KU LEUVEN

PSI

WACV 2026
Tucson, Arizona • Mar 6 - 10, 2026

Spec-Gloss Surfels and Normal-Diffuse Priors for Relightable Glossy Objects

Georgios Kouros, Minye Wu, Tinne Tuytelaars

Processing, Speech and Images (PSI) group

Department of Electrical Engineering (ESAT), KU Leuven, Belgium

What is inverse Rendering?

Inverse Rendering → **Learned 3D Scene Decomposition** → **Downstream Tasks**

Training Views

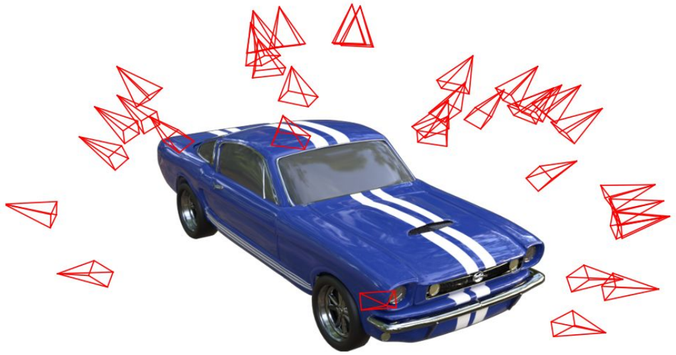

albedo rough. F0 normals

Recovered Scene Illumination

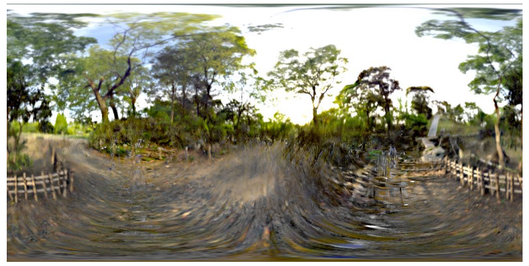

Novel View Synthesis

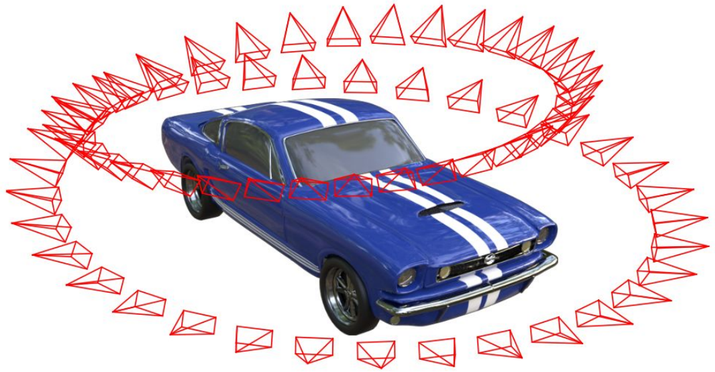

Scene Editing

Why does inverse rendering of glossy objects fail?

Inverse rendering is ill-posed because it jointly estimates

- **geometry** e.g. volumetric field, voxel grid, Gaussians
- **material properties** e.g. albedo, roughness, F_0
- **scene illumination** e.g. HDR envmap

susceptible to local minima due to ambiguity from complex view-dependent effects.

This often leads to:

- Good novel view synthesis 😀, but
- Bad relighting / scene editing 😕

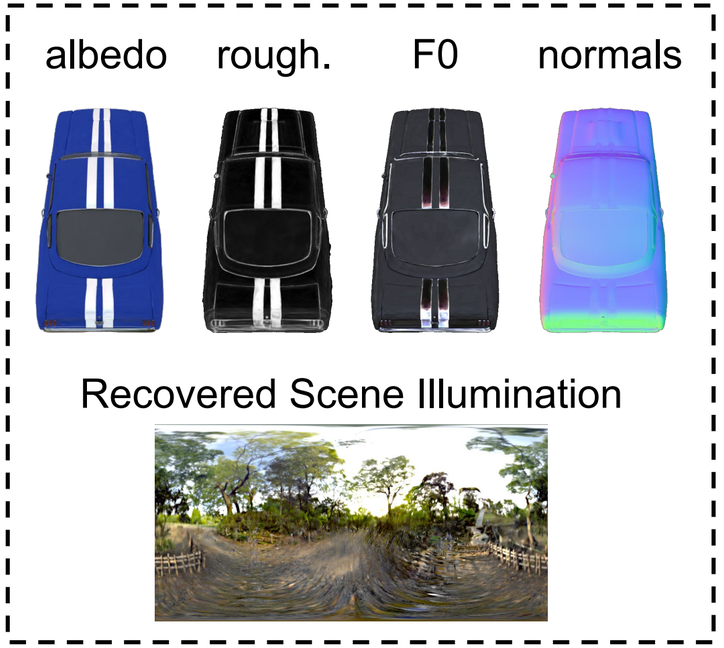

Our pipeline: 2DGS + deferred shading

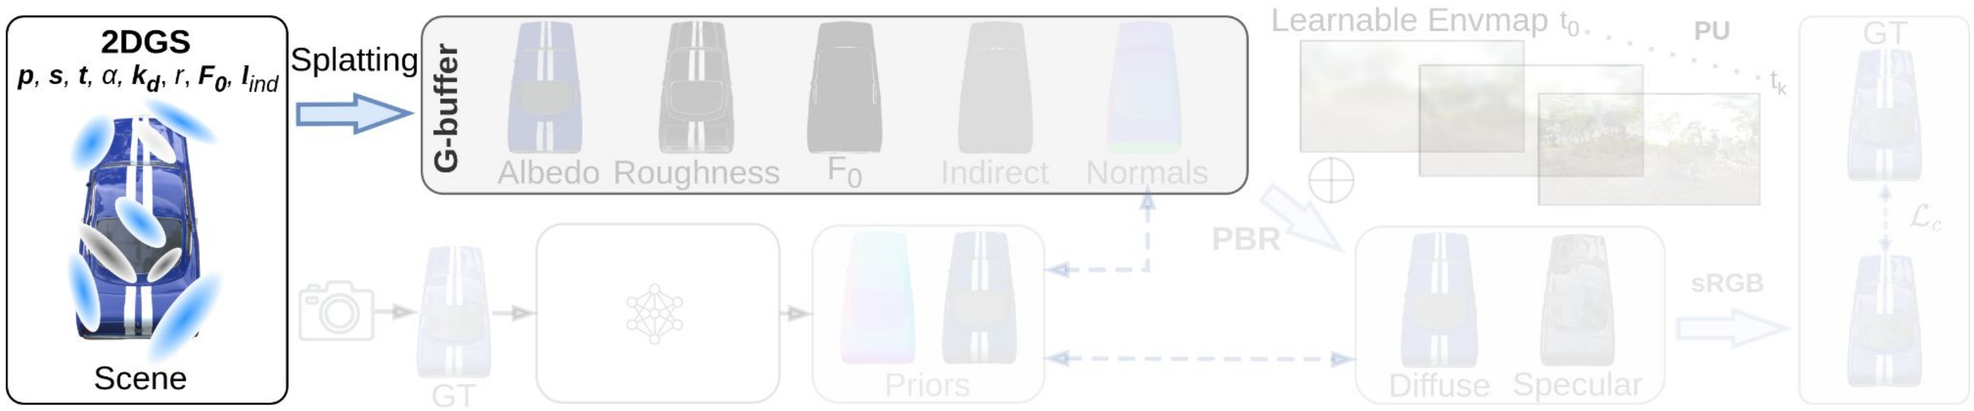

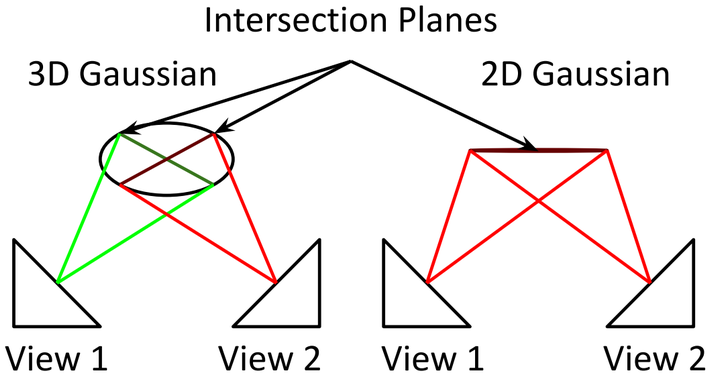

Why **2D Gaussians** and <u>not</u> **3D Gaussians**?

- View-consistent geometry
- Explicit normals as shortest axes of Gaussians

Why **deferred shading** (pixel-level) and <u>not</u> **forward shading** (Gaussian level)?

- Improved reconstruction of high-frequency reflections with less blurring
- Improved scaling on large scenes (millions of Gaussians)

Our pipeline: Spec-Gloss Material Parameterization

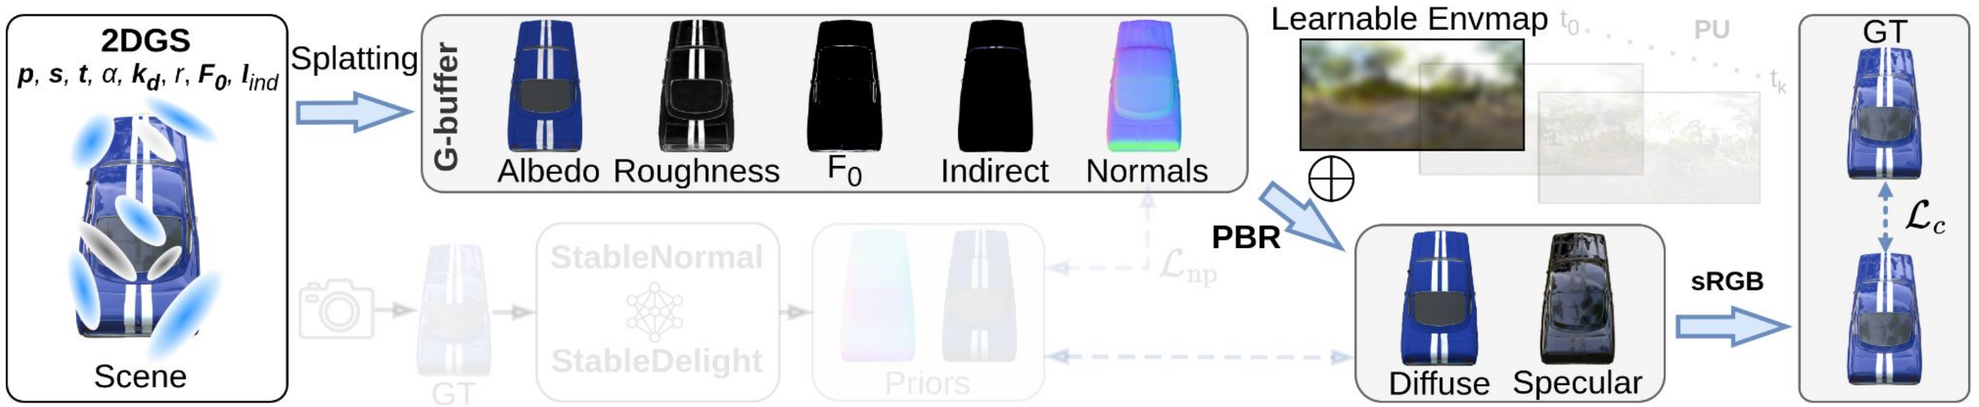

- The popular **Metallic-Roughness** (**MR**) parameterization entangles albedo / base-color (b) in specular appearance via F_0:

$$F_0(m, \boxed{b}) = (1 - m) \cdot 0.04 + m \cdot \boxed{b}$$

- The **Specular-Glossiness** (**SG**) parameterization learns F_0 as a free parameter !!!

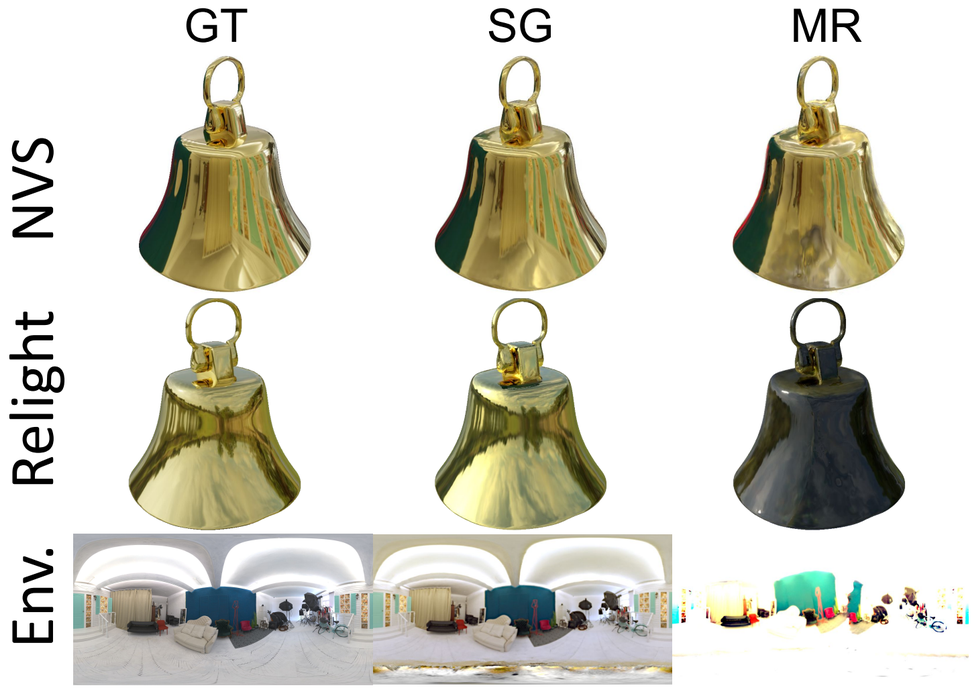

Our pipeline: Surface normal prior

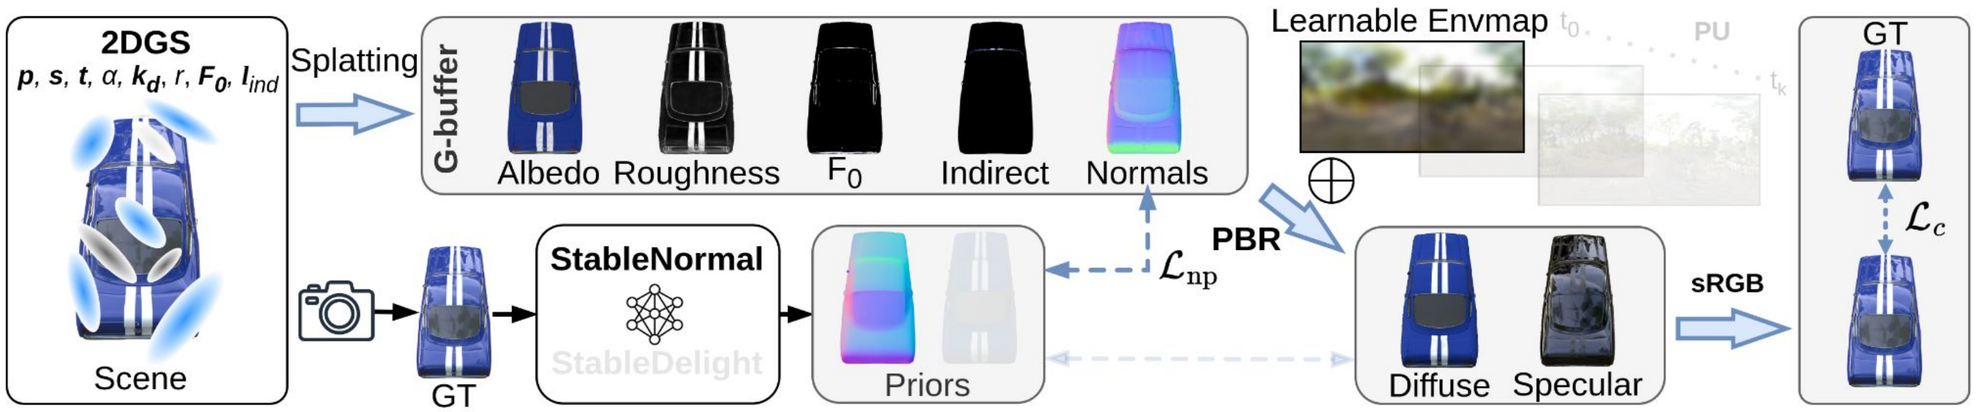

RGB input

StableNormal

Normal Prior

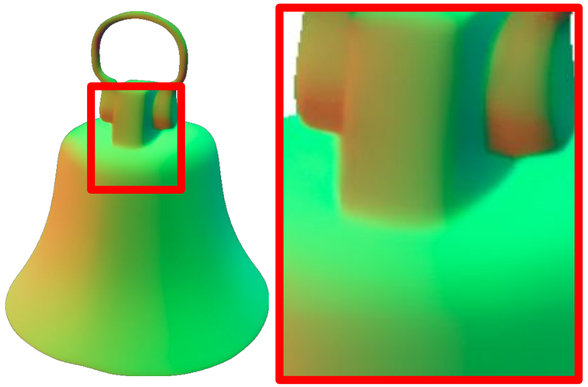

Learned Normals w/ Prior

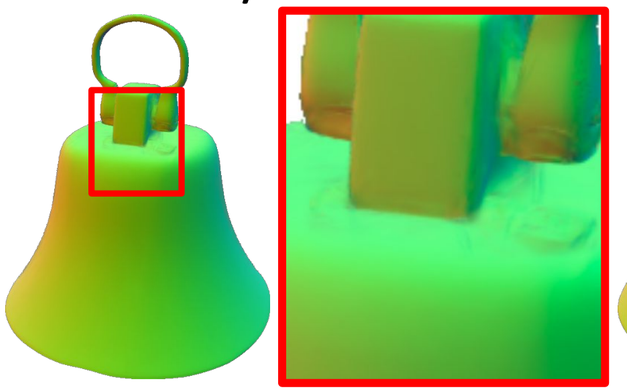

Learned Normals w/o Prior

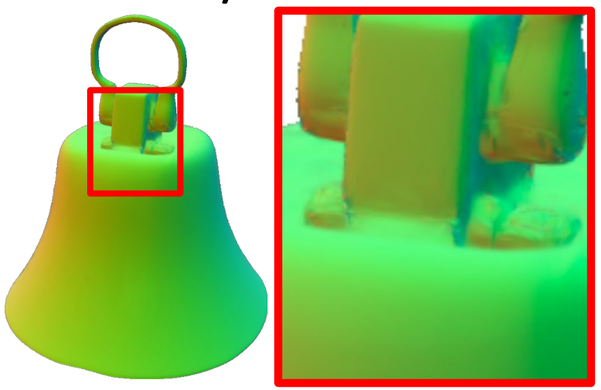

Ye, C., Qiu, L., Gu, X., Zuo, Q., Wu, Y., Dong, Z., … Han, X. (2024). StableNormal: Reducing Diffusion Variance for Stable and Sharp Normal. *ACM Transactions on Graphics (TOG)*.

Our pipeline: Diffuse color prior

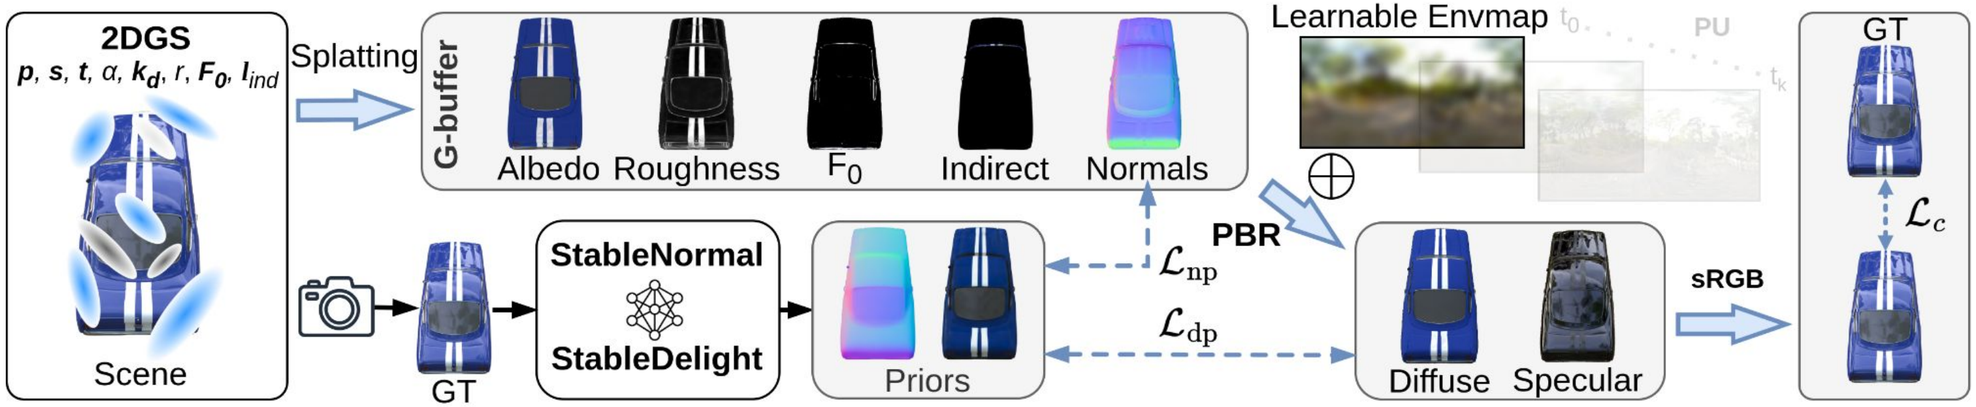

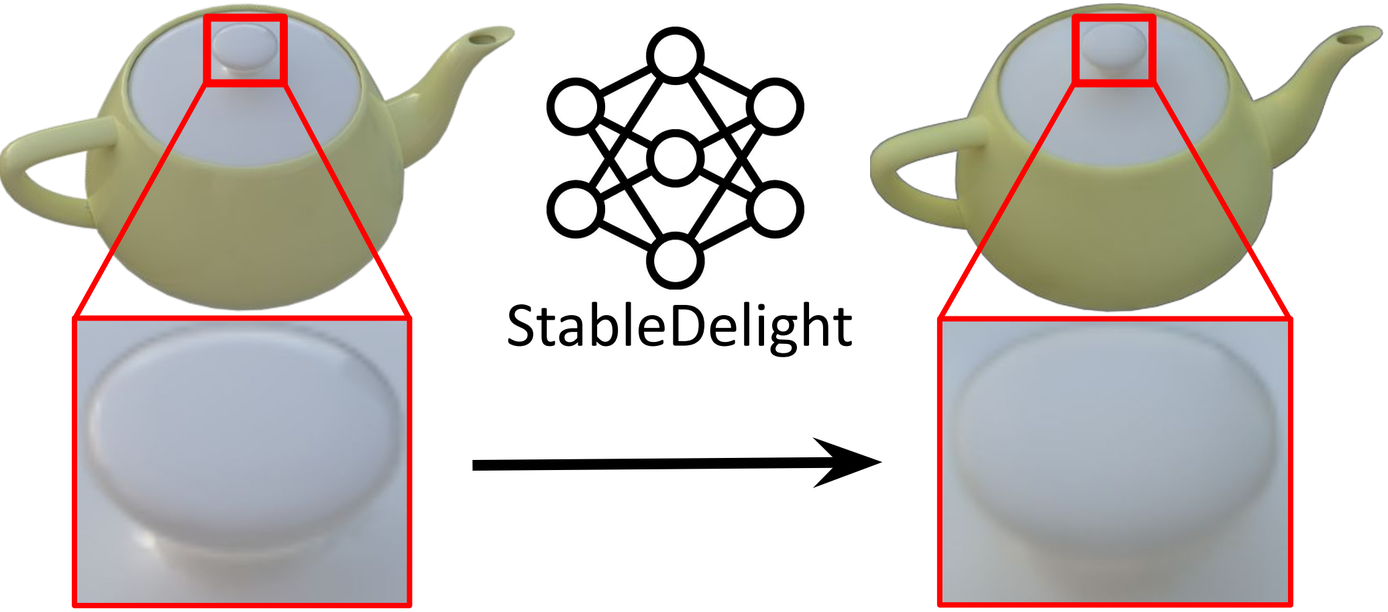

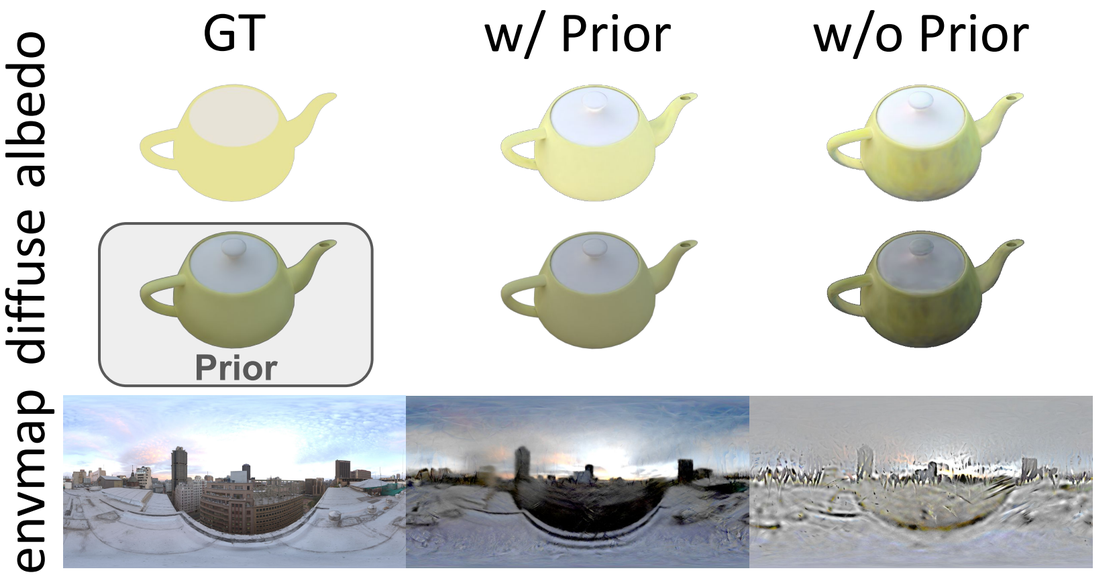

Stable-X. StableDelight: Revealing hidden textures by removing specular reflections (2025). https://github.com/Stable-X/StableDelight.

Our pipeline: Coarse-to-fine lighting optimization

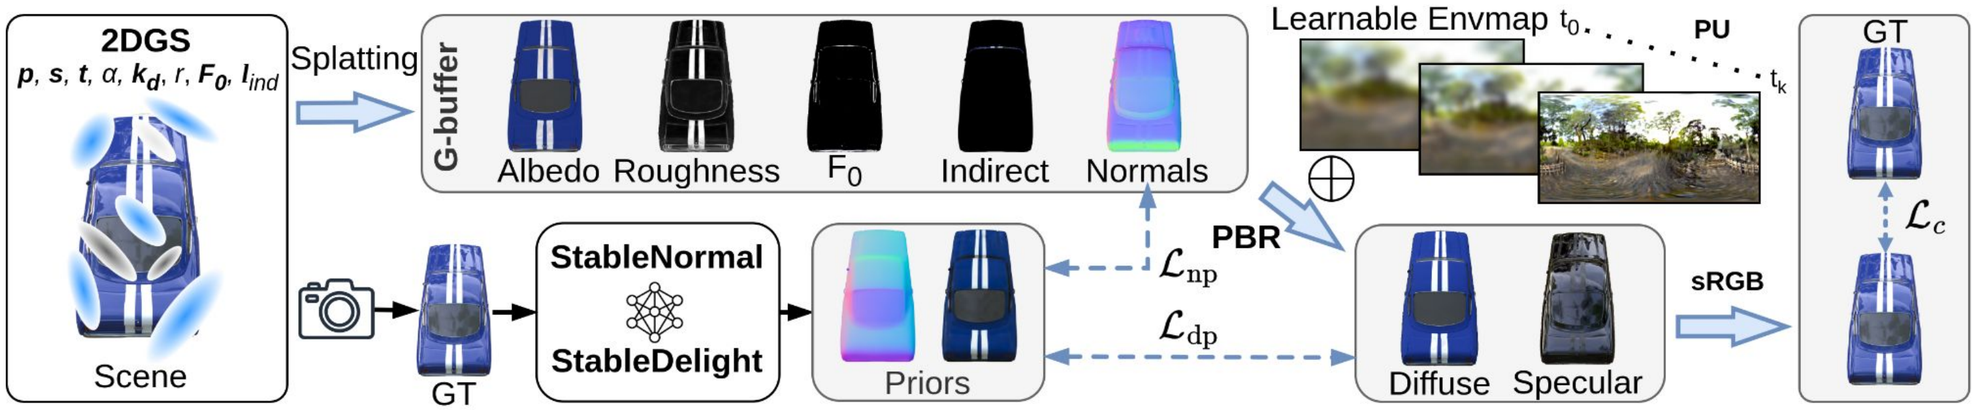

- **Envmap prefiltering** is the main bottleneck: higher res -> slower training
- Our method starts from a low resolution environment light cubemap (6×64×64) and **progressively upsamples** (**PU**) it every 15k iters until the full resolution (6×512×512)
- **PU** halves training time (2h -> 1h) without impacting envmap quality
- Facilitates recovery of **higher resolution envmaps** in reasonable time

Our pipeline: Negative-only Clipping (NOC)

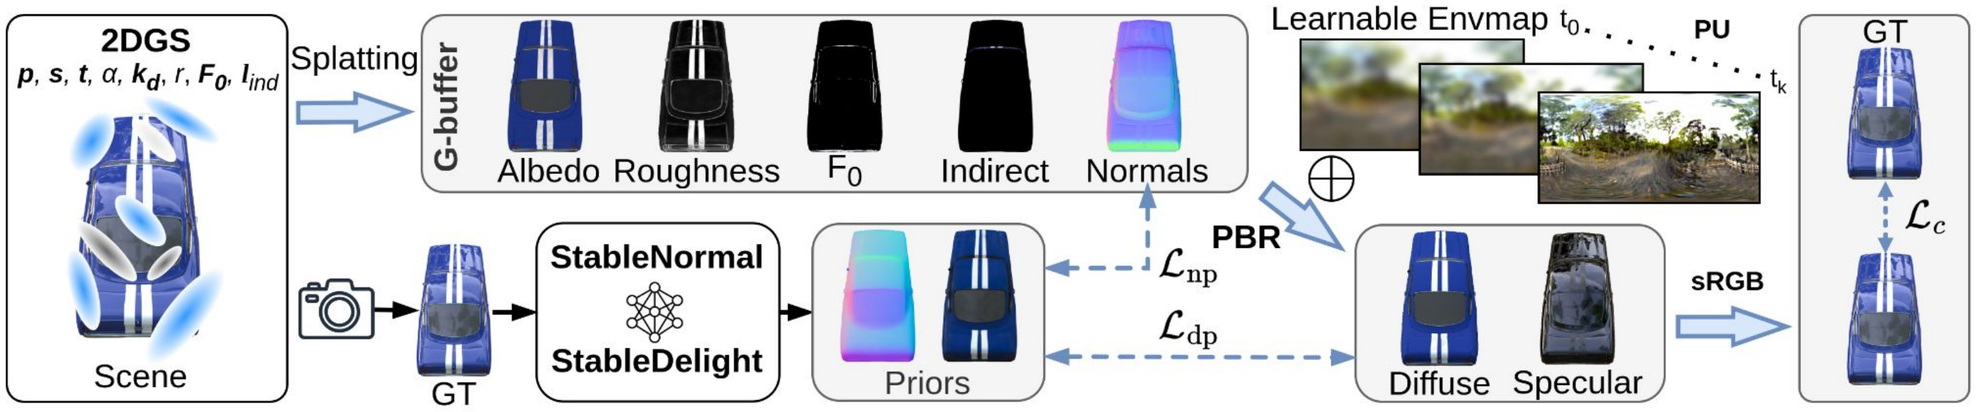

$$E \leftarrow \max(0, E_{\text{raw}})$$

Preserves sharp bright specular reflections!

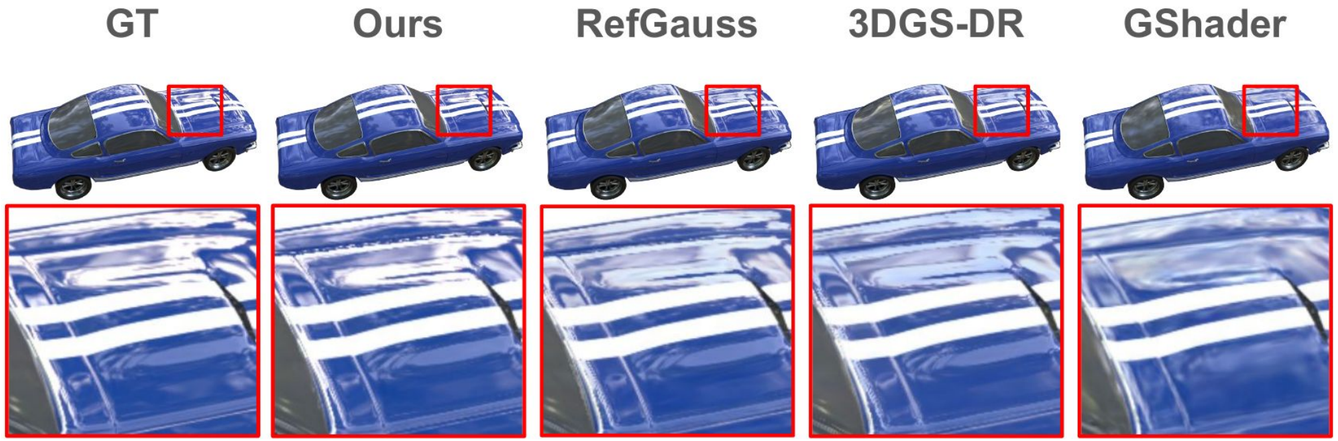

Improves albedo-lighting disentanglement!

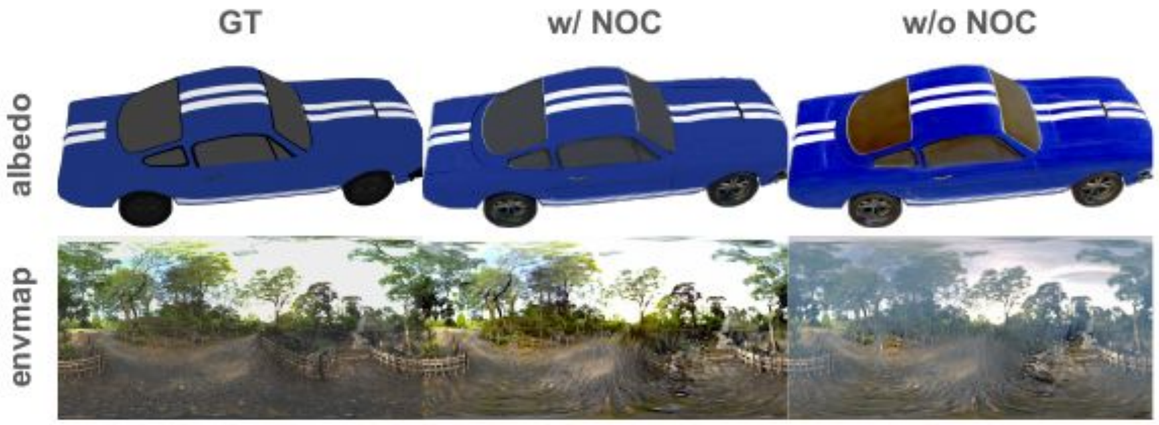

Experimental Results

Envmap Recovery

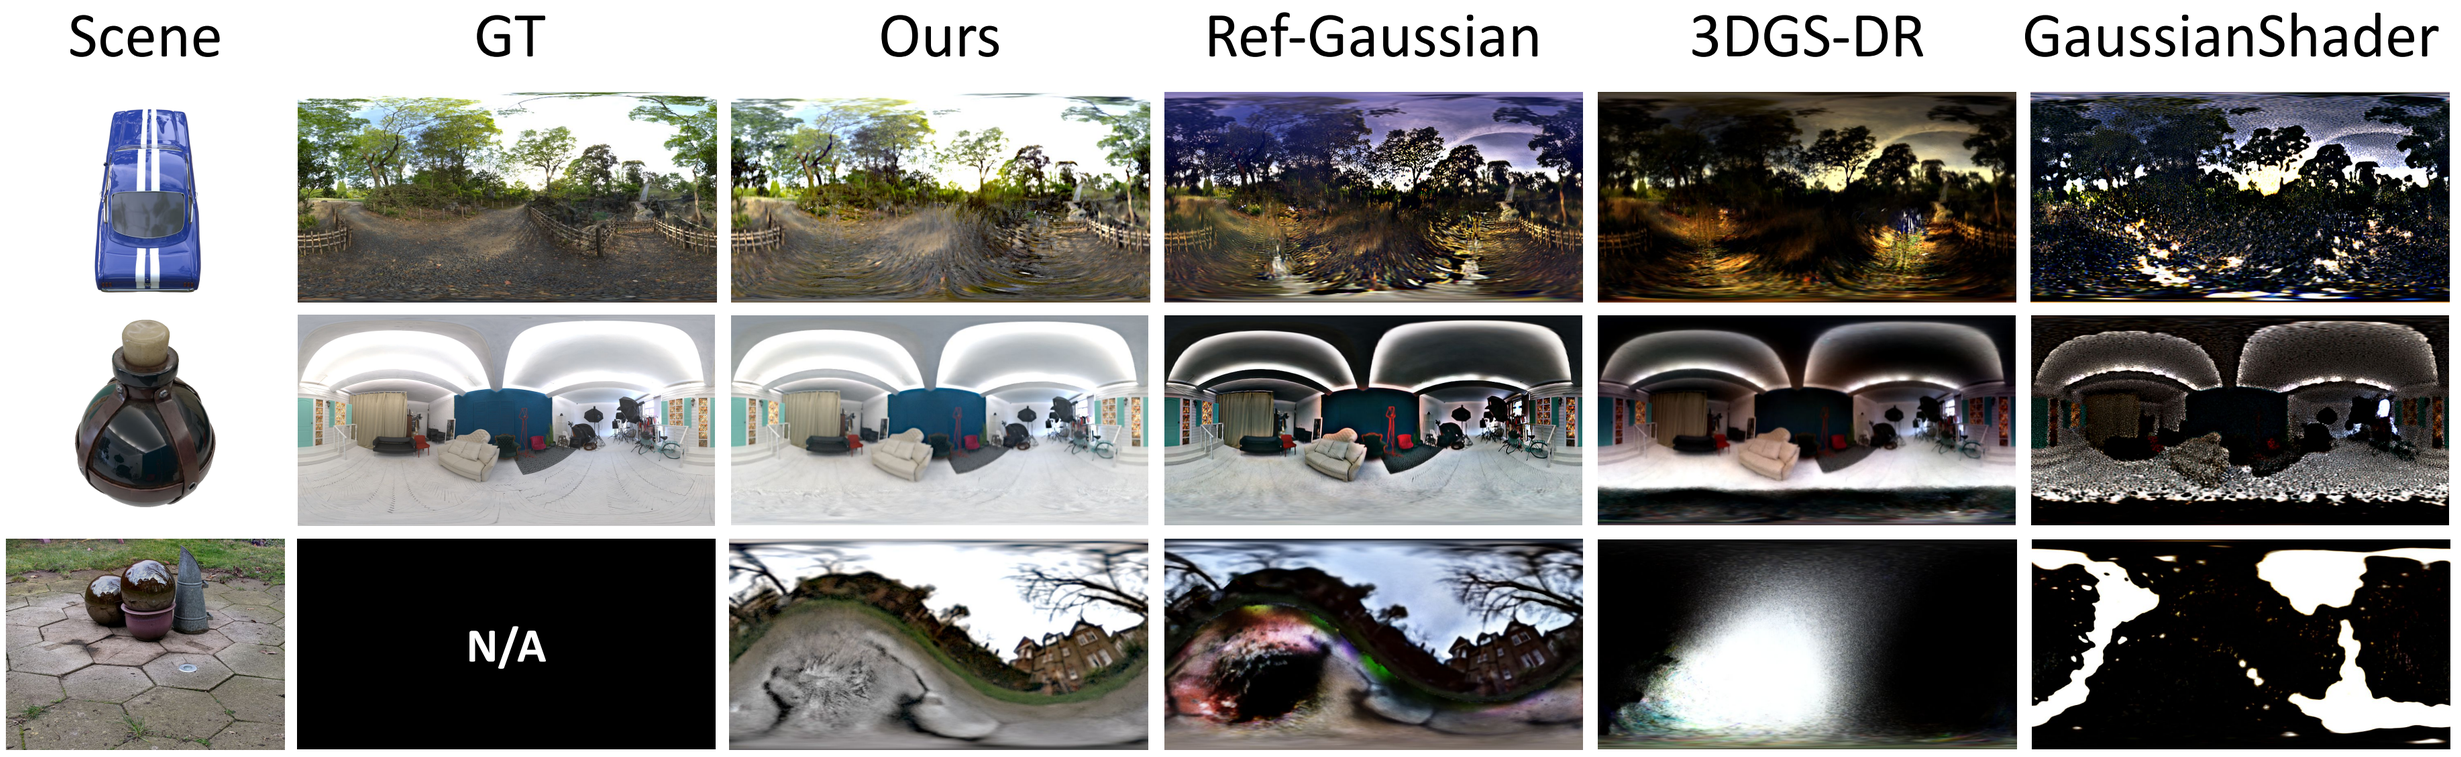

- Superior envmap recovery compared to other Gaussian Splatting methods!
- Accurate envmap → better material properties → high-quality relighting!

Relighting Results

Relighting: bell

GT

Applied Envmap

Ours

3DGS-DR

Ref-Gaussian

GaussianShader

Novel View Synthesis Results

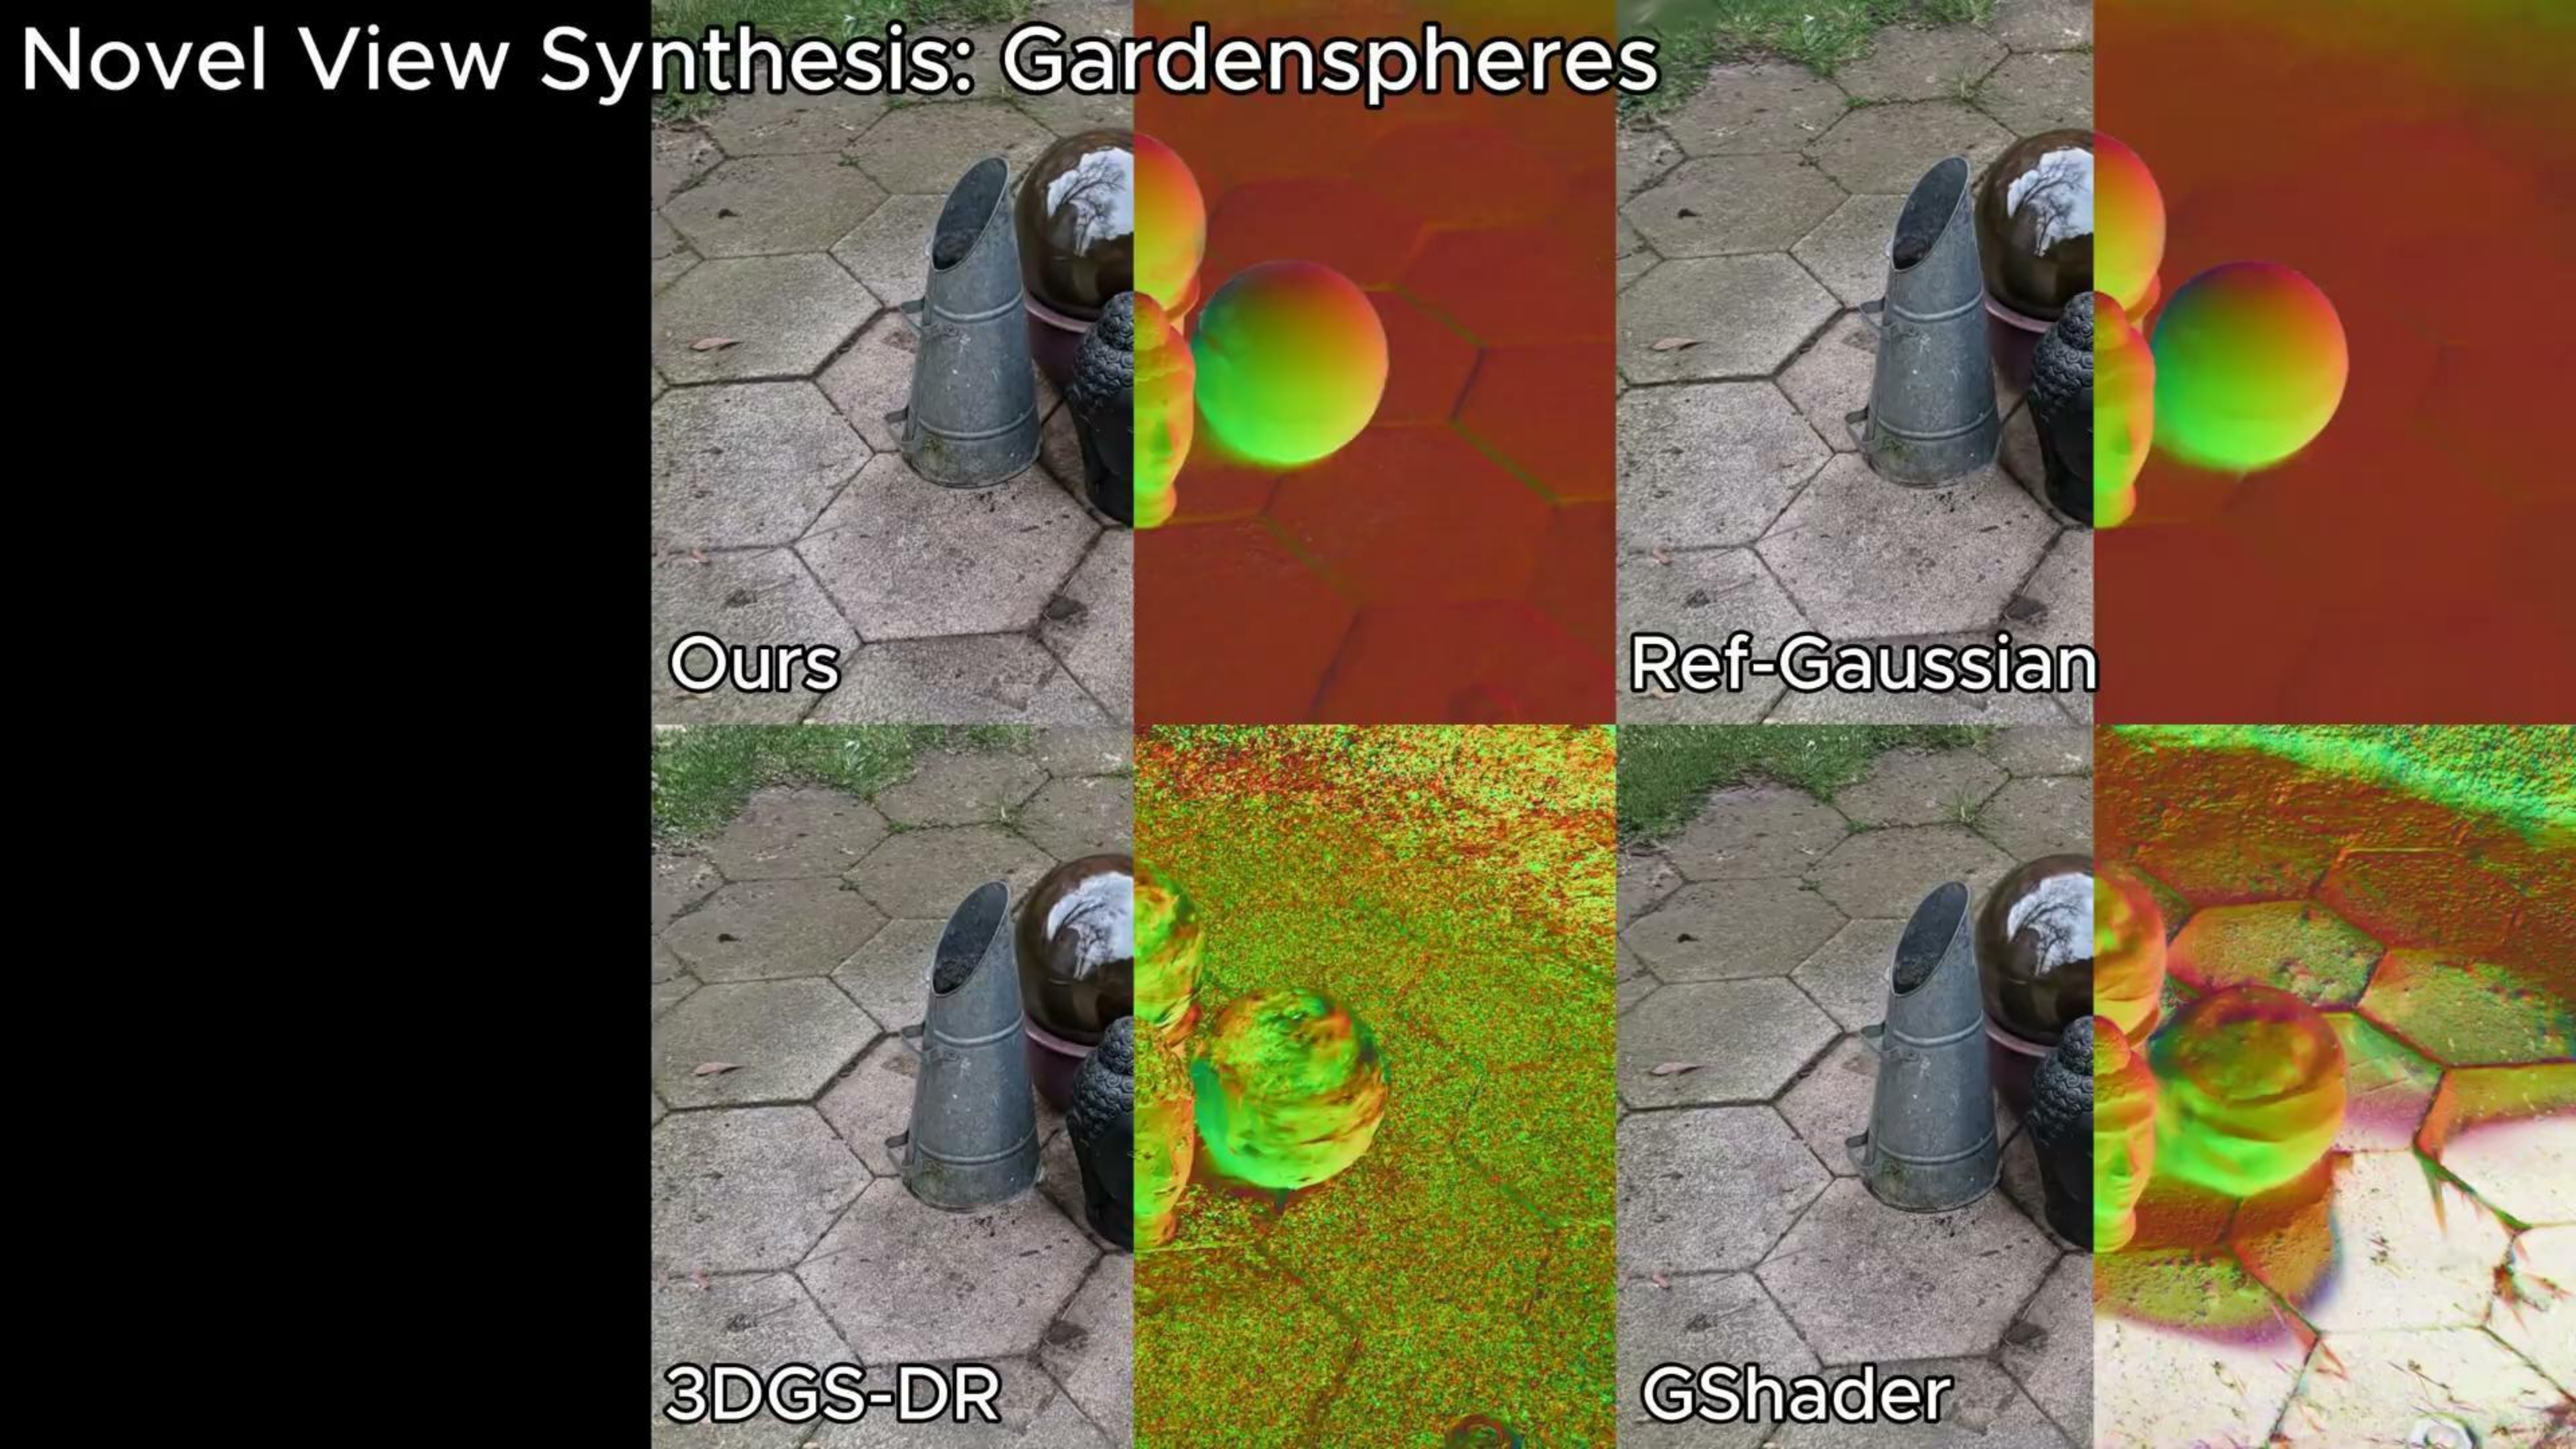
Novel View Synthesis: Gardenspheres
Ours
Ref-Gaussian
3DGS-DR
GShader

Novel View Synthesis: car

GT

Ours

Scene Editing Results

Original

Albedo Editing

KU LEUVEN

PSI

WACV 2026
Tucson, Arizona • Mar 6 - 10, 2026

Thank you for your attention!

Project Page