Computer Graphics II: Rendering

CSE 168 [Spr 26], Lecture 17: Image-Based Rendering
Ravi Ramamoorthi

http://viscomp.ucsd.edu/classes/cse168/sp26

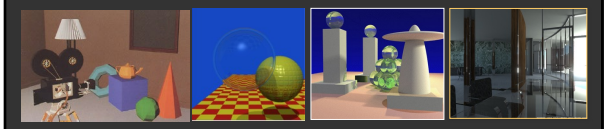

1

Motivation for Lecture

- Image-Based Rendering major new idea in graphics in past 30 years
- Many of the rendering methods, especially precomputed techniques borrow from it
- And many methods use measured data
- Also, images are an important source for rendering
- Sampled data rapidly becoming popular
- Core IBR problem of view synthesis/light fields renewed popularity (VR other applications)

2

Traditional Modeling and Rendering

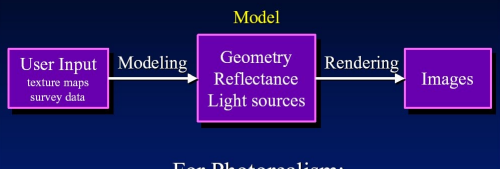

For Photorealism:

Modeling is Hard **Rendering is Slow**

Next few slides courtesy Paul Debevec: SIGGRAPH 99 course notes

3

Can we model and render this?

What do we want to do with the model?

4

Image-Based Modeling and Rendering

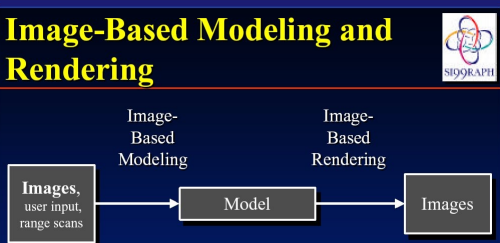

5

The Spectrum of IBMR

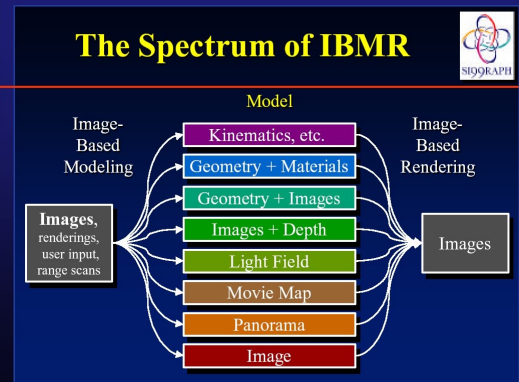

6

IBR: Pros and Cons

- Advantages
  - Easy to capture images: photorealistic by definition
  - Simple, universal representation
  - Often bypass geometry estimation?
  - Independent of scene complexity?
- Disadvantages
  - WYSIWYG but also WYSIAYG
  - Explosion of data as flexibility increased
  - Often discards intrinsic structure of model?
- Today, IBR-type methods also often used in synthetic rendering (e.g. real-time rendering PRT)
  - General concept of data-driven graphics, appearance
  - Also, data-driven geometry, animation, simulation
  - Spawned light field cameras for image capture

7

Image-Based Models: What do they allow?

Model	Movement	Geometry	Lighting
Geometry + Materials	Continuous	Global	Dynamic
Geometry + Images	Continuous	Global	Fixed
Images + Depth	Continuous	Local	Fixed
Light Field	Continuous	None	Fixed
Movie Map	Discrete	None	Fixed
Panorama	Rotation	None	Fixed
Image	None	None	Fixed

8

IBR: A brief history

- Texture maps, bump maps, environment maps [70s]
- Poggio MIT 90s: Faces, image-based analysis/synthesis
- Mid-Late 90s
  - Chen and Williams 93, View Interpolation [Images+depth]
  - Chen 95 Quicktime VR [Images from many viewpoints]
  - McMillan and Bishop 95 Plenoptic Modeling [Images w disparity]
  - Gortler et al, Levoy and Hanrahan 96 Light Fields [4D]
  - Shade et al. 98 Layered Depth Images [2.5D]
  - Debevec et al. 00 Reflectance Field [4D]
  - Inverse rendering (Marschner,Sato,Yu,Boivin,…)
- Today: IBR hasn't replaced conventional rendering, but has brought sampled and data-driven representations to graphics

9

Game #1: increase the dimensionality

2D rgb	texture
2D rgbz	range image
2.5D rgbαz	layered depth images
4D rgb	light field / Lumigraph
4D rgbz	array of range images
4.5D rgbαz	layered light fields

10

Game #2: replace the quantity represented

4D rgb	light field / Lumigraph
$\{u, v, s, t\}$	
5D rgb	plenoptic function
$\{x, y, z\} \times \{\theta, \varphi\}$	
6D ρ	free-space BRDF field
$\{u, v, s, t\} \times \{\theta_i, \varphi_i\}$	
7D ρ	BRDF volume
$\{x, y, z\} \times \{\theta_i, \varphi_i, \theta_r, \varphi_r\}$	

11

Outline

- Overview of IBR
- Basic approaches
  - *Image Warping*
    - *[2D + depth. Requires correspondence/disparity]*
  - Light Fields [4D]
  - Survey of some early work

12

To Do

- Project proposals due tomorrow (May 28)
- Final Projects due Jun 10
- PLEASE FILL OUT SET EVALUATIONS!!
- KEEP WORKING HARD

13

View Interpolation for Image Synthesis

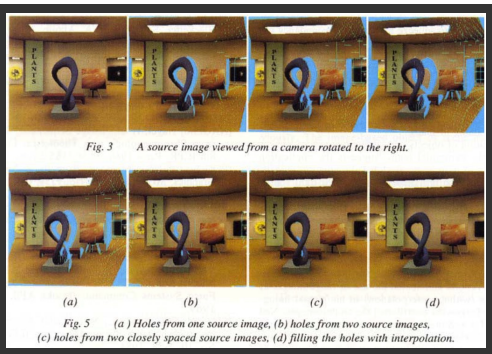

Fig. 3 A source image viewed from a camera rotated to the right.

(a) (b) (c) (d)

Fig. 5 (a) Holes from one source image, (b) holes from two source images, (c) holes from two closely spaced source images, (d) filling the holes with interpolation.

Chen and Williams. View Interpolation for Image Synthesis. SIGGRAPH 93 (Seminal Graphics Papers v1 and v2)

14

Images as a Collection of Rays

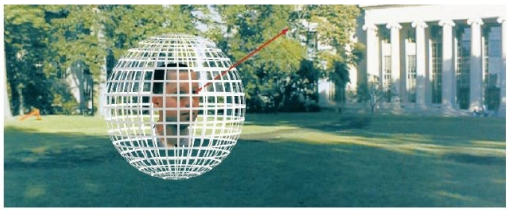

An image is a subset of the rays seen from a given point
- this "space" of rays occupies two dimensions

Warping slides courtesy Leonard McMillan, SIGGRAPH 99 course notes

15

The Plenoptic Function

✔The set of rays seen from all points …

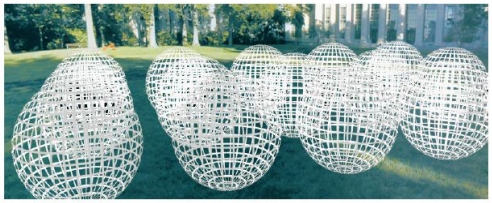

$$p = P(\theta, \phi, x, y, z, \lambda, t)$$

16

Image-based rendering is about

…reconstructing a plenoptic function from a set of samples taken from it.

✔Ignoring time, and selecting a discrete set of wavelengths gives a 5-D plenoptic function

17

Where to Begin?

✔Pinhole camera model

- Defines a mapping from image points to rays in space

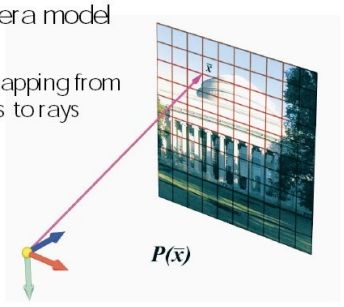

18

Mapping from Rays to Points

- Simple Derivation

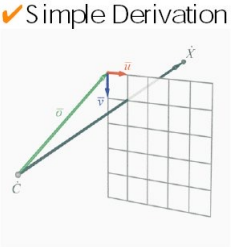

$$P = \begin{bmatrix} u_x & v_x & o_x \\ u_y & v_y & o_y \\ u_z & v_z & o_z \end{bmatrix}$$

$$\dot{X} = \dot{C} + t\,P\,\vec{x}$$

19

Correspondence

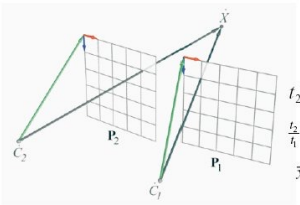

$$\dot{C}_2 + t_2 P_2 \vec{x}_2 = \dot{C}_1 + t_1 P_1 \vec{x}_1$$

$$t_2 P_2 \vec{x}_2 = \dot{C}_1 - \dot{C}_2 + t_1 P_1 \vec{x}_1$$

$$t_2 \vec{x}_2 = P_2^{-1}(\dot{C}_1 - \dot{C}_2) + t_1 P_2^{-1} P_1 \vec{x}_1$$

$$\frac{t_2}{t_1}\vec{x}_2 = \frac{1}{t_1} P_2^{-1}(\dot{C}_1 - \dot{C}_2) + P_2^{-1} P_1 \vec{x}_1$$

$$\vec{x}_2 \doteq \frac{1}{t_1}\underbrace{P_2^{-1}(\dot{C}_1 - \dot{C}_2)}_{e_{21}} + \underbrace{P_2^{-1} P_1}_{H_{21}}\vec{x}_1$$

20

Warping in Action

- A 3D Warp

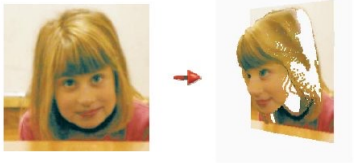

21

Visibility

- The warping equation determines where points go…

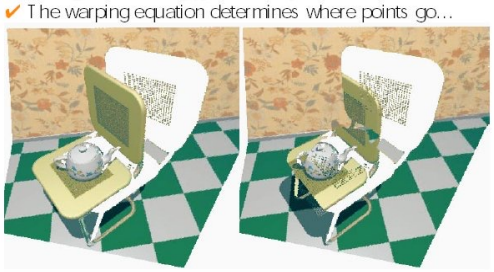

… but that is not sufficient

22

Partition Reference Image

- Project the *desired* center-of-projection onto the reference image

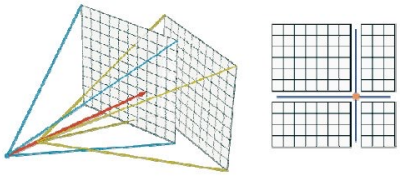

23

Enumeration

- Drawing toward the projected point guarantees an *occlusion compatible* ordering
- Ordering is consistent with a painter's algorithm
- Independent of the scene's contents
- Easily generalized to other viewing surfaces
- No auxiliary information required

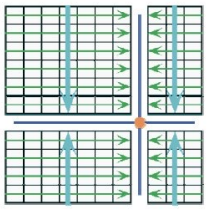

24

Reconstruction

- ✔ Typical images are discrete, not continuous
- ✔ An image can be formed by different geometries

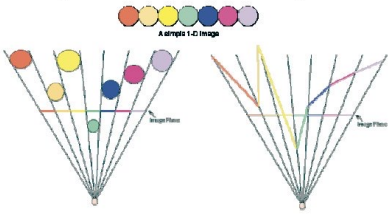

25

Outline

- Overview of IBR
- Basic approaches
  - Image Warping
    - [2D + depth. Requires correspondence/disparity]
  - *Light Fields [4D]*
  - Survey of some early work

26

Light Field Rendering

Marc Levoy *Pat Hanrahan*

Computer Science Department
Stanford University

27

Apple's QuickTime VR

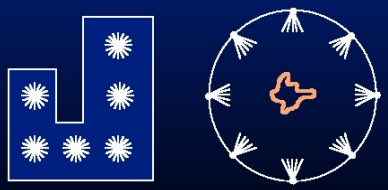

Outward Inward

28

Generating New Views

Problem: fixed vantage point/center

One Solution: view interpolation

- Interpolating between range images (Chen and Willams, 1993)
- Correspondences and epipolar analysis (McMillan and Bishop, 1995)

-> Requires depths or correspondences:

must be extracted from acquired imagery

relatively expensive and error-prone morph

29

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

30

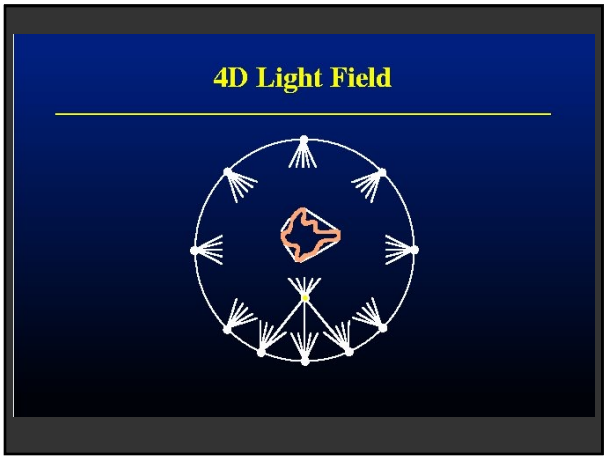

31

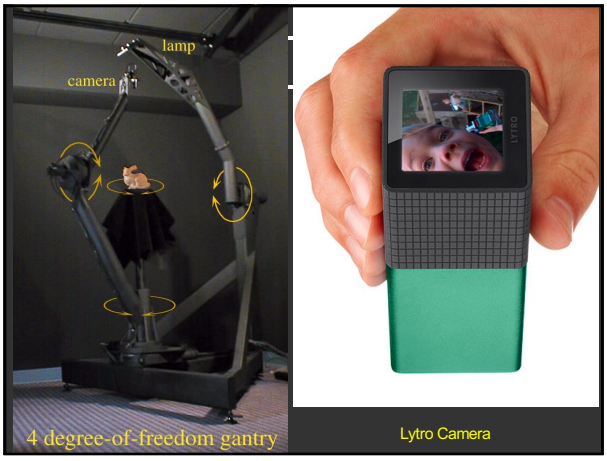

32

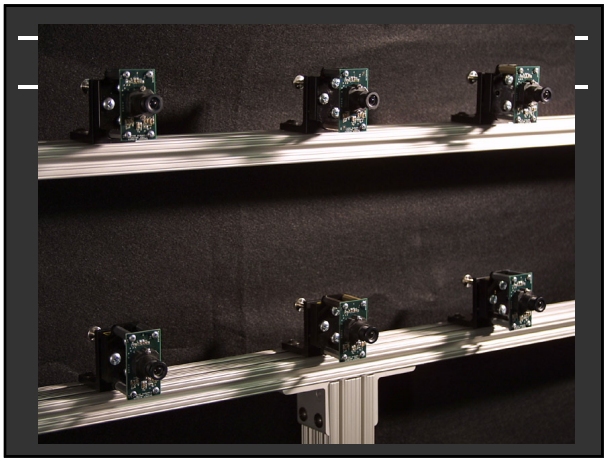

33

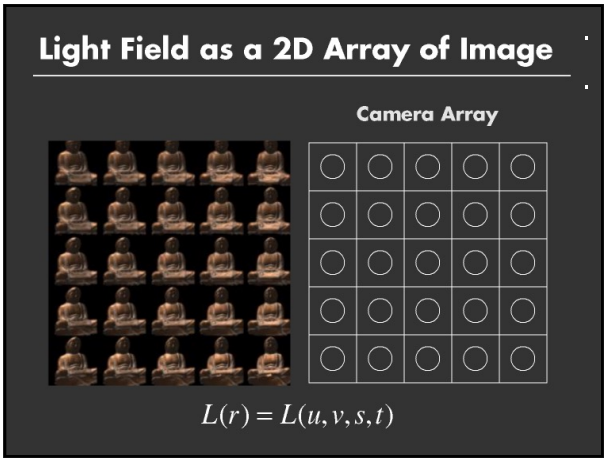

34

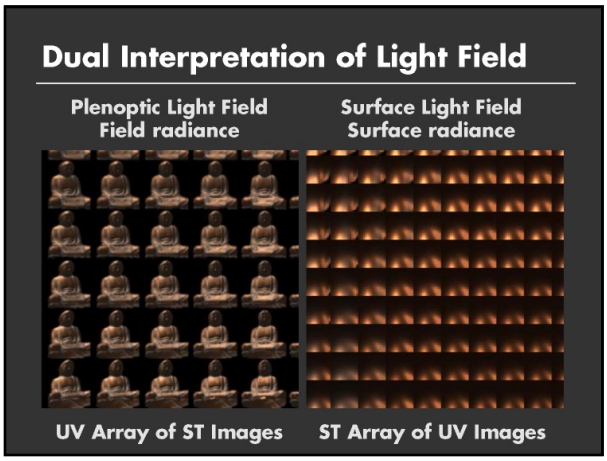

35

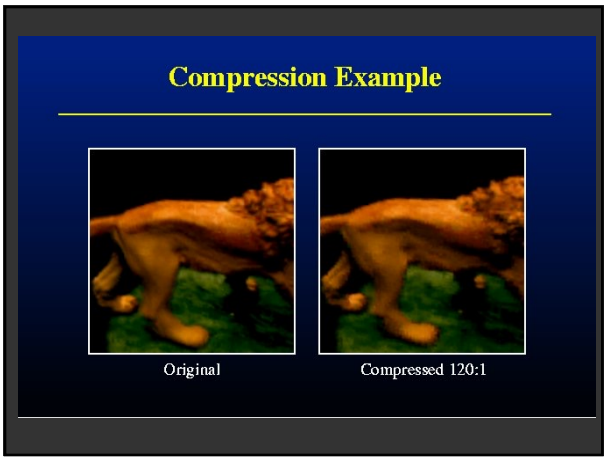

36

Outline

- Overview of IBR
- Basic approaches
  - Image Warping
    - [2D + depth. Requires correspondence/disparity]
  - Light Fields [4D]
  - *Survey of some early work*

37

Layered Depth Images [Shade 98]

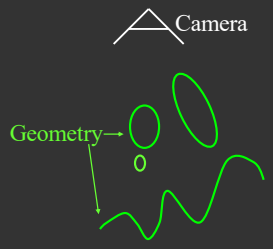

Slide from Agrawala, Ramamoorthi, Heirich, Moll, SIGGRAPH 2000

38

Layered Depth Images [Shade 98]

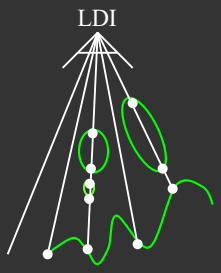

39

Layered Depth Images [Shade 98]

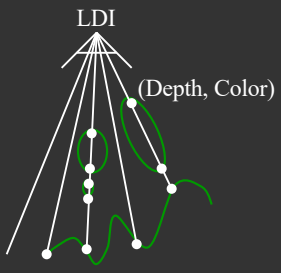

40

41

Surface Light Fields

- Miller 98, Nishino 99, Wood 00
- Reflected light field (lumisphere) on surface
- Explicit geometry as against light fields. Easier compress

42

Acquiring Reflectance Field of Human Face [Debevec et al. SIGGRAPH 00]

Illuminate subject from many incident directions

43

Example Images

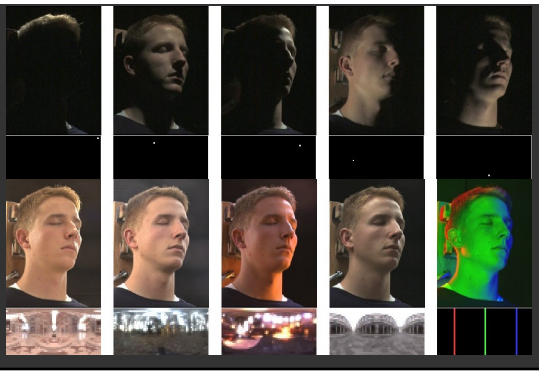

44

Outline

- Overview of IBR
- Basic approaches
  - Image Warping
    - [2D + depth. Requires correspondence/disparity]
  - Light Fields [4D]
  - *Survey of some recent work*
    - *Sampled data representations*

45

Conclusion (my views)

- IBR initially spurred great excitement: revolutionize pipeline
- But, IBR in pure form not really practical
  - WYSIAYG
  - Explosion as increase dimensions (8D transfer function)
  - Good compression, flexibility needs at least implicit geometry/BRDF
- Real future is sampled representations, data-driven method
  - Acquire (synthetic or real) data
  - Good representations for interpolation, fast rendering
  - Much of visual appearance, graphics moving in this direction
- Understand from Signal-Processing Viewpoint
  - Sampling rates, reconstruction filters
  - Factored representations, Fourier analysis
- Light Fields fundamental in many ways, including imaging
  - *Renewed interest in view synthesis (Mildenhall et al. SIG 19, NeRFs)*

46

Virtual Experiences of Real-World Scenes

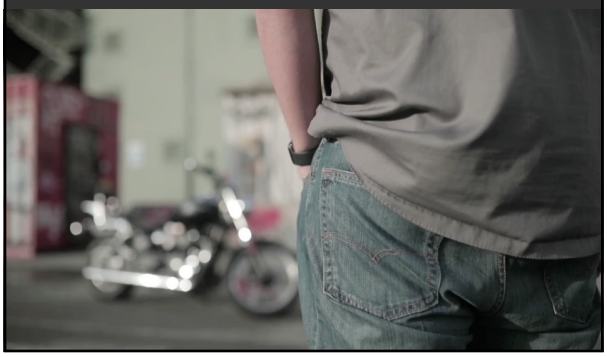

47

Input Images

48

Output Light Field

49

Local Light Field Fusion

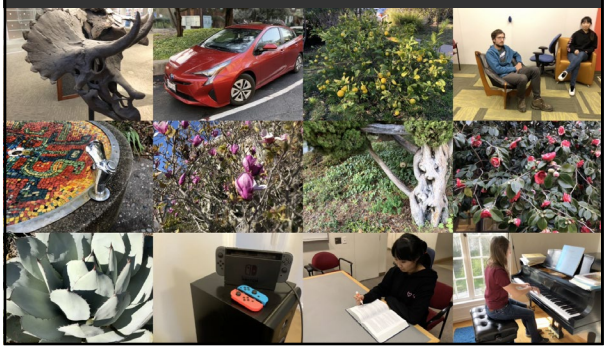

50

Neural Radiance Fields

51

Reducing compute-per-sample: learned hash grids (Instant NGP)

Müller, et al 2022, Instant Neural Graphics Primitives with a Multiresolution Hash Encoding

52

3D Gaussian Splats for Radiance Fields

Kerbl, Kopanas, Leimkuhler, Drettakis 23

53

NeRFs for Digital Twins

54

NeRFs for Landscapes

Instant NeRF (@jonstephens85 on Twitter); slide courtesy Thomas Muller

55

NeRFs for Furry Volumes

Instant NeRF (@vibrantnebula on Twitter); slide courtesy Thomas Muller

56

NeRFs + Language Models

- Text-based editing with large language models (Haque et al.)

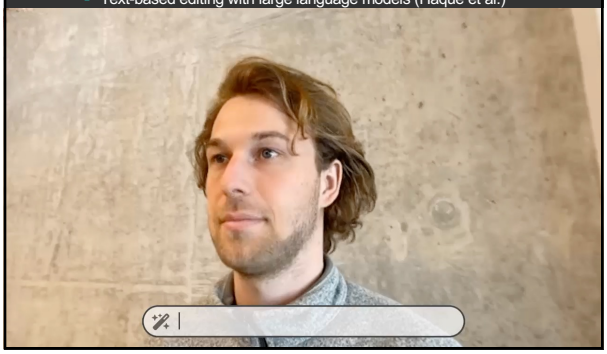

57

NeRFs in Production (Google, Luma)

Google Street View

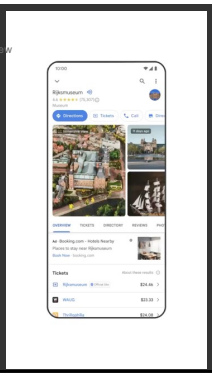

58

Personalized Avatars

Vision Pro (Apple)

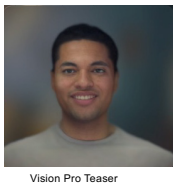

Vision Pro Teaser

Codec Avatars (Meta)

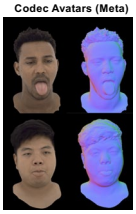

Ma et al.

Project Starline (Google)

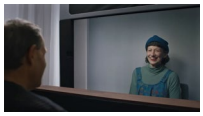

blog.google/technology/research/project-starline

Require multiple cameras, light stage capture, or depth information

59

Real-Time Radiance Fields from Single Image

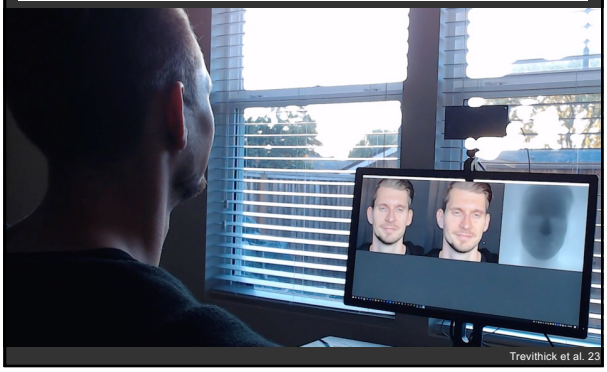

Trevithick et al. 23

60

Lifting Text-Based Avatars to 3D

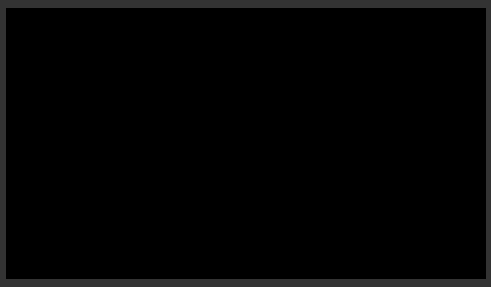

61

3D Videoconferencing

62