Intrinsic PAPR

Tackling Misattribution in 3D Intrinsic Decomposition via Proximity Attention Point Rendering

Alireza Moazeni[1]† · Shichong Peng[1] · Yanshu Zhang[1] · Chirag Vashist[1] · Ke Li[1,2,3]

[1] APEX Lab, School of Computing Science, Simon Fraser University, Canada · [2] Amii · [3] CIFAR

ECCV 2026 | MALMÖ | SEPT 8–12

Project page

ABSTRACT

Recent point-based intrinsic decomposition and inverse-rendering methods suffer from a **misattribution issue**, where individual primitives learn incorrect appearance features even though the aggregated rendering looks correct. The root cause is volume rendering, which blends translucent primitives along each ray and supervises only the final colour. Intrinsic PAPR instead uses Proximity Attention Point Rendering, which predicts appearance at ray-surface intersections and **supervises each point directly**.

CONTRIBUTIONS

- We identify the **root cause** of misattribution — the aggregation of semi-transparent primitives in volume rendering.
- The issue can be resolved by fundamentally **changing the rendering paradigm**, which is the insight behind Intrinsic PAPR.
- We supervise the decomposition robustly with a cIMLE-adapted 2D albedo prior and a **space carving loss**.

THE MISATTRIBUTION ISSUE

Volume rendering blends primitives, so **wrong per-point features** can still render correctly — PAPR **supervises each point directly**.
Volume rendering aggregates translucent primitives along a ray.
Only the final colour is supervised, not the primitives.
So a feature can be wrong though the render looks correct.

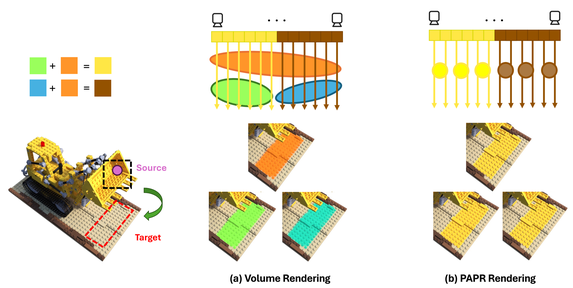

Misattribution issue. (a) Volume rendering lets primitives learn wrong features while the render stays correct. (b) PAPR supervises each point directly, so its features stay correct.

KEY RESULTS

10.3% Novel-view-synthesis PSNR improvement on synthetic scenes, over the best inverse-rendering baseline

8.0% Albedo-estimation PSNR improvement on synthetic scenes, over the best inverse-rendering baseline

>4× Reduction in albedo- and shading-transfer error, over a leading 3DGS-based method

METHOD

PAPR still aggregates — for each ray it mixes the top-K nearest points by proximity attention. But because the score grows as the perpendicular point–ray distance shrinks, the weights are sharply peaked, and each surface point dominates a cone of rays on which its weight is near 1, rather than volume rendering's translucent blend. Therefore the photometric residual backpropagates primarily to that point, which — with PAPR's parsimonious point set — makes the features of every point identifiable and mitigates misattribution.

- Per-point albedo and shading features are aggregated by proximity attention into two feature maps, trained end-to-end.
- A 2D albedo prior adapted with conditional Implicit Maximum Likelihood Estimation (**cIMLE**) handles monocular ambiguity.
- A space carving loss keeps the chosen albedo **consistent across views**.

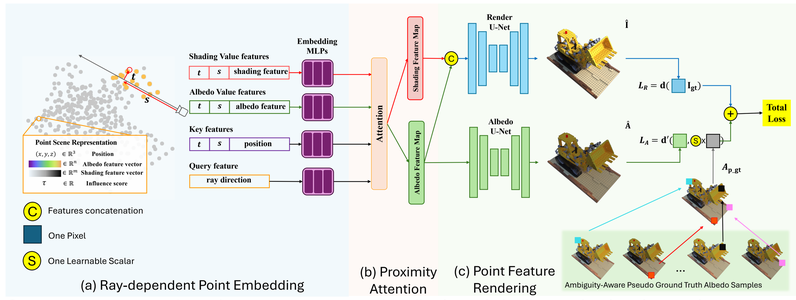

One pass, two maps. Attention aggregates per-point albedo and shading, trained end-to-end.

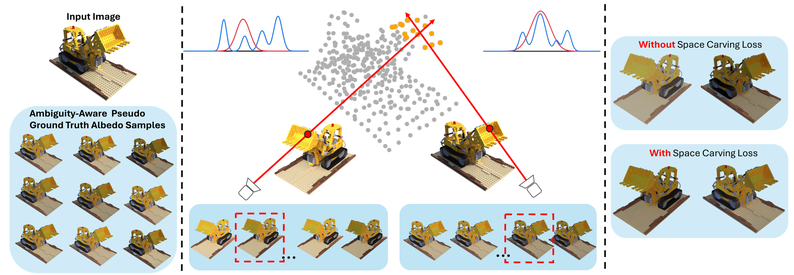

Multi-view consensus. Space carving selects the albedo mode that agrees across views.

QUANTITATIVE RESULTS

Novel View Synthesis (Image Quality Metrics)

Method	Synthetic			Real-world		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
DPIR	26.47	0.867	0.133	20.17	0.756	0.290
MaterialFusion	28.59	0.902	0.103	25.18	0.801	0.237
Intrinsic-NeRF	29.51	0.936	0.051	23.79	0.768	0.252
GS-IR	30.40	0.945	0.055	26.44	0.851	0.185
IRGS	31.35	0.949	0.048	27.26	0.878	0.157
DiscretizedSDF	31.72	0.959	0.038	27.86	0.876	0.147
PAPR	32.84	0.973	0.035	28.03	0.865	0.159
▸ **Intrinsic PAPR (Ours)**	**35.00**	**0.978**	**0.021**	**29.99**	**0.912**	**0.094**

Albedo Reconstruction (Synthetic)

Method	PSNR ↑	SSIM ↑	LPIPS ↓
DPIR	22.89	0.854	0.215
Intrinsic-NeRF	24.86	0.880	0.108
MaterialFusion	25.36	0.899	0.097
GS-IR	26.33	0.877	0.127
IRGS	28.44	0.913	0.095
DiscretizedSDF	28.89	0.921	0.097
▸ **Intrinsic PAPR (Ours)**	**31.21**	**0.949**	**0.056**

ALBEDO AND SHADING TRANSFER

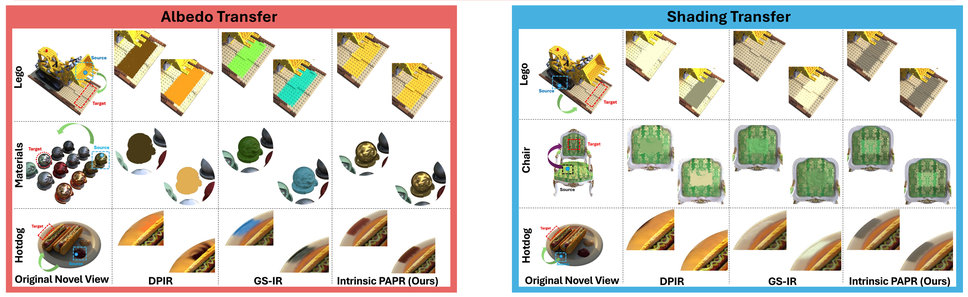

QUALITATIVE ALBEDO COMPARISON

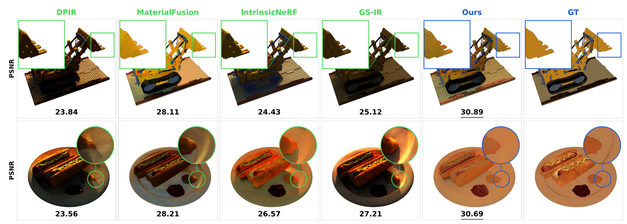

APPLICATIONS

Freeform Albedo Editing

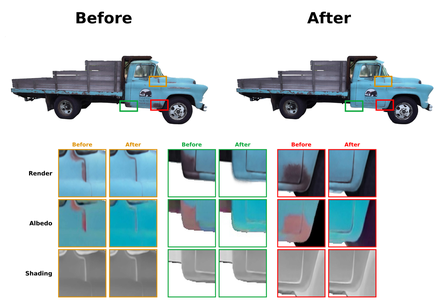

Geometry-aware shading manipulation — shadow adding

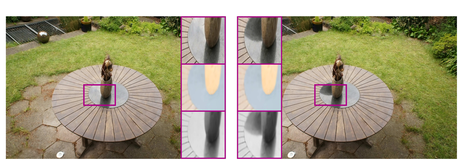

Geometry-aware shading manipulation — shadow removal

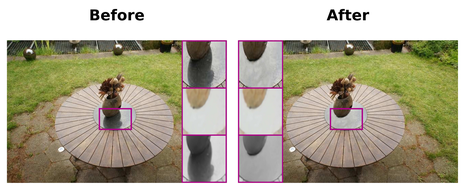

Scene-level/part-level shading manipulation

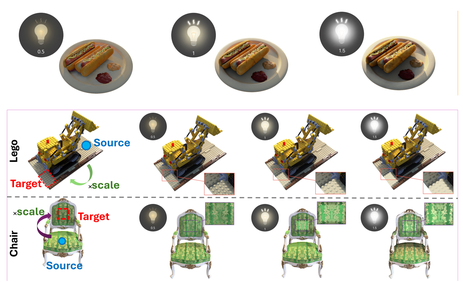

This research was enabled in part by NSERC, the Canada Foundation for Innovation, the Canada CIFAR AI Chairs program, the BC DRI Group and the Digital Research Alliance of Canada.

Our key design choice of reducing the number of free parameters for more accurate recovery also means that some effects, such as full relighting and subsurface scattering, cannot be modelled; and because we assume surface opacity, highly semi-transparent phenomena (e.g., smoke or clear glass) are likewise not handled.

[1]APEX Lab, SFU · [2]Amii · [3]CIFAR · †Corresponding author