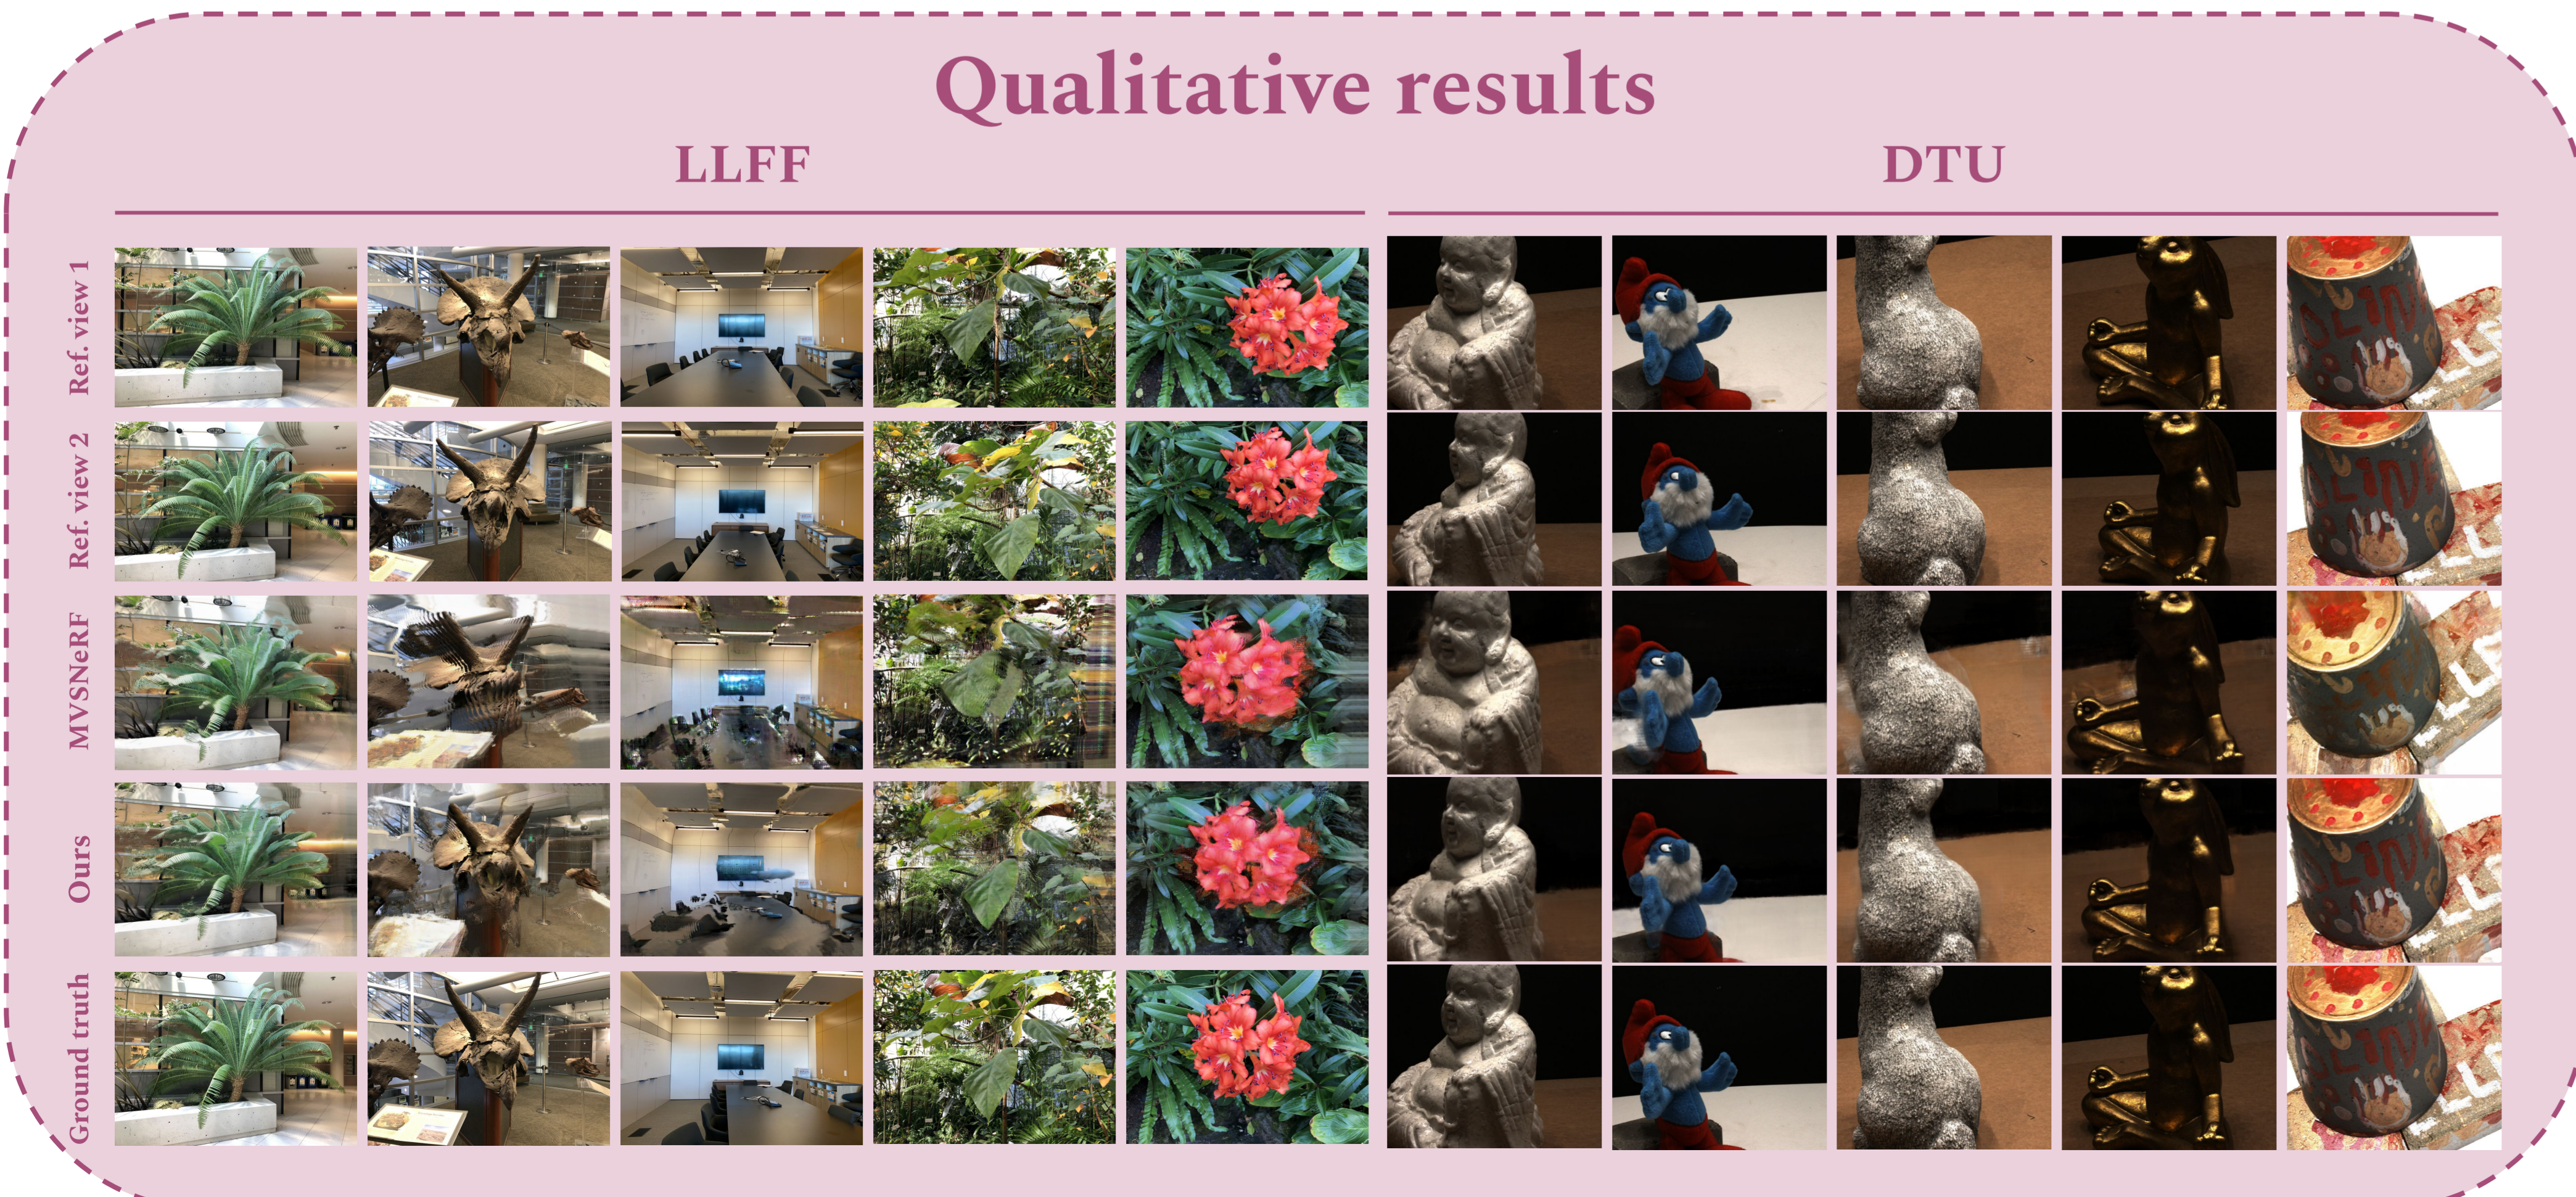
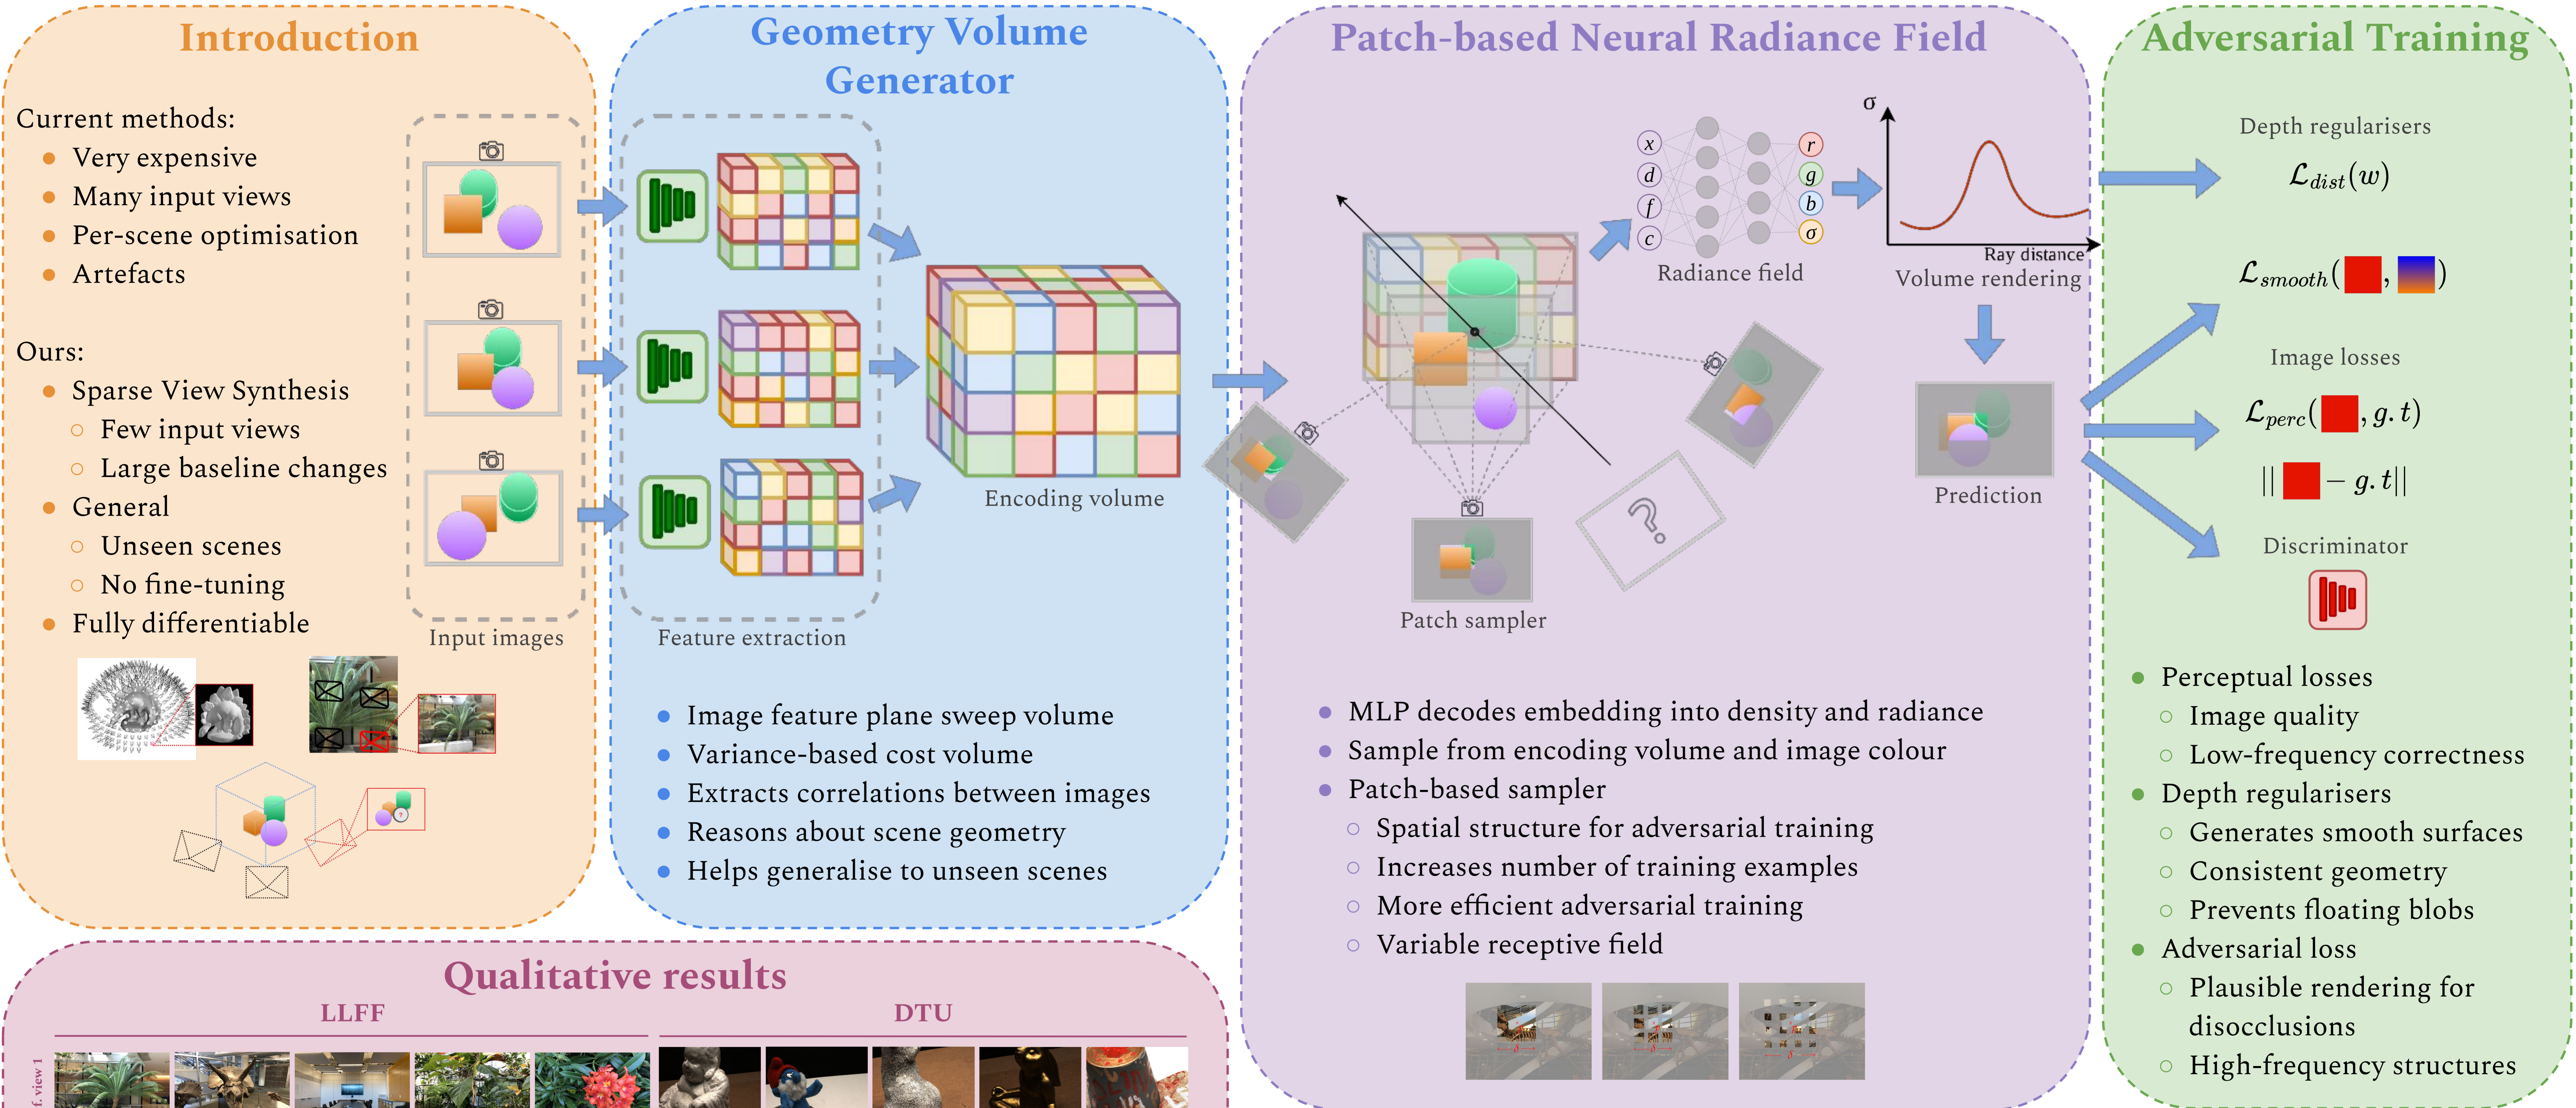


SVS: Adversarial refinement for sparse novel view synthesis

Violeta Menéndez González^{1,2}, Andrew Gilbert¹, Graeme Phillipson², Stephen Jolly², Simon Hadfield¹
 (v.menendezgonzalez@surrey.ac.uk); 1) Centre for Vision, Speech and Signal Processing, University of Surrey; 2) BBC R&D

Download the code!
github.com/violetamenendez/svs-sparse-novel-view



Quantitative results

Table 1: Quantitative evaluation. We evaluate our model over the DTU and Forward Facing datasets. **Bold** is best result, *italic* is second best. * results taken from paper.

| Model | Experiment | DTU | | | Forward facing | | |
|--------------|------------|--------------|---------------|---------------|----------------|---------------|---------------|
| | | PSNR↑ | SSIM↑ | LPIPS↓ | PSNR↑ | SSIM↑ | LPIPS↓ |
| RegNeRF* [4] | Optimised | 18.89 | 0.745 | 0.190 | 19.08 | 0.587 | 0.336 |
| IBRNet [3] | | 12.71 | 0.4772 | 0.5678 | 16.40 | 0.5230 | 0.4986 |
| MVSNeRF [1] | Unseen | 18.92 | <i>0.6831</i> | <i>0.2580</i> | 16.98 | 0.5839 | <i>0.3853</i> |
| Ours | | 19.03 | 0.6929 | 0.2066 | <i>16.55</i> | <i>0.5534</i> | 0.3441 |

Table 2: Ablation study. We study the effect of each addition to the model on the Forward facing dataset. **Bold** is best result, *italic* is second best.

| Depth Smooth | Distortion | Adversarial | PSNR↑ | SSIM↑ | LPIPS↓ |
|--------------|------------|-------------|--------------|---------------|---------------|
| × | × | × | 16.11 | 0.5318 | 0.4791 |
| ✓ | × | × | <i>16.14</i> | 0.5361 | 0.4709 |
| ✓ | ✓ | × | 16.10 | <i>0.5419</i> | <i>0.4611</i> |
| ✓ | ✓ | ✓ | 16.55 | 0.5534 | 0.3441 |

Conclusion

- Generalisable
- Sparse input views
- Larger baseline changes
- Perceptual quality improvements up to 60%
- Training instability
- Some lack of fidelity

References

- [1] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. MVSNeRF: Fast Generalizable Radiance Field Reconstruction from Multi-View Stereo. CVPR, 2021.
- [2] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tanaka, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. ECCV, 2020.
- [3] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P. Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. IBRNet: Learning Multi-View Image-Based Rendering. CVPR, 2021.
- [4] Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. RegNeRF: Regularizing Neural Radiance Fields for View Synthesis from Sparse Inputs. CVPR, 2022.
- [5] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. GRAF: Generative Radiance Fields for 3D-Aware Image Synthesis. NeurIPS, 2020.
- [6] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. CVPR, 2022.
- [7] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised Monocular Depth Estimation with Left-Right Consistency. CVPR, 2017.