Image-Based Rendering

CSE 274, Lecture 8: Deep Learning for IBR

Ravi Ramamoorthi
http://www.cs.ucsd.edu/~ravir

1

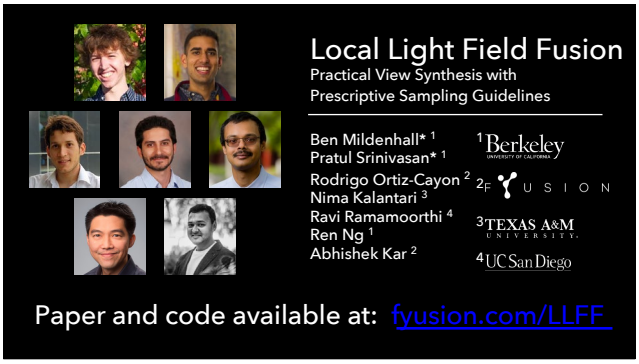

2

Virtual Experiences of Real-World Scenes

3

Input Images

4

Output Virtual Experience

5

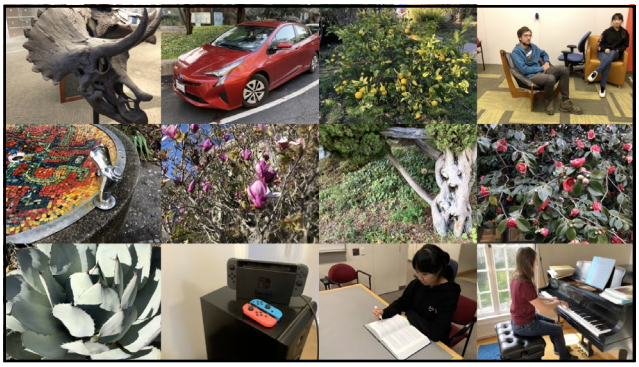

6

Novel View Synthesis Problem Setup

Inputs: Sampled Images of a Scene

Outputs: Rendered Novel Views

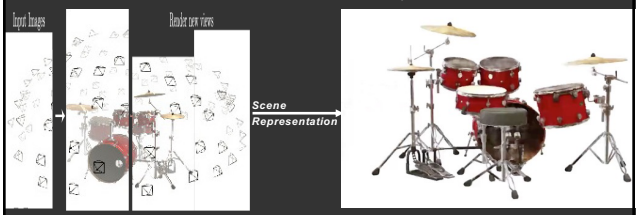

7

Light Fields

Gershun's and Moon's idea of a light field:

Radiance as a function of a ray or line: L (x, y, z, θ, ϕ)

- In "free space" (no occluders) 5D reduces to 4D
  - Exterior of the convex hull of an object
  - Interior of an environment
- Images are 2D slices
  - Insert acquired imagery
  - Extract image from a given viewpoint

From Levoy and Hanrahan. Light Field Rendering. SIGGRAPH 96

8

4D Light Field

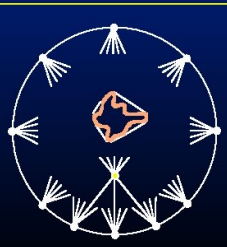

9

Very dense sampling yields robust solutions

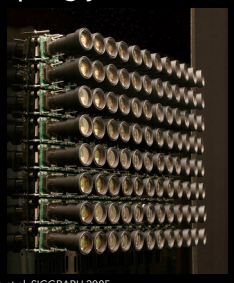

Stanford Multi-Camera Array. Wilburn et al. SIGGRAPH 2005

10

Very dense sampling yields robust solutions

Google "Welcome to Light Fields". Overbeck et al. SIGGRAPH Asia 2018

11

How do we achieve robust high-quality performance with sparser casual input view sampling?

12

13

14

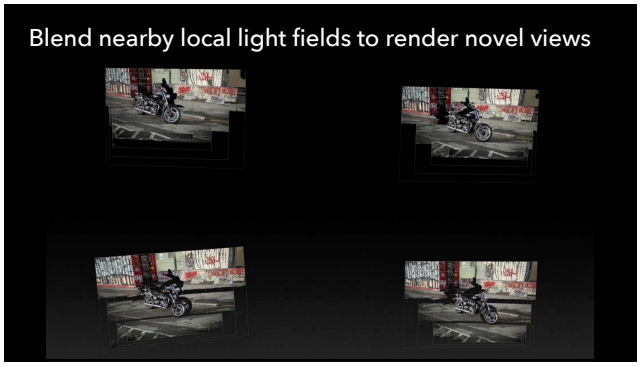

15

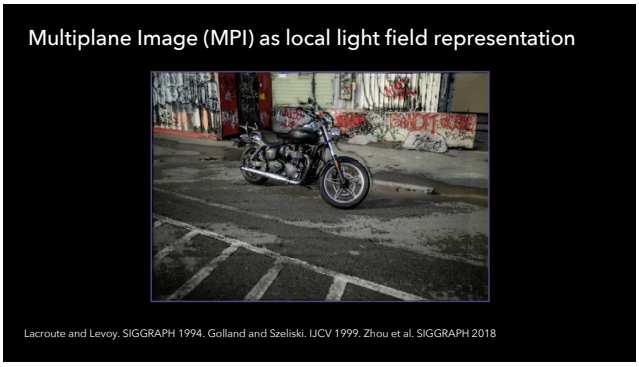

16

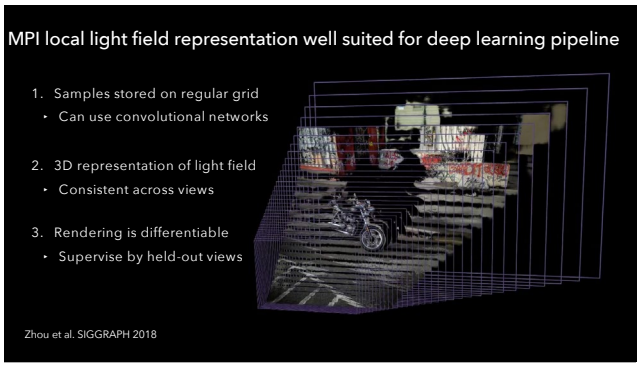

17

18

19

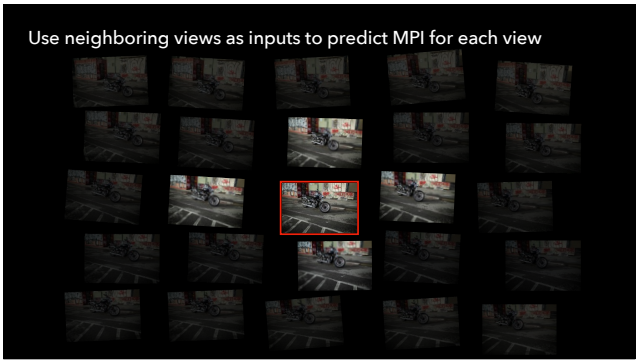

20

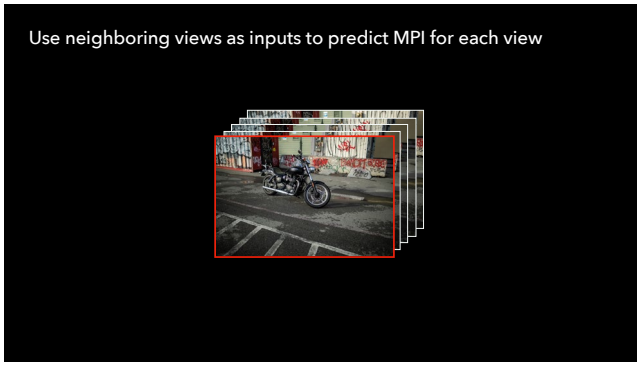

21

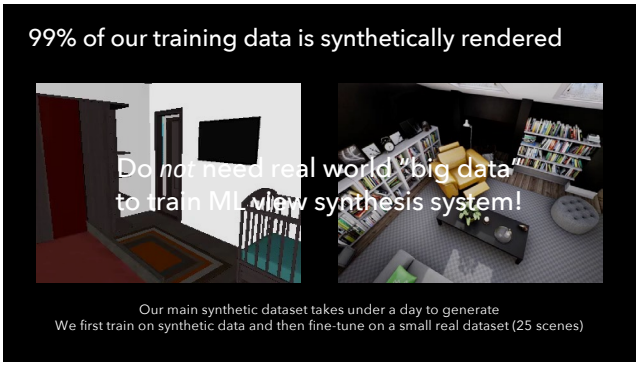

22

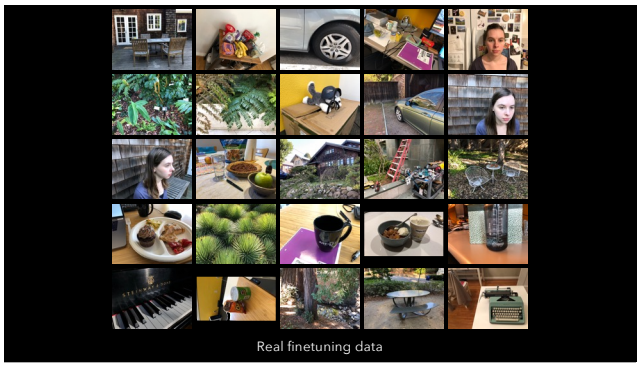

23

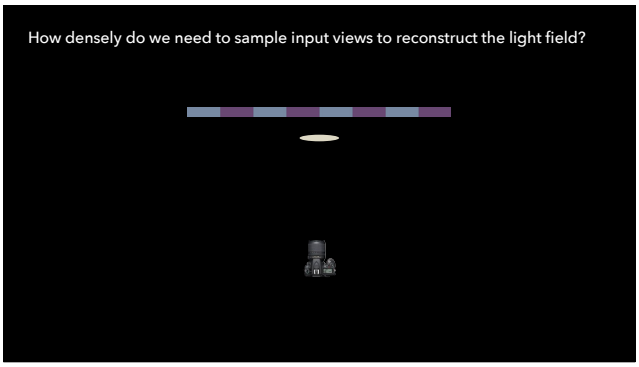

24

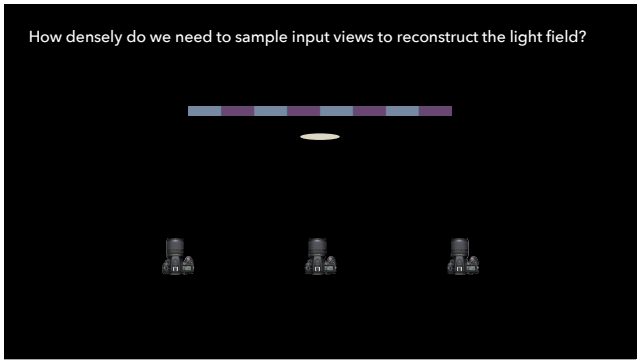

25

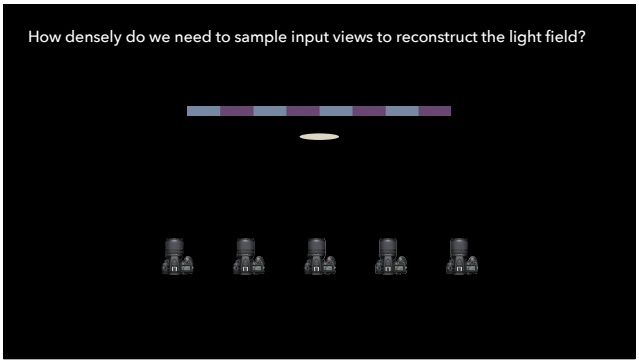

26

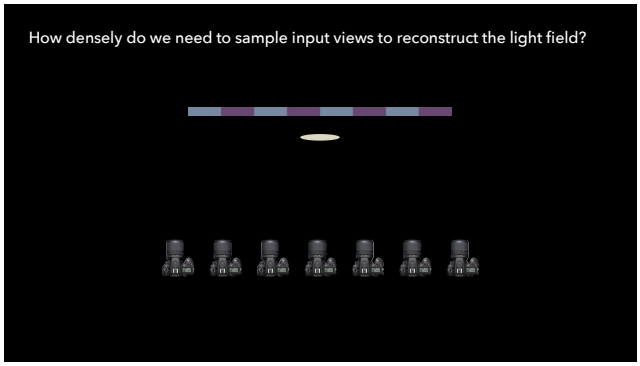

27

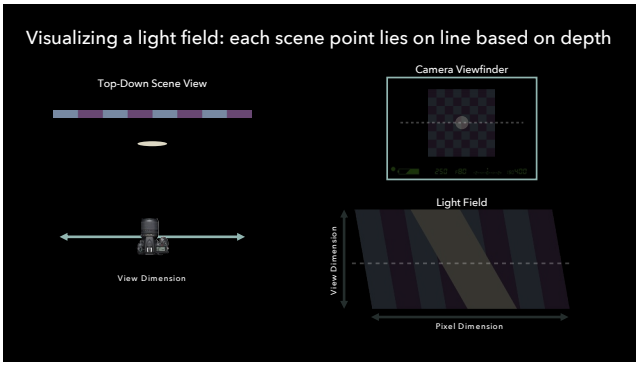

28

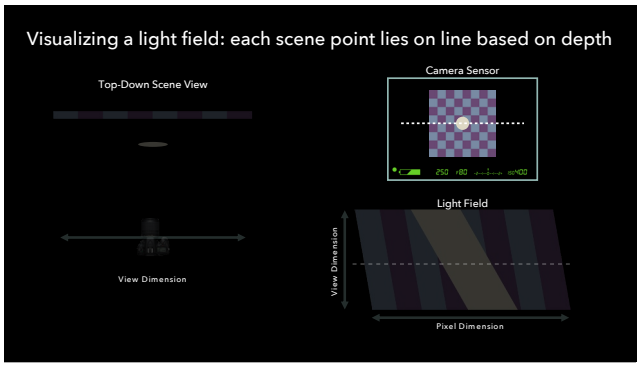

29

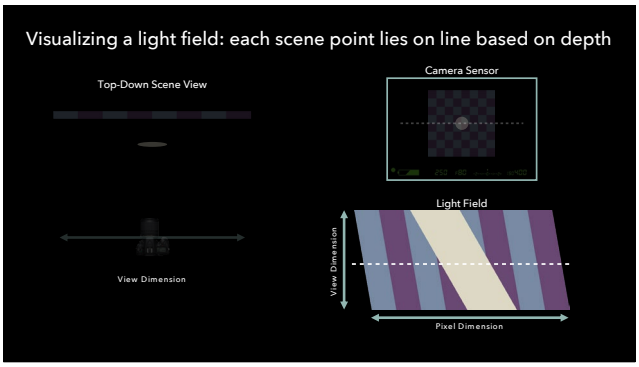

30

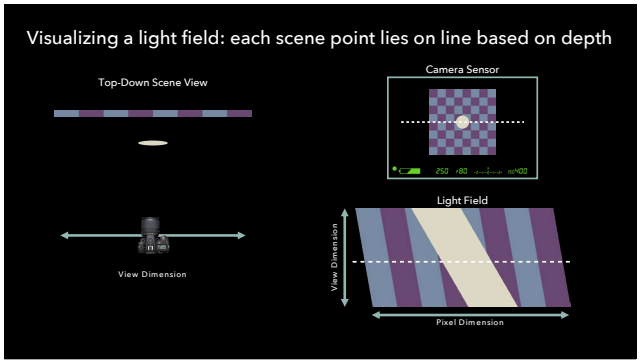

31

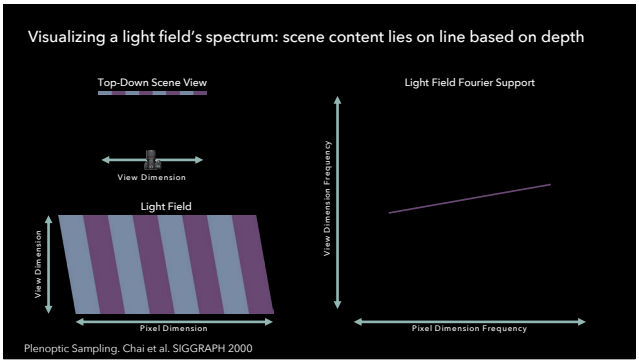

32

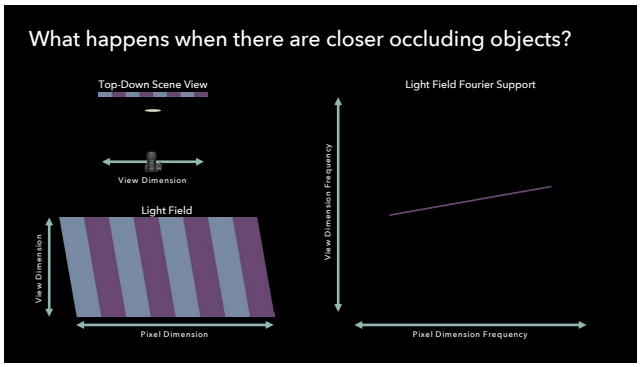

33

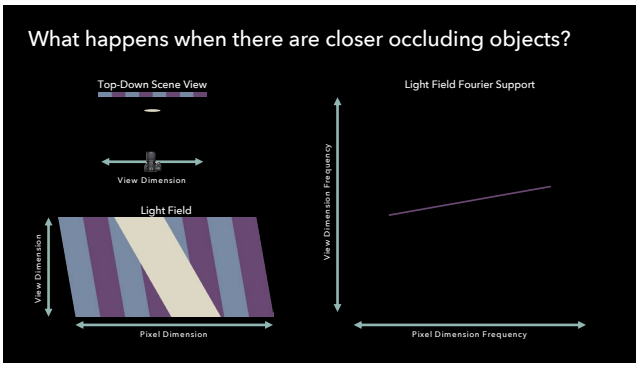

34

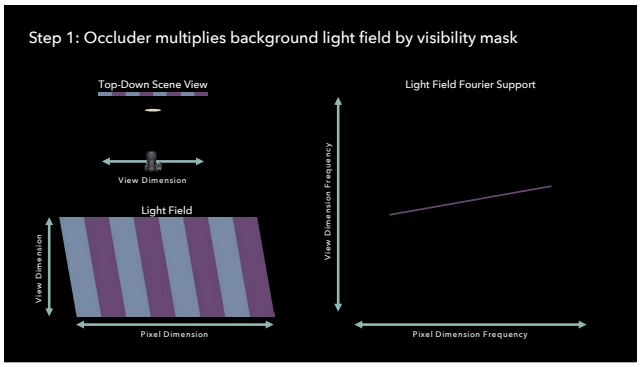

35

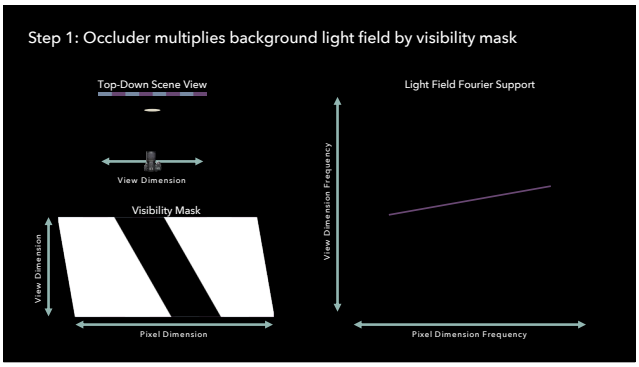

36

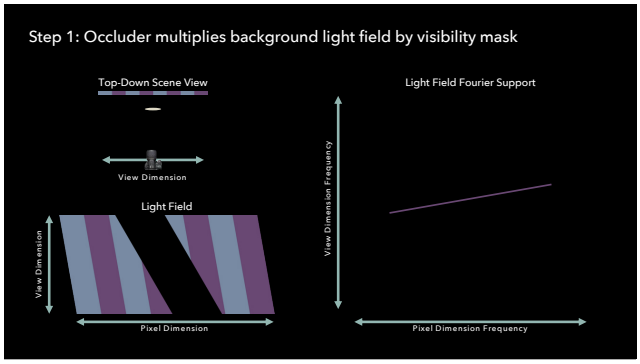

37

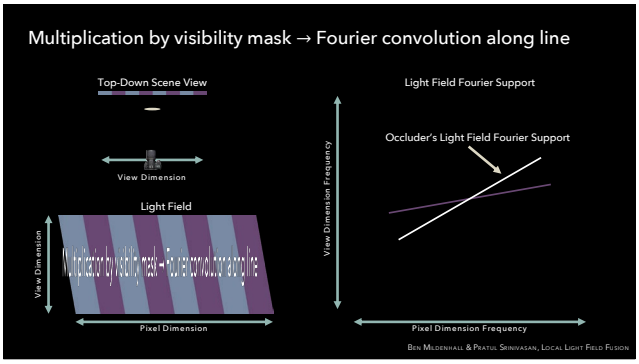

38

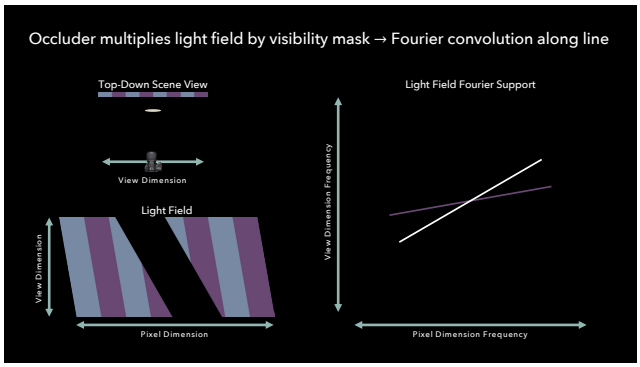

39

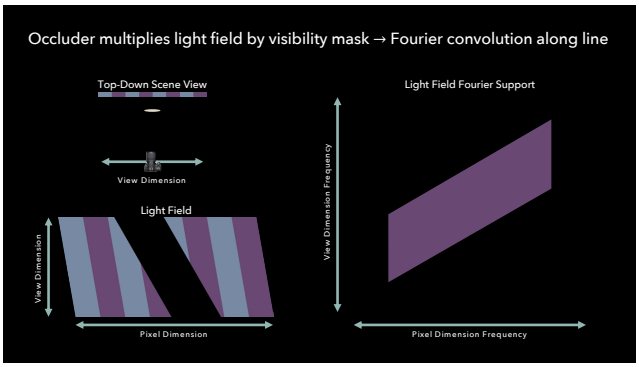

40

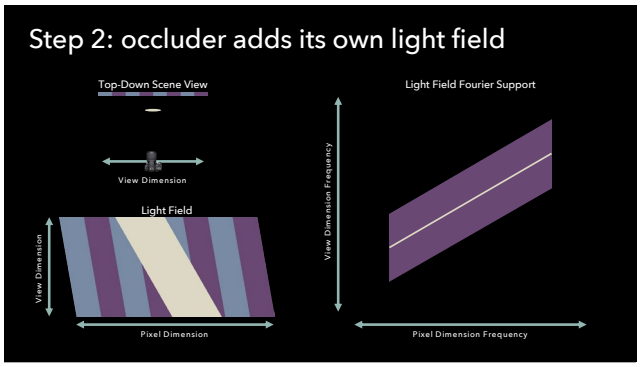

41

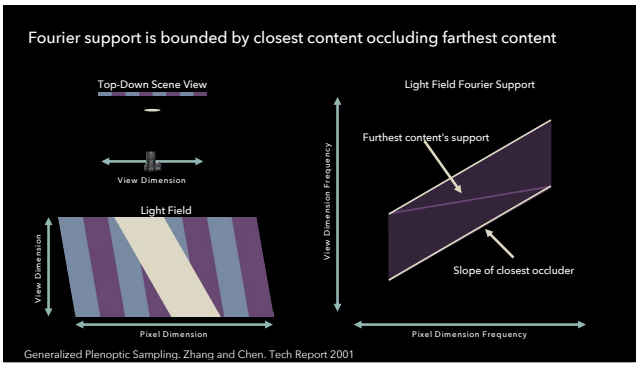

42

Nyquist rate view sampling means packing parallelograms as closely as possible

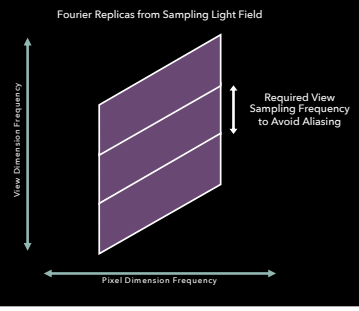

43

Nyquist rate means 1 pixel max disparity between adjacent views

Capturing a scene with closest content 0.5 m away for viewing on current VR headsets requires an image every millimeter!

44

What if we have *ideal* MPI layers for each view?

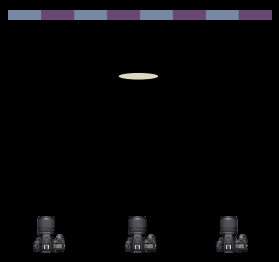

45

What if we have *ideal* MPI layers for each view?

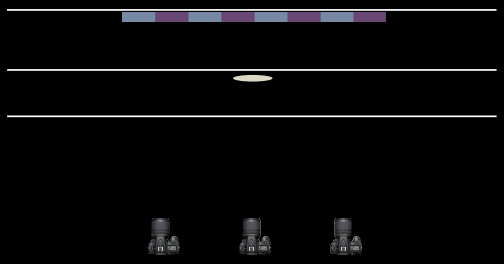

46

What if we have *ideal* MPI layers for each view?

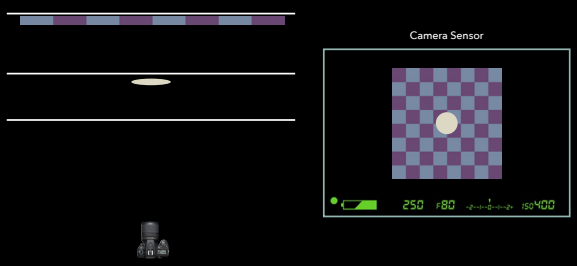

47

What if we have *ideal* MPI layers for each view?

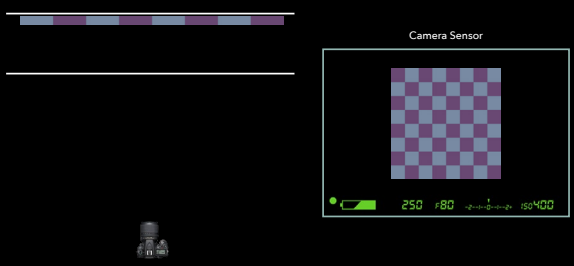

48

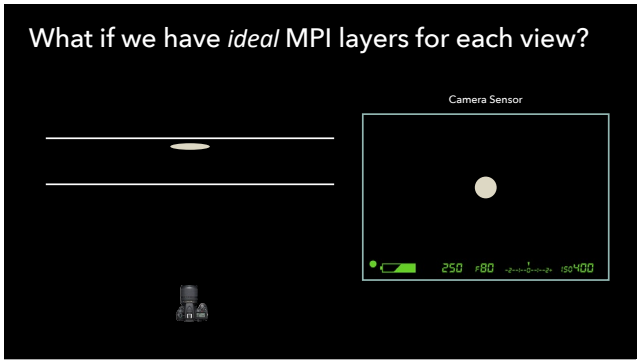

49

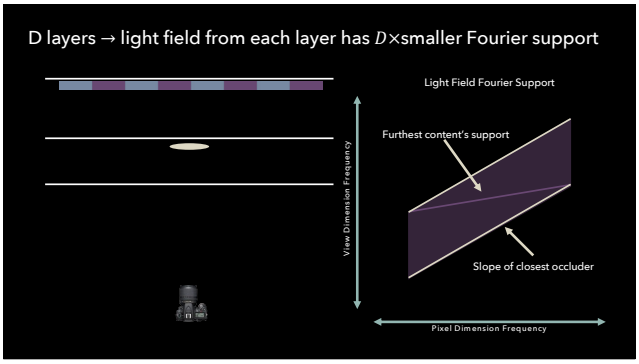

50

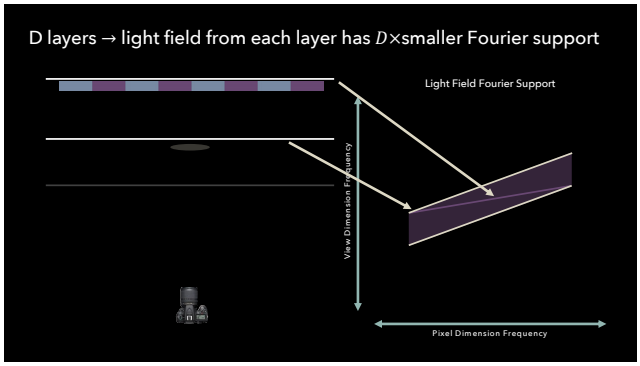

51

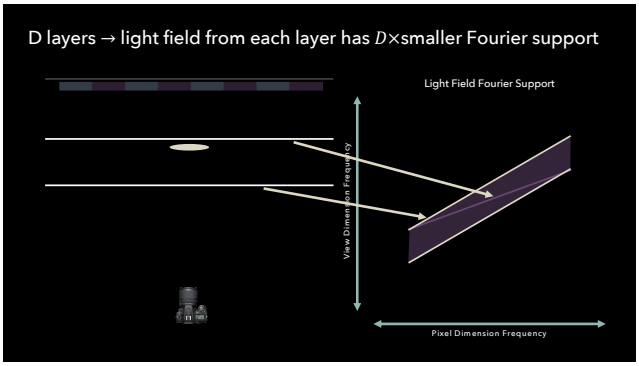

52

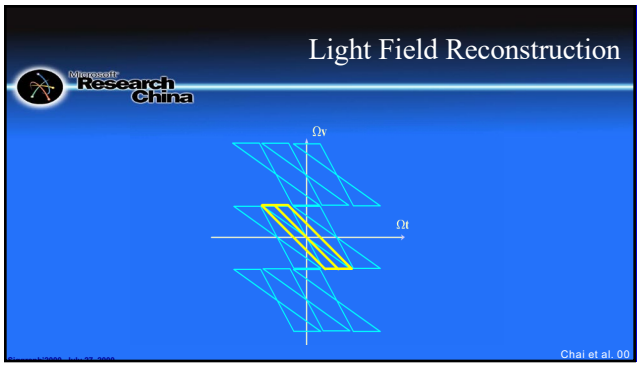

53

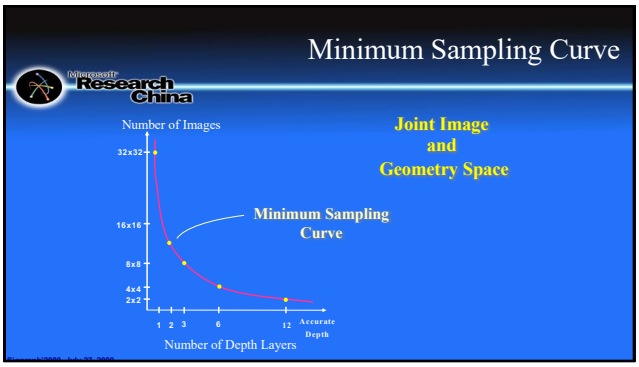

54

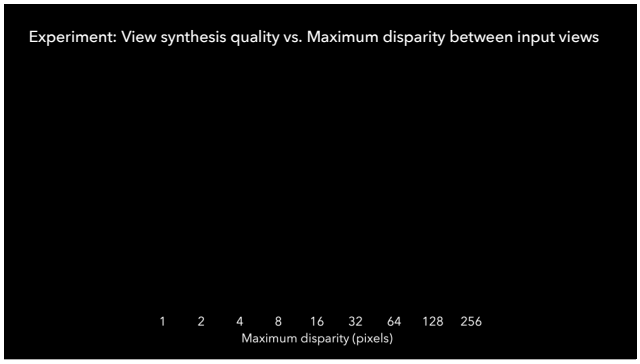

55

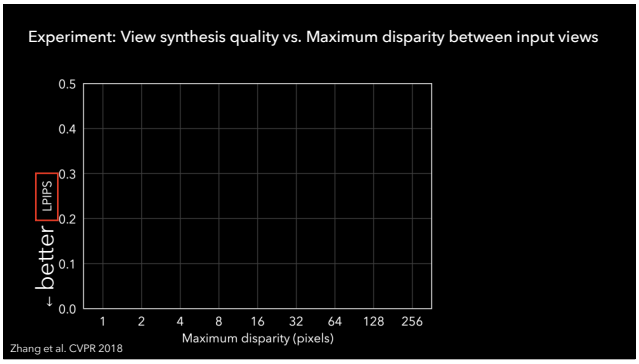

56

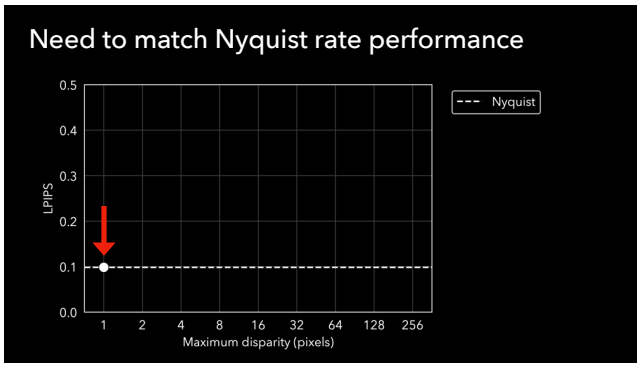

57

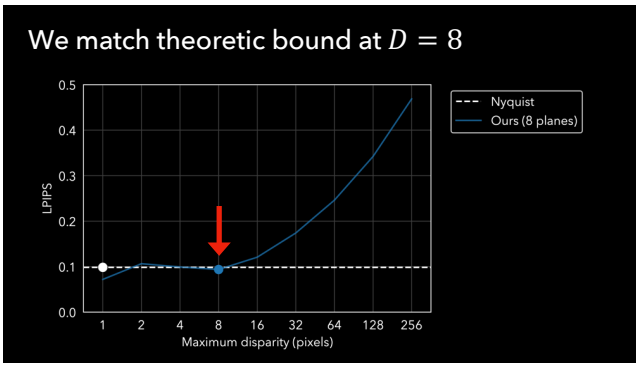

58

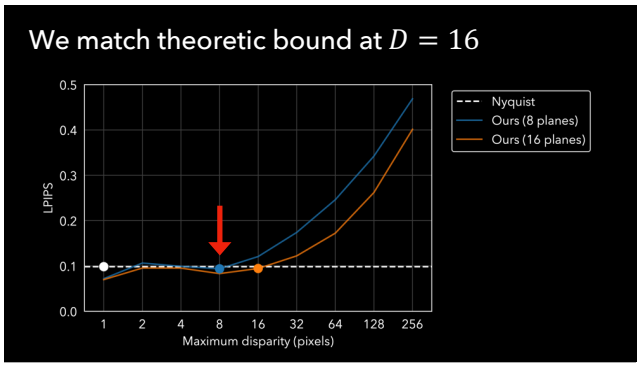

59

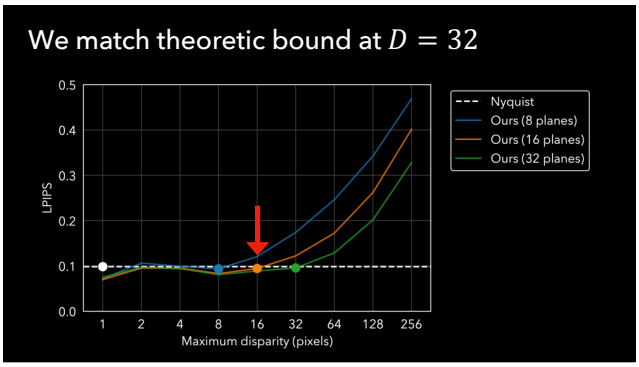

60

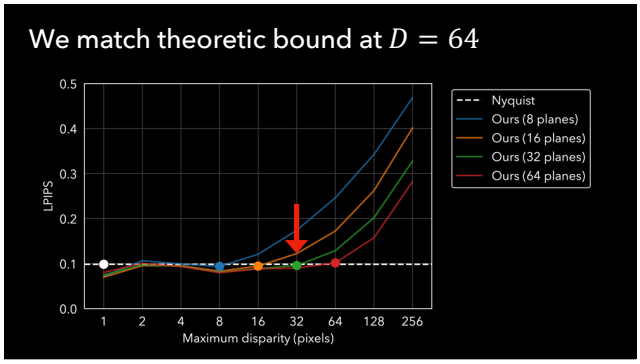
We match theoretic bound at $D = 64$
LPIPS
0.5
0.4
0.3
0.2
0.1
0.0
1
2
4
8
16
32
64
128
256
Maximum disparity (pixels)
Nyquist
Ours (8 planes)
Ours (16 planes)
Ours (32 planes)
Ours (64 planes)

61

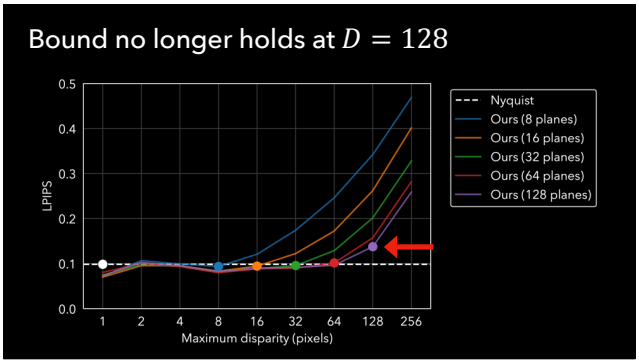
Bound no longer holds at $D = 128$
LPIPS
0.5
0.4
0.3
0.2
0.1
0.0
1
2
4
8
16
32
64
128
256
Maximum disparity (pixels)
Nyquist
Ours (8 planes)
Ours (16 planes)
Ours (32 planes)
Ours (64 planes)
Ours (128 planes)

62

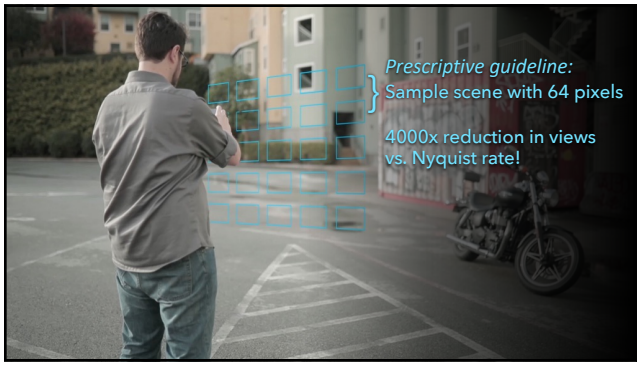
Prescriptive guideline:
Sample scene with 64 pixels
4000x reduction in views
vs. Nyquist rate!

63

64

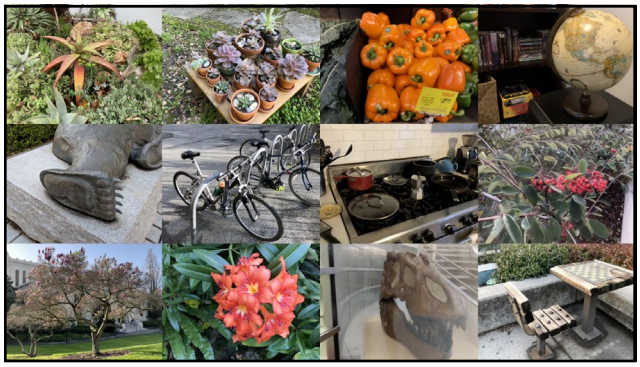

65