas, combined through the standard linear operator. This provides exact paired ground truth for both the recovered illumination and the downstream portrait appearance—ground truth that is generally unavailable for arbitrary in-the-wild portraits, where one cannot simultaneously capture a photograph and its complete co-located 360° HDR illumination. It is therefore not an unrestricted paired in-the-wild evaluation. Training and testing are strictly separated: 100 subjects and 3,000 HDR maps are used for training, while 20 unseen subjects and 527 unseen HDR maps are reserved for testing. This scope does not undermine the reported generalization across unseen subjects and illuminations. Qualitative results on uncontrolled in-the-wild photographs are given in Appendix E.

Rendering protocol. To accurately visualize absolute radiometric differences, all virtual-sphere renderings (and the quantitative results in Table 1) are produced directly from linear HDR under a single fixed exposure shared across all methods and the ground truth within each scene. We intentionally bypass the percentile-based per-image tone mapping commonly used in standard evaluation pipelines. Because percentile normalization derives its exposure gain dynamically from each image's own intensity distribution, it is mathematically invariant to global radiance scales. Consequently, it artificially masks absolute energy differences—which our OLAT-GRPO specifically calibrates—and disproportionately amplifies dark background values. By employing a fixed-exposure linear rendering protocol, we ensure that spatial, chromatic, and absolute-energy differences remain

physically legible. We note that this protocol applies exclusively to the virtual-sphere visualizations and their corresponding metrics; panorama-level evaluations and downstream OLAT portrait relighting operate natively in linear HDR space and are inherently unaffected.

C Evaluation Fairness: Unobserved-Only Metrics, Matched Input and Retrieval Baselines

Because FaceProbe conditions on the visible background, a natural question is whether the quantitative margins in Table 2 are inflated by accurately reconstructing the observed pixels. To investigate this alongside input-fairness and retrieval baselines under a unified computational budget, we conduct paired diagnostic evaluations on a representative 200-sample test subset (evaluated directly in 128×256 linear panorama space) and explicitly compute metrics restricted exclusively to the unobserved region (the masked outpainting area).

As detailed in Table 5, restricting the evaluation strictly to unobserved pixels actually **increases** our margin over LuxDiT: the Angular Error margin expands from 10.597° to 10.648° (with FaceProbe outperforming LuxDiT on 99.5% of samples, $p = 8.2 \times 10^{-35}$), and the si-RMSE margin widens from 0.729 to 0.782. This confirms that our advantage is entirely driven by superior outpainting and lighting recovery rather than background reproduction.

Table 5: **Comparison of Full-Sphere vs. Unobserved-Only Metrics** on the 200-sample paired diagnostic test subset.

Method	Angular Error (Full)	Angular Error (Unseen)	si-RMSE (Full)	si-RMSE (Unseen)
LuxDiT (Background only)	20.760°	20.706°	16.622	16.145
LuxDiT (+ Person input)	20.597°	20.527°	16.443	15.960
FaceProbe (Ours)	$\mathbf{10.163}^\circ$	$\mathbf{10.058}^\circ$	**15.893**	**15.363**

Matched-Input Fairness. To guarantee input information fairness, we reran the official LuxDiT pipeline feeding it the identical person-containing photograph available to FaceProbe (Table 5). Providing LuxDiT with the person changes its angular error by only 0.163° ($p = 0.294$). Furthermore, evaluating dominant source properties (computed as the solid-angle-weighted centroid of the top 1% HDR radiance and normalized chromaticity difference):

- LuxDiT (Background only): Source Direction Error = 84.96°, Chromaticity Error = 0.0829.
- LuxDiT (+ Person input): Source Direction Error = 55.27°, Chromaticity Error = 0.0824.
- **FaceProbe (Ours)**: Source Direction Error = $\mathbf{37.82}^\circ$, Chromaticity Error = $\mathbf{0.0398}$.

While providing the portrait improves LuxDiT's direction estimation, FaceProbe further reduces direction error by 17.45° ($p = 9.9 \times 10^{-9}$) and cuts chromaticity error by more than half ($p = 1.5 \times 10^{-16}$). This demonstrates that our dedicated reference pathway and physical alignment harness facial cues far more effectively than unconstrained diffusion architectures. We note that we deliberately do *not* insert a calibrated probe object into the scene: an inserted mirror ball would directly image the environment and thereby add information that is absent from the original photograph, whereas an uncalibrated face is indirect, subject-dependent evidence already present in the input. Supplying both methods with the same person-containing photograph is therefore the appropriate matched-information test for our setting.

Context-Retrieval Baseline. We additionally test whether the accuracy can be explained by retrieving, or memorizing, a training panorama from the visible context. On the same split: (i) retrieving the training panorama whose visible region best matches the test input gives 19.006° full-sphere angular error; (ii) returning the mean training panorama gives 17.502°; and (iii) an oracle permitted to select its best training panorama using the complete GT sphere still reaches only 15.355°. FaceProbe attains $\mathbf{10.163}^\circ$. This baseline deliberately favours the retrieval explanation, since (iii) uses ground-truth information no method has access to, yet it remains far worse than FaceProbe—evidence that the model performs conditional recovery rather than recalling a training environment.

Statistical Protocol. All comparisons above are paired: identical samples, masks, resolutions and evaluation code, with no sample selected according to outcome. p-values are from paired tests over

samples, and the third-party benchmark additionally reports bootstrap 95% confidence intervals of the mean ($10{,}000$ resamples).

D Isolating Facial Probes via Facial Perturbations

To rigorously test whether the model actively extracts illumination information from facial cues or merely relies on background outpainting priors, we performed facial perturbation experiments on a paired subset.

Matched-Background Perturbations. Holding the background texture and visibility mask pixel-for-pixel identical, we systematically varied only the reference portrait condition. As shown in Table 6, weakening the cue (blurring) causes significant degradation ($+0.206°$, $p = 1.6 \times 10^{-8}$); completely removing the face (blanking) further worsens the error ($+0.311°$, $p = 2.7 \times 10^{-10}$); and providing an inconsistent portrait from another subject causes the most severe degradation ($+0.578°$, $p = 3.3 \times 10^{-18}$). Scale-invariant RMSE degrades in an identical monotonic sequence ($+0.102/+0.132/+0.198$, all $p < 5 \times 10^{-7}$). A model ignoring its portrait pathway would show no such ordered effect.

Table 6: **Matched-Background Facial Perturbation Ablation.** Performance degradation when modifying only the facial reference while keeping the background context pixel-identical.

Reference Condition	Δ Angular Error	Samples Degraded (%)	p-value
Face Heavily Blurred	$+0.206°$	66.0%	1.6×10^{-8}
Face Blanked	$+0.311°$	68.0%	2.7×10^{-10}
Face Swapped (Different Subject)	$+0.578°$	77.0%	3.3×10^{-18}

Separating Face Evidence from the Outpainting Prior. We further isolate the portrait without background context by evaluating the dominant-source direction error. As shown in Table 7, when no background and no portrait are provided, the direction error is $104.58°$. Conditioning purely on the isolated portrait (portrait only on a pure black background) reduces the error by **49.44°** down to $55.14°$ (75.0% of samples improve, $p = 4.3 \times 10^{-16}$). This directly supports our intended conclusion: the portrait provides illumination evidence independent of background outpainting. The full input condition performs best ($37.82°$), showing that portrait evidence and visible context are complementary. We do not claim that an uncalibrated face uniquely determines the complete HDR map; rather, it constrains the solution jointly with the context and the HDR prior.

Table 7: **Dominant Light Source Direction Error under Varying Conditioning.**

Condition	Dominant Source Direction Error ↓
No Background and No Portrait	$104.58°$
Portrait Only (on Black Background)	$55.14°$
Full Input (Portrait + Background Context)	$\mathbf{37.82}°$

E In-the-Wild Qualitative Evaluation and Out-of-Distribution Generalization

Furthermore, to comprehensively evaluate real-world generalization on non-laboratory natural scenes, we include qualitative relighting evaluations on arbitrary in-the-wild portrait photographs. Because one cannot simultaneously capture an unconstrained photograph and its complete co-located $360°$ HDR illumination, these in-the-wild samples inherently lack ground truth. Consequently, they were strictly excluded from our training phase and represent an entirely distinct data distribution from our LightStage-based train and test sets, serving as a rigorous test of true out-of-distribution generalization. As illustrated in Figure 8, even under complex, unconstrained natural capture conditions, FaceProbe consistently infers highly plausible lighting geometries. In contrast to baselines like LuxDiT and DiffusionLight, which often hallucinate arbitrary backgrounds or fail to capture the correct color temperature and directionality, our method robustly extracts the illumination cues directly from the real-world faces.

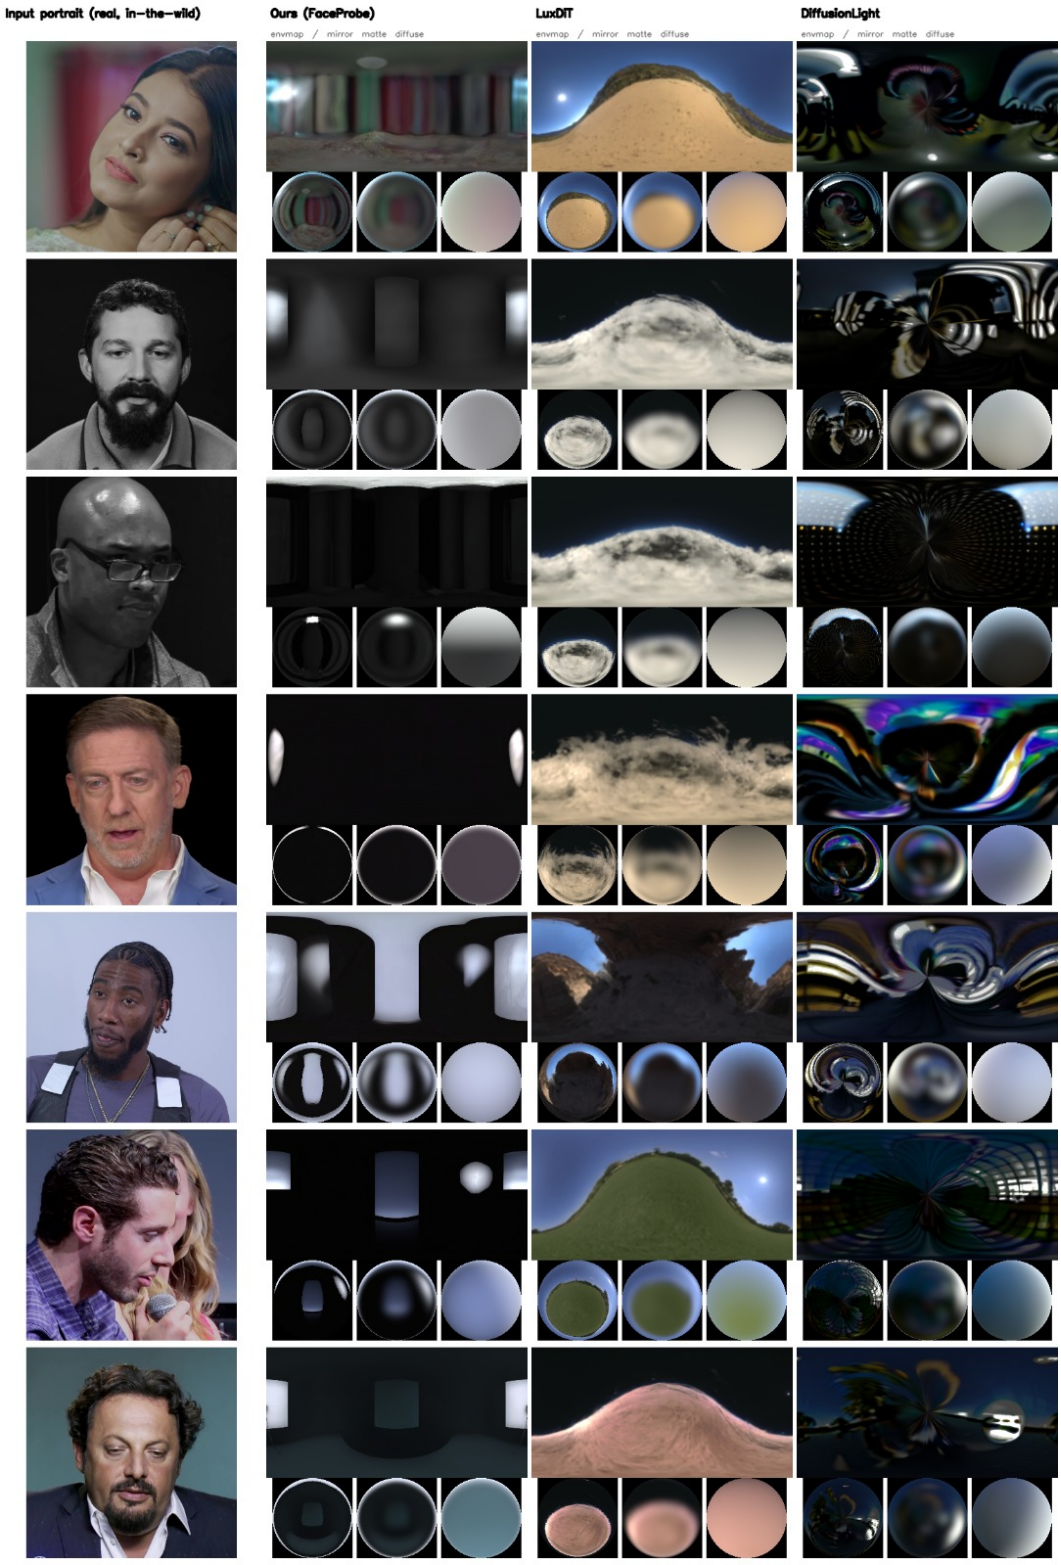

Figure 8: **Qualitative comparison on real-world in-the-wild portraits.** Given unconstrained natural portraits (left column), our FaceProbe robustly infers plausible HDR environment maps and physically consistent lighting distributions. Baselines (LuxDiT and DiffusionLight) struggle with severe color shifts and unconstrained hallucinations. High-resolution versions are provided in the supplementary material.

F Visual Results and Supplementary Comparisons

In this section, we present representative samples from our self-constructed dataset and provide additional visual comparisons with previous state-of-the-art methods. Furthermore, we include relighting results obtained using our estimated environment maps to further validate the performance of our framework.

F.1 Visual Results and Comparisons of Environment Map Estimation

In Figures 9 and 10, we provide extensive visual comparisons of panoramic environment map predictions against the strongest baseline, LuxDiT. As demonstrated, LuxDiT suffers from severe erroneous radiance distributions and inaccurate radiance energy estimation. In contrast, by leveraging explicit face-cues, our framework effectively anchors the global radiance distribution and maintains fundamental energy conservation, overcoming the probabilistic hallucinations inherent in unconstrained models. Furthermore, a close inspection comparing our base model (w/o GRPO) with the full pipeline reveals that the integration of OLAT-GRPO significantly refines the energy conservation of specular highlights. This physical alignment ensures that the recovered high-frequency lighting structures are not only visually sharper but also photometrically more accurate, yielding superior overall fidelity.

F.2 Additional Qualitative Comparisons on Portrait Relighting

Figures 11 and 12 provide extensive visual comparisons between our full pipeline and LuxDiT on the downstream portrait relighting task. As observed, portraits relit by LuxDiT frequently suffer from unnatural skin exposures, misplaced facial shadows, and distorted specularities. These rendering artifacts are the direct downstream consequences of its unconstrained environment map generation, which inherently struggles with spatial misalignments and absolute radiance energy errors. In profound contrast, our relit portraits exhibit exceptional physical fidelity, closely mirroring the ground truth across diverse complex illuminations. Fundamentally, this superiority stems from our physics-grounded generation paradigm. By explicitly coupling the generative process with the physical rendering equation via OLAT-GRPO, our framework ensures that the predicted environment maps are not merely visually plausible textures, but strictly energy-conserving HDR environment maps. Consequently, our method flawlessly reproduces accurate shading gradients and high-frequency facial specularities that strictly obey physical light transport laws.

F.3 Additional Qualitative Comparisons on Virtual Spheres

In Figure 13, we provide enlarged high-resolution views and additional qualitative comparisons of virtual sphere relighting across diverse indoor and outdoor scenes, extending Figure 3 of the main paper. All results are rendered directly from linear HDR predictions under a single fixed exposure shared with the ground truth, demonstrating that our full pipeline consistently recovers accurate light source directions, chromaticities, and global radiance energy compared to existing baselines.

F.4 Enlarged High-Resolution View of Portrait Relighting

To facilitate closer inspection of the fine-grained facial shading, specular highlights, and exposure calibration discussed in Section 4.3, we provide an enlarged, full-page version of the portrait relighting comparison (Figure 7 of the main paper) in Figure 14. As shown in the enlarged panels, incorporating OLAT-GRPO effectively rectifies radiance scale deviations and produces realistic skin reflectance that closely matches the ground truth.

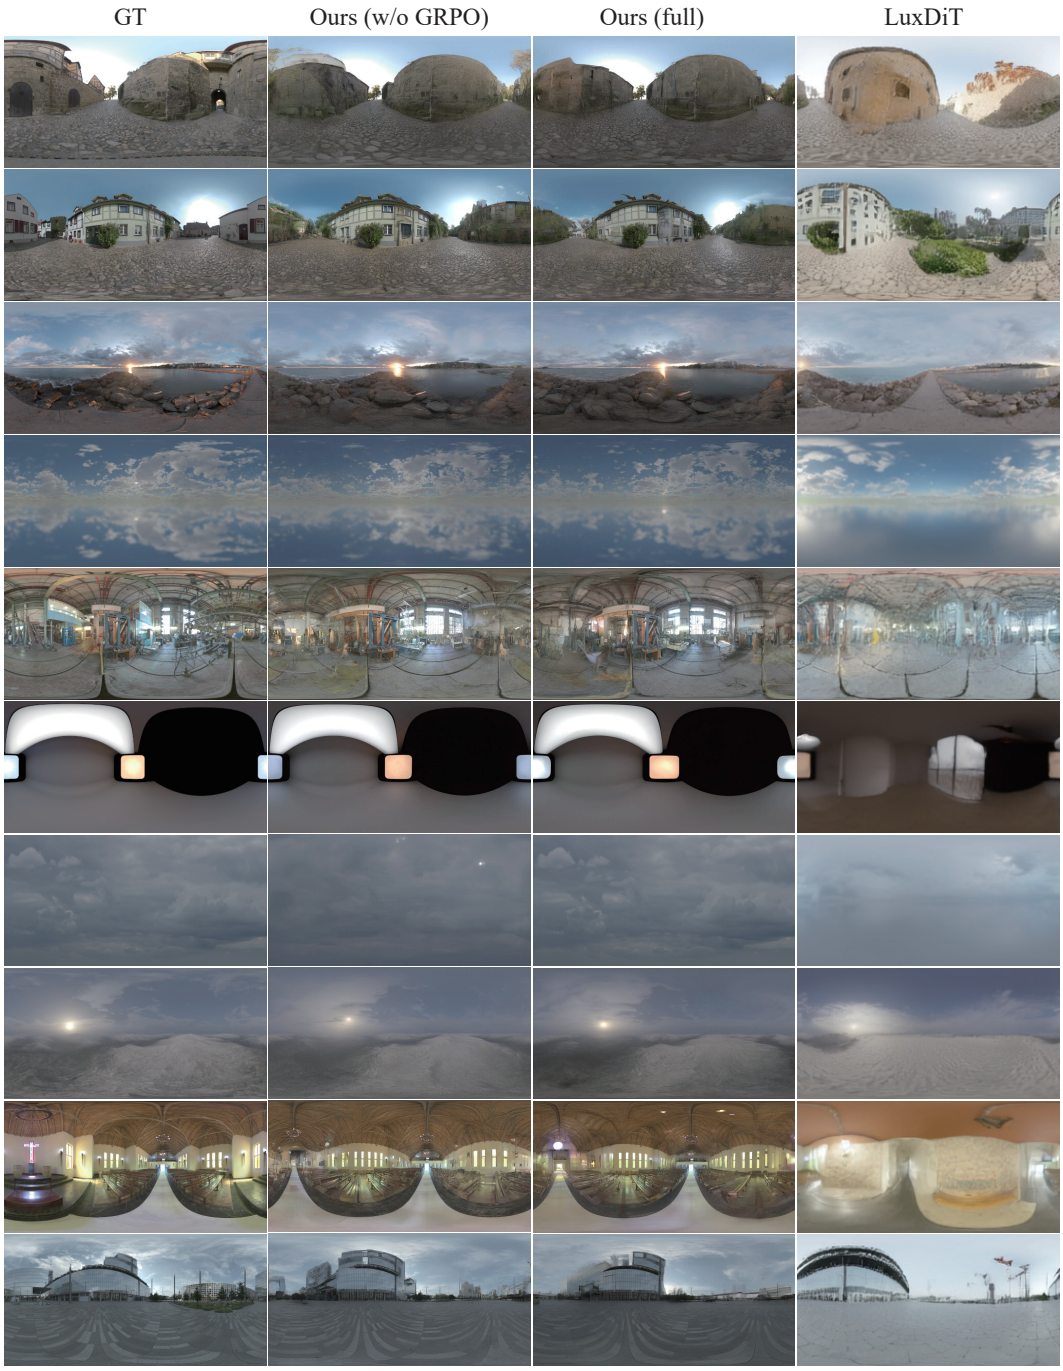

Figure 9: Additional visual comparisons with LuxDiT. As observed, our generated environment maps exhibit a higher resemblance to the GT, particularly regarding the distribution of specular highlights. Furthermore, close inspection reveals that our model equipped with OLAT-GRPO demonstrates superior fidelity and sensitivity in highlight recovery. High-resolution versions are provided in the supplementary material.

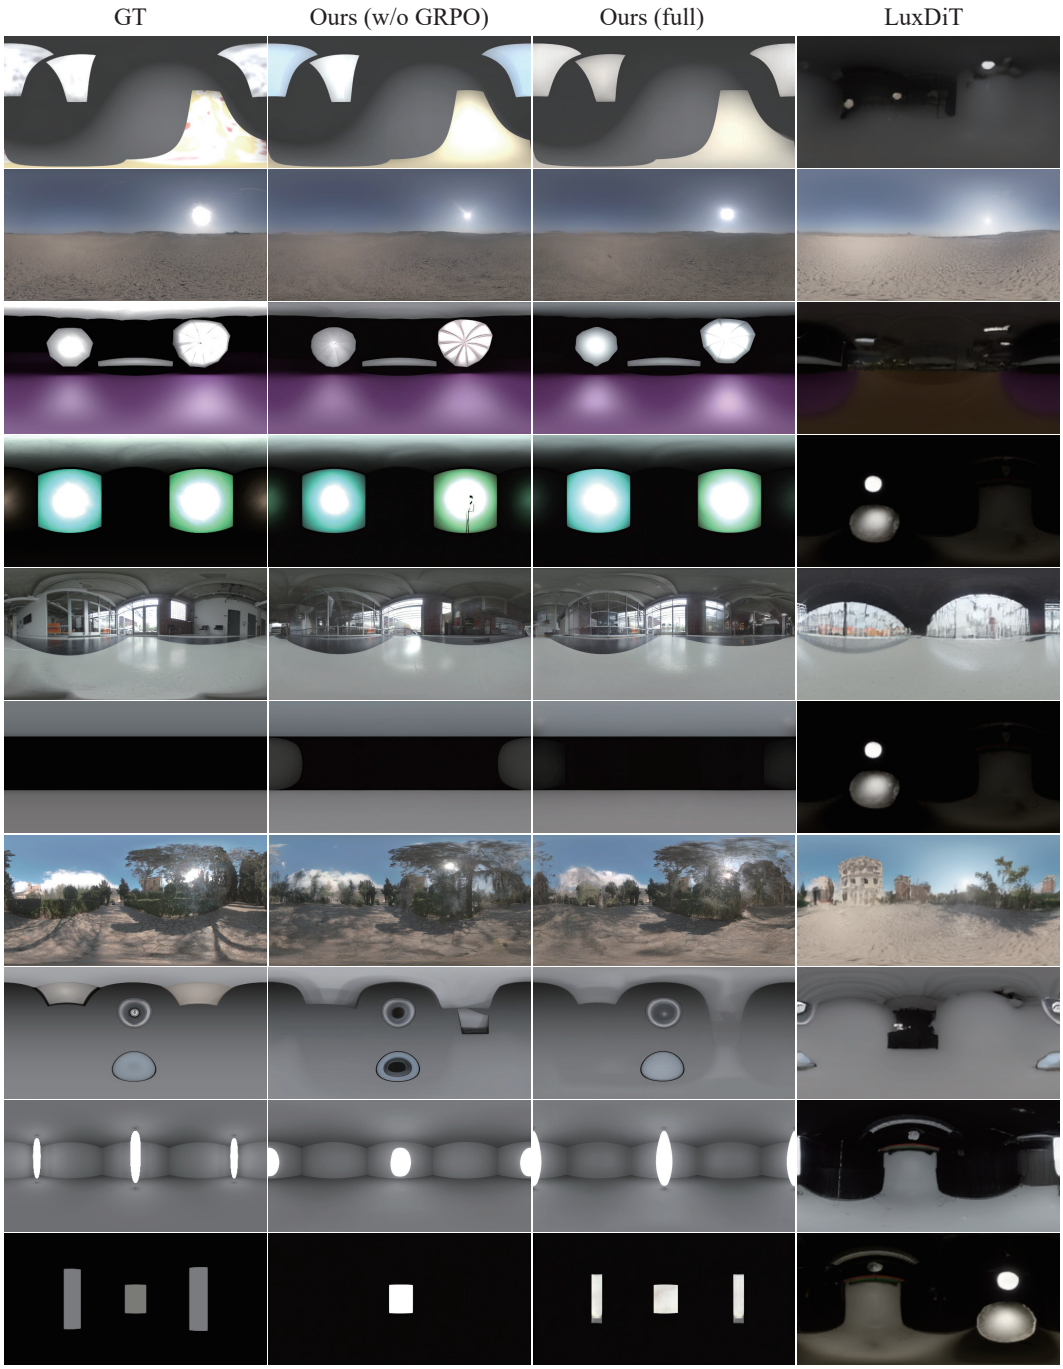

Figure 10: Additional visual comparisons with LuxDiT. As observed, our generated environment maps exhibit a higher resemblance to the GT, particularly regarding the distribution of specular highlights. Furthermore, close inspection reveals that our model equipped with OLAT-GRPO demonstrates superior fidelity and sensitivity in highlight recovery. High-resolution versions are provided in the supplementary material.

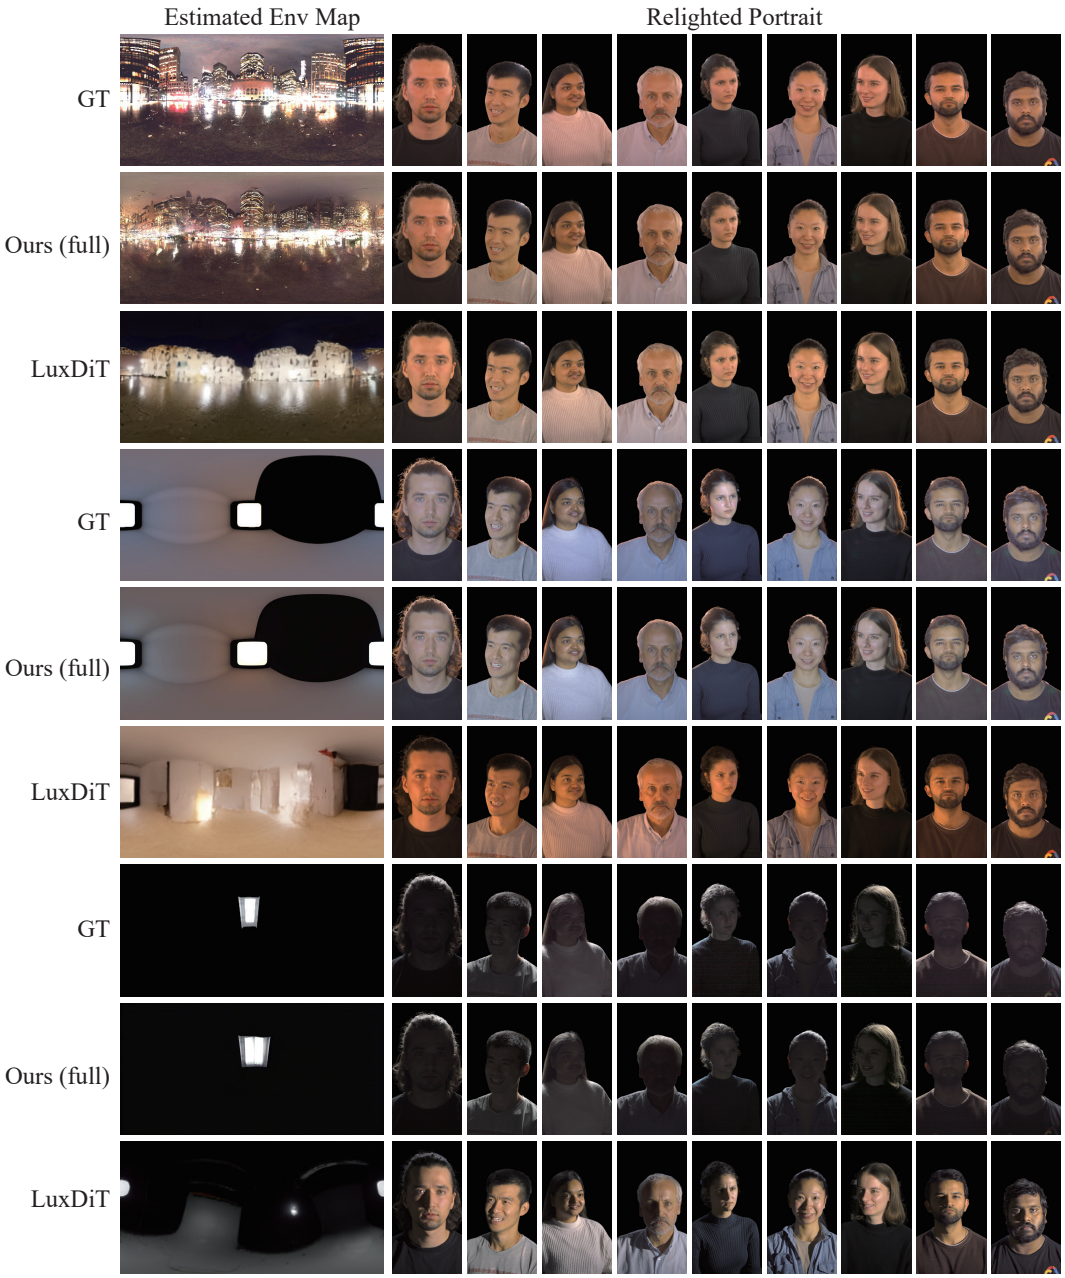

Figure 11: Additional qualitative comparisons on downstream portrait relighting against LuxDiT. While LuxDiT yields erroneous facial shadowing and unnatural skin exposures due to its physically unconstrained radiance estimation, our full pipeline seamlessly reproduces accurate shading gradients and faithful facial specularities. This demonstrates that the environment maps generated by our approach act as rigorous, energy-conserving HDR environment maps suitable for physical light transport. High-resolution versions are provided in the supplementary material.

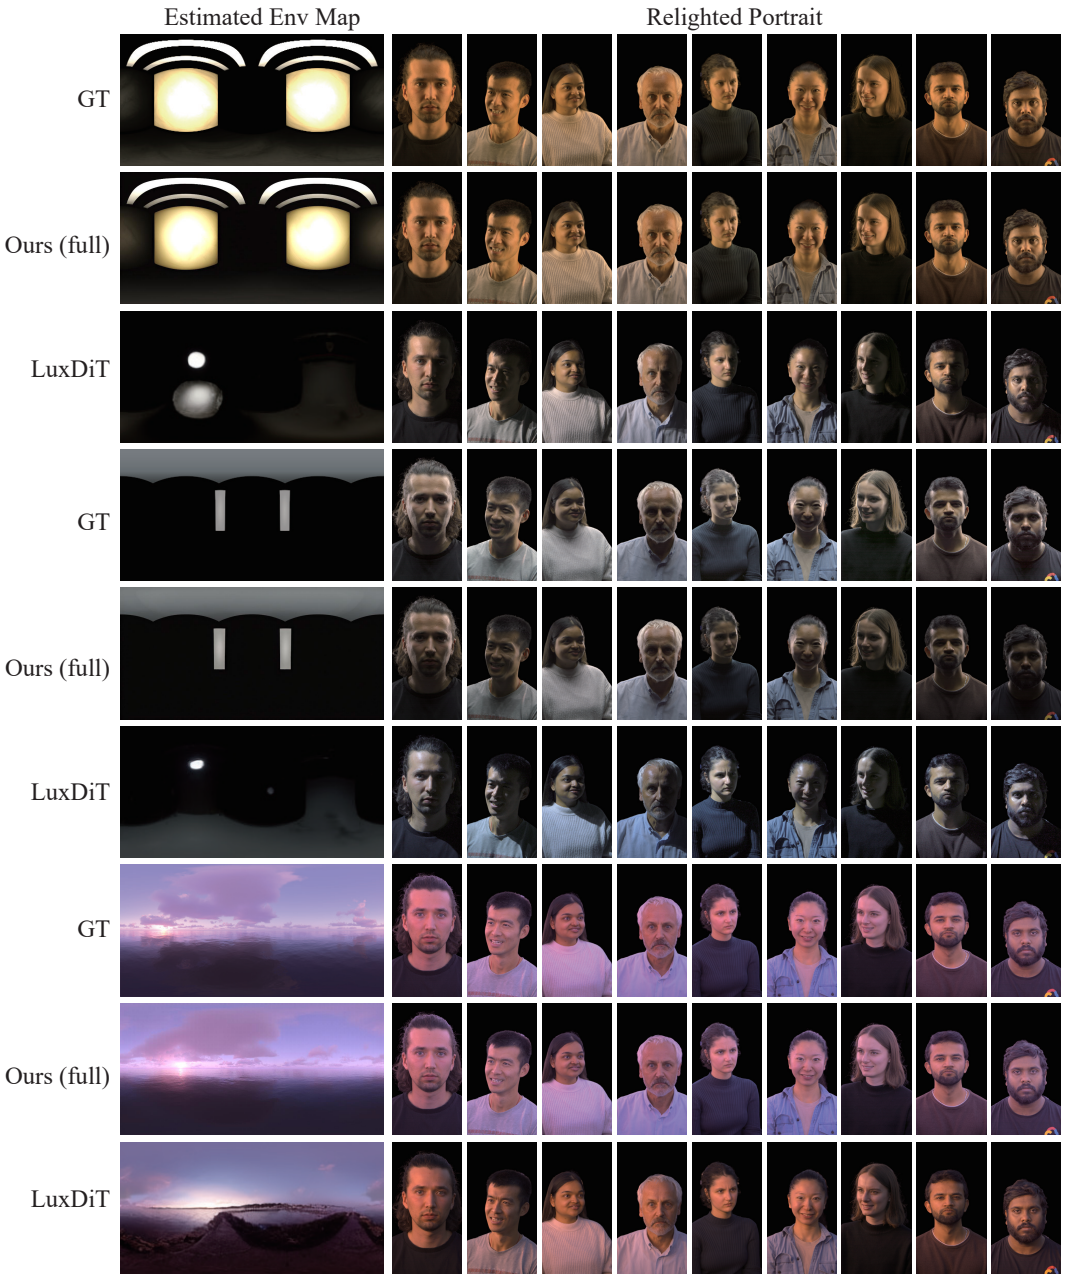

Figure 12: Additional qualitative comparisons on downstream portrait relighting against LuxDiT. While LuxDiT yields erroneous facial shadowing and unnatural skin exposures due to its physically unconstrained radiance estimation, our full pipeline seamlessly reproduces accurate shading gradients and faithful facial specularities. This demonstrates that the environment maps generated by our approach act as rigorous, energy-conserving HDR environment maps suitable for physical light transport. High-resolution versions are provided in the supplementary material.

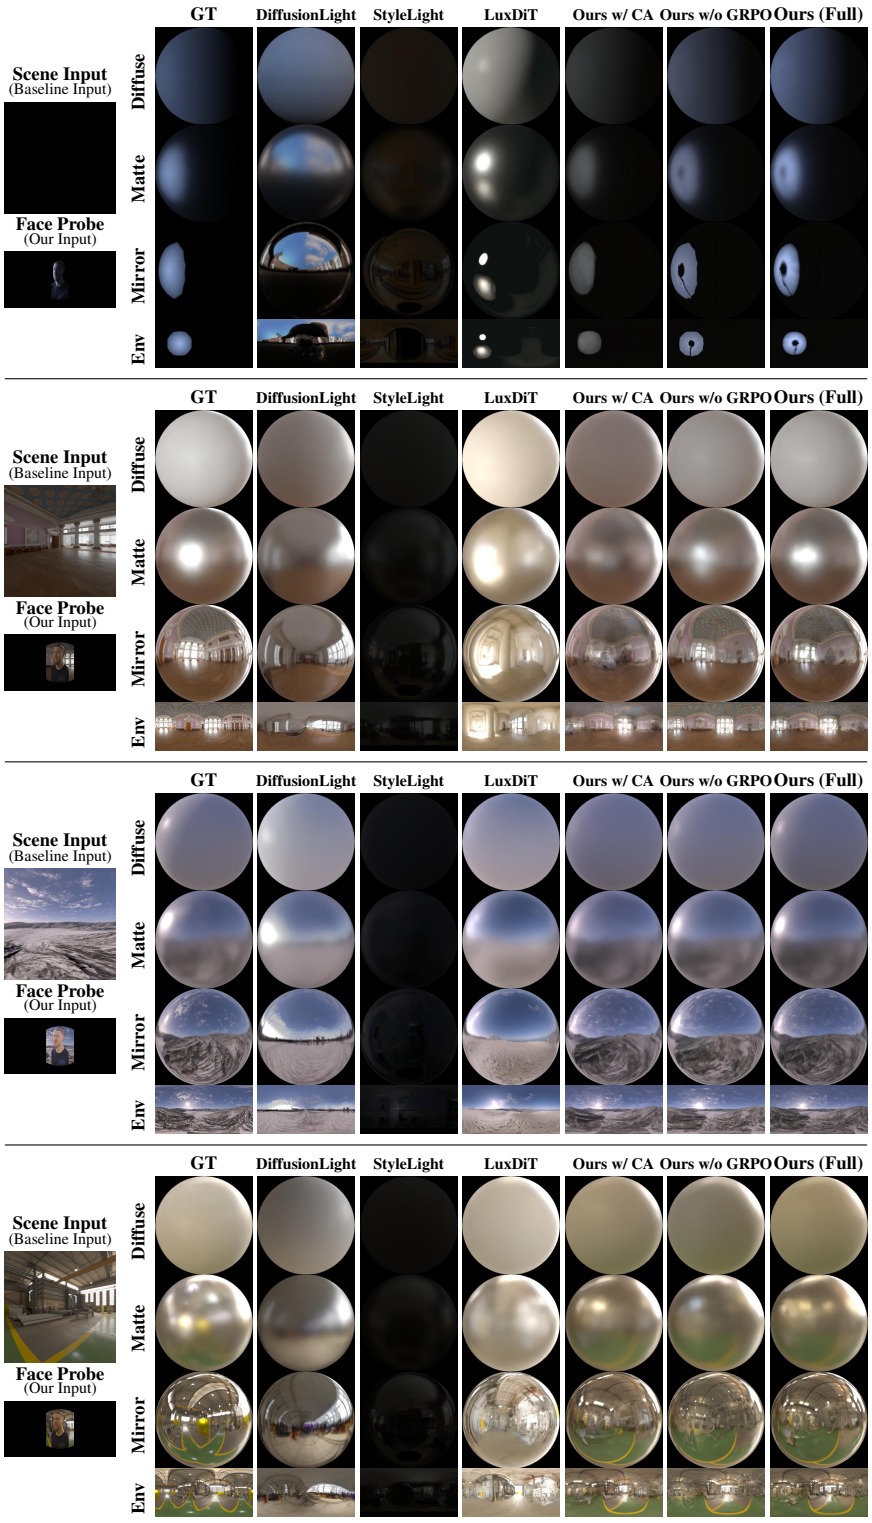

Figure 13: Extended qualitative comparisons of relighting on virtual spheres under a unified fixed-exposure linear HDR rendering protocol (extending Figure 3 of the main text). Our method accurately captures both zero-context facial lighting cues and complex indoor/outdoor illumination distributions while maintaining strict energy consistency.

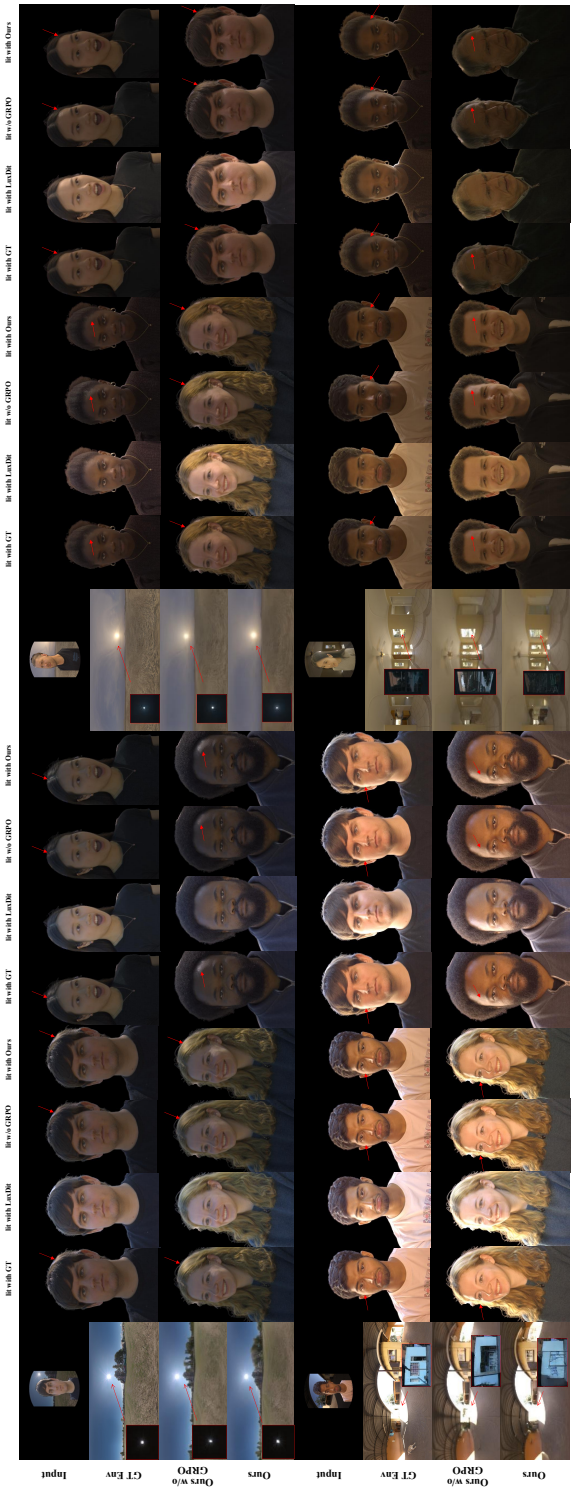

Figure 14: Enlarged, full-page high-resolution view of the portrait relighting comparison from Figure 7 of the main paper.

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 main claims regarding cue-driven masked panoramic completion and physics-aware alignment via OLAT-GRPO are clearly stated in the Abstract and Section 1, and are comprehensively supported by experimental evaluations in Section 4 and Appendices B–F.

   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: We explicitly discuss the methodological limitations (such as reconstructing infinitely sharp synthetic light source boundaries) and future directions in Section 5.

   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: This paper proposes an empirical and methodological framework for inverse rendering and lighting estimation, and does not include formal theoretical theorems or 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: We provide complete architectural details, mathematical formulations, reward definitions, dataset split protocols, and training hyperparameters in Section 3, Section 4, and Appendix A, along with evaluation and rendering protocols in Appendices B–D.

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 full dataset includes privately captured LightStage OLAT subject bases and commercially sourced third-party HDR scenes that cannot be publicly redistributed due to privacy and licensing restrictions; however, we will make the code publicly available on GitHub in the future.

   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: All experimental settings, including strictly disjoint subject/HDR data splits (Section 4), LoRA and DiT configurations (Section 3.1), OLAT-GRPO reward weights and optimization hyperparameters (Section 3.3 and Appendix A), and evaluation protocols (Appendices B–D), are specified in detail.

   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: We report 95% bootstrap confidence intervals of the mean (10,000 resamples) on the independent third-party benchmark in Section 4.4 (Table 4) and paired statistical significance tests (p-values) for controlled ablations and matched-input comparisons in Section 4.5, Appendix C, and Appendix D.

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: [Yes]

   Justification: We specify the compute hardware (two NVIDIA A100 GPUs), numerical precision (bfloat16), optimizer settings, and exact training step counts for both task-specific adaptation (30K steps) and OLAT-GRPO alignment (2K steps) in Appendix A.

   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 conducted in this paper conforms in every respect with the NeurIPS Code of Ethics.

   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: [N/A]

    Justification: This work focuses on foundational inverse rendering and 360° HDR environment lighting estimation for computer graphics and virtual production; predicting scene illumination maps does not introduce direct or unique negative societal impacts beyond general rendering technologies.

    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: [N/A]

    Justification: Our framework estimates 360° HDR environment illumination maps rather than releasing unrestricted human-identity generators or uncurated scraped datasets, and thus poses no high risk for misuse requiring specialized safeguards.

    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: [Yes]

Justification: All existing open-source models and evaluation codebases (Wan2.1, DiffusionLight, StyleLight, LuxDiT) as well as public HDR environment assets (PolyHaven) are properly credited and cited in Section 3, Section 4, and the References in compliance with their respective licenses.

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: This paper does not publicly release new standalone datasets or human-subject asset collections.

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: [N/A]

Justification: This research does not involve crowdsourcing or new experiments with human subjects requiring participant instructions or 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: [N/A]

    Justification: This work conducts computational experiments using pre-existing OLAT basis captures and HDR environment maps and does not involve new human-subject recruitment or procedures requiring IRB approval.

    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: Large Language Models (LLMs) were not used as any component of the core methodology, architecture, or experimental development in this research.

    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.