EYELINE LABS

SIGGRAPH 2026
Los Angeles

Promoting LDR Panoramas to HDRI with Rendering-Space Losses

Zhengming Yu[1,3], Li Ma[1], Mingming He[1], Jingdong Zhang[3], Jionghao Wang[3], Ning Yu[1,2], Paul Debevec[1,2]

1. Eyeline Labs · 2. Netflix · 3. Texas A&M University

Problem

- HDRIs are standard light sources for image-based lighting, but casual LDR panoramas **clip at the saturated area**.
- Clipped highlights destroy **diffuse** irradiance and **specular** reflections.
- Pixel losses treat all errors similarly, even though rendering quality is dominated by lighting integrals.

Related Work

- Single-image HDR reconstruction improves local pixels, but does not directly supervise IBL behavior.
- DiffHDR shows the value of video diffusion priors for exposure recovery.

Missing: losses aligned with how HDRIs are rendered.

Our Approach

- Fine-tune Wan 2.1 VACE 14B with LoRA for clipped-HDRI promotion.
- Use mask-conditioned inpainting of saturated highlight regions.
- Add diffuse SH and glossy GGX rendering-space losses for IBL fidelity.
- Train in linear radiance with log encoding for VAE compatibility.

Method

Clipped LDR panorama + mask → **HDRI-Finetuned Wan 2.1 VACE** → **Predicted linear HDRI panorama**

Rendering-space supervision

Diffuse SH loss	Glossy GGX loss
low-frequency irradiance	specular IBL response

$$\mathcal{L}_{\text{diffuse}} = \frac{1}{27}\sum_{c,k}\left(\log(1+c_k^{\text{pred}}) - \log(1+c_k^{\text{gt}})\right)^2 \qquad \mathcal{L}_{\text{glossy}} = \left\|\log(1+\mathbf{E}_{\text{glossy}}^{\text{pred}}) - \log(1+\mathbf{E}_{\text{glossy}}^{\text{gt}})\right\|_2^2$$

$$\mathcal{L} = \mathcal{L}_{\text{noise}} + \lambda_d\,\mathcal{L}_{\text{diffuse}} + \lambda_g\,\mathcal{L}_{\text{glossy}}, \qquad \lambda_d = \lambda_g = 0.05$$

- The LDR panorama is encoded as a single-frame video input. Clipped regions are mask-conditioned for inpainting.
- Diffuse SH compares 27 second-order coefficients in log space, directly supervising irradiance.
- Glossy GGX compares prefiltered specular lobes over 64 × 128 HDRI maps.

Results

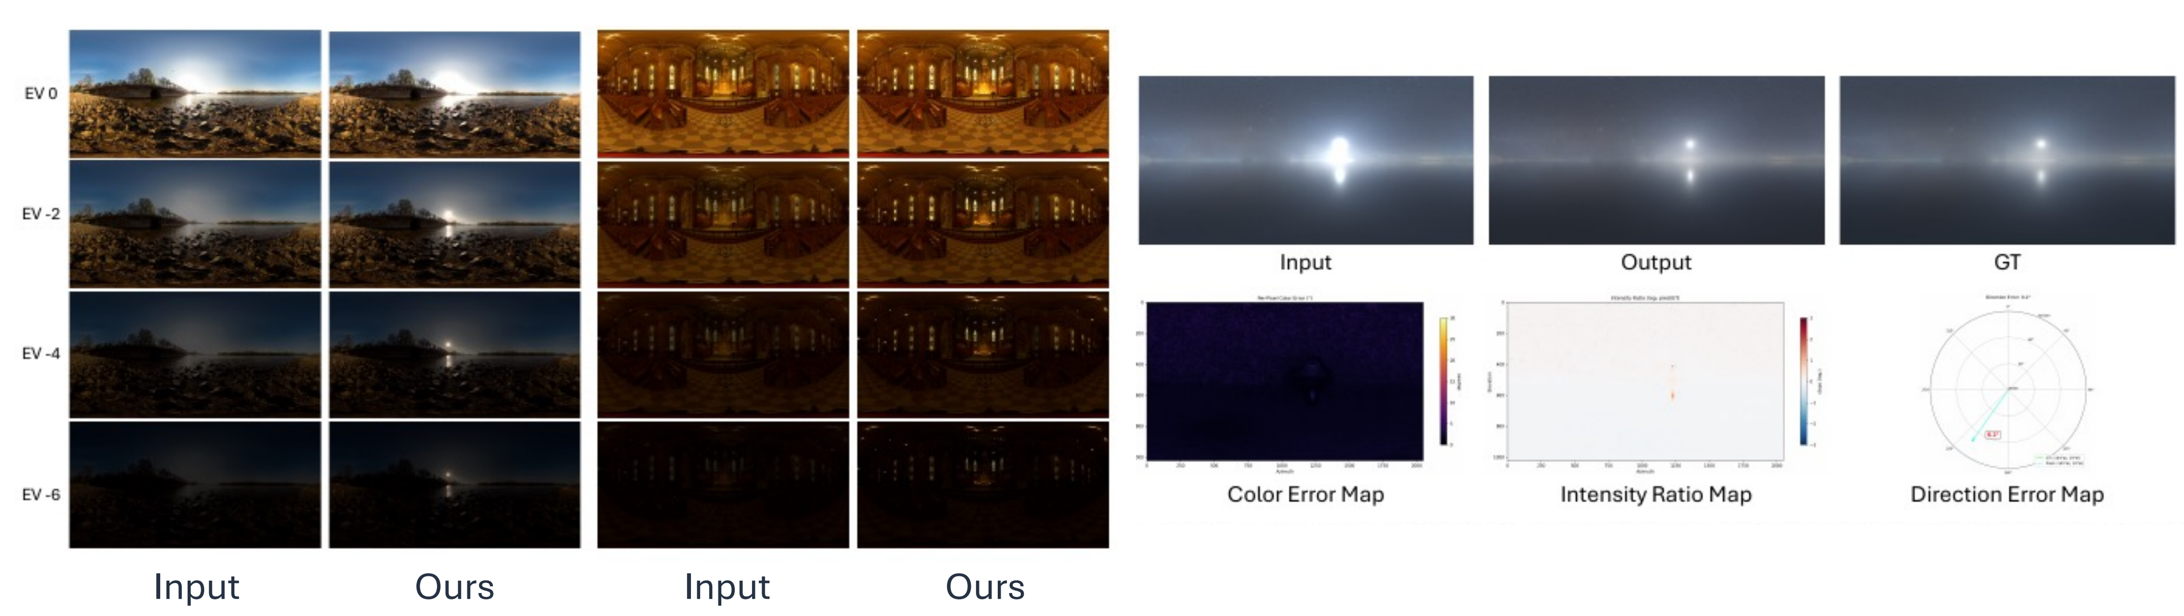

Fig.1 The model recovers clipped highlight details from LDR panorama input and achieve low color error, direction error, and close intensity ratio with the Ground Truth HDRI.

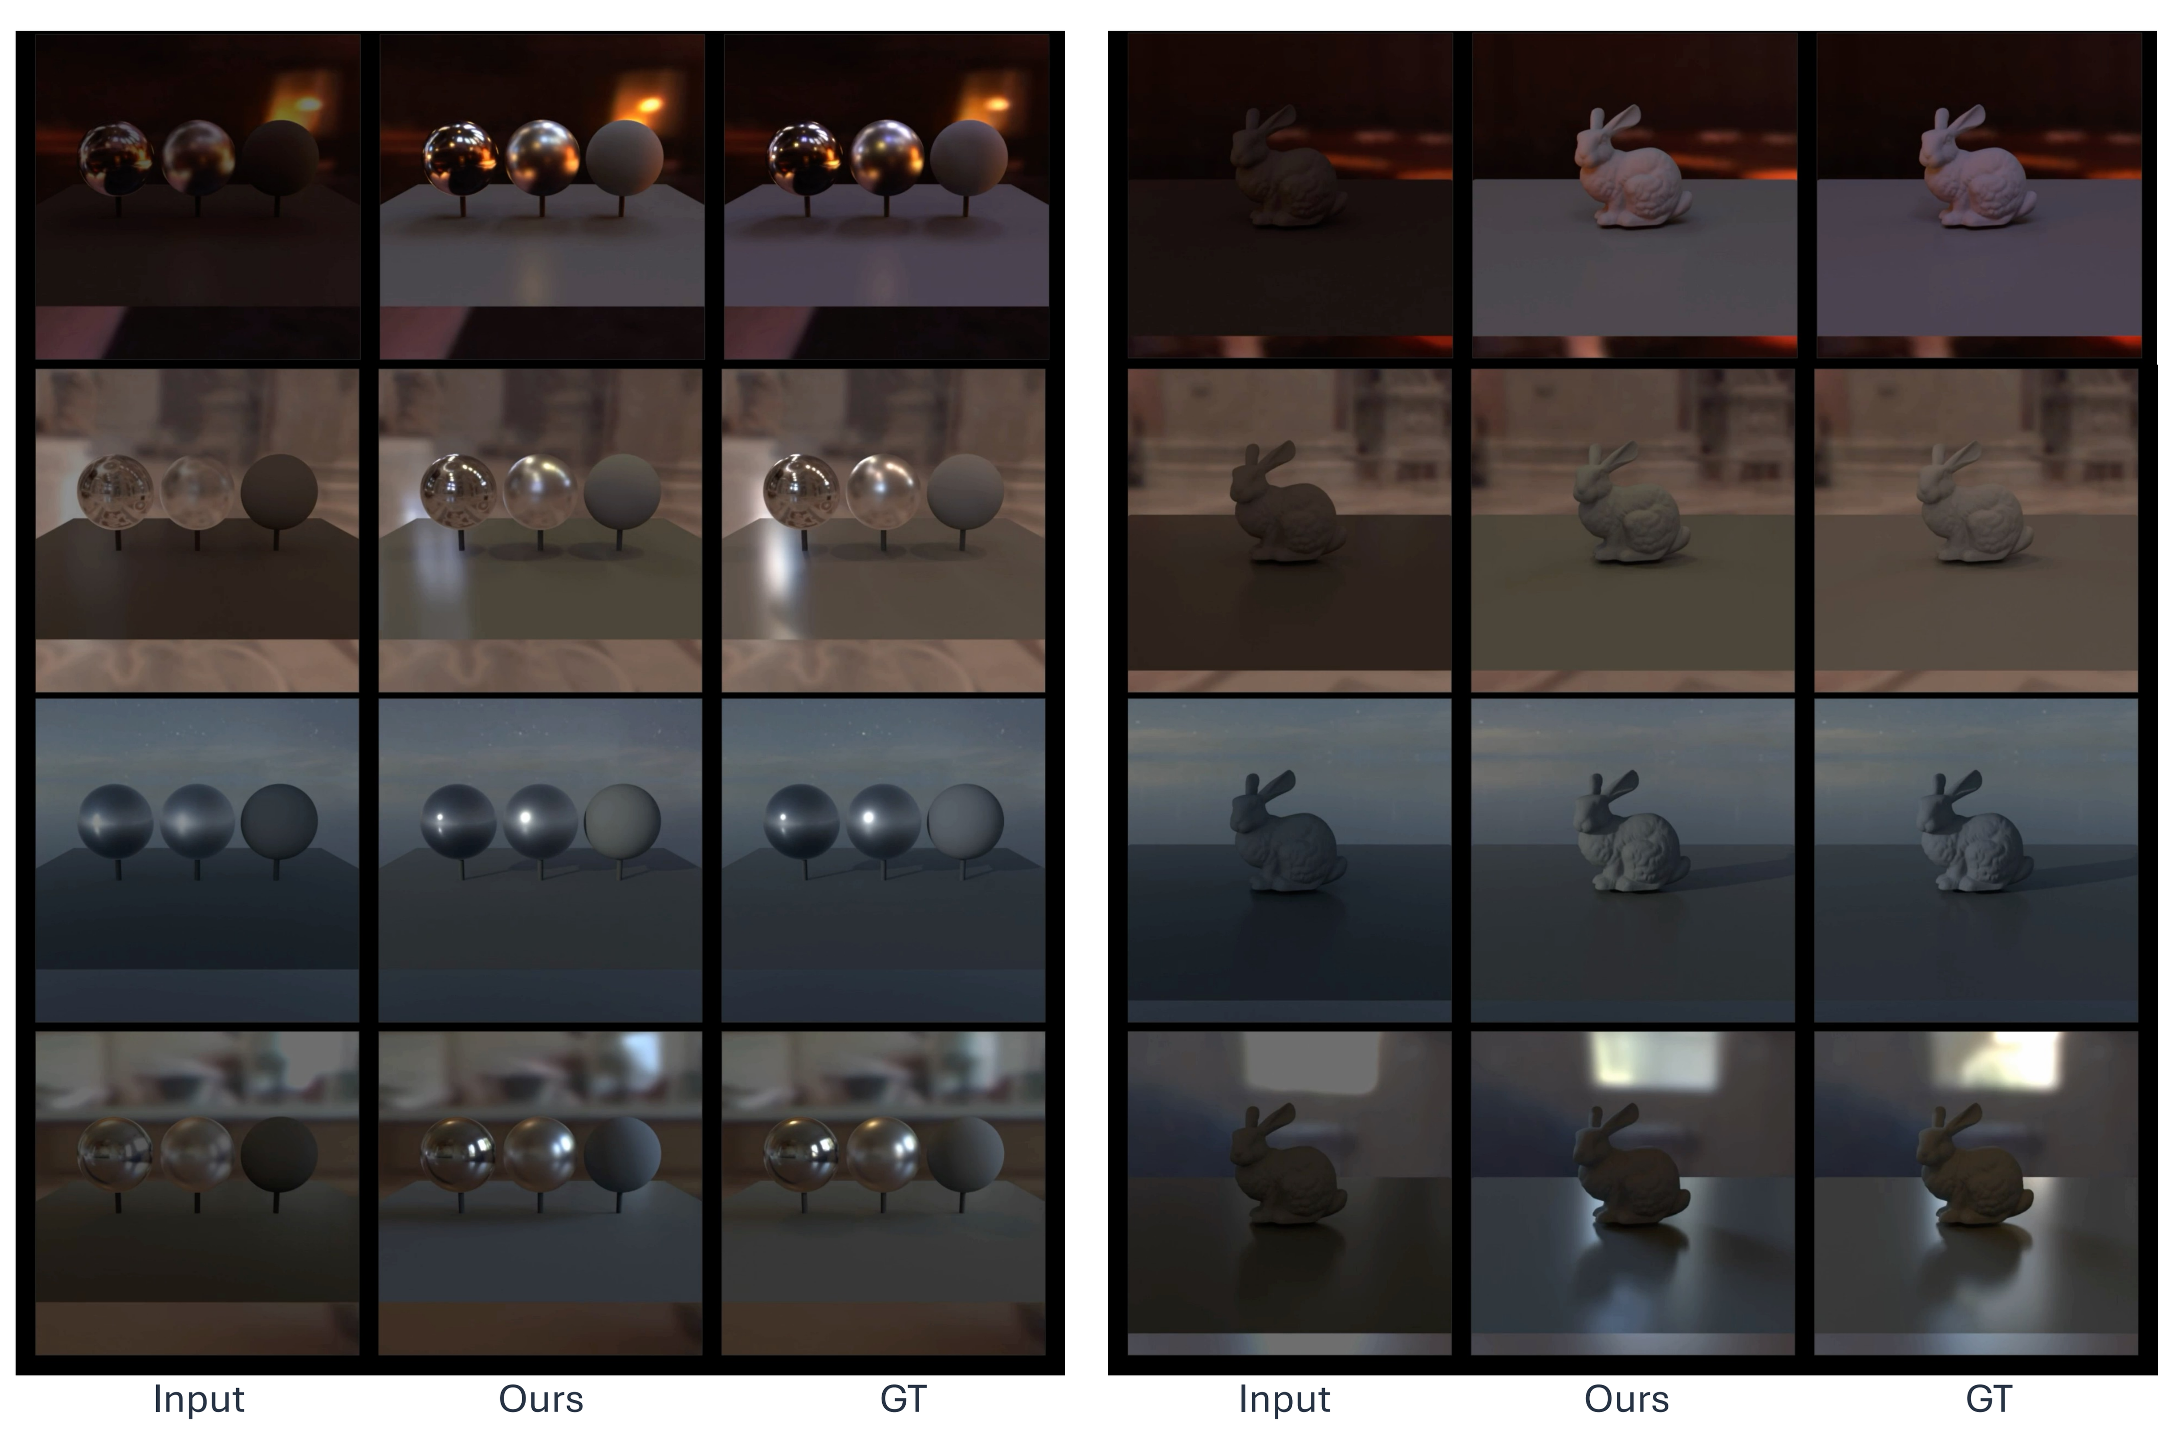

Fig.2 Relighting comparison. Our reconstruction recovers the illumination intensity, color, shadowing, and specular reflections needed for realistic relighting.

References

[1] Poly Haven. 2024. https://polyhaven.com/
[2] Debevec. Rendering synthetic objects into real scenes. SIGGRAPH 2008 Classes.
[3] Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In Seminal Graphics Papers: Pushing the Boundaries, Volume 2. 643–652.
[4] Eilertsen et al. HDR image reconstruction from a single exposure. TOG 2017.
[5] Hu et al. Single-image HDR reconstruction by learning reverse camera pipeline. CVPR 2020.
[6] Hu et al. LoRA: Low-rank adaptation of large language models. ICLR 2022.
[7] Wan Team. Wan: Open and advanced large-scale video generative models. 2025.
[8] Ramamoorthi and Hanrahan. Efficient irradiance environment maps. SIGGRAPH 2001.
[9] Walter et al. Microfacet models for rough surfaces. Rendering Techniques 2007.
[10] Jiang et al. Vace: All-in-one video creation and editing. ICCV 2025. 17191–17202.
[11] Yu et al. DiffHDR: Re-Exposing LDR Videos with Video Diffusion Models. ECCV 2026.