Through the Looking Glass

Refractive Gaussian Splatting for Transparent Container Reconstruction

TUDelft

CSE3000 Research Project

Joan Gayán Llanos jgayan@tudelft.nl Supervisor: Michael Weinmann Examiner: Emir Demirović EEMCS, Delft University of Technology

Physically modelling refraction lets 3D Gaussian Splatting reconstruct objects inside transparent containers, where standard 3DGS collapses into floater artifacts.

M Method at a glance

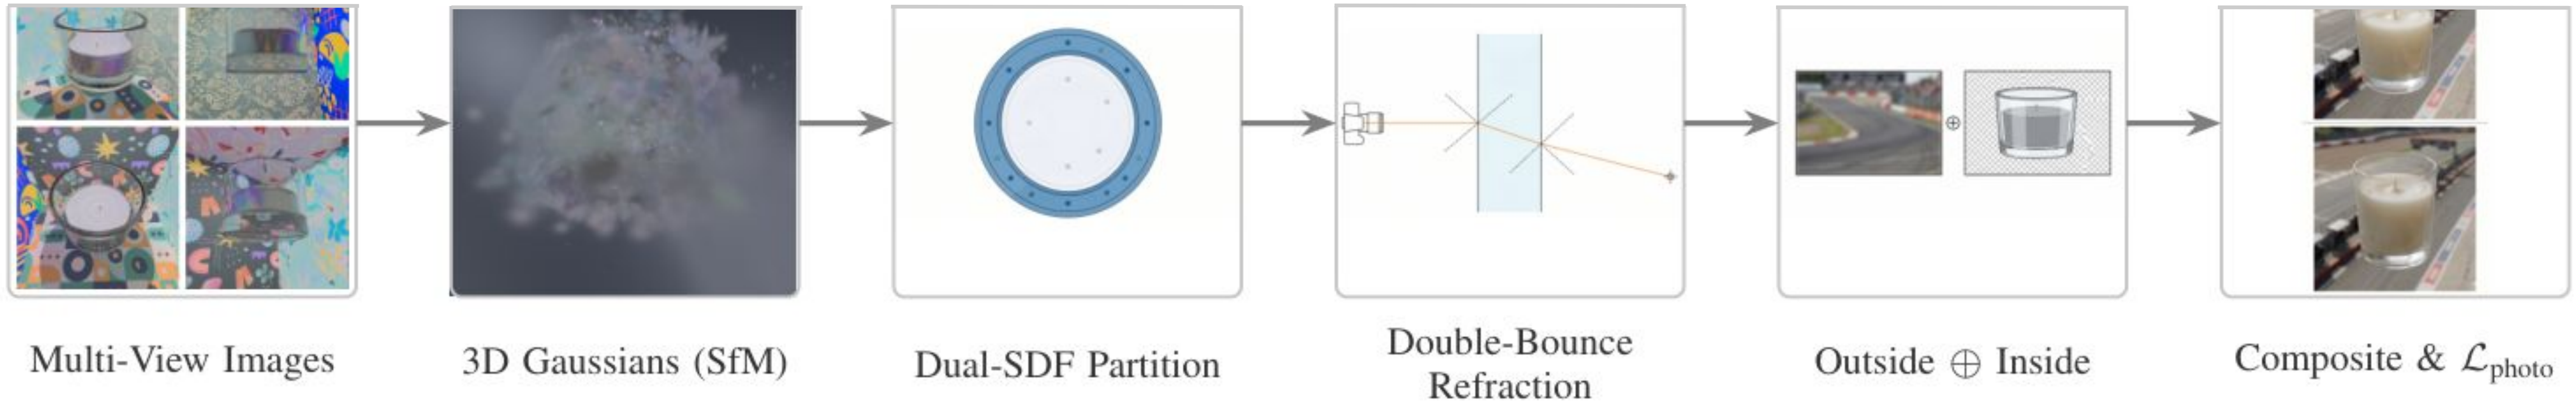

Pipeline. A dual Signed Distance Field (SDF) partitioning sorts every Gaussian into exterior, glass-wall, or cavity, routes the non-exterior ones through a double-bounce refraction engine with Jacobian covariance warping ($J\,\Sigma\,J^T$), then composites against the ground truth and back-propagates the photometric loss.

1 The problem

3DGS assumes straight rays, but glass bends them at each interface, shifting a point's apparent position per view.
The optimiser masks this conflict with dense floaters, losing the interior object.

Fig. 1. spherebg view. Left: Splatfacto, refraction hidden behind floaters. Right: ours, coherent glass.

Takeaway Straight-ray 3DGS cannot see through glass, so it masks refraction with floaters.

Research question

How can objects inside transparent containers be reconstructed efficiently and accurately with an explicit 3D Gaussian framework that physically models refraction, volumetric decay, and double-bounce light transport?

2 Dual-SDF partitioning

One Signed Distance Field (SDF) cannot separate solid glass from the hollow air cavity.
A second cavity SDF labels each Gaussian glass-wall, hollow-interior, or exterior.
Exterior takes the normal rasteriser, the rest enter the refraction engine, so background never bleeds inside.

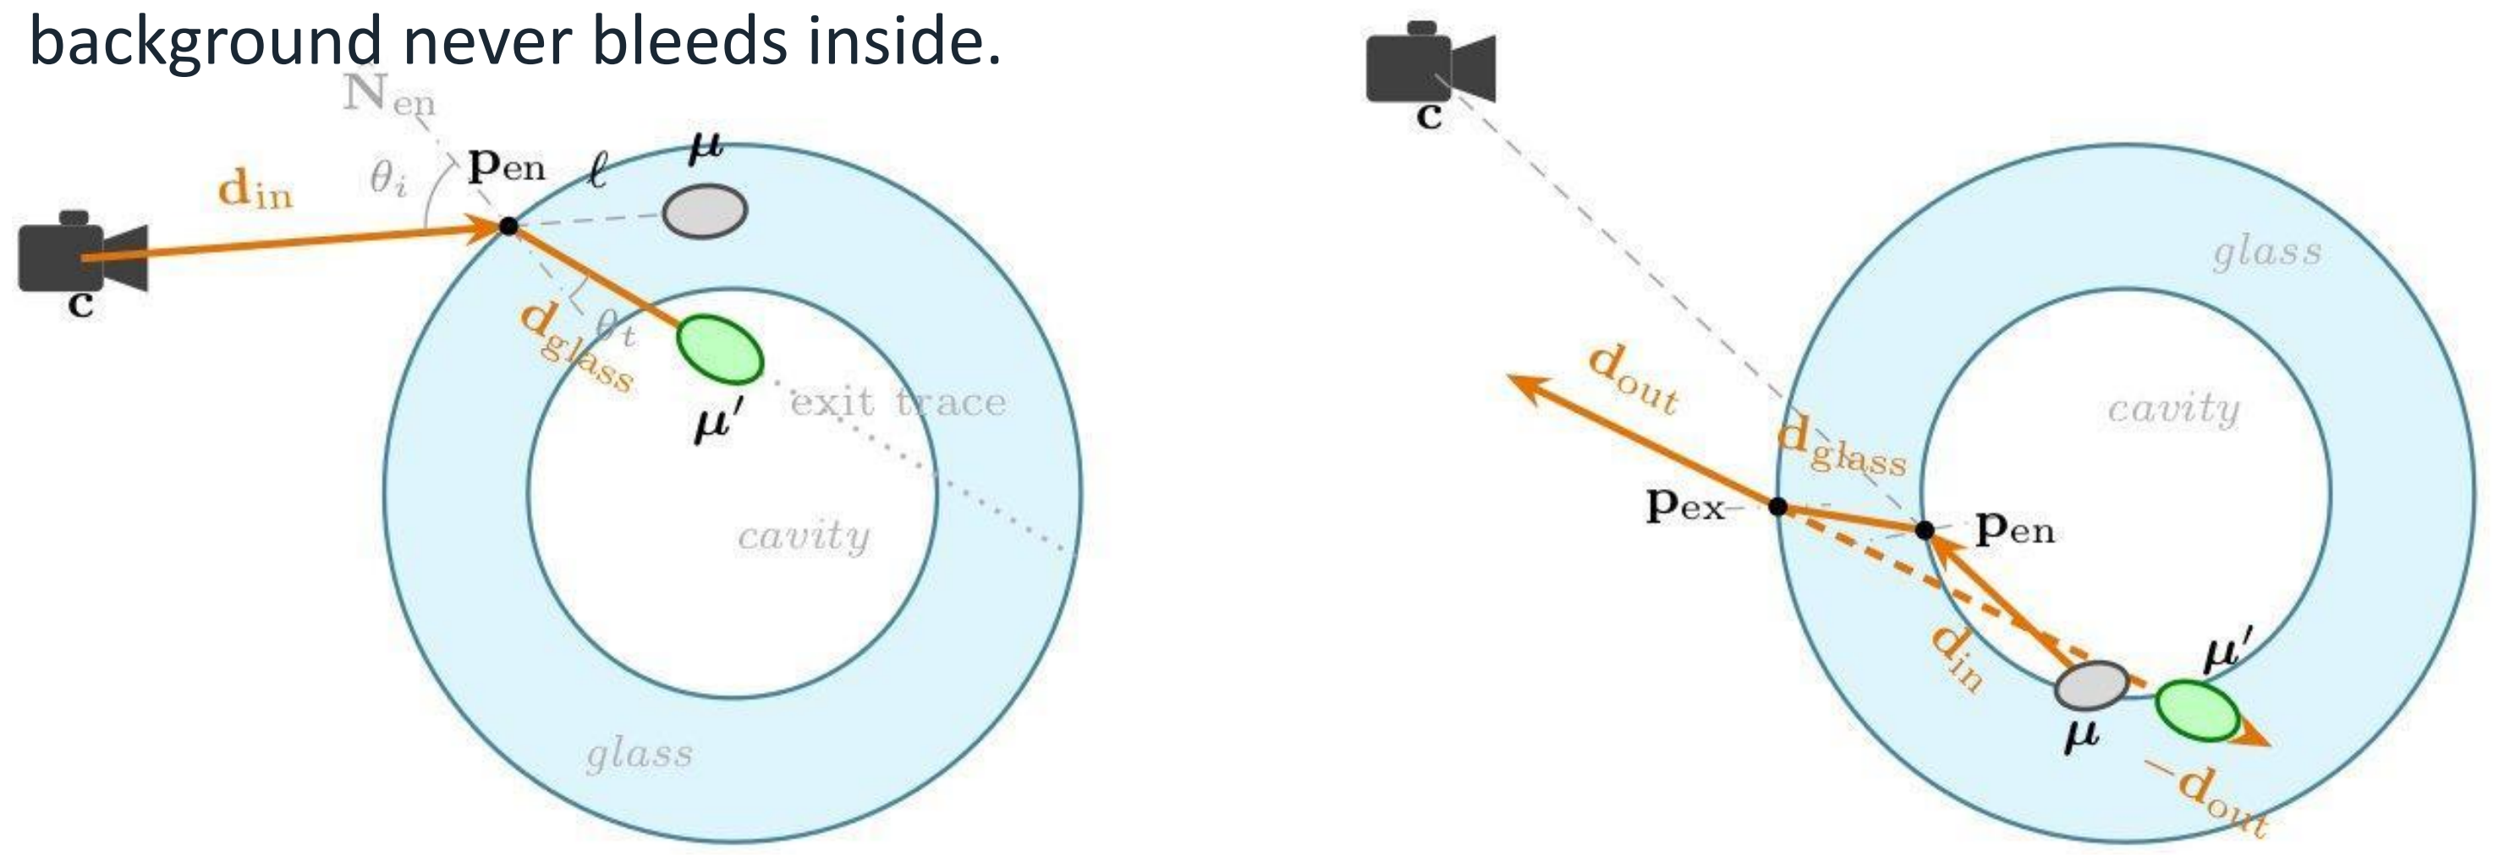

Fig. 2. Glass-wall: one bounce. Hollow-interior: two bounces.

Takeaway A dual SDF routes each Gaussian to the correct optical rule before rendering.

3 Refraction engine + stability

Snell at both interfaces moves each Gaussian to its apparent position μ', and a forward-difference Jacobian warps its covariance, $\Sigma' = J\,\Sigma\,J^\top$.
Reset only exterior Gaussians, while glass-wall Gaussians fetch the refracted background.
Fresnel transmittance and Beer-Lambert decay add physical attenuation.

Takeaway Double-bounce tracing and a Jacobian warp place and shade each Gaussian, while selective updates keep interior detail alive.

4 Results

Beats the straight-ray baseline on every metric and background. Three RefRef scenes, Splatfacto reproduced under the Oracle protocol.

	cube	sphere	natural
Splatfacto	13.02	12.76	14.93
Ours	**21.19**	**19.67**	**18.44**
Gain	**+8.2**	**+6.9**	**+3.5**

Table 1. PSNR in dB (higher is better). SSIM and LPIPS improve on every scene too.

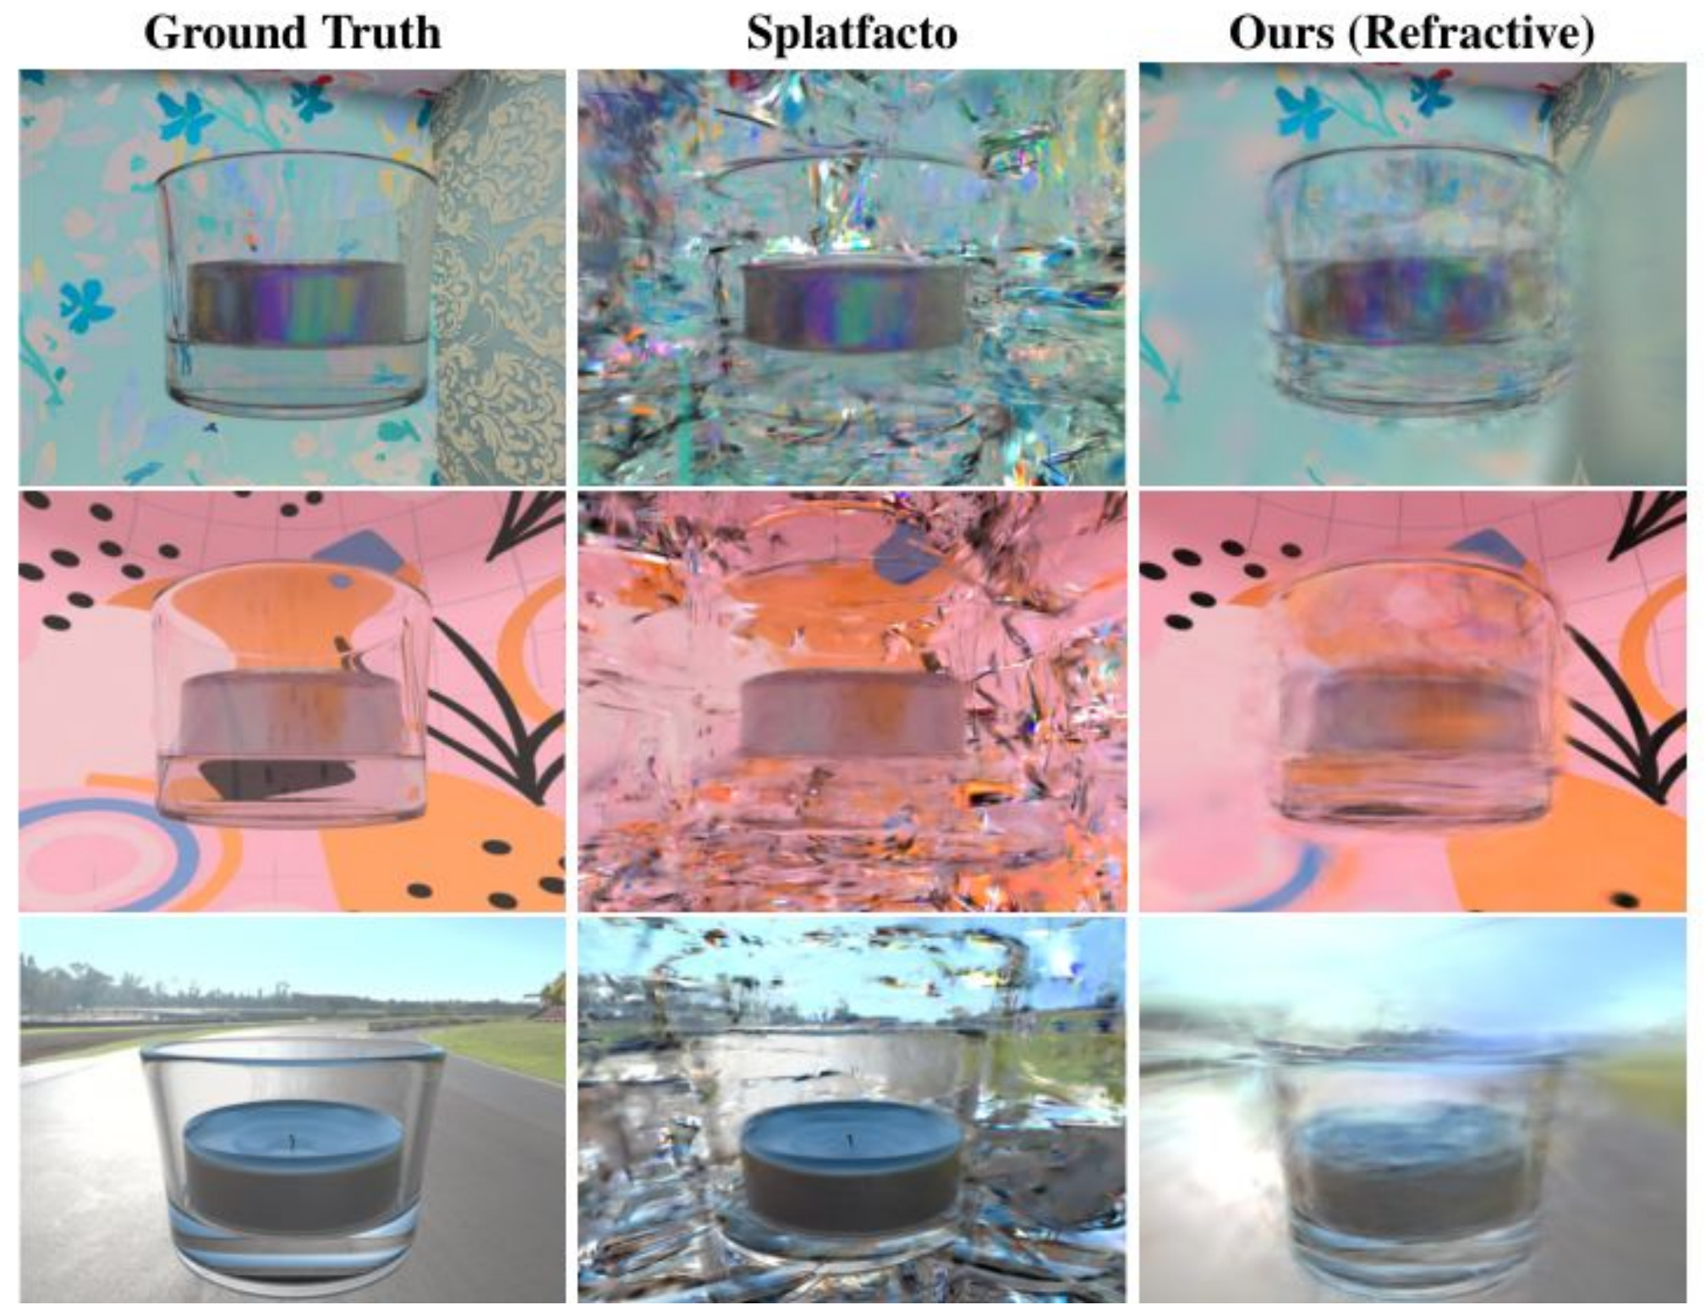

Fig. 3. GT / Splatfacto / Ours across three scenes.

Takeaway +8.2 / 6.9 / 3.5 dB PSNR, recovering glass geometry the baseline destroys.

5 Ablation, limits, conclusion

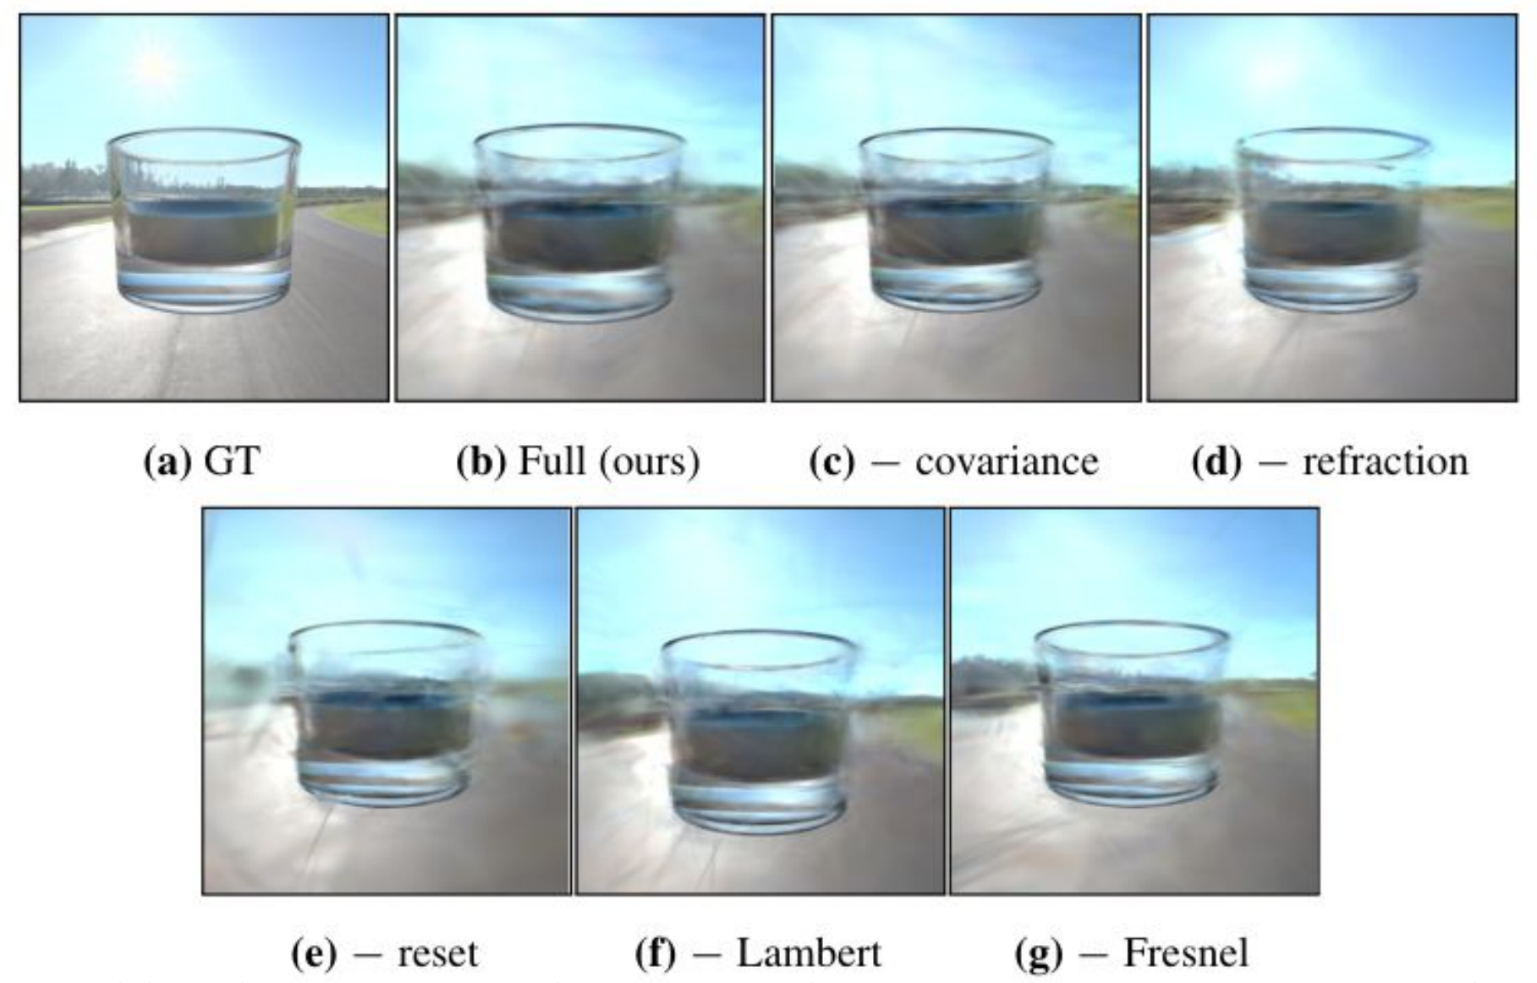

(a) GT (b) Full (ours) (c) − covariance (d) − refraction

(e) − reset (f) − Lambert (g) − Fresnel

Fig. 4. Ablation on natural background, isolating each component. Removing the refraction fetch causes the largest quality drop.

Refraction fetch is the most important component, with the largest drop when removed.
Covariance warp sharpens fine detail, and the **selective reset** preserves interior density.
Limitation: the exterior background is sampled as a 2D image, so a ray bent toward an unobserved direction has no pixel to fetch. This, not the optical model, is the residual error on the natural scene.

Takeaway Accurate refraction in an explicit Gaussian framework. The remaining gap is the 2D background, not the geometry.

Responsible research. Metrics are single runs of fixed configurations with no noticeable divergence on repeat. Code, RefRef meshes, and all training commands ship with a Kaggle notebook. Fixed glass index n = 1.5.

References. Kerbl 2023 (3DGS). Yin 2025 (RefRef). Tong 2023 (ReNeuS). Wu 2025 (3DGUT). Shao 2026 (RefracGS). Cao 2025 (GlassGaussian). Zwicker 2002 (EWA).