Course Wrap-Up

Computer Graphics
CMU 15-462/15-662

Where is graphics headed?

- **Extremely diverse field**
  - **Anything related to synthesis of perceptual phenomena**
  - **Images, sound, physical objects, …**
- **Some very "hot" topics right now:**
  - **Virtual Reality / Augmented Reality (VR/AR)**
  - **3D Fabrication (e.g., 3D printing)**
  - **Computational Photography**

Virtual reality (VR) vs augmented reality (AR)

VR = virtual reality

User is completely immersed in virtual world (sees only light emitted by display

AR = augmented reality

Display is an overlay that augments user's normal view of the real world (e.g., terminator)

Image credit: Terminator 2 (naturally)

VR headsets

Oculus Rift (Crescent Bay Prototype)

Sony Morpheus

HTV Vive

Google Cardboard

AR headset: Microsoft Hololens

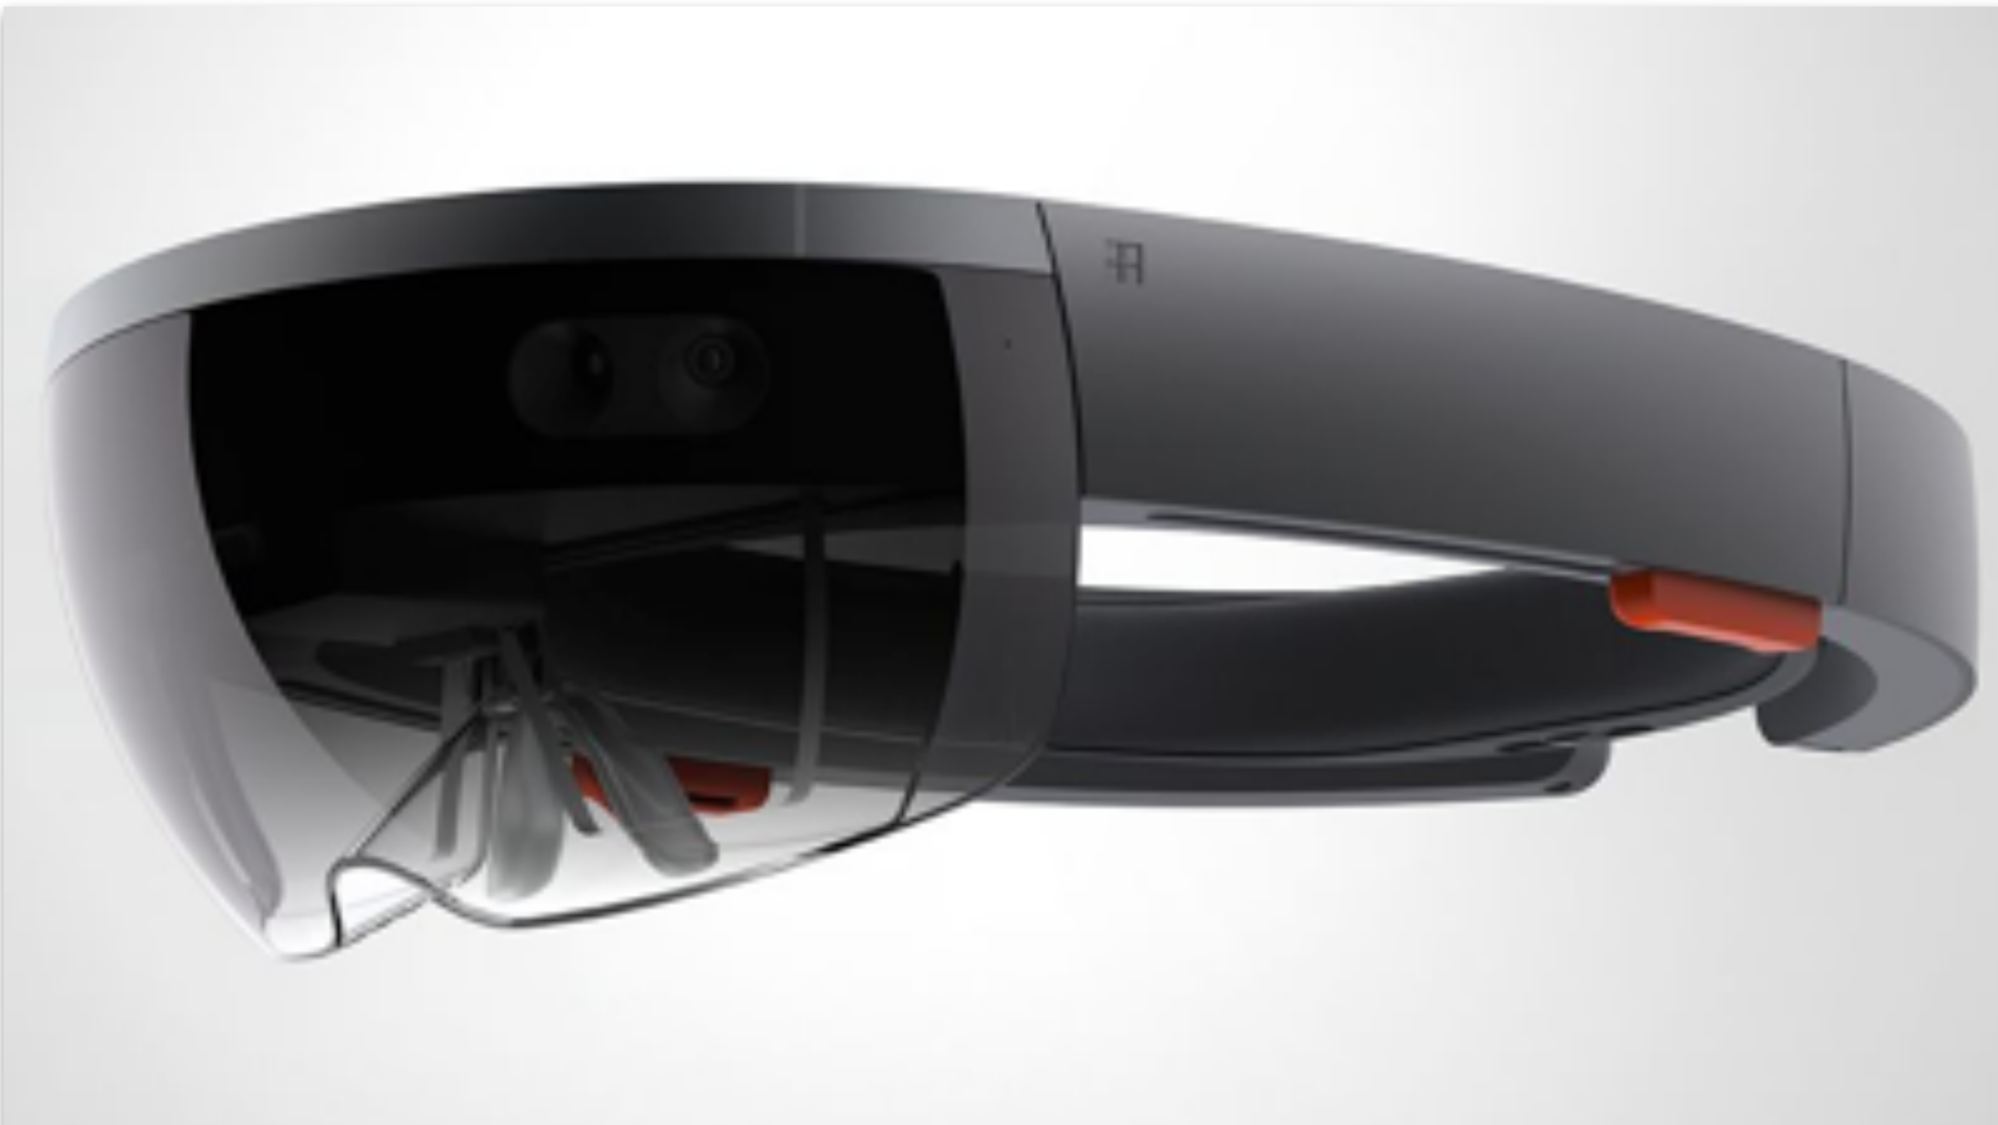

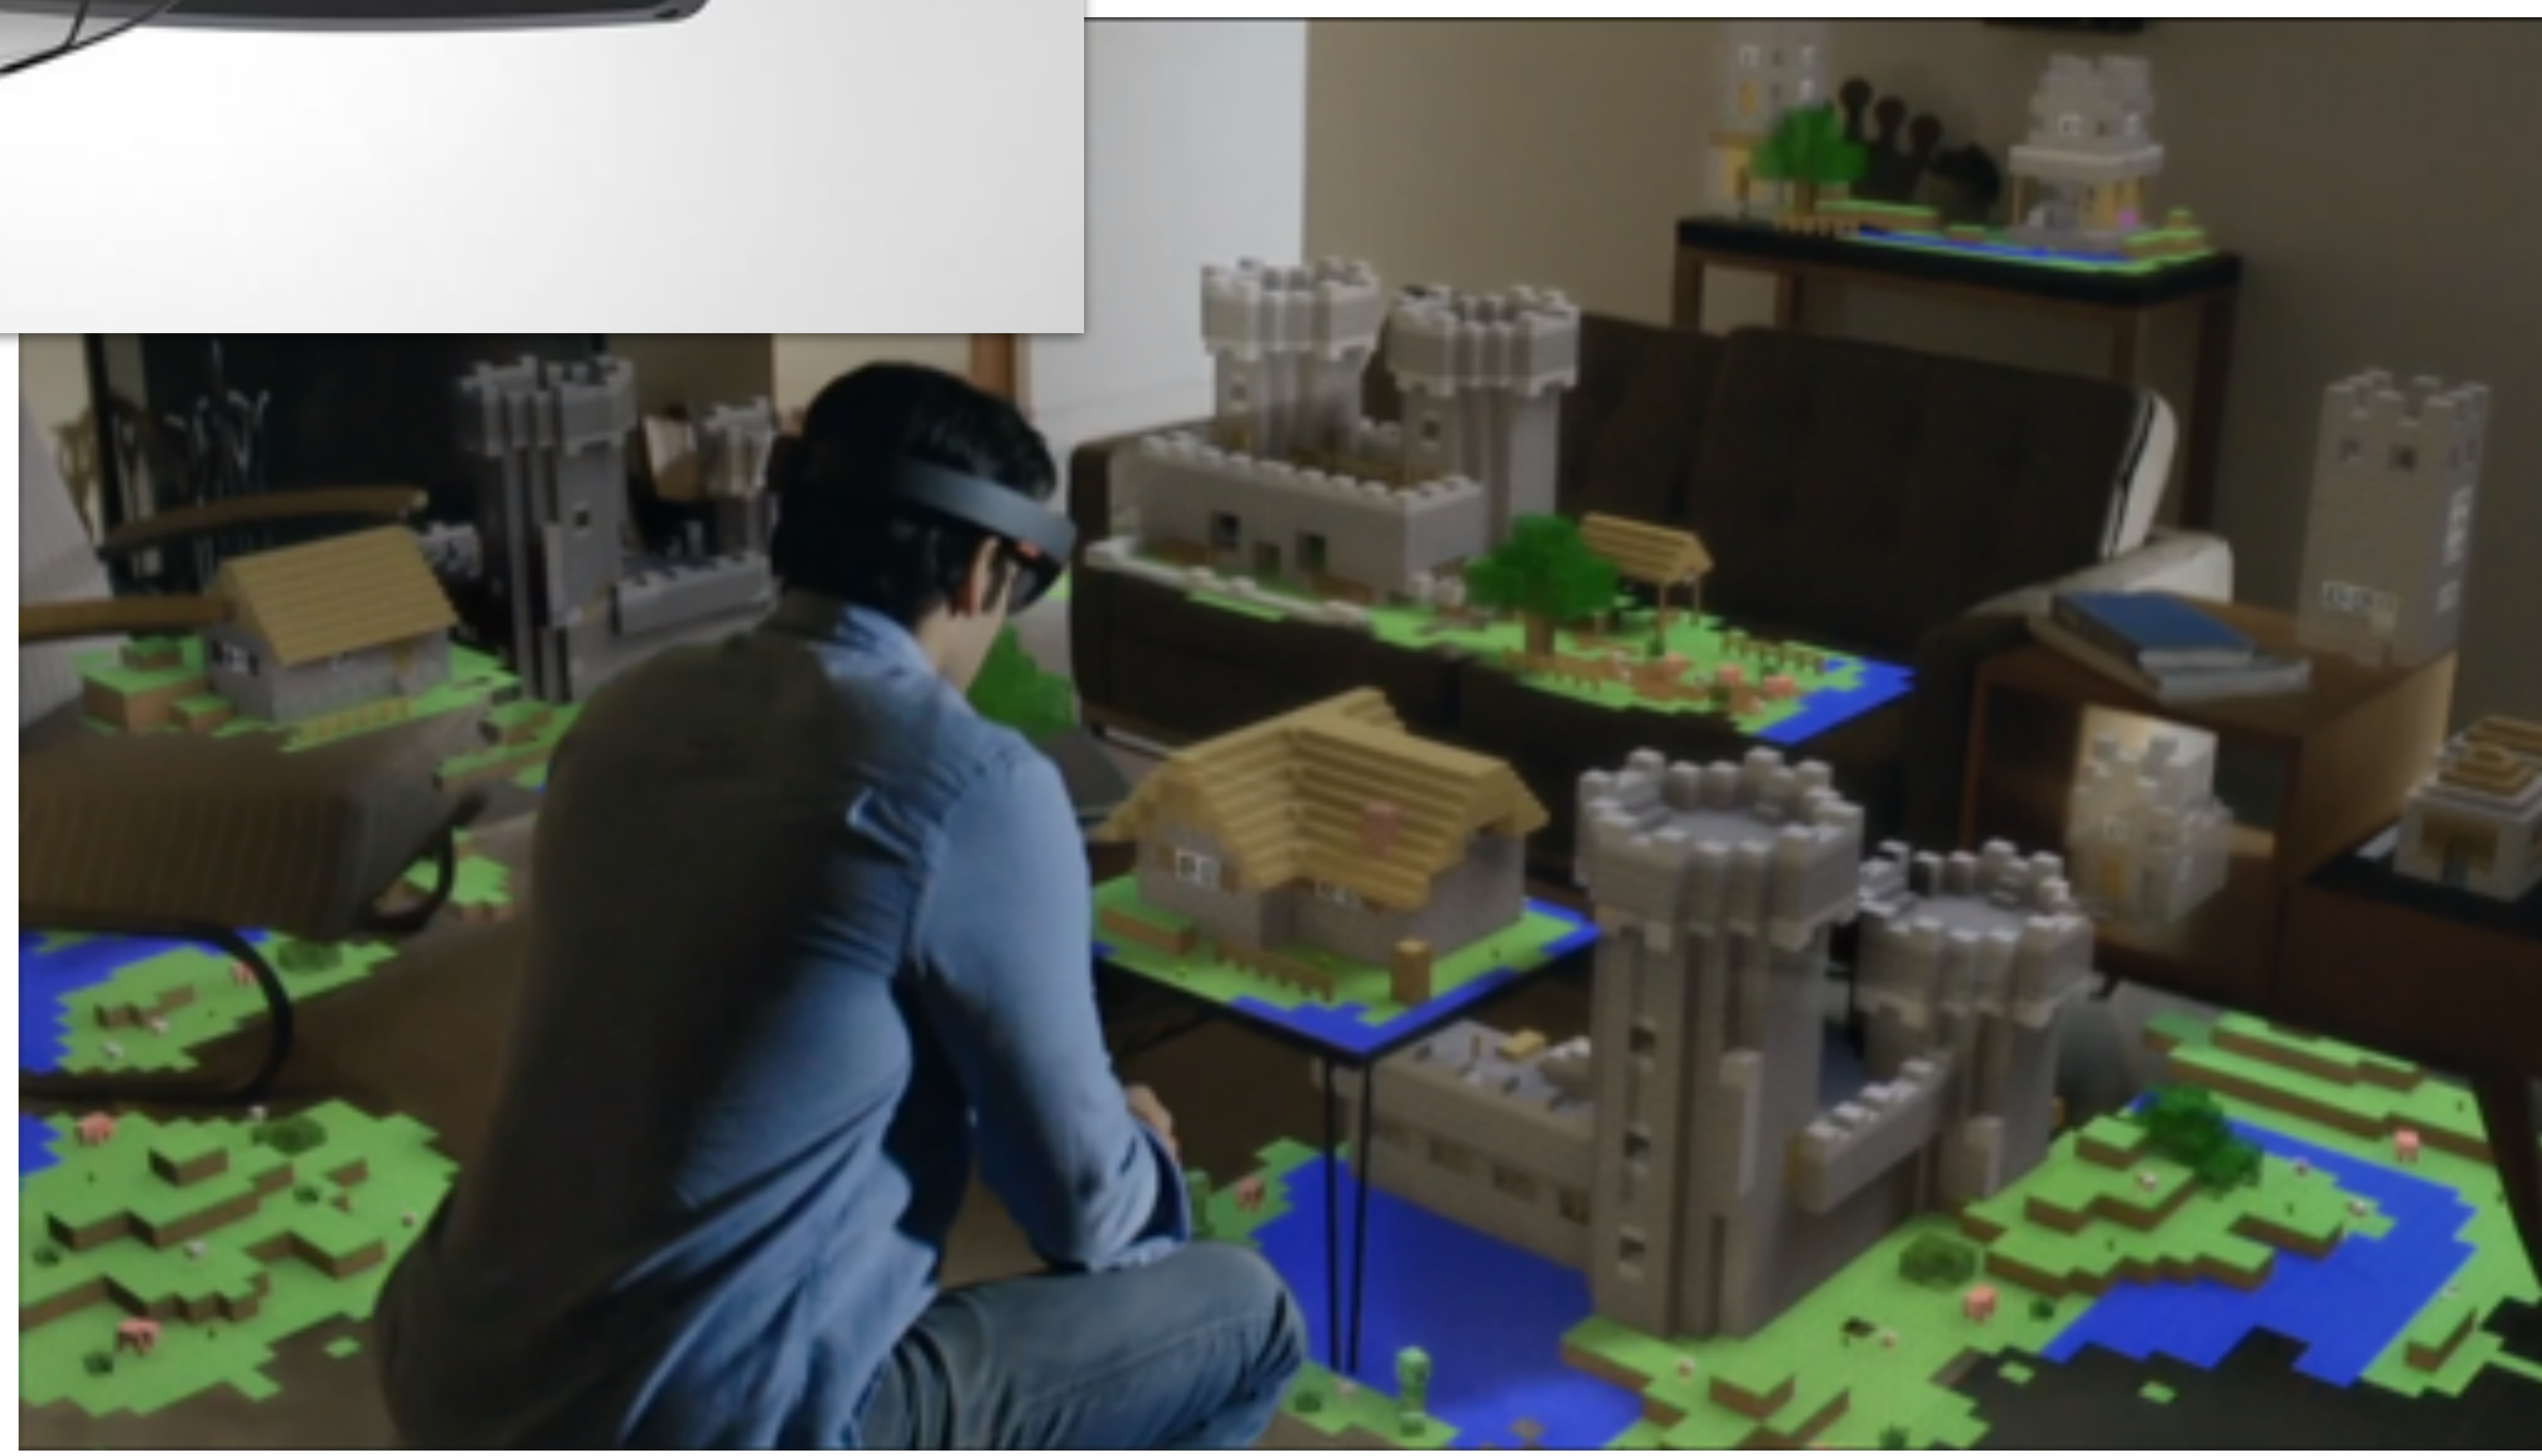

Today: rendering challenges of VR

- Since you are now all experts in rendering, you can appreciate some of the unique challenges of rendering in the context of modern VR headsets

- VR also presents many other difficult technical challenges
  - display technologies
  - accurate tracking of face, head, and body position
  - haptics (simulation of touch)
  - sound synthesis
  - user interface challenges (inability of user to walk around environment, how to manipulate objects in virtual world)
  - content creation challenges
  - and on and on…

VR gaming

Bullet Train Demo (Epic)

VR video

VR teleconference / video chat

http://vrchat.com/

Oculus Rift DK2

Rift DK2 is best documented of modern headsets, so I'll use it for discussion here

Oculus Rift DK2

Oculus Rift DK2 headset

Image credit: ifixit.com

Oculus Rift DK2 display

5.7" 1920 x 1080 OLED display
75 Hz refresh rate

(Same display as Galaxy Note 3)

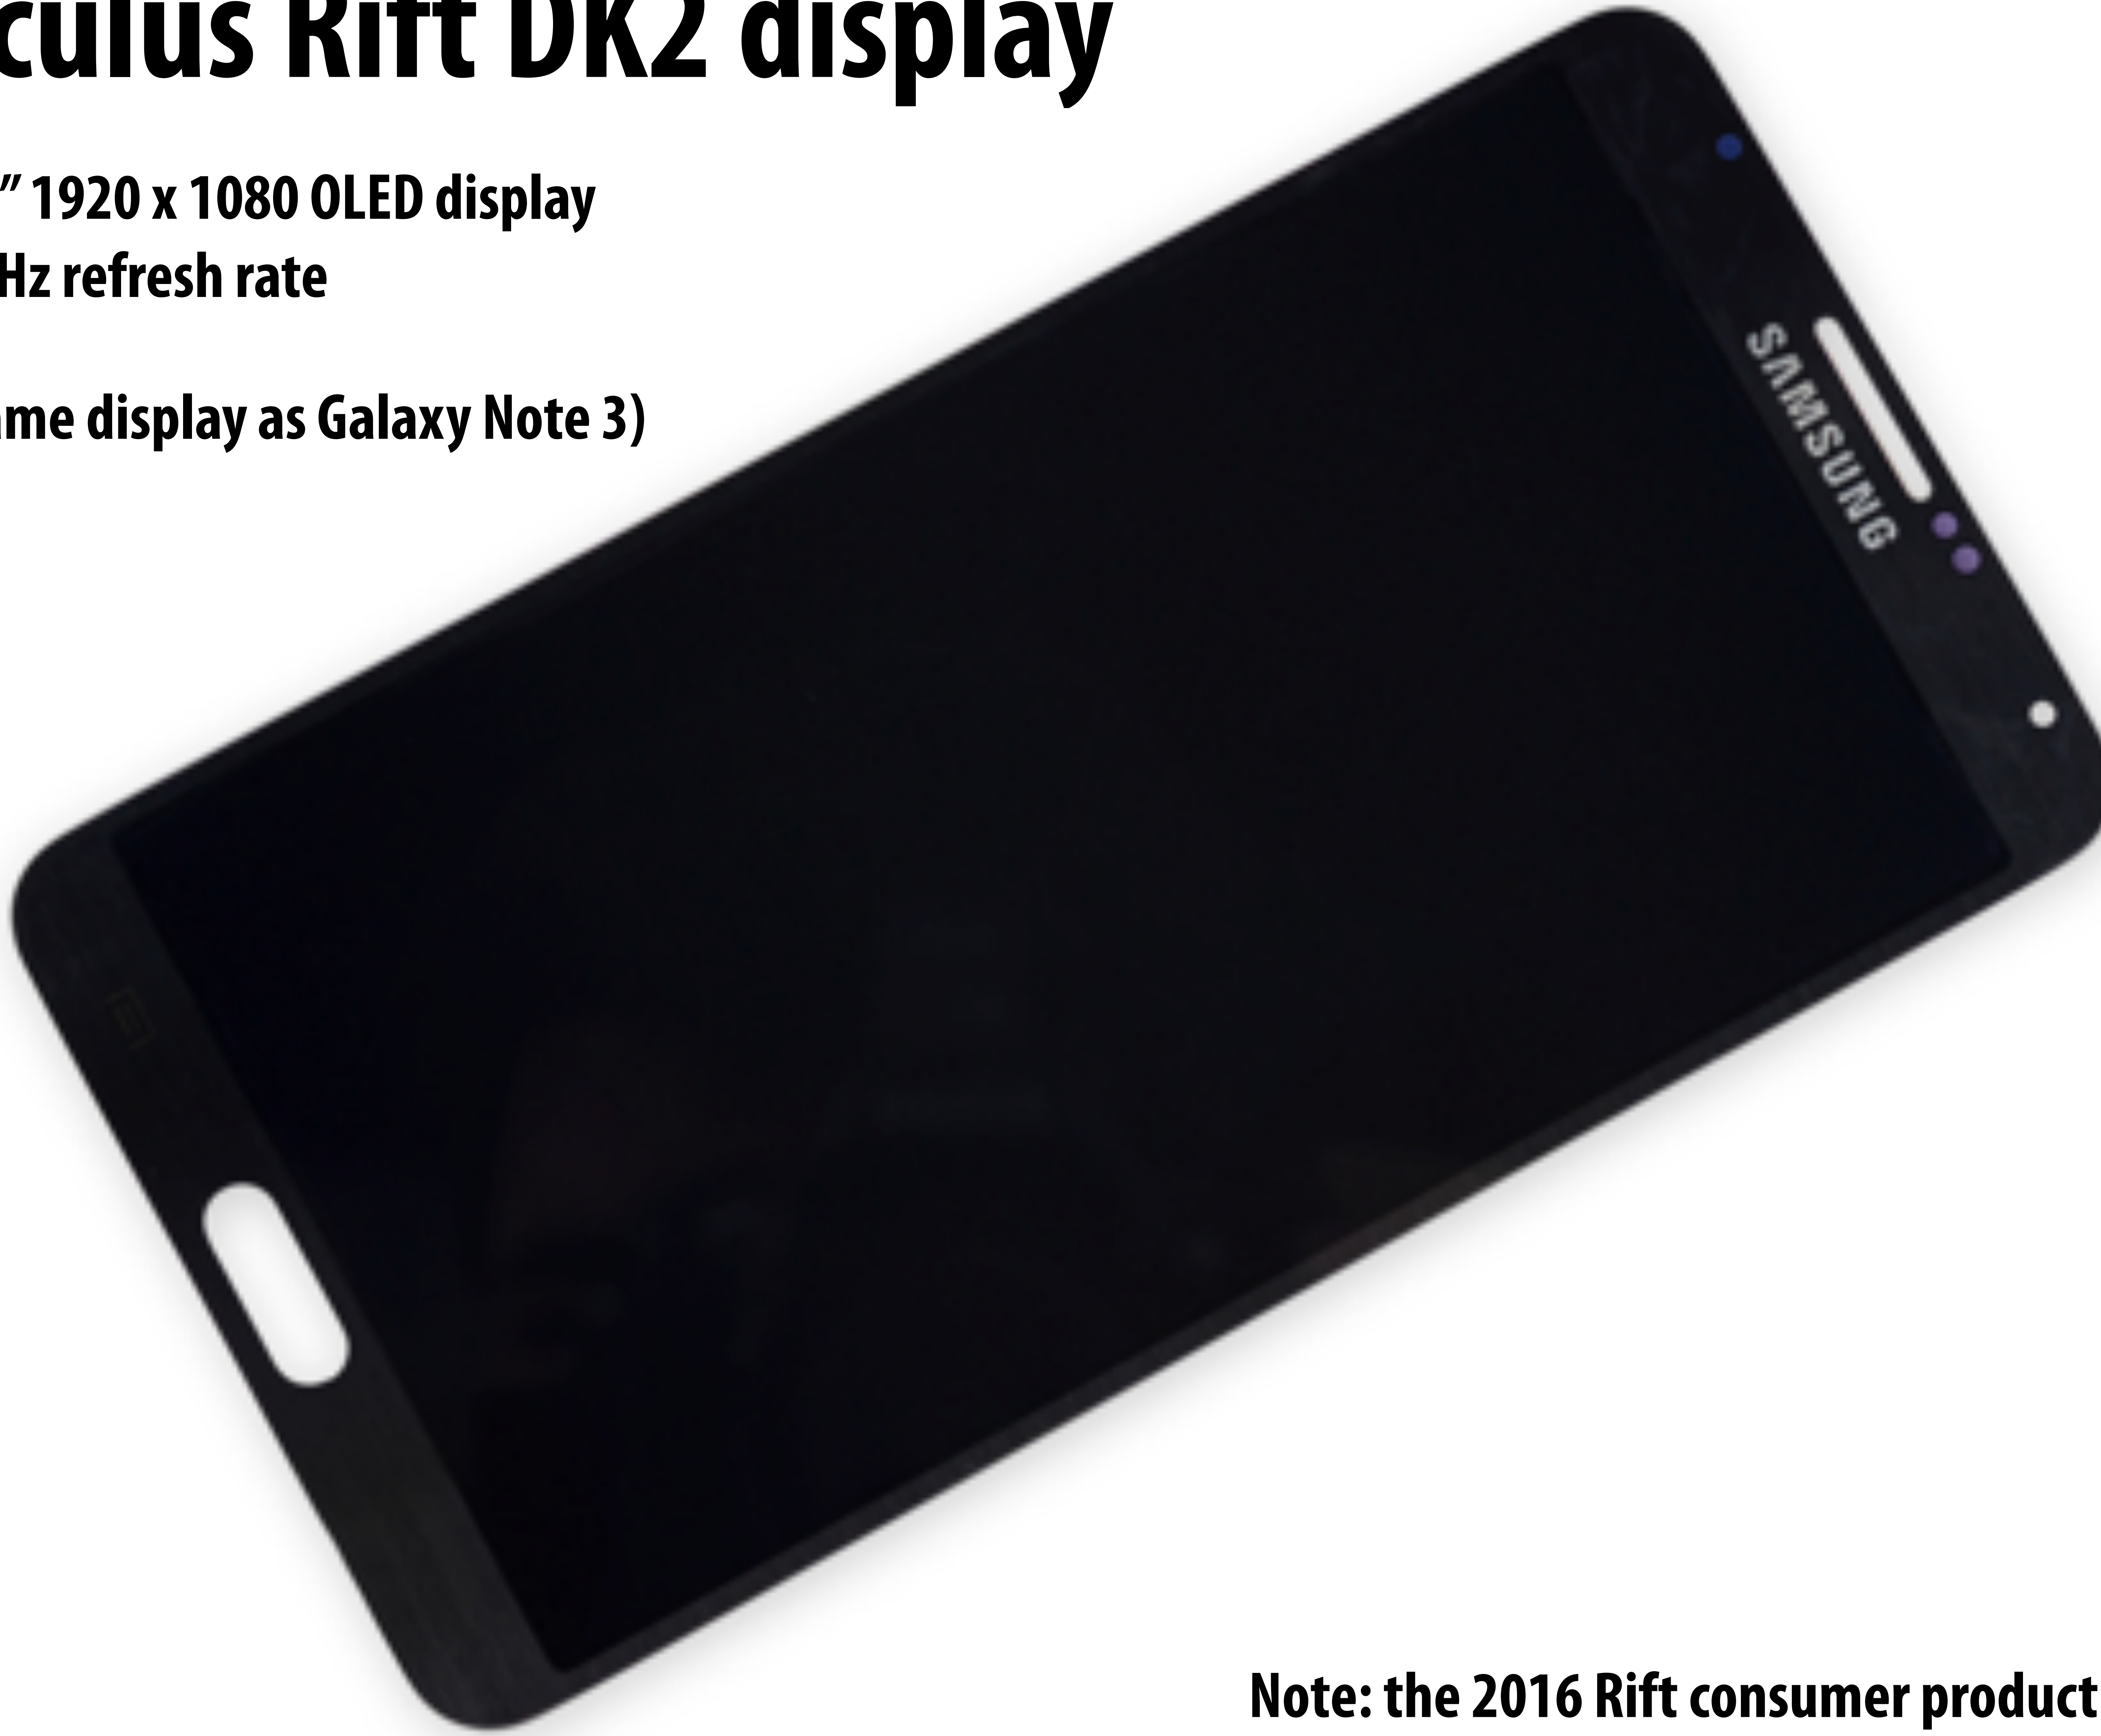

Note: the 2016 Rift consumer product features two 1080×1200 displays at 90Hz.

Image credit: ifixit.com

Oculus Rift DK2 headset

Image credit: ifixit.com

Requirement: wide field of view

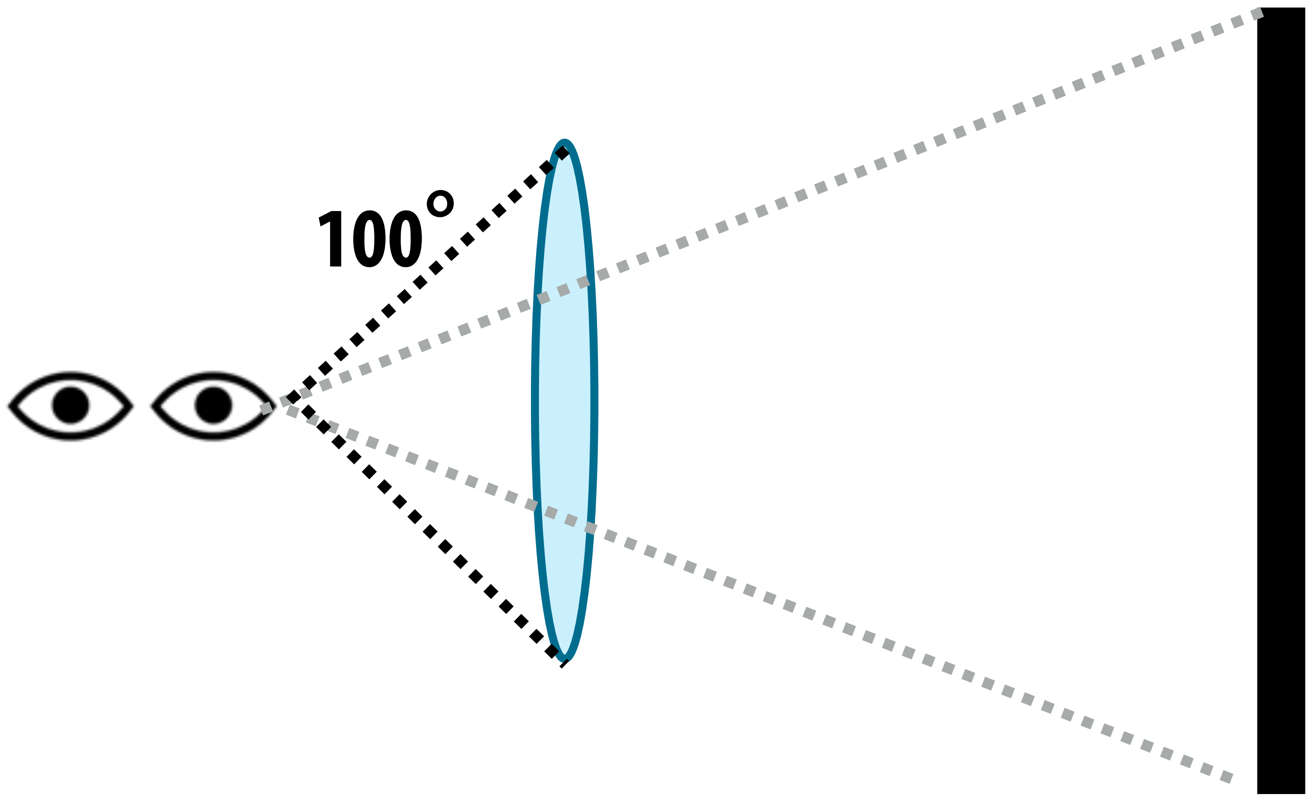

View of checkerboard through Oculus Rift lens

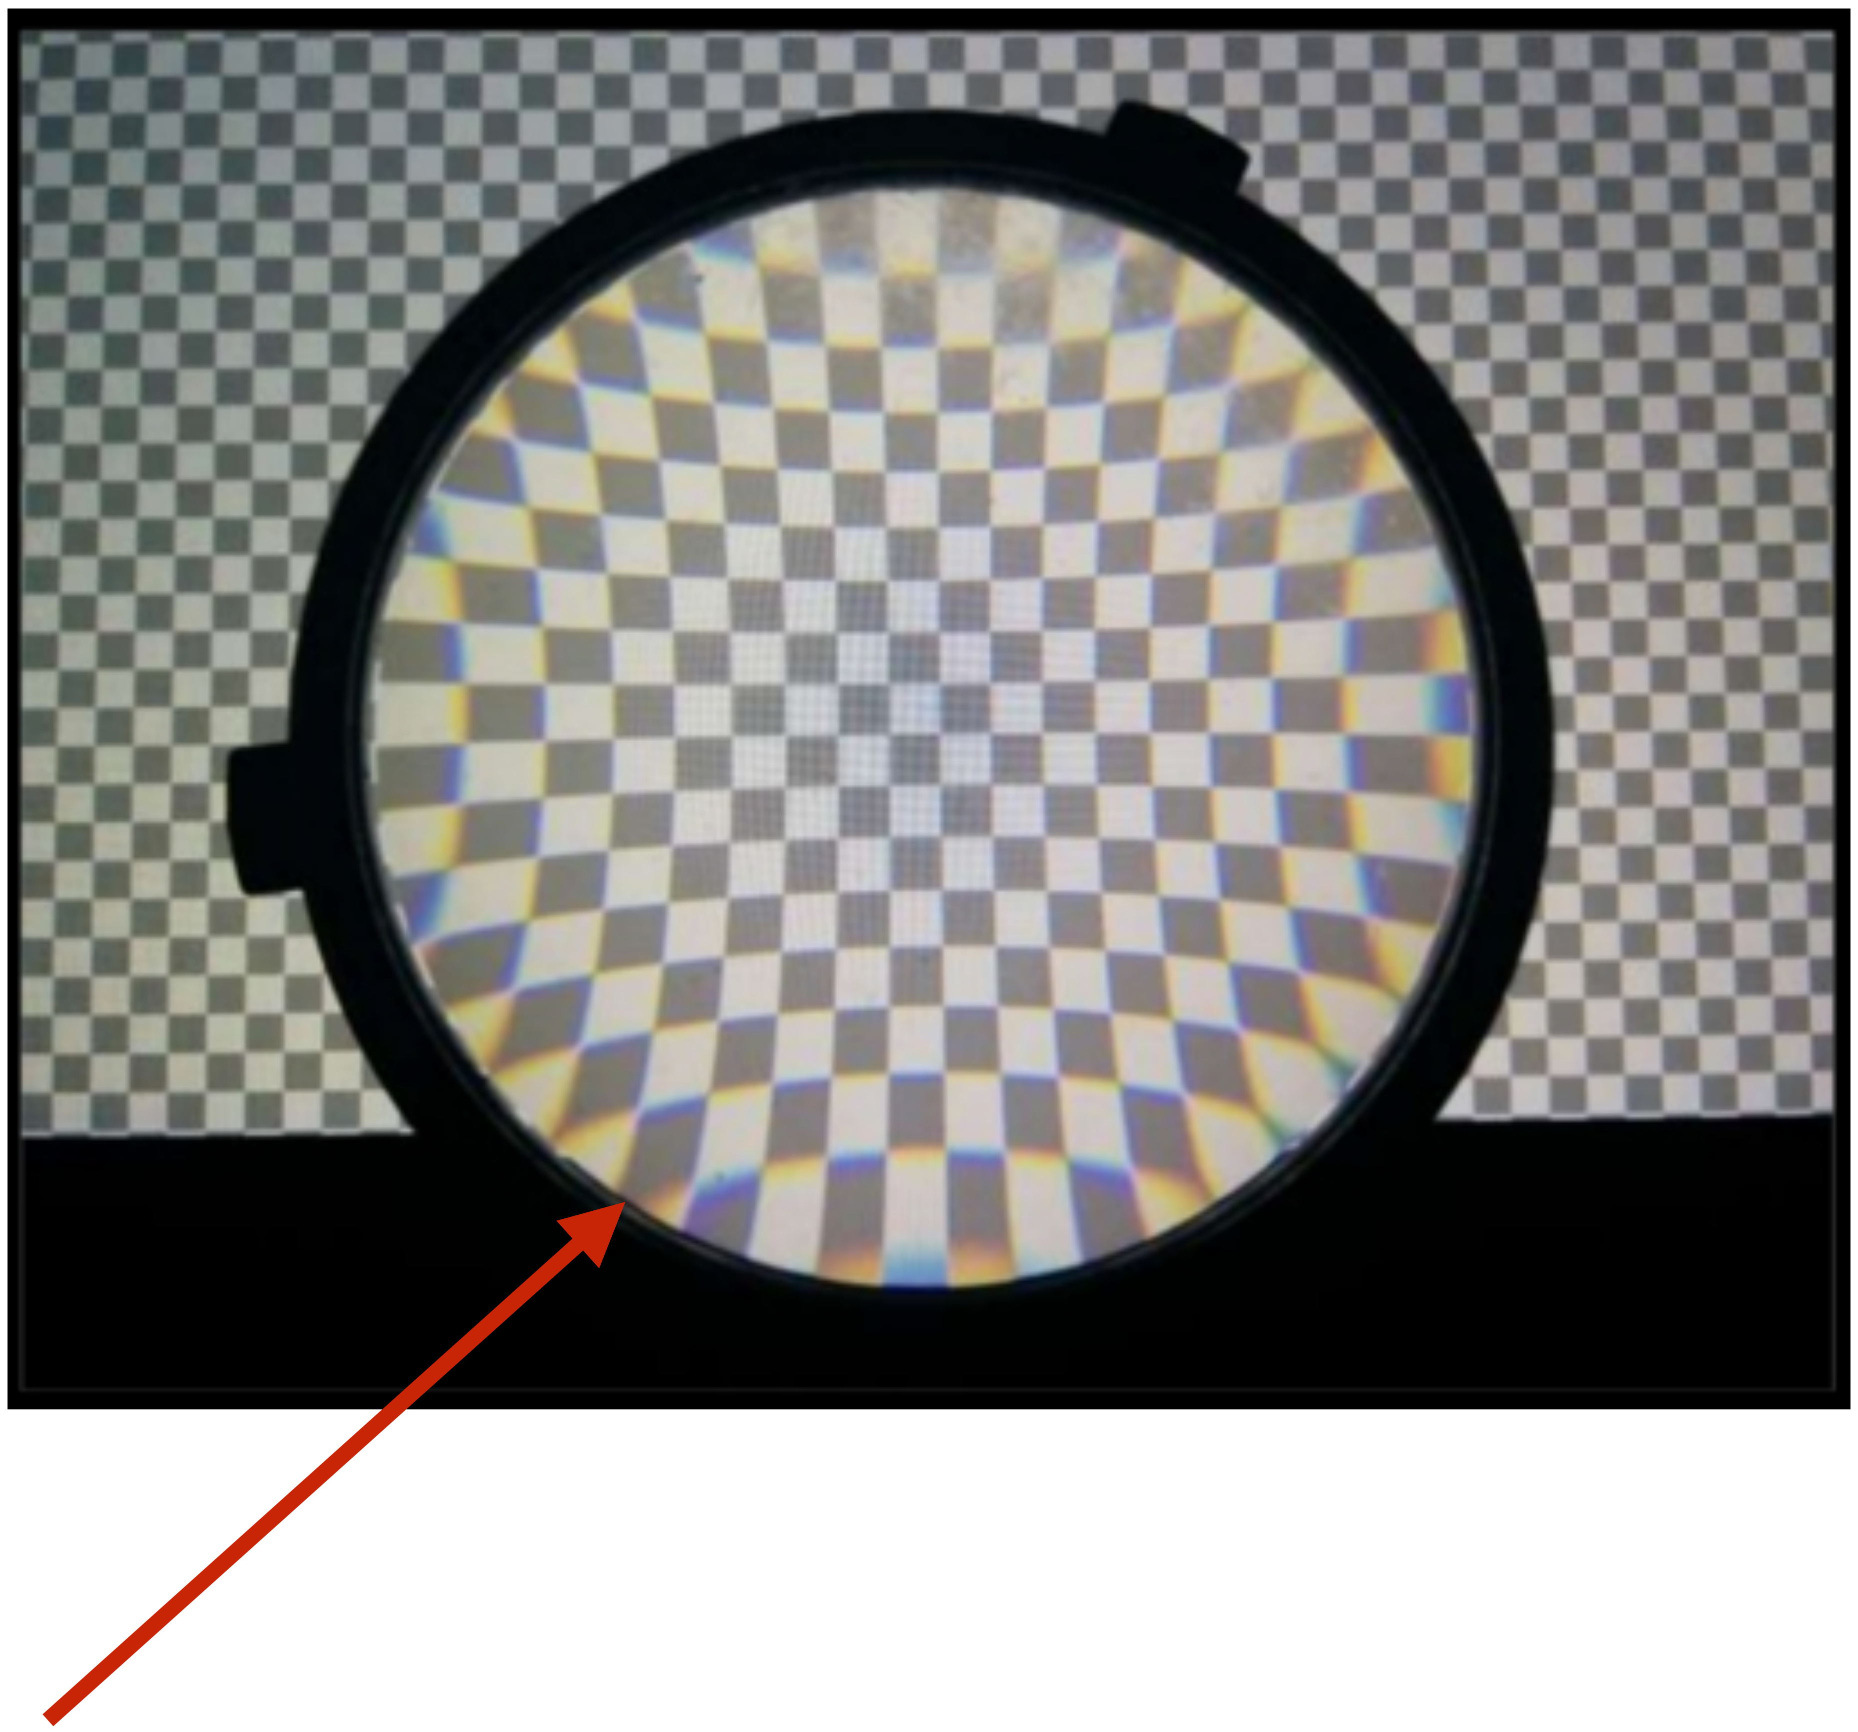

Lens introduces distortion

- Pincushion distortion
- Chromatic aberration (different wavelengths of light refract by different amount)

Icon credit: Eyes designed by SuperAtic LABS from the thenounproject.com
Image credit: Cass Everitt

Role of optics

1. Create wide field of view
2. Place focal plane at several meters away from eye (close to infinity)

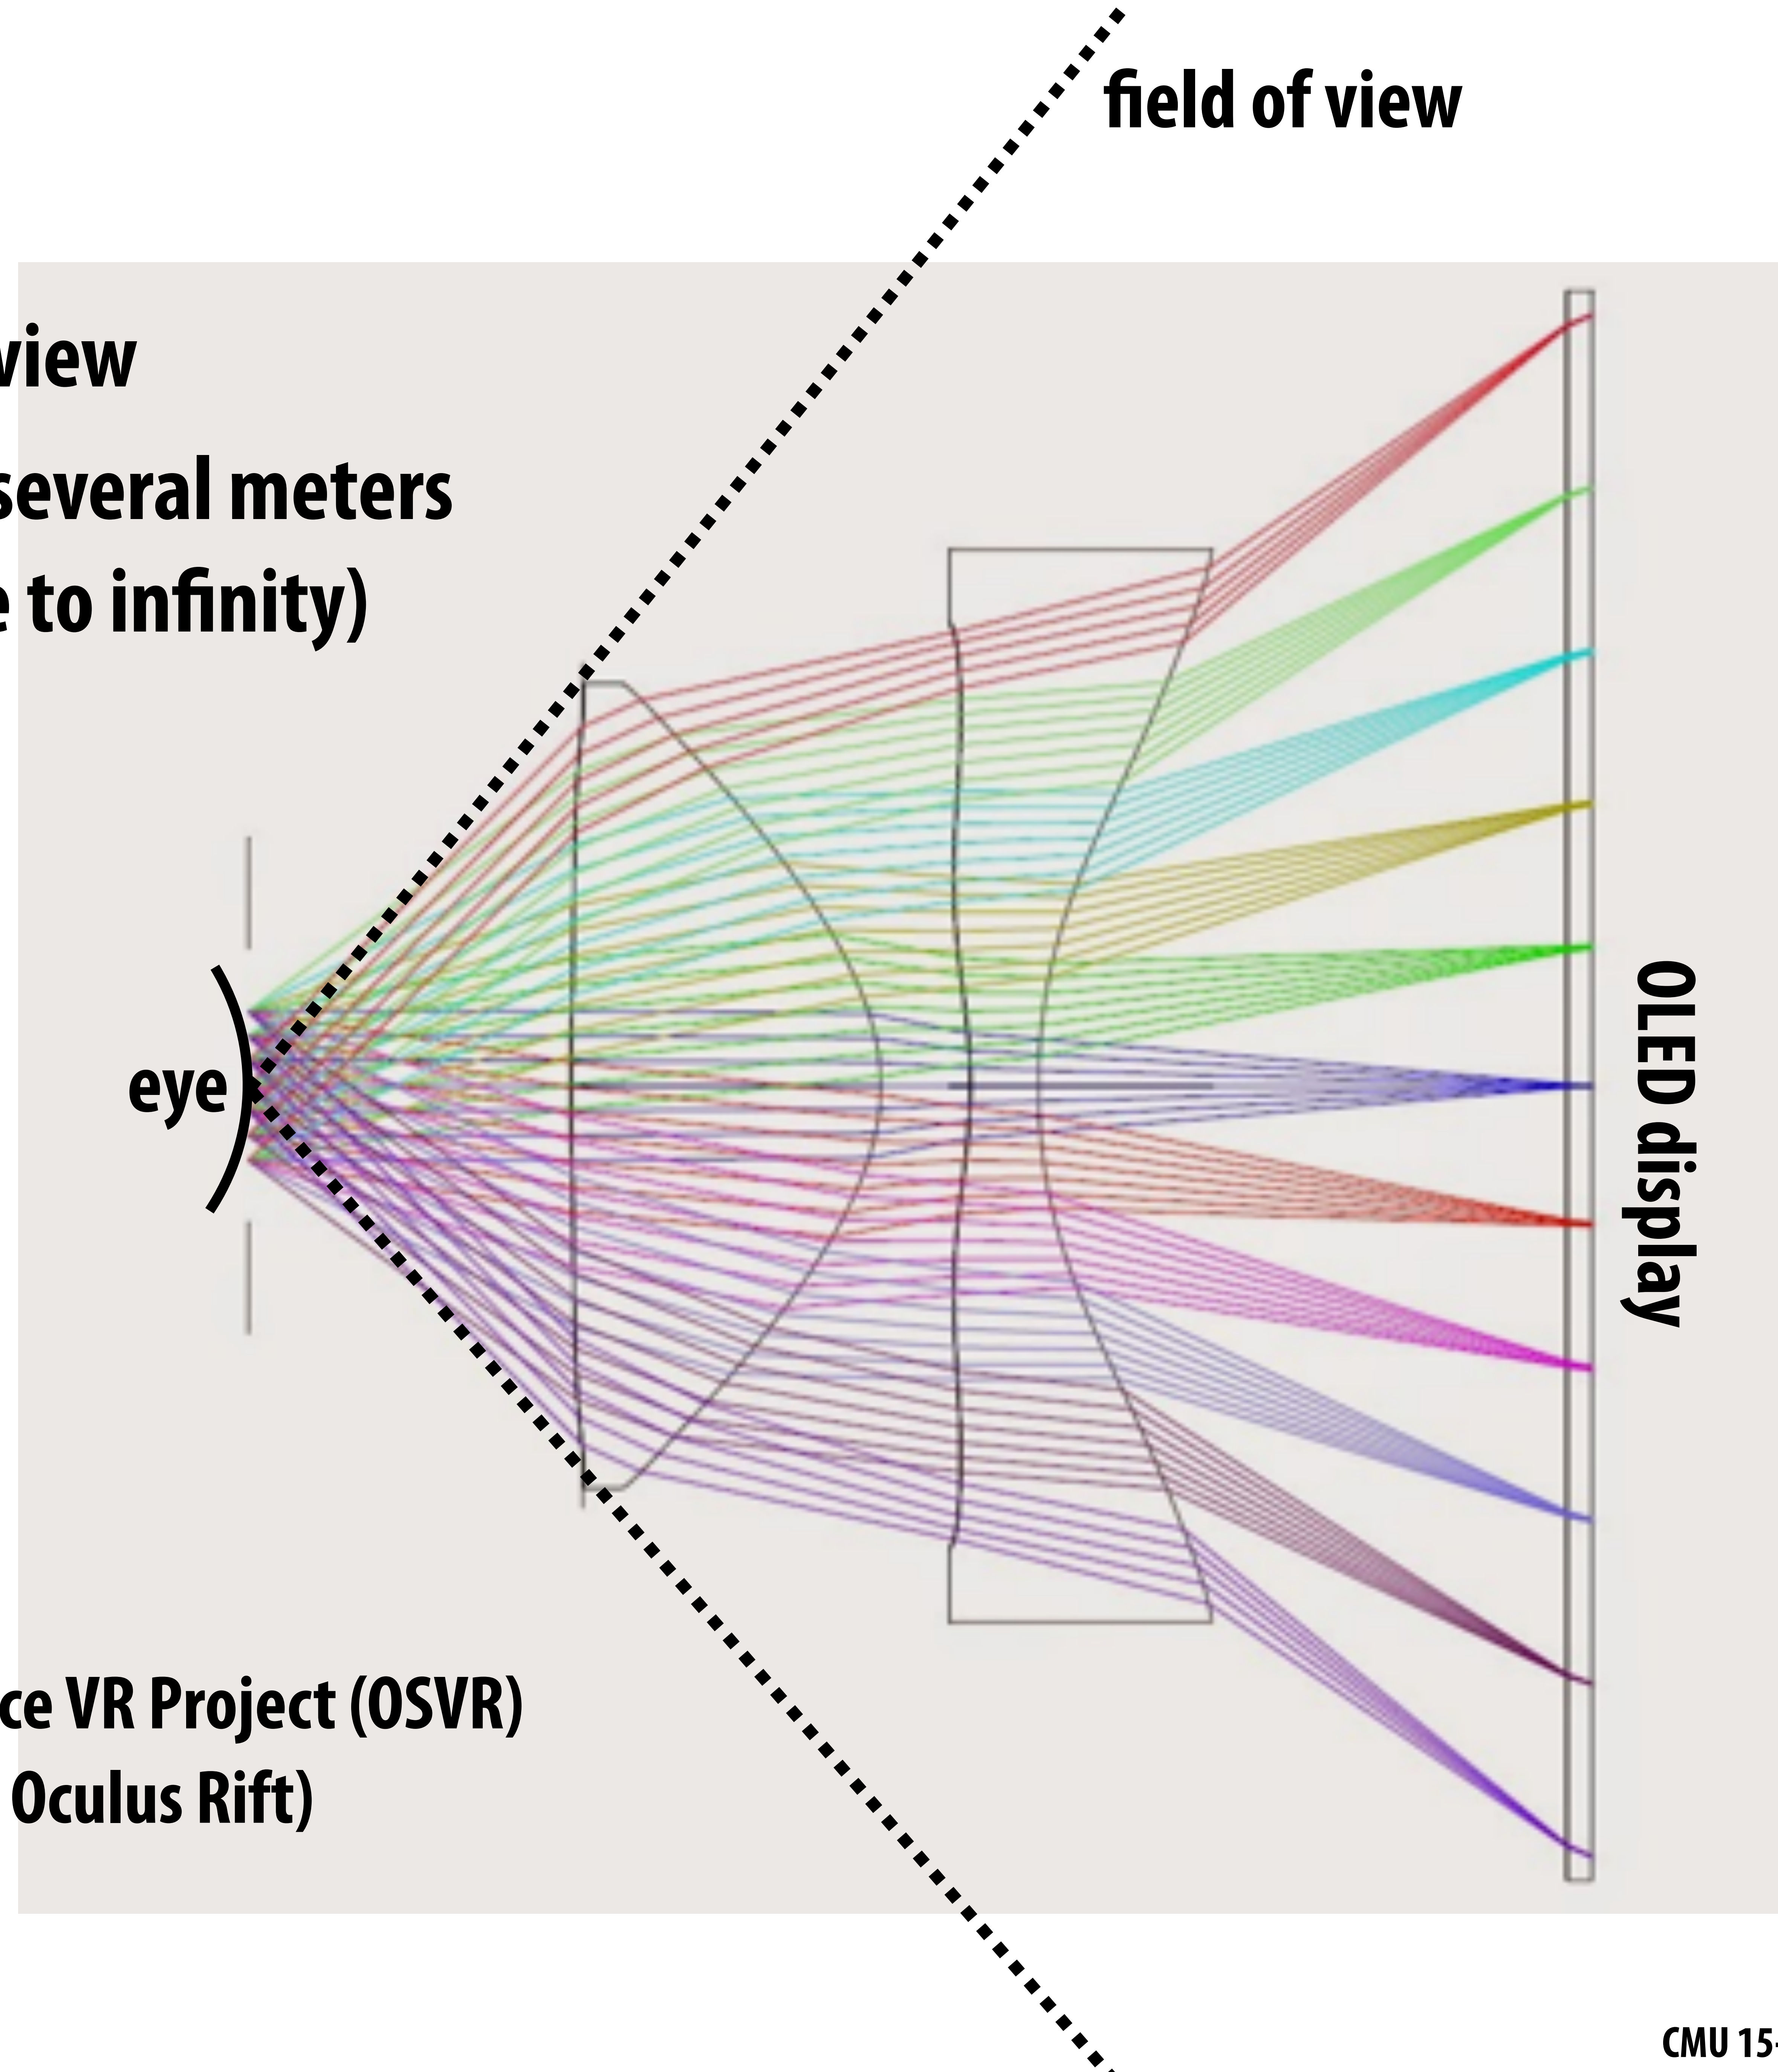

Lens diagram from Open Source VR Project (OSVR)
(Not the lens system from the Oculus Rift)
http://www.osvr.org/

Rendered output must compensate for distortion of lens in front of display

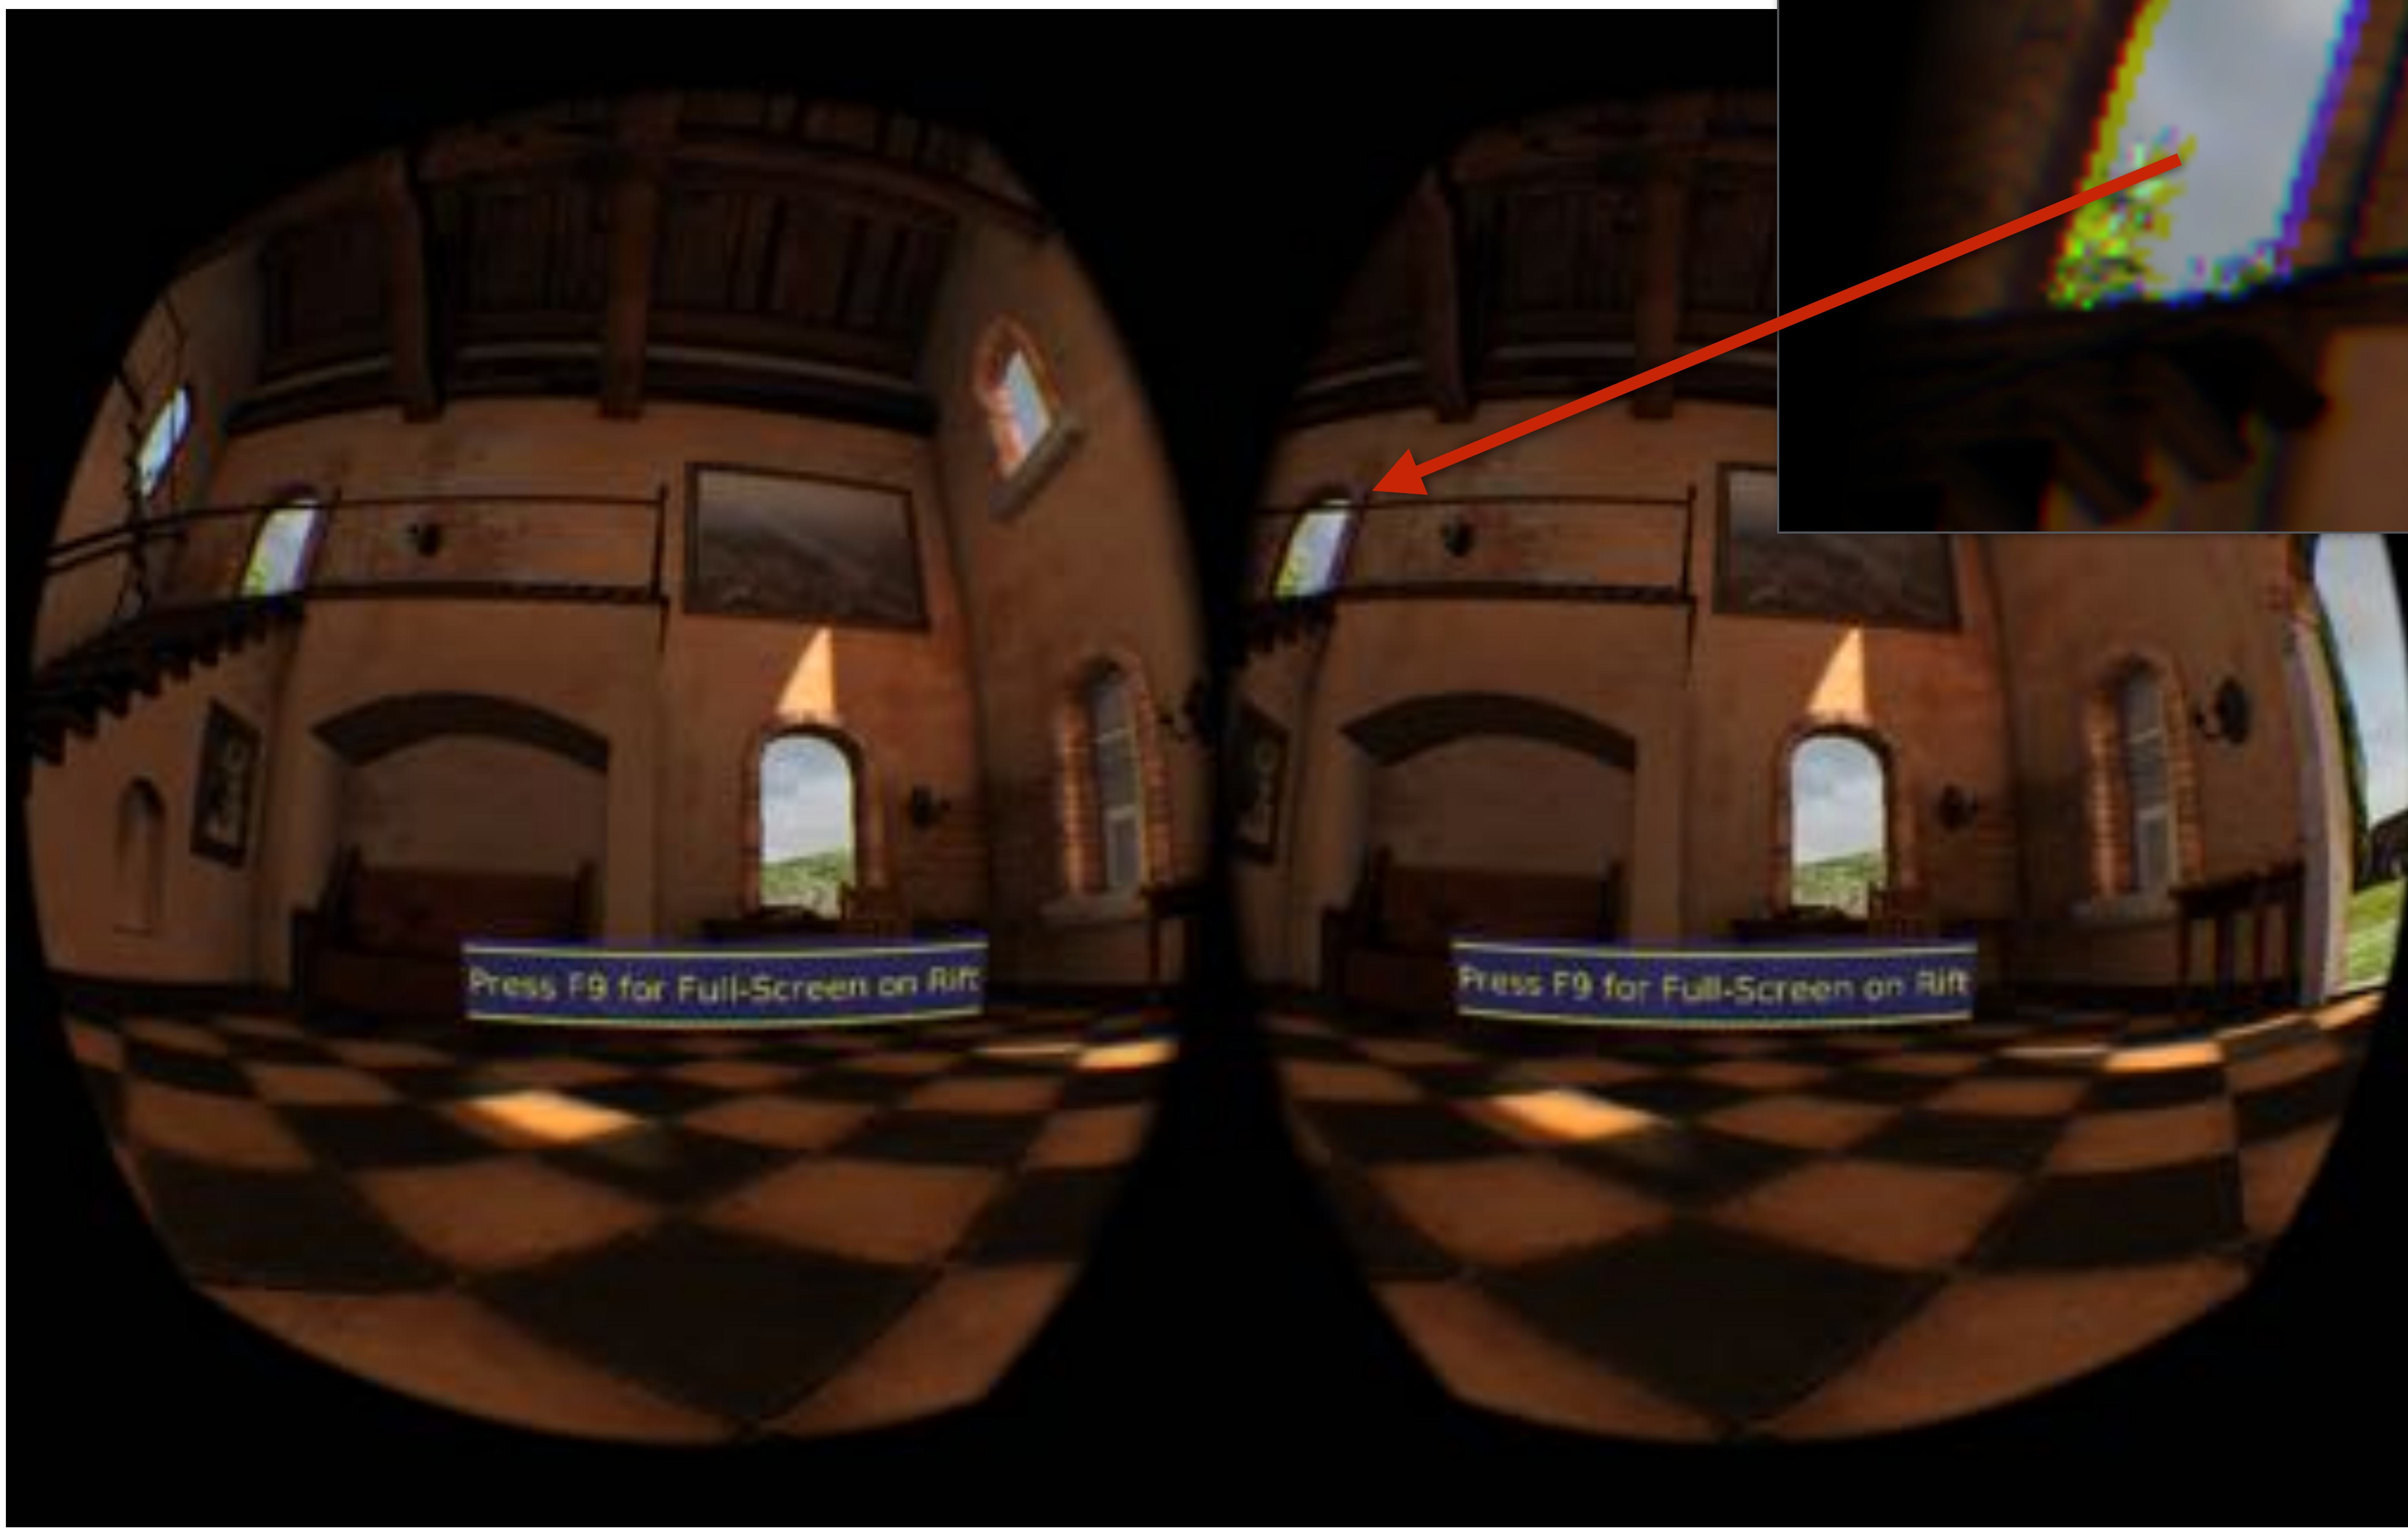

Step 1: render scene using traditional graphics pipeline at full resolution for each eye

Step 2: warp images and composite into frame rendering is viewed correctly after lens distortion

(Can apply unique distortion to R, G, B to approximate correction for chromatic aberration)

Image credit: Oculus VR developer guide

Still have vergence conflict…

- **Given design of current VR displays, consider what happens when objects are up-close to eye in virtual scene**
  - Eyes must remain accommodated to near infinity (otherwise image on screen won't be in focus)
  - But eyes must converge in attempt to fuse stereoscopic images of object up close
  - Brain receives conflicting depth clues… (discomfort, fatigue, nausea)

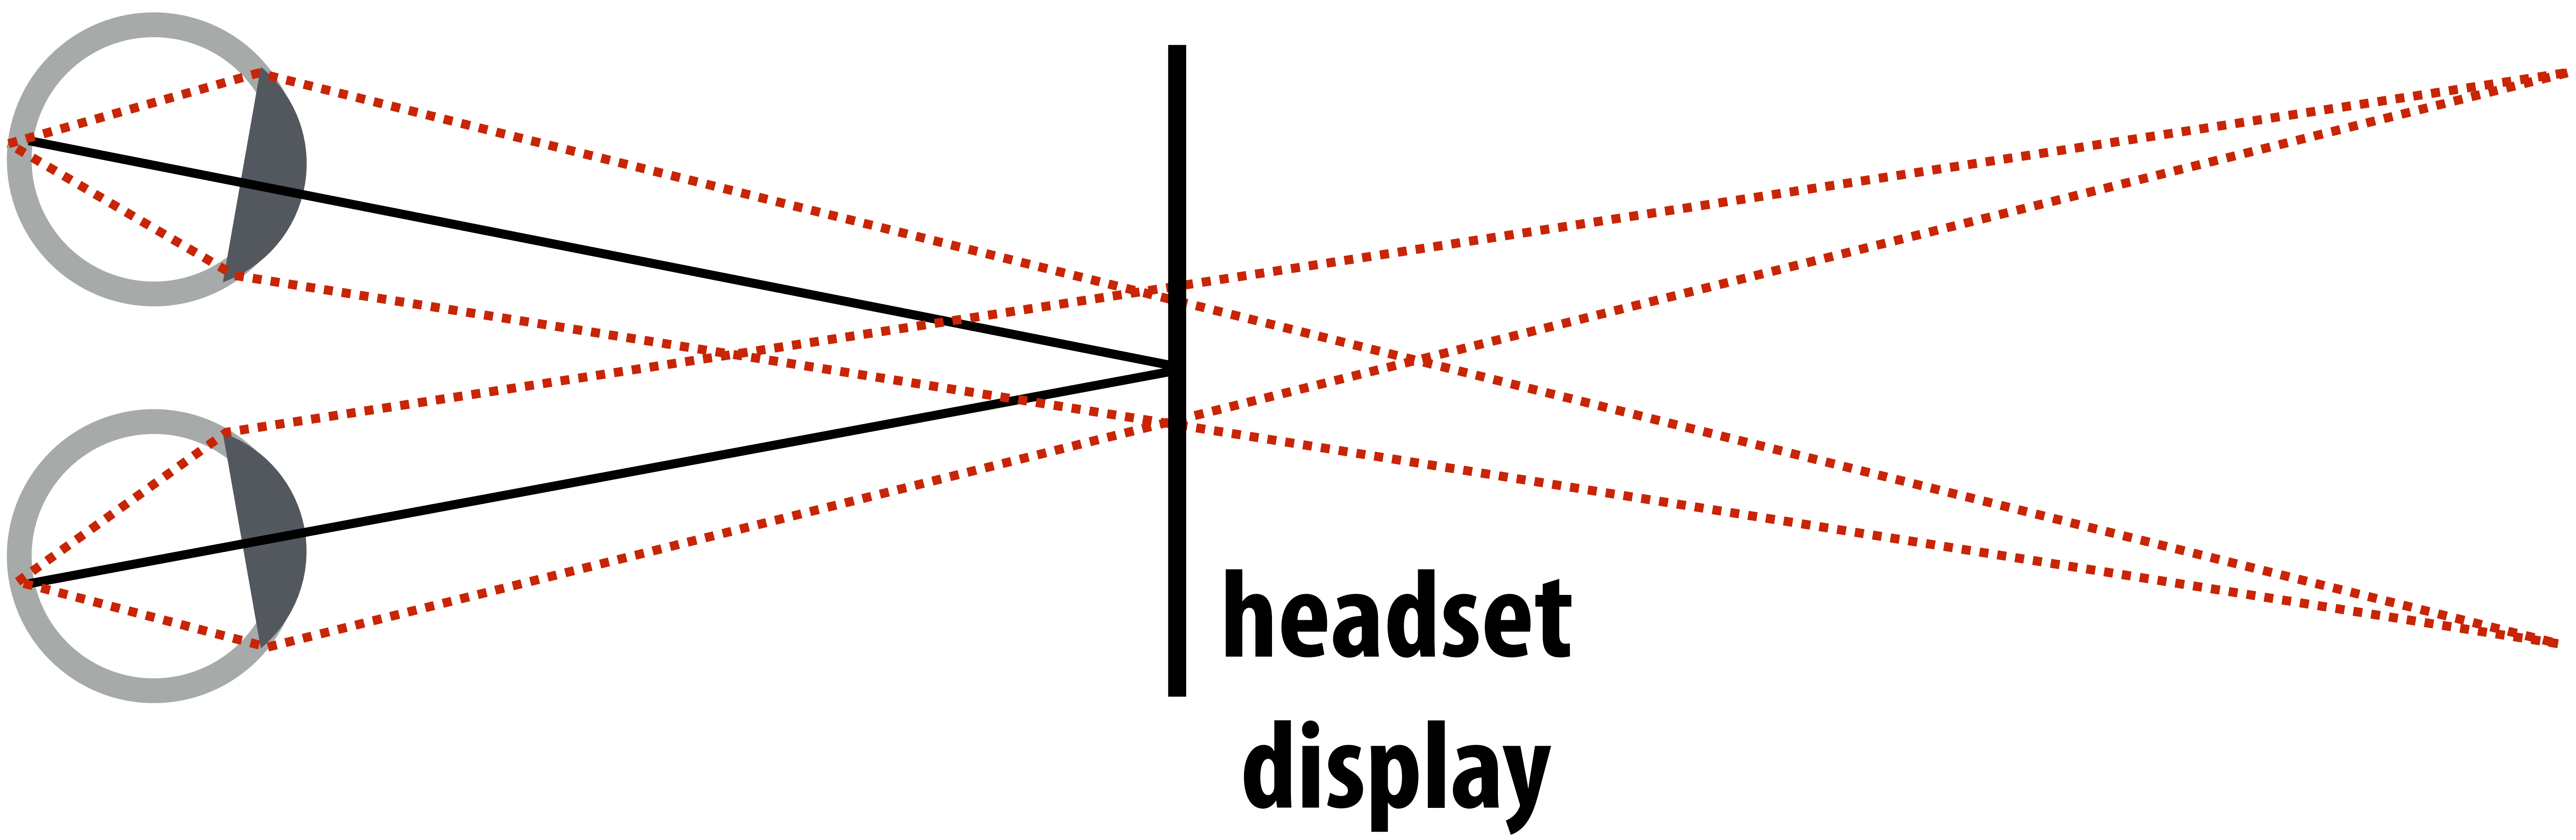

Have to be cross-eyed!

Aside: near-eye light field displays

If only you could recreate the true light field, you wouldn't have this problem…

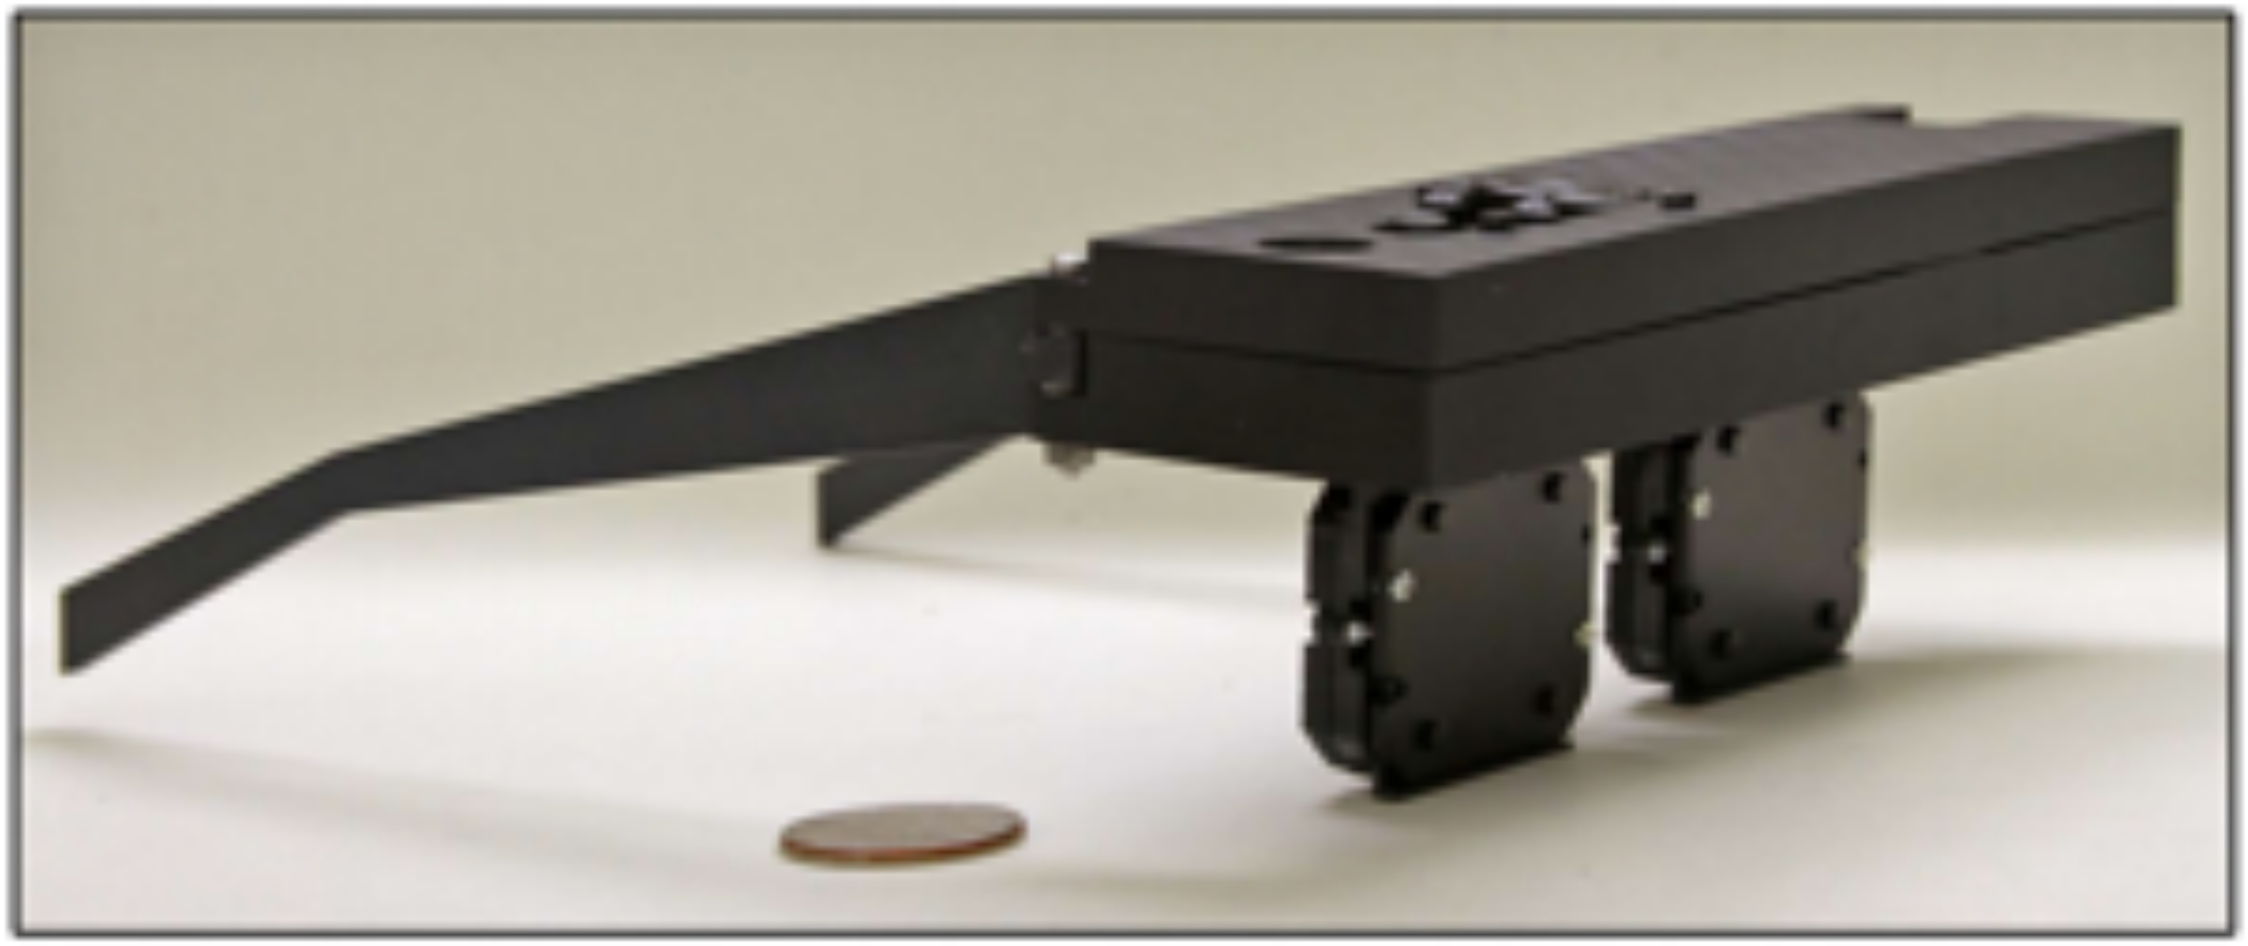

Challenge: rendering via planar projection

Recall: rasterization-based graphics is based on perspective projection to plane

- Reasonable for modest FOV, but distorts image under high FOV
- Recall: VR rendering spans wide FOV

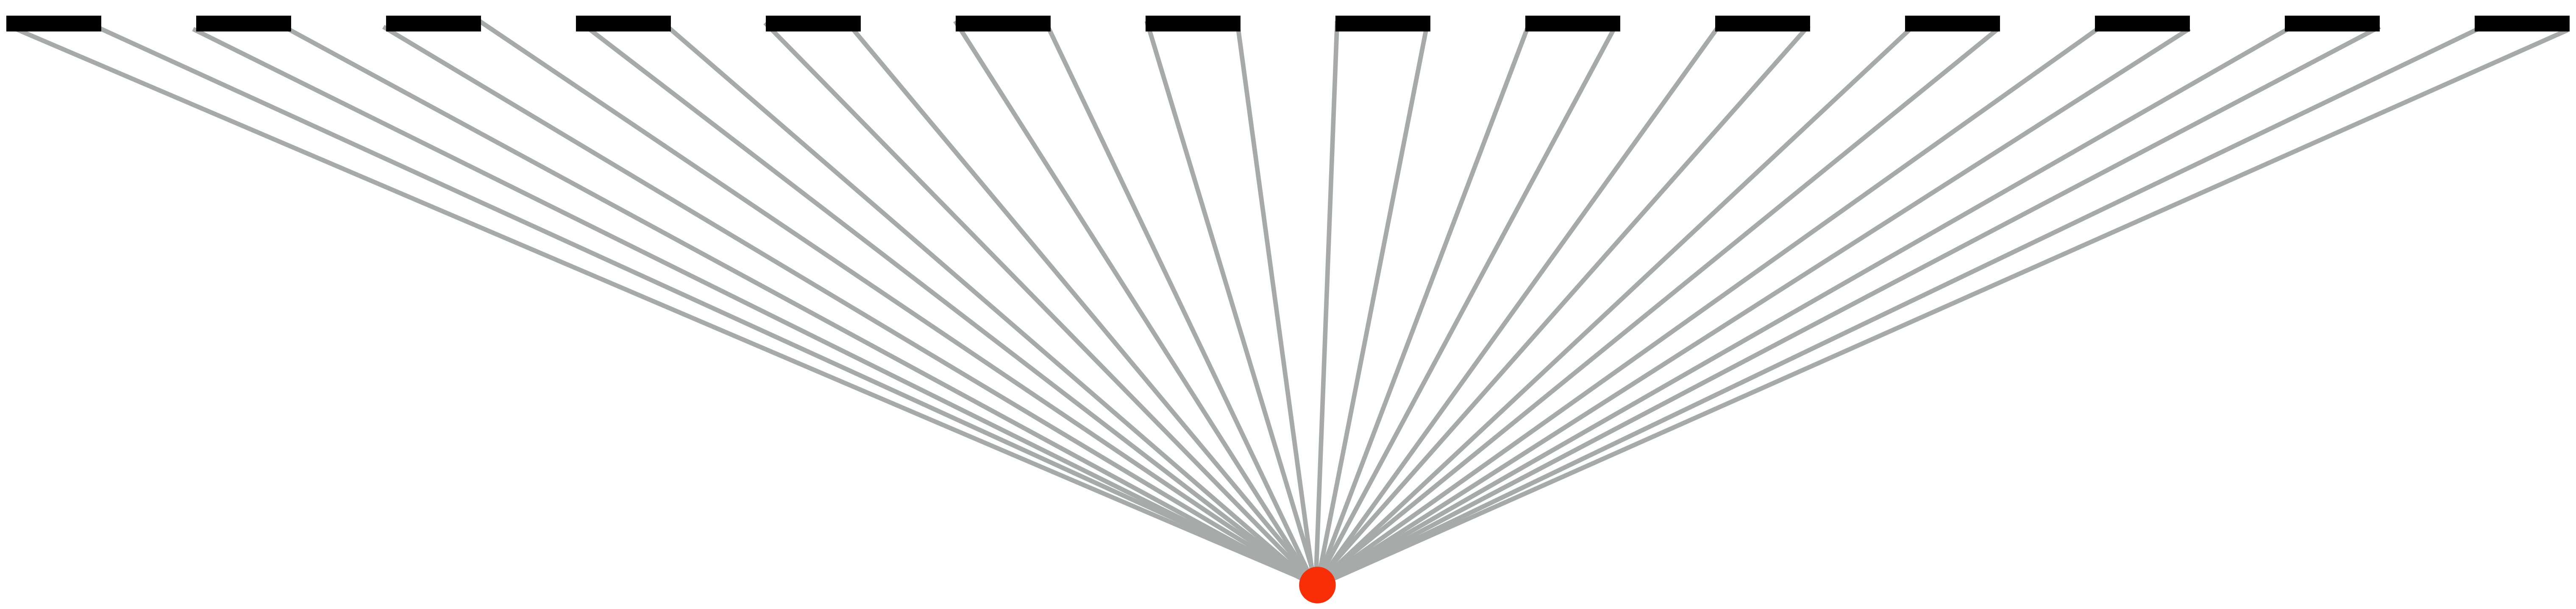

Pixels span larger angle in center of image
(lowest angular resolution in center)

Future investigations may consider: curved displays, ray casting to achieve uniform angular resolution, rendering with piecewise linear projection plane (different plane per tile of screen)

Image credit: Cass Everitt

Foveated rendering

Idea: track user's gaze, render with increasingly lower resolution farther away from gaze point

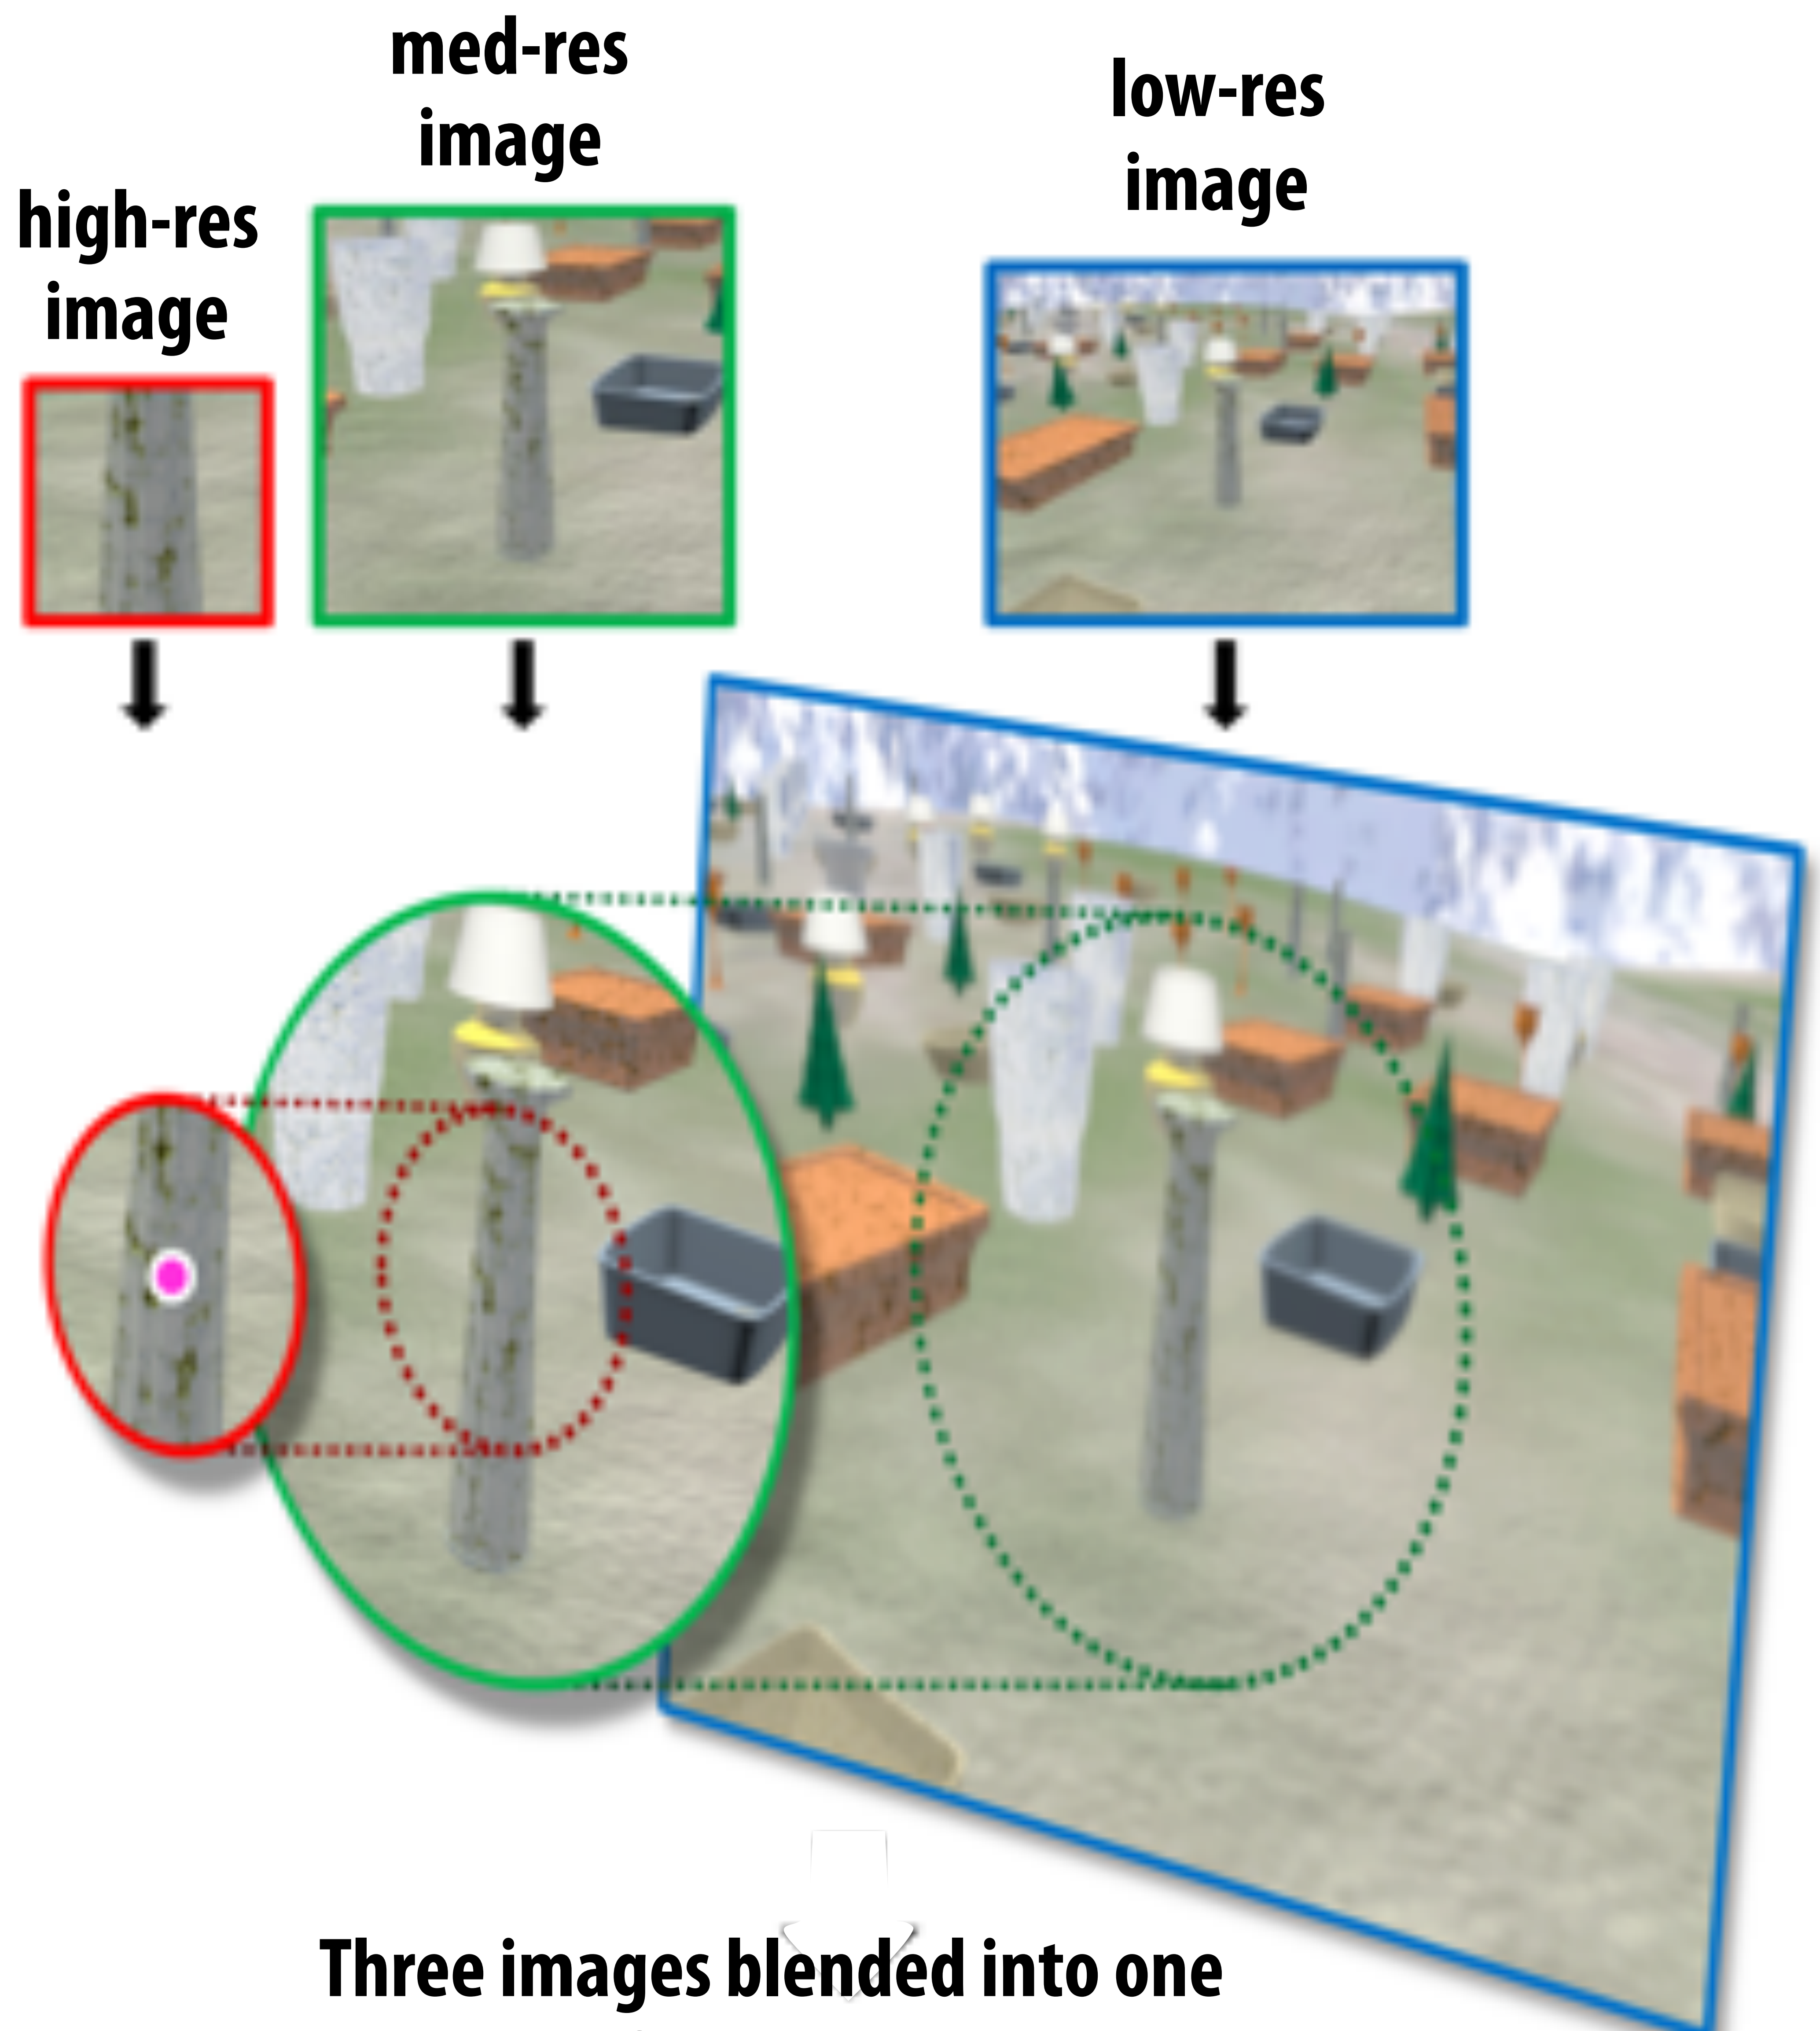

Three images blended into one for display

More recent/near future VR system

Low-latency image processing for subject tracking

High-resolution, high-frame rate, wide-field of view display

Massive parallel computation for high-resolution rendering

In headset motion/accel sensors + eye tracker

Exceptionally high bandwidth connection between renderer and display: e.g., 4K x 4K per eye at 90 fps!

On headset graphics processor for sensor processing and re-projection

Summary: virtual reality presents many new challenges for graphics systems developers

- **Major goal: minimize latency of head movement to photons**
  - Requires low latency tracking (not discussed today)
    - Combination of external camera image processing (vision) and high rate headset sensors
    - Heavy use of prediction
  - Requires high-performance rendering
    - High-resolution, wide field-of-view output
    - High frame-rate
    - Rendering must compensate for constraints of display system:
      - Optical distortion (geometric, chromatic)
      - Temporal offsets in rows of pixels
- **Significant research interest in display technologies that are alternatives to flat screens with lenses in front of them**

Interest in acquiring VR content

Google's JumpVR video:
16 4K GoPro cameras

Consider challenge of:
Registering/3D align video stream (on site)
Broadcast encoded video stream across the country to 50 million viewers

Lytro Immerge
(leveraging light field camera technology to acquire VR content)

Stanford Camera Array

Wilburn et al. 2005

640 x 480 tightly synchronized, repositionable cameras

Custom processing board per camera

Tethered to PCs for additional processing/storage

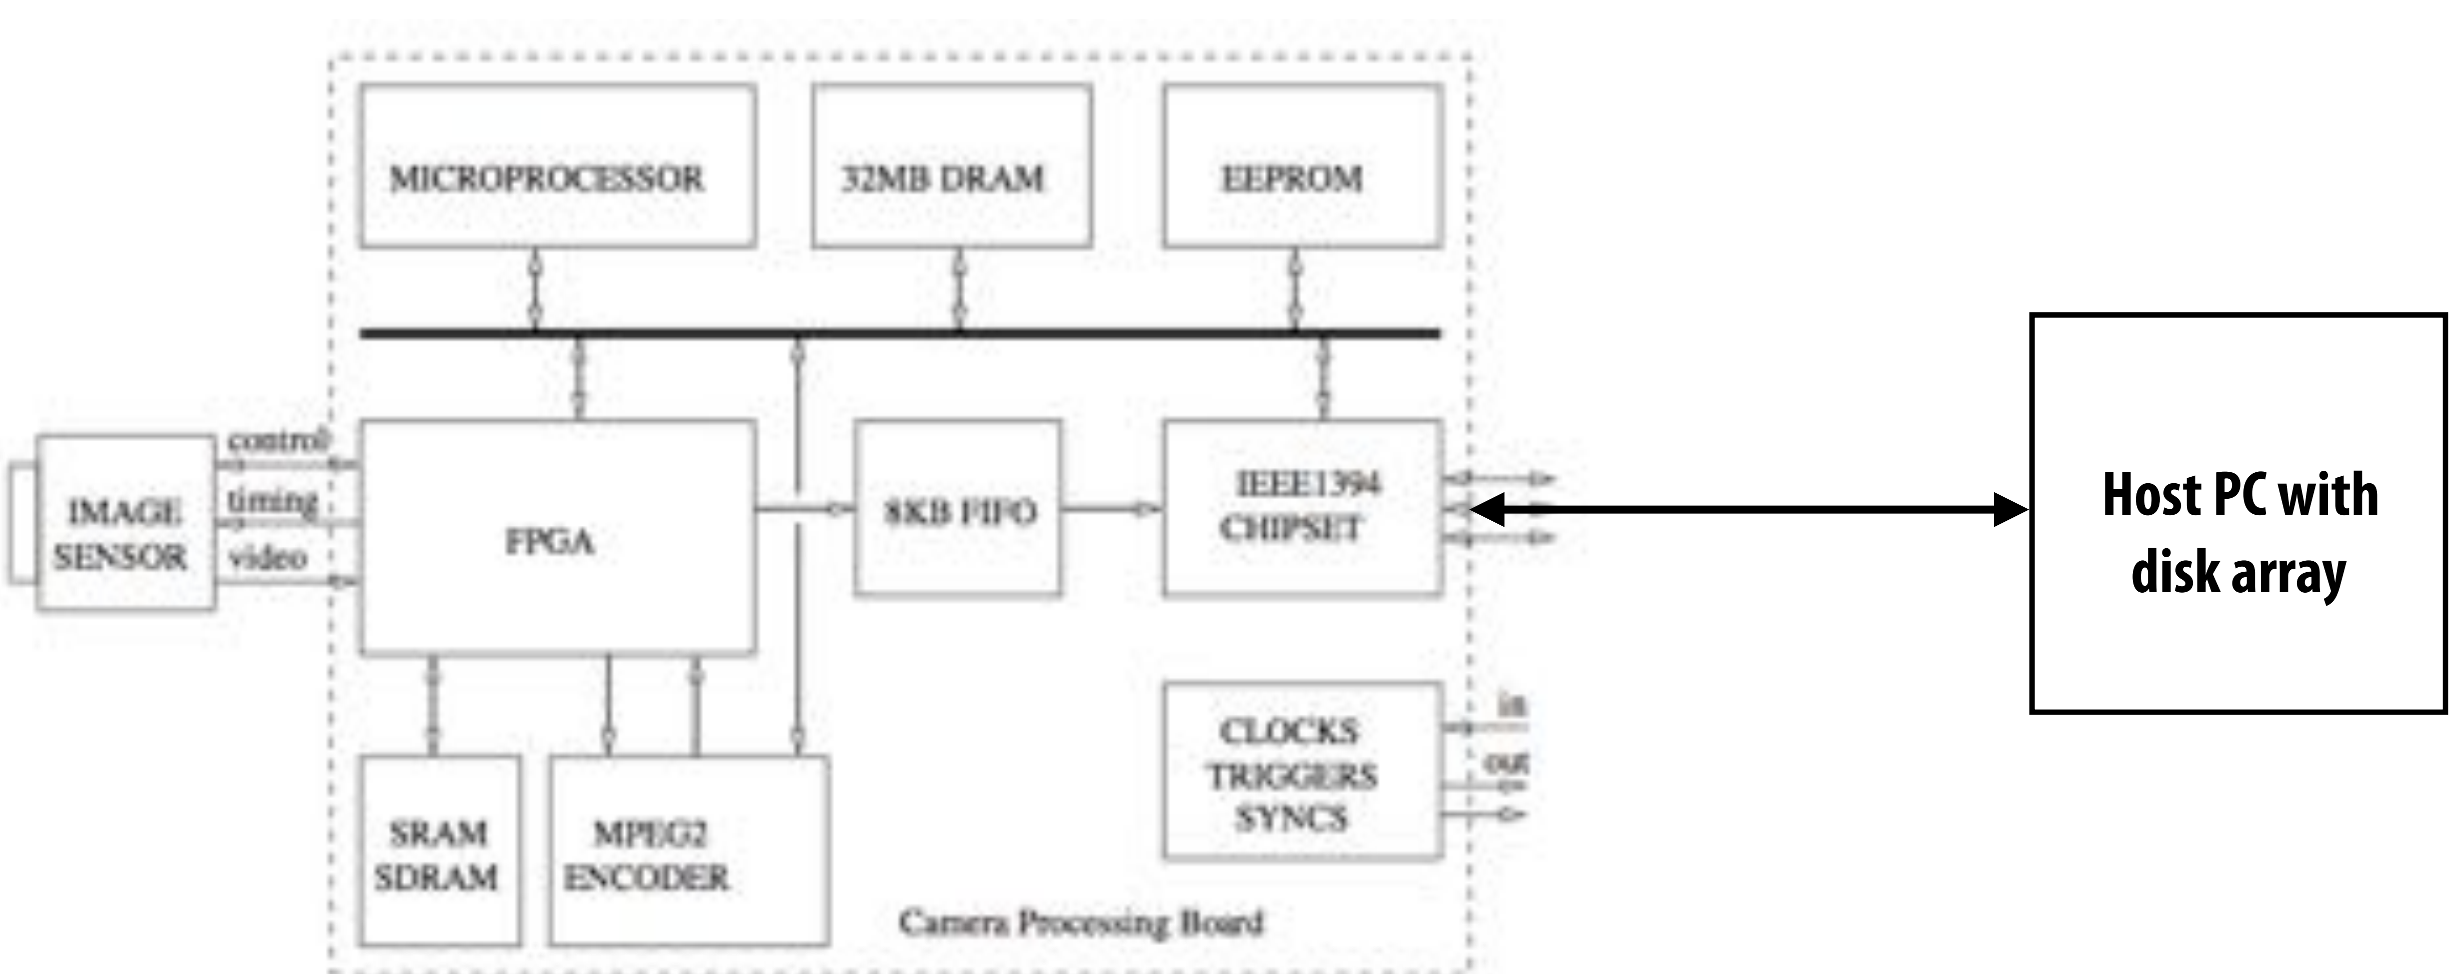

Light field storage layouts

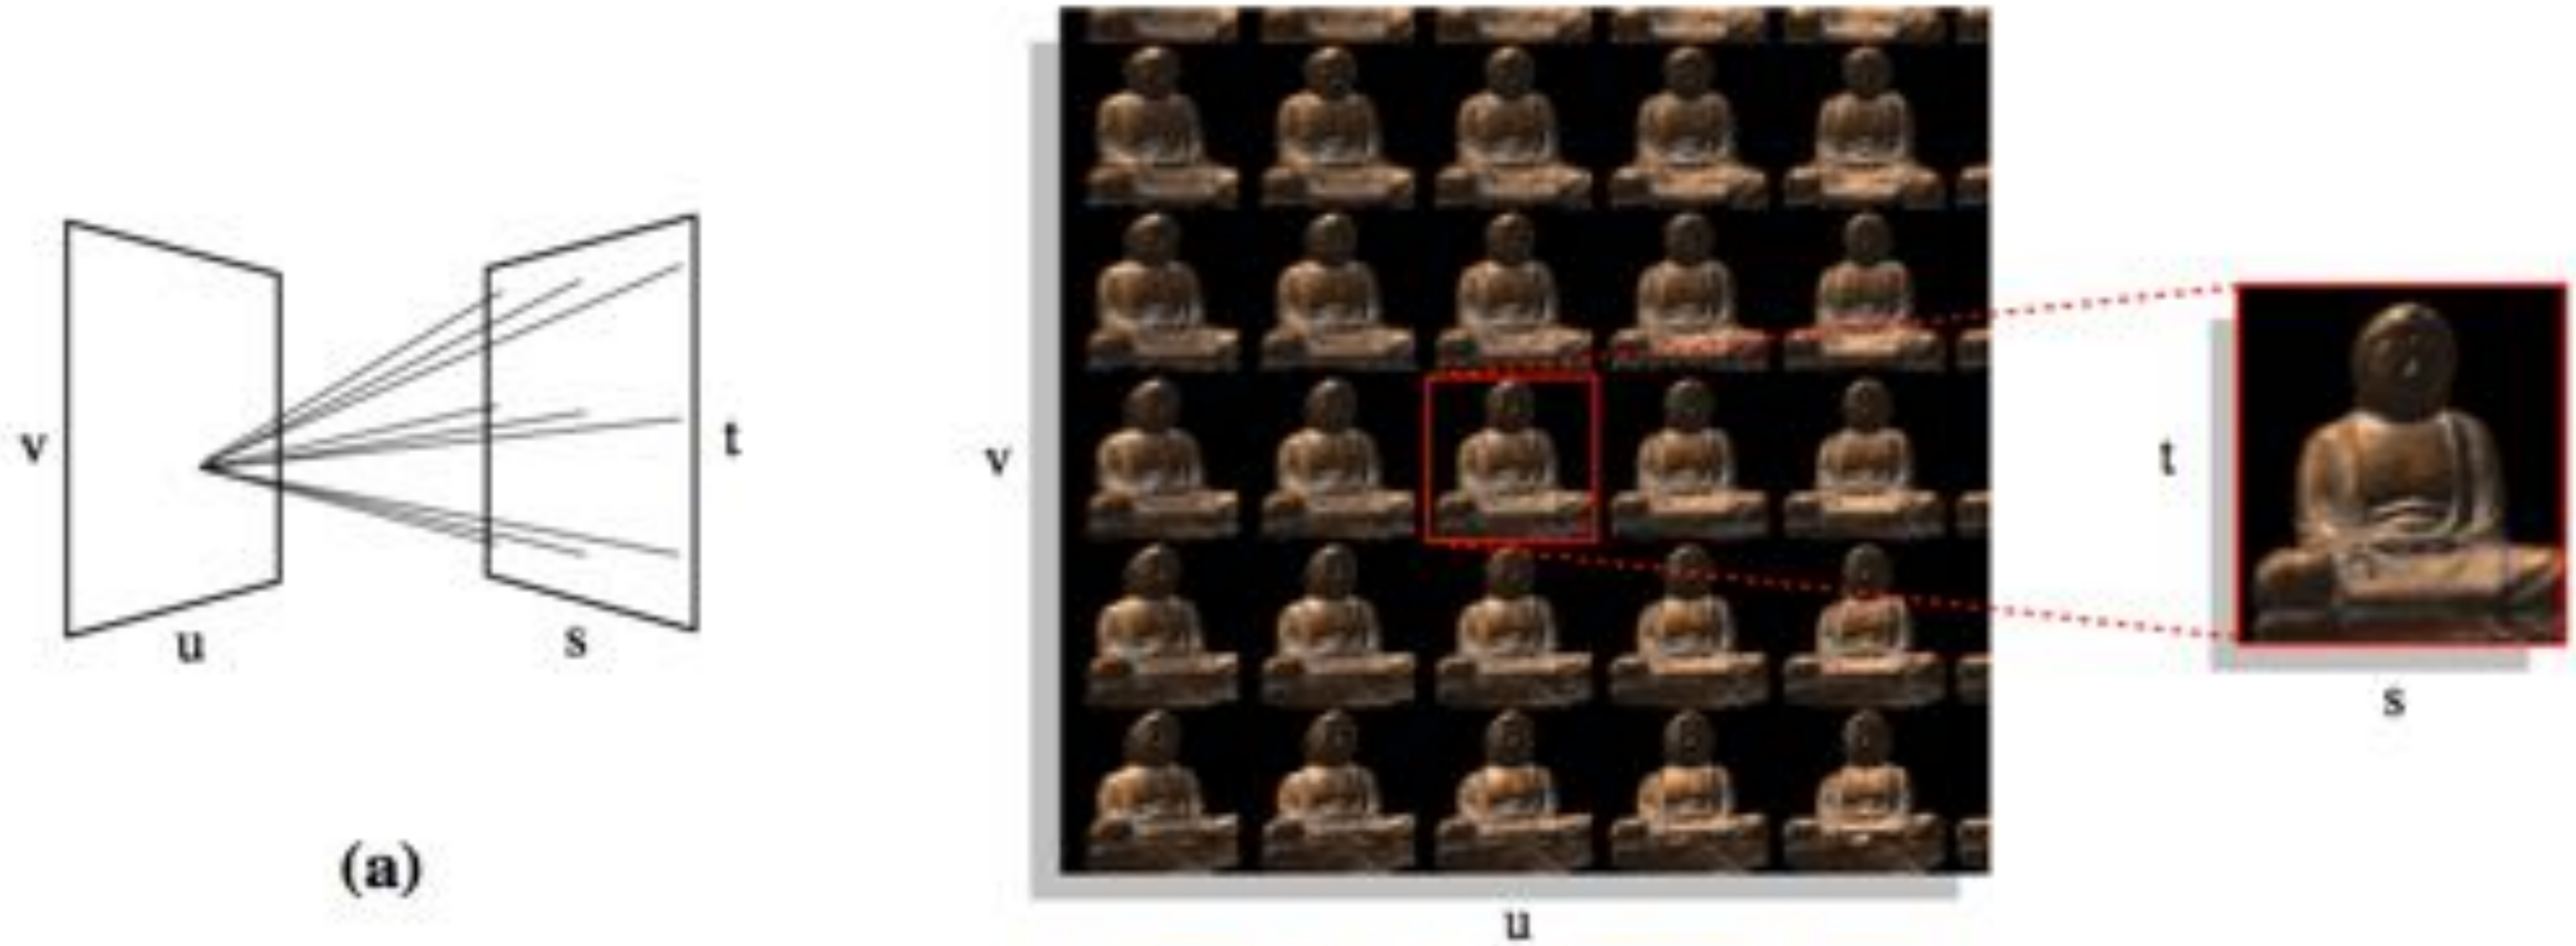

(a)

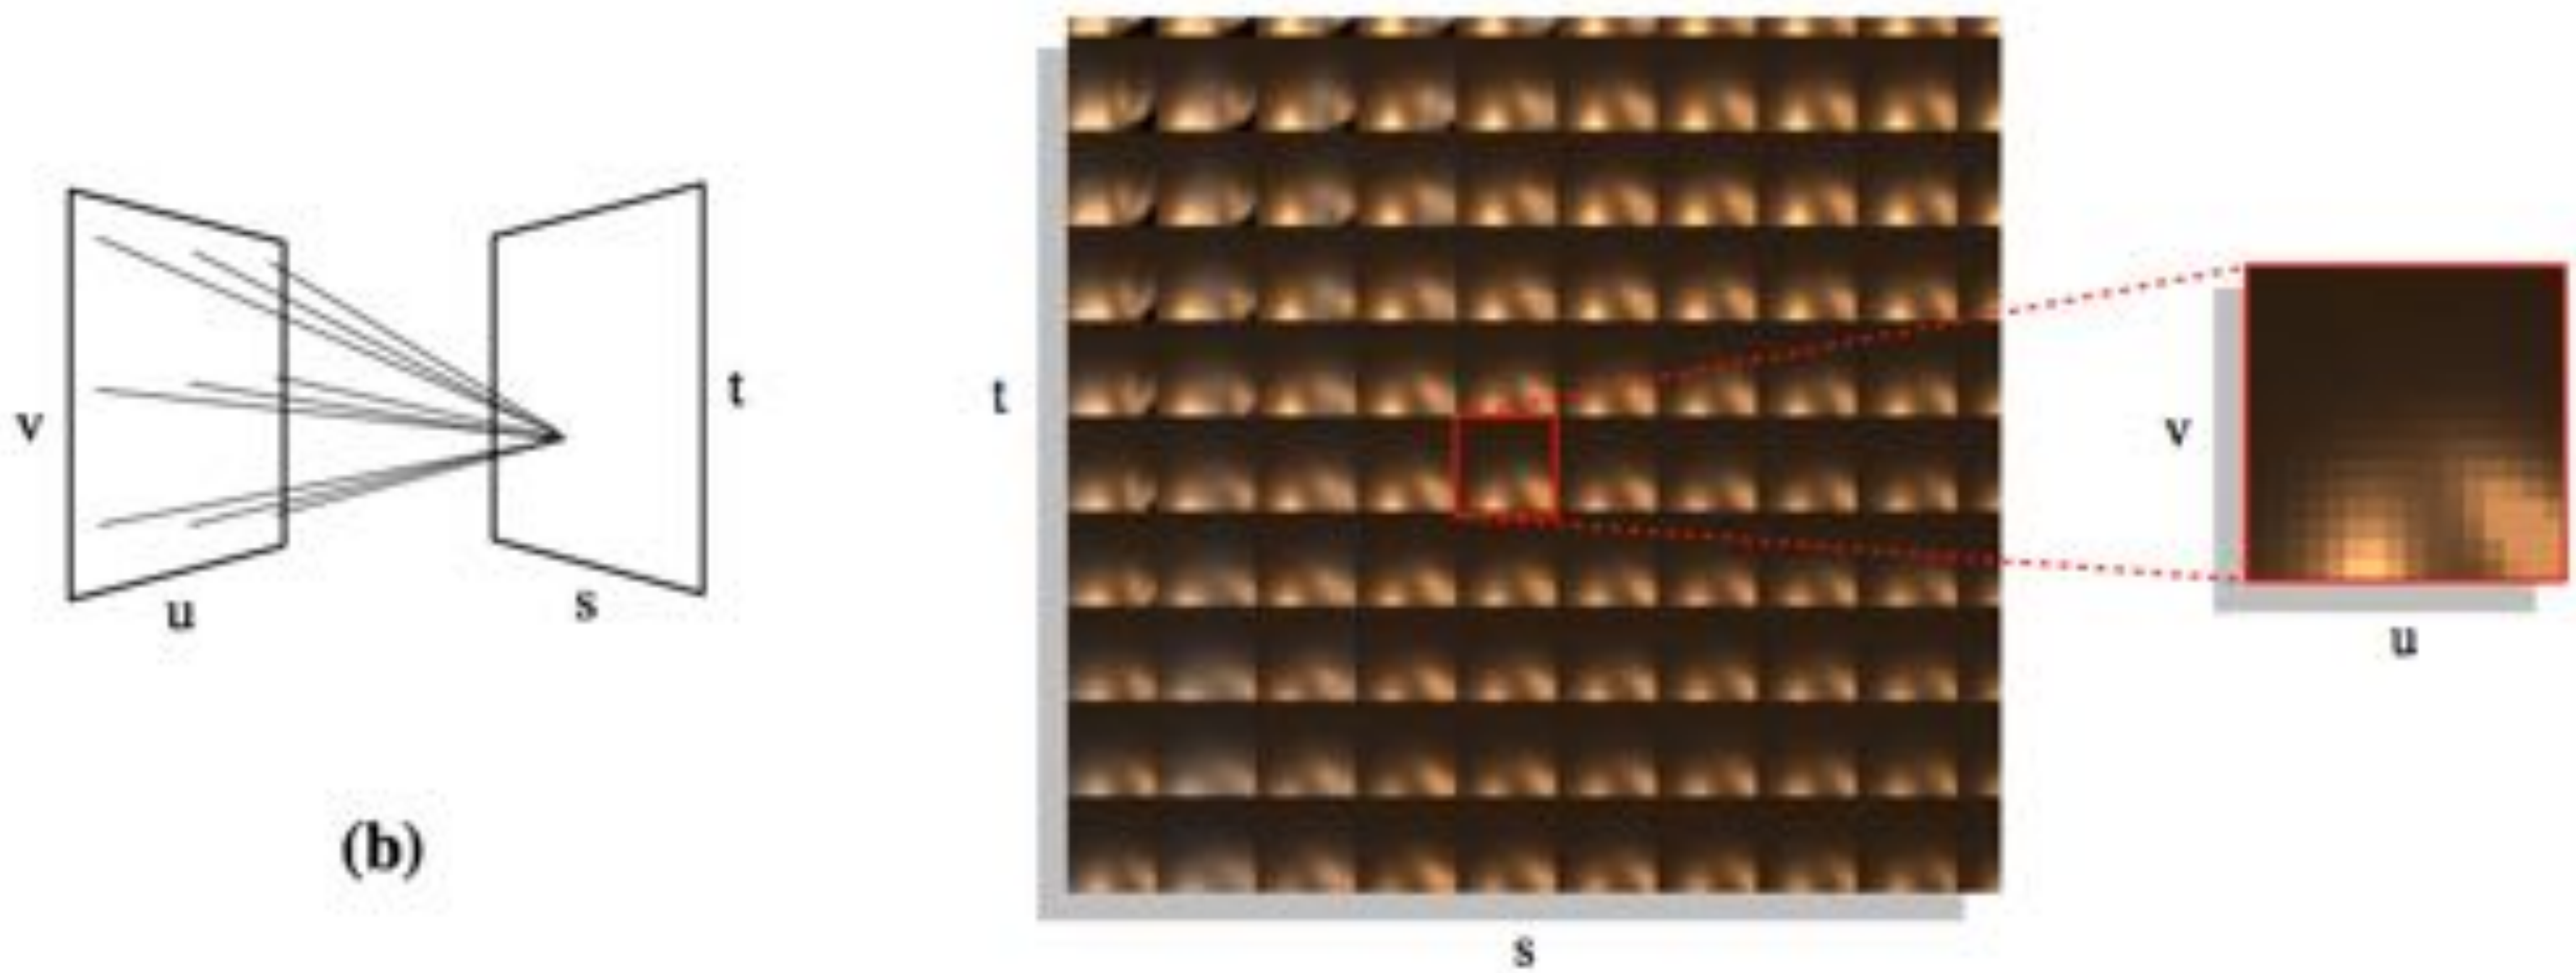

(b)

[Image credit: Levoy and Hanrahan 96]

Microlens Array

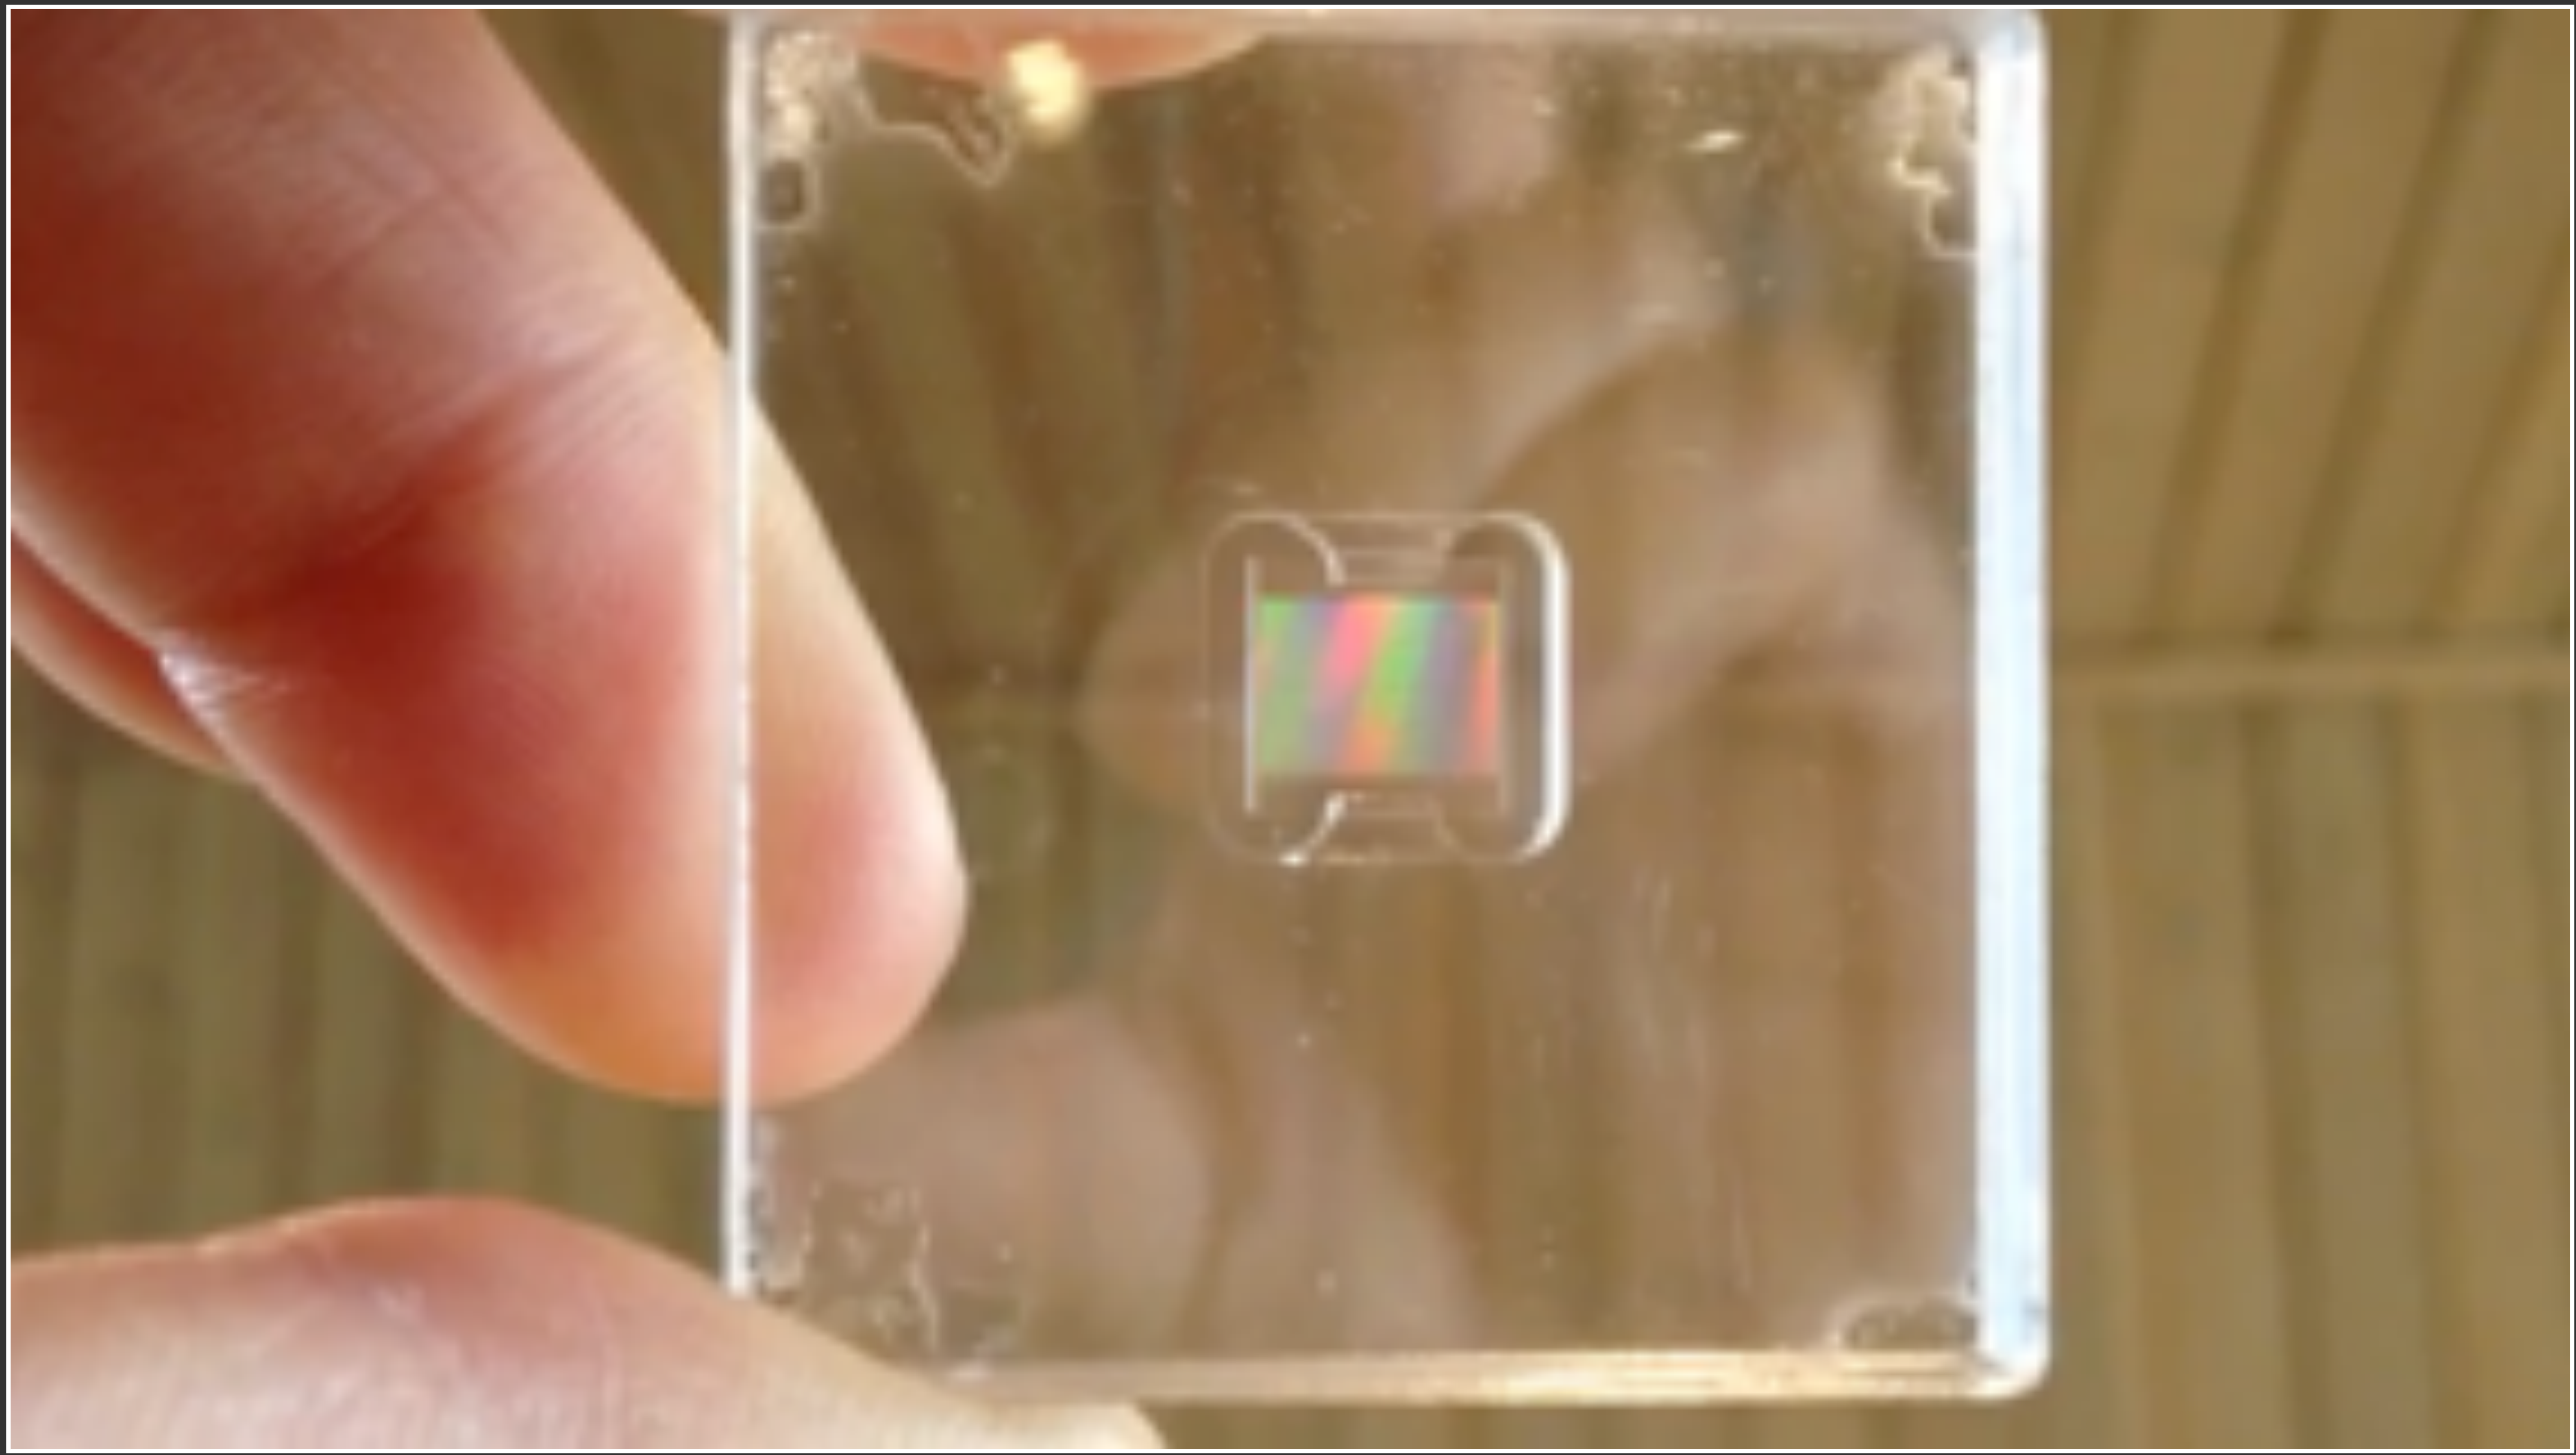

Raw Data From Light Field Sensor

Raw Data From Light Field Sensor

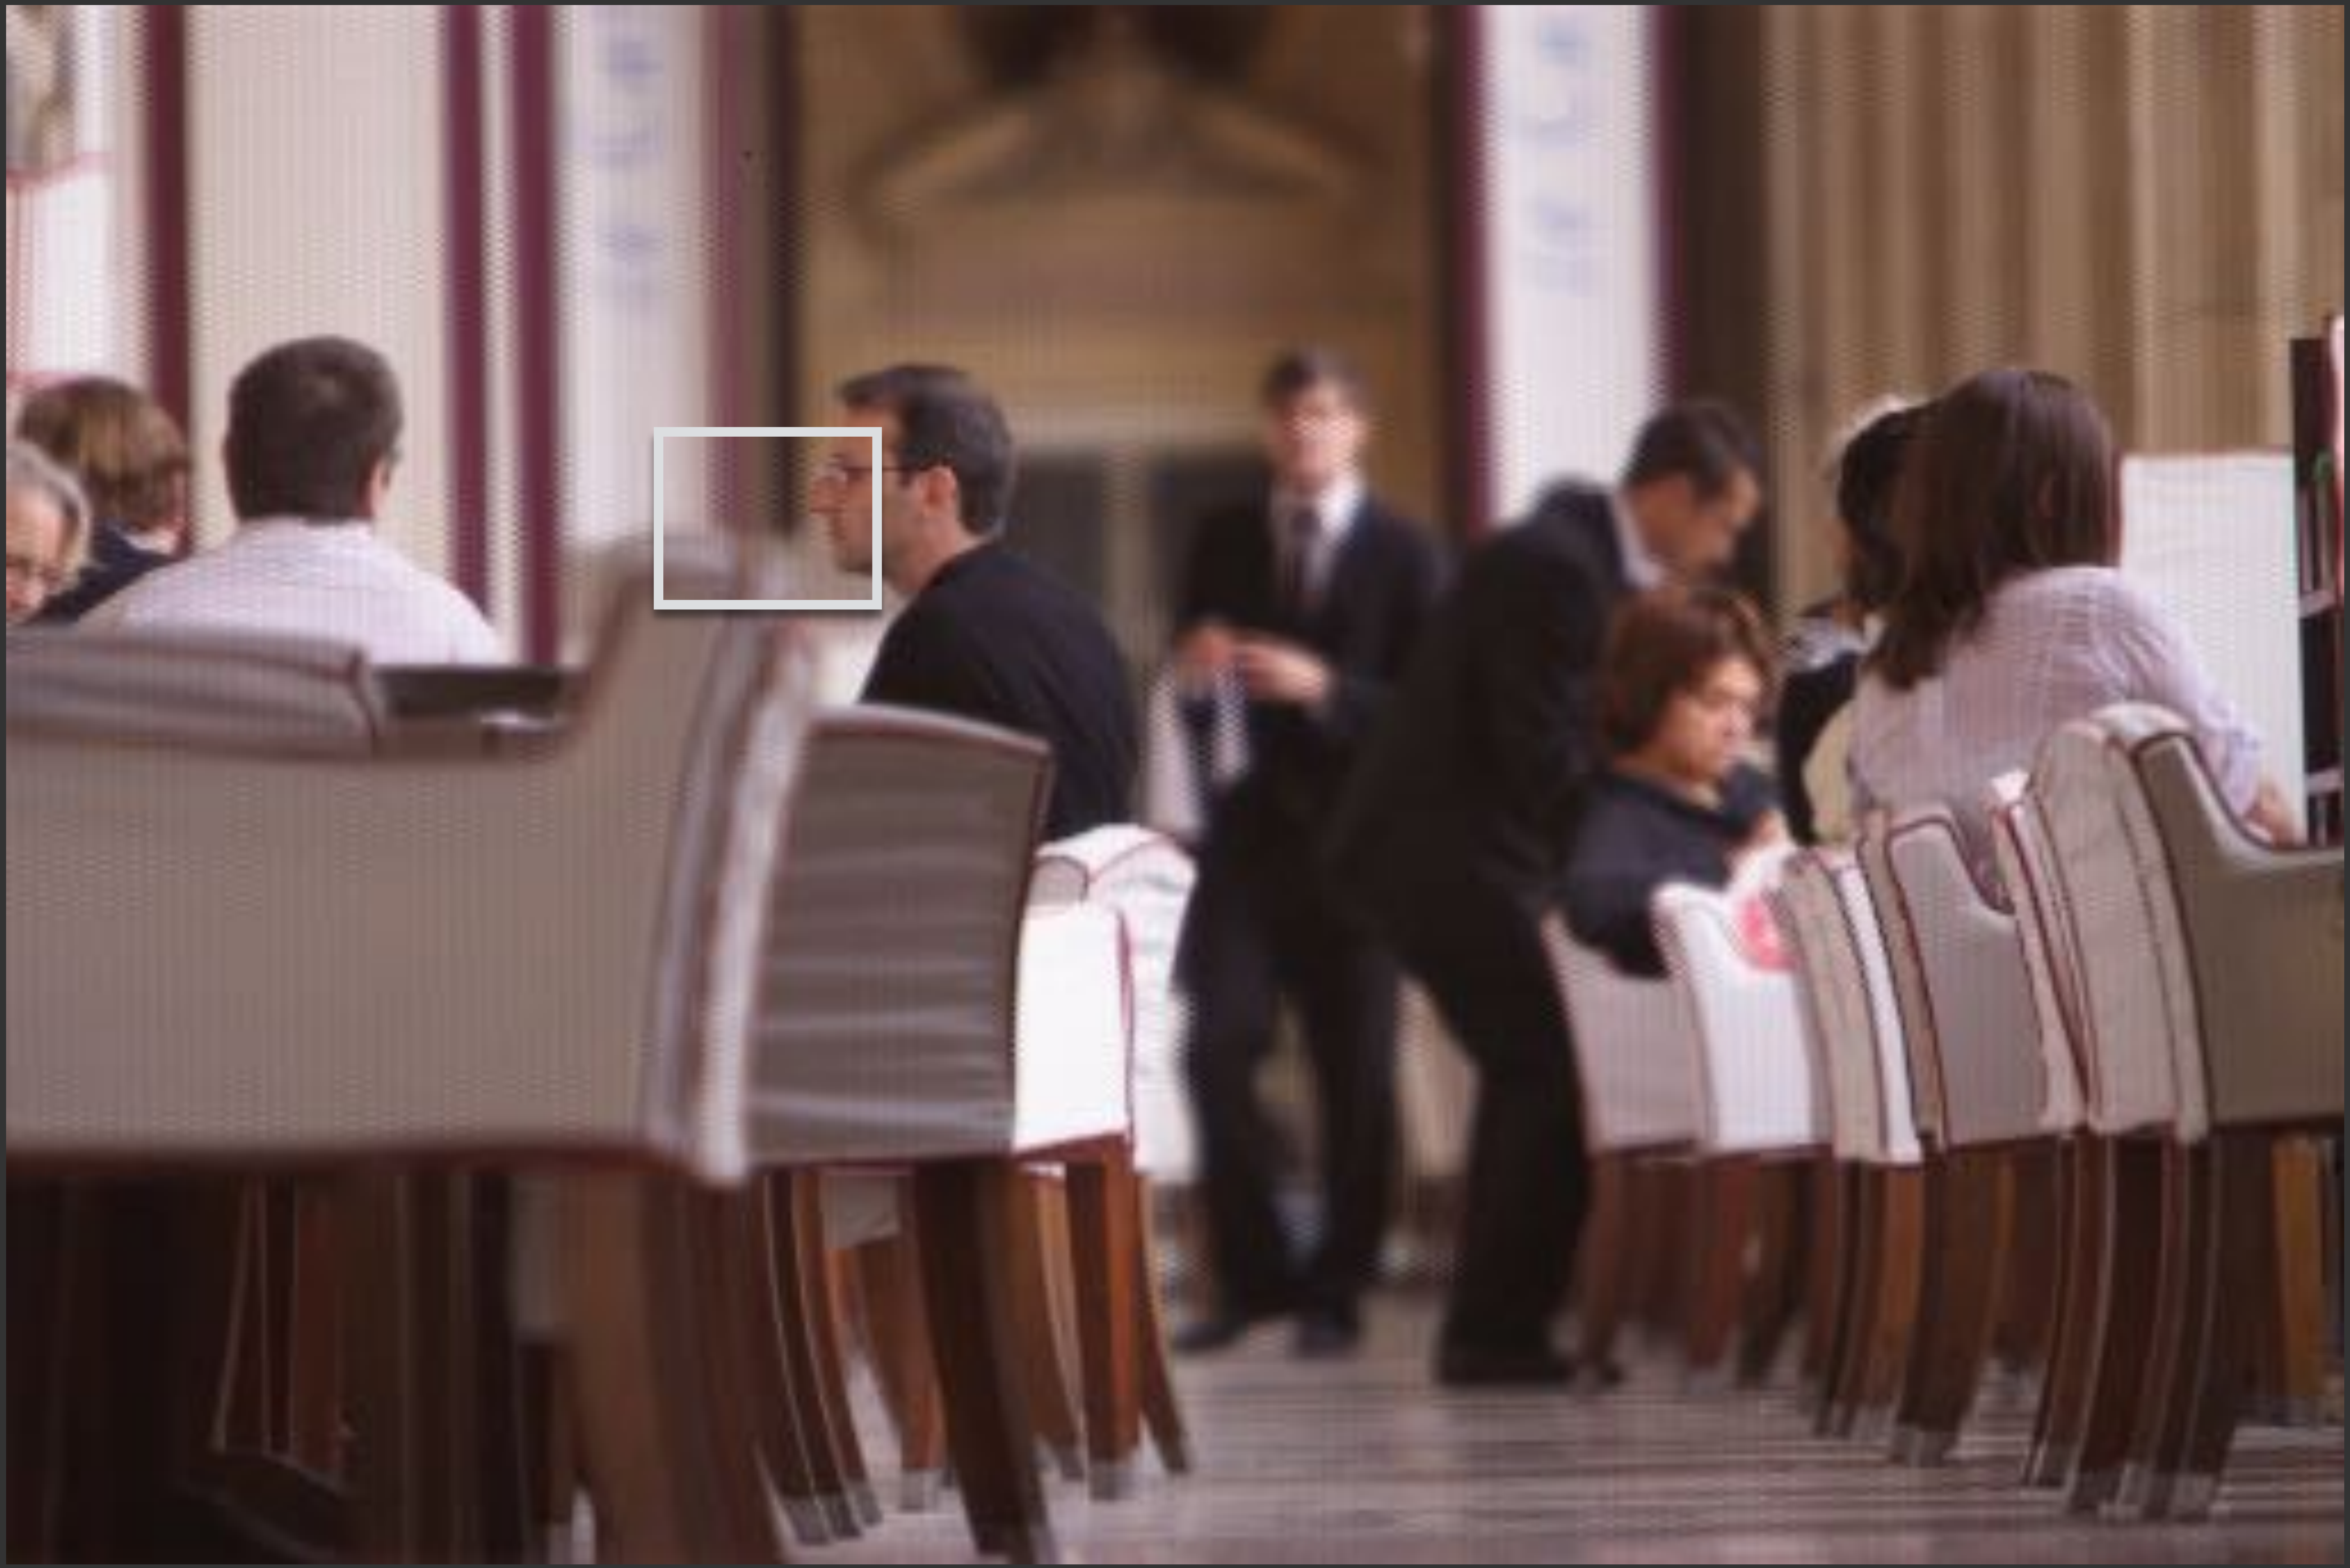

Raw Data From Light Field Sensor

Really captures “many different images” (rays)

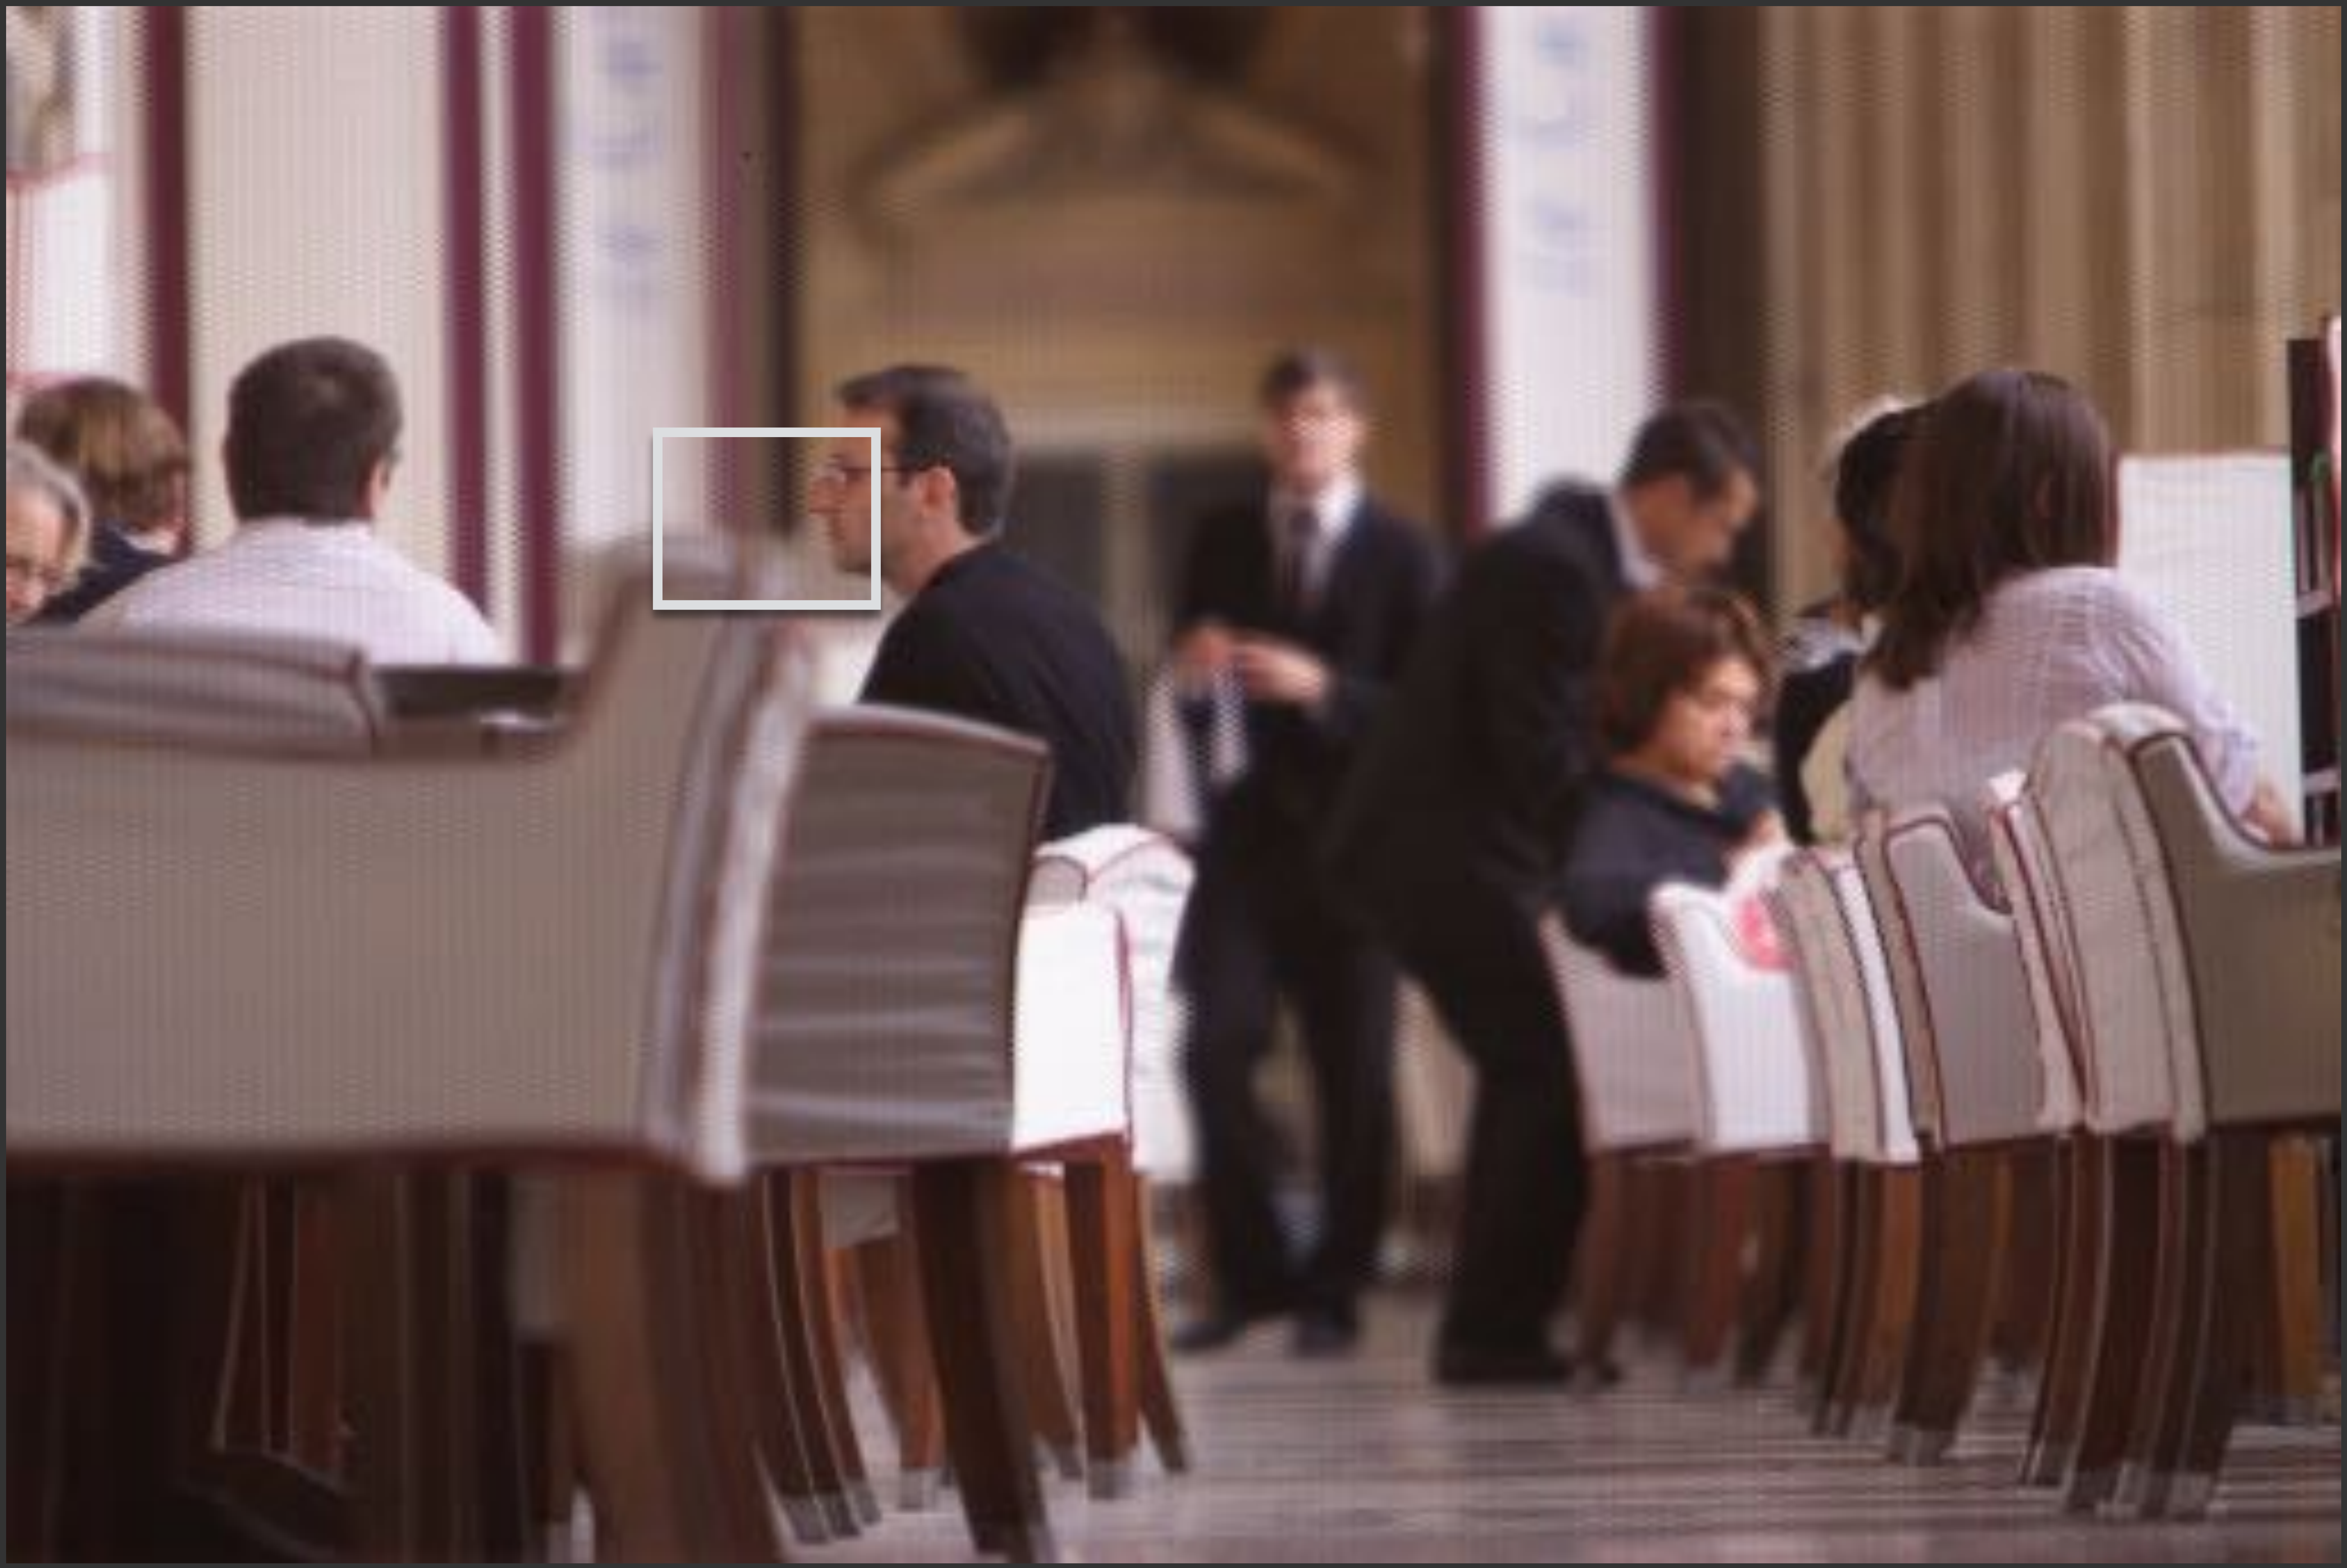

Application: computational Change of Viewpoint

Lateral movement (left)

Application: computational Change of Viewpoint

Lateral movement (right)

Computational Photography

Even with standard digital camera, the values of pixels in photograph you see on screen are quite different than the values output by the photosensor in the original camera

<u>Computation</u> has become a fundamental aspect of producing high-quality pictures

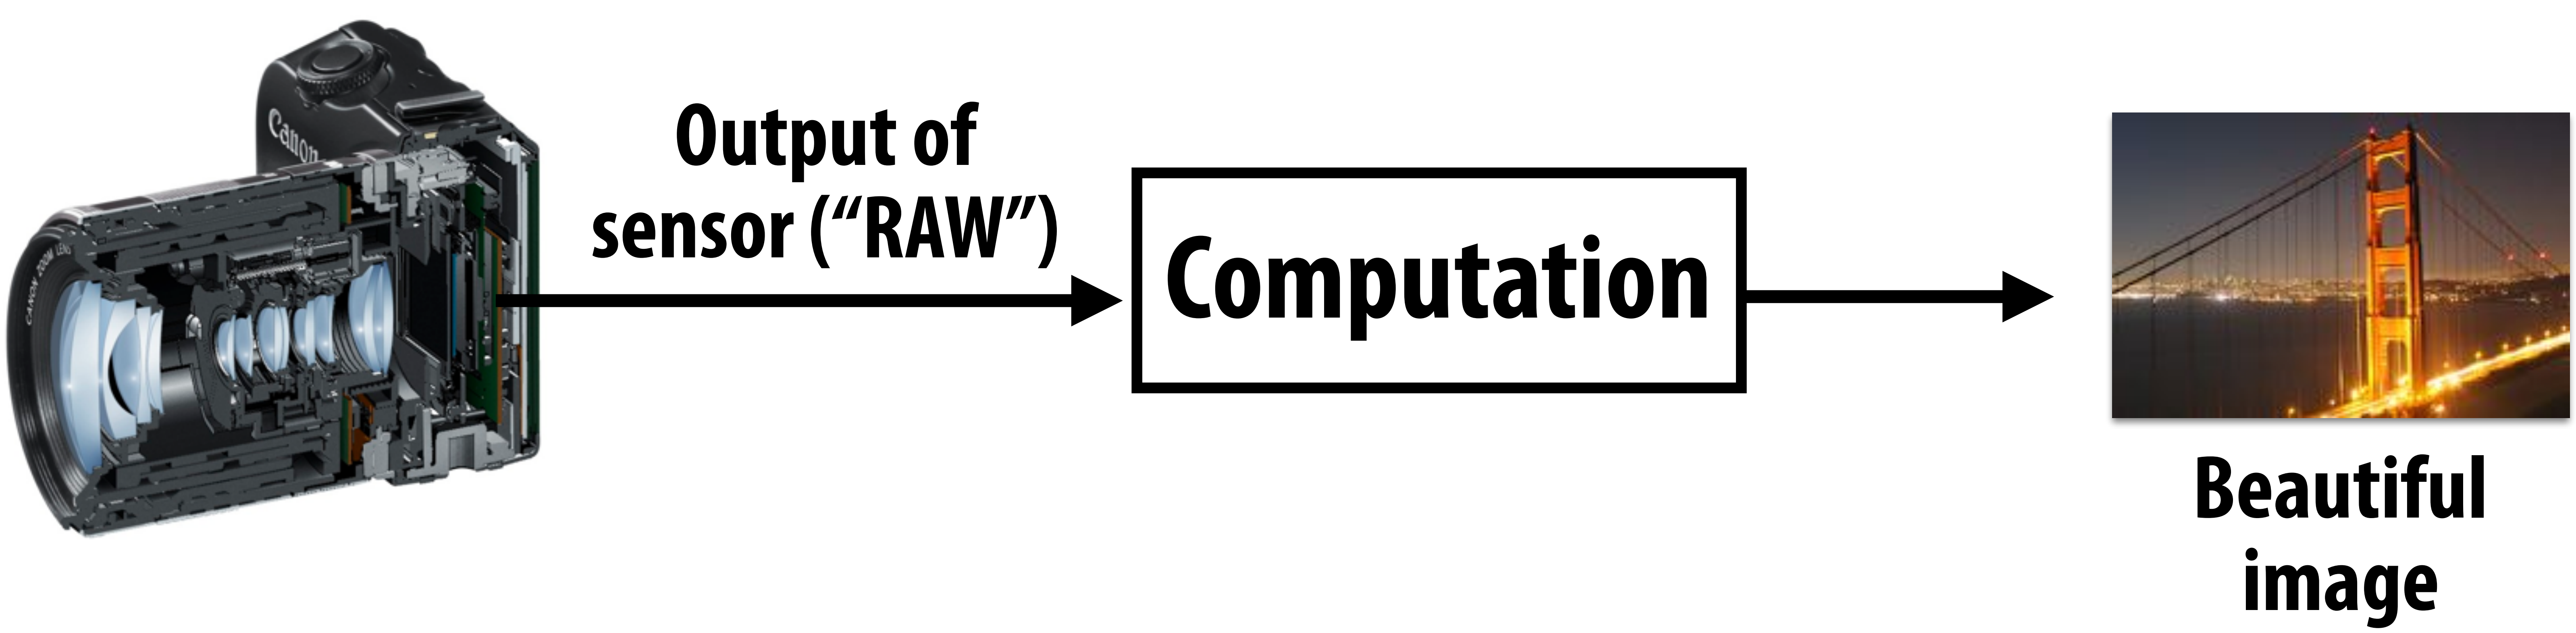

Computational Photography

- Since we're processing images anyway, why not take it even further?

Example: HDR Tone Mapping

- **For any real camera, single photo has limits on intensity range**

Example: HDR Tone Mapping

- **For any real camera, single photo has limits on intensity range**

Example: HDR Tone Mapping

- For any real camera, single photo has limits on intensity range

Example: HDR Tone Mapping

- Use algorithms to combine images, convey original "feeling"

Example: Image Completion

- **Fill in missing parts of image**

Example: Image Completion

- **Use data to “fill in” missing/undesirable parts of images**

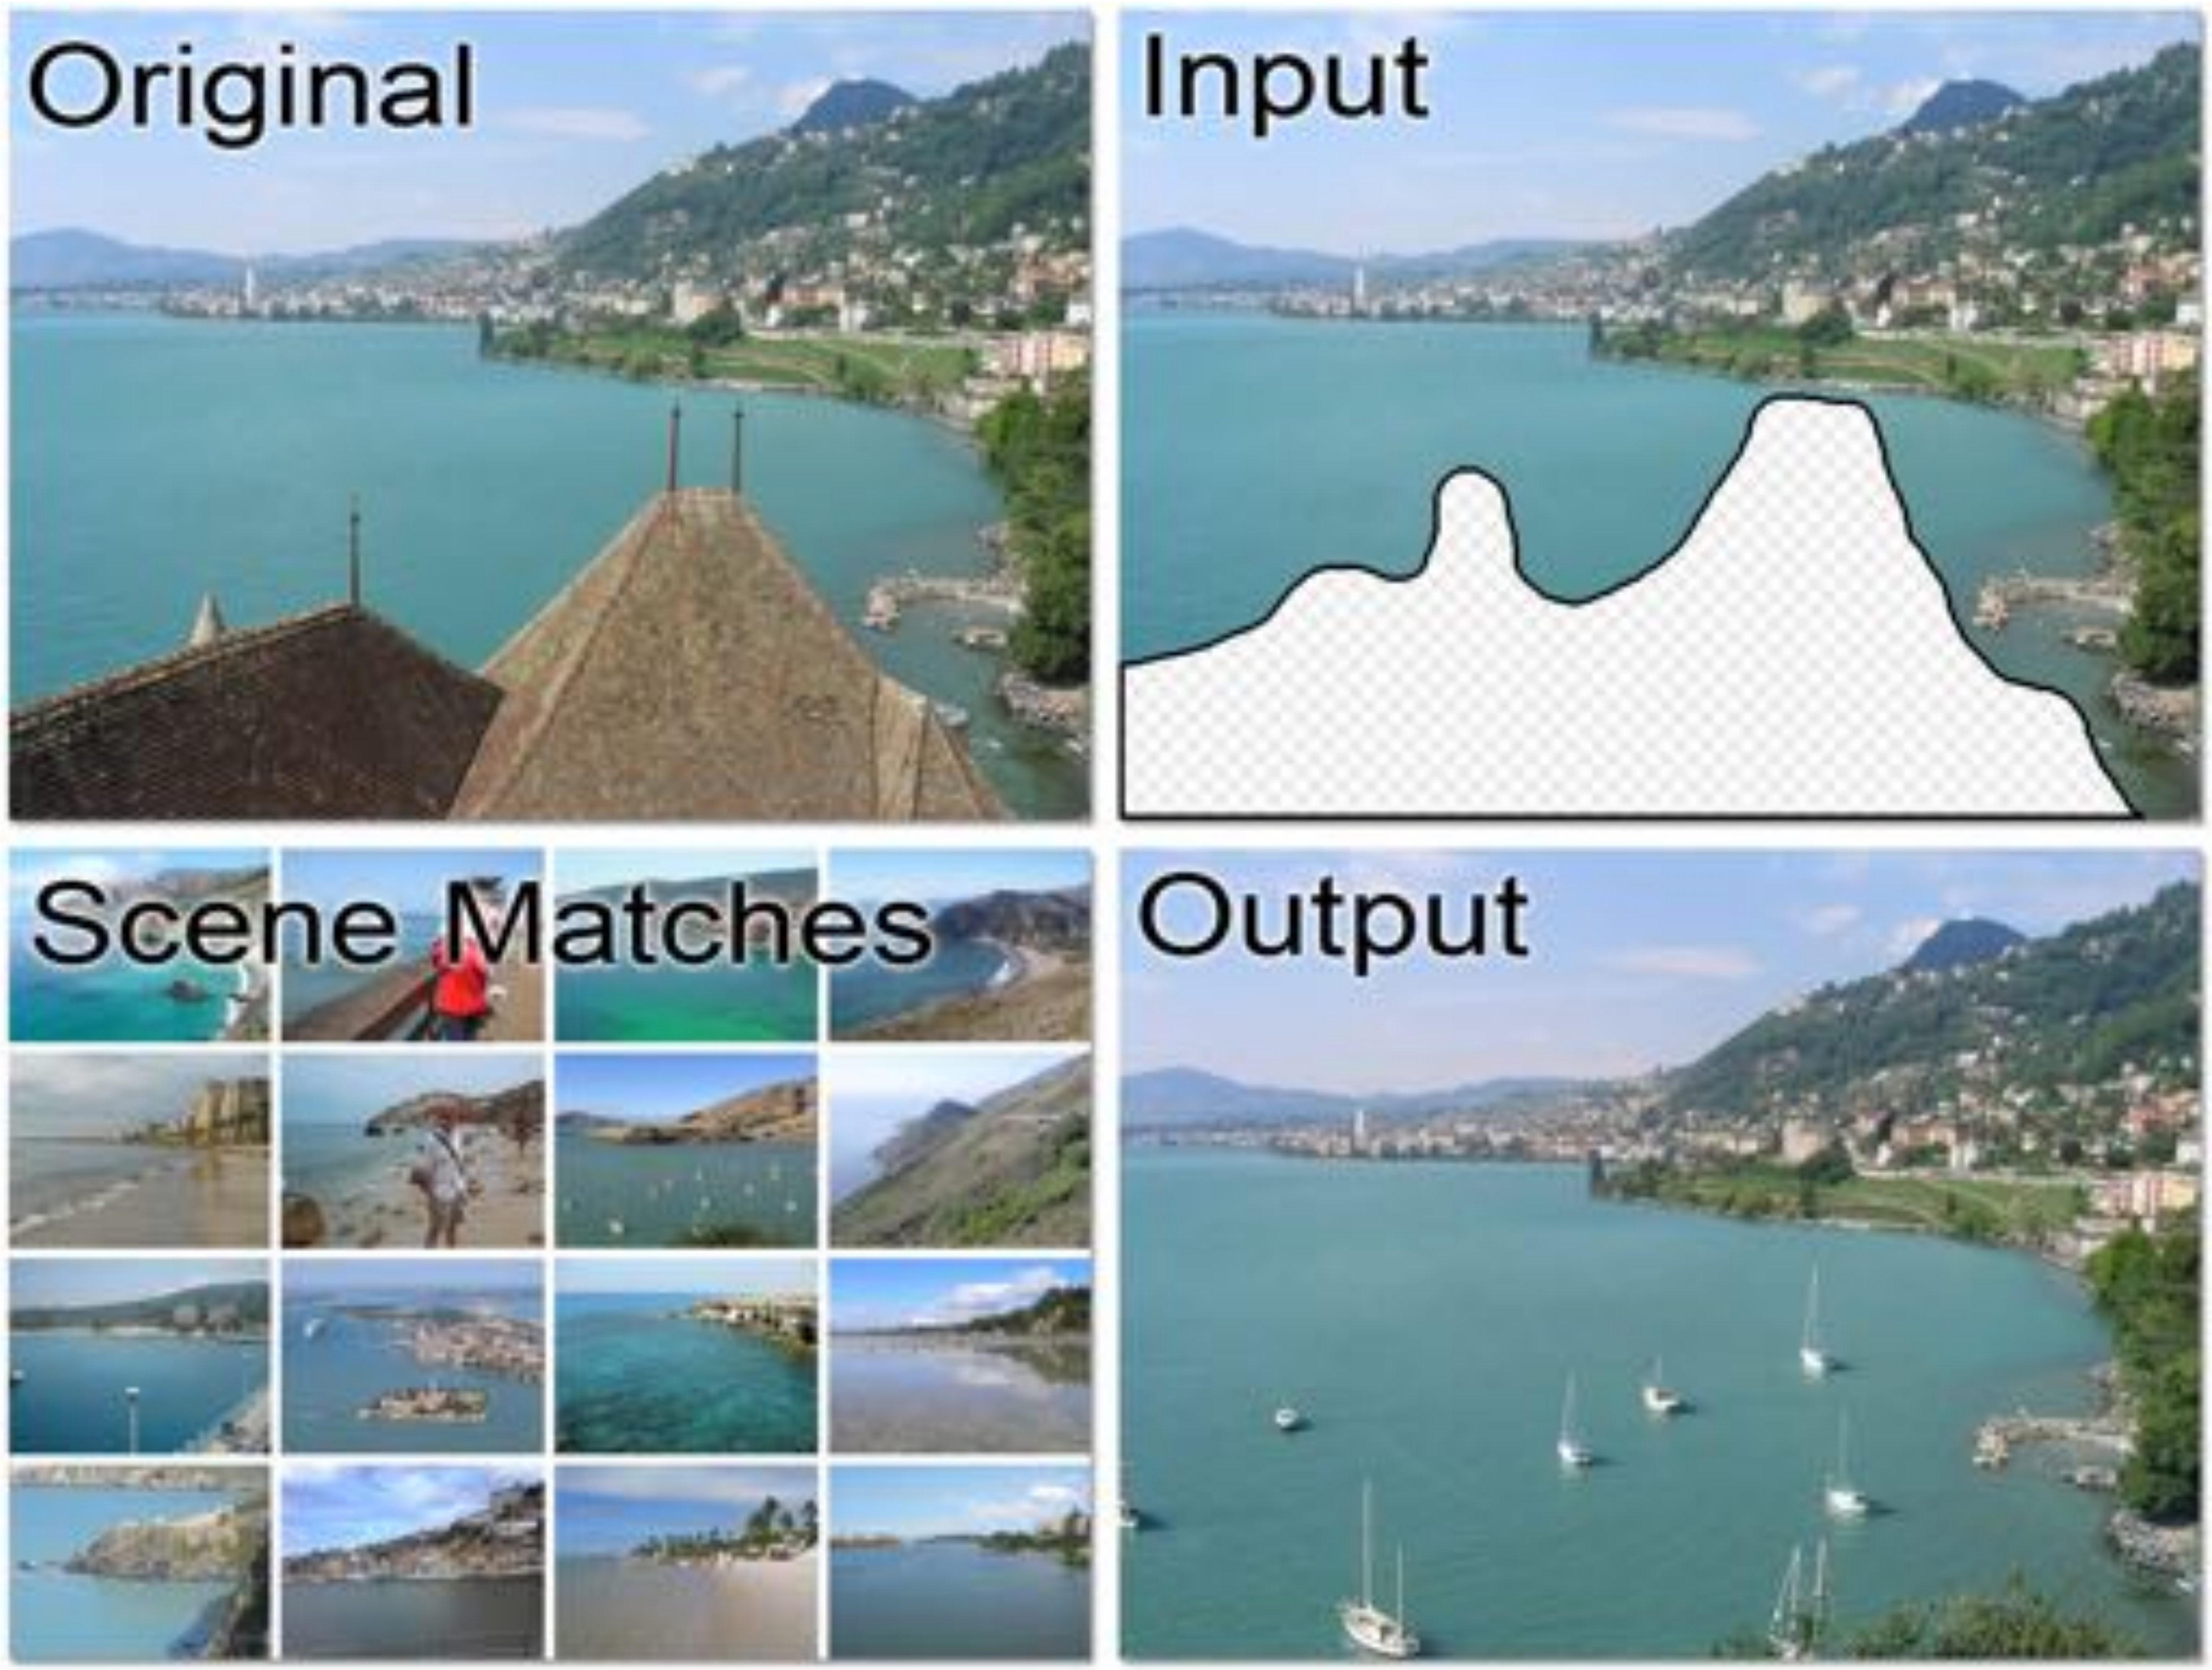

Hays & Efros, “Scene Completion Using Millions of Photographs” (SIGGRAPH 2007)

Example: Facial Reenactment

- **Map motion from one video to other video (use 3D model)**

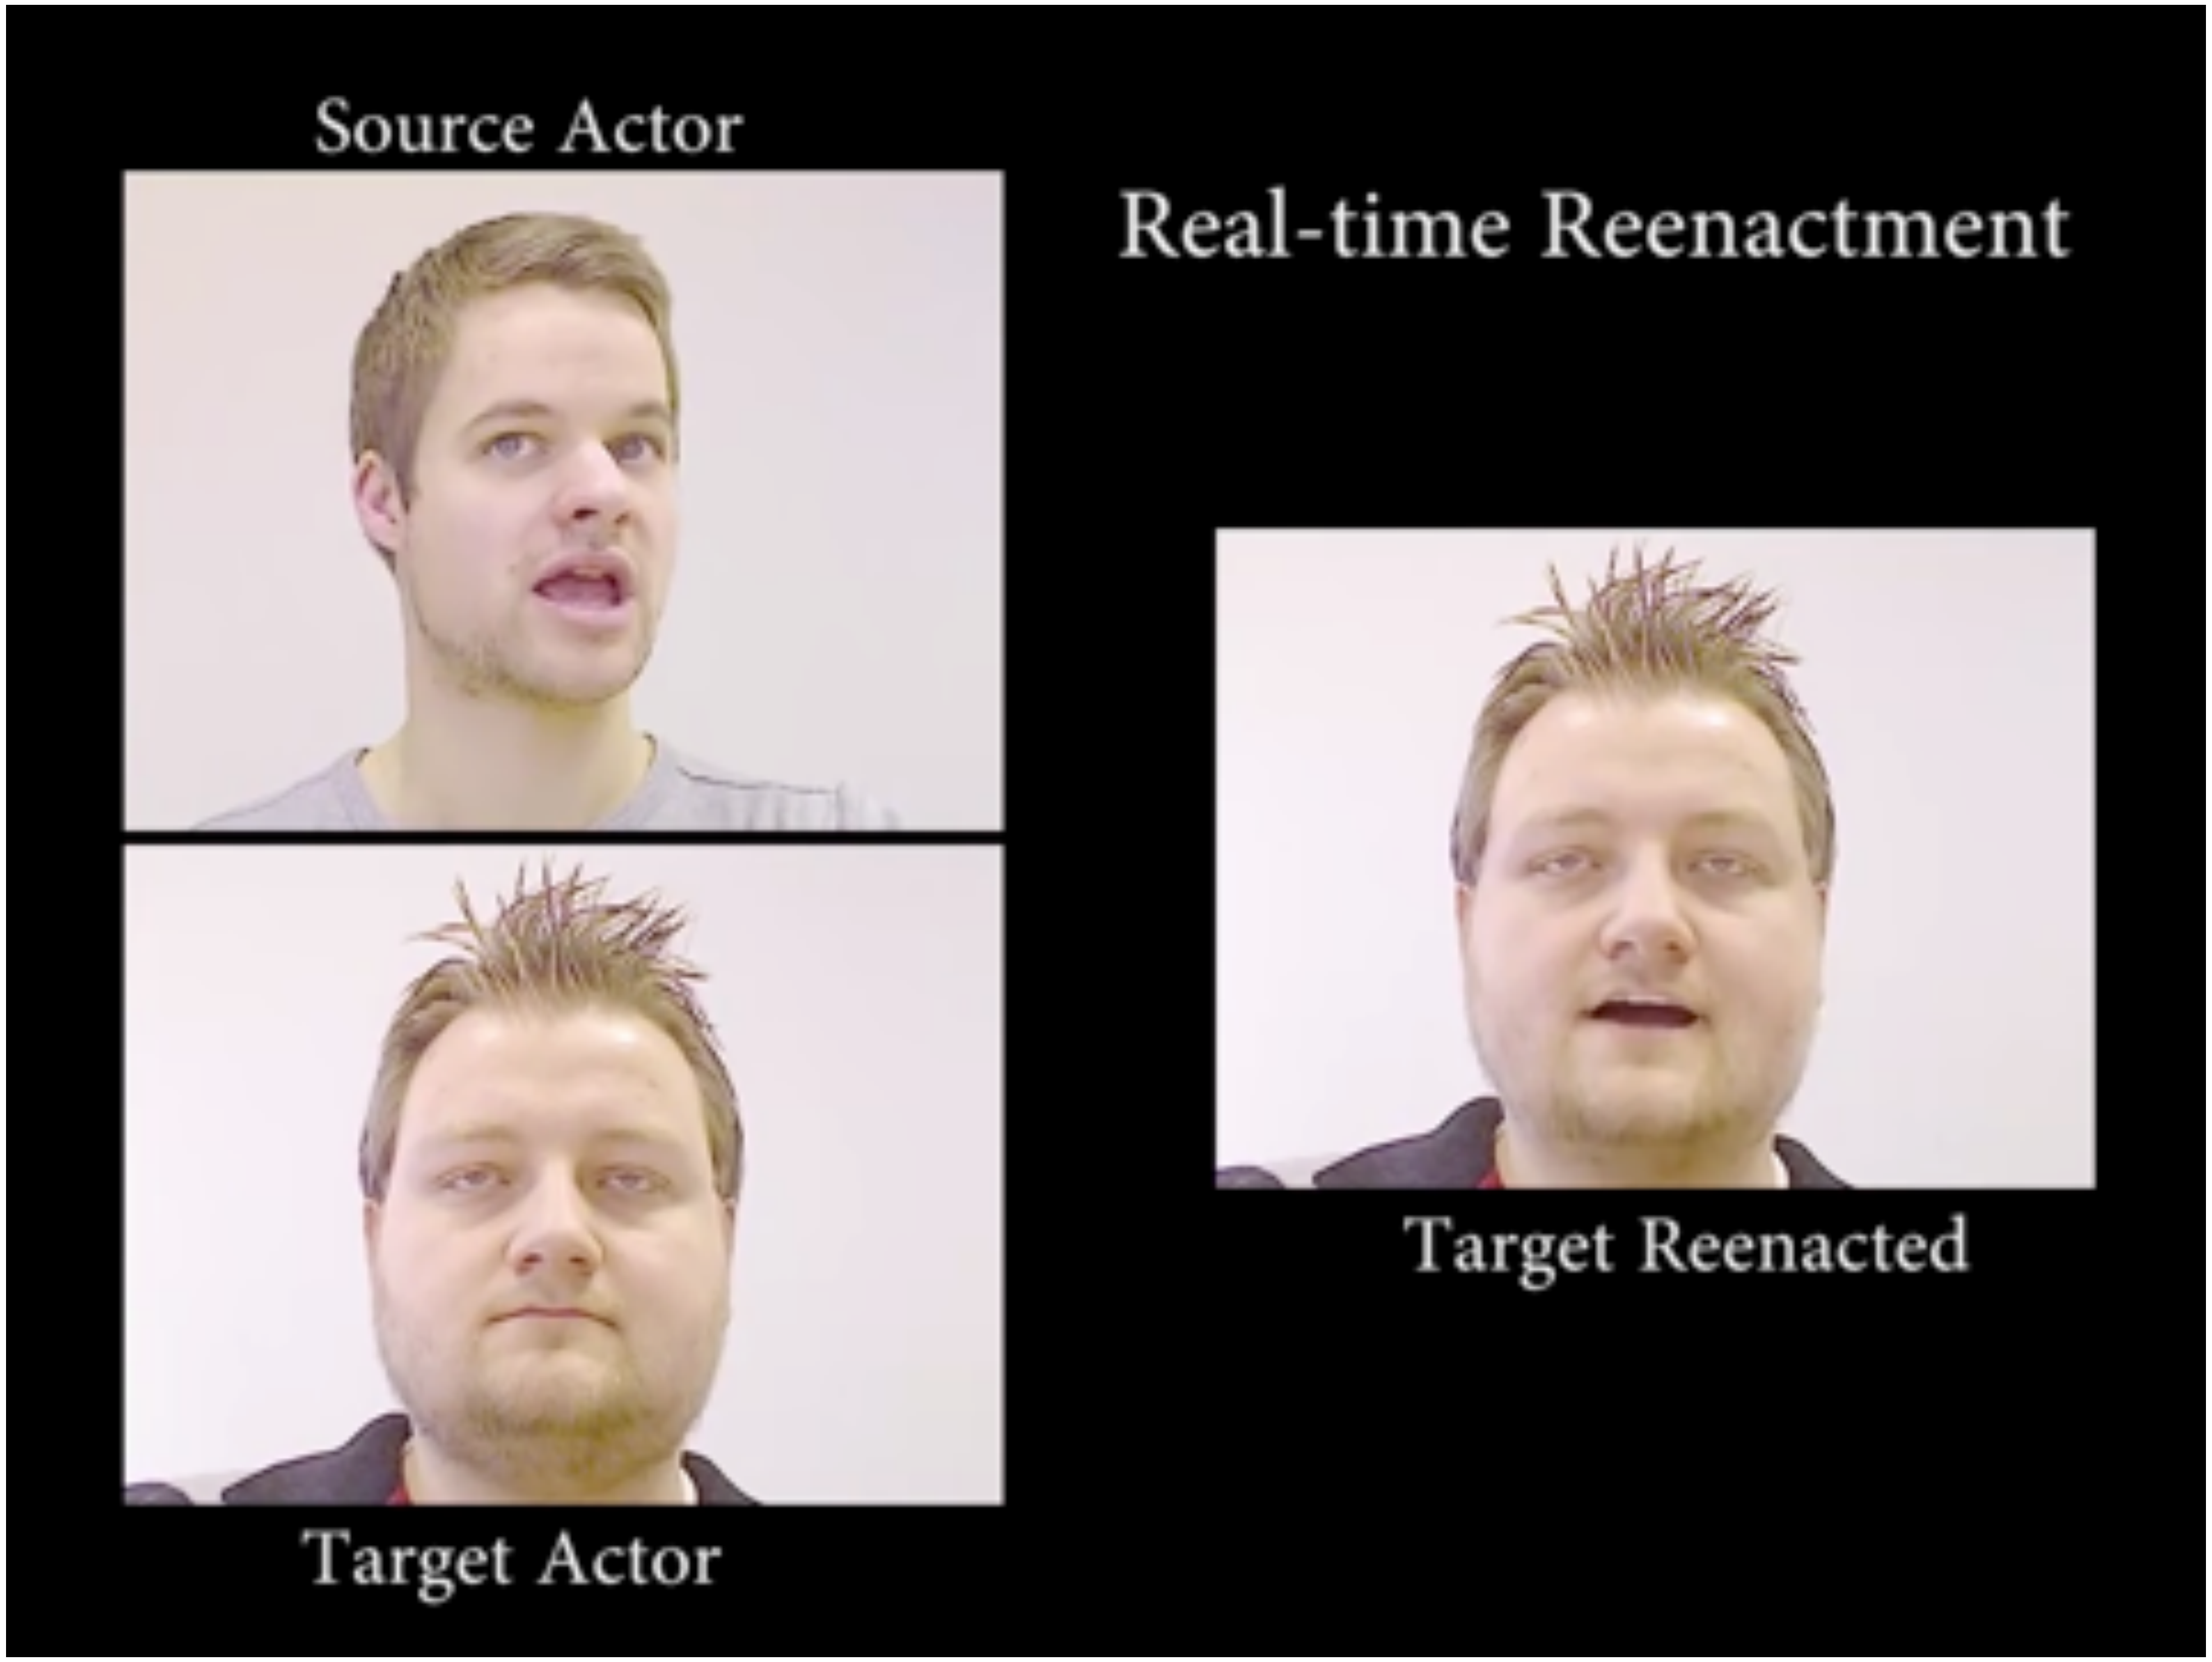

Thies et al, "Real-Time Expression Transfer for Facial Reenactment" (Transactions on Graphics 2015)

Ethical Issues?

https://www.radiolab.org/story/breaking-news/

Digital Forensics

- **Also algorithms to detect whether a photo is fake!**

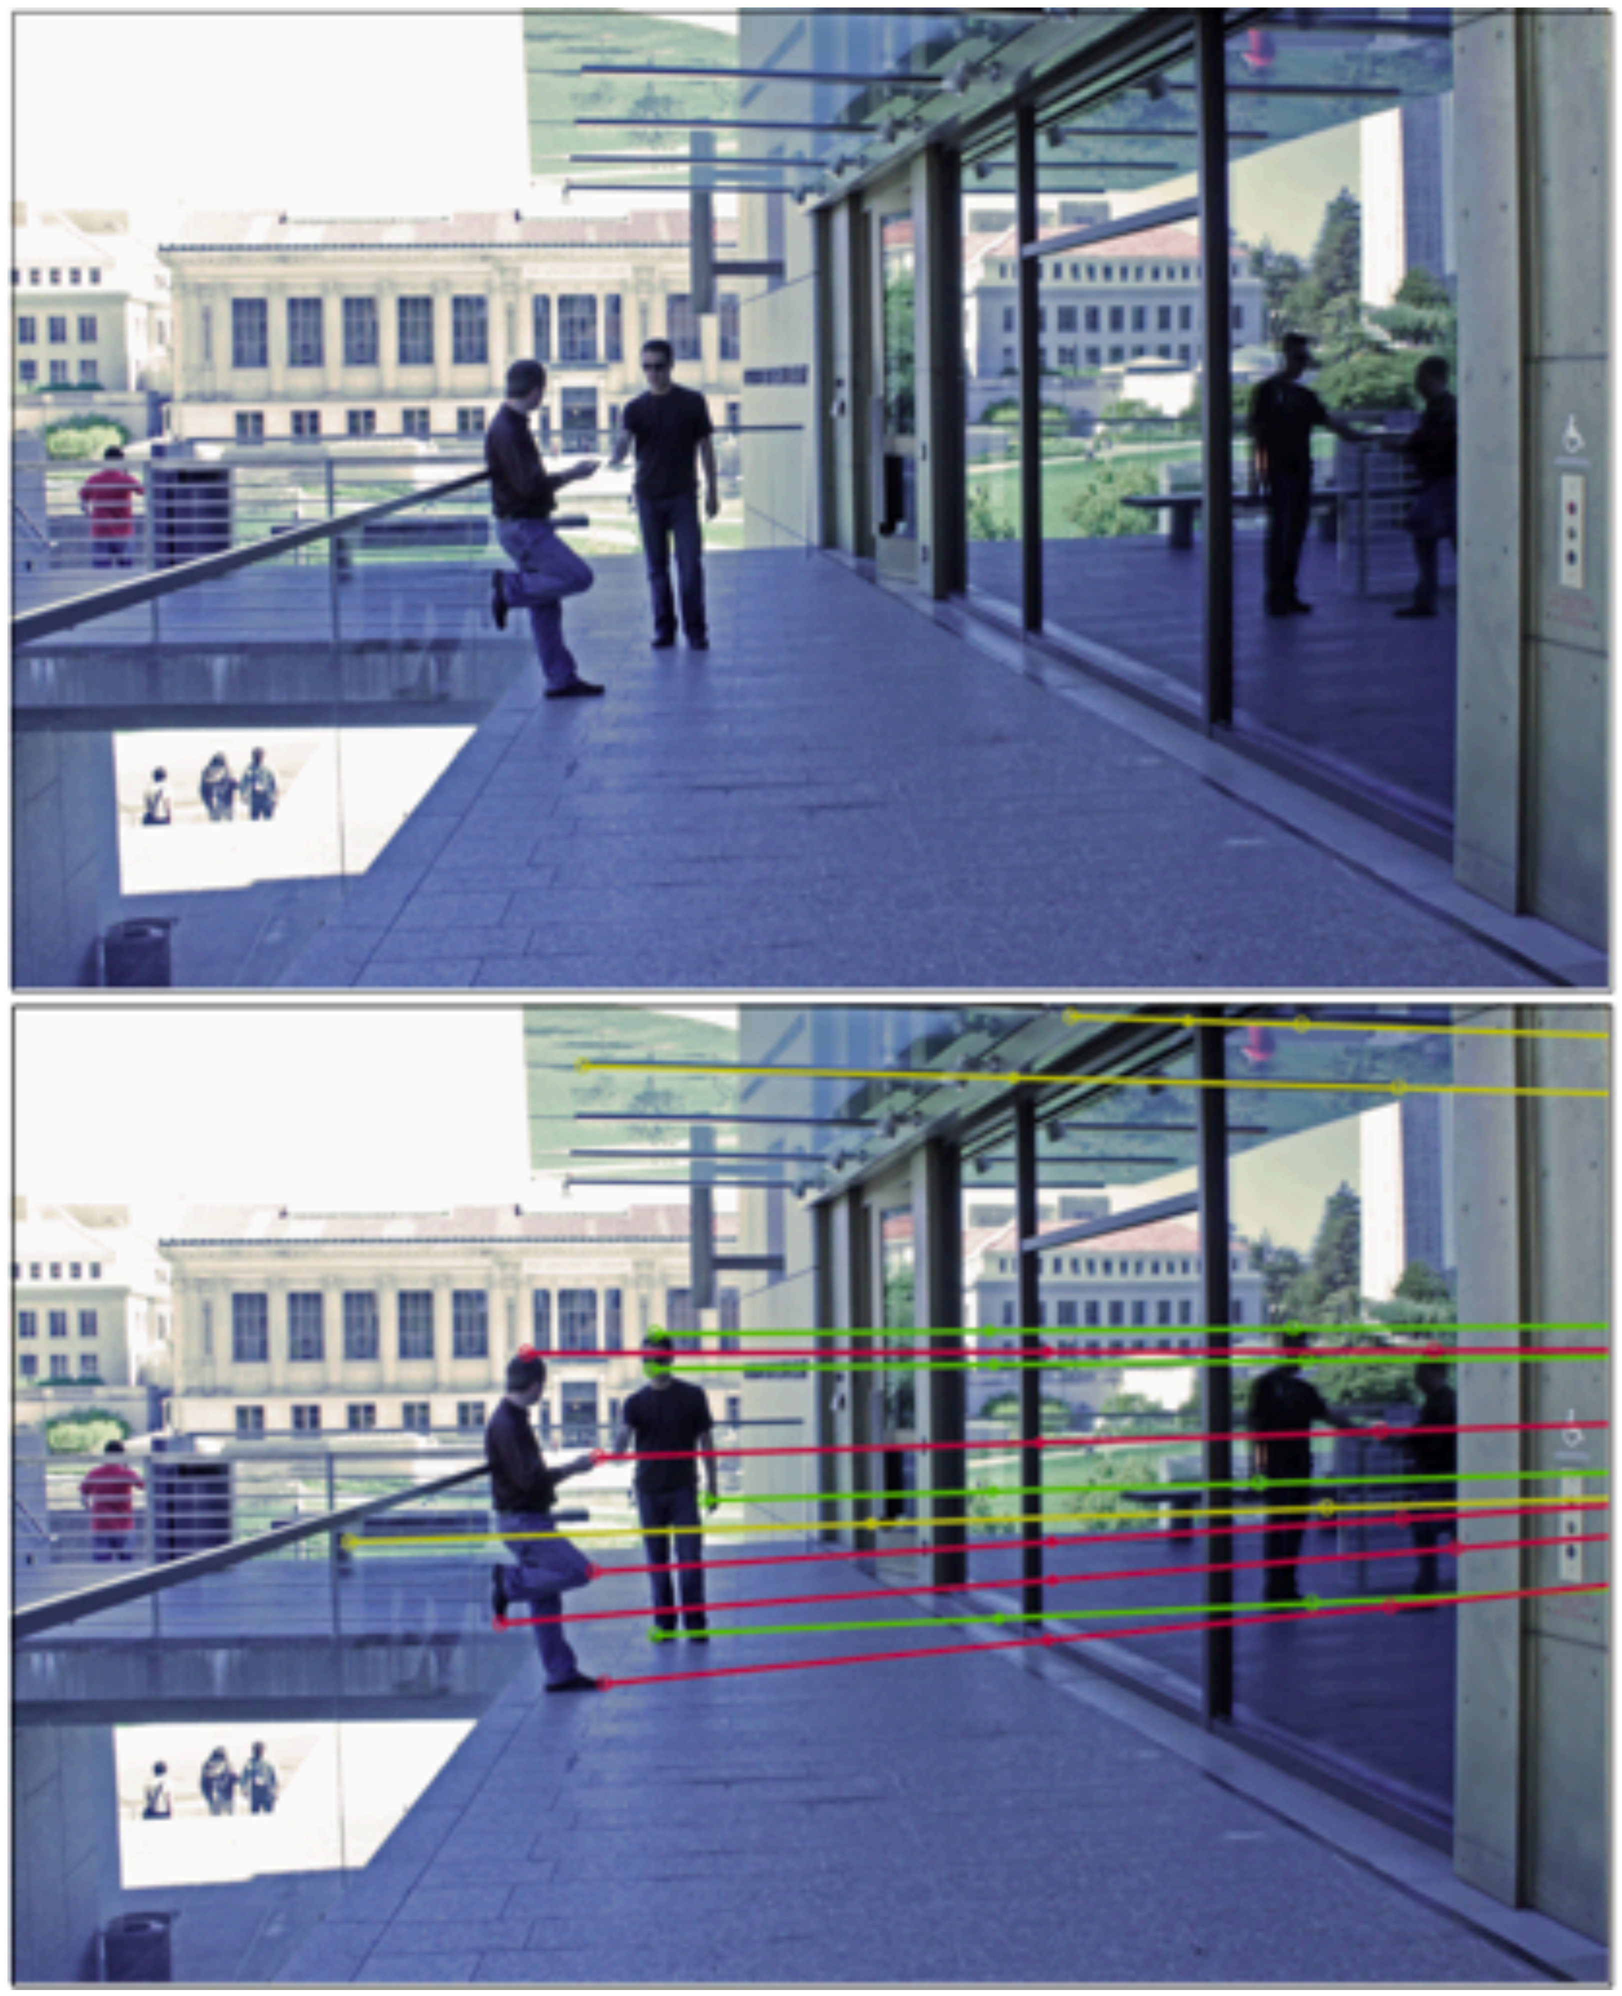

Kee et al, "Exposing Photo Manipulation with Inconsistent Shadows" (Transactions on Graphics 2013)

3D Fabrication

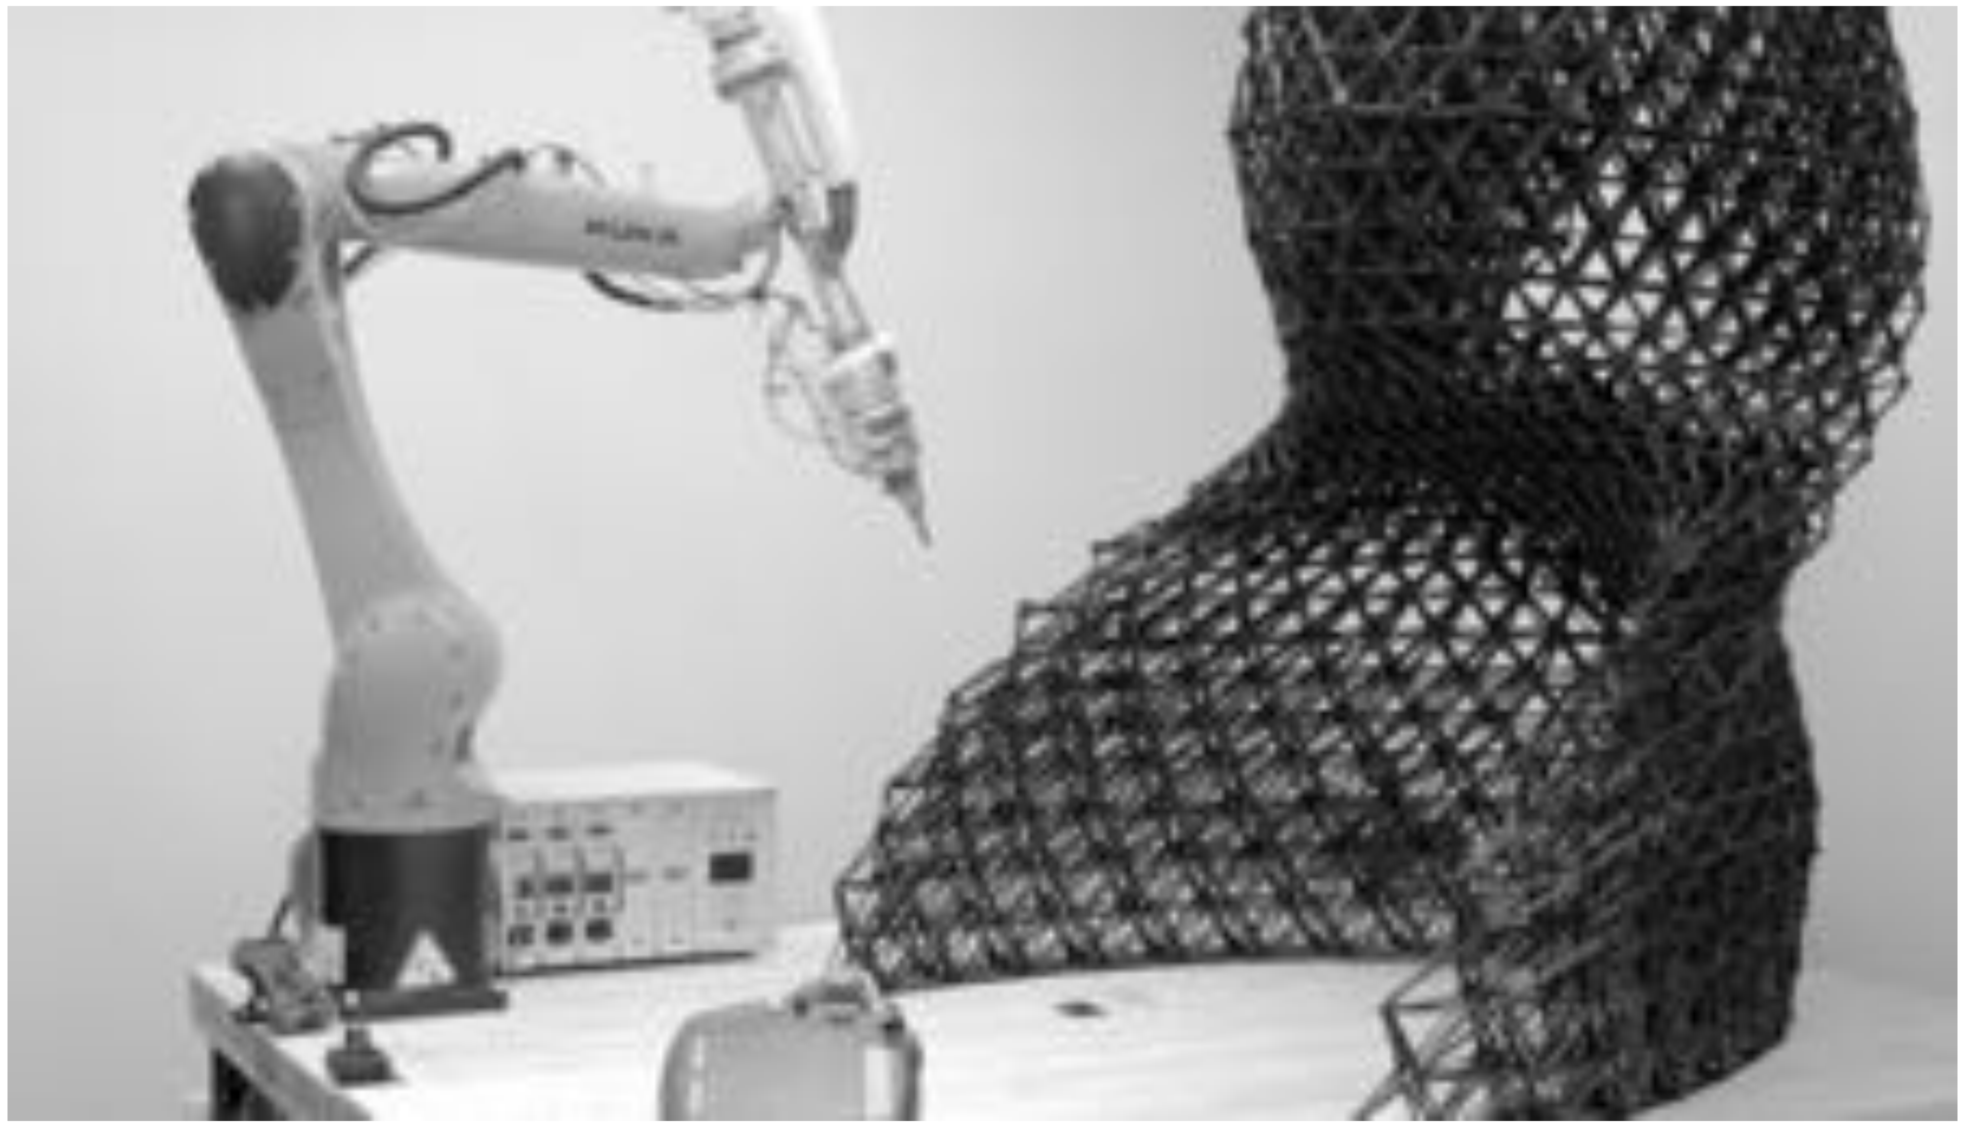

And much much more!

- Check out SIGGRAPH for latest & greatest graphics stuff
- You can (start to) understand these papers!

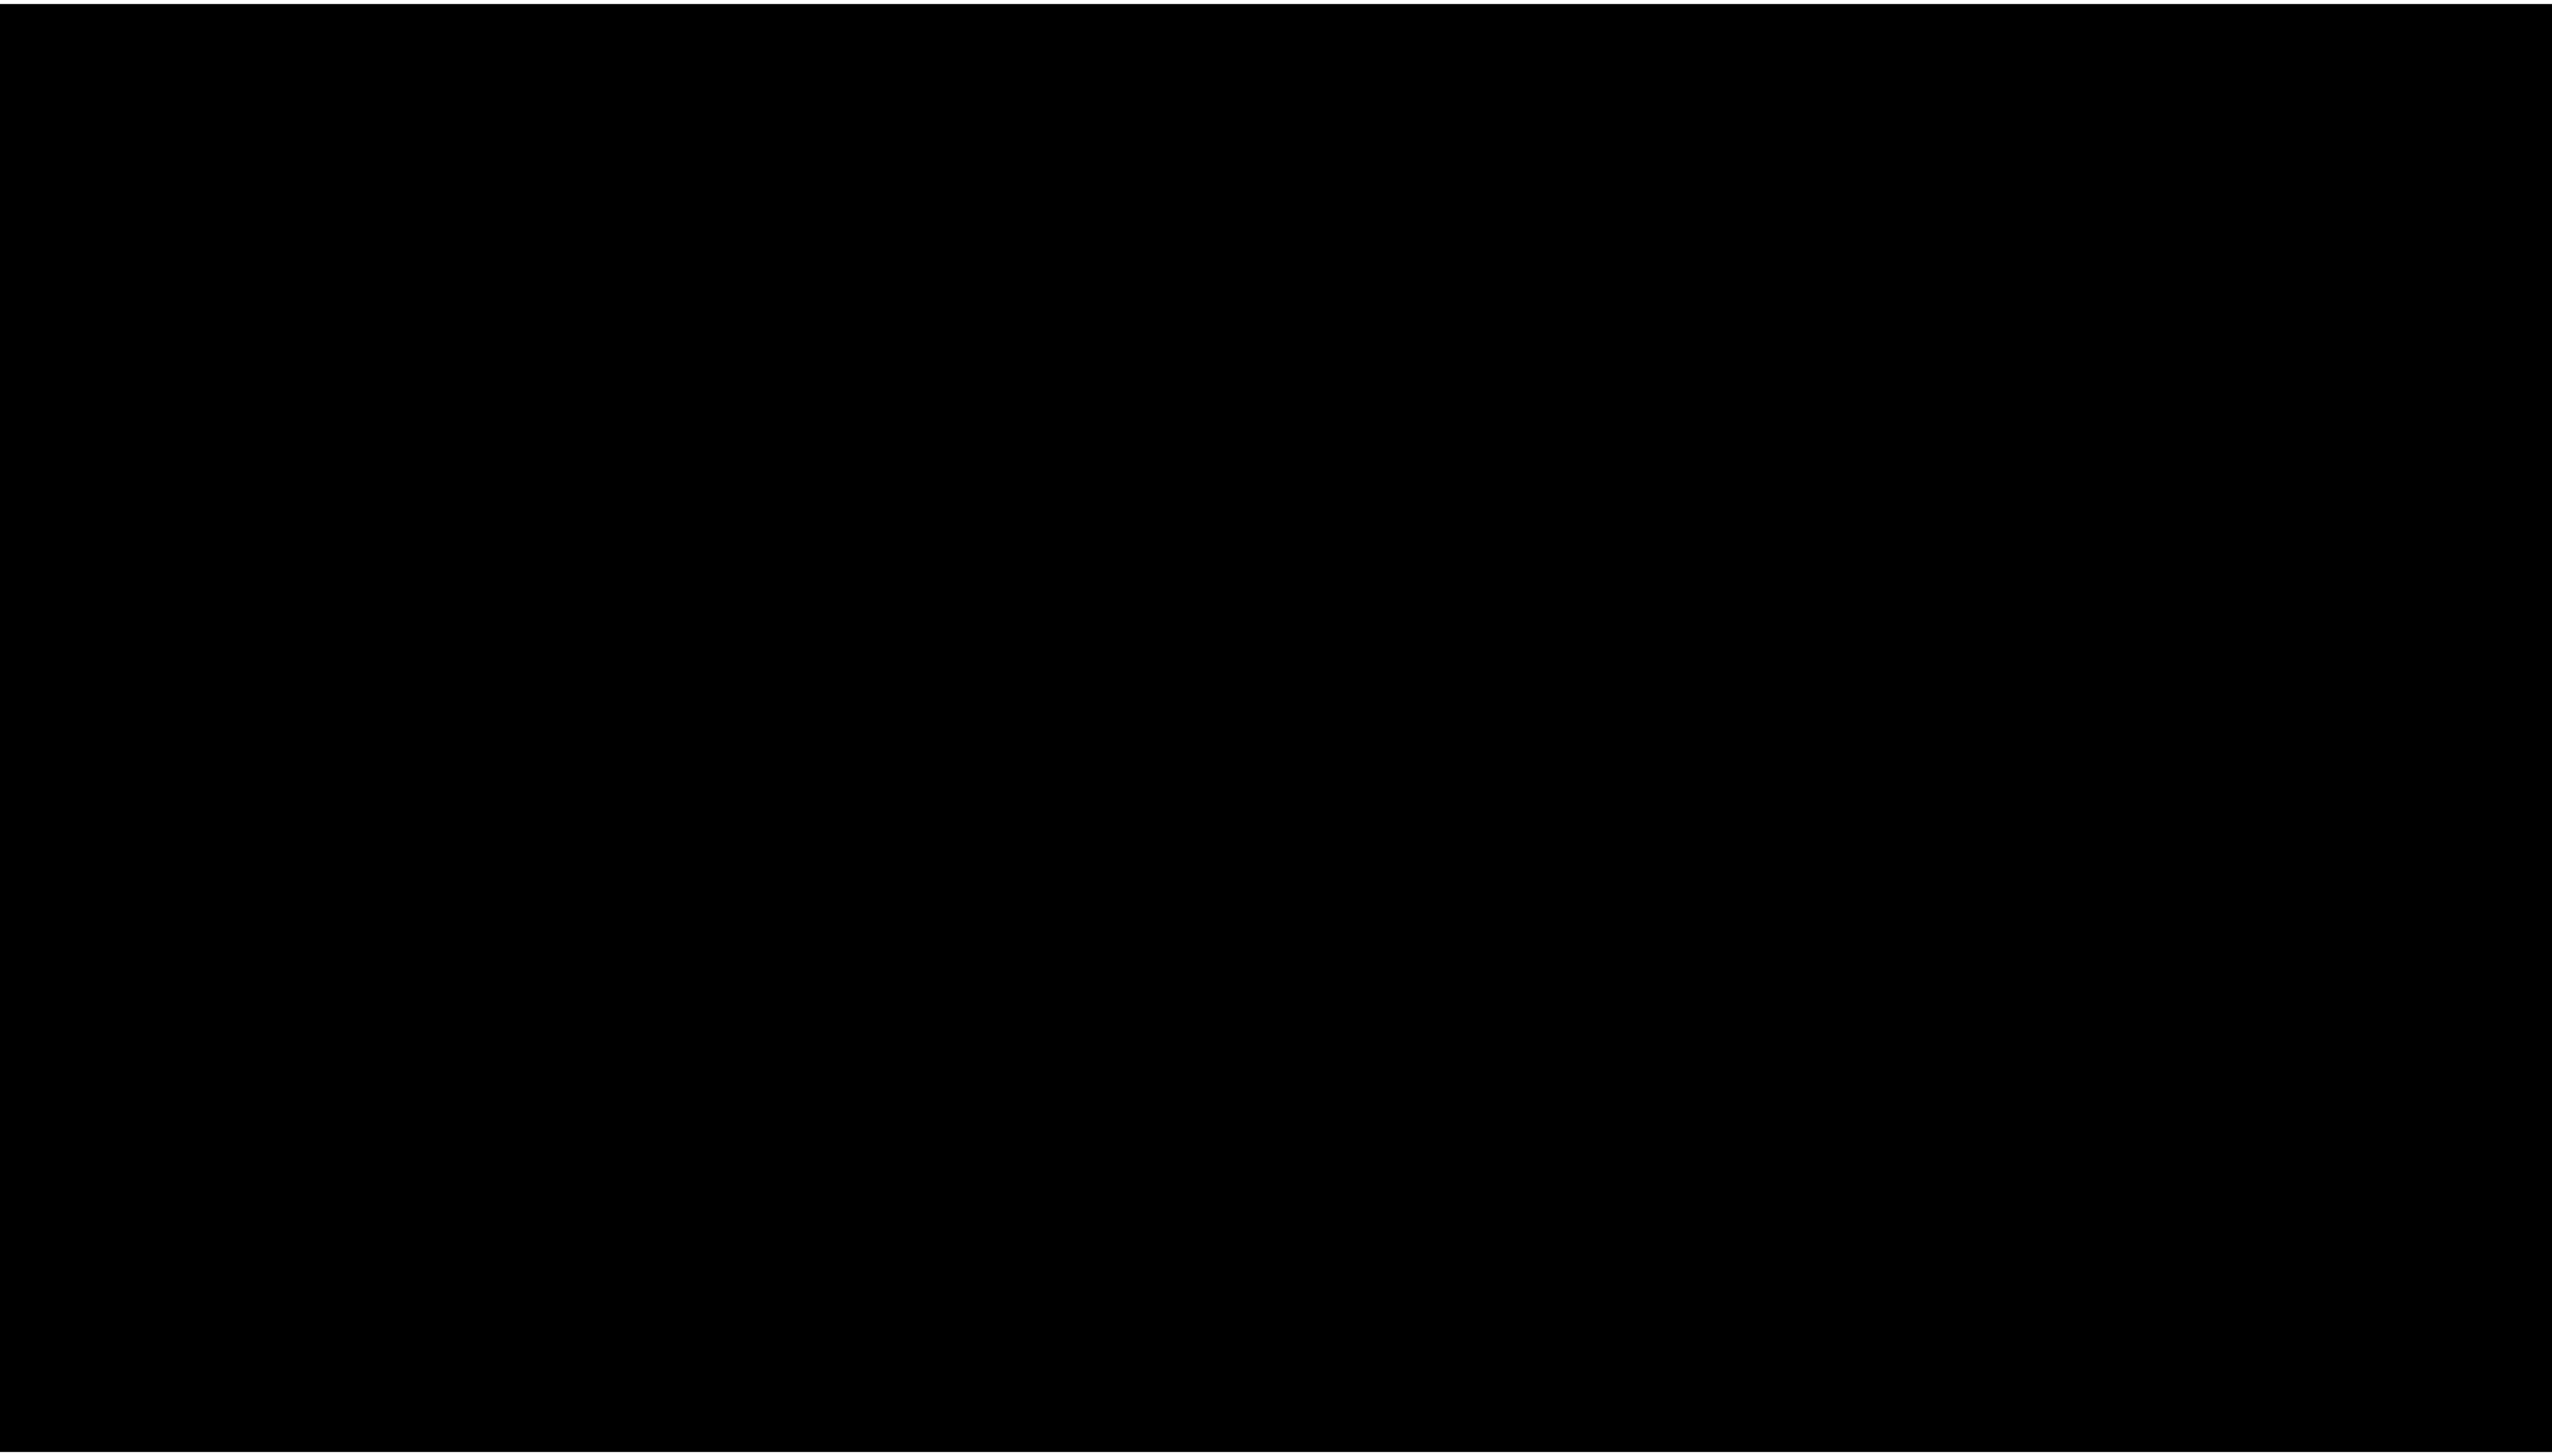

Course wrap up

Final Exam

- **Same format as midterm**
  - **In-class, longer exam**
  - **Can bring one "sticky note" of information**
  - **Cumulative (anything in the semester is fair game)**

Student project demos!

Other Cool Graphics-Related Courses (Fall ’19)

- **Computational Photography 15-463/663/862**
  - **Ioannis Gkioulekas**
- **Computer Game Programming 15-466/666**
  - **Jim McCann**

15-463/15-663/15-862 Computational Photography

Learn about scientific and unconventional cameras – and build your own!

cameras that capture video at the speed of light

cameras that measure depth in real time

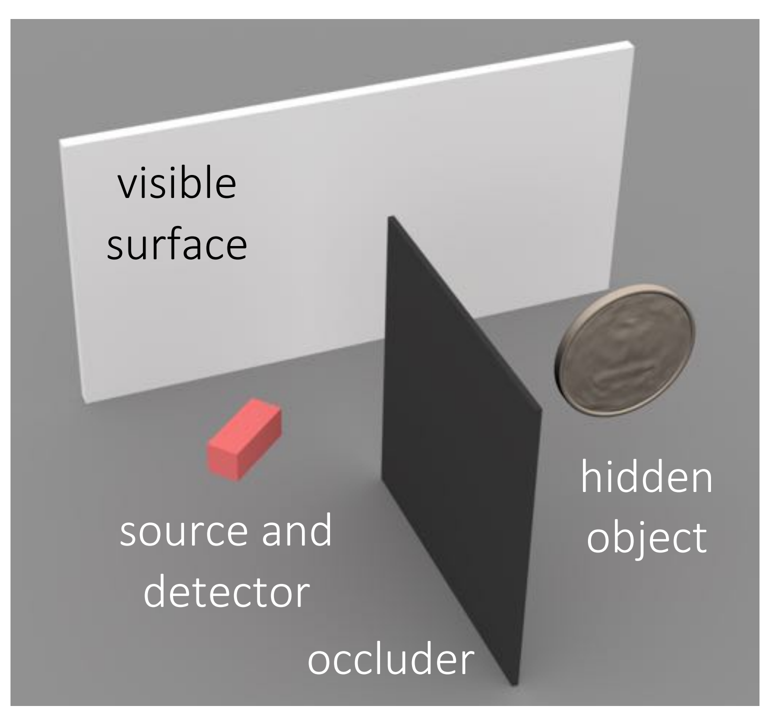

cameras that see around corners

cameras that measure entire lightfields

http://graphics.cs.cmu.edu/courses/15-463/

Other Cool Graphics-Related Courses (Sp '20)

- Technical Animation 15-464/664
  - Nancy Pollard
- Animation Art and Technology 15-465/60-414
  - Jessica Hodgins and James Duesing
- Discrete Differential Geometry 15-458/858
  - Keenan Crane
- Hands: Design/Control for Dexterous Manipulation 16-848
  - Nancy Pollard

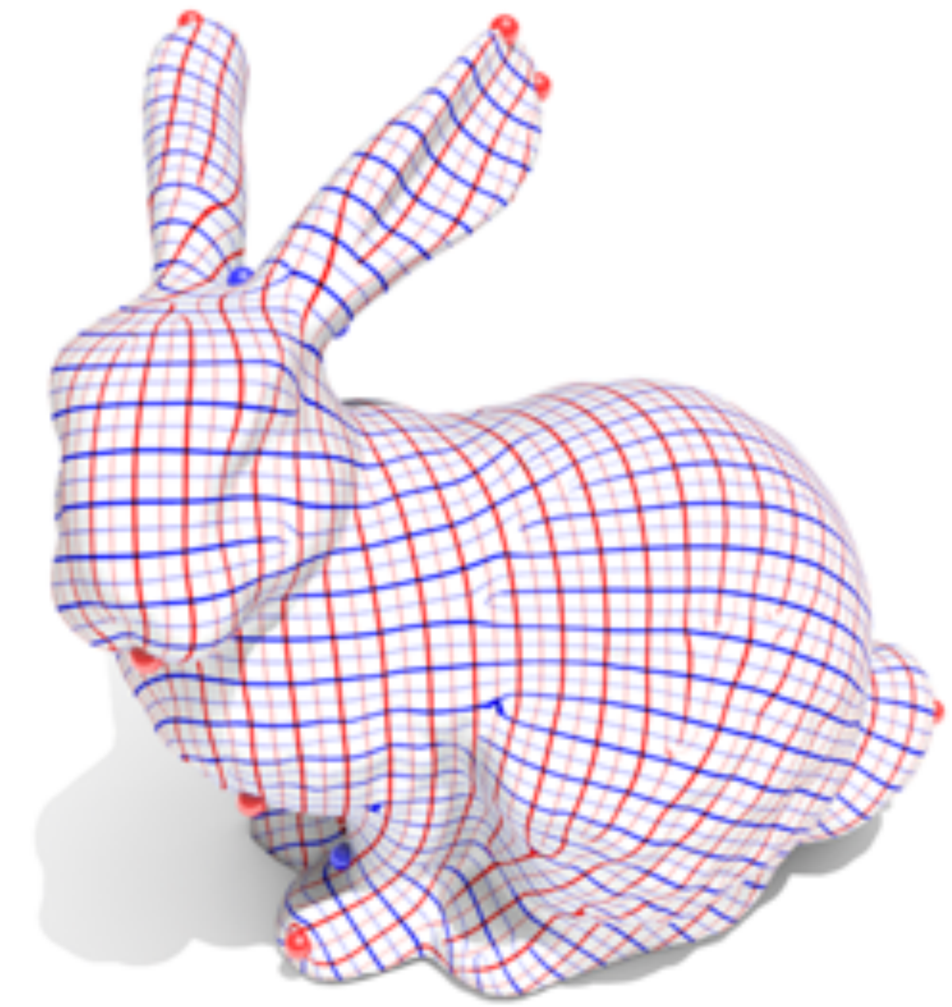

TAs and independent study!

- 15-462 next semester (<u>and next Spring</u>) is looking for TAs!
  - Email us if interested, and we'll direct you to Profs. McCann and O'Toole

- Students that did well in 462 have a great foundation for moving on to independent study or research in graphics
  - independent study
  - senior thesis
  - improve Scotty3D :-)
  - …
  - Come talk to us / email us

Thanks for being a great class!

See you at the final! (study hard, but don't stress too much)

Credit: Inside Out (Pixar)