Rendering for Hyper-Realistic Displays: A Benchmark Study

Akshay Jindal and Rafał K. Mantiuk

UNIVERSITY OF CAMBRIDGE

CVMP 2022

Motivation

- Computational displays are getting closer to reproducing reality by delivering almost a complete set of perceptual realism cues [5].
- **The rendering methods for these displays are required to process high-resolution HDR images in real-time to generate novel views from arbitrary viewpoints while delivering correct depth cues.**
- Image-based rendering (IBR) methods, a natural choice for such displays, come with their own set of trade-offs. It is unclear how their artefacts translate to perceived quality.

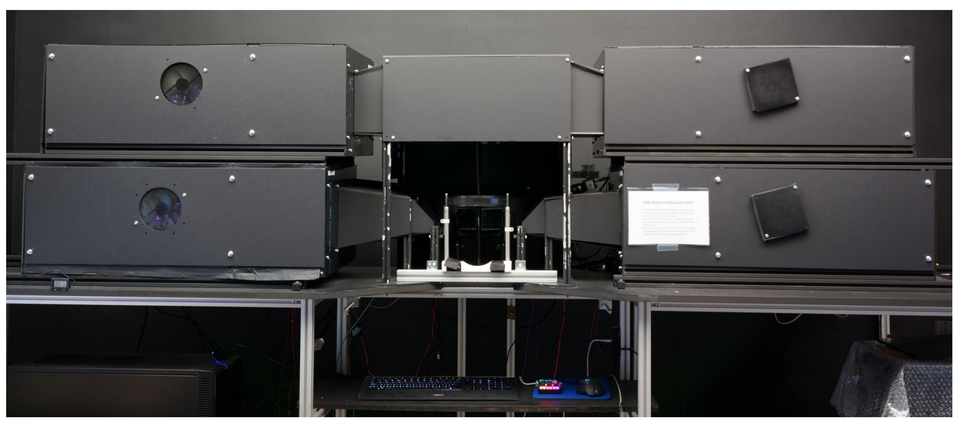

High-Dynamic-Range Multi-Focal Stereo Display

Aim

- In this work, we aim to understand the performance and quality trade-offs of real-time IBR techniques by comparing 3 IBR methods:
  - **Dynamically reparameterised lightfield (DRLF) [2]**
  - **Lumigraph implemented as a mesh with view-dependent textures [5]**
  - **NeX, a neural multi-plane images method [4]**.

Methods

- We compared the three methods on 4 forward-facing synthetic scenes (2 geometries × 2 materials).

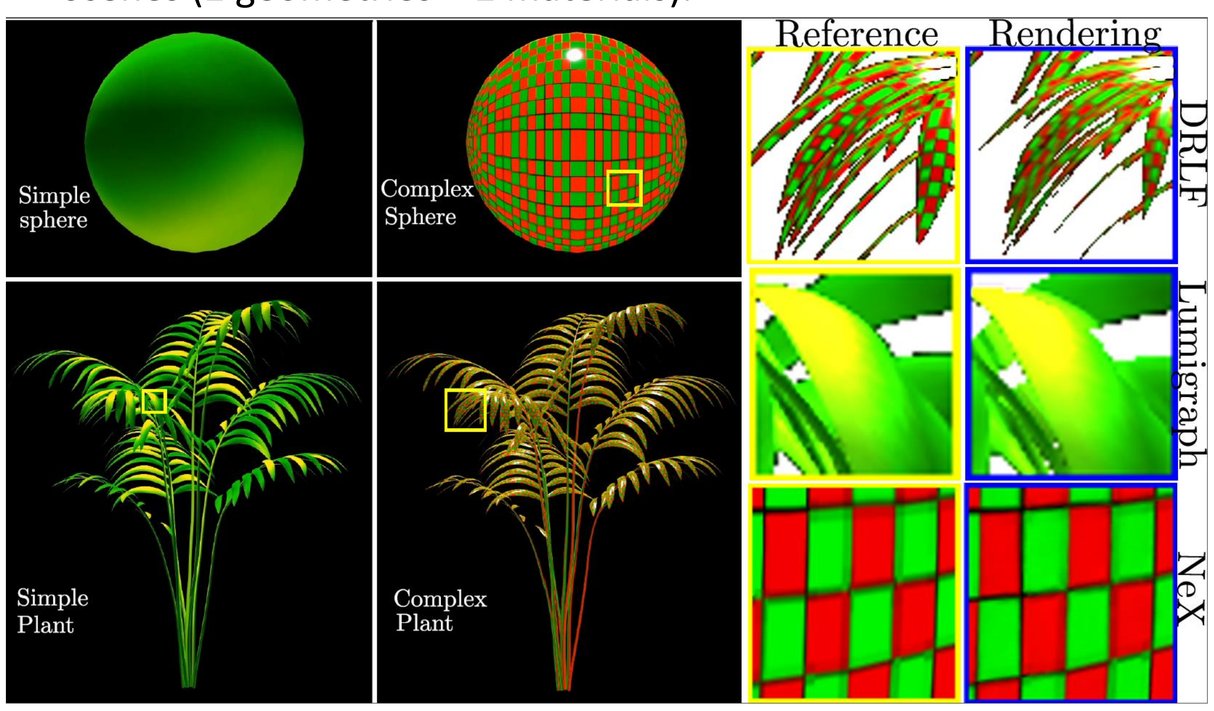

3D scenes used for comparing different IBR methods and sample artefacts caused by the three rendering methods

- The geometries used were a sphere mesh and a plant mesh.
- The meshes were mapped to two materials: a Lambertian material with a low-frequency gradient image as its diffuse component and a specular material (Phong shading) with a high-frequency checkerboard as its diffuse component.
- We rendered 20 images of 2160×1440 resolution and 16-bit per channel with a baseline of 100 mm for all 4 scenes.

Results

- The DRLF method can achieve rendering quality similar to NeX at a much lower performance cost.
- Mesh size is not a bottleneck on performance for lumigraph as there is no noticeable effect of the number of triangles between sphere (960 triangles) and plant scene (84,748 triangles).
- **The three methods produce very different types of artefacts**: DRLF leads to blurring of regions away from focal plane, Lumigraph leads to ghosting near thin edges, and NeX leads to changes in texture appearance. Artefacts become more prominent in videos.
- While lumigraph got the highest quality scores on PU-PSNR, PU-SSIM, HDRVDP3, and MDTVSFA, **there is no consensus on the rating and ranking of the methods across 4 metrics.**

What metric to use to evaluate and train IBR methods?

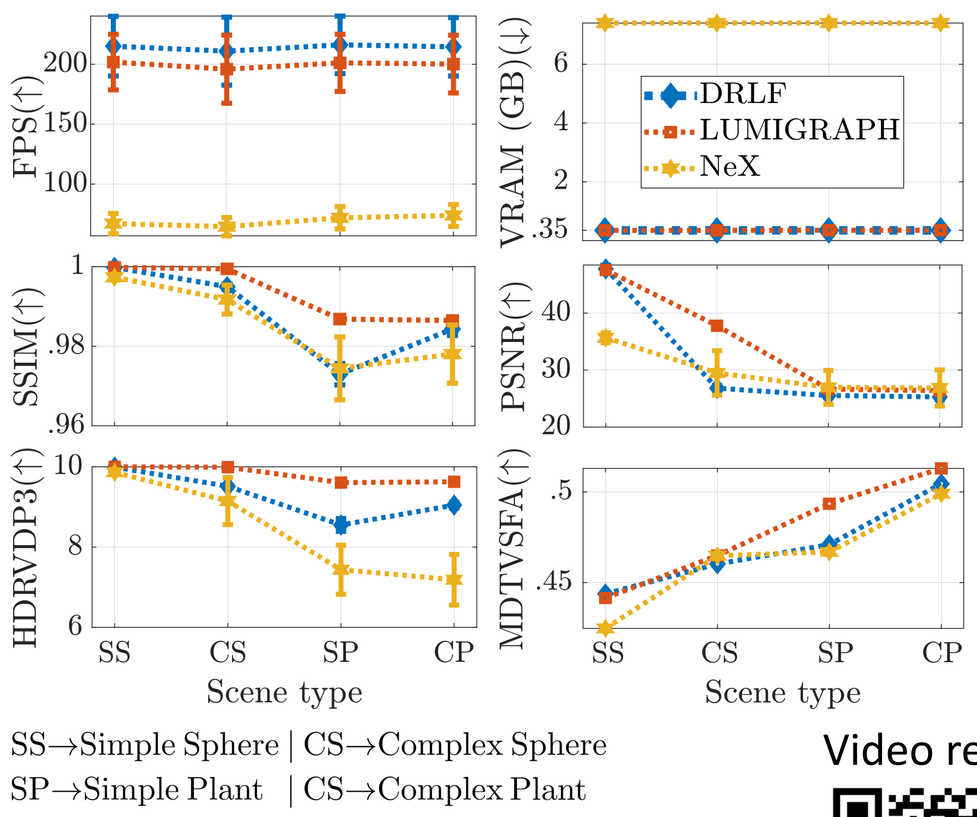

Video results

References

[1] M. Azimi and R. K. Mantiuk. Pu21: A novel perceptually uniform encoding for adapting existing quality metrics for hdr. In 2021 Picture Coding Symposium (PCS). IEEE, 2021.
[2] Isaksen, L. McMillan, and S. J. Gortler. Dynamically reparameterized light fields. In Proceedings of the 27th annual conference on Computer graphics and interactive techniques, 2000.
[3] D. Li, T. Jiang, and M. Jiang. Unified quality assessment of in-the-wild videos with mixed datasets training. International Journal of Computer Vision, 2021.
[4] S. Wizadwongsa, P. Phongthawee, J. Yenphraphai, and S. Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion. In Proc. of Computer Vision and Pattern Recognition, 2021.
[5] F. Zhong, A. Jindal, A.O. Yöntem, P. Hanji, S. J. Watt, and R. K. Mantiuk. Reproducing reality with a high-dynamic-range multi-focal stereo display. ACM Trans. Graph., 2021.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie-Skoldowska-Curie Grant Agreement No. 765911

realVision

www.realvision-itn.eu