Surf-NeRF: Surface Regularised Neural Radiance Fields

Jack Naylor, Viorela Ila, Donald G. Dansereau
The University of Sydney
{firstname.lastname}@sydney.edu.au

CVPR JUNE 3-7, 2026 DENVER COLORADO

Project Page · Jack Naylor · Lab Website

Motivation & Contributions

- Radiance-based representations produce realistic novel views of a scene, but have poor geometric representation around non-Lambertian scene elements.

- Reflection parameterisations [1] improve geometry, but still suffer from shape-radiance ambiguity.

- Other models of radiance (e.g. surface light fields) have priors on model of appearance AND geometry [2]

- We propose Surf-NeRF consisting of: 1) representation driven regularisation towards a surface light field model, 2) improved positional encoding for non-planar geometry which we benchmark on a 3) new captured dataset.

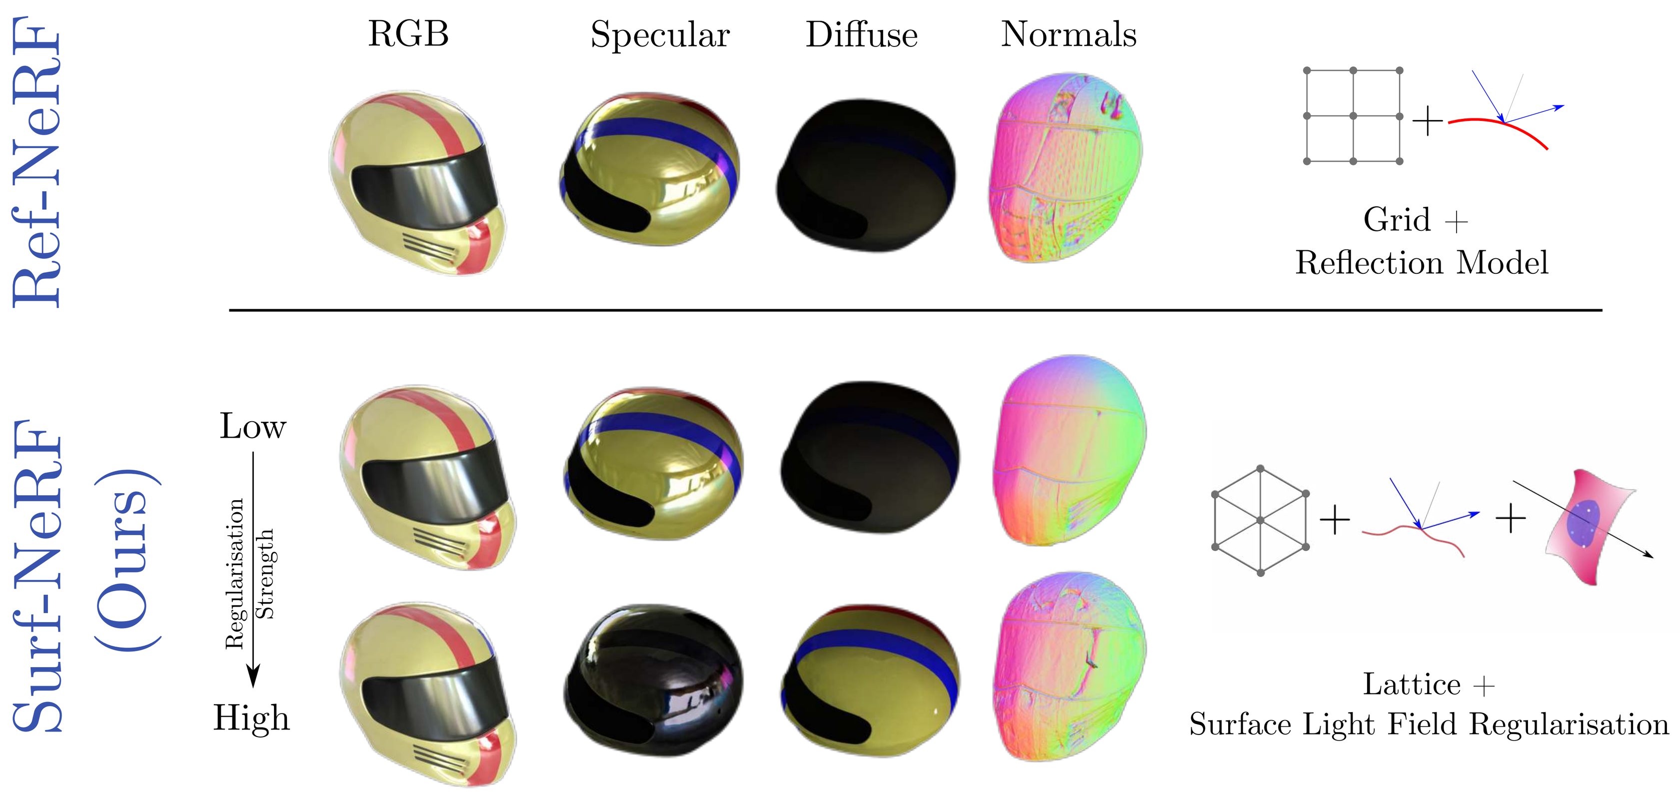

Methodology

- Introduce permutohedral lattice hash encoding suitable for smooth surfaces

- Tetrahedrons can better approximate non-planar surfaces compared to voxels

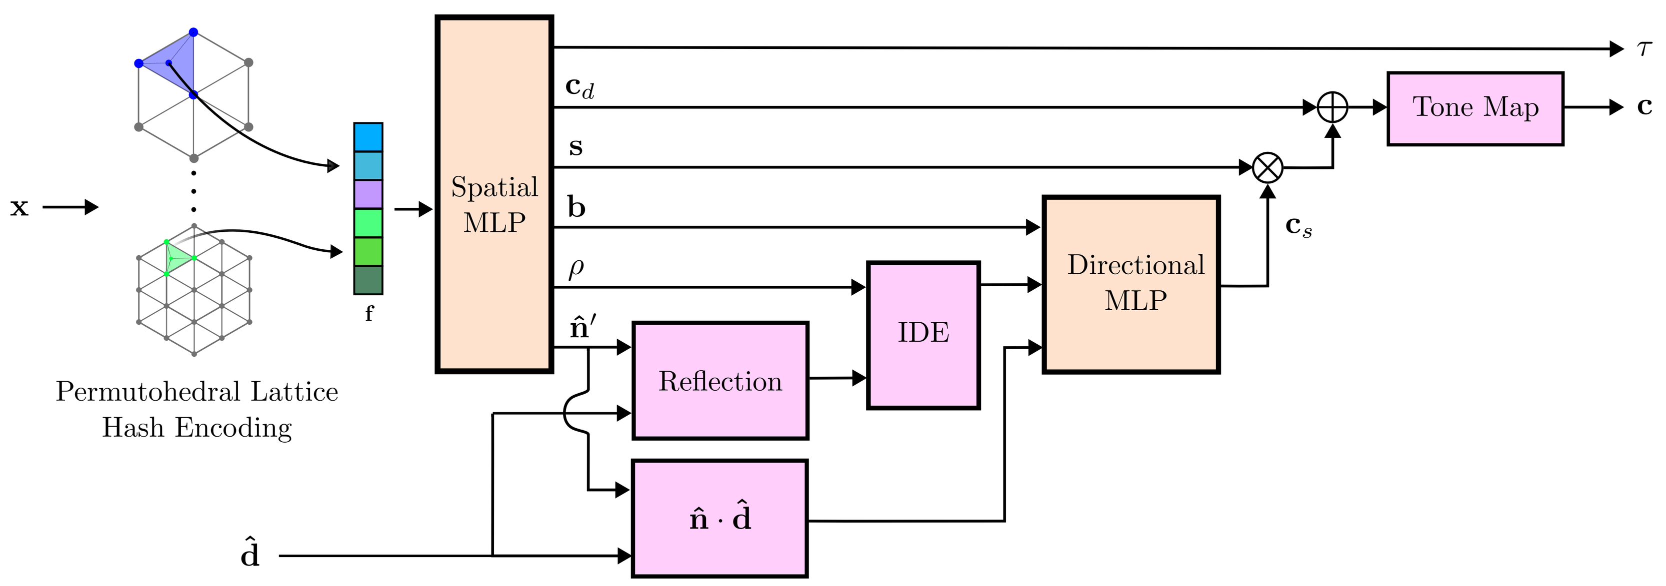

- Curriculum learning of surface: apply regularisation with second pass through network with increasing frequency during training

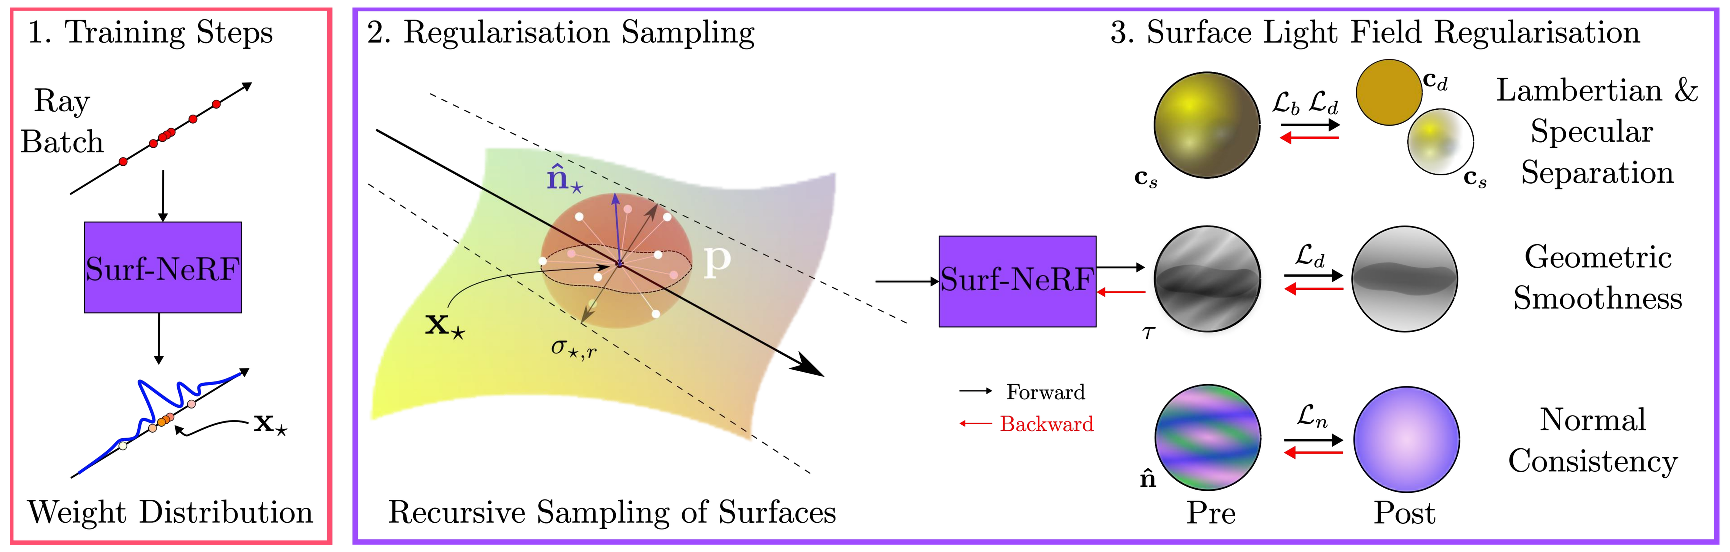

- Sample surface spatially and directionally

- Regularise towards a surface light-field model of appearance and geometry

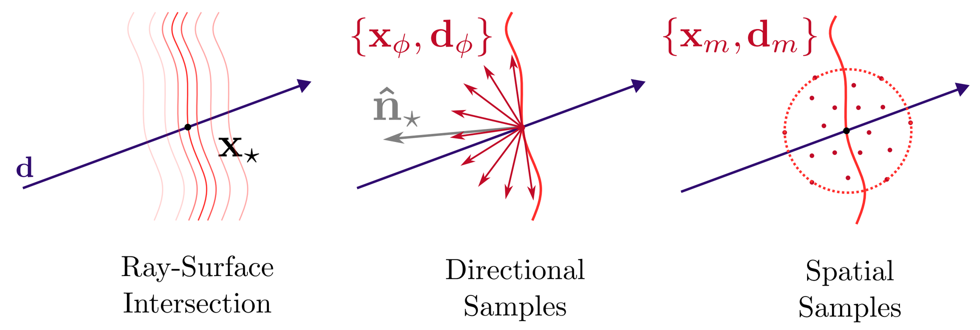

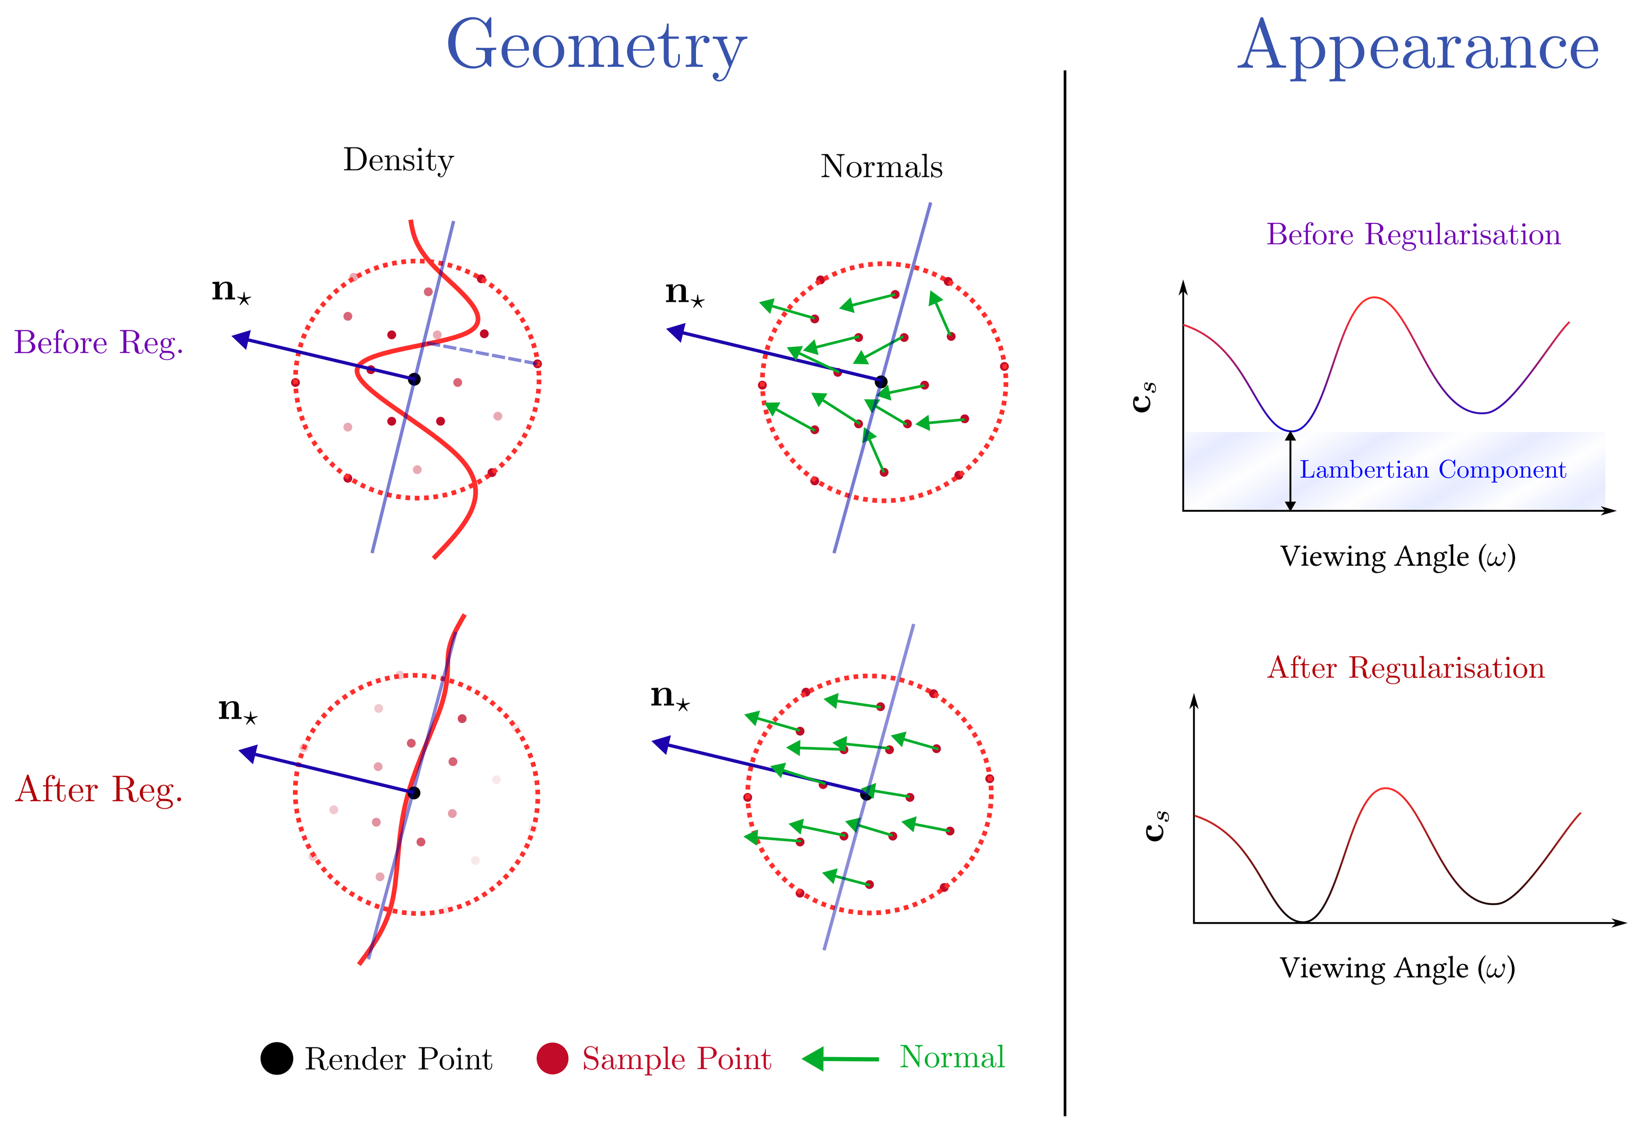

Results

- Significantly improved geometric representation, normal consistency

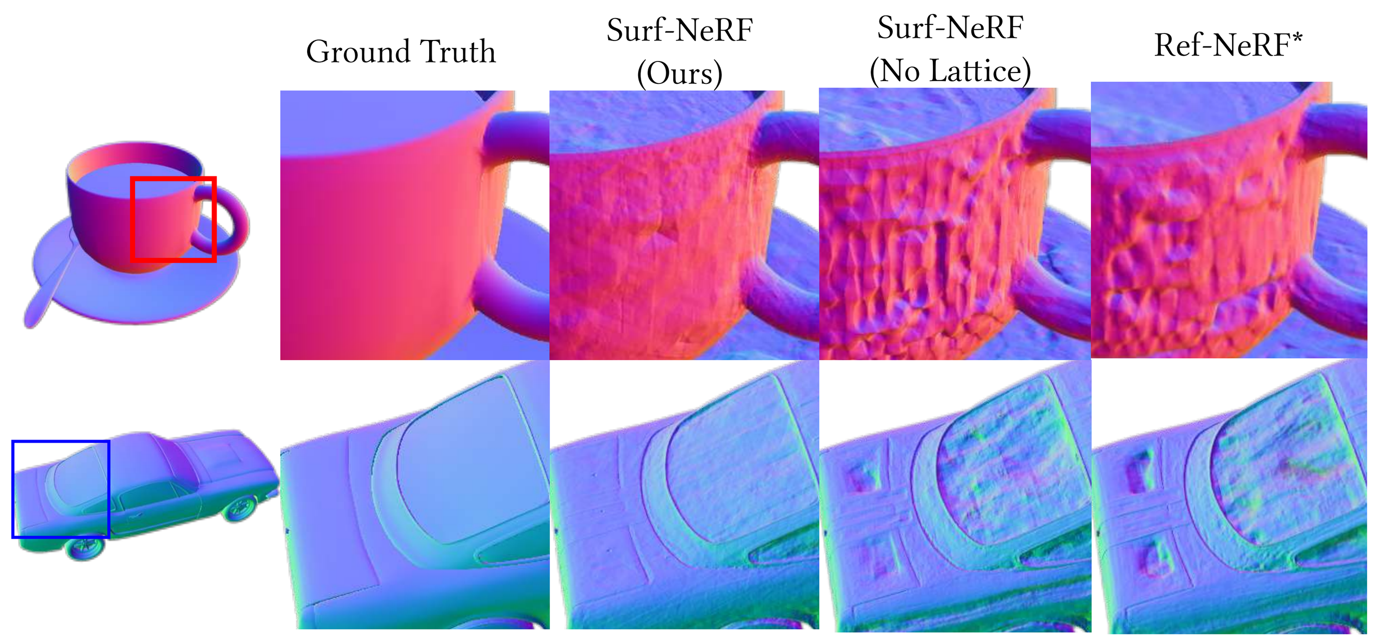

- More consistent rendering of specularities

- Separation of diffuse + specular appearance

- 28% improvement to surface normals over hash-based approaches

	Model	PSNR ↑	SSIM ↑	MAE ↓	RMSE ↓
PE	MipNeRF	31.29	0.943	56.17	0.125
	MipNeRF+Diff	30.84	0.936	60.00	0.122
	Ref-NeRF	33.21	0.971	25.89	0.117
	Surf-NeRF (Reg. Only)	33.01	0.967	22.15	0.117
Grid	ZipNeRF	31.02	0.952	19.47	0.209
	Zip+Ref-NeRF	33.14	0.964	14.70	0.195
	Lattice+Ref-NeRF	33.21	0.964	13.76	0.202
	Surf-NeRF (No Lattice)	32.13	0.954	13.35	0.202
	Surf-NeRF (Ours)	32.82	0.955	10.60	0.198

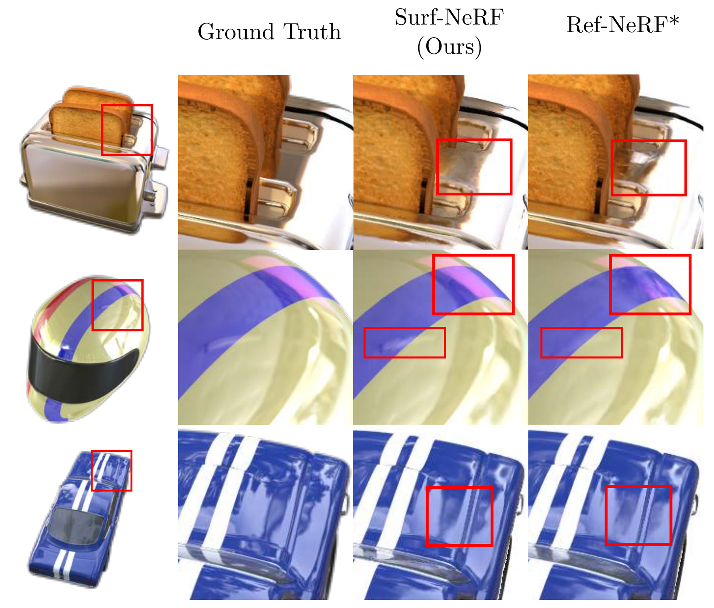

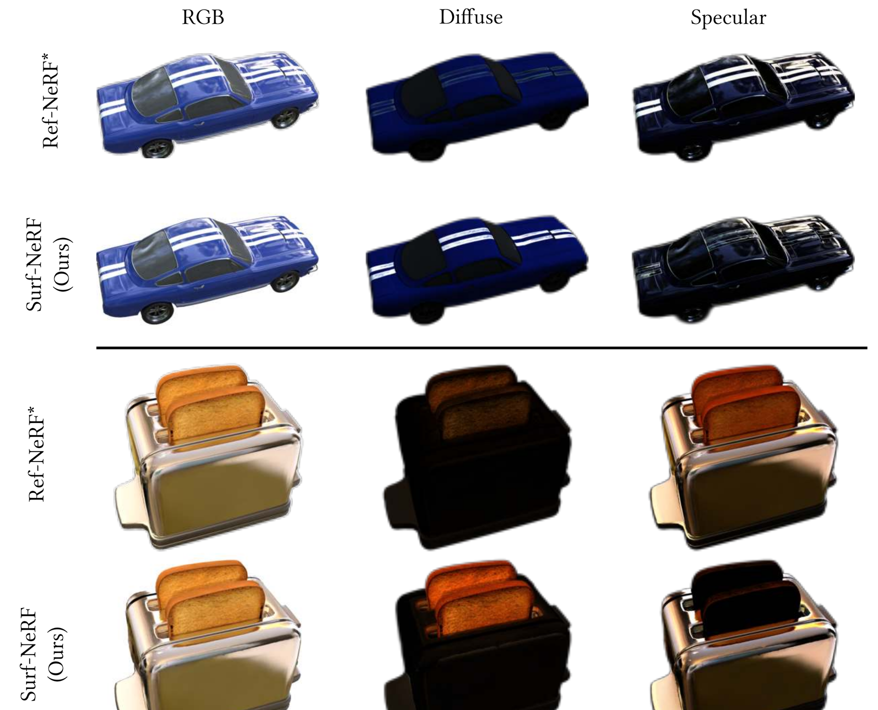

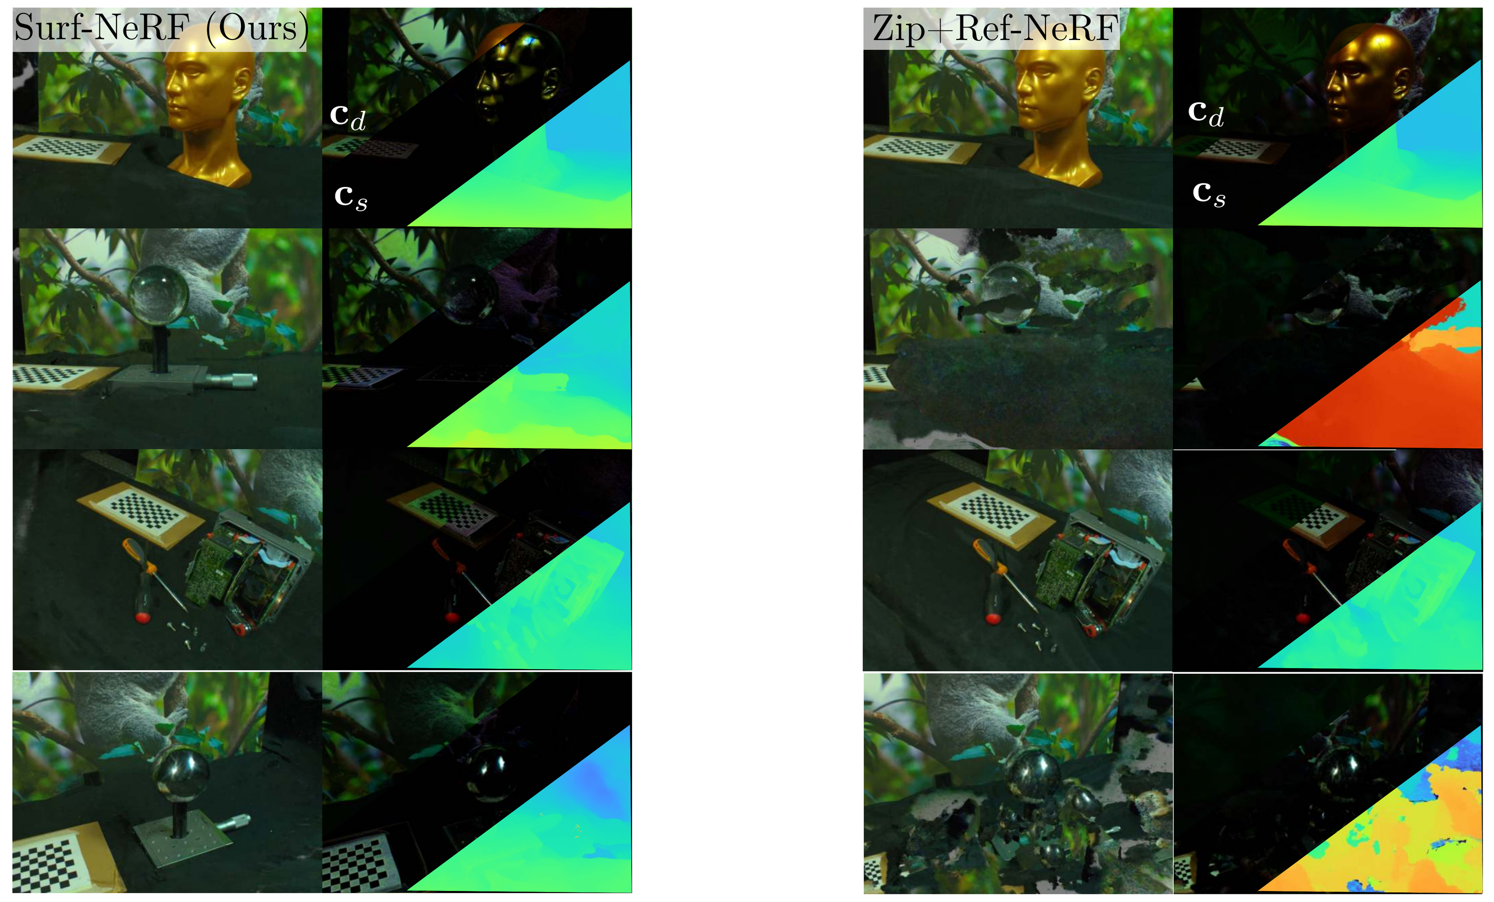

- Improved separation of diffuse + specular and surface geometry on sparse captured multi-view data (Koala Dataset)

Conclusions

- Regularisation towards a surface light field model substantially improves geometric representation while maintaining volumetric rendering

- Limitations: 2x steps through network, possible to ablate fine details which are better represented volumetrically, slightly reduced rendering quality

- Future work: we anticipate similar regularisation can be applied to 3DGS/ Radiant Foam + other scene representations.

References & Acknowledgements

[1] Verbin et al., *Ref-NeRF: Structured View-Dependent Appearance For Neural Radiance Fields*, CVPR, 2022
[2] Wood et al., *Surface Light Fields for 3D Photography*, SIGGRAPH, 2000

The authors wish to acknowledge Ford Motor Company and ARIAM Hub. Computational aspects of this work were supported by Australia's NCI Gadi, an NVIDIA Academic Grant and Google TPU Research Cloud.