Light field photography and videography

Marc Levoy

Computer Science Department
Stanford University

List of projects

- high performance imaging
  using large camera arrays
- light field photography
  using a handheld plenoptic camera
- dual photography

High performance imaging using large camera arrays

Bennett Wilburn, Neel Joshi, Vaibhav Vaish, Eino-Ville Talvala, Emilio Antunez, Adam Barth, Andrew Adams, Mark Horowitz, Marc Levoy

(Proc. SIGGRAPH 2005)

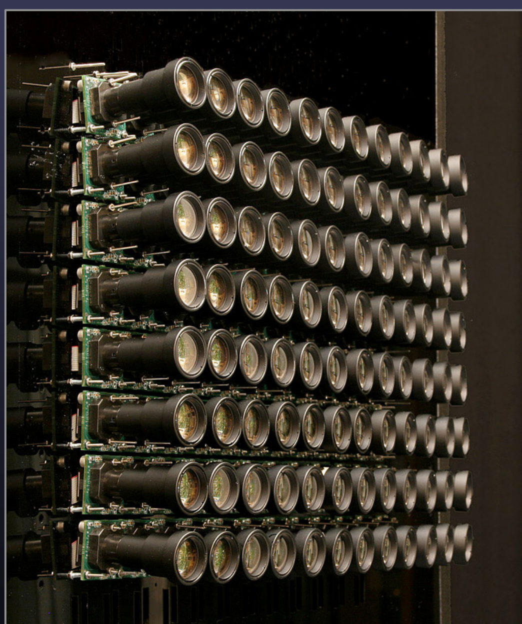

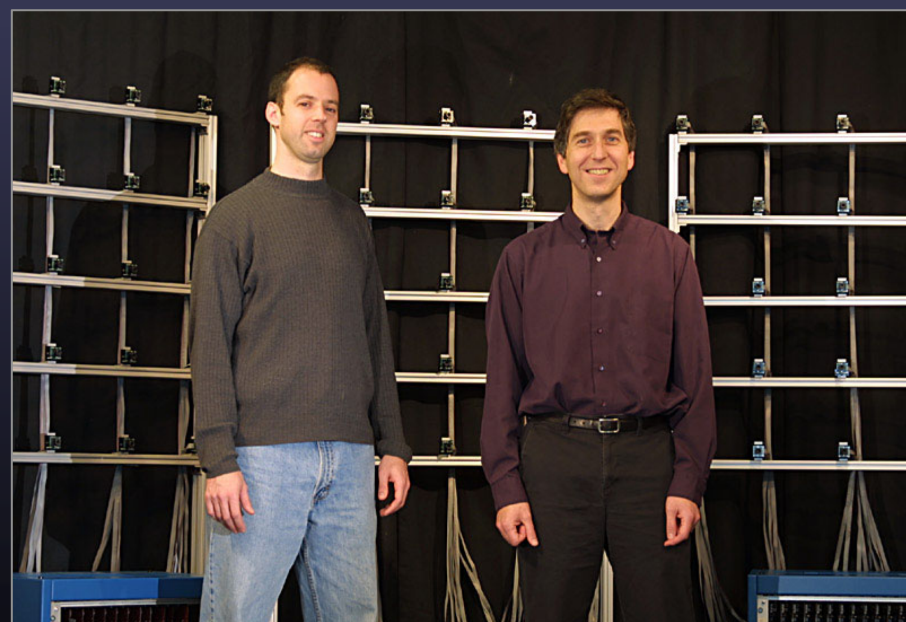

Stanford multi-camera array

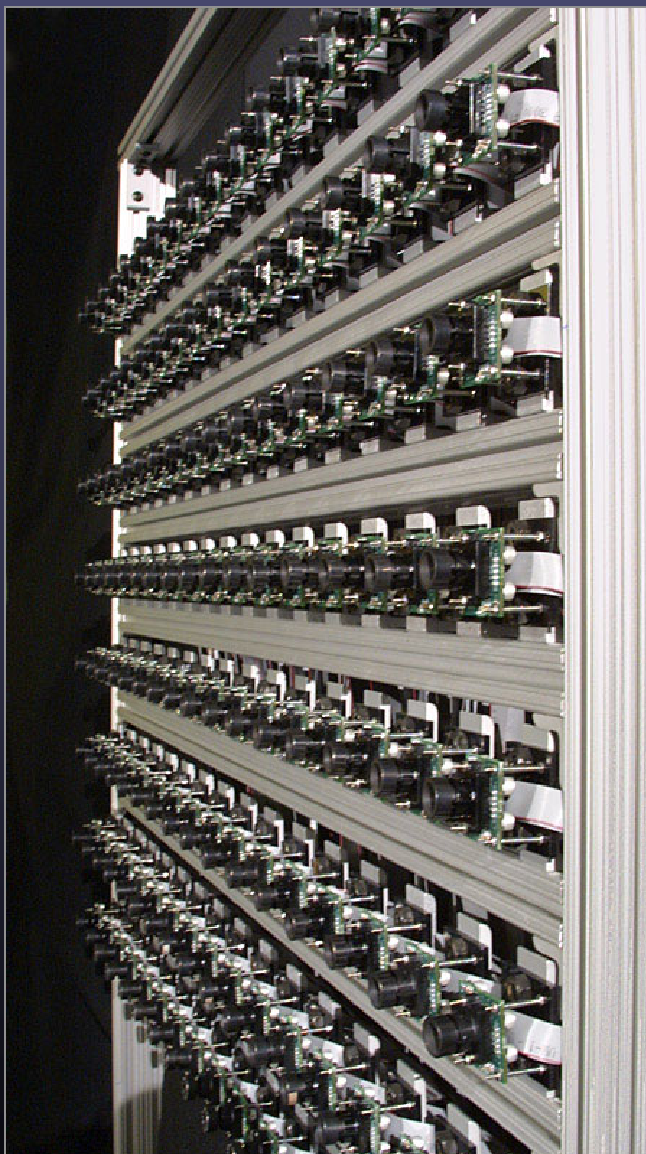

- 640 × 480 pixels ×
  30 fps × 128 cameras

- synchronized timing
- continuous streaming
- flexible arrangement

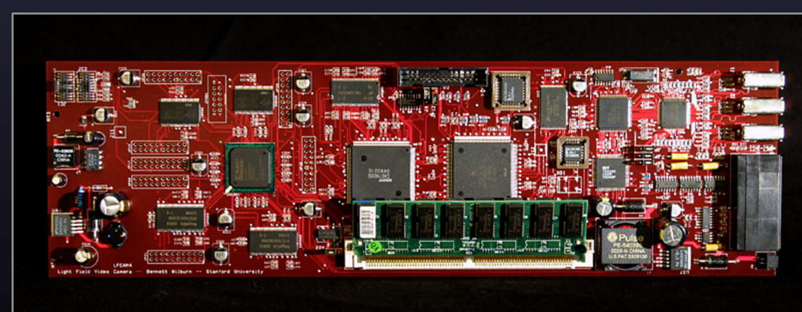

Ways to use large camera arrays

- widely spaced → light field capture
- tightly packed → high-performance imaging
- intermediate spacing → synthetic aperture photography

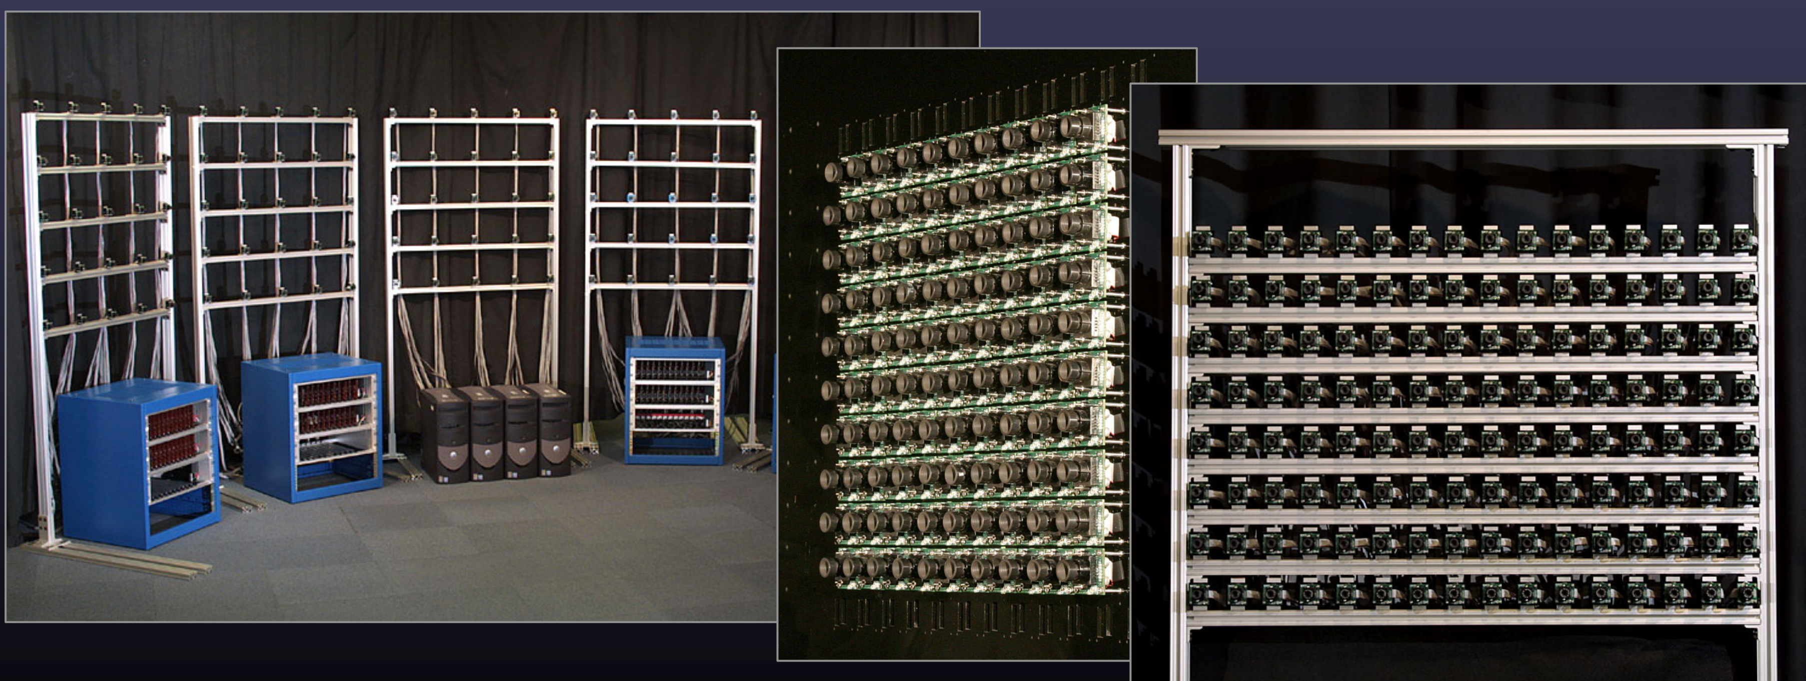

Intermediate camera spacing: synthetic aperture photography

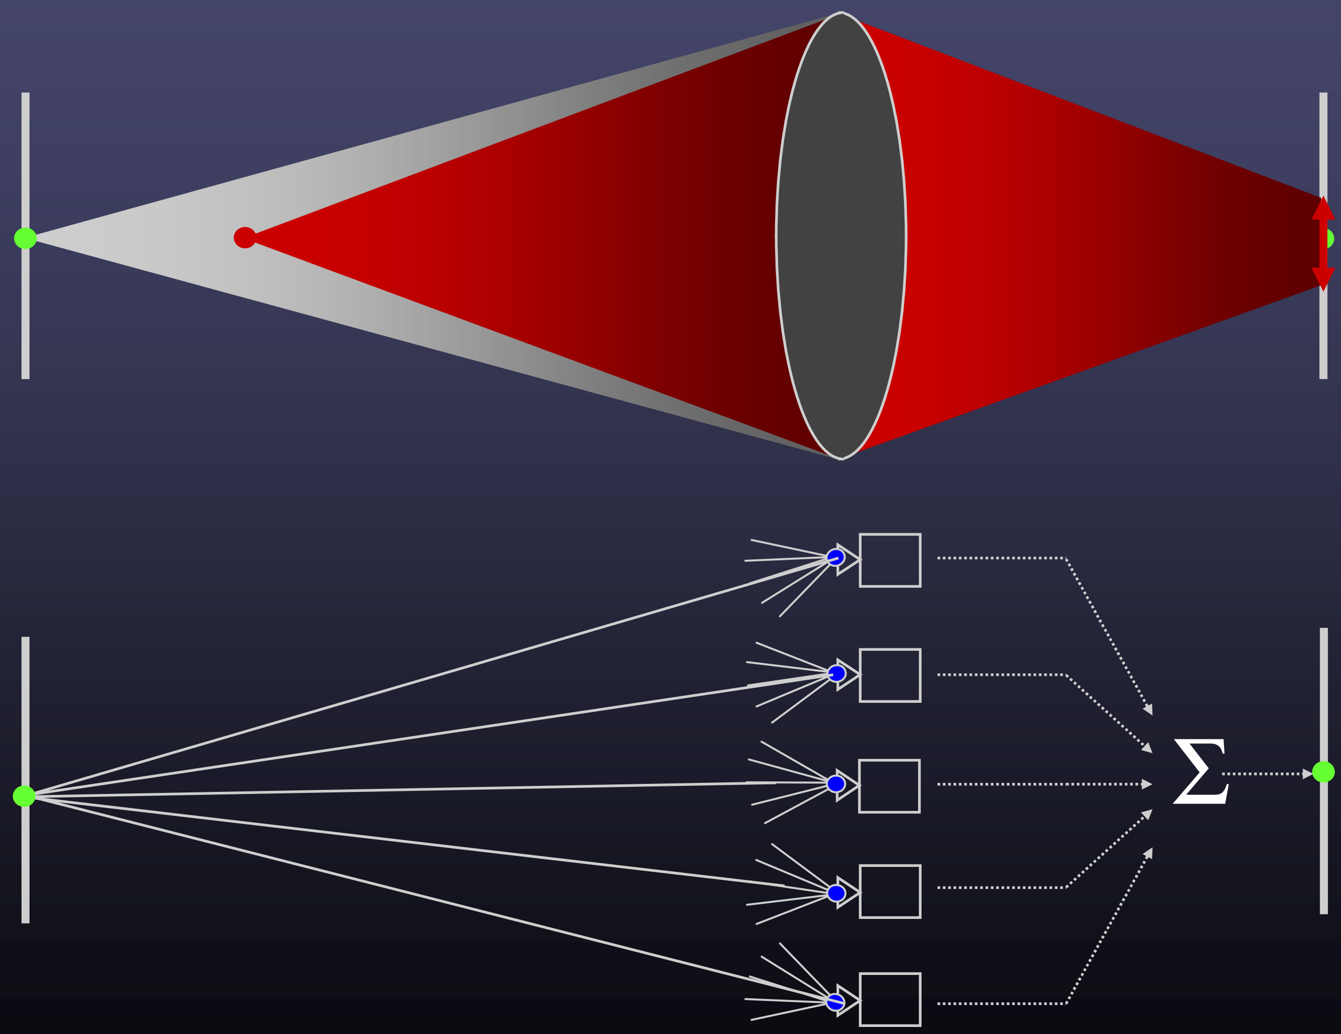

Example using 45 cameras
[Vaish CVPR 2004]

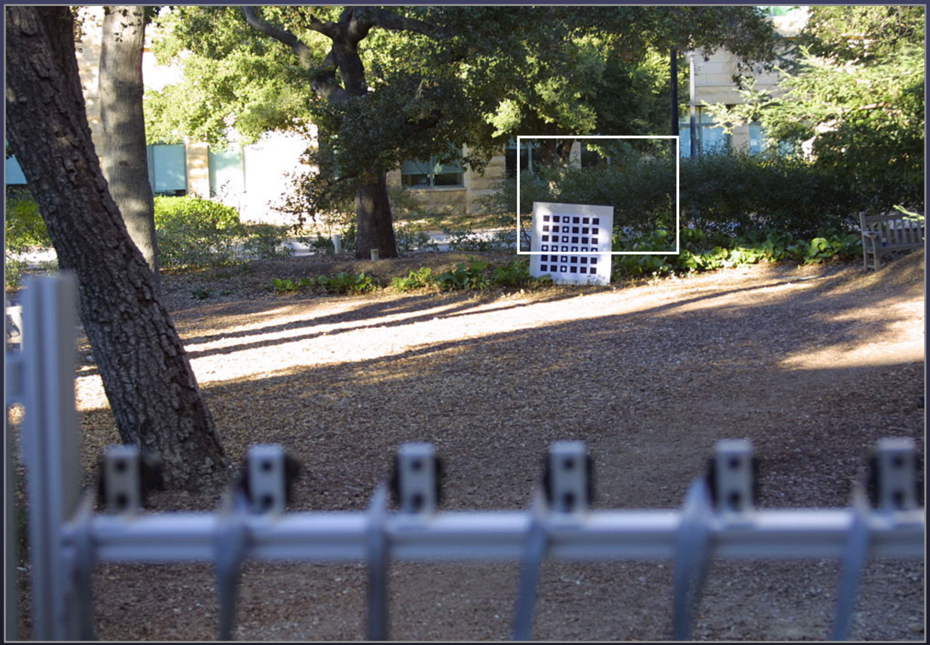


Tiled camera array

Can we match the image quality of a cinema camera?

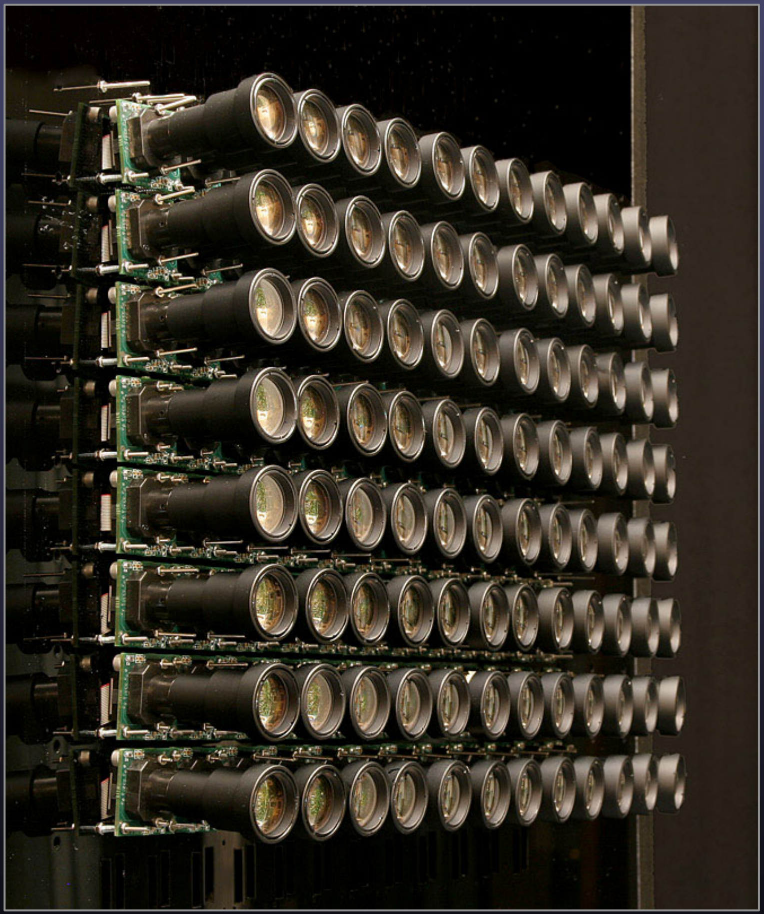

- world's largest video camera
- no parallax for distant objects
- poor lenses limit image quality
- seamless mosaicing isn't hard

Tiled panoramic image
(before geometric or color calibration)

Tiled panoramic image (after calibration and blending)

Tiled camera array

Can we match the image quality of a cinema camera?

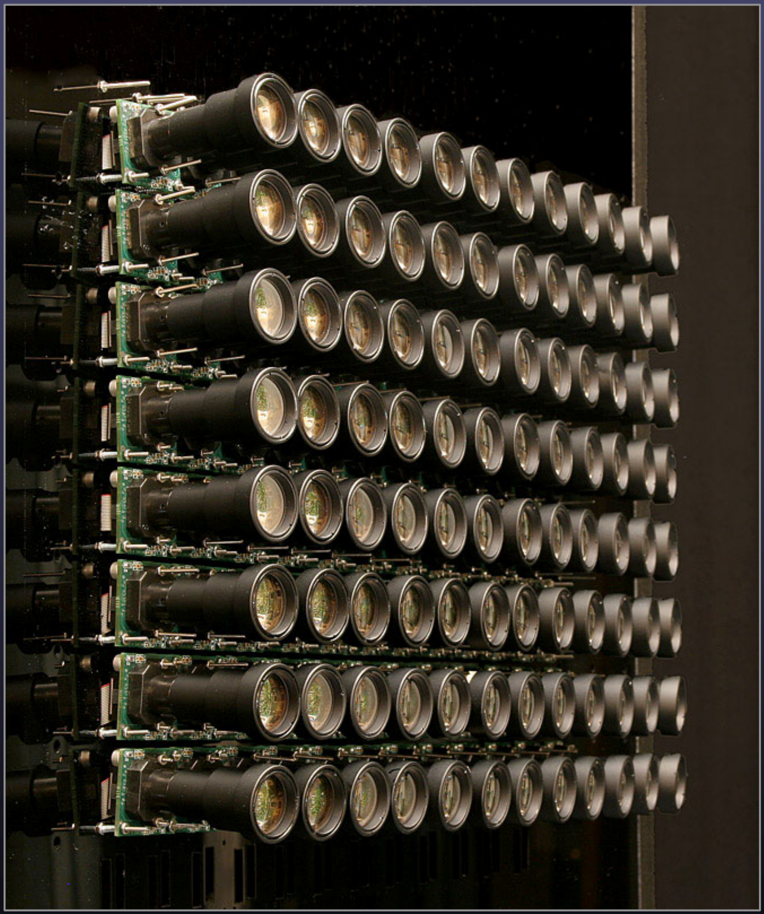

- world's largest video camera
- no parallax for distant objects
- poor lenses limit image quality
- seamless mosaicing isn't hard
- per-camera exposure metering
- HDR within and between tiles

same exposure
in all cameras

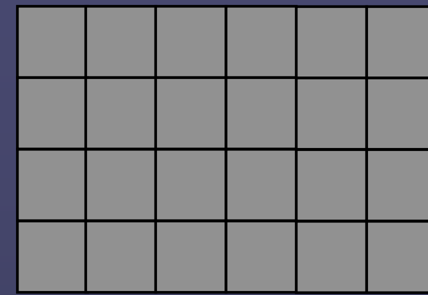

individually
metered

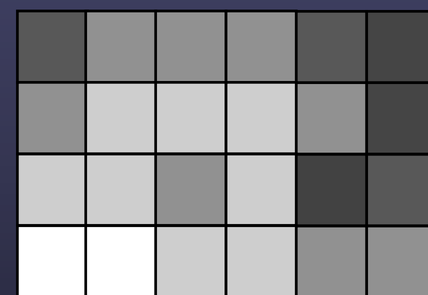

checkerboard
of exposures

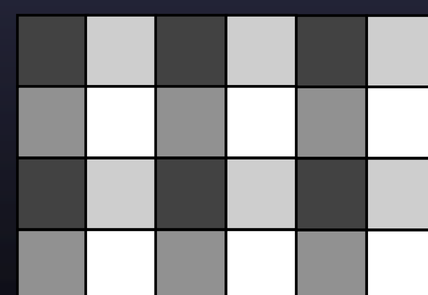

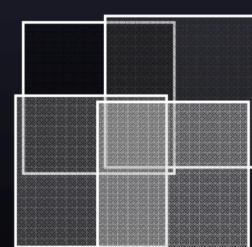

High-performance photography as multi-dimensional sampling

- spatial resolution
- field of view
- frame rate
- dynamic range
- bits of precision
- depth of field
- focus setting
- color sensitivity

Spacetime aperture shaping

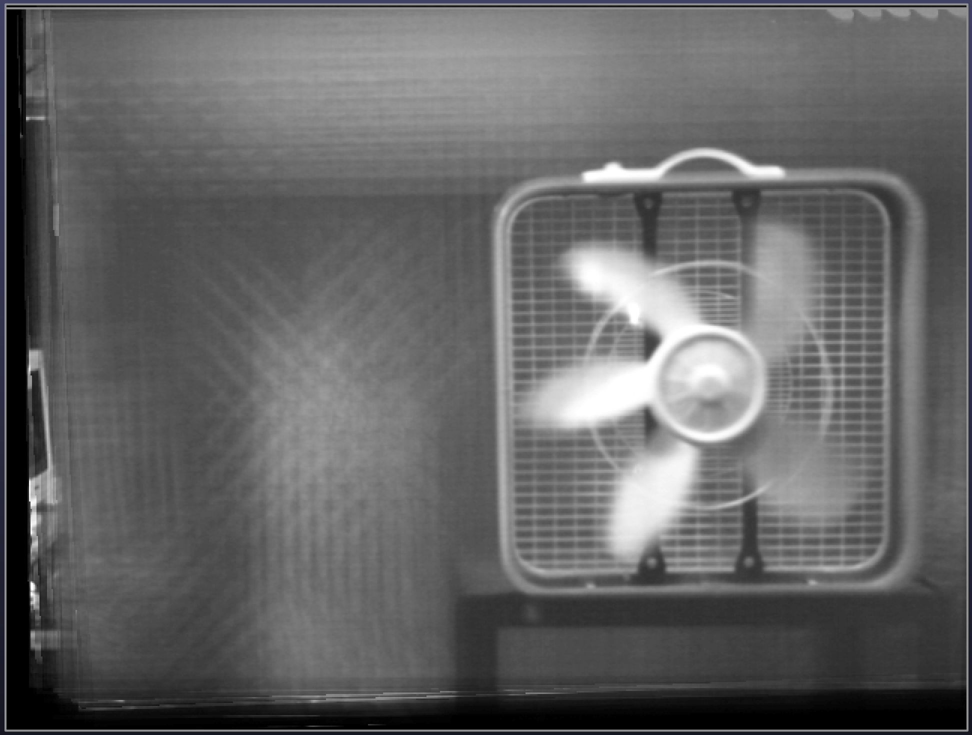

- shorten exposure time to freeze motion → dark
- stretch contrast to restore level → noisy
- increase (synthetic) aperture to capture more light → decreases depth of field

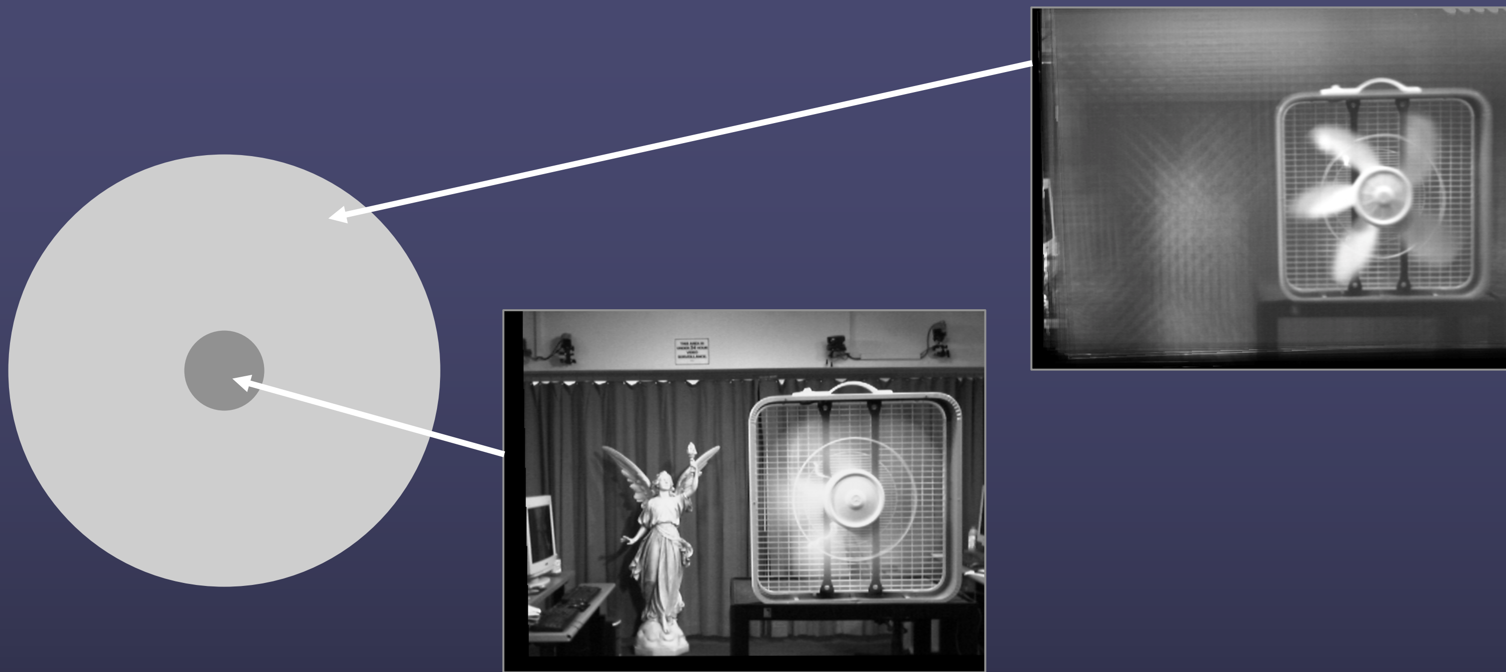

- center of aperture: few cameras, long exposure → high depth of field, low noise, but action is blurred
- periphery of aperture: many cameras, short exposure → freezes action, low noise, but low depth of field

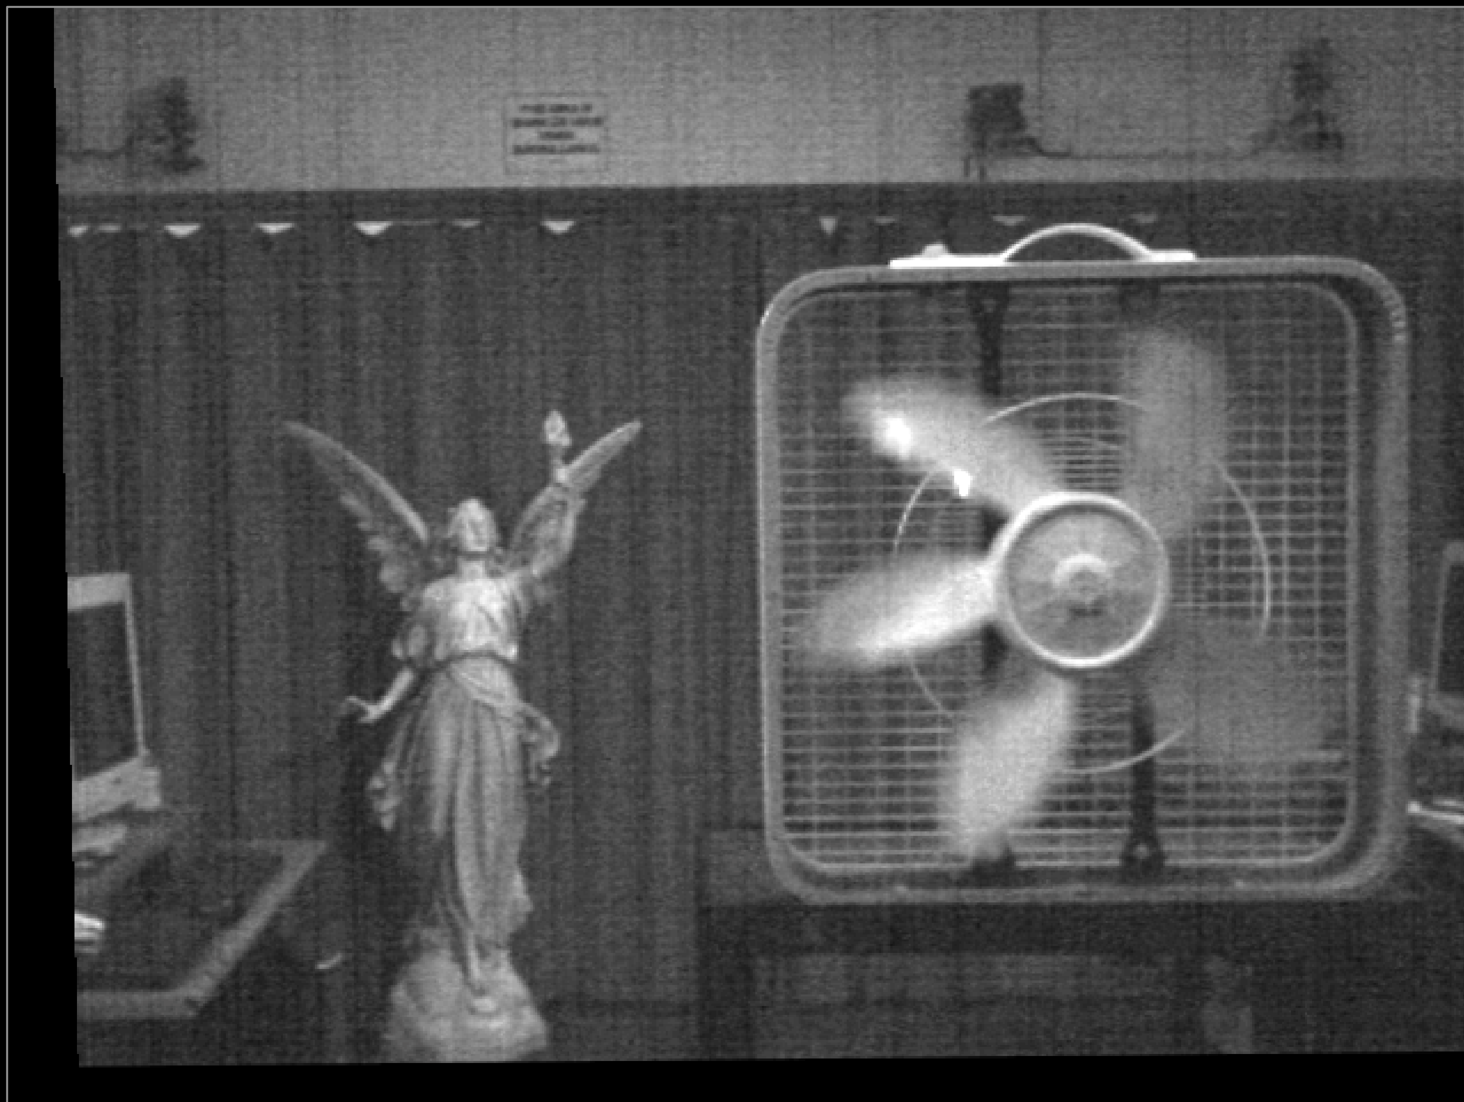

Light field photography using a handheld plenoptic camera

Ren Ng, Marc Levoy, Mathieu Brédif,
Gene Duval, Mark Horowitz and Pat Hanrahan

(Proc. SIGGRAPH 2005
and TR 2005-02)

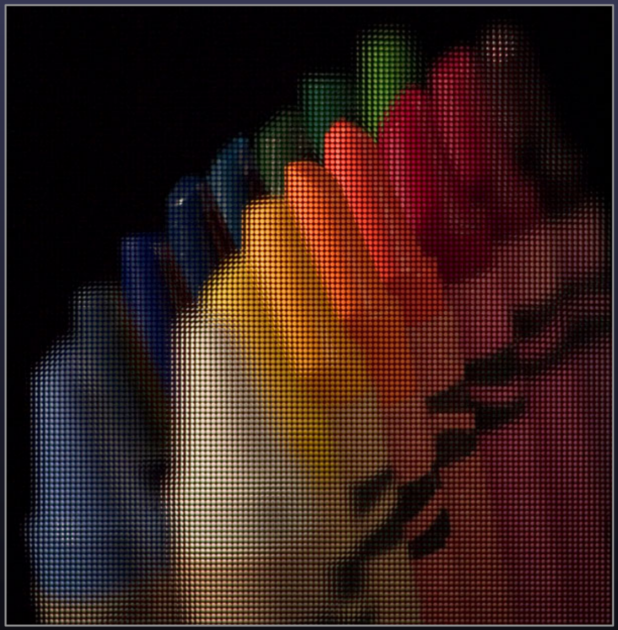

Conventional versus light field camera

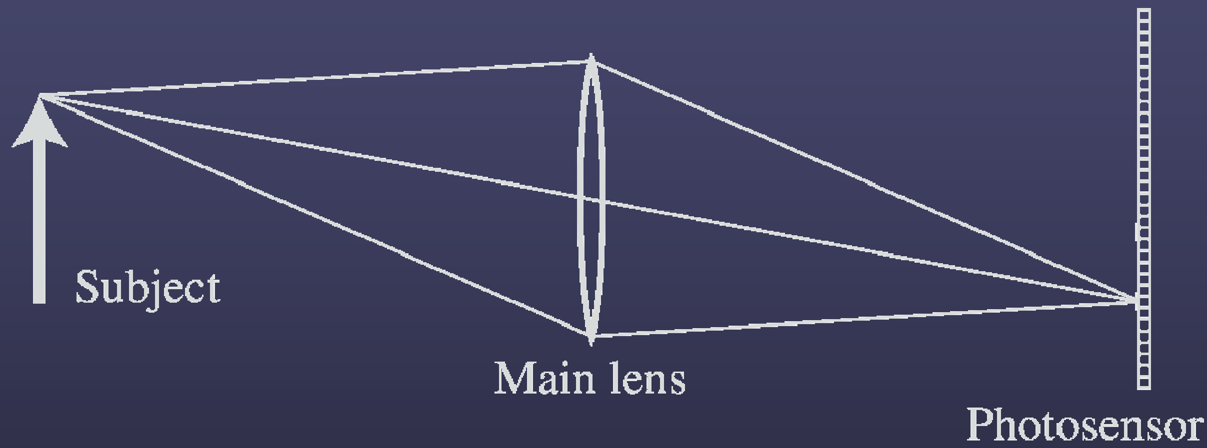

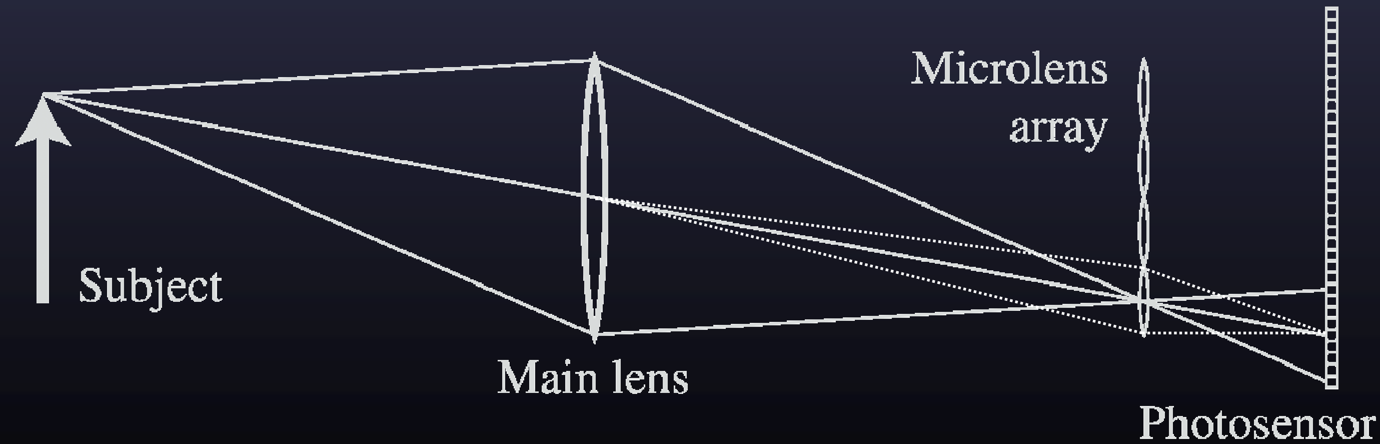

Conventional versus light field camera

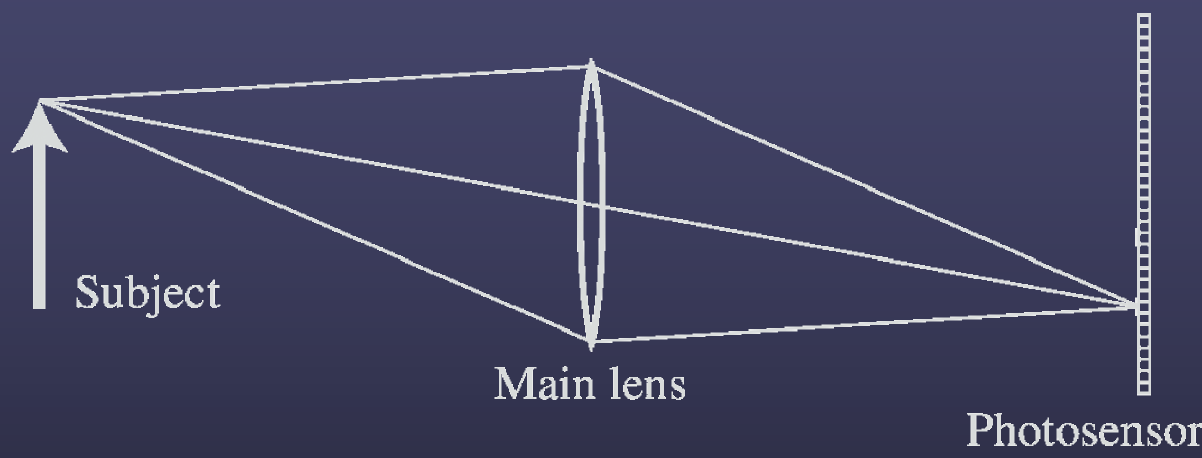

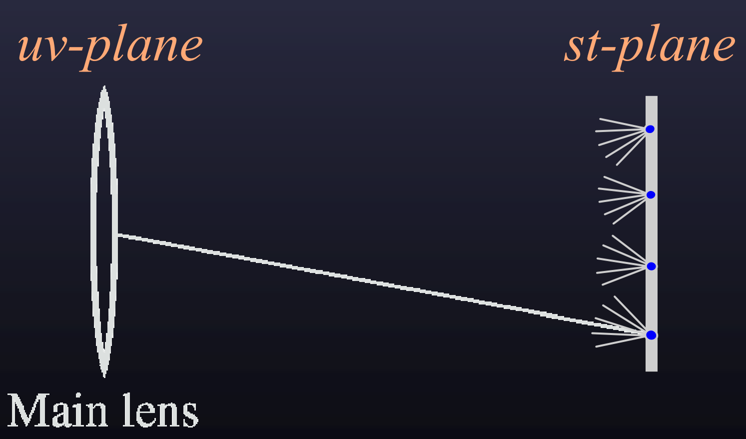

Conventional versus light field camera

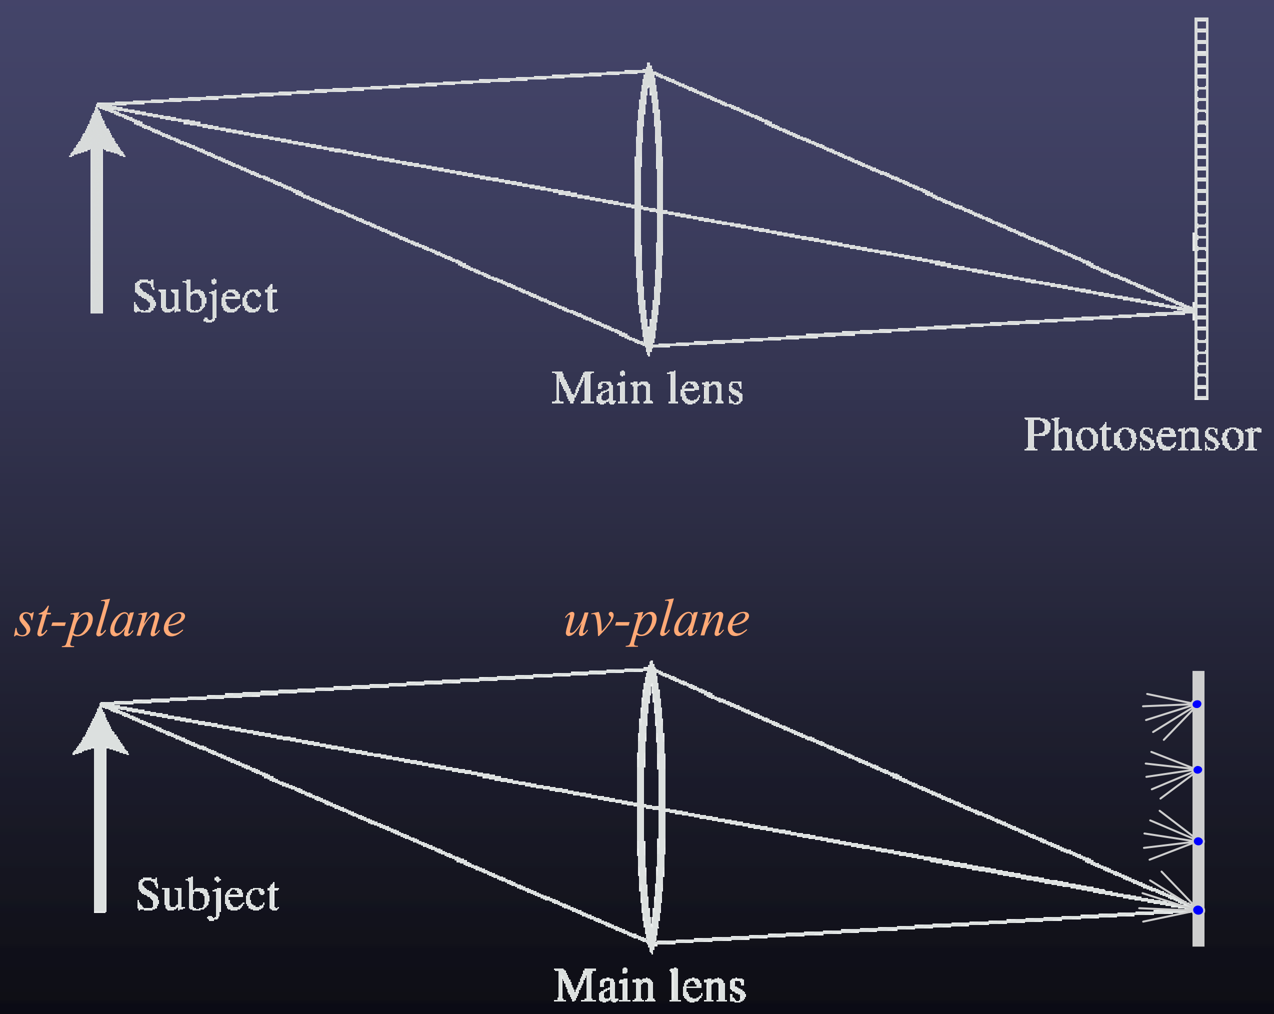

Prototype camera

Contax medium format camera

Kodak 16-megapixel sensor

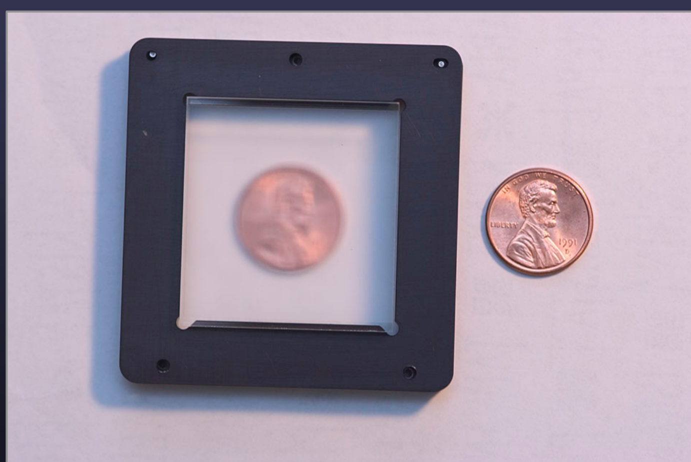

Adaptive Optics microlens array

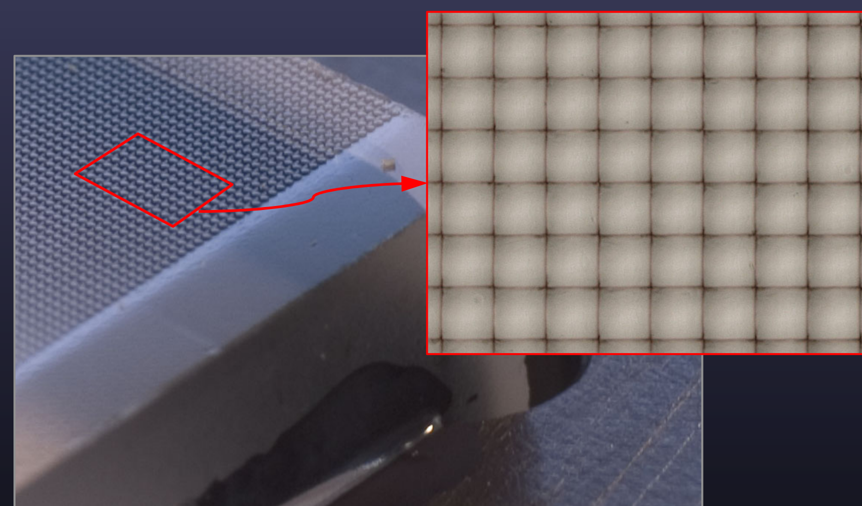

125μ square-sided microlenses

$$4000 \times 4000 \text{ pixels} \div 292 \times 292 \text{ lenses} = 14 \times 14 \text{ pixels per lens}$$

Mechanical design

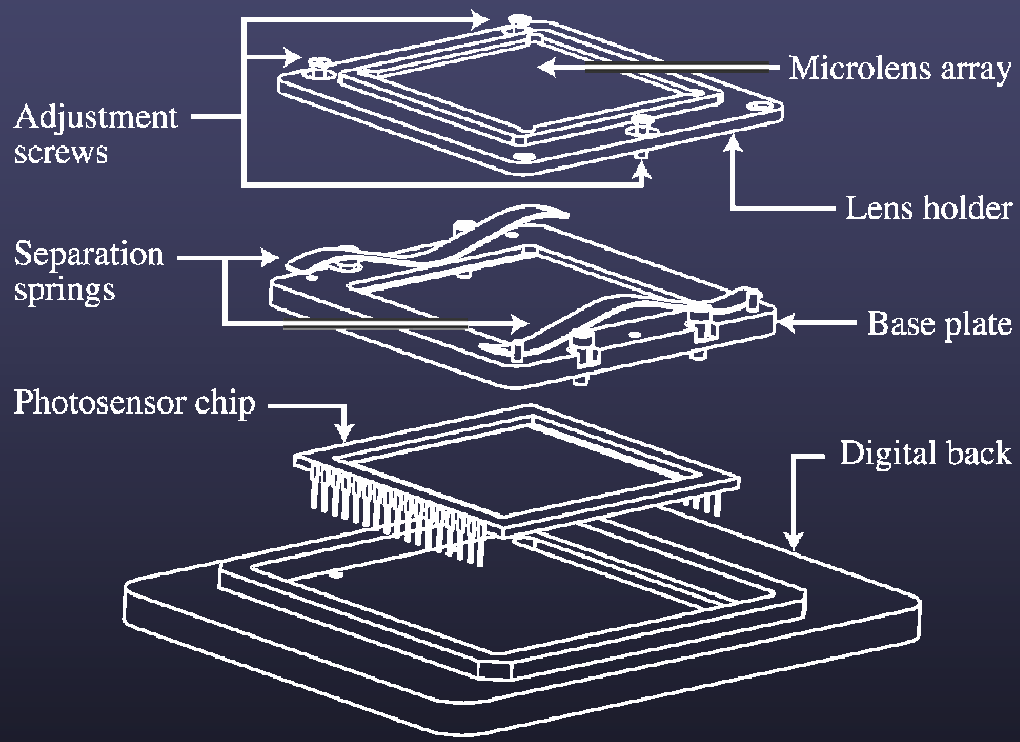

- microlenses float 500µ above sensor
- focused using 3 precision screws

Prior work

- integral photography
  - microlens array + film
  - application is autostereoscopic effect
- [Adelson 1992]
  - proposed this camera
  - built an optical bench prototype using relay lenses
  - application was stereo vision, not photography

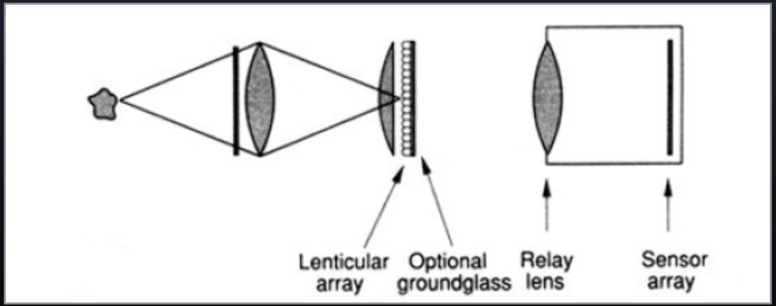

Digitally stopping-down

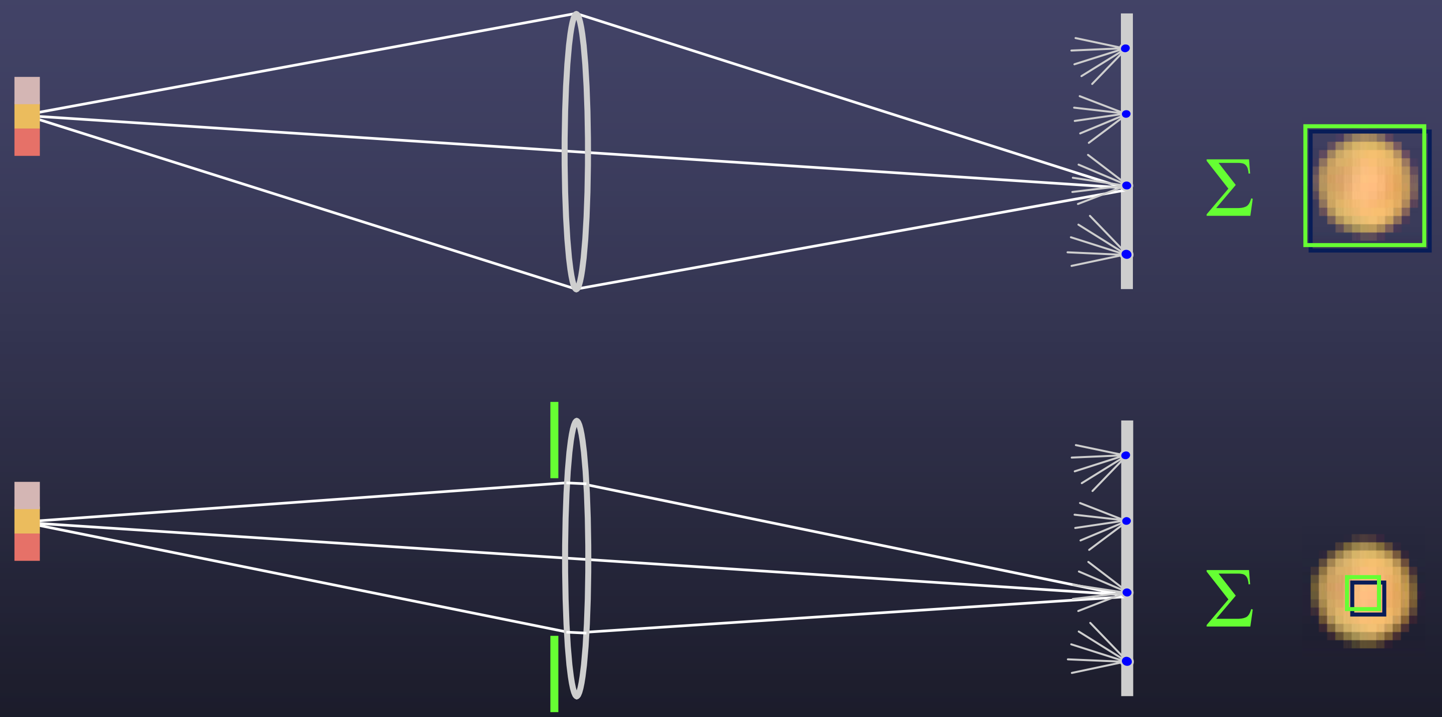

- stopping down = summing only the central portion of each microlens

Digital refocusing

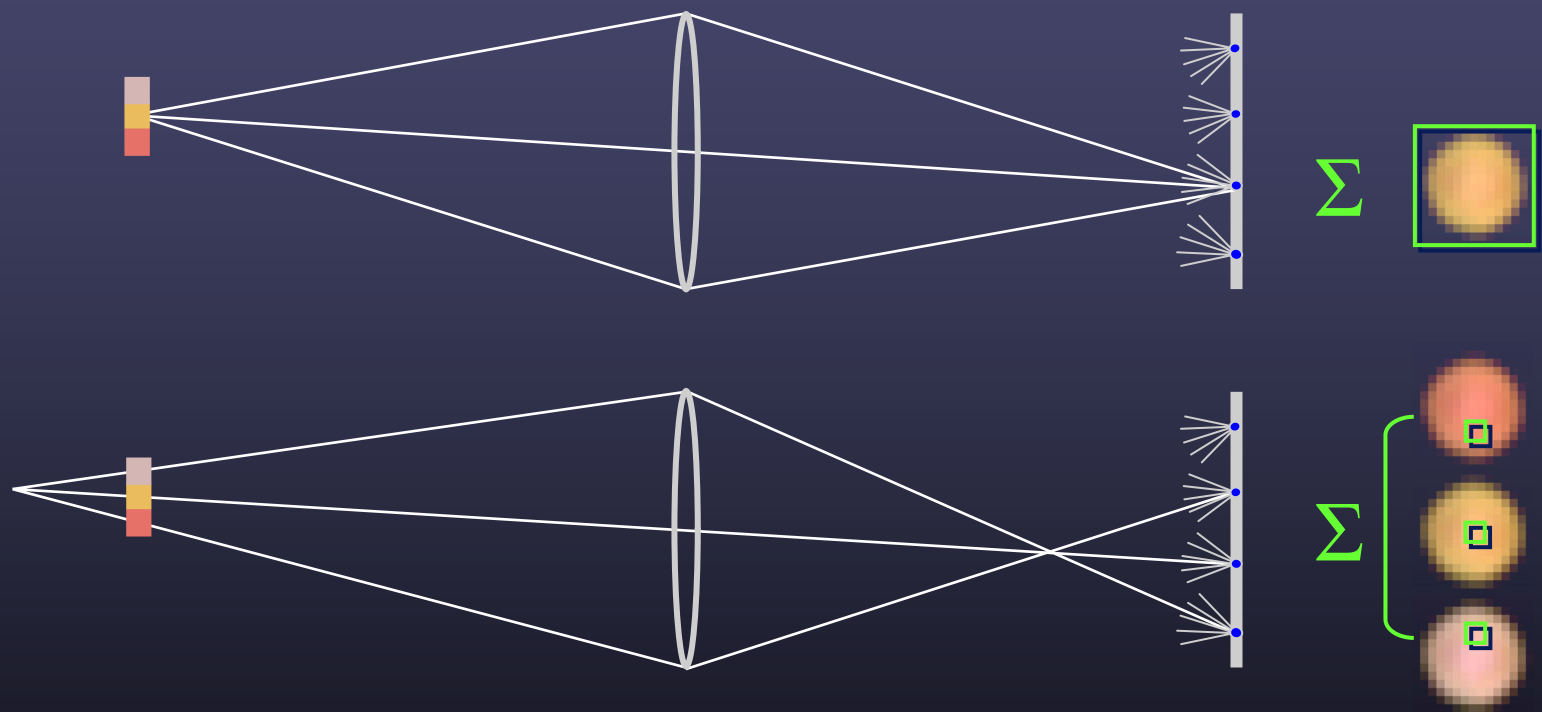

- refocusing = summing windows extracted from several microlenses

A digital refocusing theorem

- an f / N light field camera, with P × P pixels under each microlens, can produce views as sharp as an f / (N × P) conventional camera

– *or* –

- it can produce views with a shallow depth of field (f / N) focused anywhere within the depth of field of an f / (N × P) camera

Example of digital refocusing

Refocusing portraits

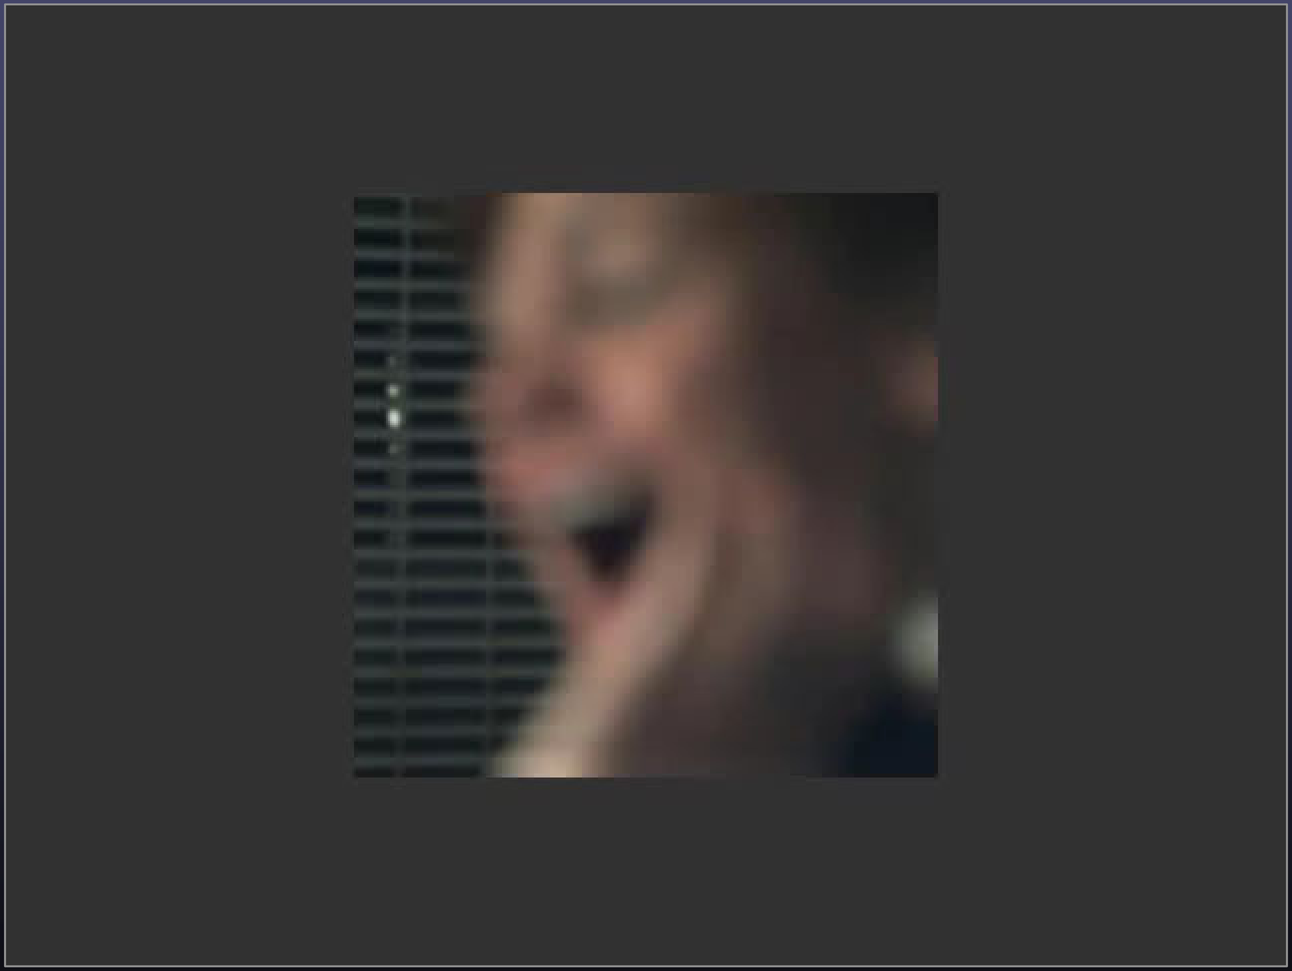

Action photography

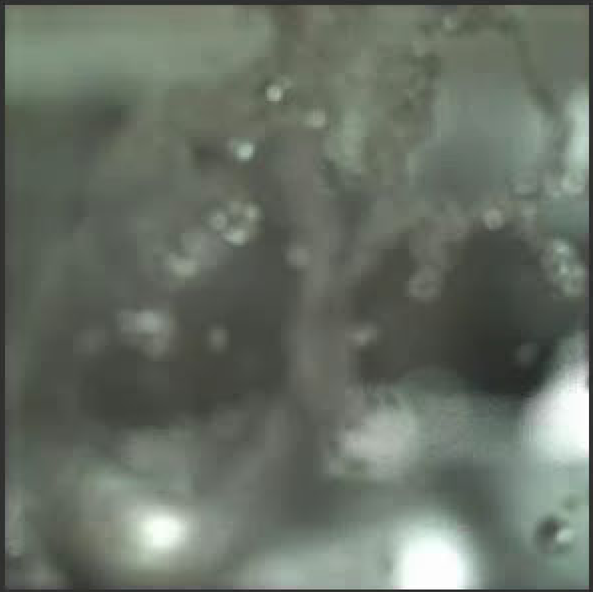

Focusing through a splash of water

Extending the depth of field

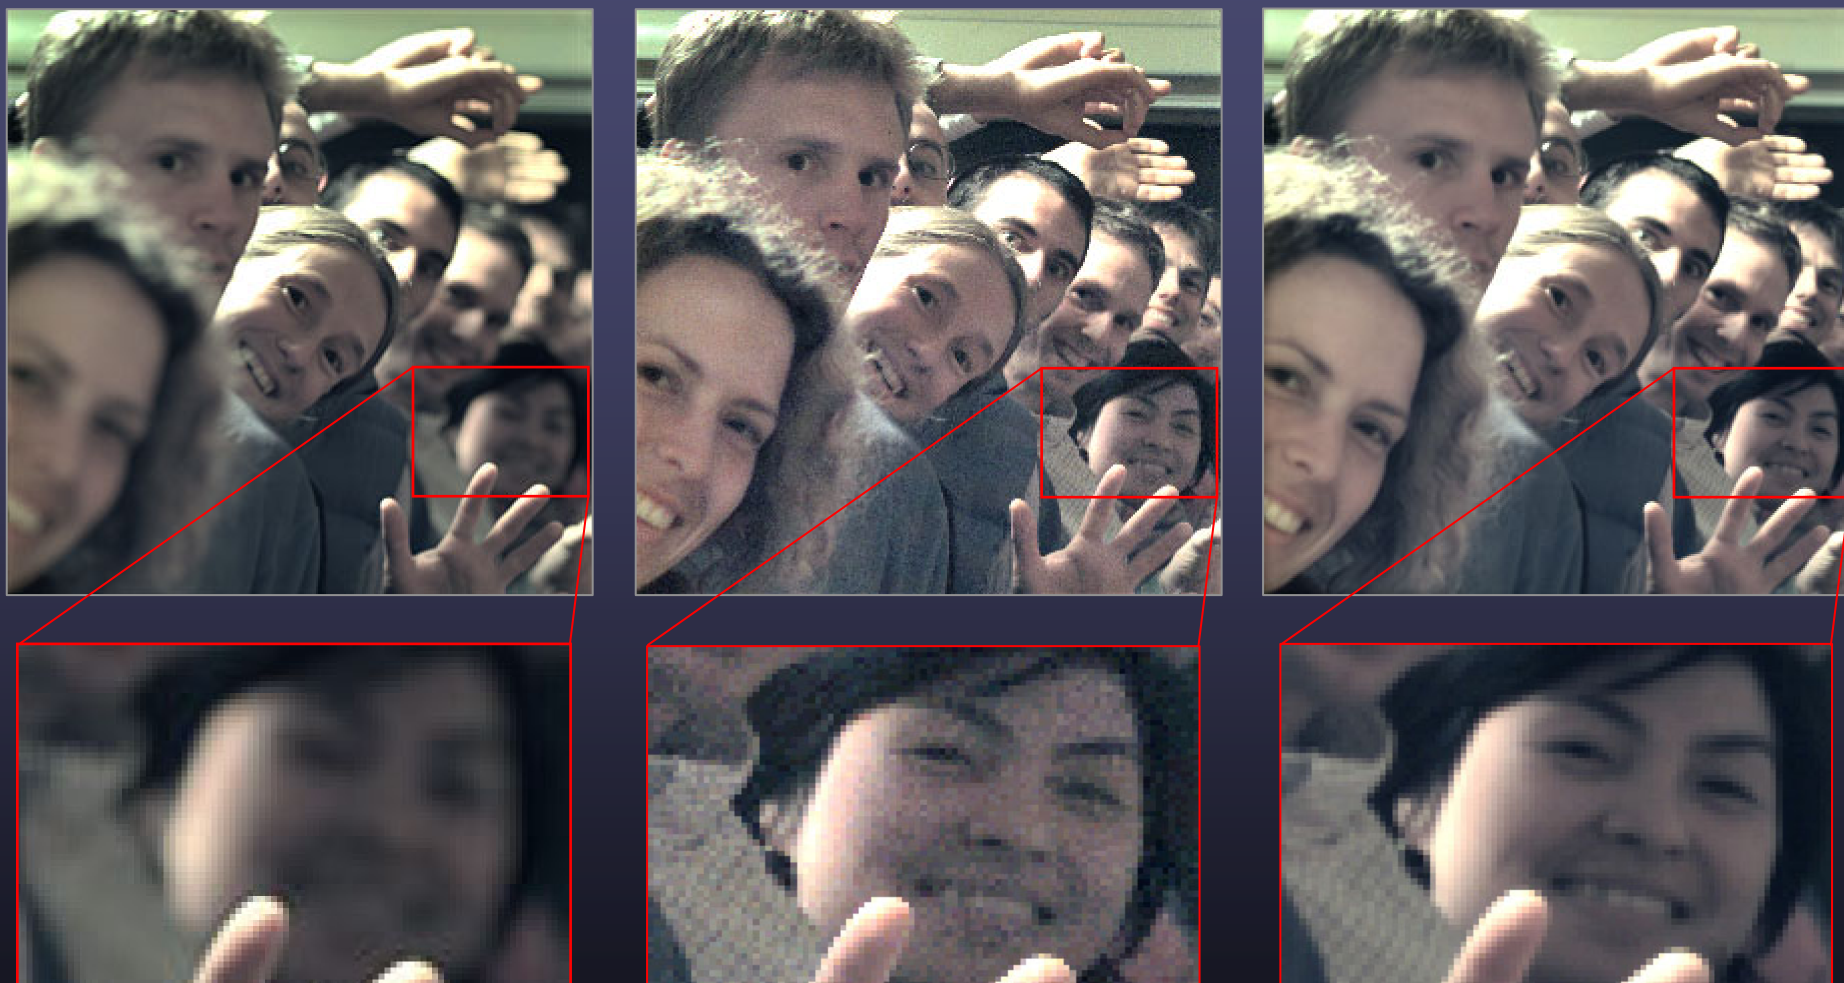

conventional photograph,
main lens at *f* / 4

conventional photograph,
main lens at *f* / 22

light field, main lens at *f* / 4,
after all-focus algorithm
[Agarwala 2004]

Macrophotography

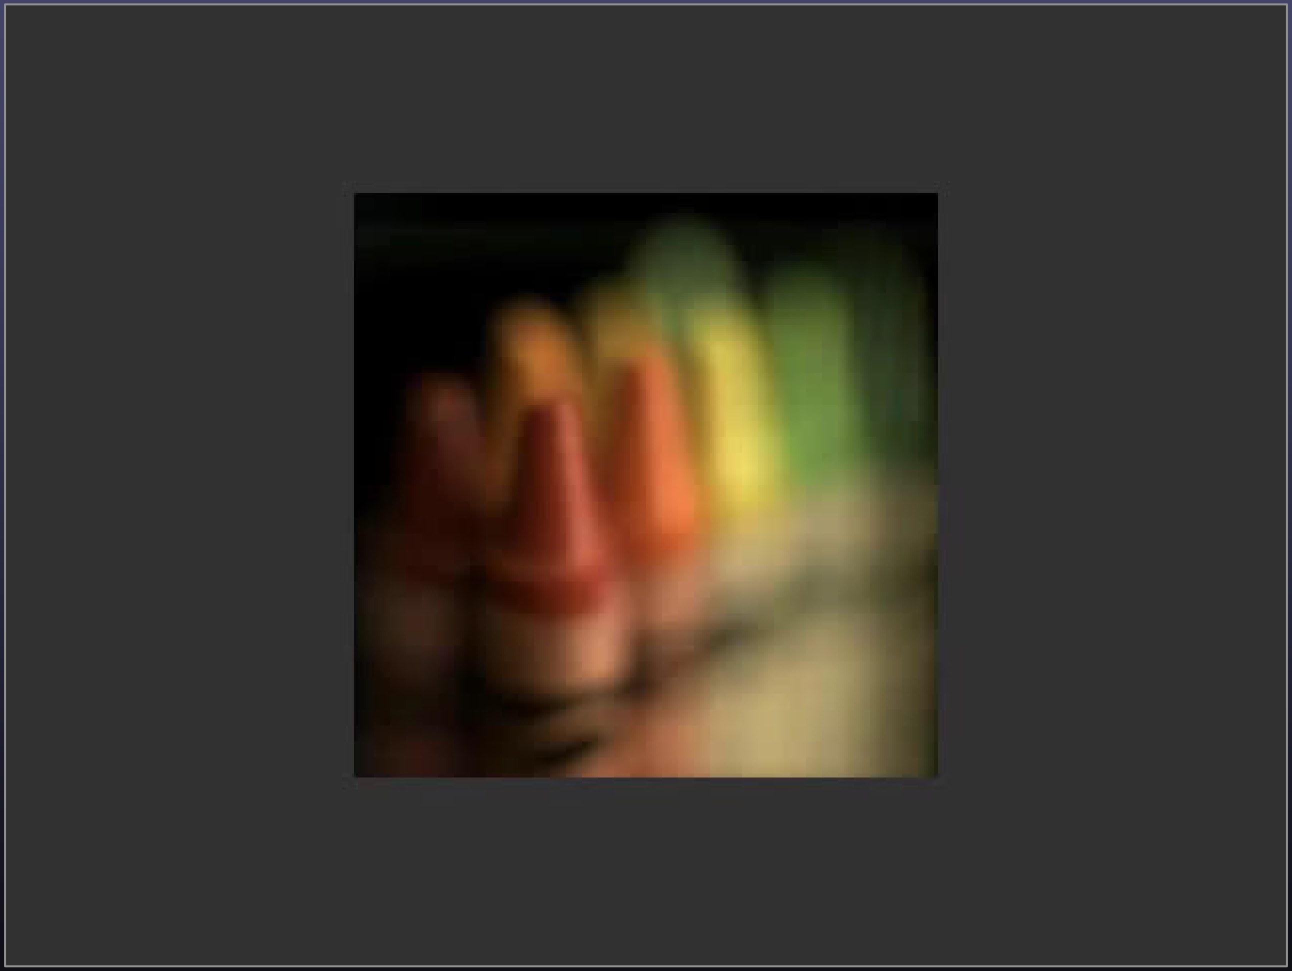

Digitally moving the observer

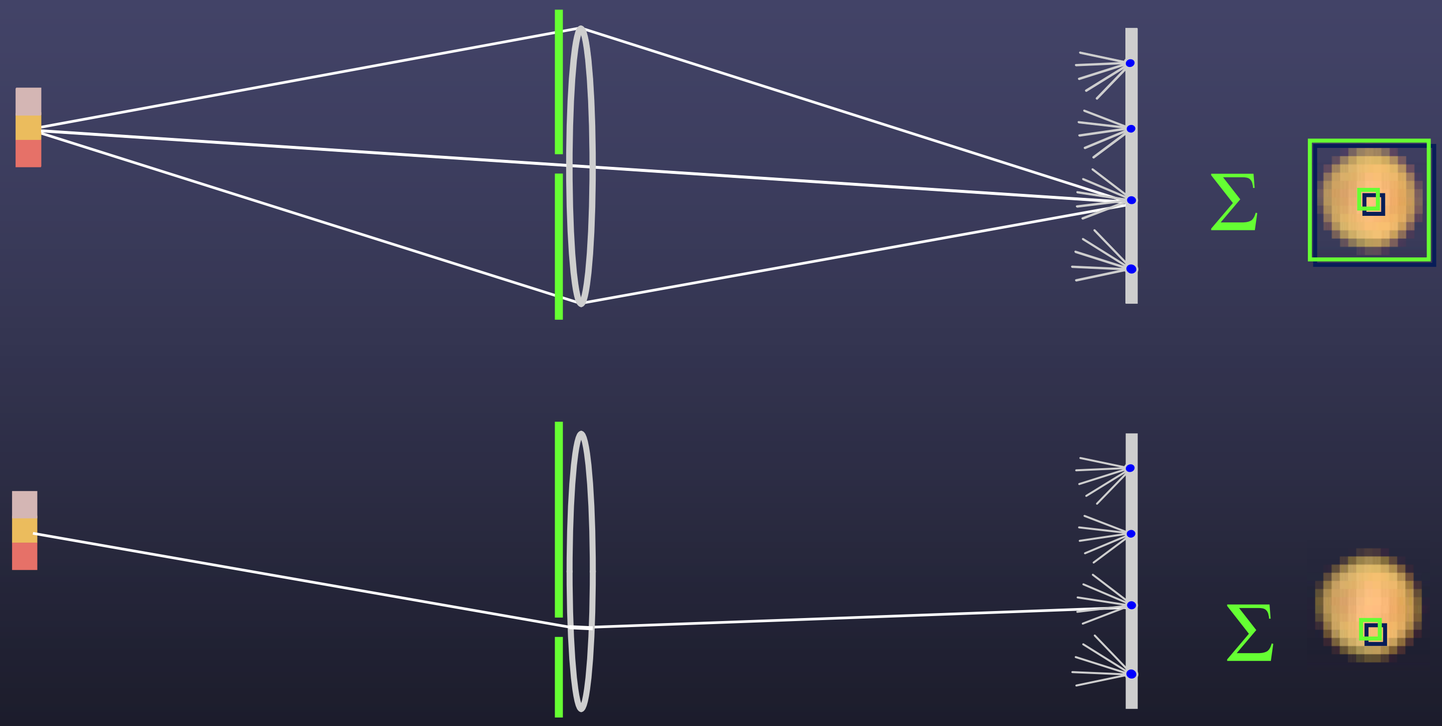

- moving the observer = moving the window we extract from the microlenses

Example of moving the observer

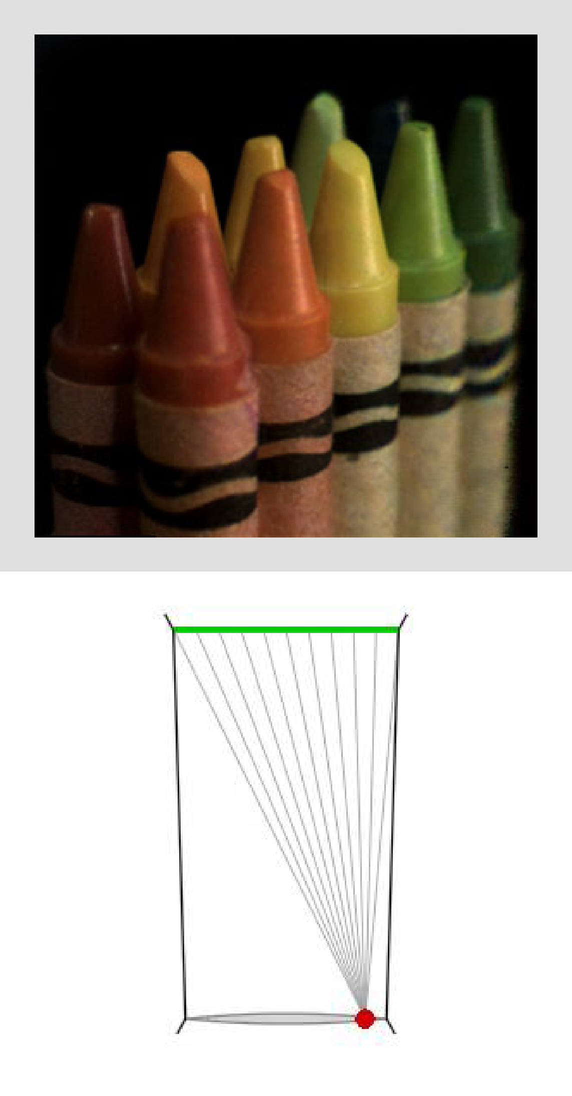

Moving backward and forward

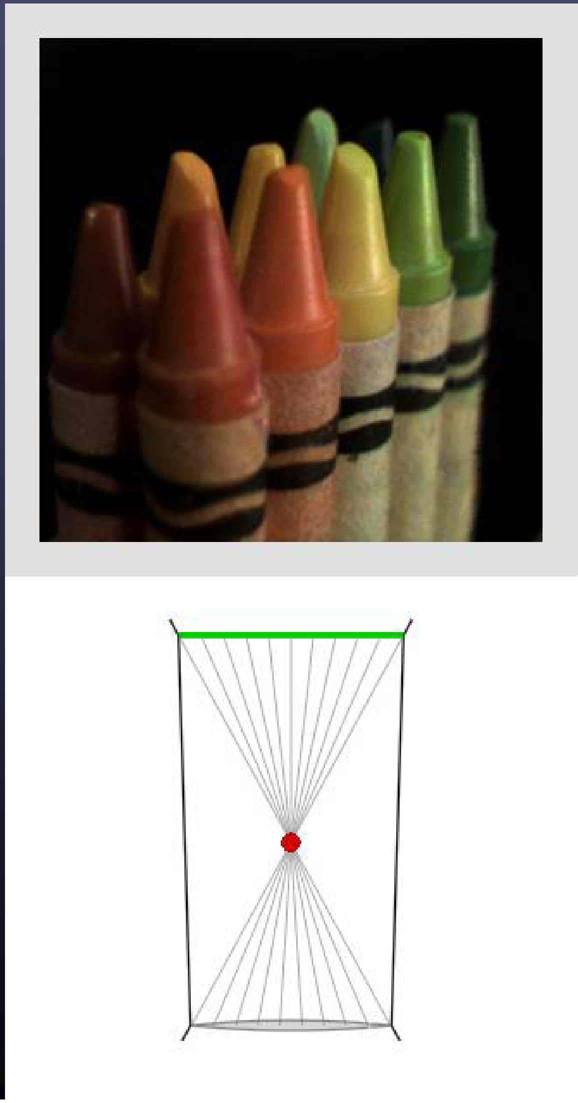

Implications

- cuts the unwanted link between exposure (due to the aperture) and depth of field
- trades off (excess) spatial resolution for ability to refocus and adjust the perspective
- sensor pixels should be made even smaller, subject to the diffraction limit

  36mm × 24mm ÷ 2.5μ pixels = 266 megapixels

  20K × 13K pixels

  4000 × 2666 pixels × 20 × 20 rays per pixel

Can we build a light field microscope?

- ability to photograph moving specimens
- digital refocusing → focal stack → deconvolution microscopy → volume data

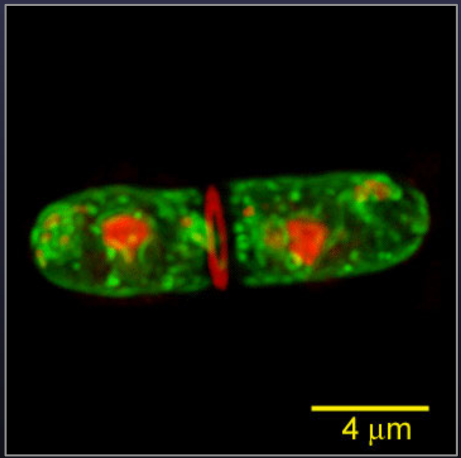

Dual Photography

*Pradeep Sen, Billy Chen, Gaurav Garg, Steve Marschner,
Mark Horowitz, Marc Levoy, Hendrik Lensch*

(Proc. SIGGRAPH 2005)

Helmholtz reciprocity

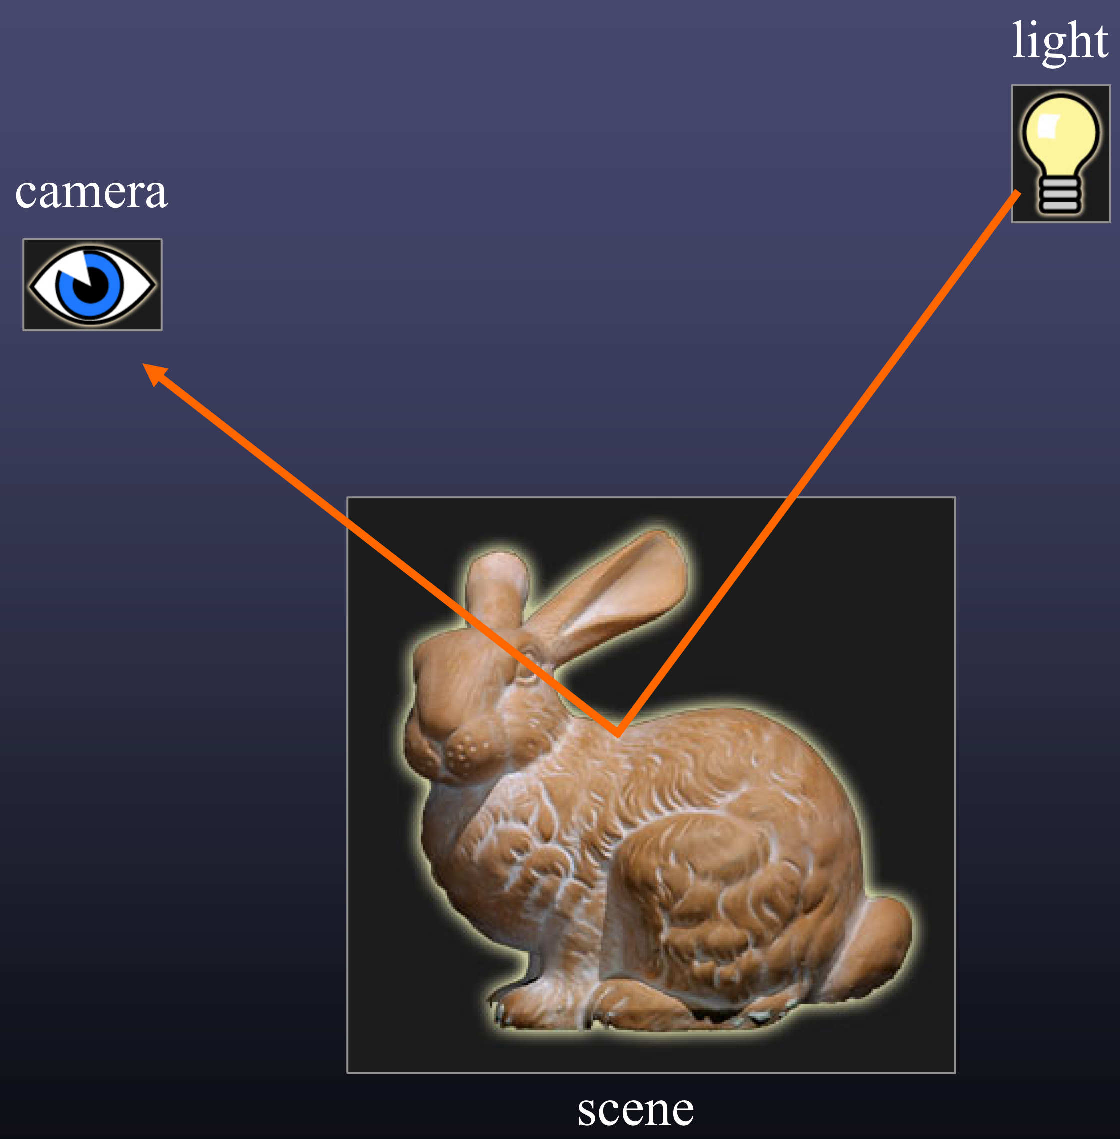

Helmholtz reciprocity

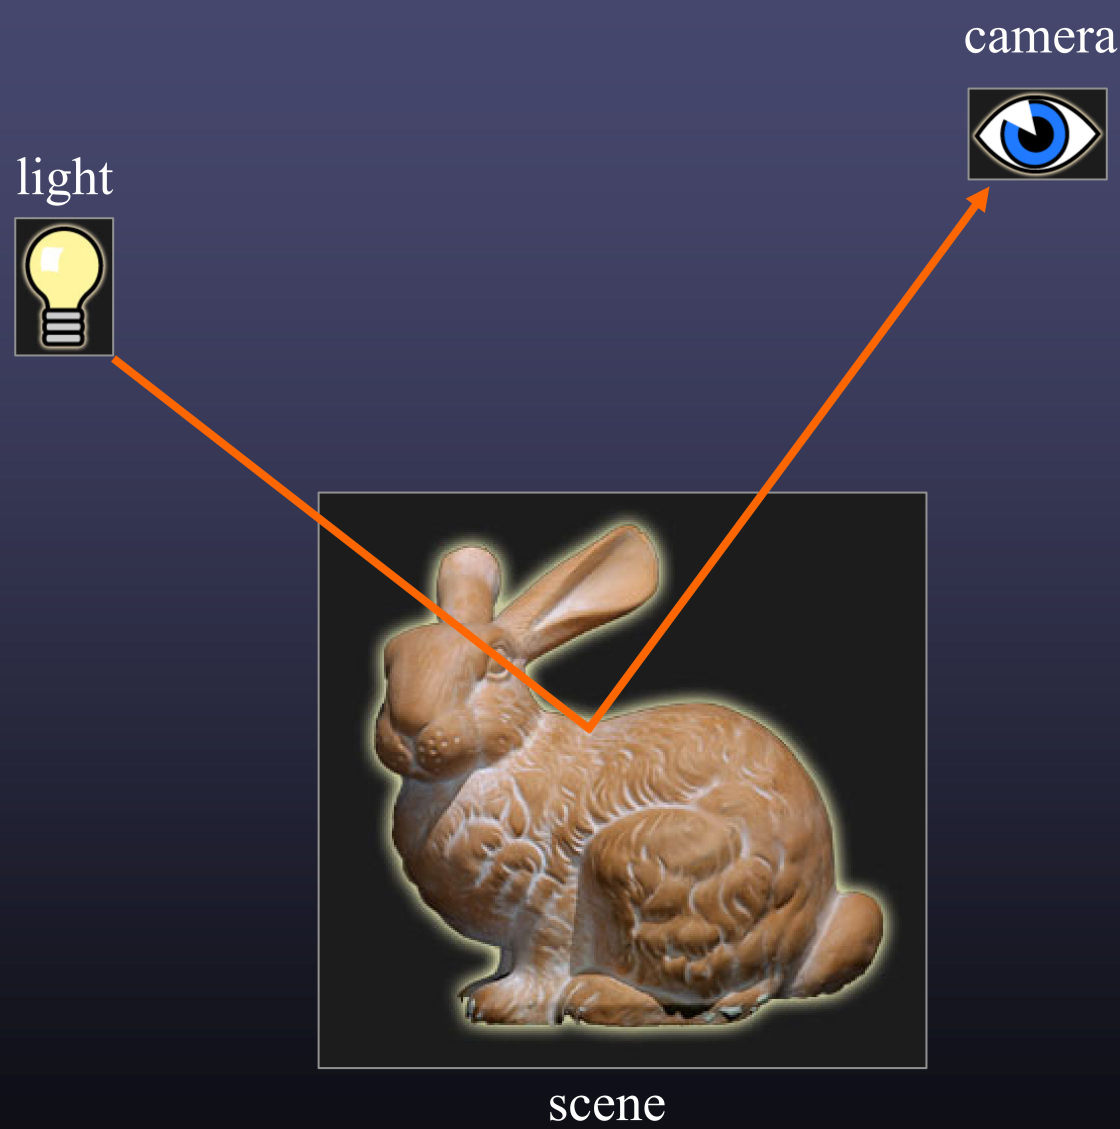

Measuring transport along a set of paths

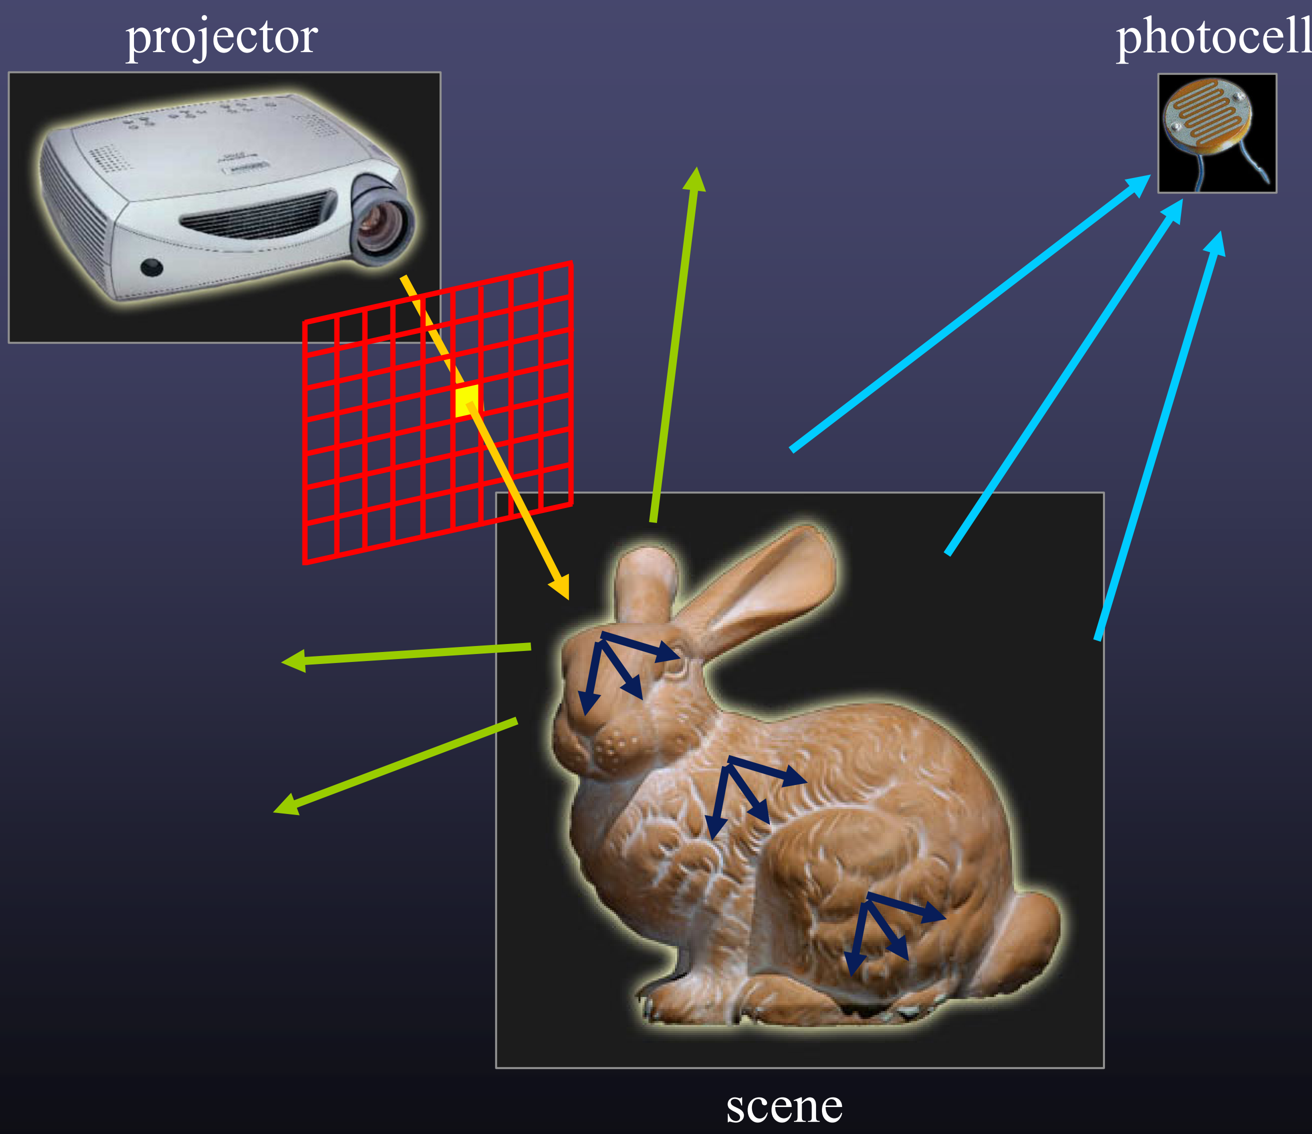

Reversing the paths

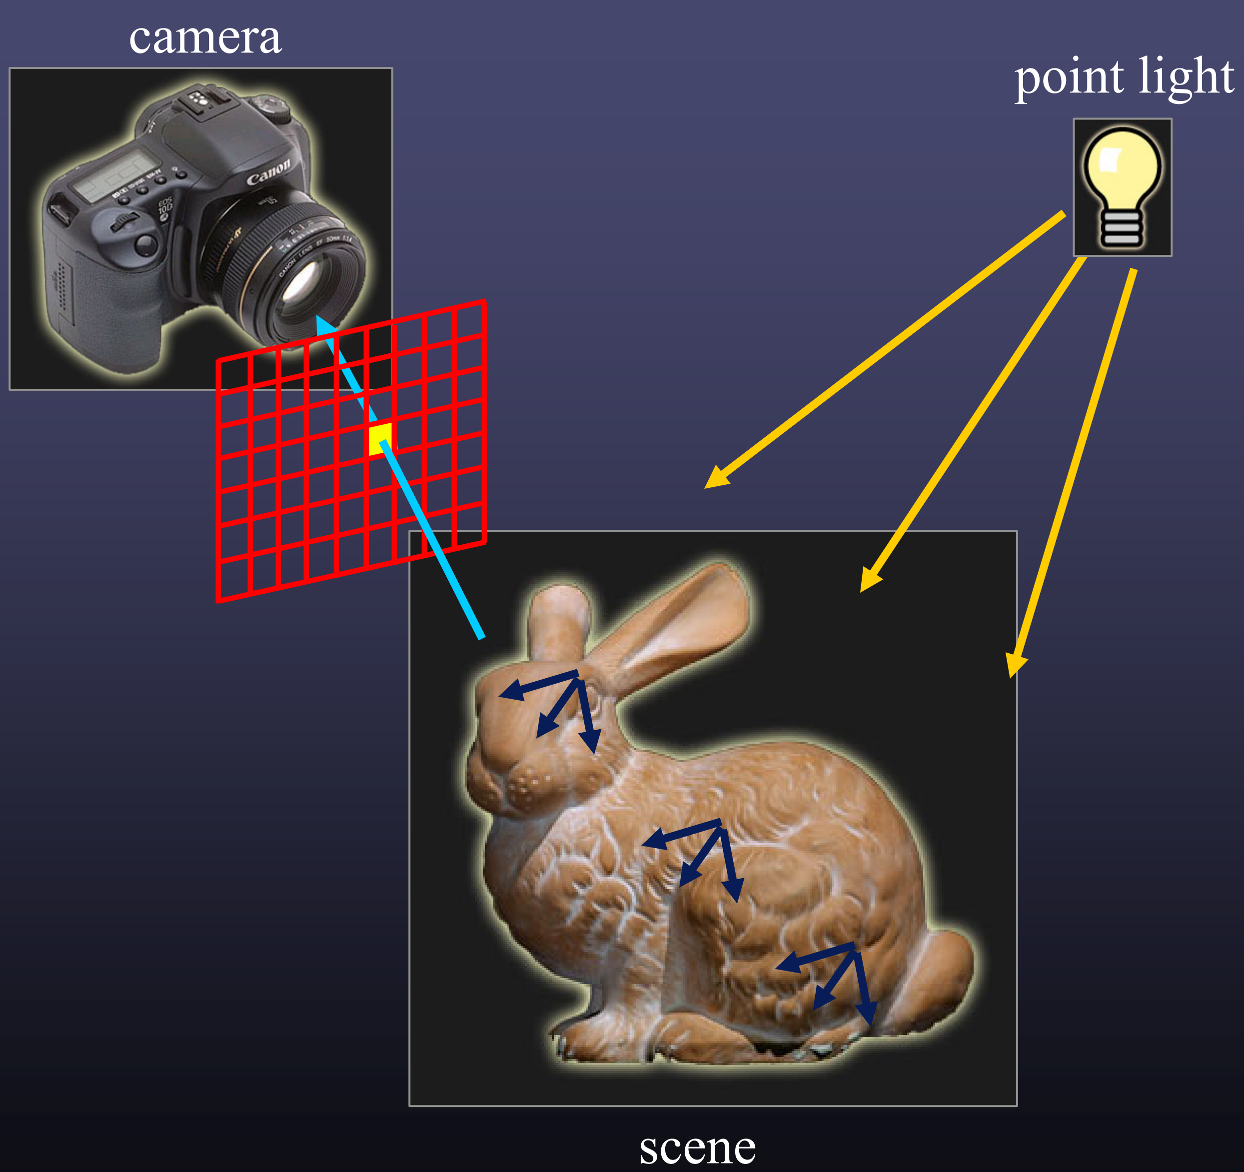

Forming a dual photograph

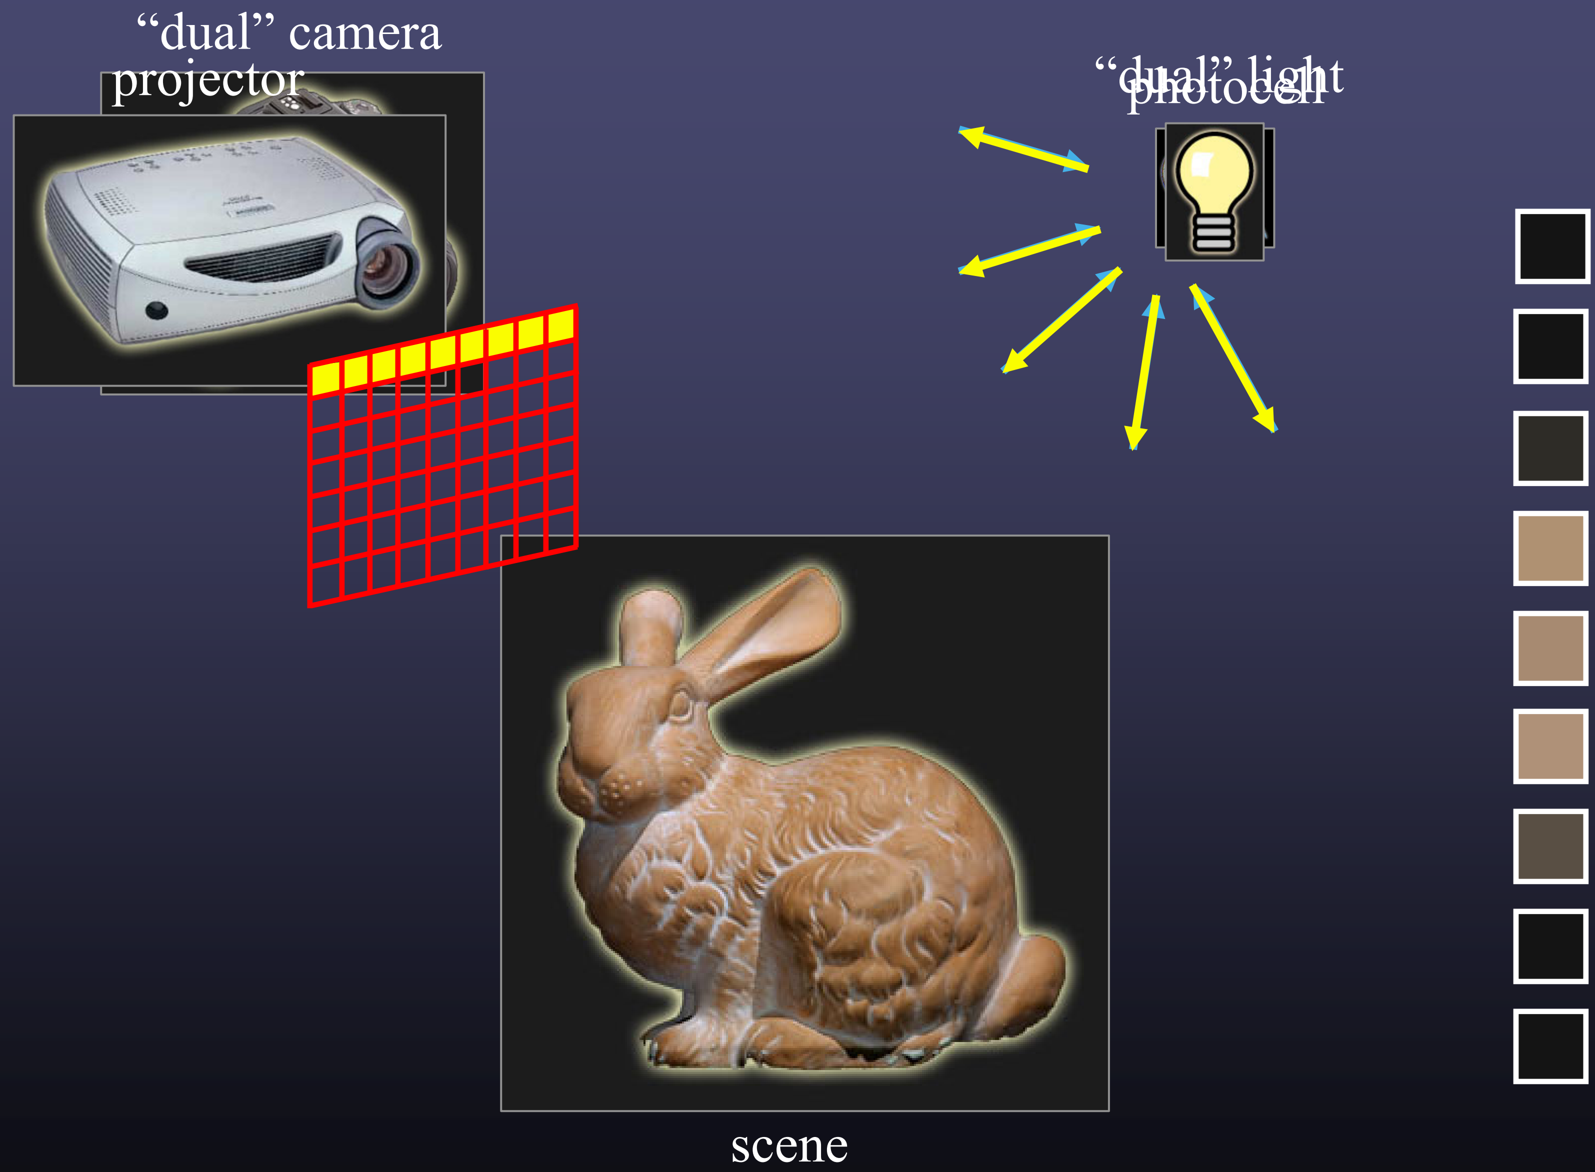

Forming a dual photograph

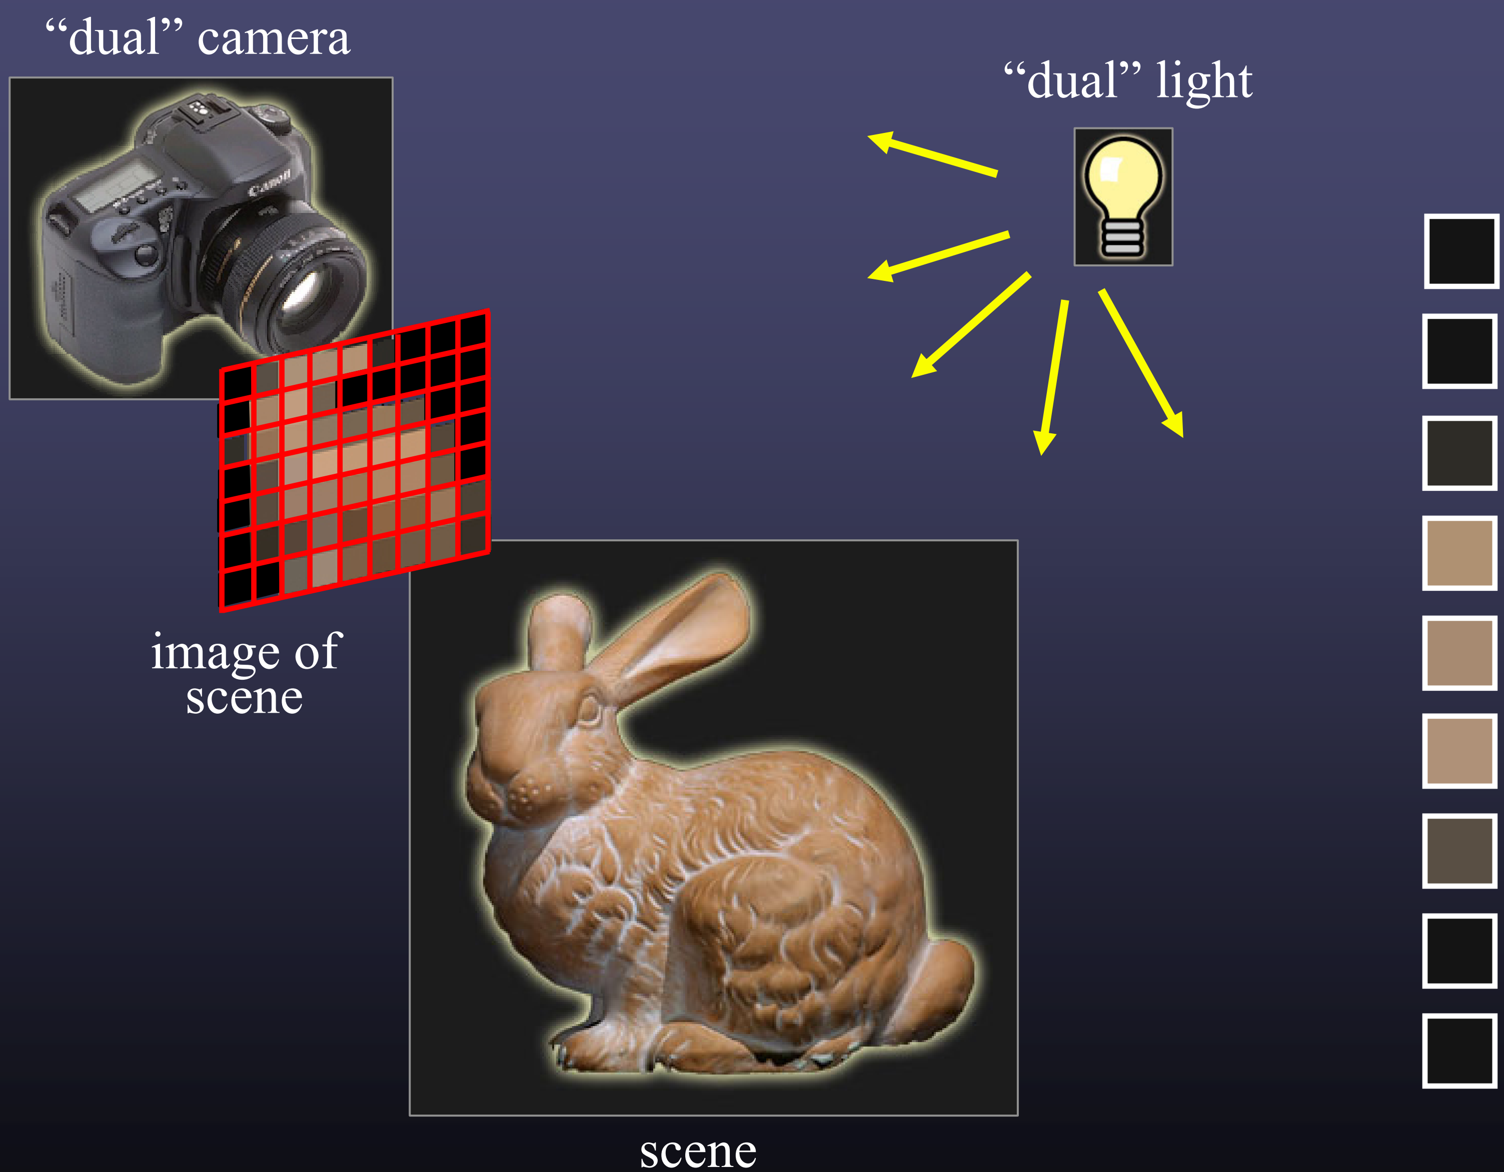

Physical demonstration

- light replaced with projector
- camera replaced with photocell
- projector scanned across the scene

conventional photograph,
with light coming from right

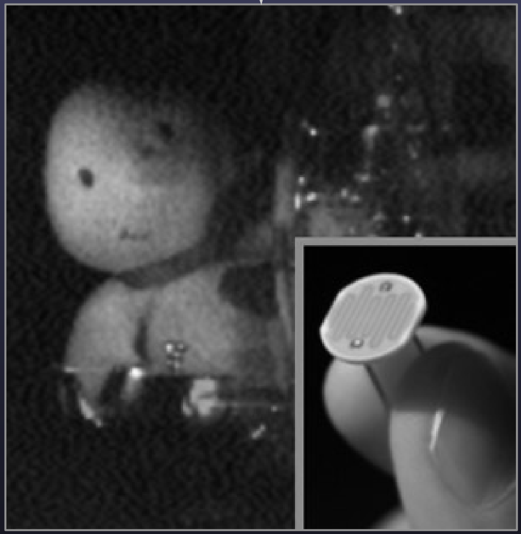

dual photograph,
as seen from projector's position
and as illuminated from photocell's position

Related imaging methods

- time-of-flight scanner
  - if they return reflectance as well as range
  - but their light source and sensor are typically coaxial
- scanning electron microscope

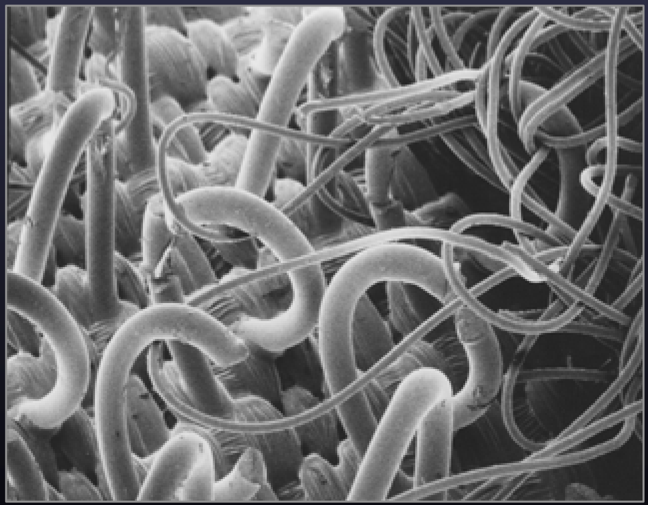

Velcro® at 35x magnification,
Museum of Science, Boston

The 4D transport matrix

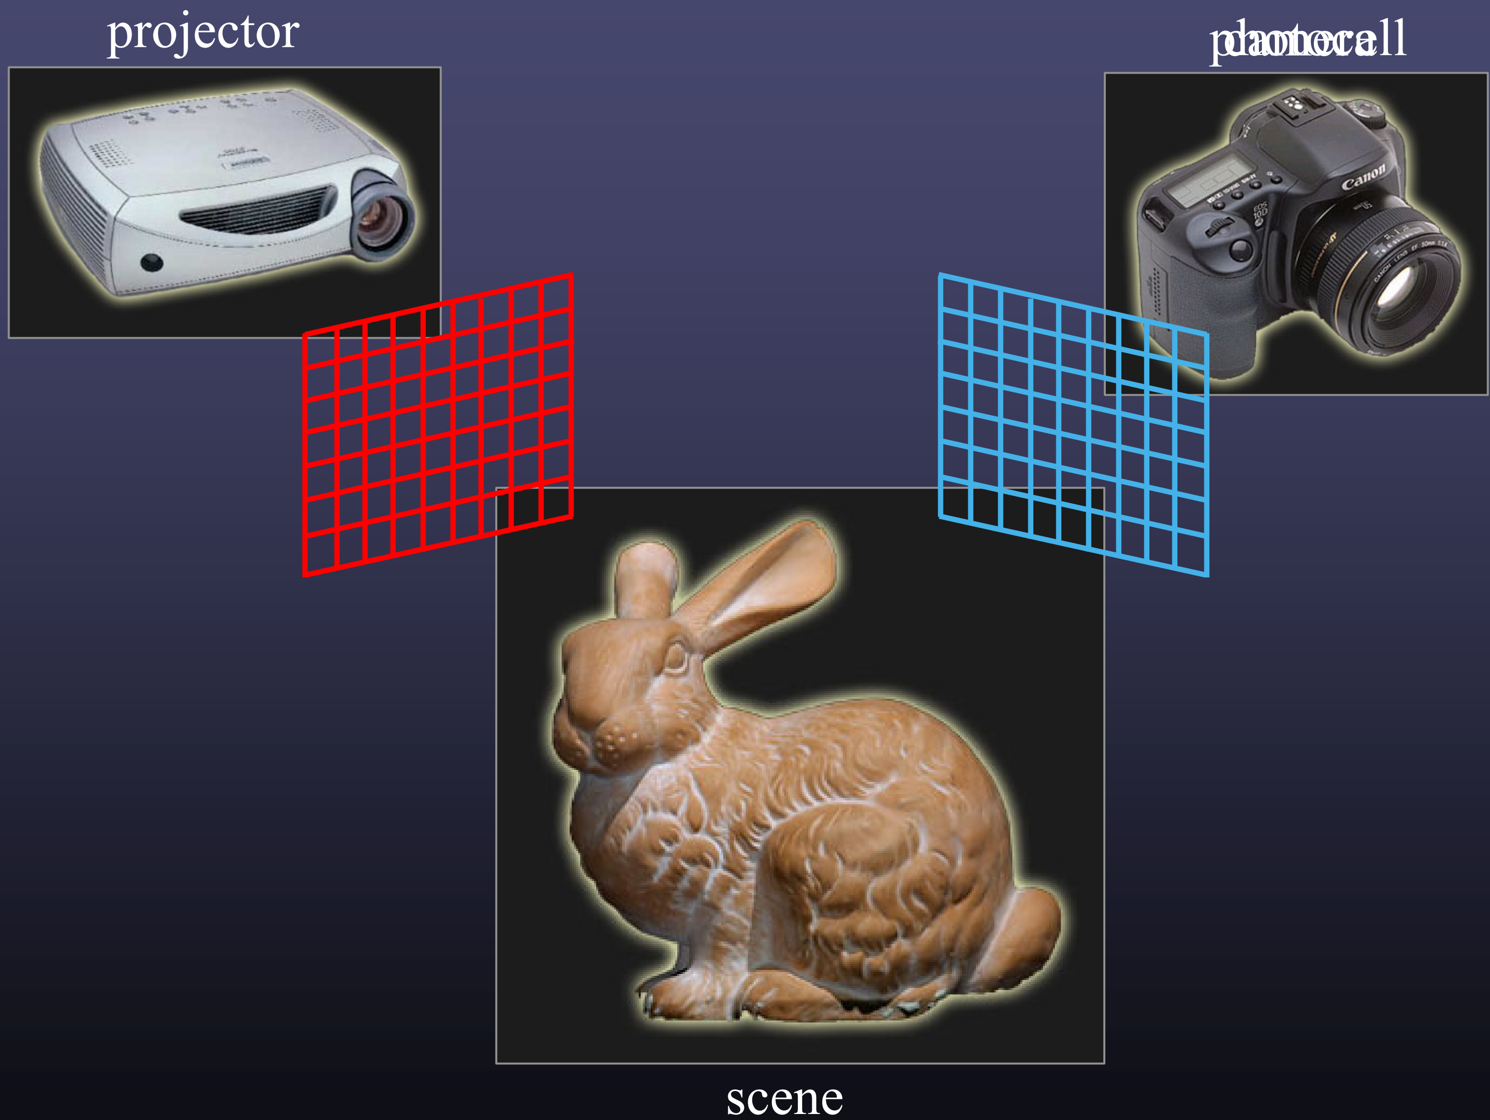

The 4D transport matrix

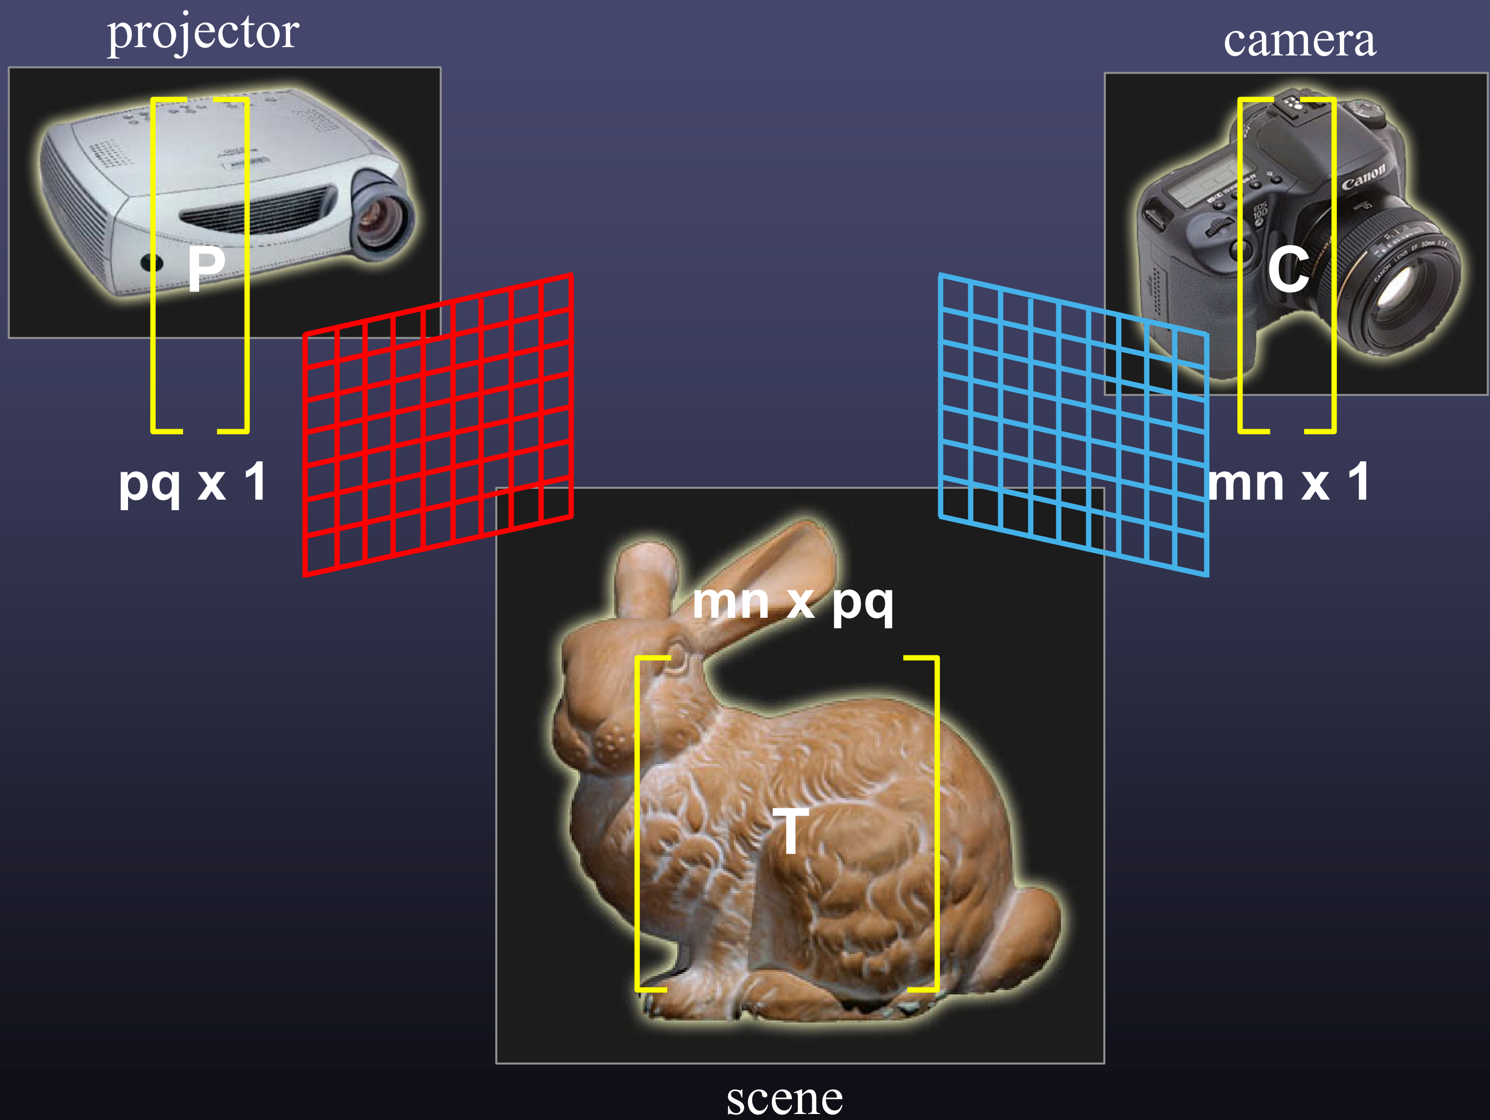

The 4D transport matrix

mn x pq

$$\begin{bmatrix} C \end{bmatrix} = \begin{bmatrix} \quad T \quad \end{bmatrix} \begin{bmatrix} P \end{bmatrix}$$

mn x 1 pq x 1

The 4D transport matrix

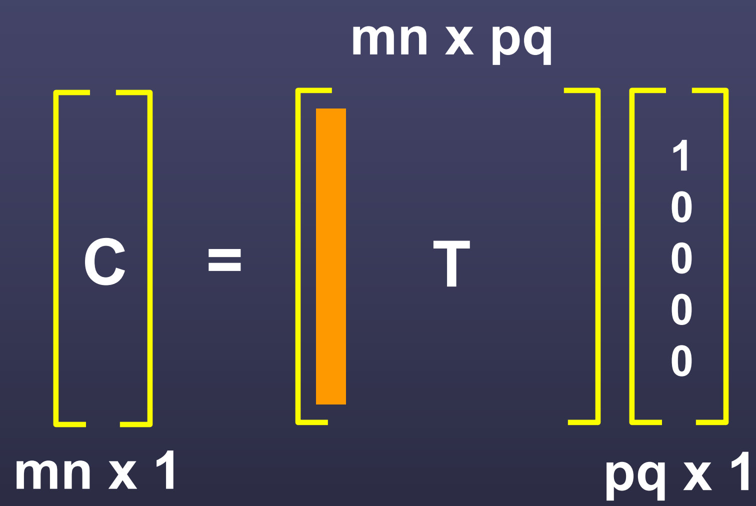

The 4D transport matrix

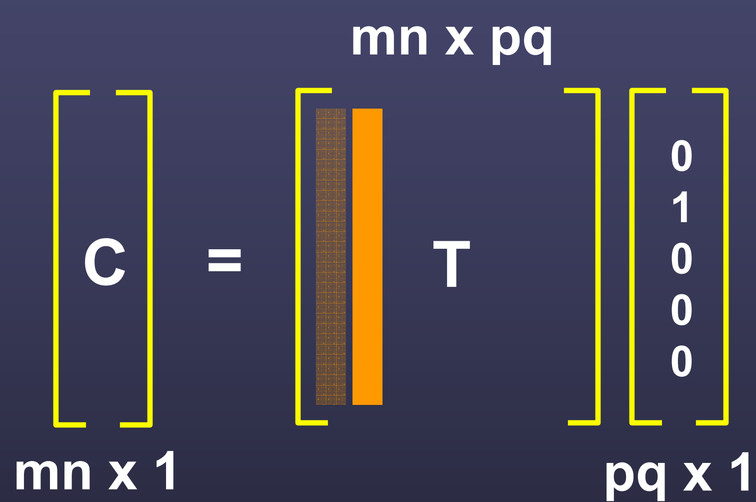

The 4D transport matrix

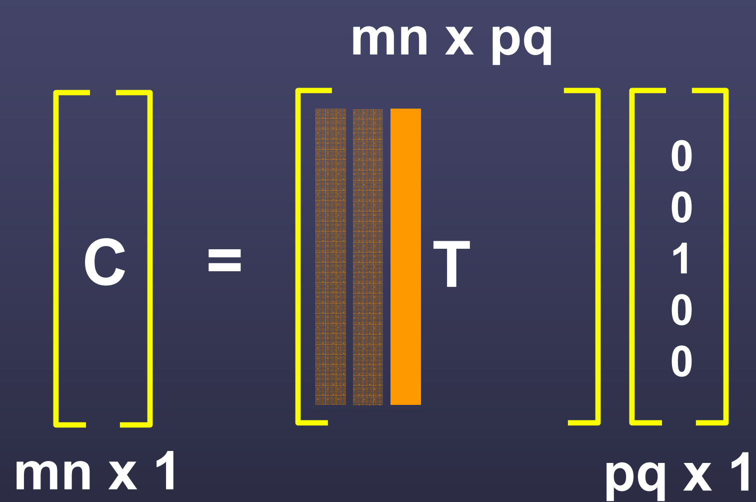

The 4D transport matrix

$$\underset{mn \times 1}{\left[\, C \,\right]} = \underset{mn \times pq}{\left[\quad T \quad\right]} \underset{pq \times 1}{\left[\, P \,\right]}$$

The 4D transport matrix

$$\underset{mn \times 1}{\begin{bmatrix} C \end{bmatrix}} = \overset{mn \times pq}{\begin{bmatrix} \quad T \quad \end{bmatrix}} \underset{pq \times 1}{\begin{bmatrix} P \end{bmatrix}}$$

applying Helmholtz reciprocity...

$$\underset{pq \times 1}{\begin{bmatrix} C' \end{bmatrix}} = \overset{pq \times mn}{\begin{bmatrix} \quad T^T \quad \end{bmatrix}} \underset{mn \times 1}{\begin{bmatrix} P' \end{bmatrix}}$$

Example

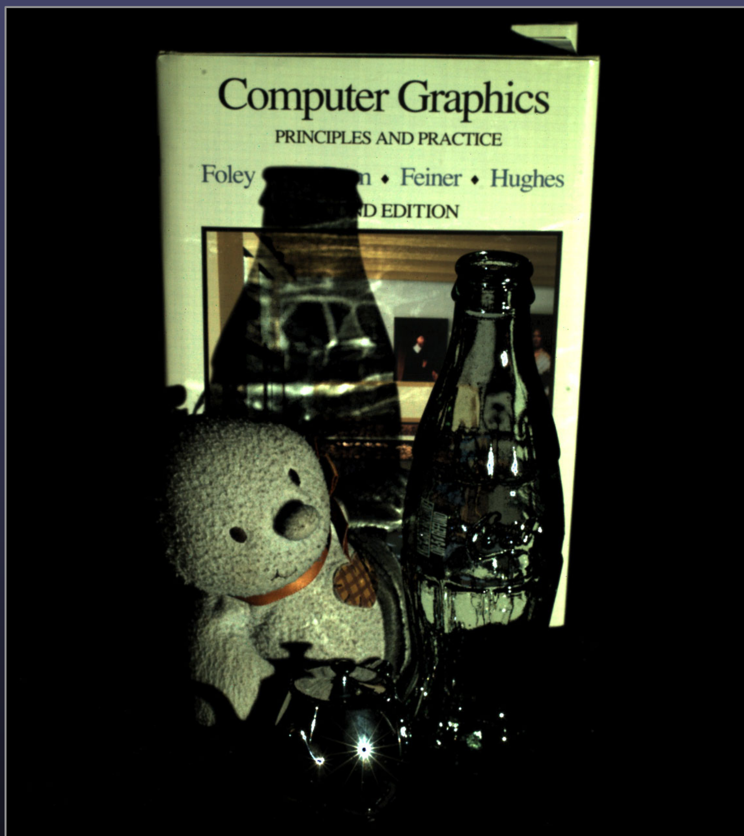

conventional photograph
with light coming from right

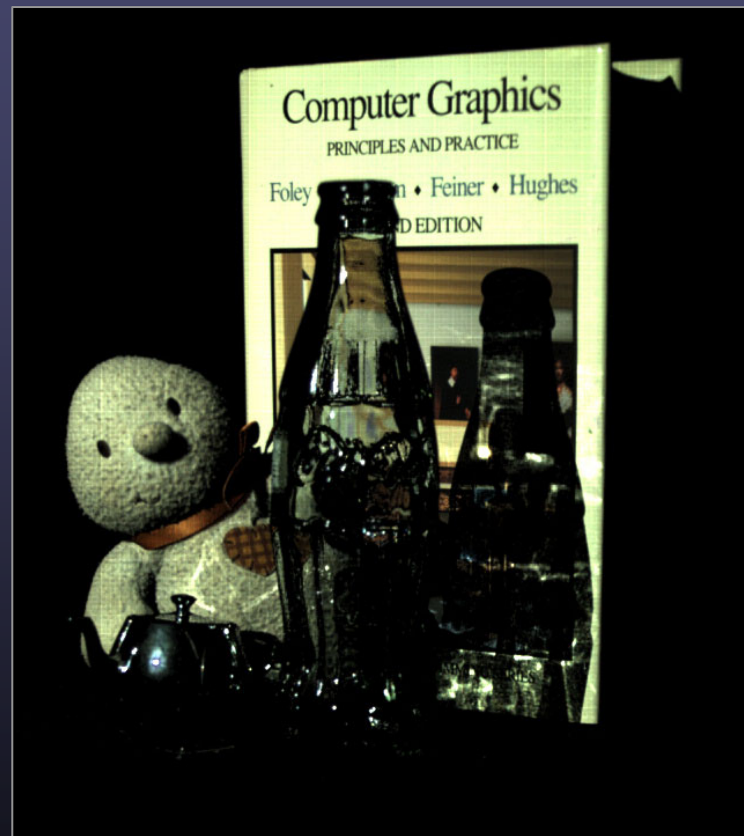

dual photograph
as seen from projector's position

Properties of the transport matrix

- little interreflection → sparse matrix
- many interreflections → dense matrix
- convex object → diagonal matrix
- concave object → full matrix

Can we create a dual photograph entirely from diffuse reflections?

Dual photography from diffuse reflections

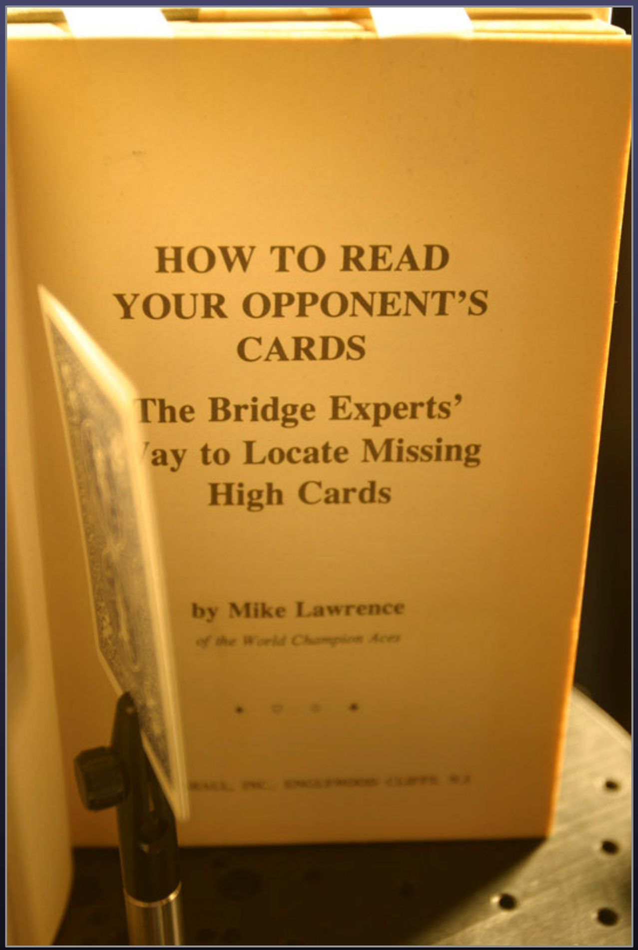

the camera's view

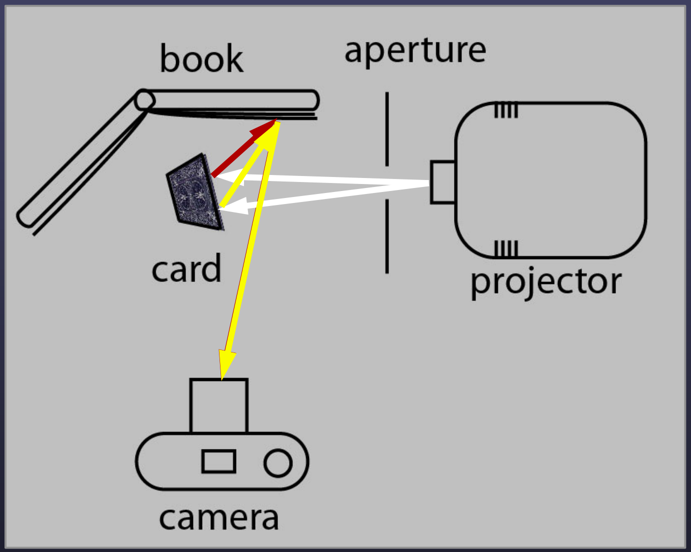

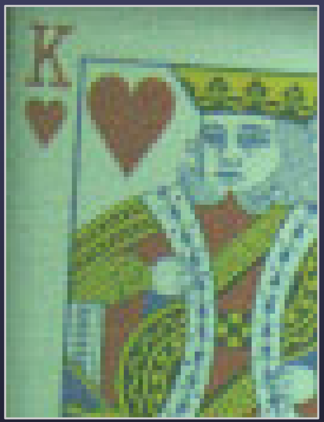

The relighting problem

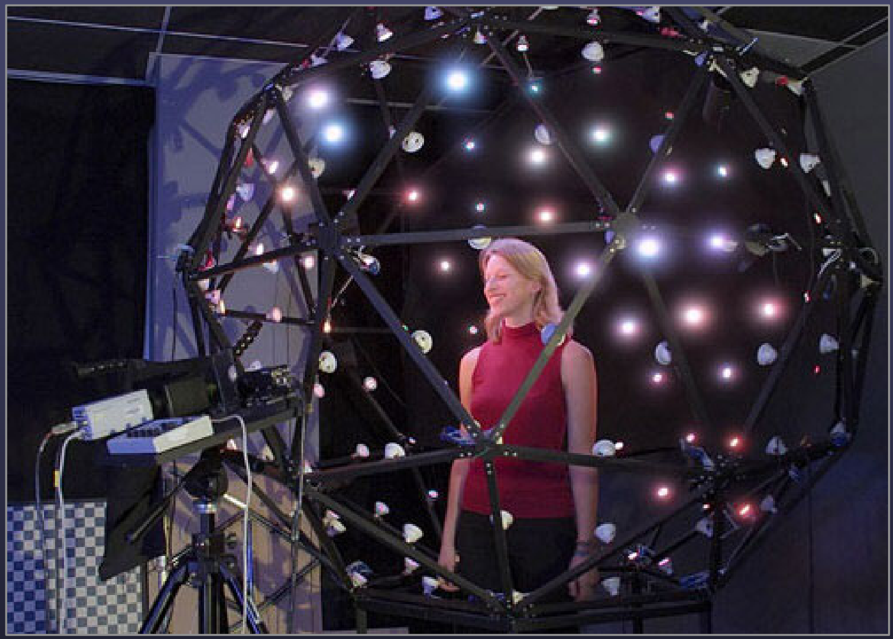

Paul Debevec's
Light Stage 3

- subject captured under multiple lights
- one light at a time, so subject must hold still
- point lights are used, so can't relight with cast shadows

The 6D transport matrix

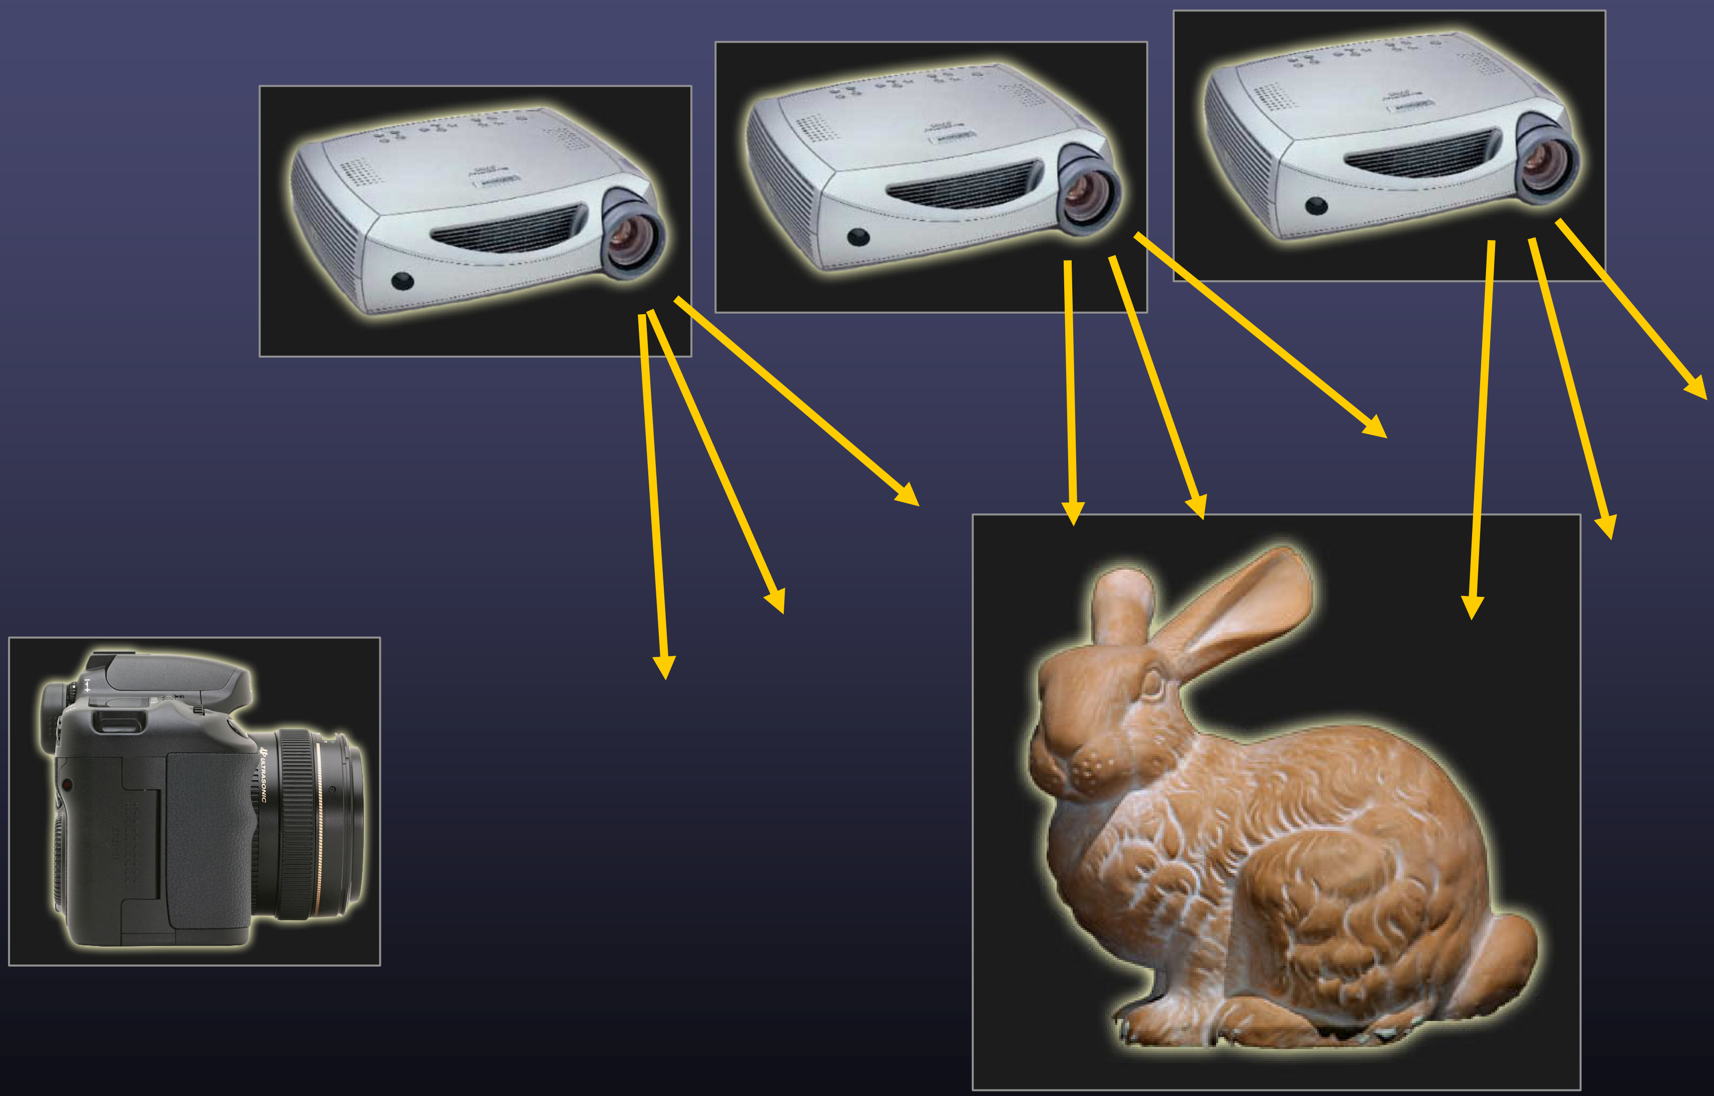

The 6D transport matrix

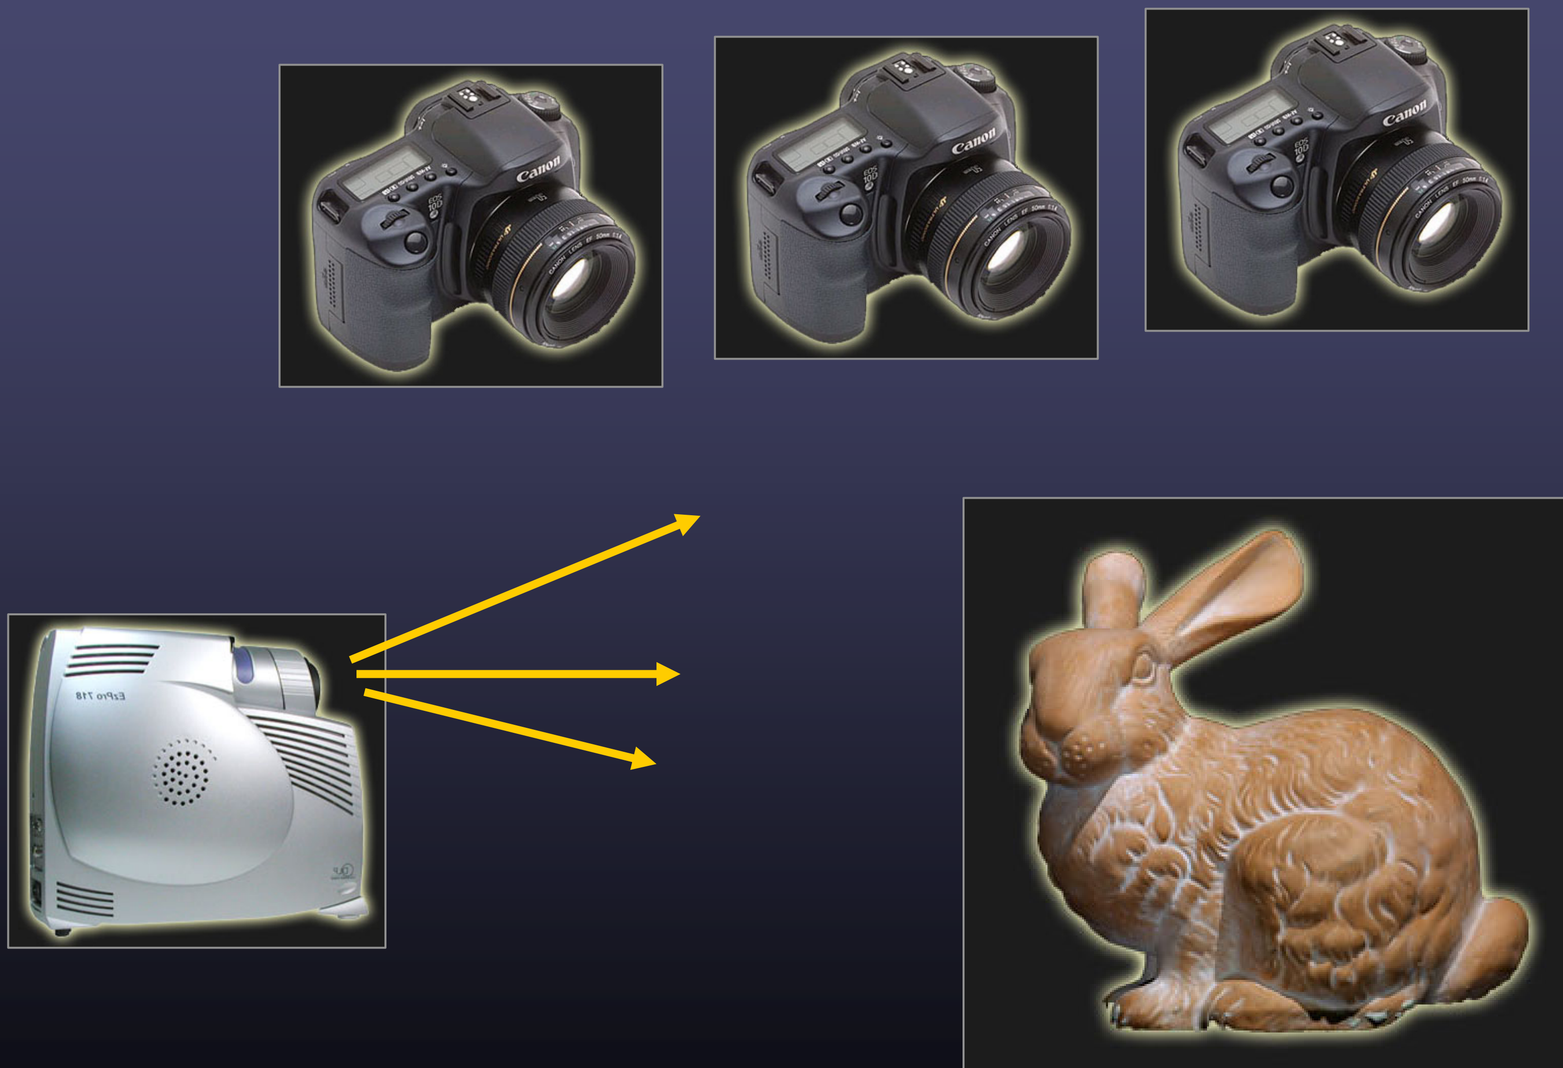

The advantage of dual photography

- capture of a scene as illuminated by different lights cannot be parallelized
- capture of a scene as viewed by different cameras can be parallelized

Measuring the 6D transport matrix

projector

camera array

scene

Relighting with complex illumination

projector

camera array

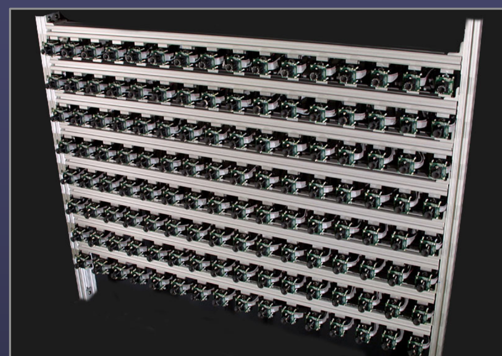

scene

$$\underset{pq \times 1}{\left[\, C' \,\right]} = \underset{pq \times mn \times uv}{\left[\quad T^T \quad\right]} \underset{mn \times uv \times 1}{\left[\, P' \,\right]}$$

- step 1: measure 6D transport matrix T
- step 2: capture a 4D light field
- step 3: relight scene using captured light field

Running time

- the different rays within a projector can in fact be parallelized to some extent

- this parallelism can be discovered using a coarse-to-fine adaptive scan

- can measure a 6D transport matrix in 5 minutes

Can we measure an 8D transport matrix?

projector array

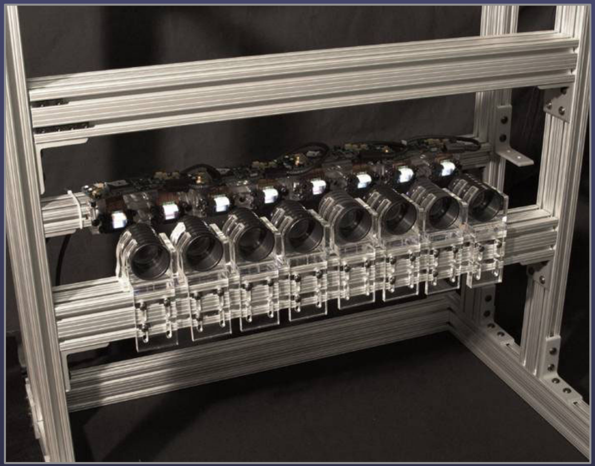

camera array

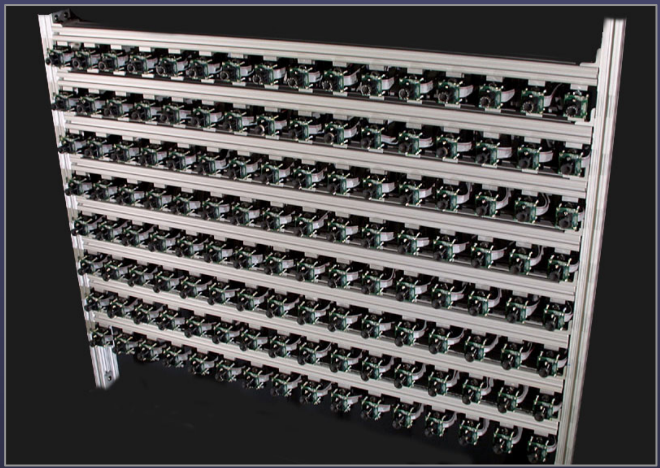

scene

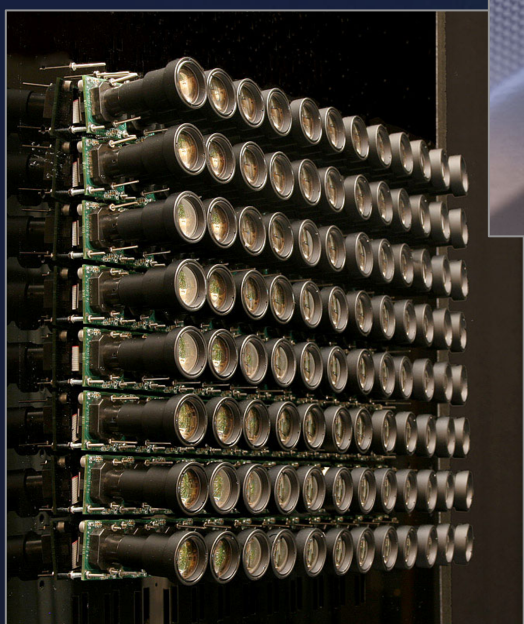

$$\underset{mn \times 1}{\begin{bmatrix} C \end{bmatrix}} = \underset{mn \times pq}{\begin{bmatrix} \quad T \quad \end{bmatrix}} \underset{pq \times 1}{\begin{bmatrix} P \end{bmatrix}}$$

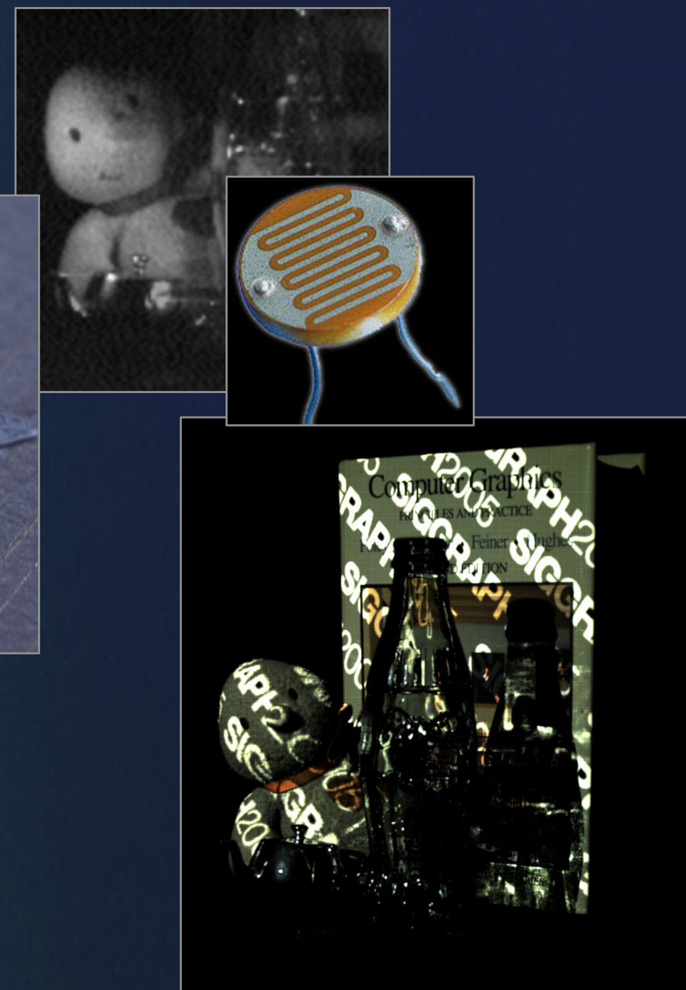

http://graphics.stanford.edu