Continuous Levels of Detail for Light Field Networks

David Li, Brandon Y. Feng, and Amitabh Varshney

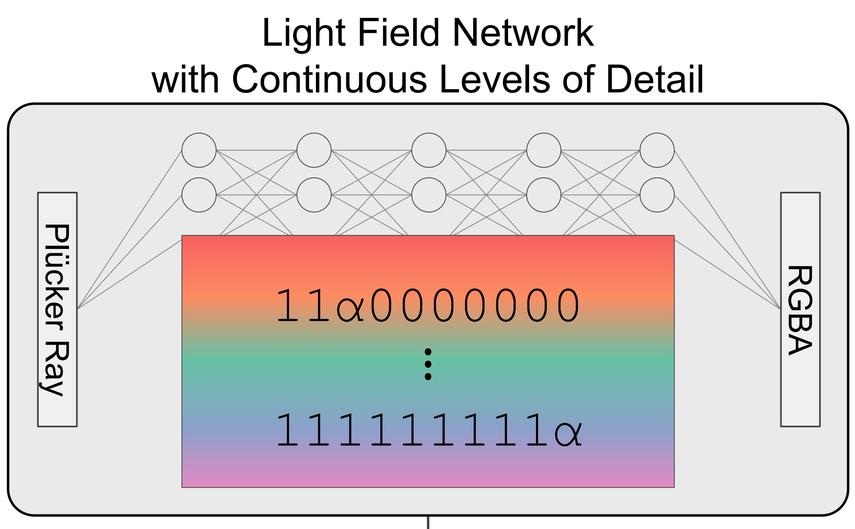

Introduction

- Neural fields enable a compact photo-realistic representation of 3D scenes by encoding them into neural networks.
- Light Field Networks (LFNs), proposed by Sitzmann et al. in 2021, can represent 3D objects and render them in real-time without volume rendering by directly predicting the color for each ray or pixel.

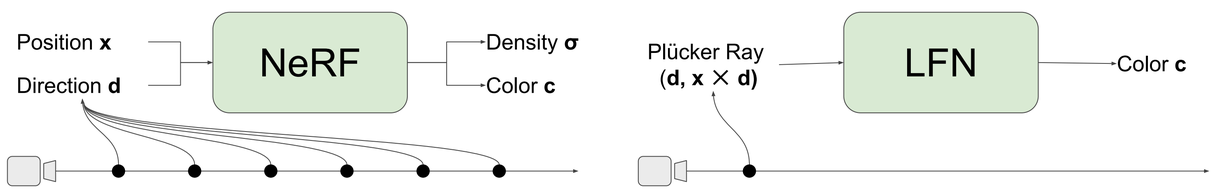

- For computer graphics applications, levels of detail (LODs) provide anti-aliasing and more efficient rendering of objects at different scales.
- In prior work, discrete LODs enable LFNs to be progressively streamed and to render at four scales: 1, ½, ¼, and ⅛.

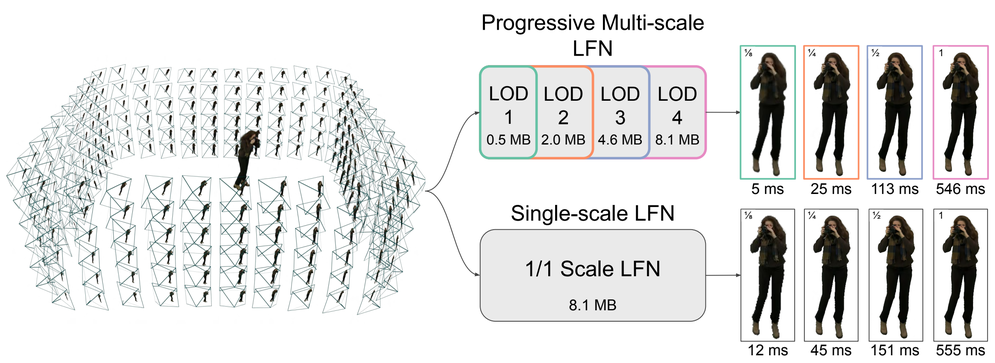

Progressive Multi-Scale Light Field Networks (3DV 2022)

Method

We propose to combine the following techniques to achieve continuous levels of detail:

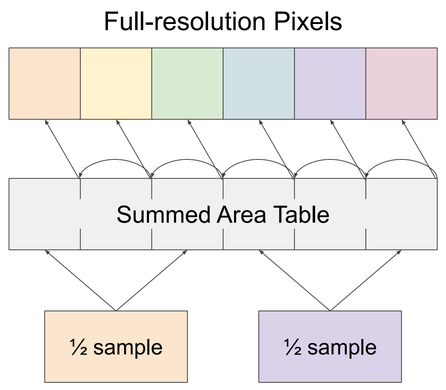

Summed-area tables allow for arbitrary scale and position sampling at training time.

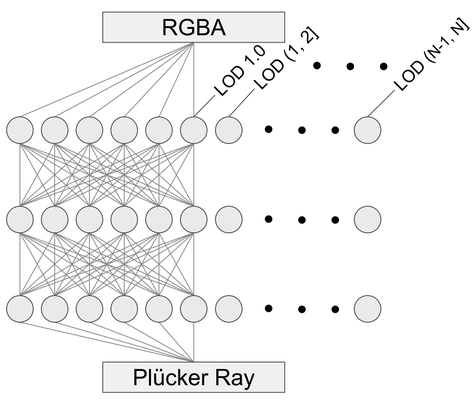

Variable-size layers enable arbitrary size execution with hundreds of performance levels.

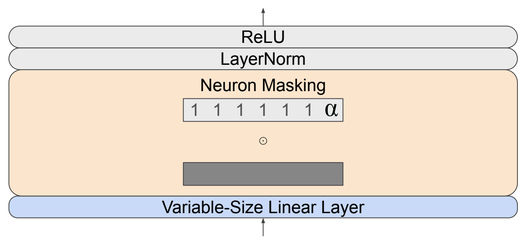

Neuron masking is used to continuously interpolate between neural network sizes.

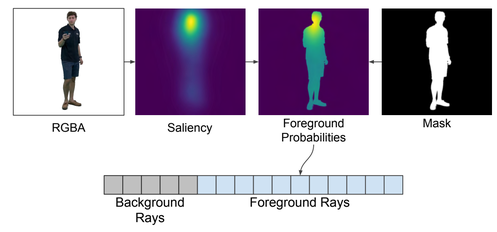

Saliency-based importance sampling helps salient regions resolve at earlier LODs.

Experimental Results

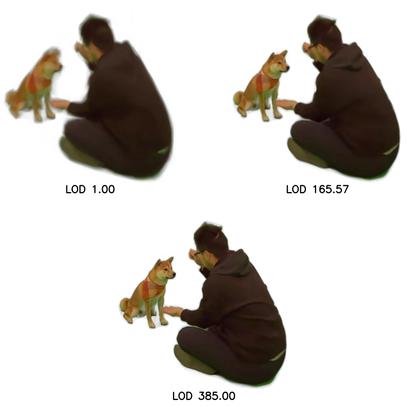

We can now render at any arbitrary LOD.

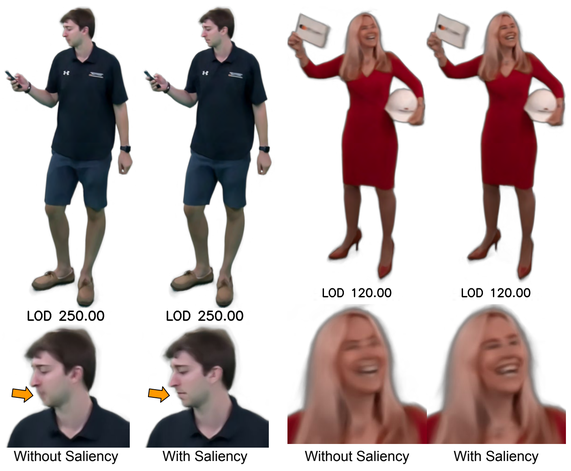

Face details resolve at lower LODs with saliency-based sampling.

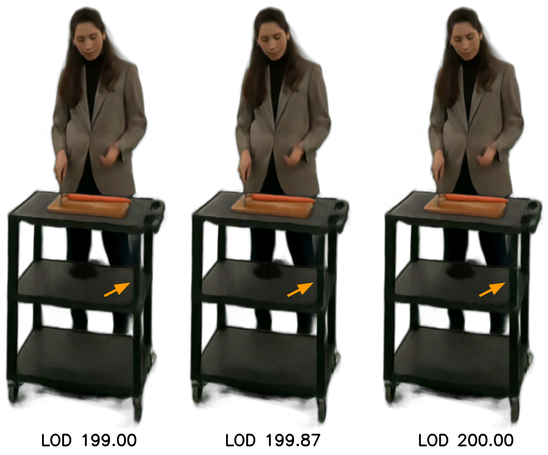

Details emerge gradually with fractional LODs.

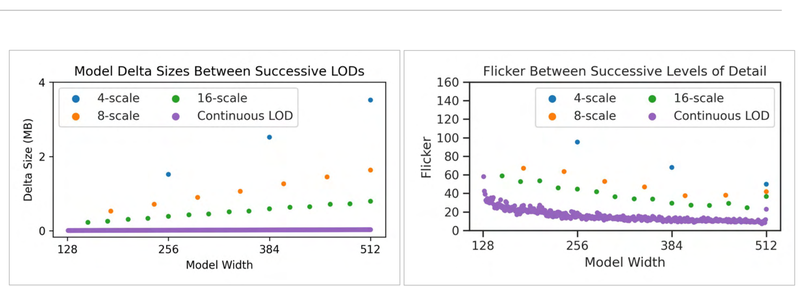

We observe smaller model delta sizes and less flickering during transitions.

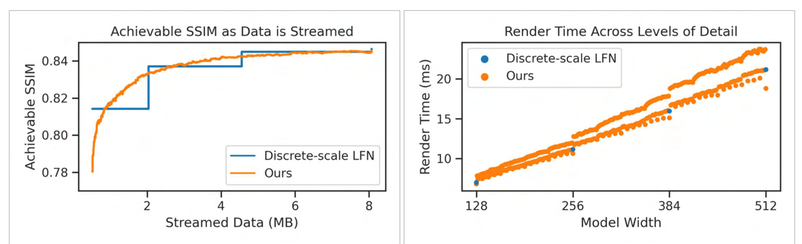

We observe smoother scaling between available data, performance, and quality.

BMVC 2023

University of Maryland Department of Computer Science
UMIACS University of Maryland Institute for Advanced Computer Studies