Progressive Multi-Scale Light Field Networks

David Li and Amitabh Varshney
augmentariumlab.github.io/multiscale-lfn/

Background

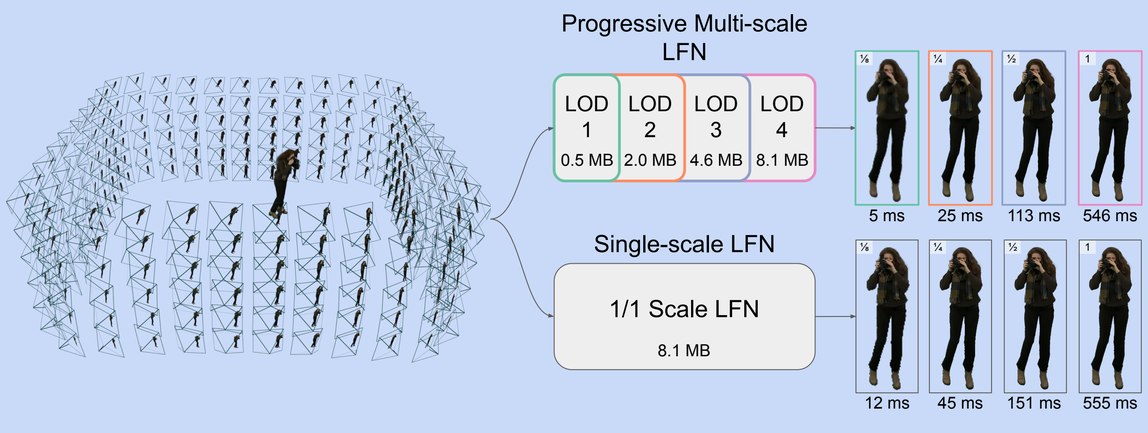

- Neural fields have become popular due to their ability to efficiently encode light fields.
- Light Field Networks (LFNs), proposed by Sitzmann et al. in 2021, can render in real-time by directly predicting the color for each input ray.
- Rendering with a high-resolution light field network at smaller scales can lead to aliasing due to the subsampling, similar to nearest downsampling.

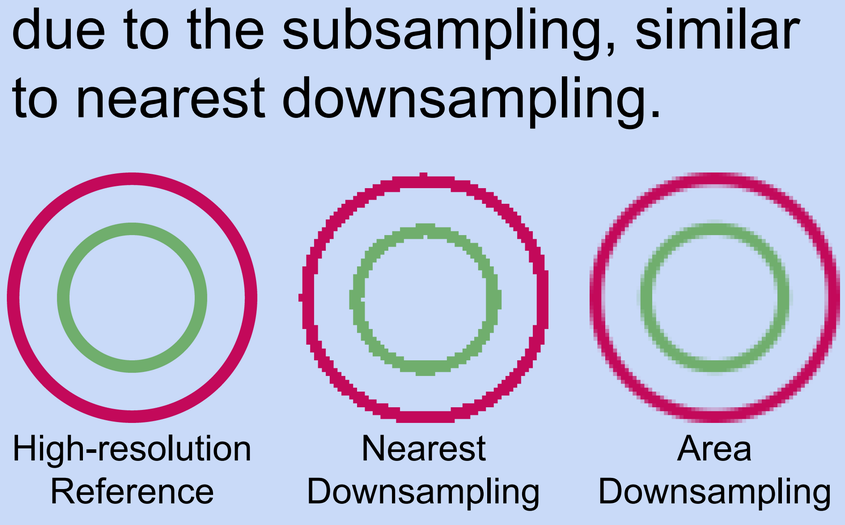

- LFNs do not have a progressive representation, so the entire model must be downloaded before any rendering can begin.

Method

- We propose encoding light fields at several scales, similar to a mipmap.
- Specifically, we use 1/1, ½, ¼, and ⅛ scale light fields.
- Then we can render using an appropriate resolution light field to reduce aliasing and flickering.

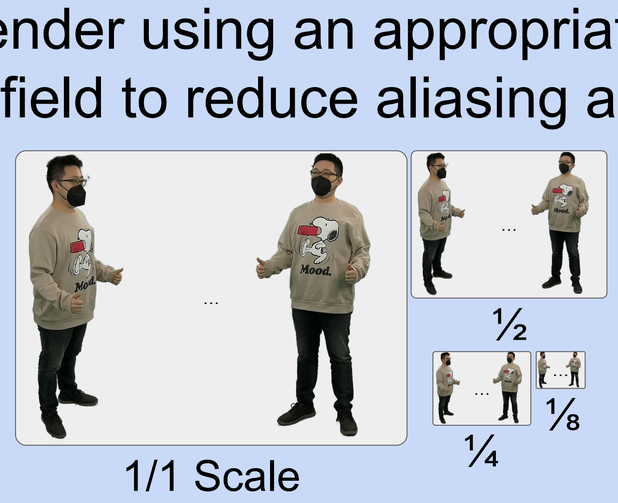

- We encode each scale of light field using a subset of neurons at each layer to get a progressive model with multiple levels of detail (LODs).

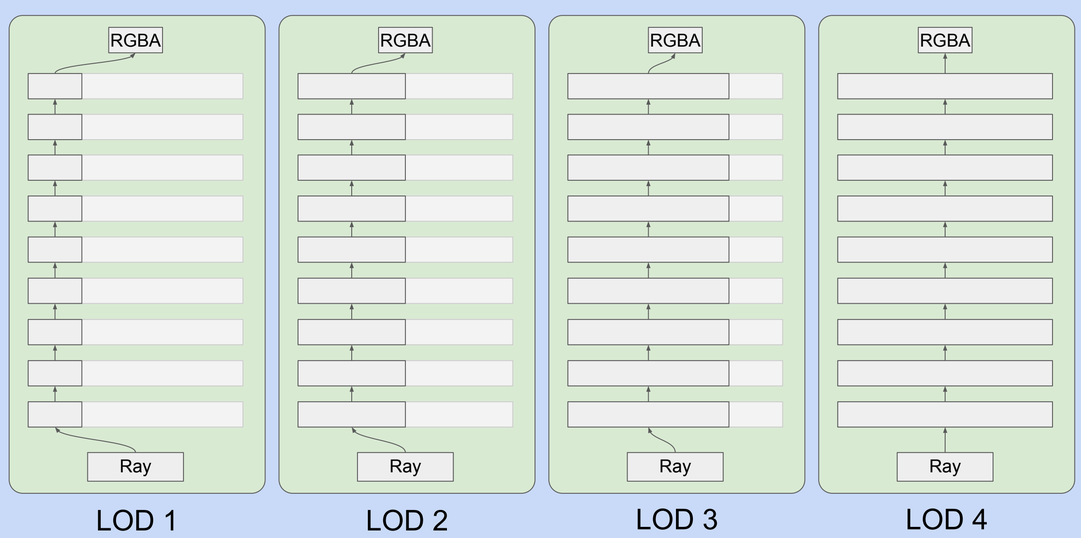

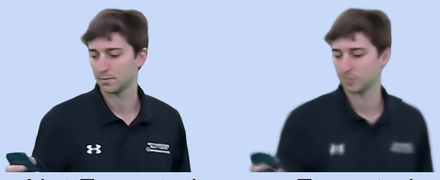

Experiments

- Our method achieves better results and reduces the flickering and aliasing artifacts at smaller sizes compared to a single-scale LFN.
- Our method has a smaller total model size and takes less time to train compared to using multiple LFNs.
- Using lower LODs when rendering at lower resolutions can help improve performance.

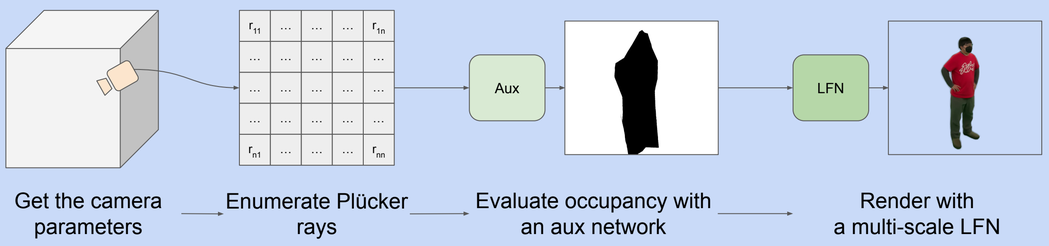

Average Rendering Quality Over All Datasets

Model	LOD 1	LOD 2	LOD 3	LOD 4
Single-scale LFN	26.95	28.05	28.21	27.75
Multiple LFNs	29.13	29.88	29.27	27.75
Multi-scale LFN	29.37	29.88	29.01	28.12

(a) PSNR at ⅛, ¼, ½, and 1⁄1 scale

Model	LOD 1	LOD 2	LOD 3	LOD 4
Single-scale LFN	0.8584	0.8662	0.8527	0.8480
Multiple LFNs	0.8133	0.8572	0.8532	0.8480
Multi-scale LFN	0.8834	0.8819	0.8626	0.8570

(b) SSIM at ⅛, ¼, ½, and 1⁄1 scale

Multiple LFNs *vs* Progressive Multi-scale LFN

Model	LOD 1	LOD 2	LOD 3	LOD 4	Total
Multiple LFNs	0.518	2.036	4.554	8.072	15.180
Multi-scale LFN	——	8.072	——		8.072

(a) Model Size (MB)

Model	LOD 1	LOD 2	LOD 3	LOD 4	Total
Multiple LFNs	3.51	6.36	4.09	11.35	25.31
Multi-scale LFN	——	17.78	——		17.78

(b) Average Training Time Over All Datasets (hours)

Average Rendering Performance

LOD 1	LOD 2	LOD 3	LOD 4
3.7	3.9	4.7	5.9

(a) Rendering at ⅛ scale (ms per frame)

LOD 1	LOD 2	LOD 3	LOD 4
4	11	58	305

(b) Rendering at ⅛, ¼, ½, 1⁄1 scale (ms per frame)

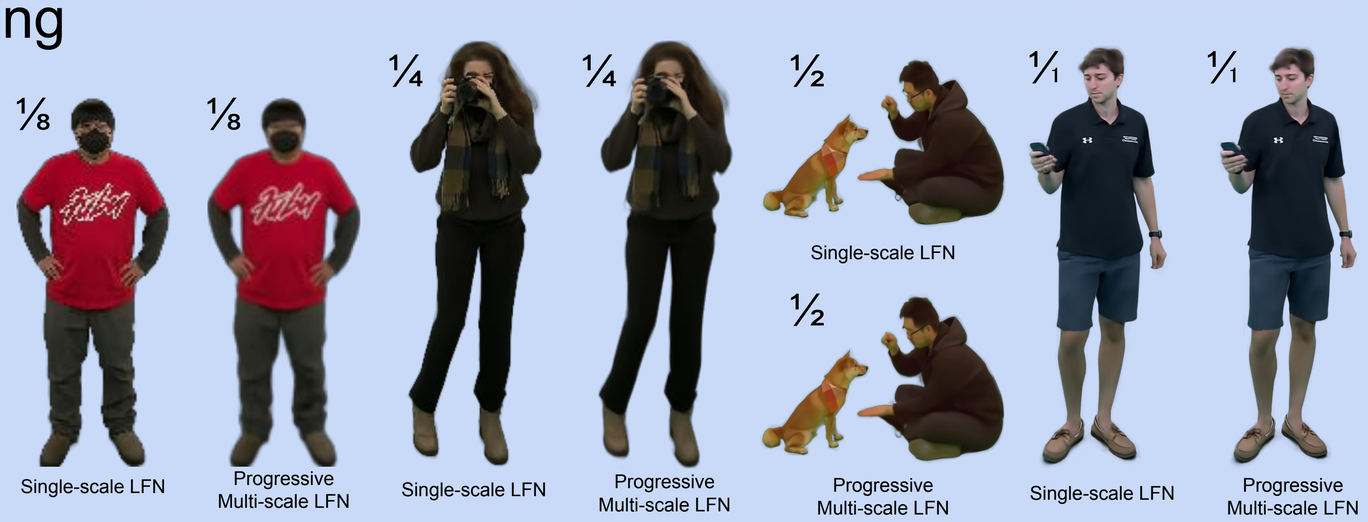

3DV 2022

UNIVERSITY OF MARYLAND

UMIACS
University of Maryland Institute for Advanced Computer Studies