Light Field

Modeling a desktop

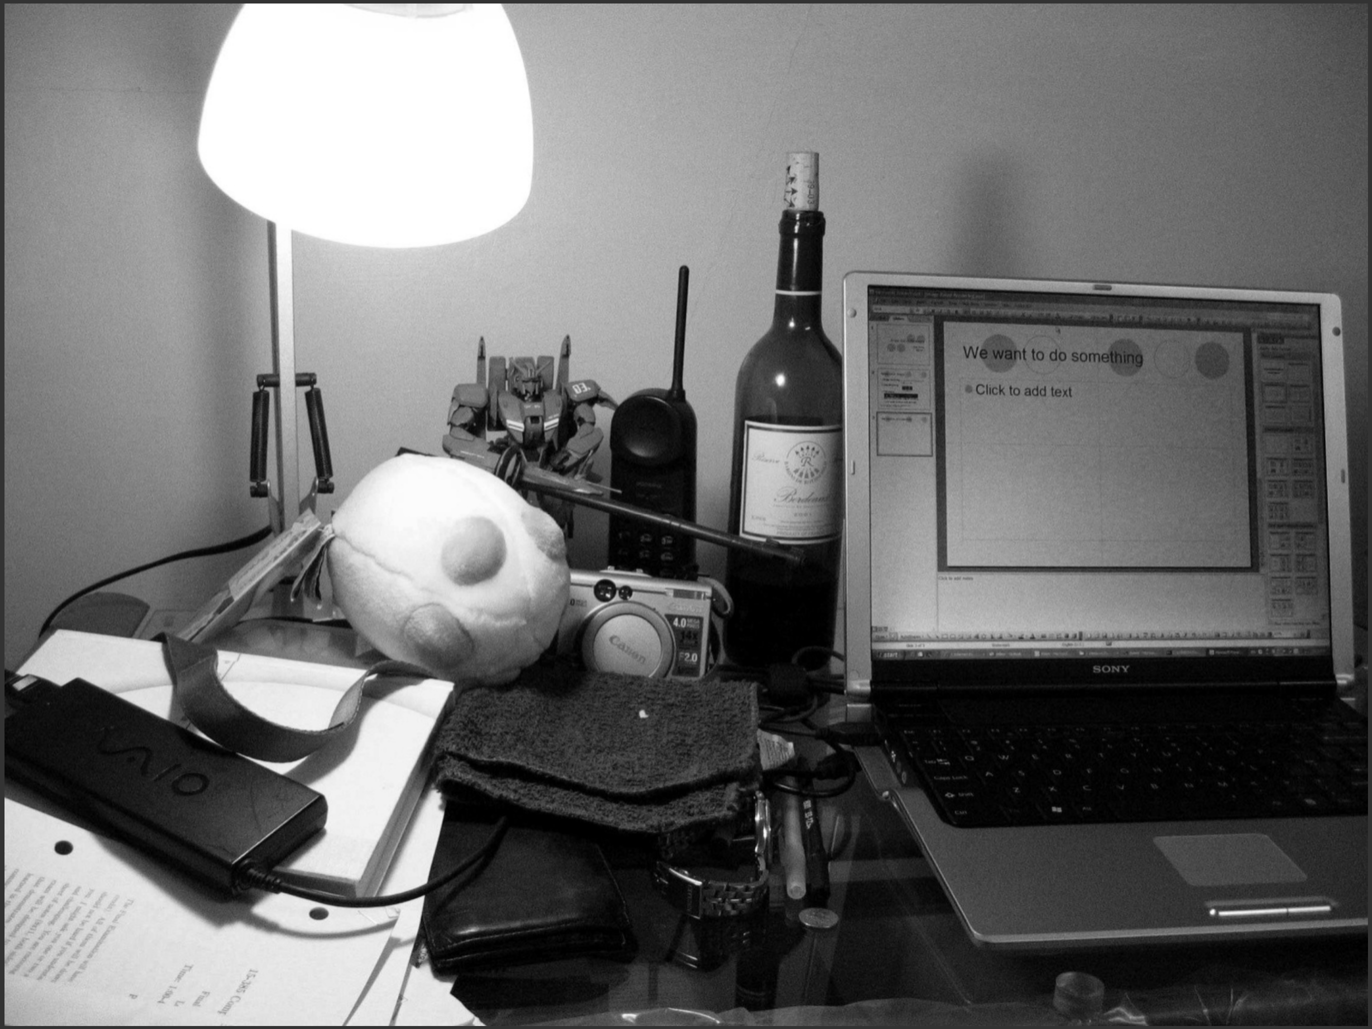


Bordeaux
2001
PRODUCT OF FRANCE


Backspace
Enter

Image Based Rendering

- Fast Realistic Rendering without 3D models

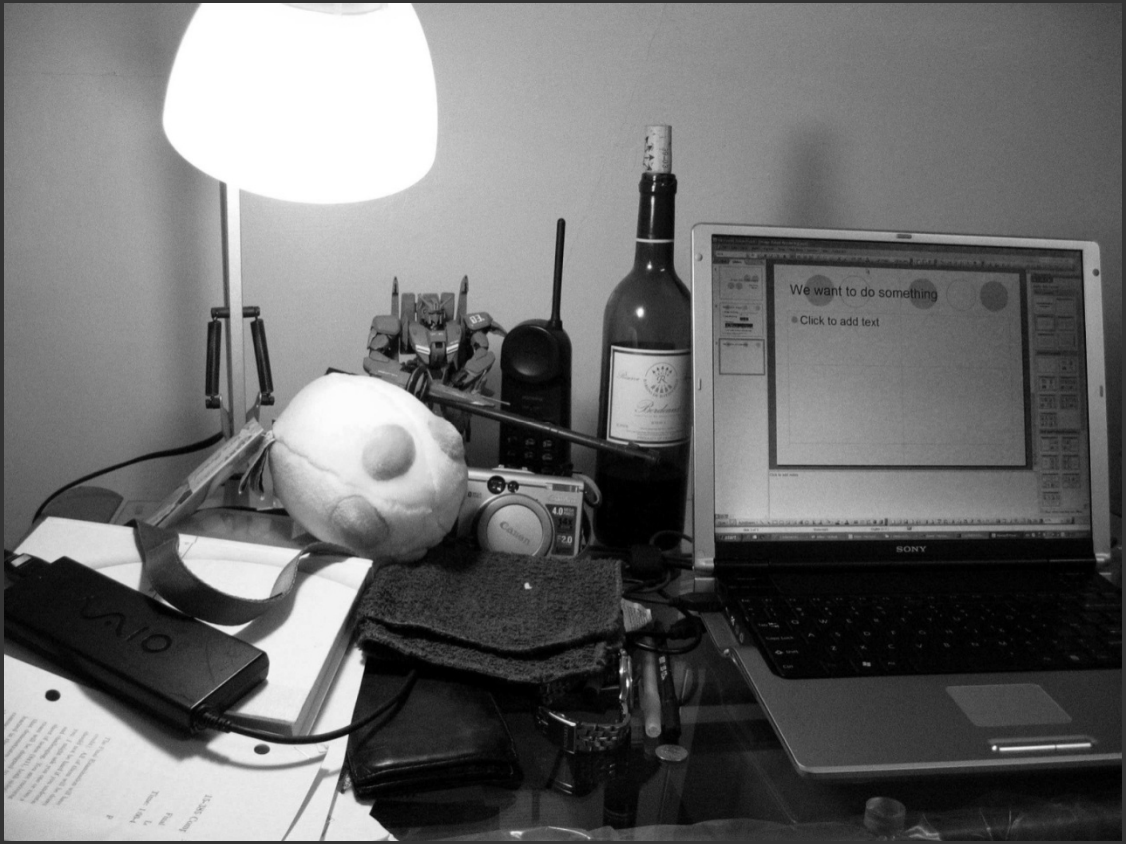

Start from Ray Tracing

- Rendering is about computing color along each ray

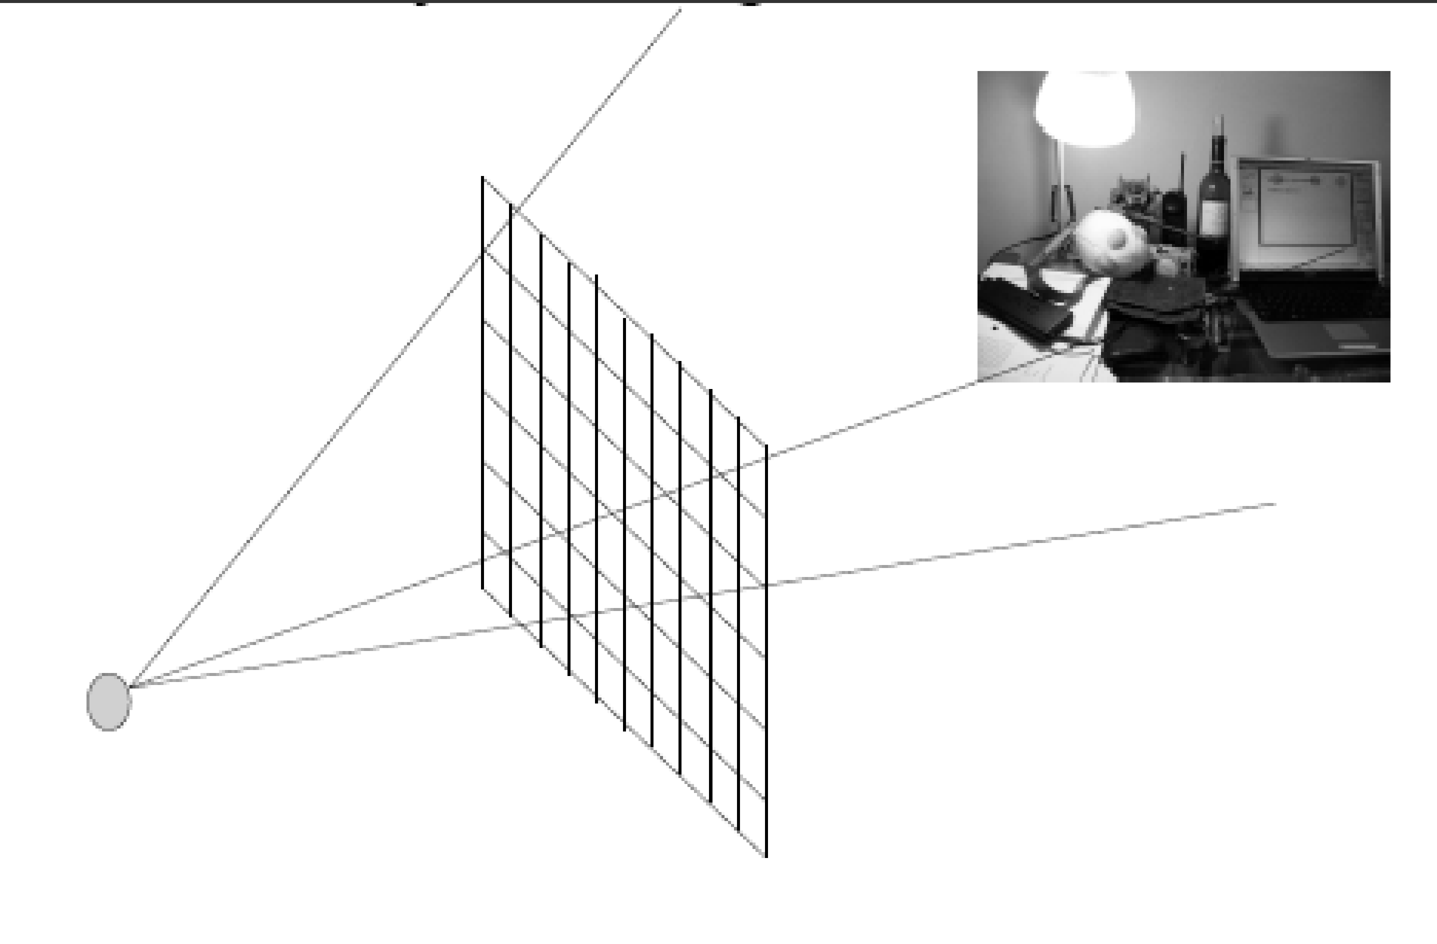

Sampling Rays

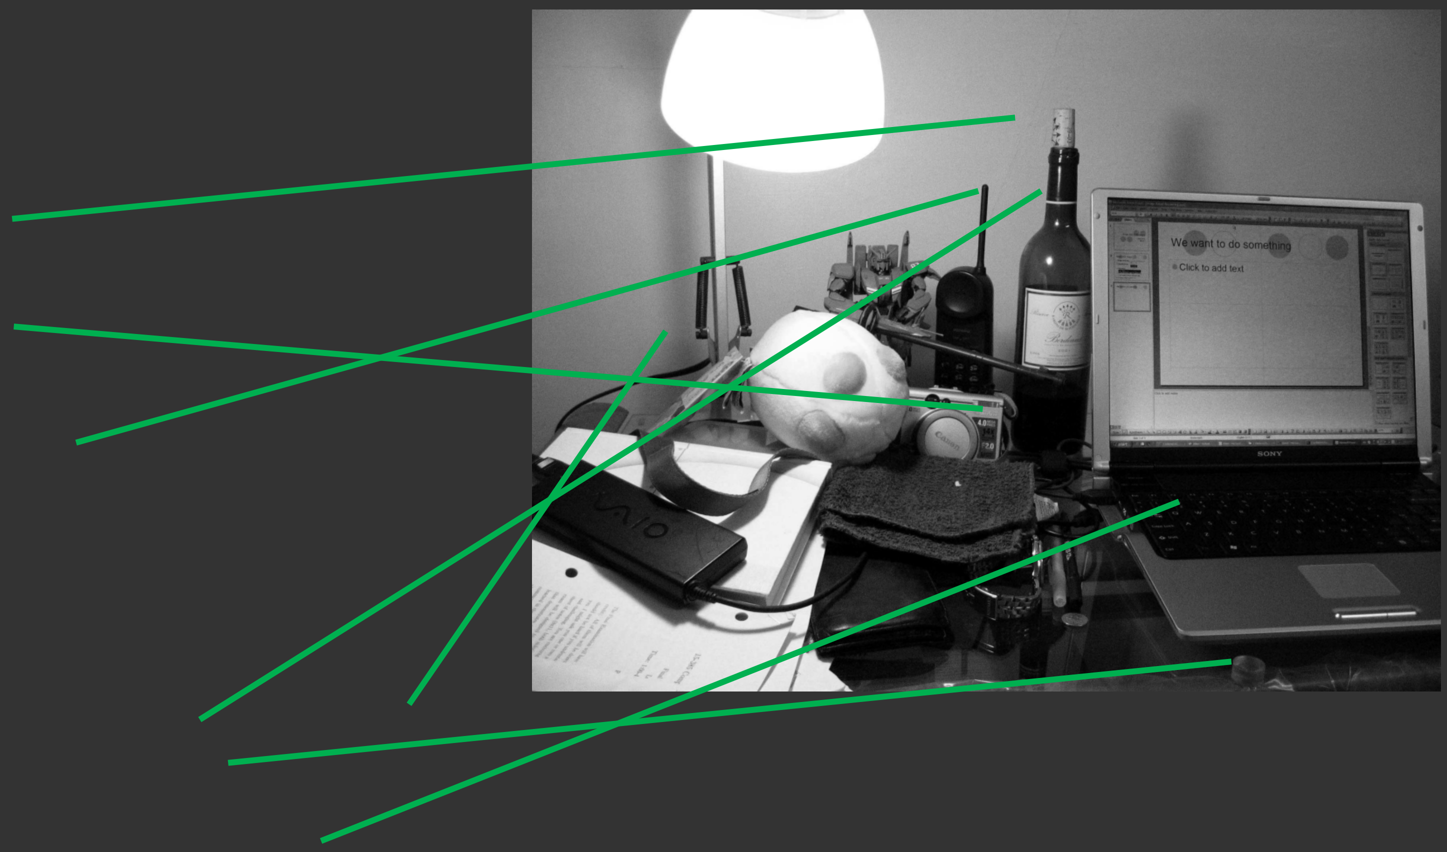

Sampling Rays by Taking Pictures

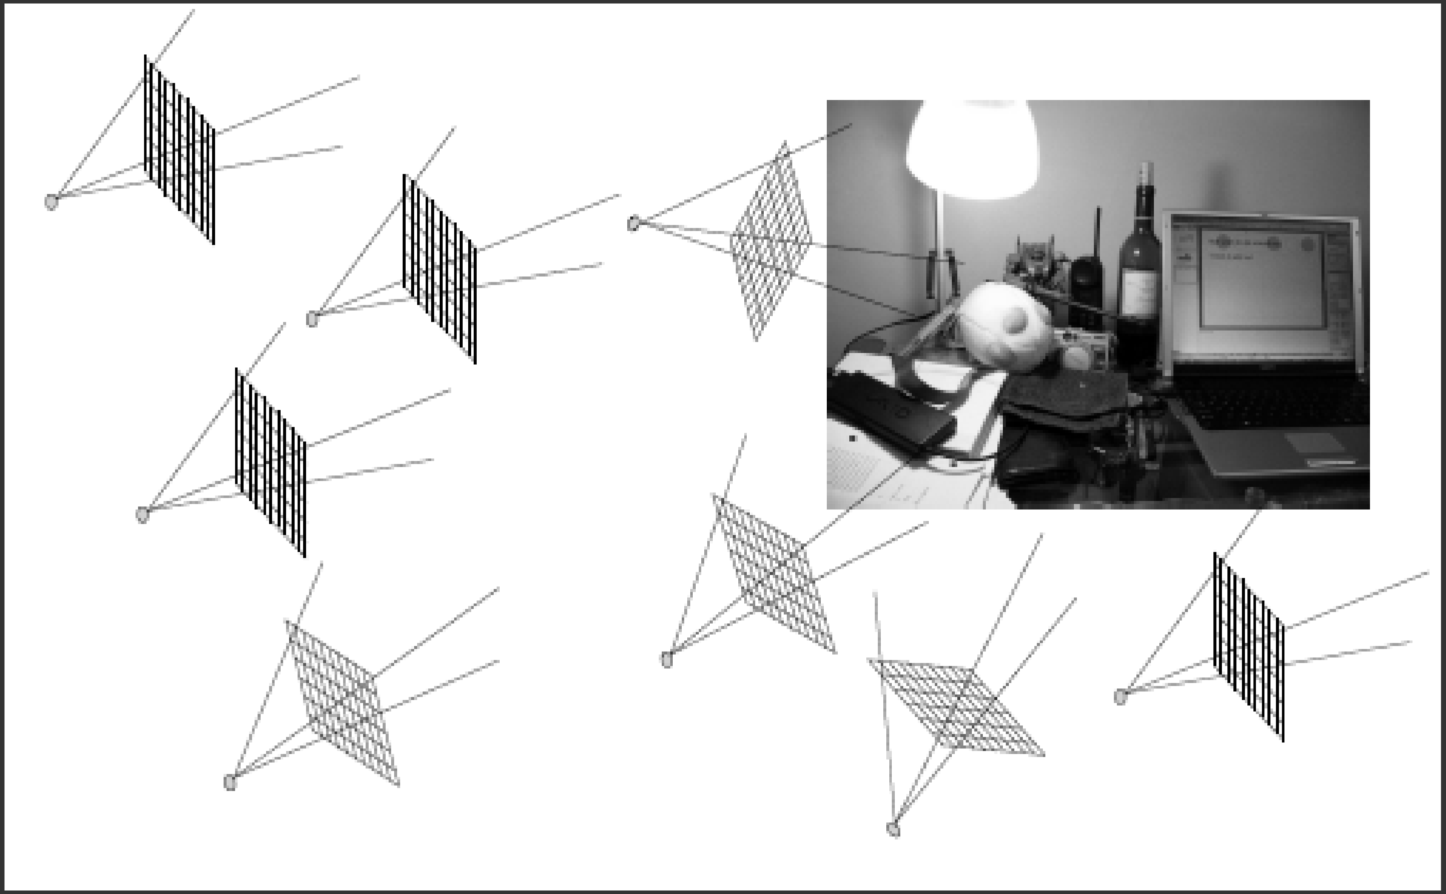

Rendering as Ray Resampling

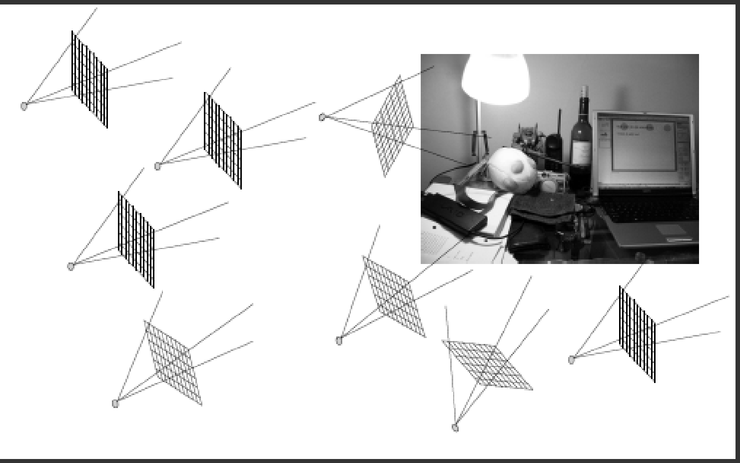

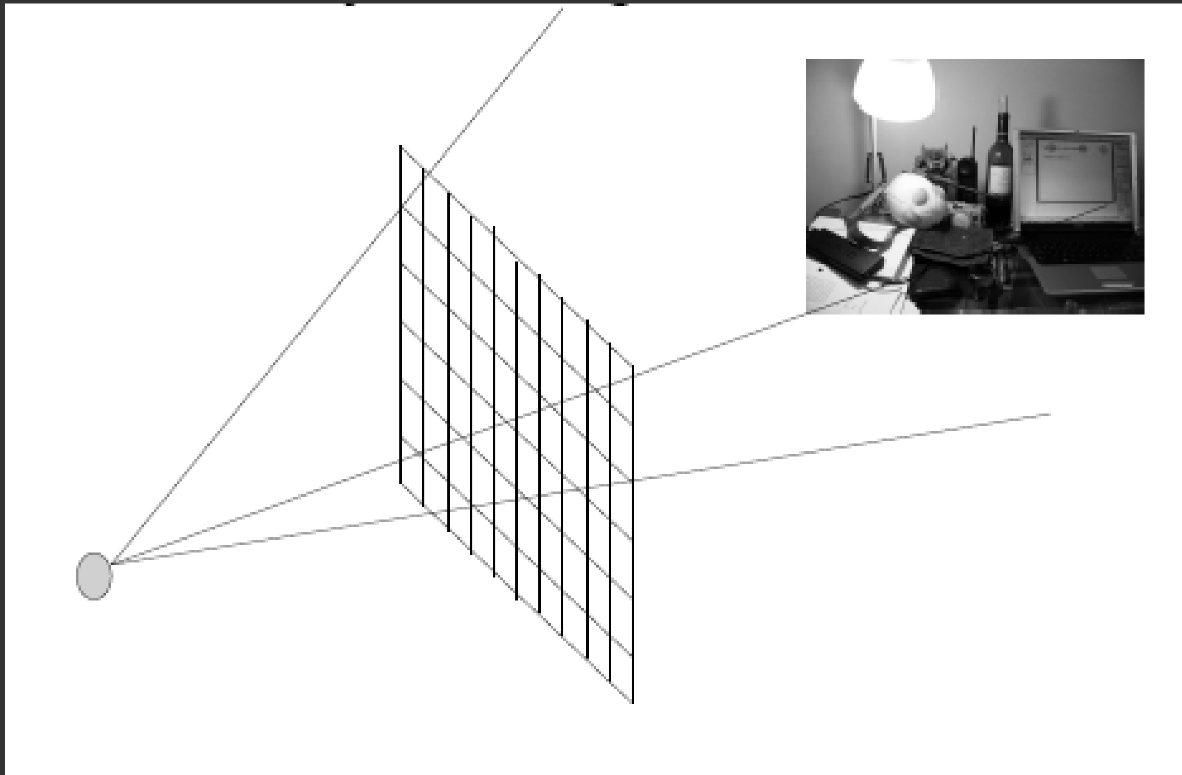

Ray space

- How to parameterize the ray space
- How to sample and resample rays

Two Plane Parameterization

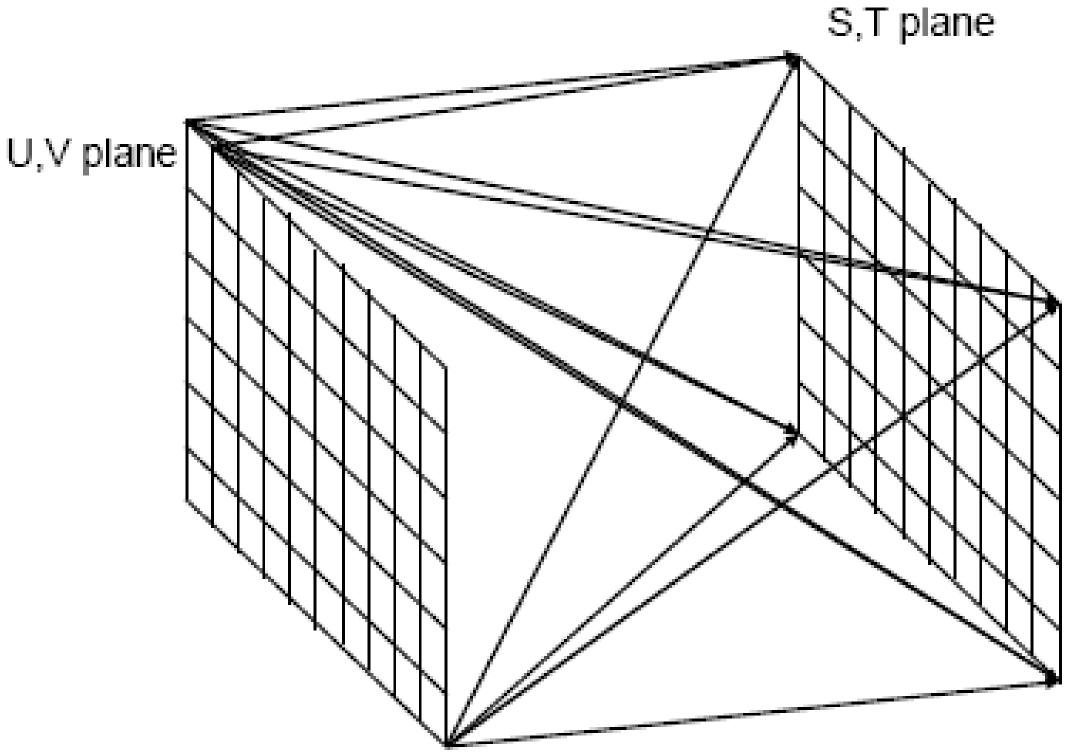

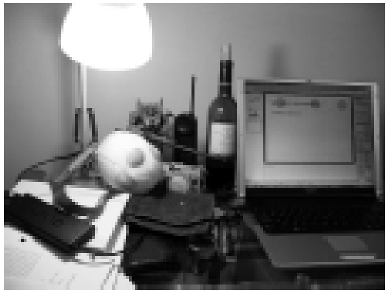

Stanford Camera Array

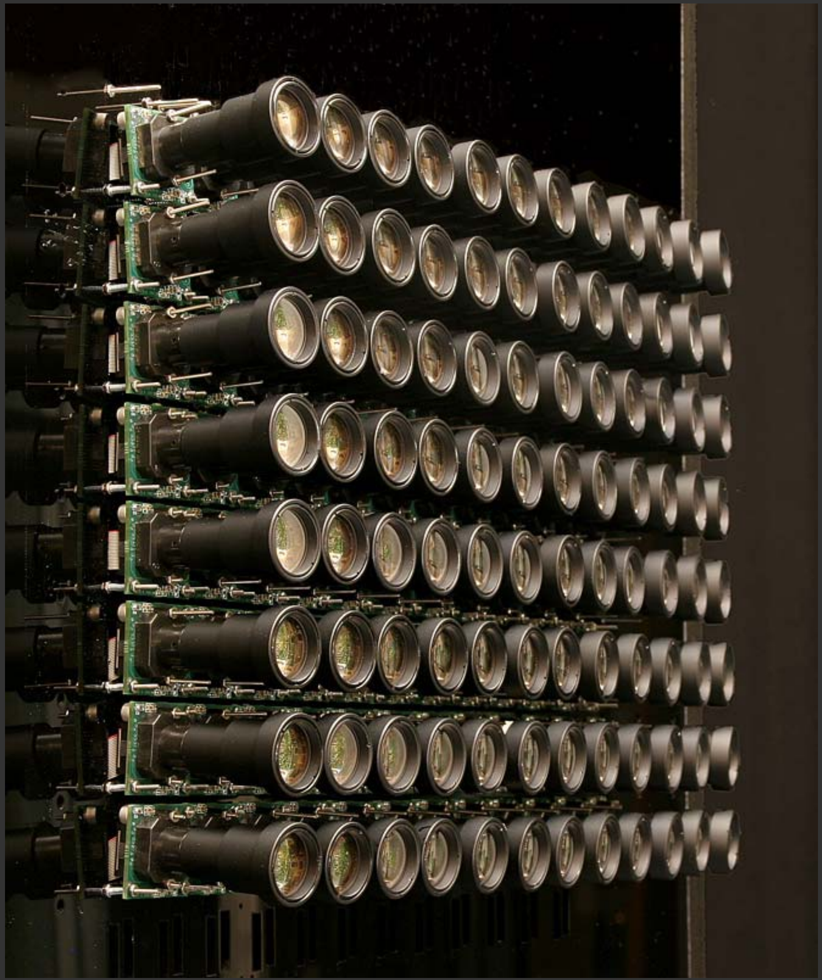

Light Field Rendering

- Very Fast

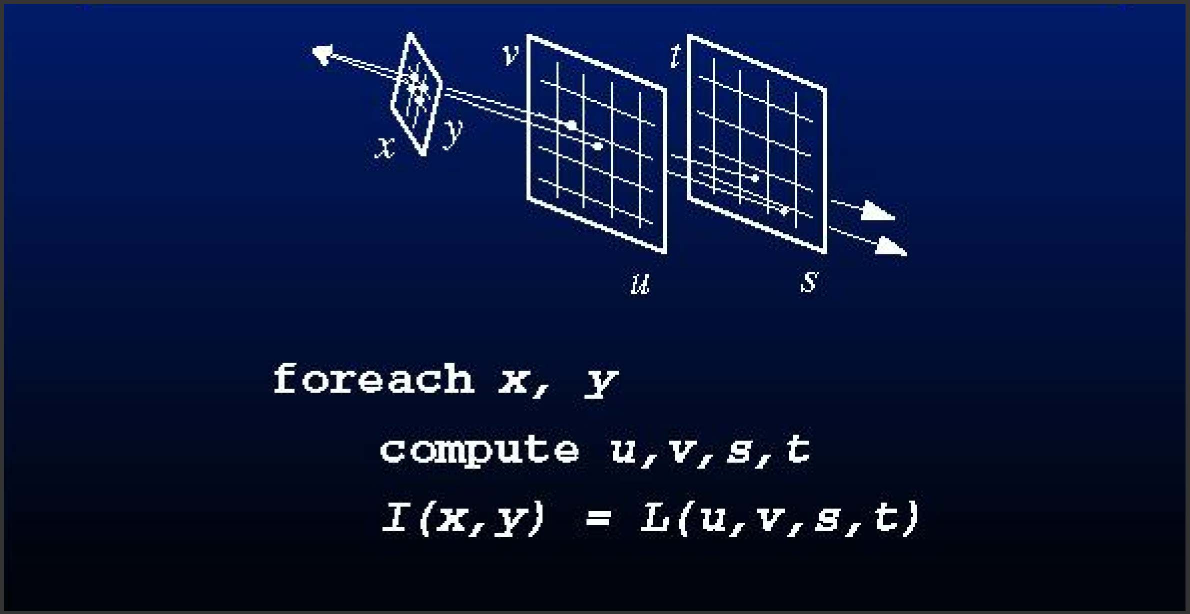

Light Field Rendering

- 4D interpolation

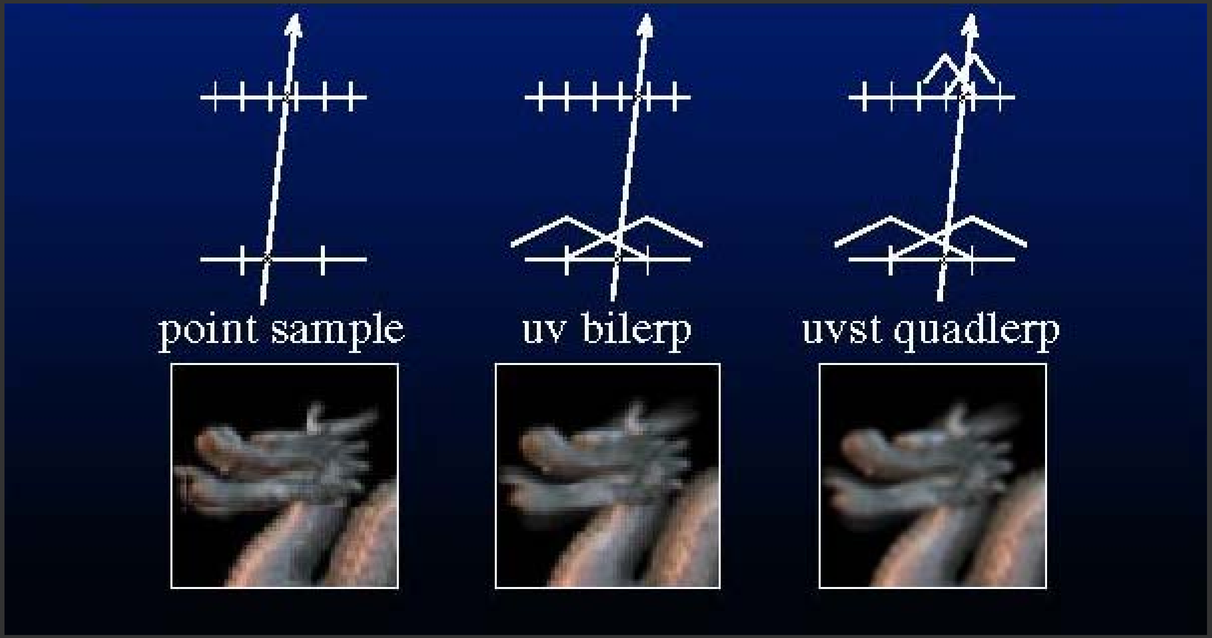

Dynamic Reparameterized Light Fields

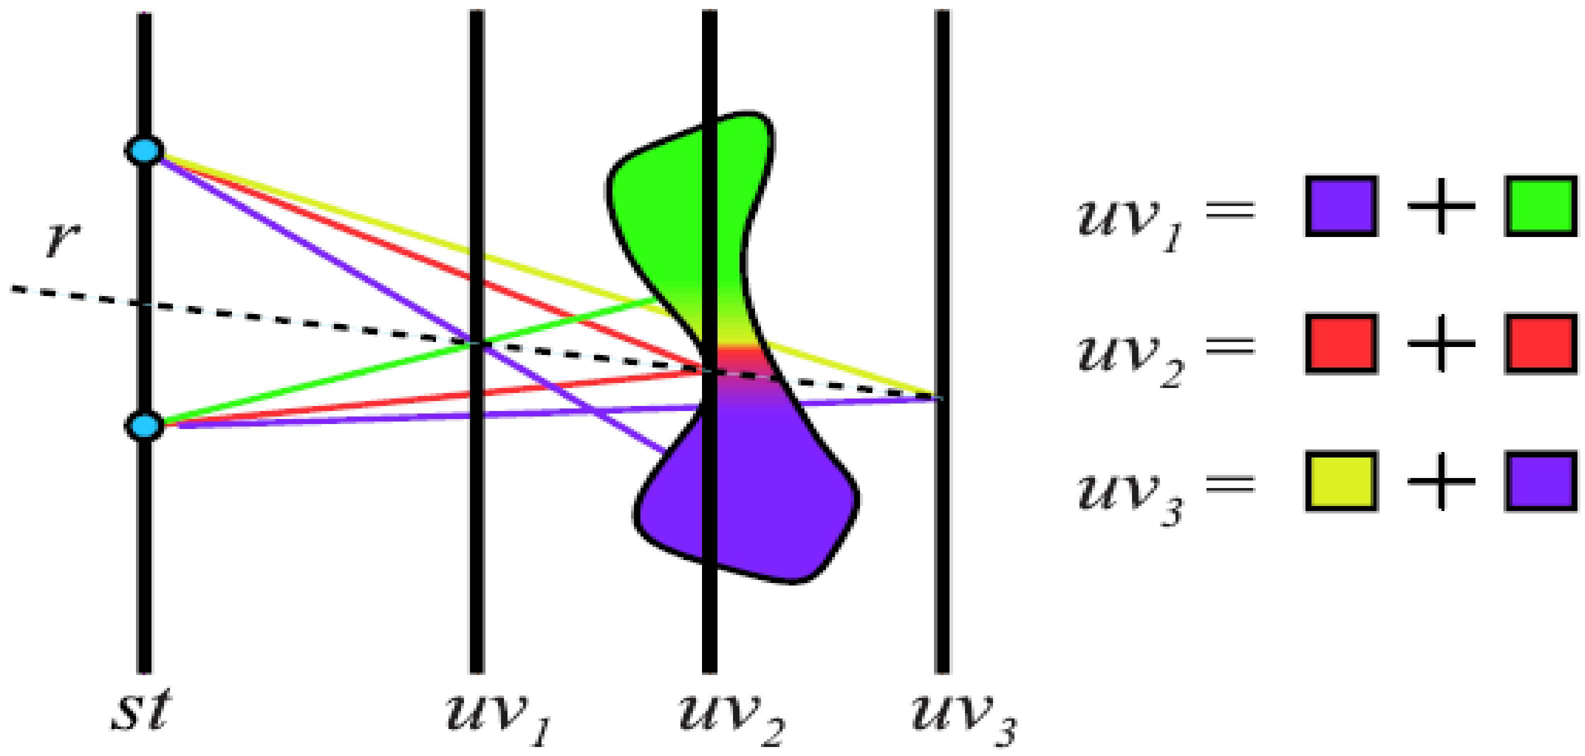

Dynamic Reparameterized Light Fields

- Move to desired new focal surface
- Create a new 4D space with new focal surface
- Recove ray with Reparameterization
- (u, v, s, t) => $(u, v, f, g)_F$

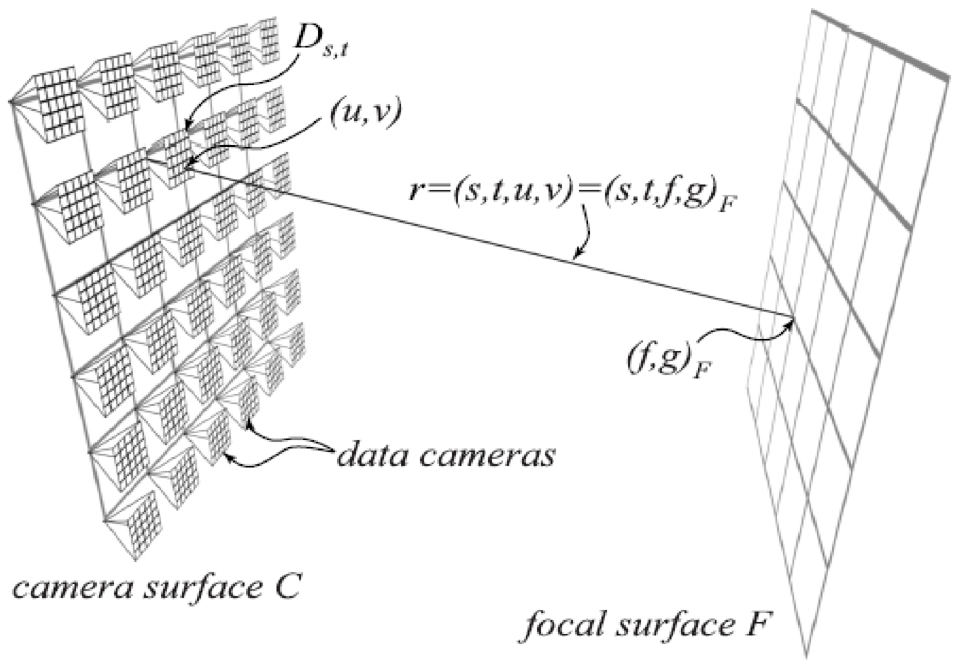

Dynamic Reparameterized Light Fields

- Recover ray r
- Resample from ray (s’, t’, f, g) and (s’’, t’’, f, g)
- Interpolation, reconstruction with filter, … , etc

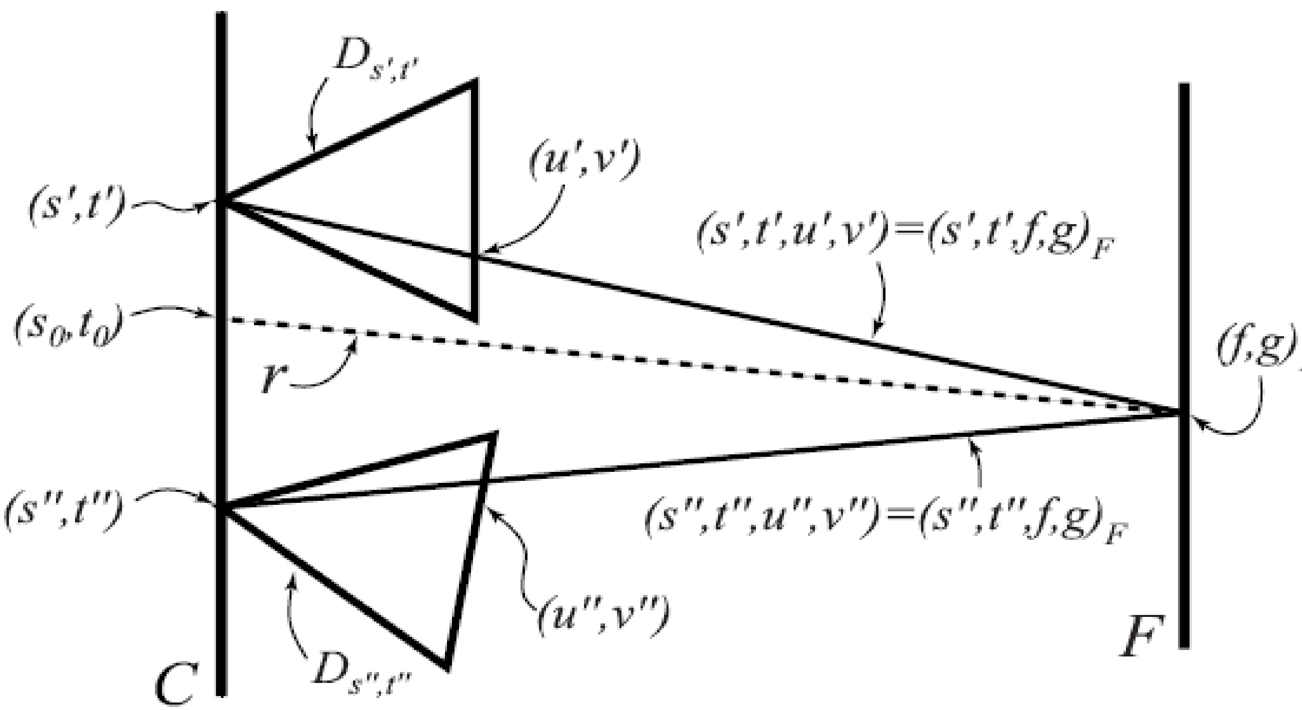

Dynamic Reparameterized Light Fields

- Change the shape of focal surface
- Gives focus on 3D object rather than planes

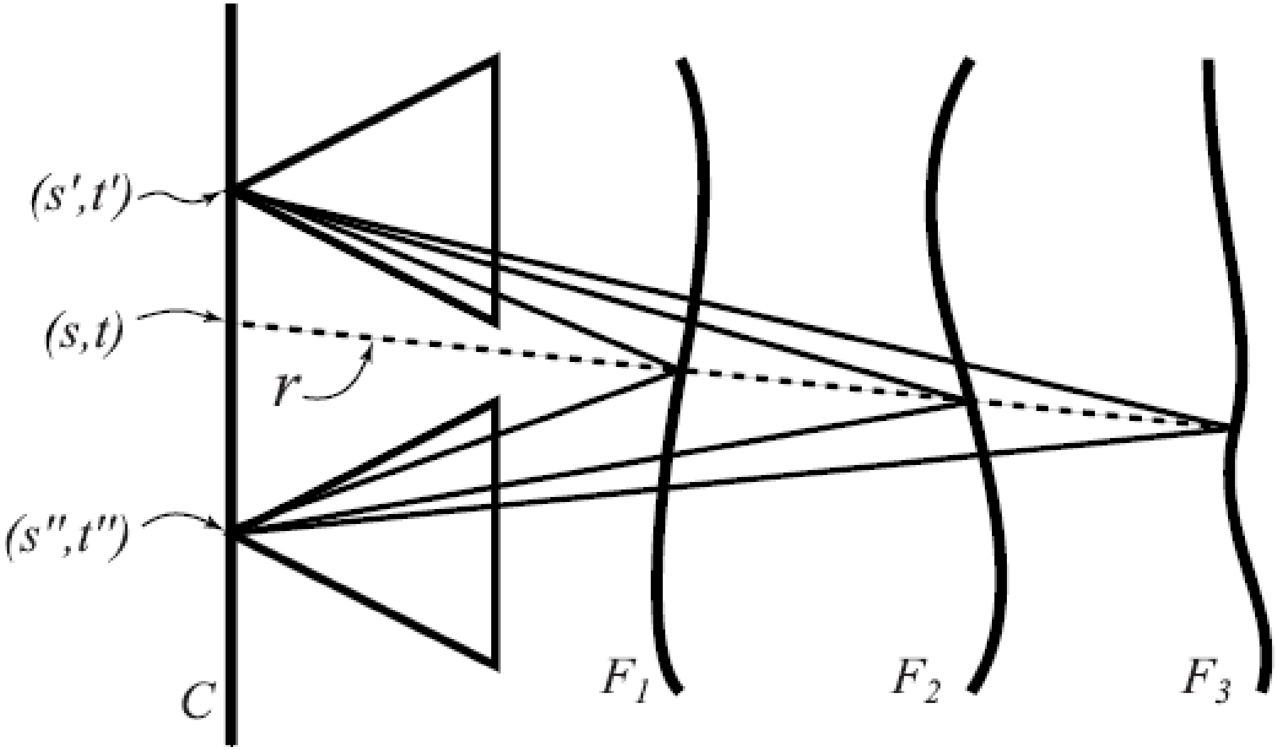

Dynamic Reparameterized Light Fields

Dynamic Reparameterized Light Fields

Variable Apertures

- Also can generate variable aperture
- Aperture
  - Control amount of light
  - Control depth of fields
- Aperture Filter:
  - Control how many cameras are used to resample a required ray
  - Larger apertures produce images with narrow range of focus

Aperture Filters

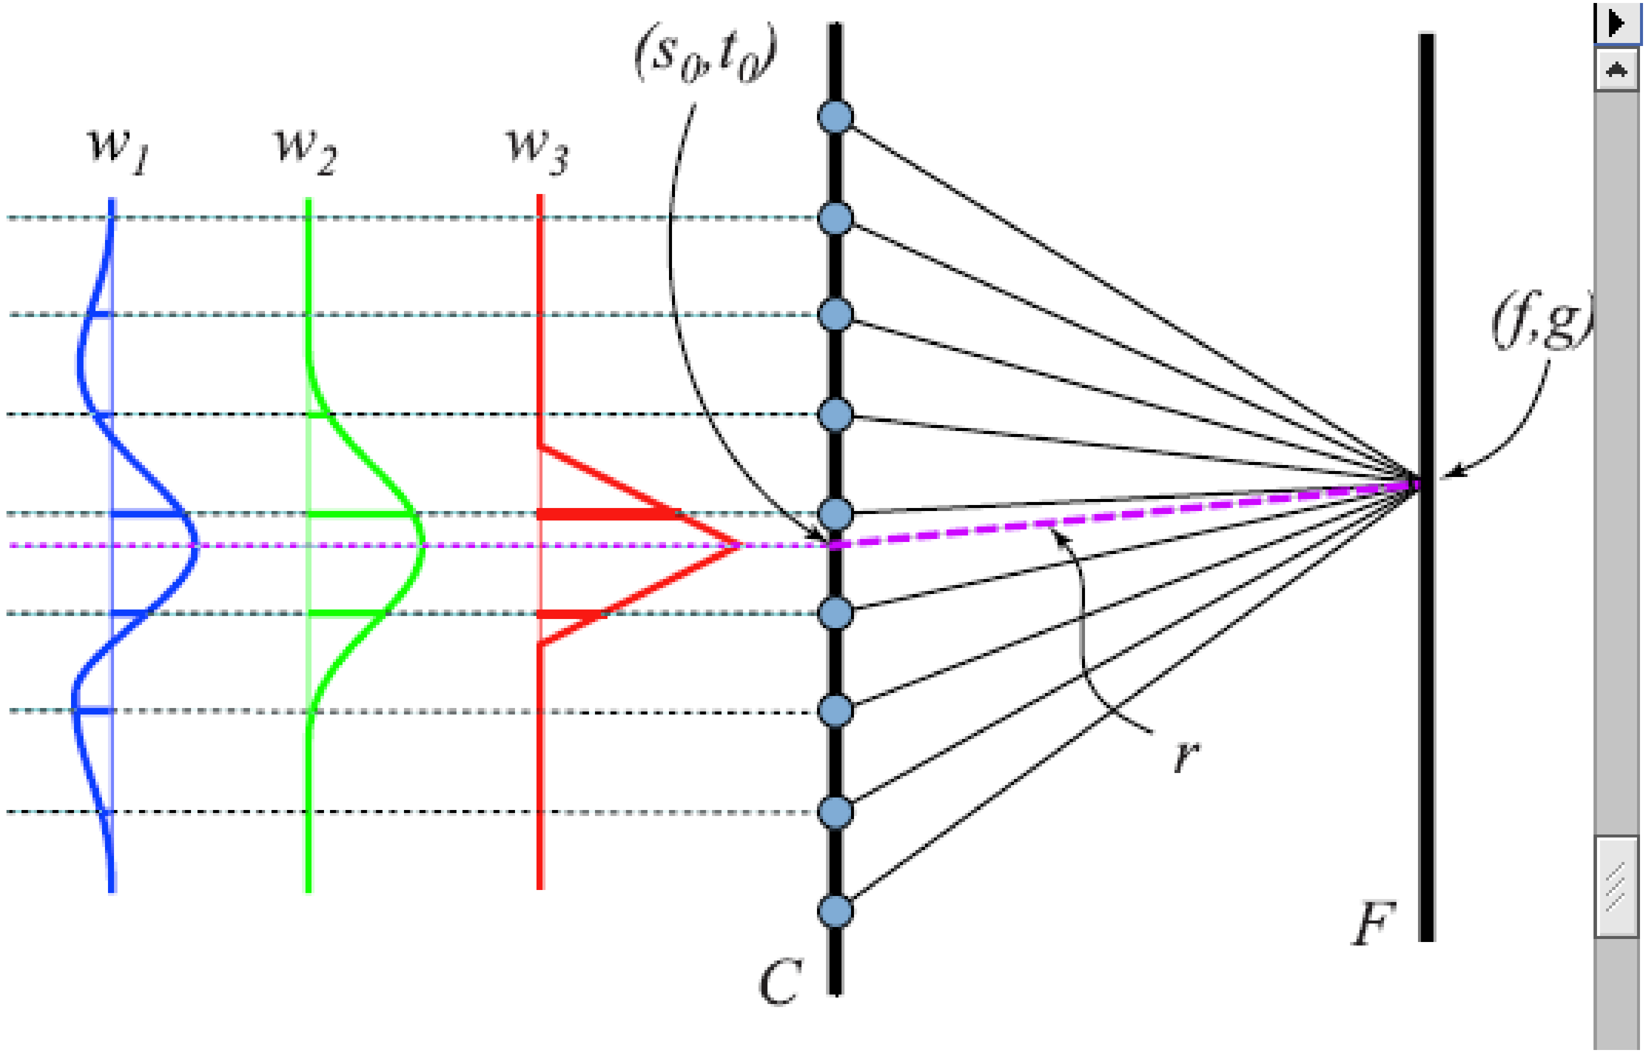

Variable Apertures

Variable Apertures

Stanford multi-camera array

- 640 × 480 pixels ×
  30 fps × 128 cameras

- synchronized timing
- continuous streaming
- flexible arrangement

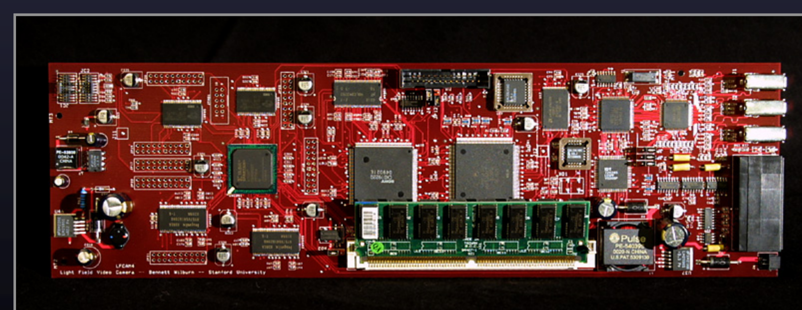

© 2005 Marc Levoy

Ways to use large camera arrays

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

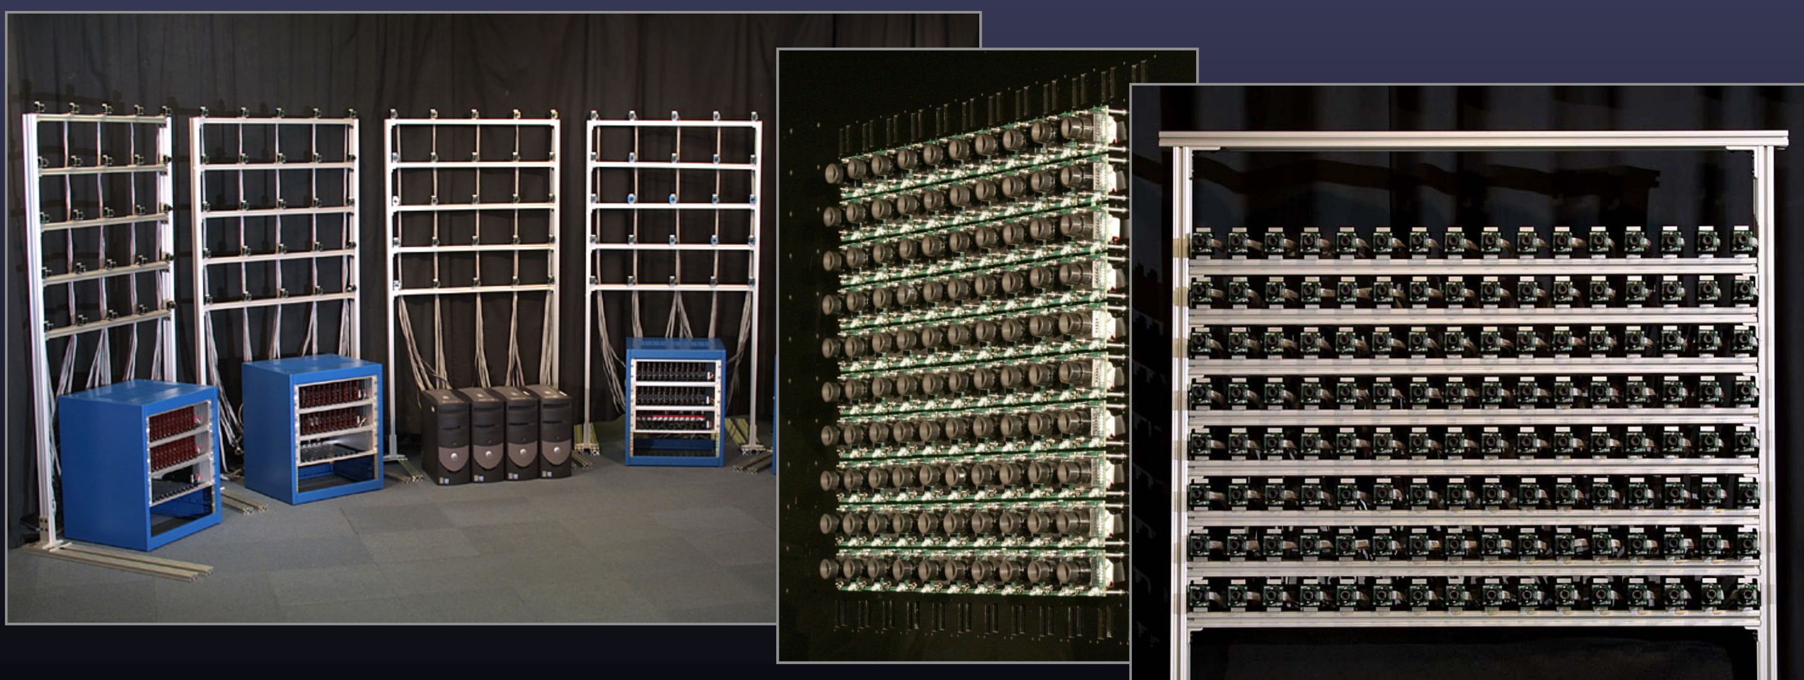

© 2005 Marc Levoy

Intermediate camera spacing: synthetic aperture photography

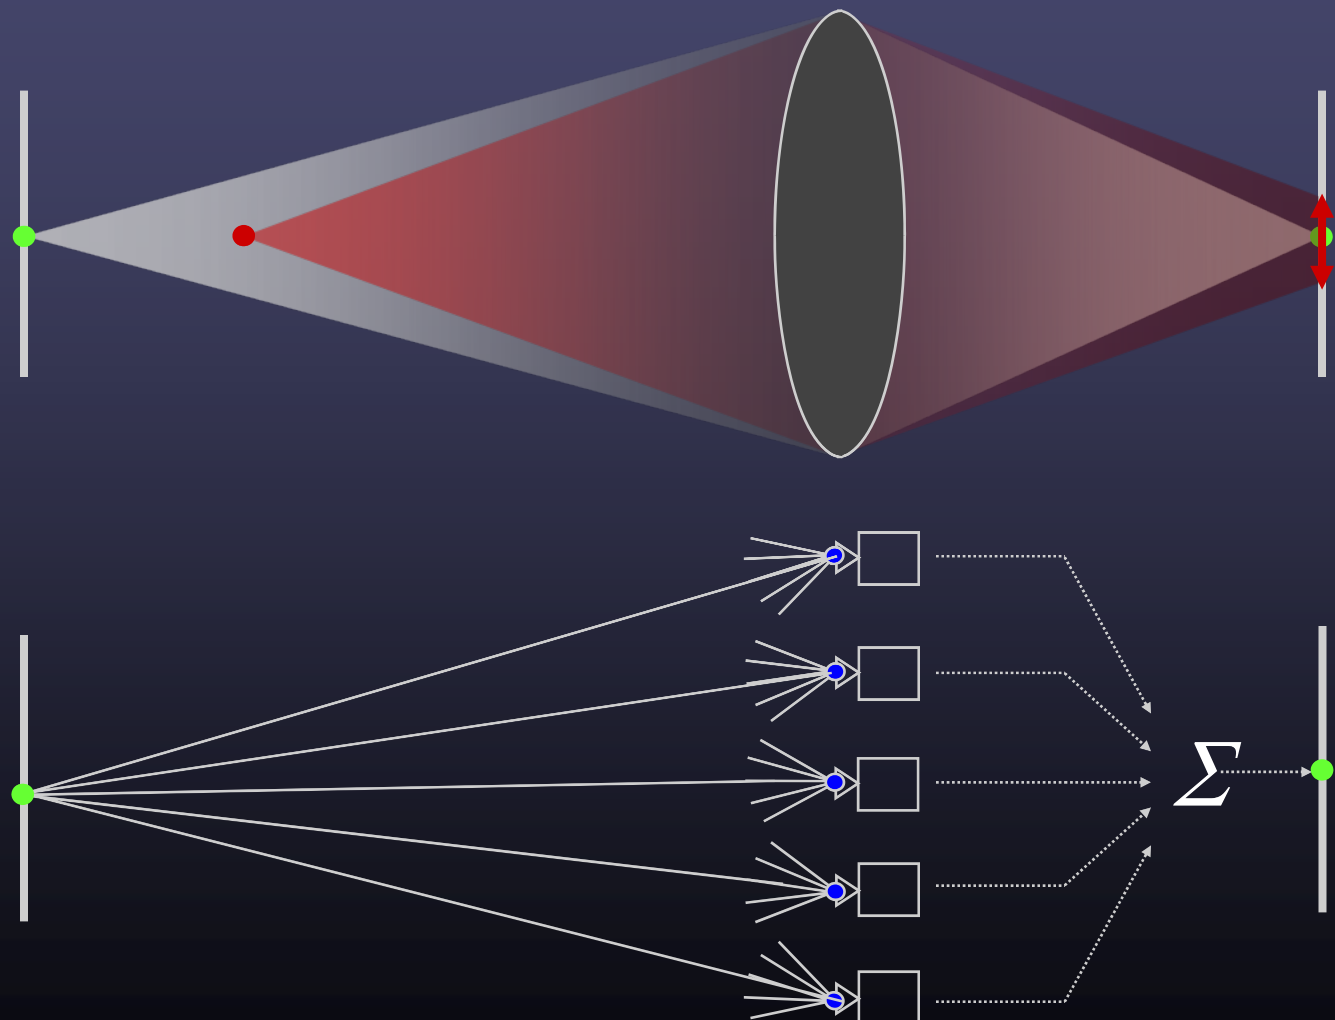

© 2005 Marc Levoy

Example using 45 cameras
[Vaish CVPR 2004]

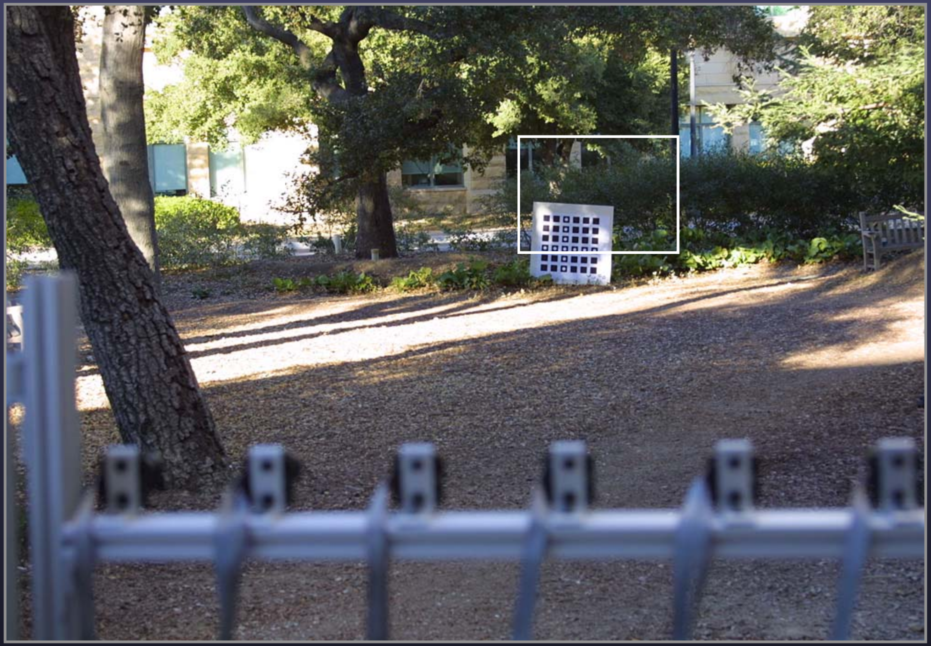

© 2005 Marc Levoy


Tiled camera array

Can we match the image quality of a cinema camera?

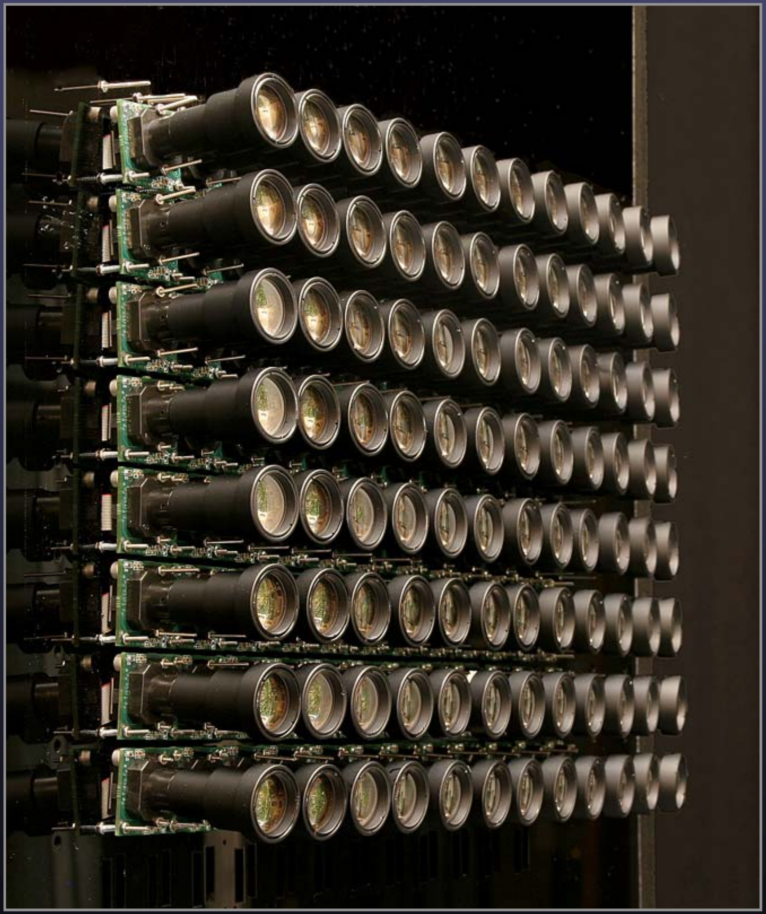

- world's largest video camera
- no parallax for distant objects
- poor lenses limit image quality
- seamless mosaicing isn't hard

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?

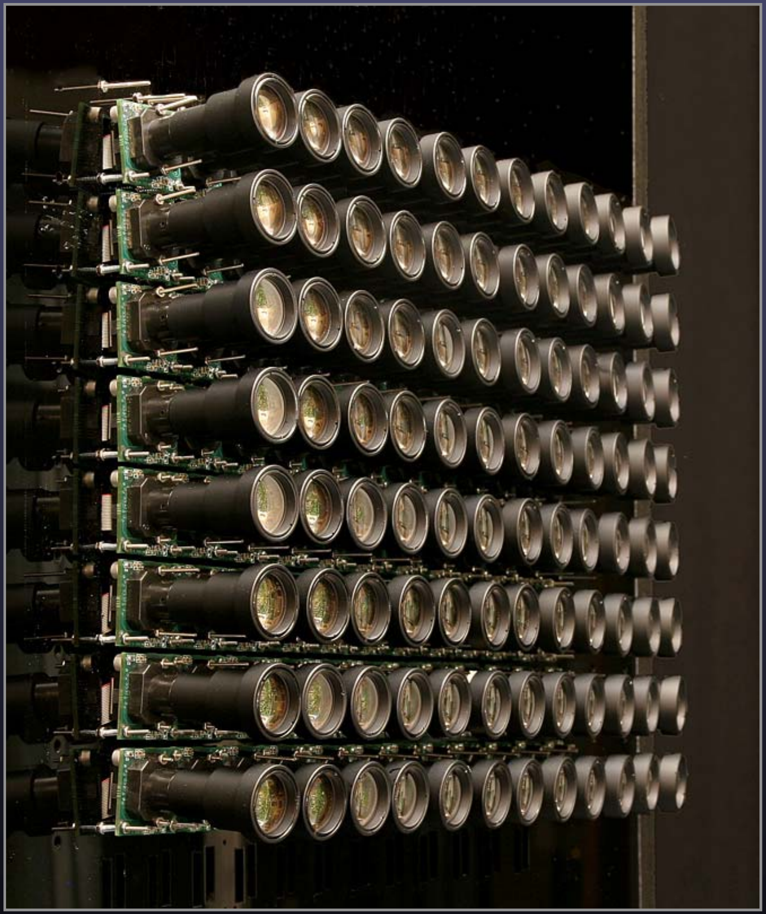

- world's largest video camera
- no parallax for distant objects
- poor lenses limit image quality
- seamless mosaicing isn't hard
- per-camera exposure metering
- HDR within and between tiles

same exposure in all cameras

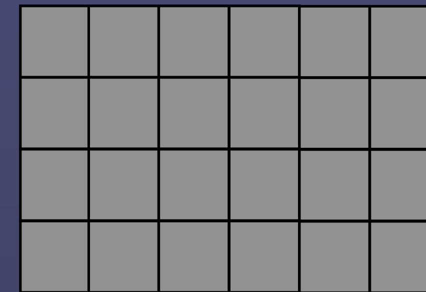

individually metered

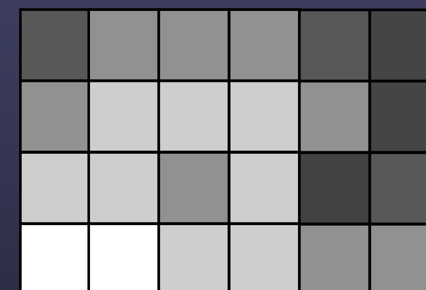

checkerboard of exposures

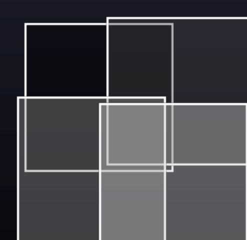

High-performance photography as multi-dimensional sampling

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

© 2005 Marc Levoy

Light field photography using a handheld plenoptic camera

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

What’s wrong with conventional cameras?

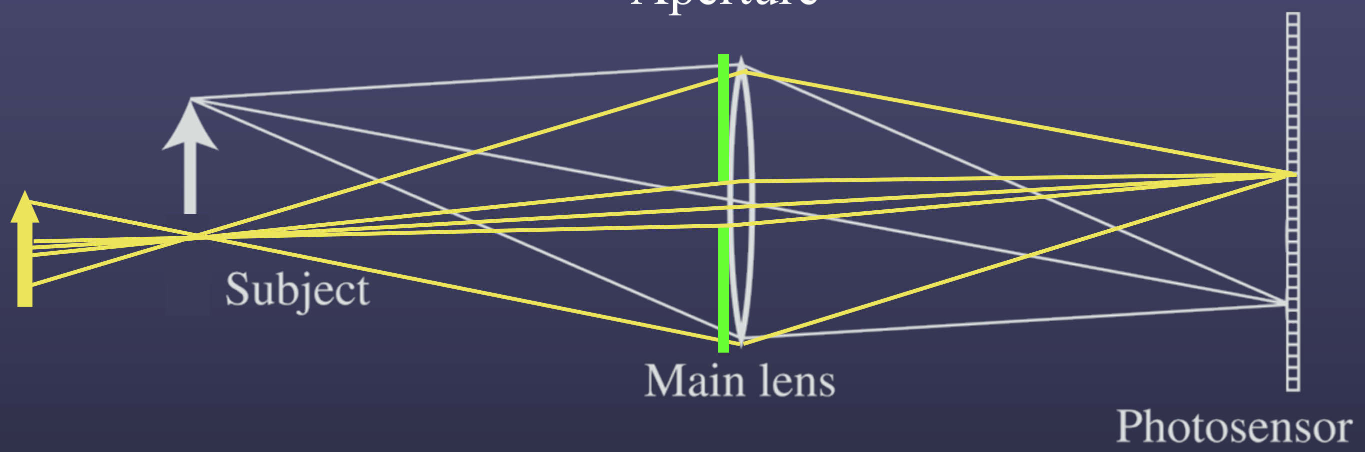

© 2005 Marc Levoy

Capture the light field inside a camera

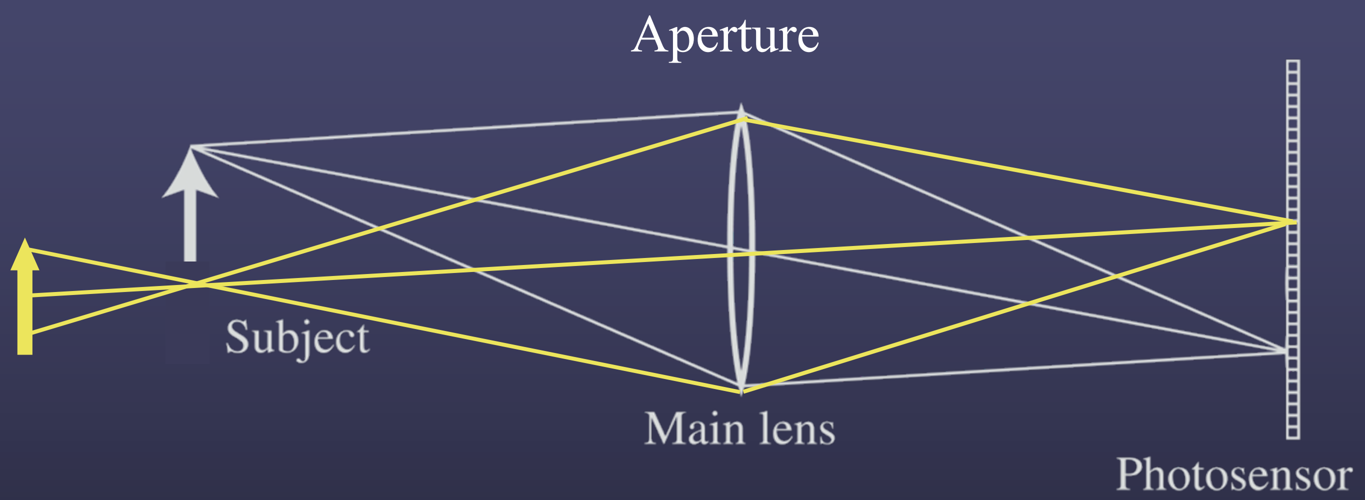

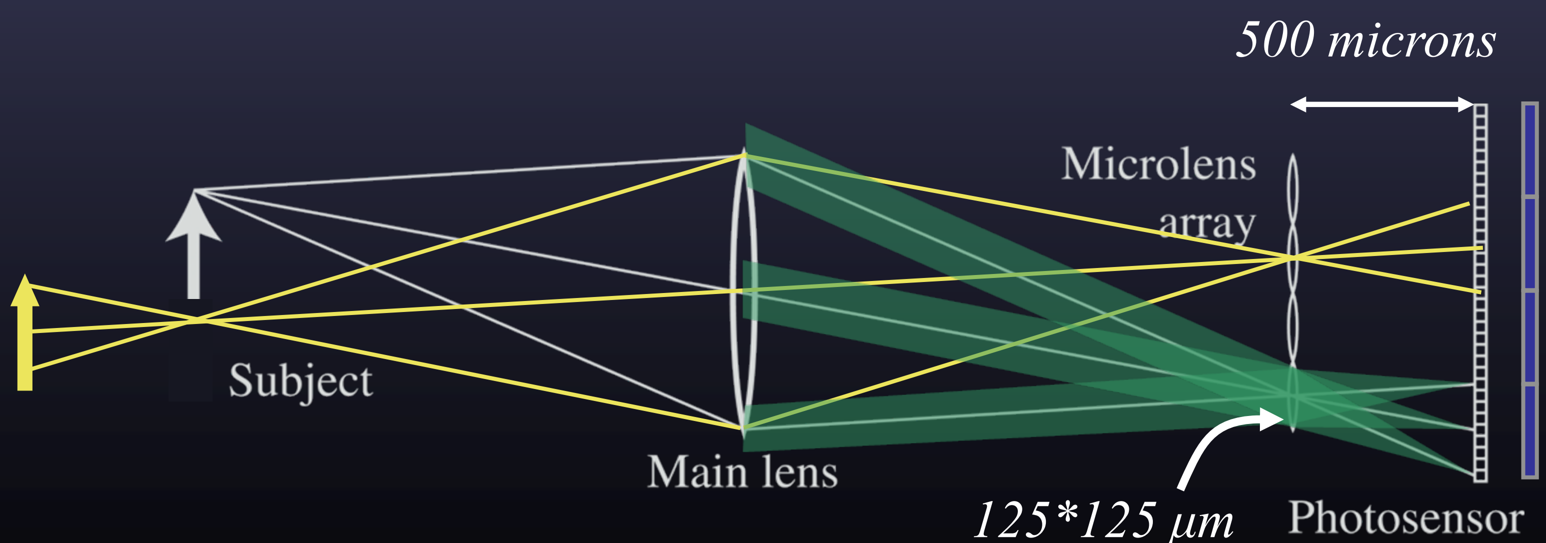

© 2005 Marc Levoy

Conventional versus light field camera

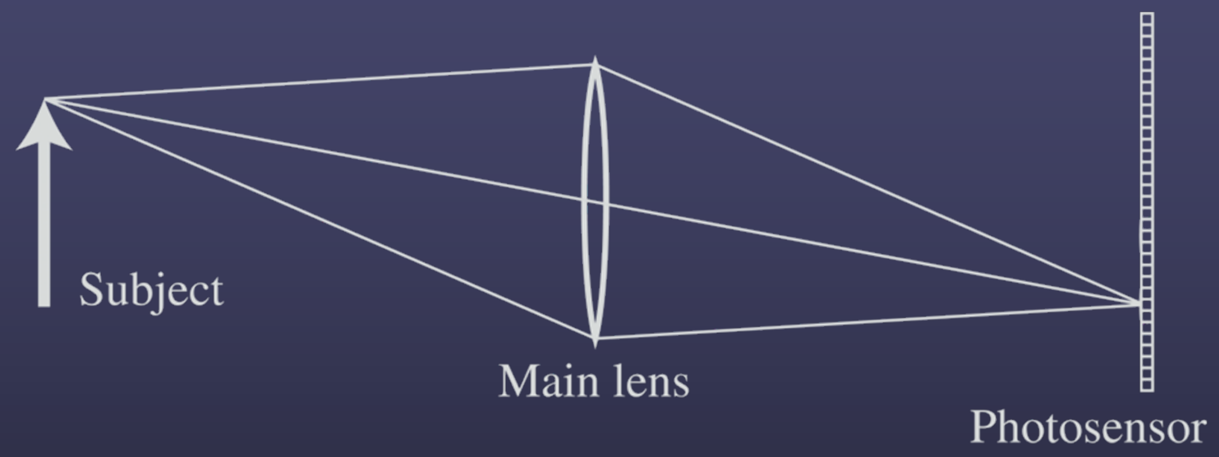

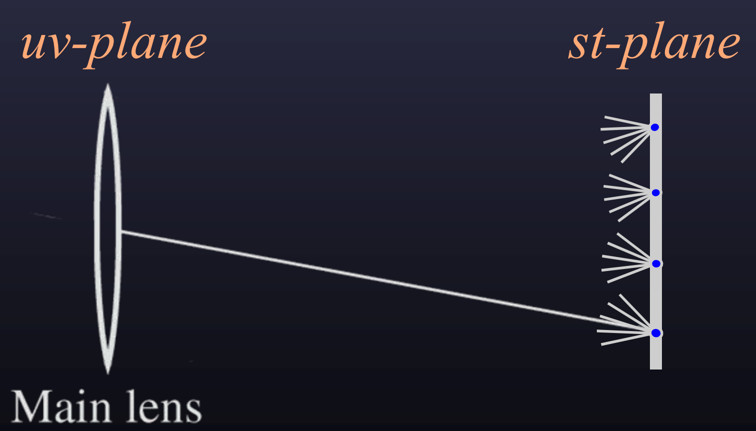

© 2005 Marc Levoy

Conventional versus light field camera

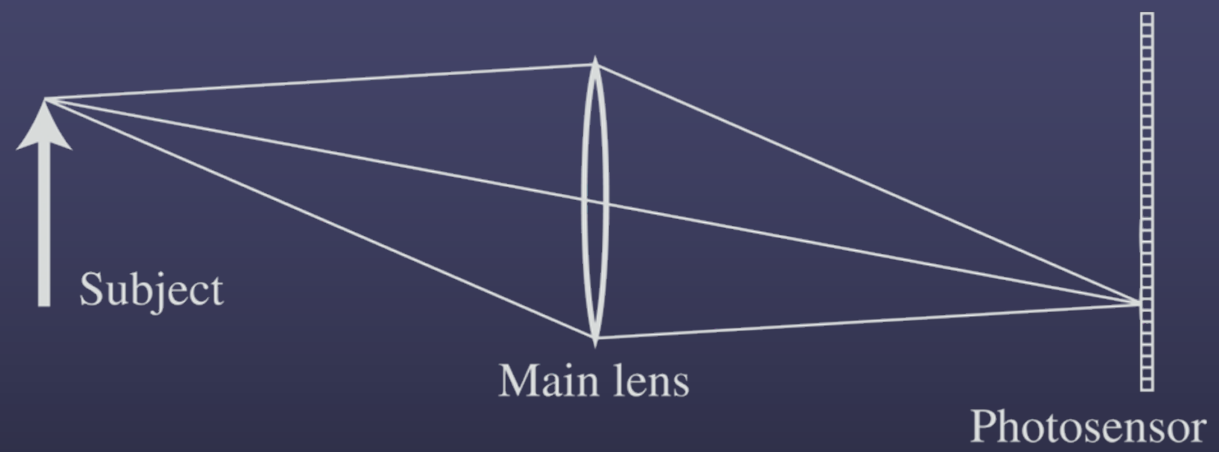

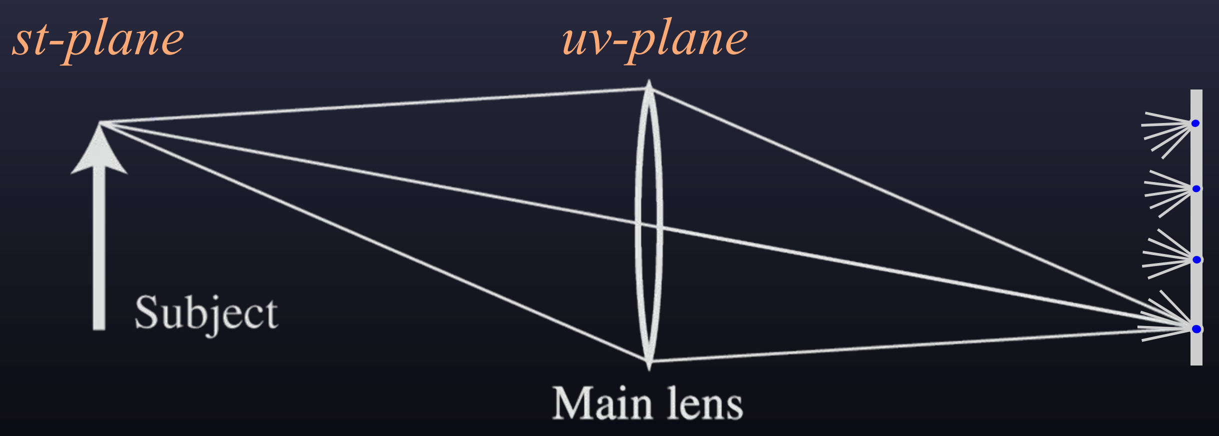

© 2005 Marc Levoy

Prototype camera

Contax medium format camera

Kodak 16-megapixel sensor

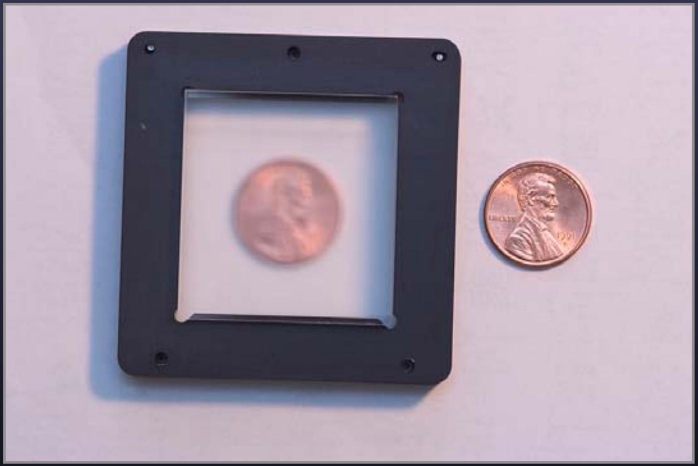

Adaptive Optics microlens array

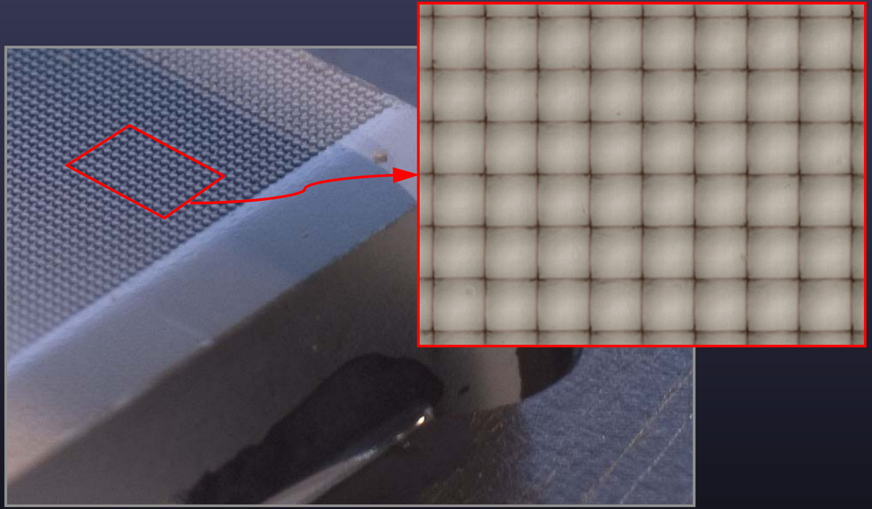

125μ square-sided microlenses

4000 × 4000 pixels ÷ 292 × 292 lenses = 14 × 14 pixels per lens

Light Field in a Single Exposure

© 2005 Marc Levoy

Light Field in a Single Exposure

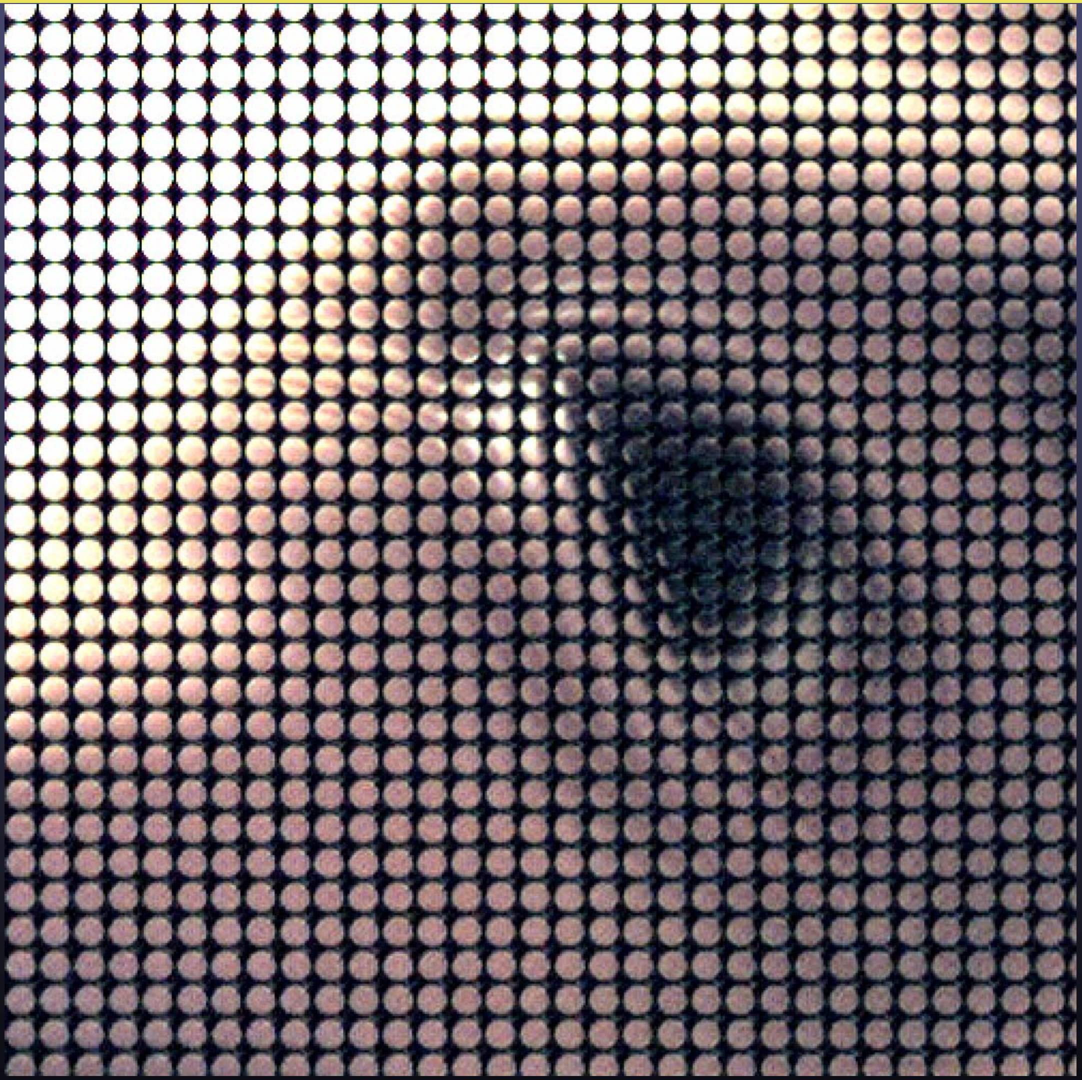

© 2005 Marc Levoy

Light field inside a camera body

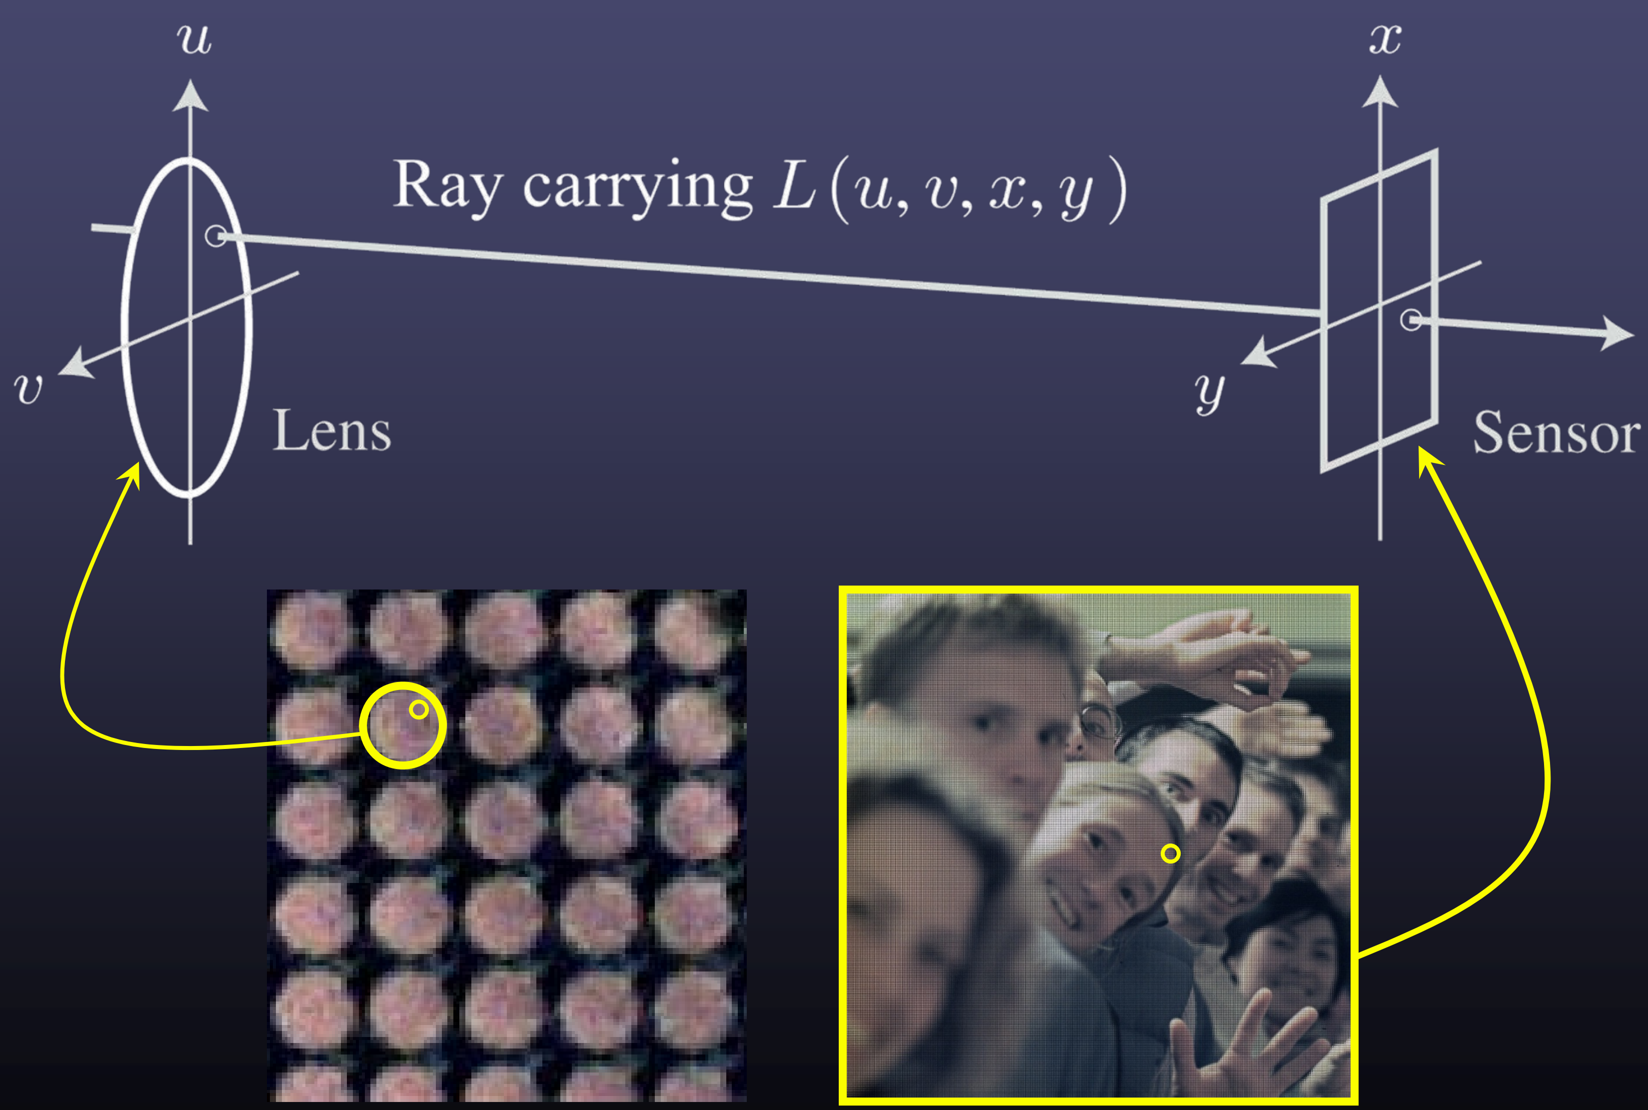

© 2005 Marc Levoy

Digitally stopping-down

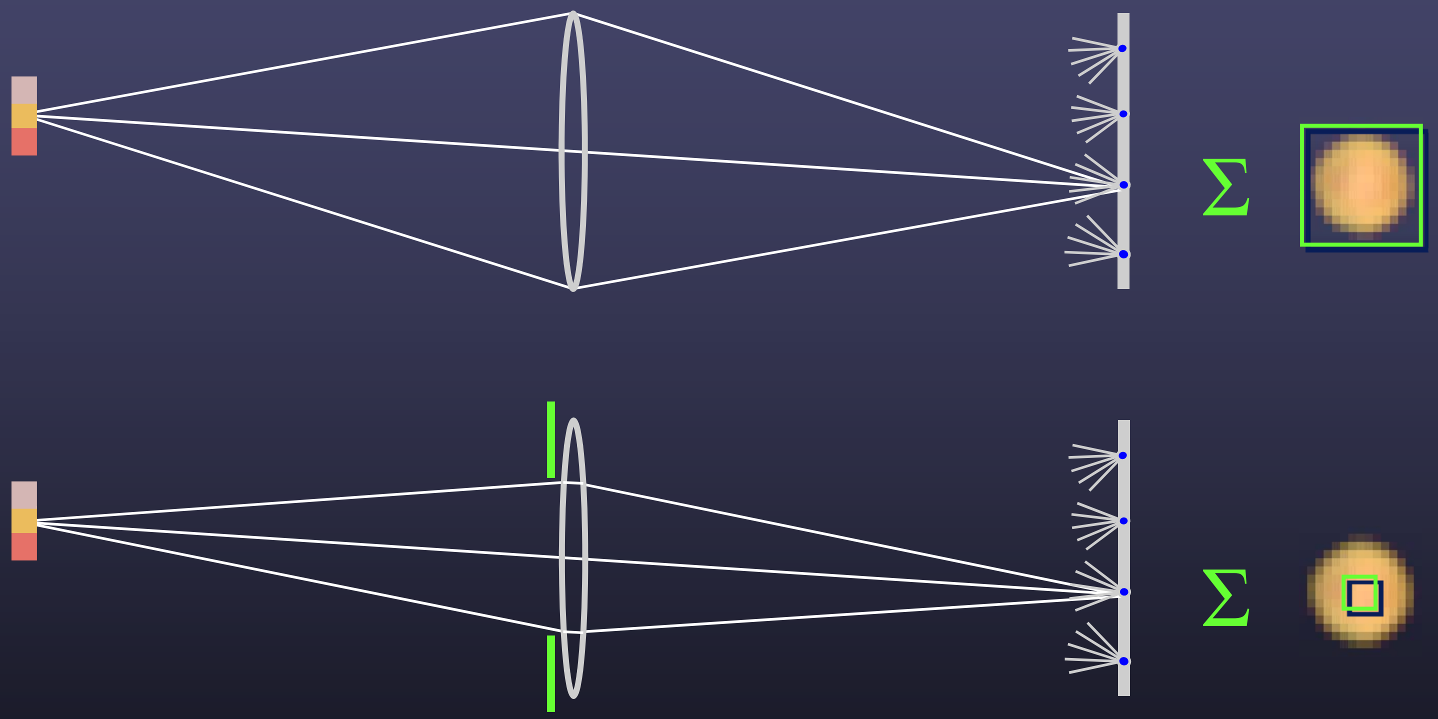

- stopping down = summing only the central portion of each microlens

© 2005 Marc Levoy

Digital refocusing

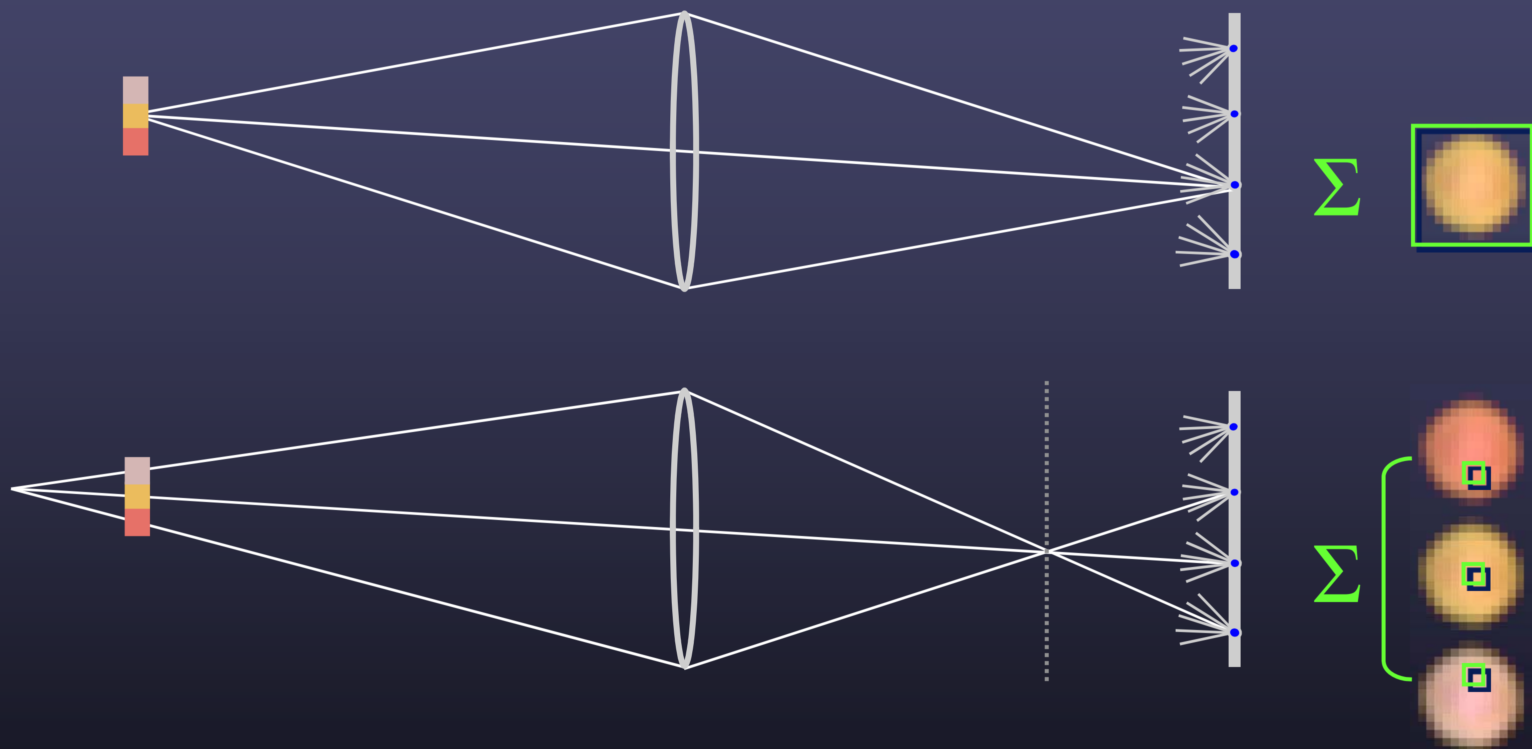

- refocusing = summing windows extracted from several microlenses

© 2005 Marc Levoy

Example of digital refocusing

© 2005 Marc Levoy

Refocusing portraits

© 2005 Marc Levoy

Action photography

Focusing through a splash of water

© 2005 Marc Levoy

Extending the depth of field

conventional photograph,
main lens at *f* / 4

conventional photograph,
main lens at *f* / 22

Scene-dependent focal plane

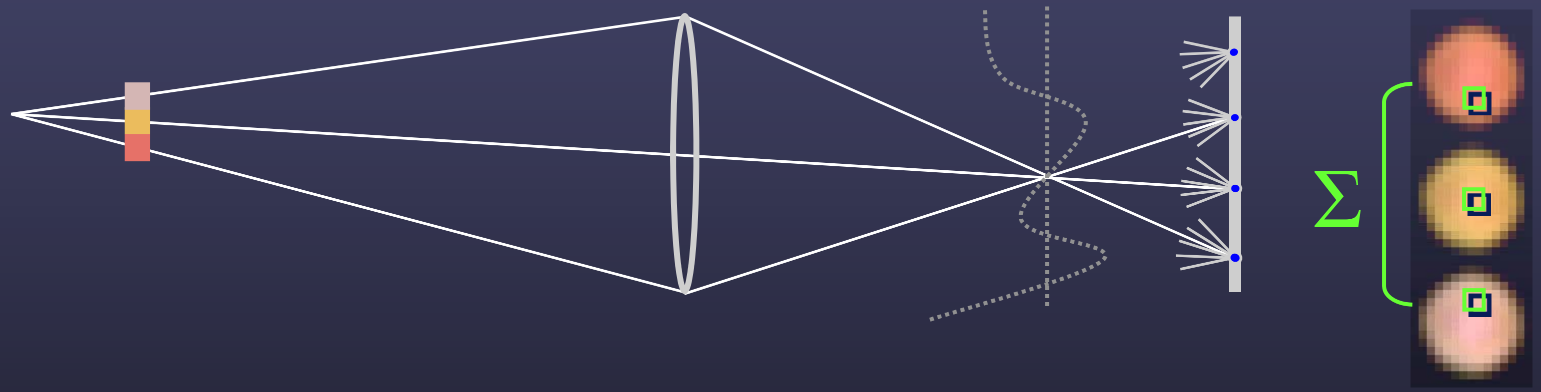

Depth from focus problem

Interactive solution [Agarwala 2004]

© 2005 Marc Levoy

Extending the depth of field

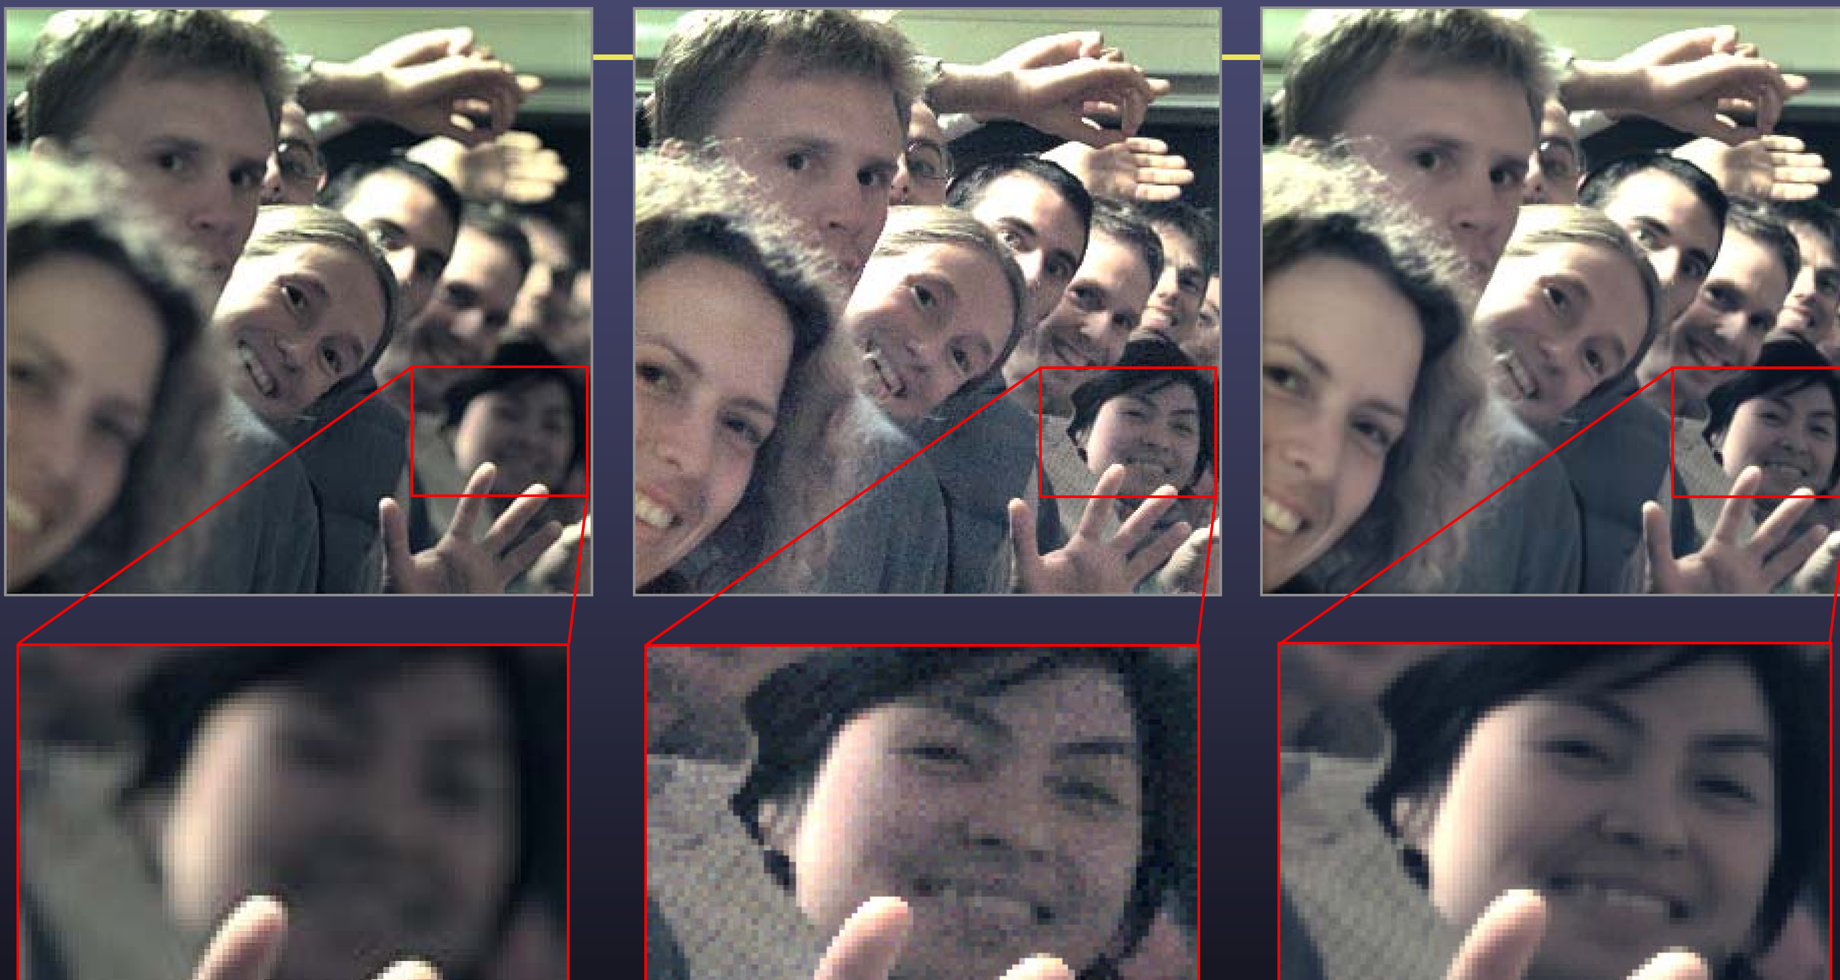

conventional photograph, main lens at $f/4$

conventional photograph, main lens at $f/22$

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

© 2005 Marc Levoy

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

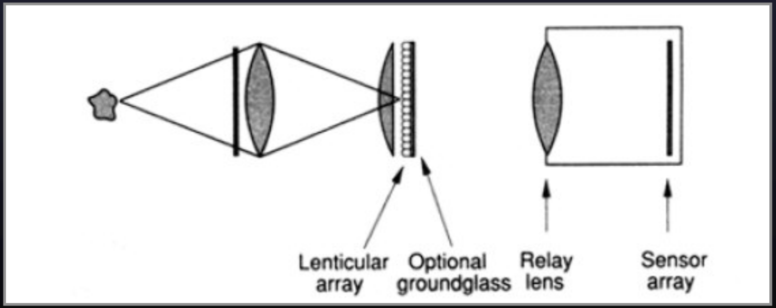

© 2005 Marc Levoy

Digitally moving the observer

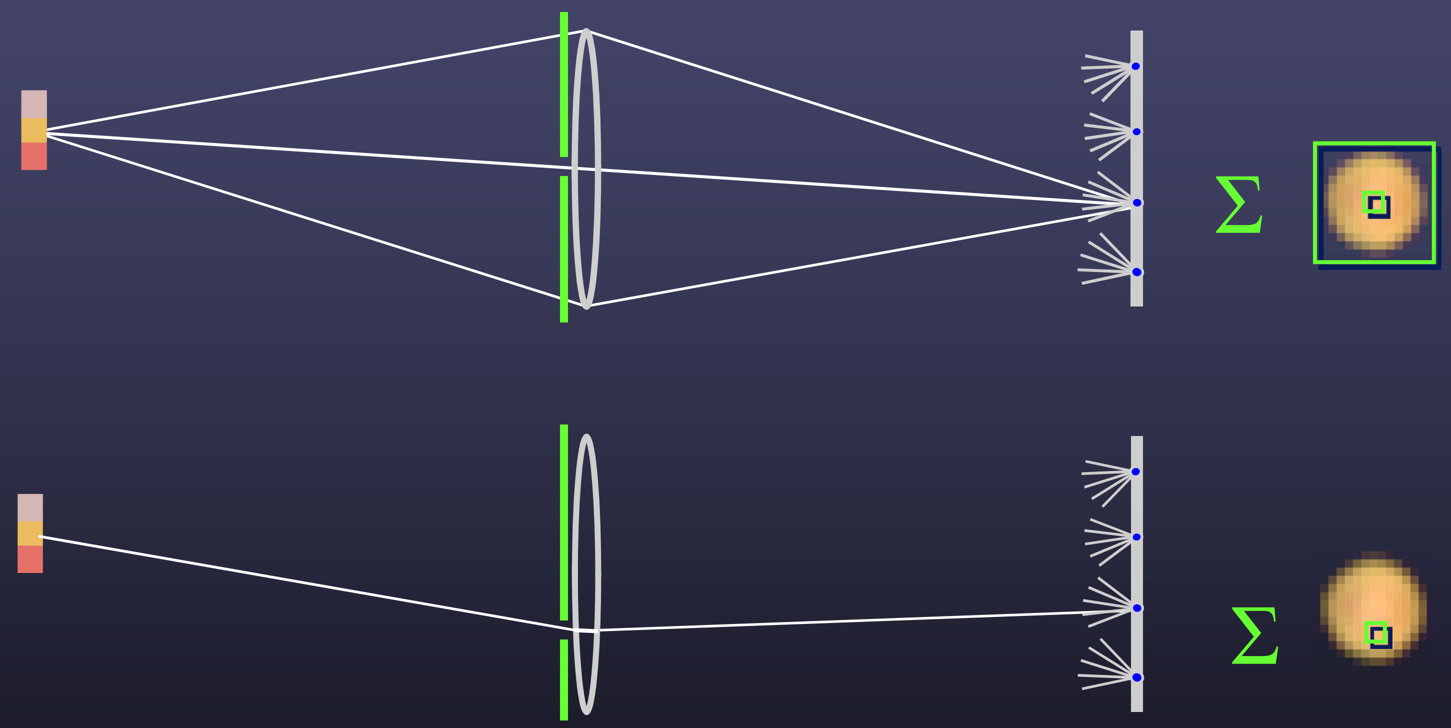

- moving the observer = moving the window we extract from the microlenses

© 2005 Marc Levoy

Example of moving the observer

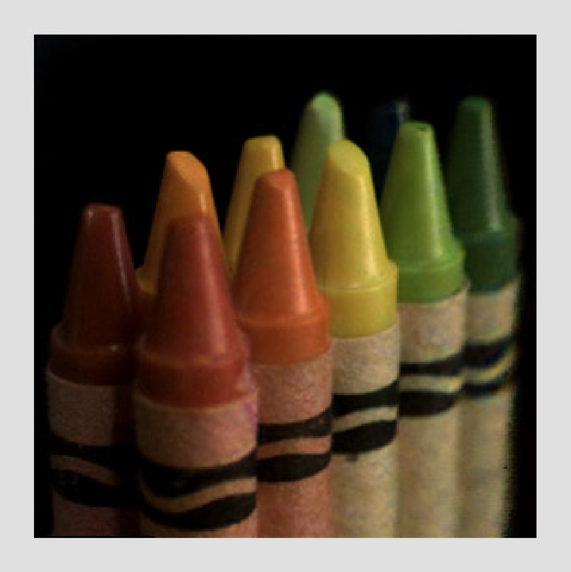

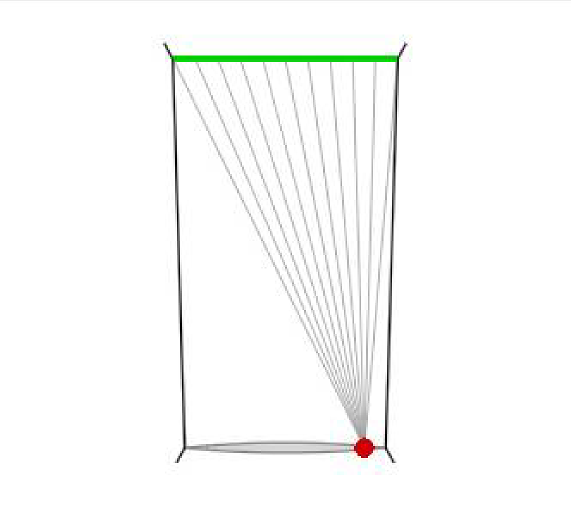

© 2005 Marc Levoy

Moving backward and forward

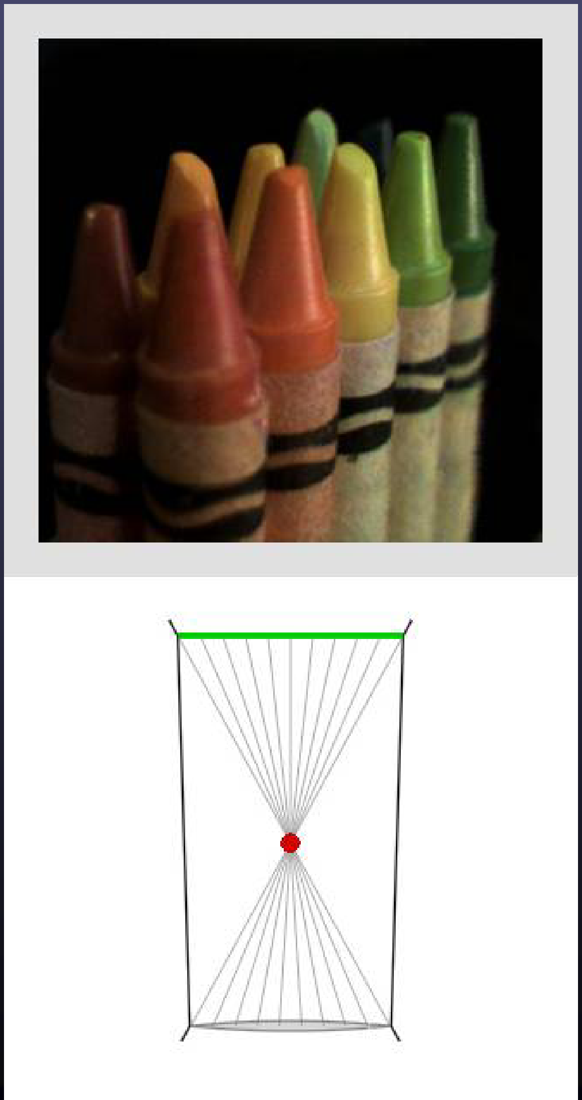

© 2005 Marc Levoy

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
- Application in microscope

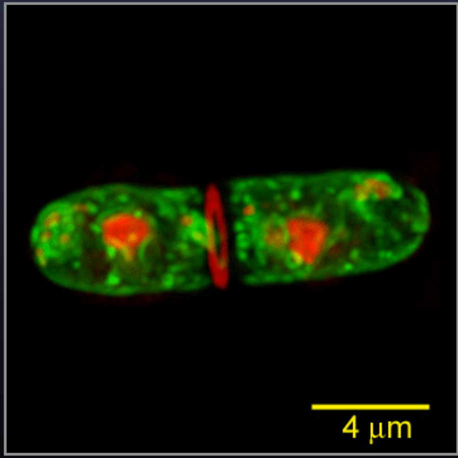

© 2005 Marc Levoy

Vision Sensing

Multi-View Stereo for Community Photo Collections
Michael Goesele, et al, ICCV 2007

Venus de Milo

The Digital Michelangelo Project, Stanford

How to sense 3D very accurately?

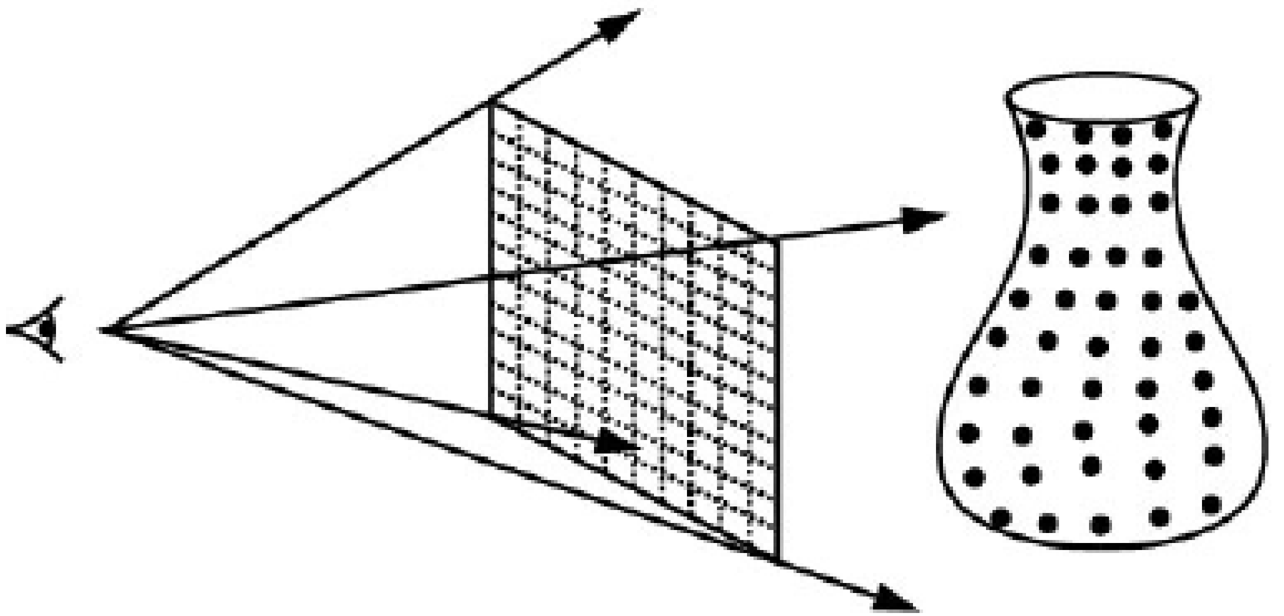

Range image

How to sense 3D very accurately?

- *range acquisition*
  - *contact*
    - *mechanical (CMM, jointed arm)*
  - *transmissive*
    - *industrial CT*
    - *MRI*
  - *reflective*
    - *non-optical*
      - *radar*
      - *ultrasound*
      - *sonar*
    - *optical*

- optical methods
  - passive
    - shape from X:
      - stereo
      - motion
      - shading
      - texture
      - focus
      - defocus
  - active
    - active variants of passive methods
      - Stereo w. projected texture
      - Active depth from defocus
      - Photometric stereo
    - time of flight
    - triangulation

Triangulation

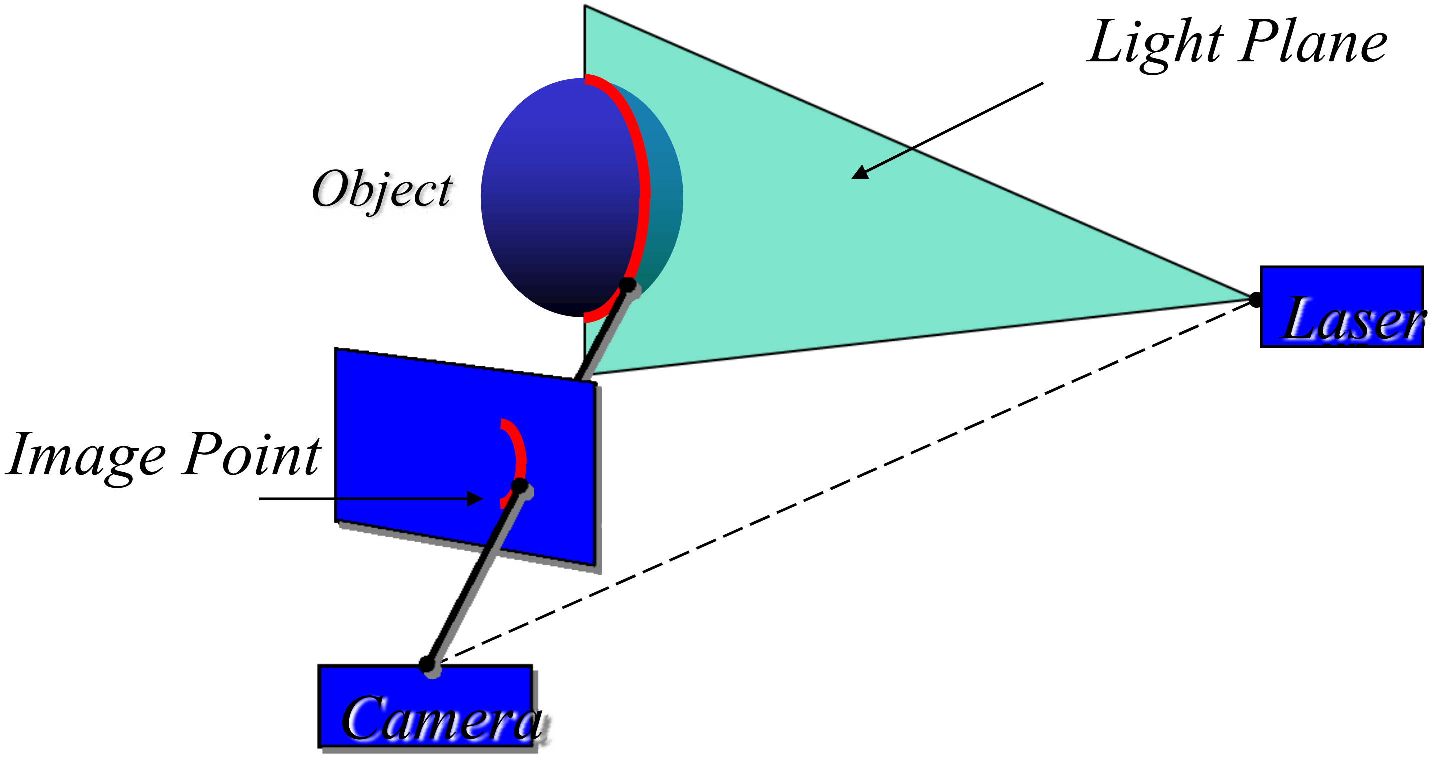

- Depth from ray-plane triangulation:
  - Intersect camera ray with light plane

Example: Laser scanner

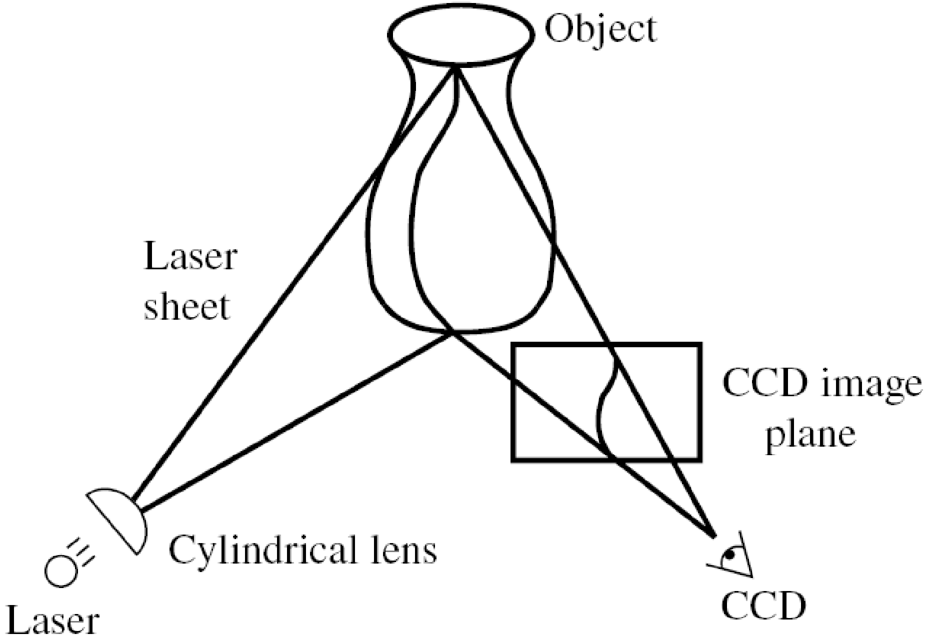

Cyberware® face and head scanner

+ *very accurate < 0.01 mm*
− *more than 10sec per scan*

Example: Laser scanner

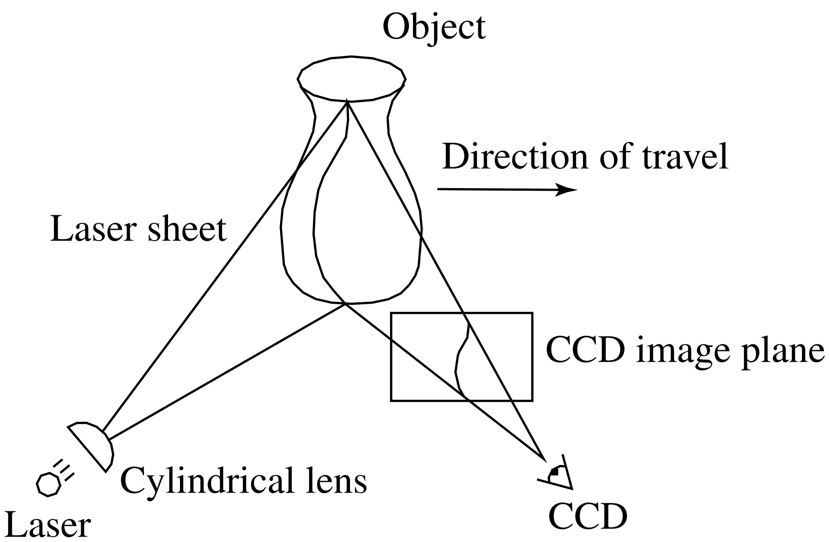

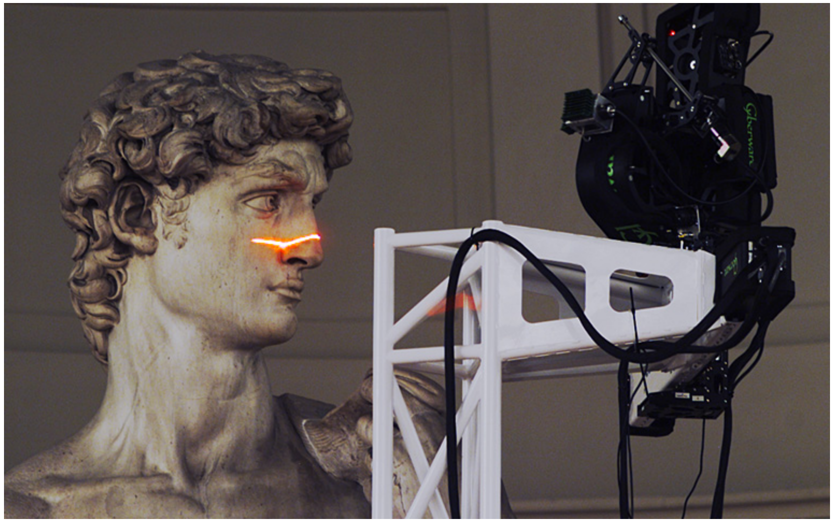

Digital Michelangelo Project

[*http://graphics.stanford.edu/projects/mich/*](http://graphics.stanford.edu/projects/mich/)

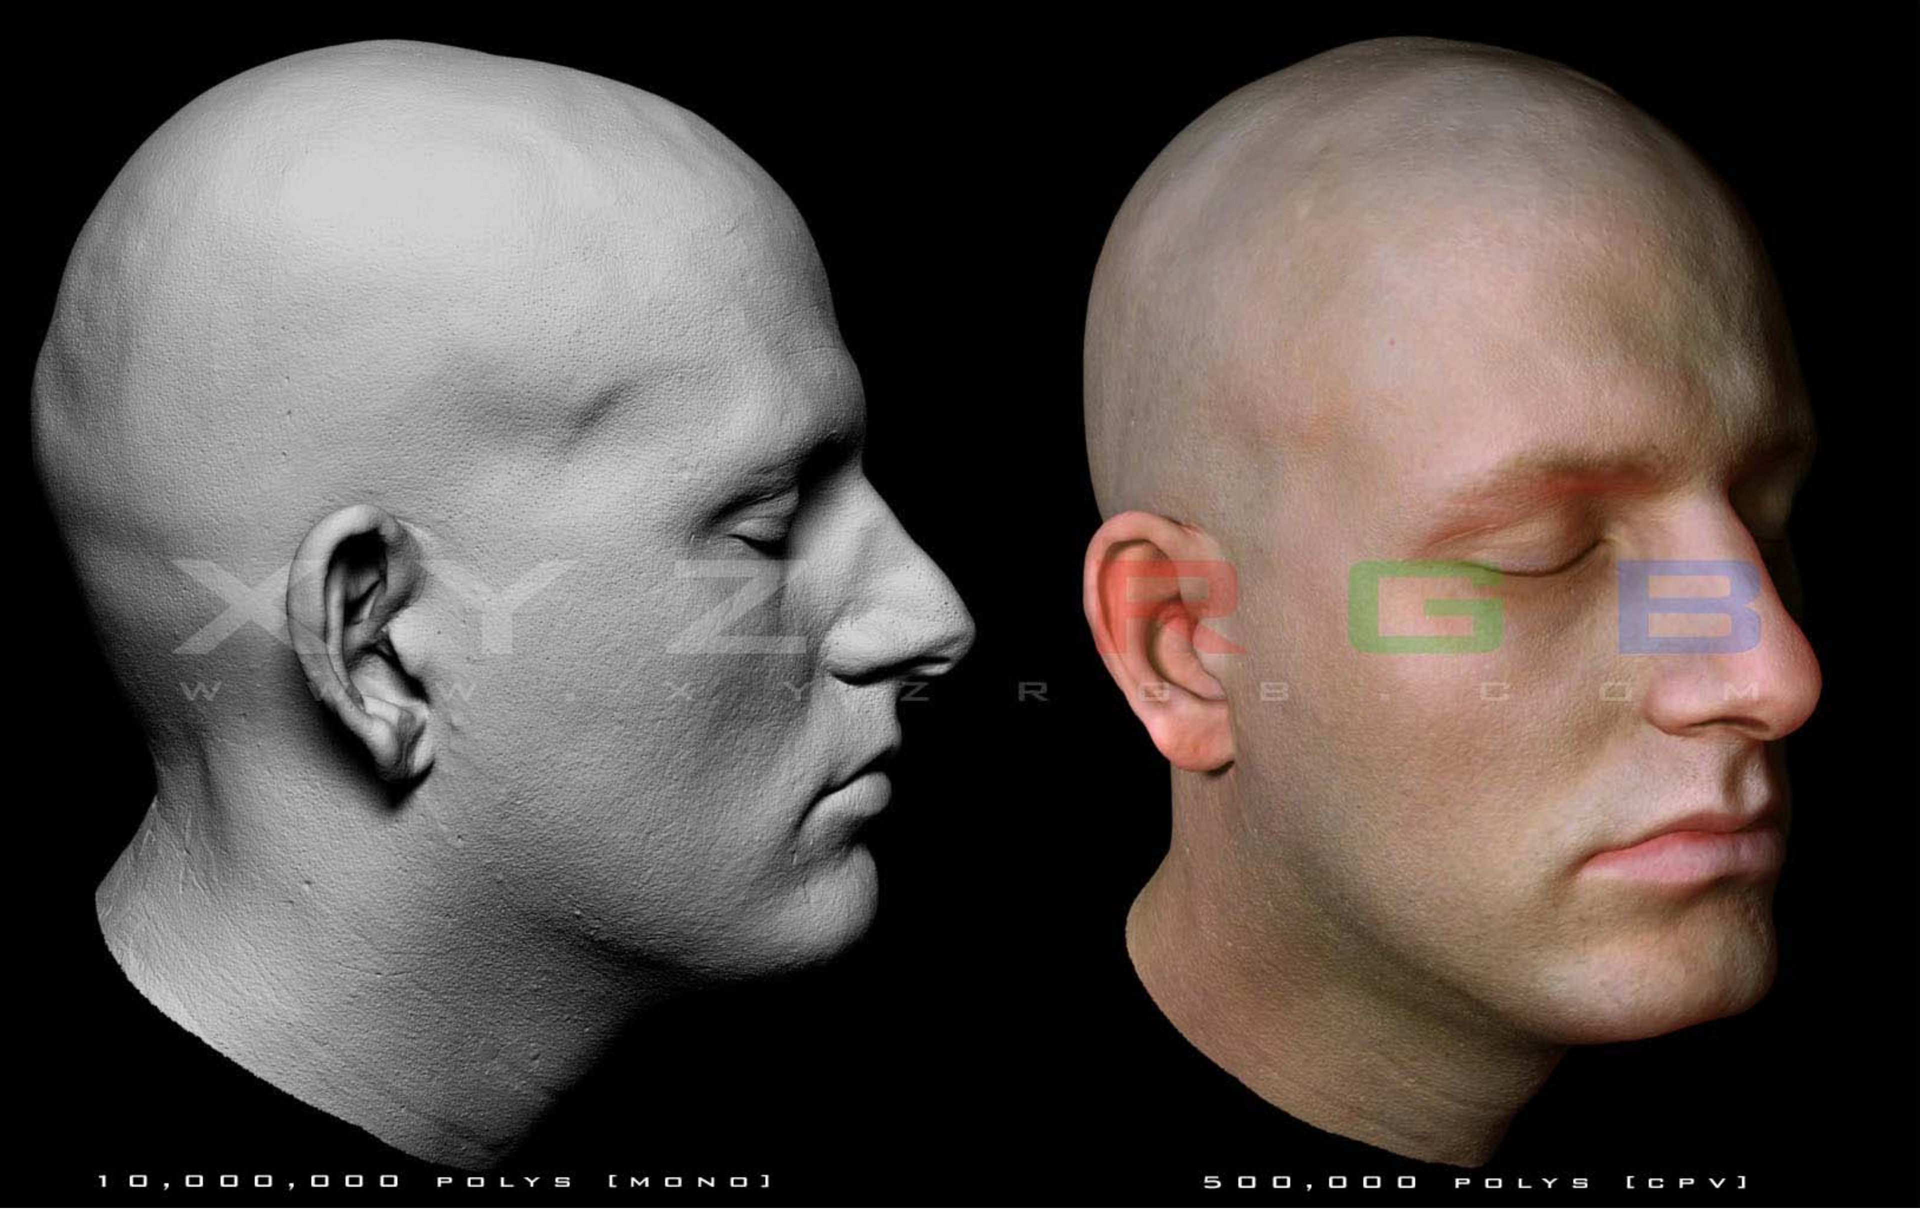

XYZRGB

Shadow scanning

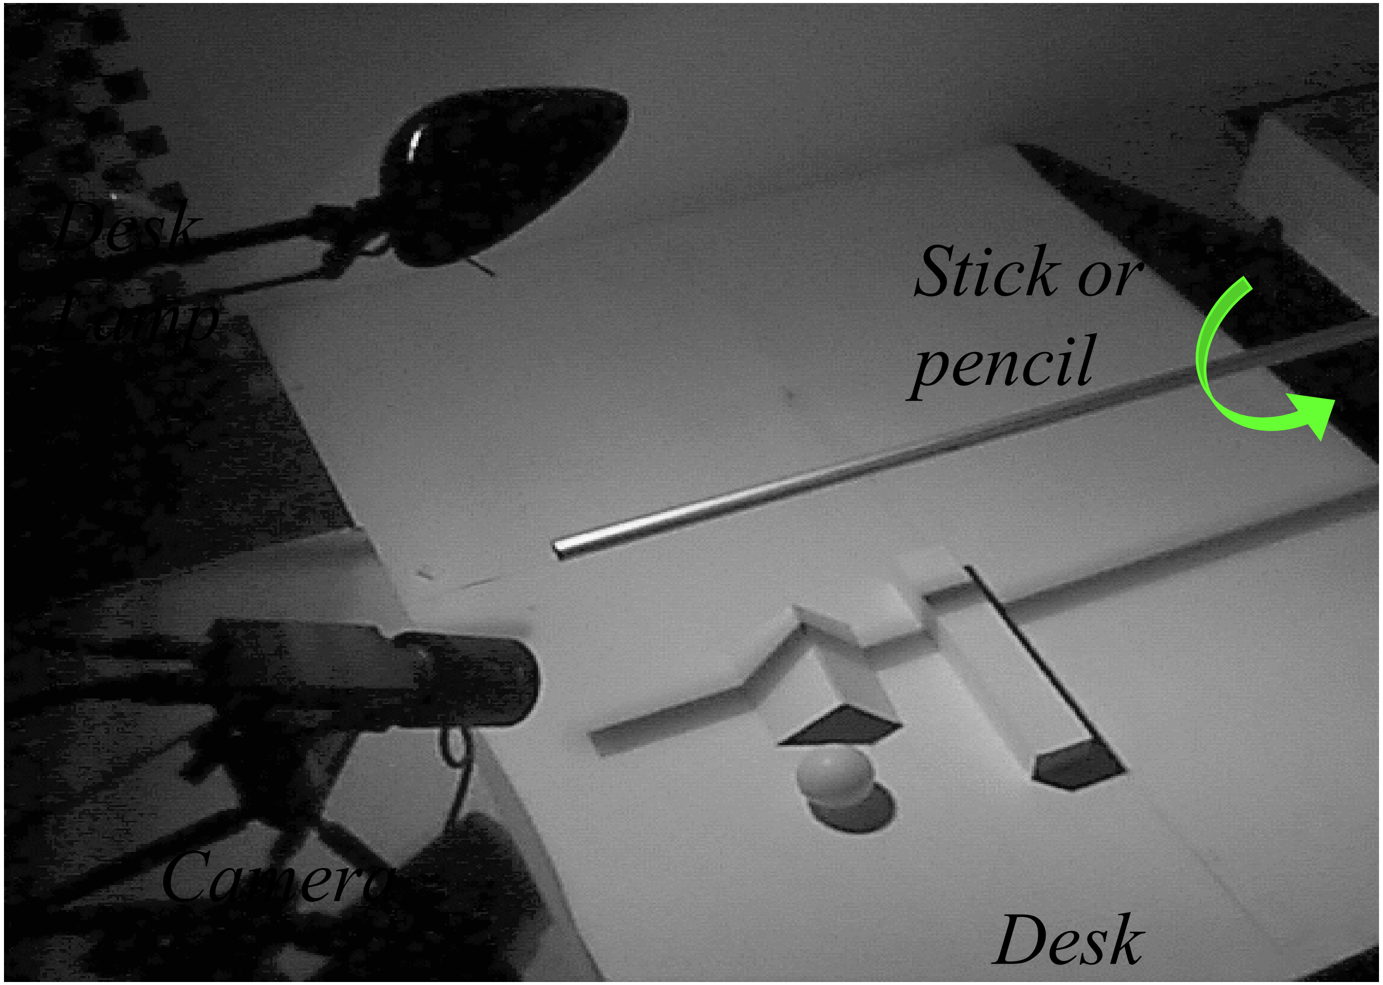

http://www.vision.caltech.edu/bouguetj/ICCV98/

Basic idea

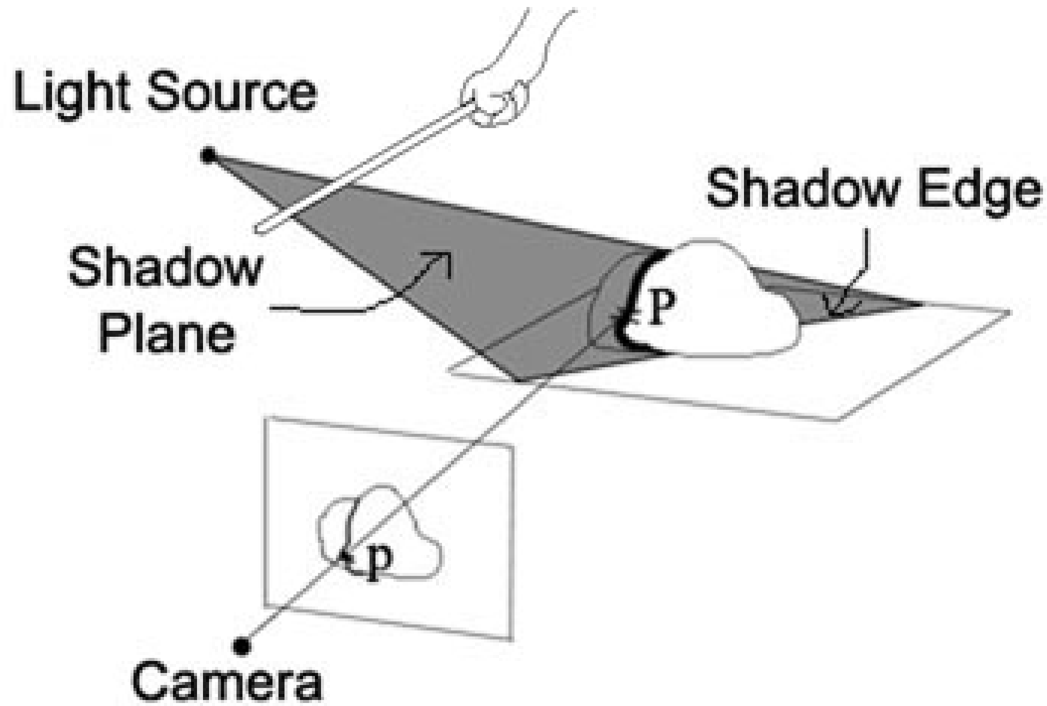

- Calibration issues:
  - where's the camera wrt. ground plane?
  - where's the shadow plane?
    - depends on light source position, shadow edge

Two Plane Version

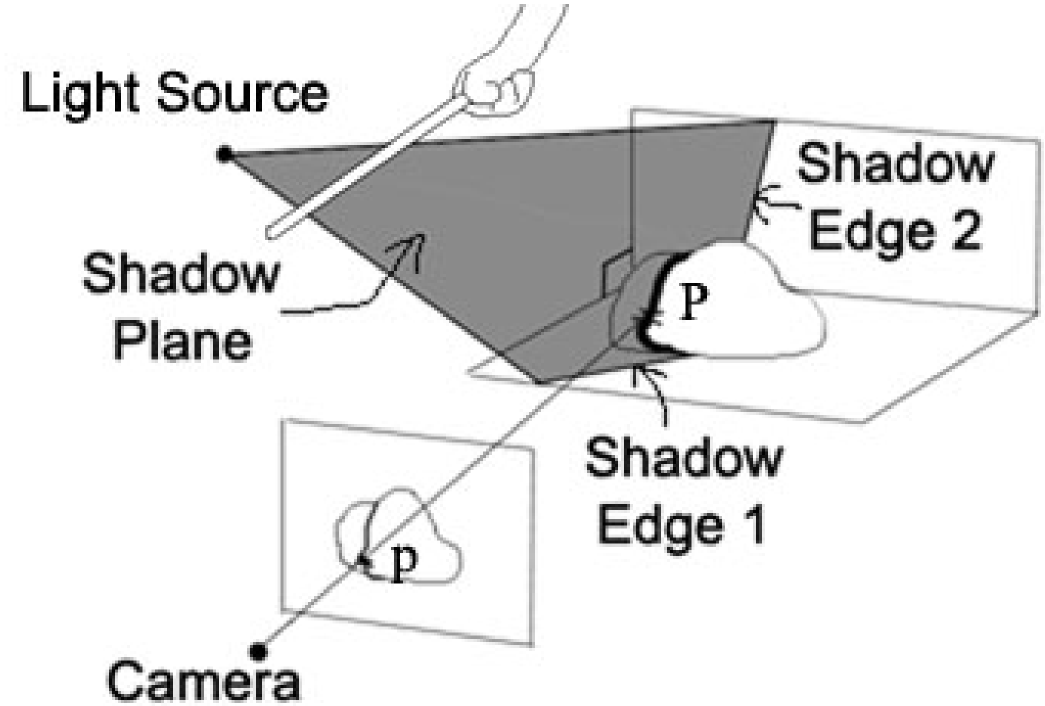

- Advantages
  - don’t need to pre-calibrate the light source
  - shadow plane determined from two shadow edges

Estimating shadow lines

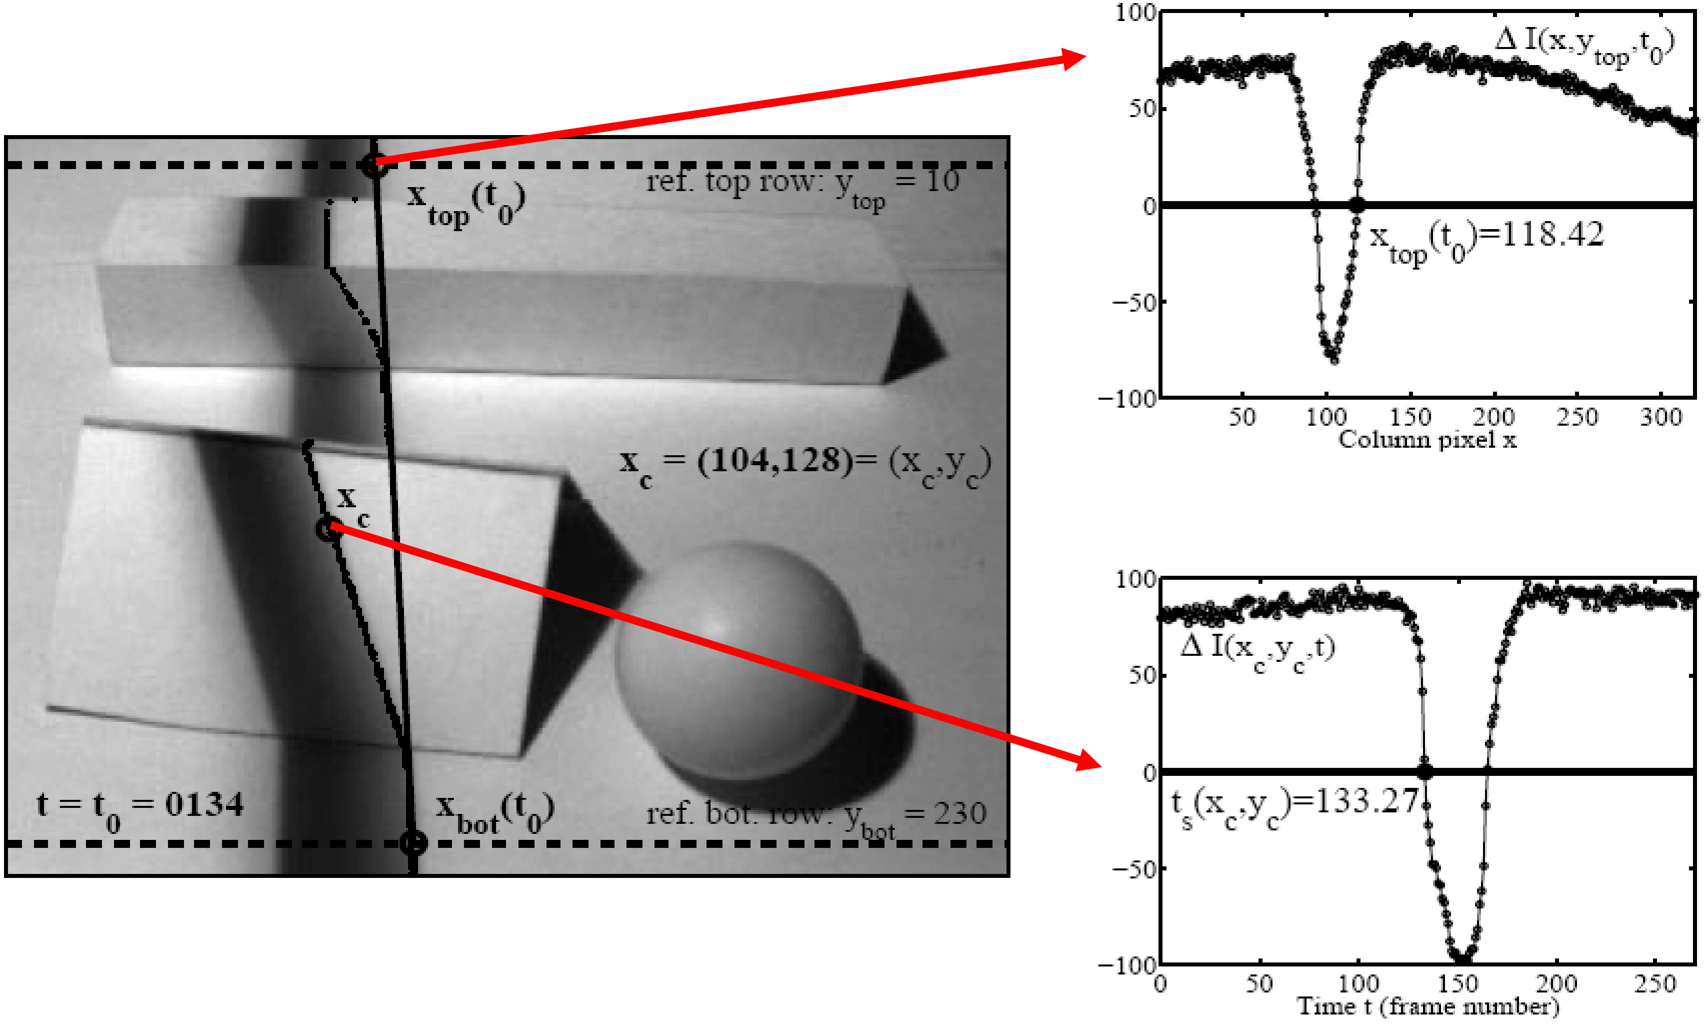

Shadow scanning in action

Results

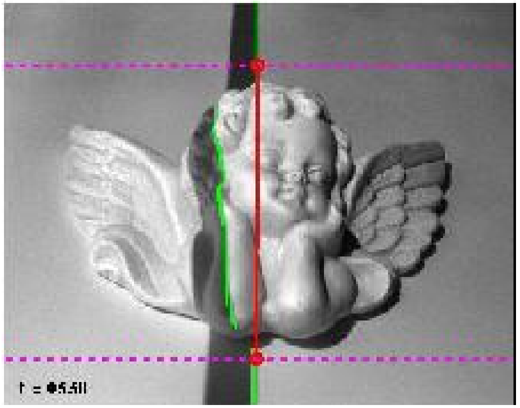

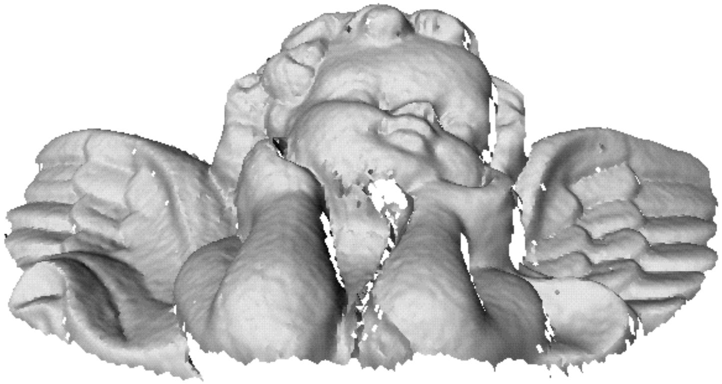

accuracy: 0.1mm over 10cm → *~ 0.1% error*

Textured objects

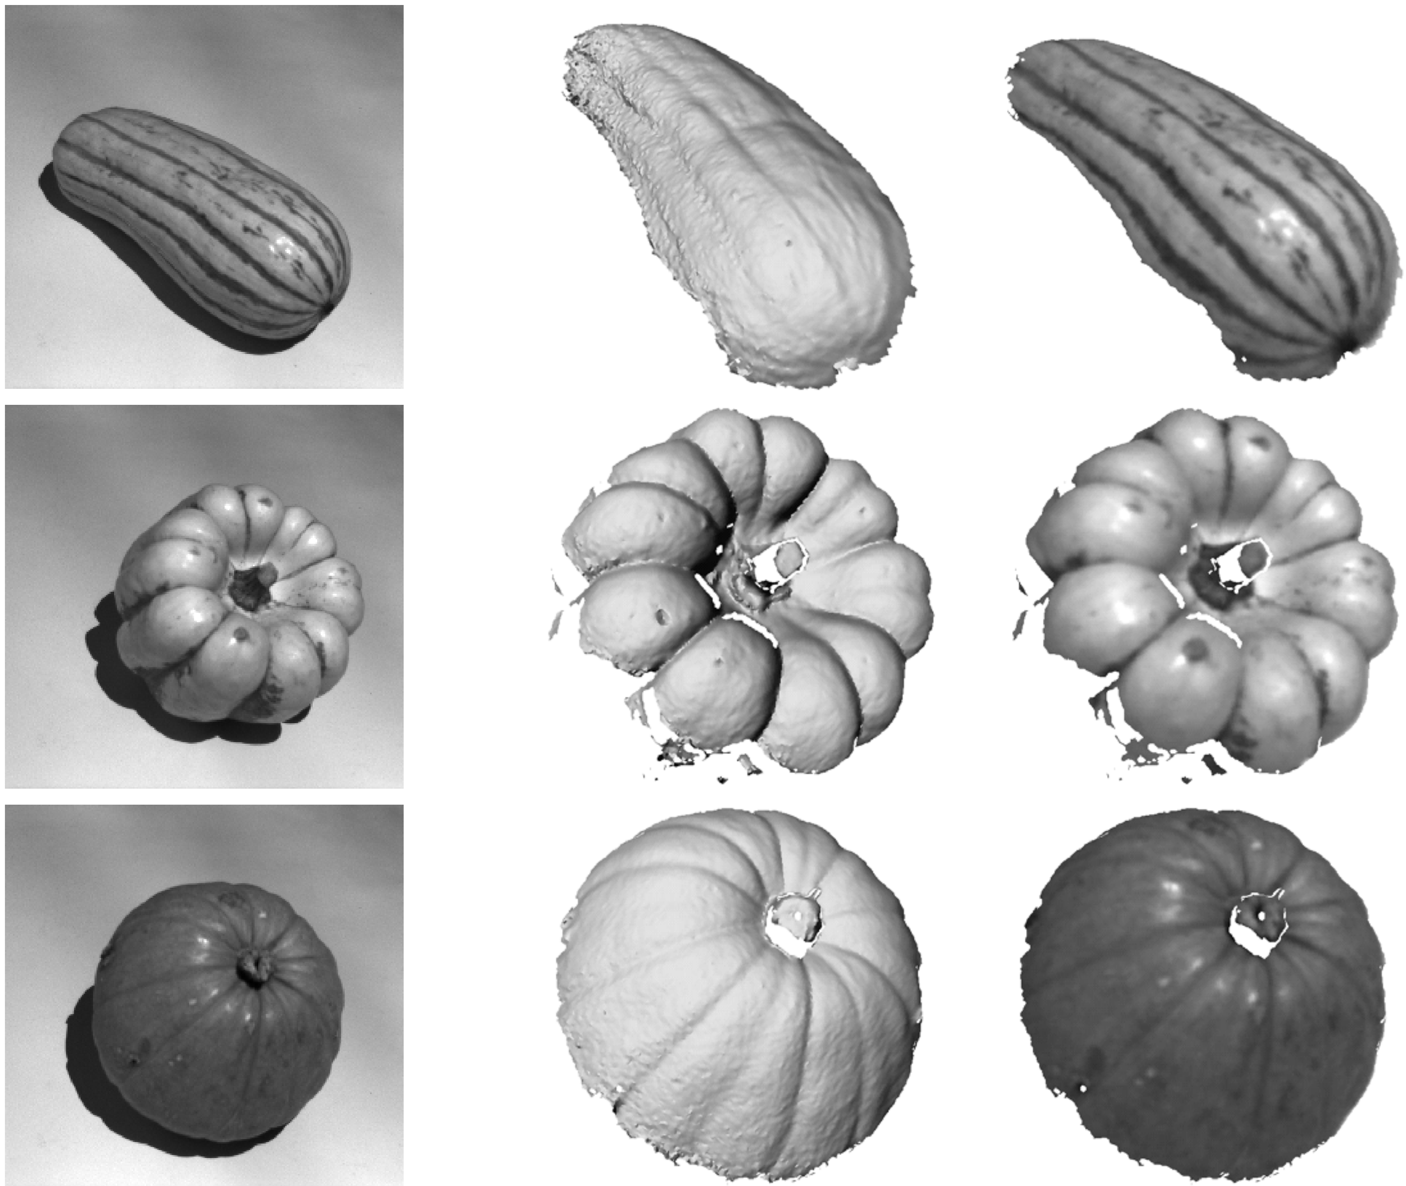

Scanning with the sun

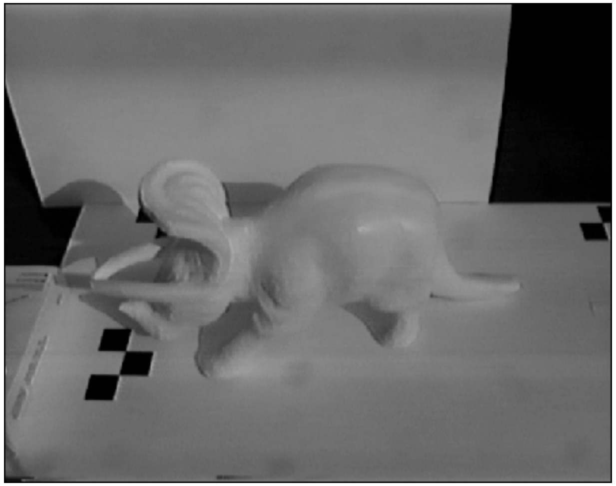

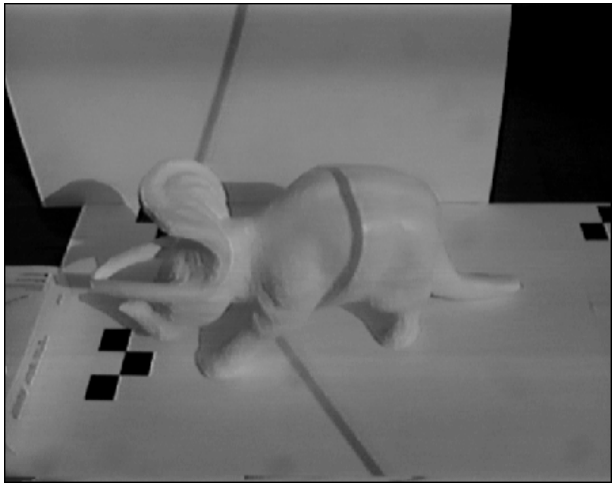

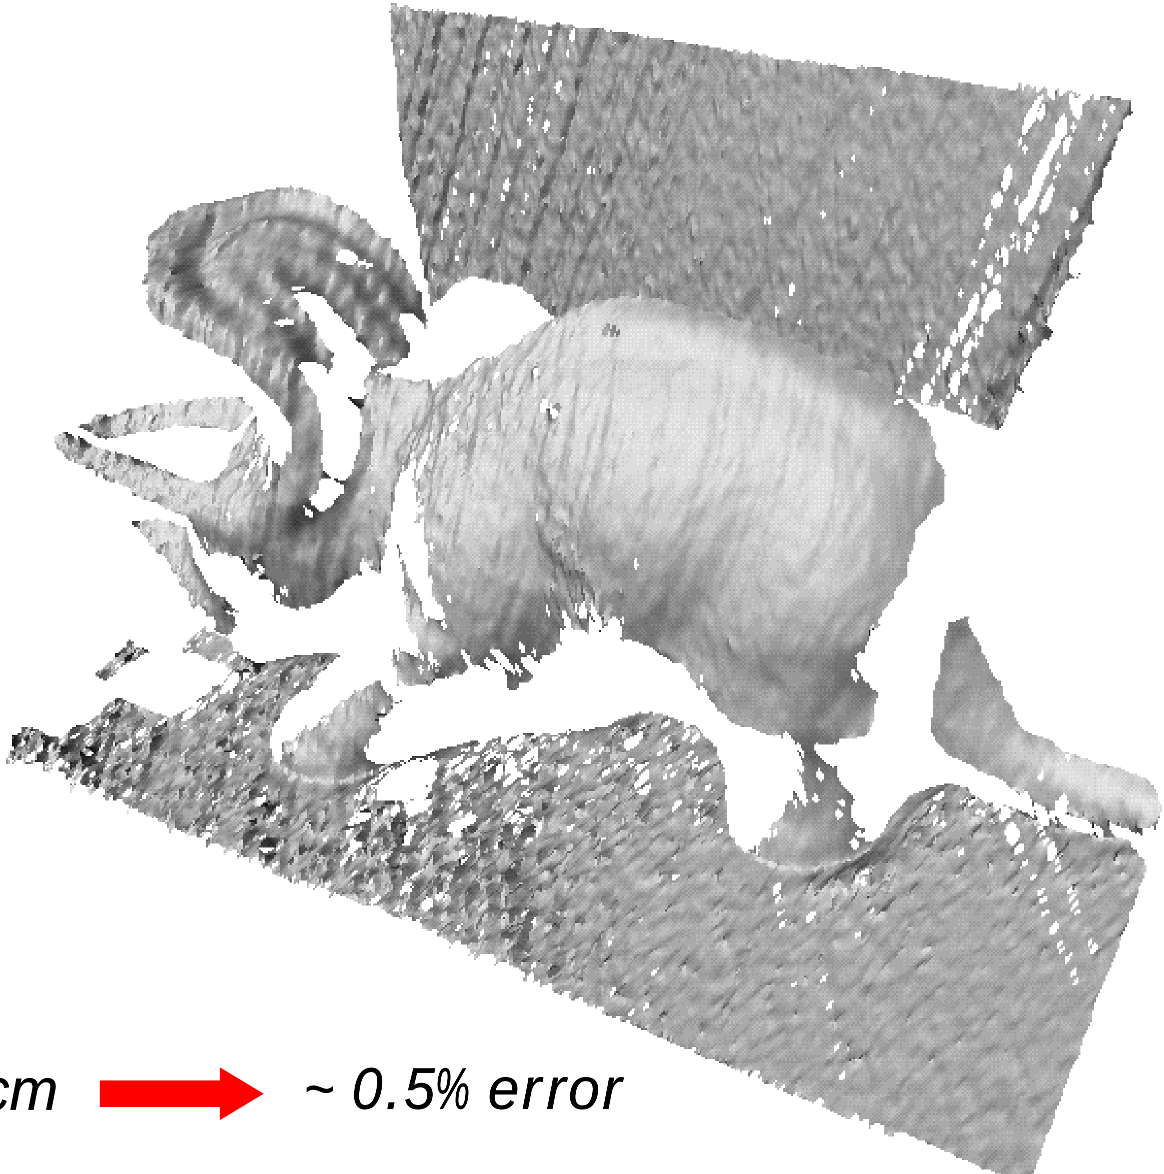

accuracy: 1mm over 50cm → *~ 0.5% error*

Scanning with the sun

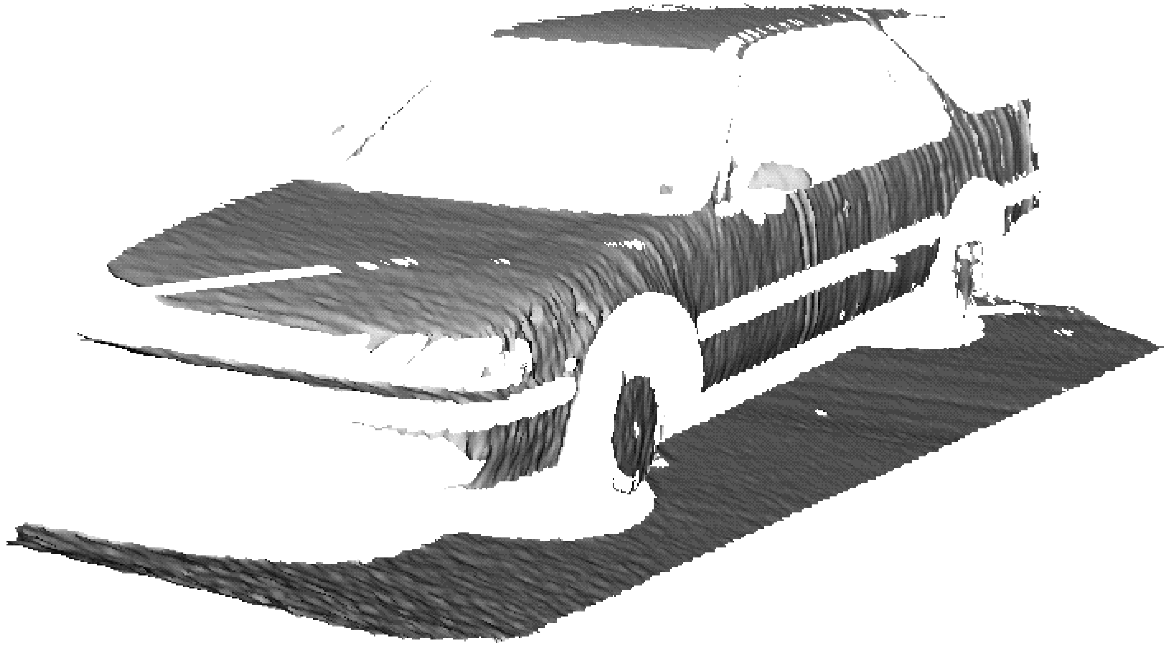

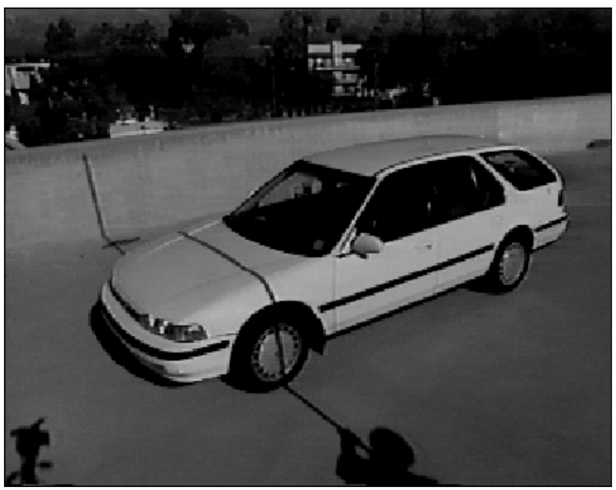

accuracy: 1cm over 2m

→ *~ 0.5% error*

Faster Acquisition?

- Project multiple stripes simultaneously
- Correspondence problem: which stripe is which?

- Common types of patterns:
  - *Binary coded light striping*
  - *Gray/color coded light striping*

Binary Coding

Faster:

$2^n - 1$ *stripes in* n *images.*

Example:

3 binary-encoded patterns which allows the measuring surface to be divided in 8 sub-regions

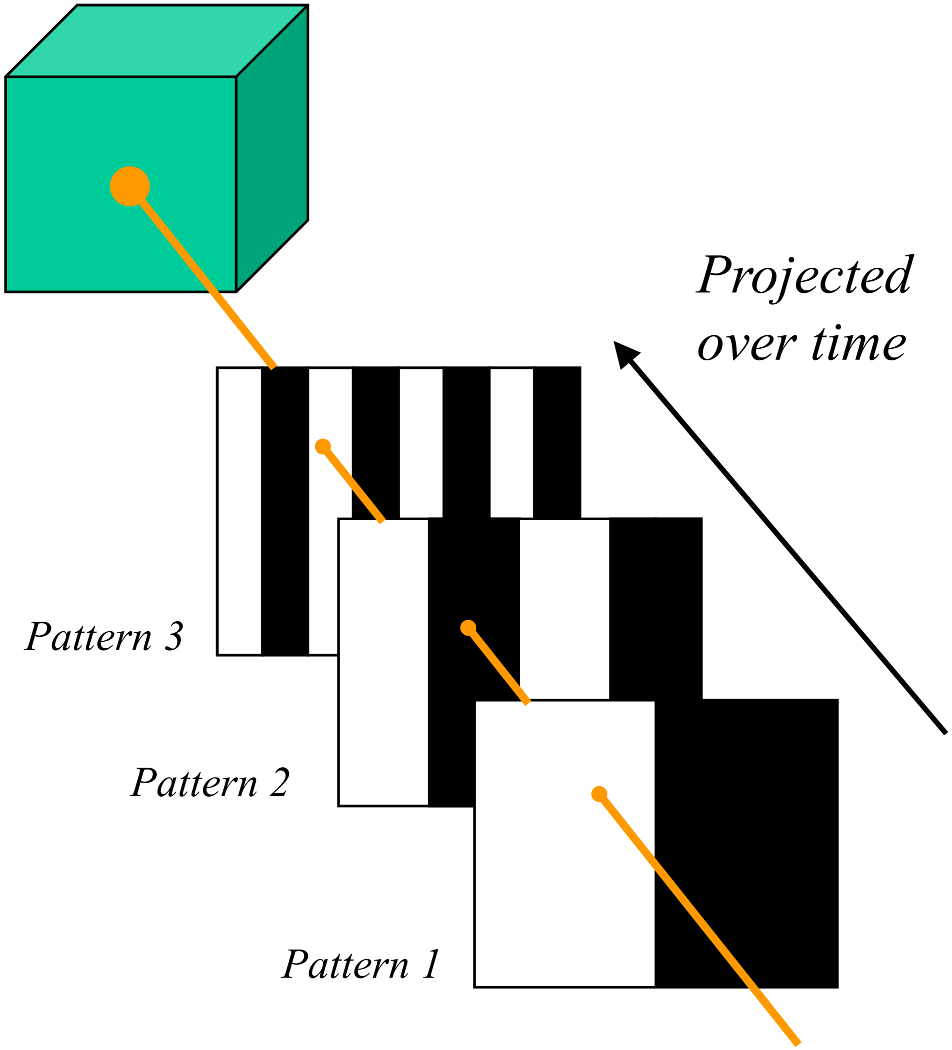

Binary Coding

- Assign each stripe a unique illumination code over time [Posdamer 82]

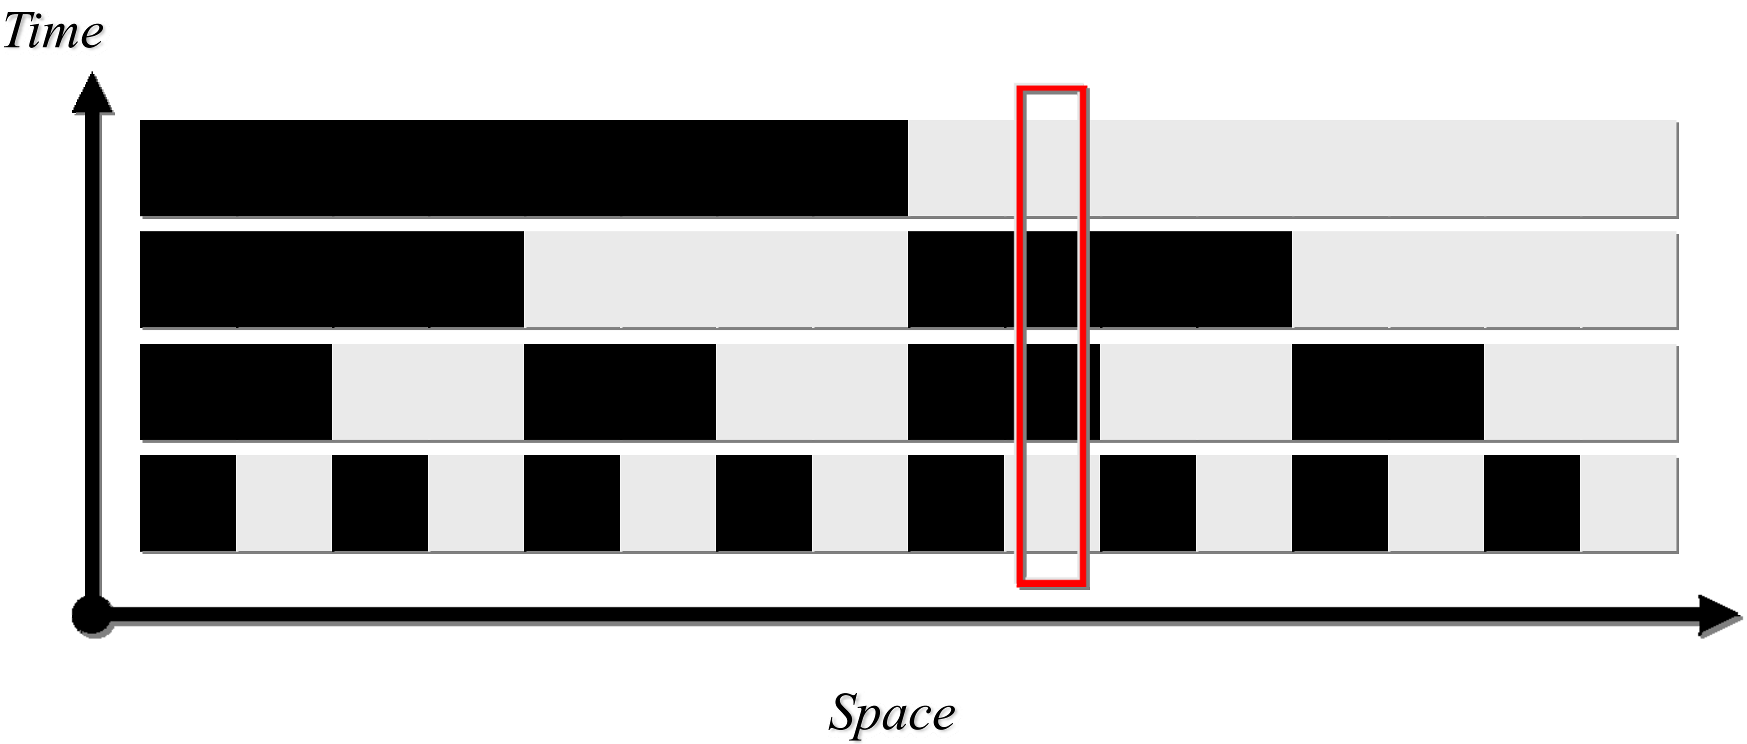

Binary Coding

Example: 7 binary patterns proposed by Posdamer & Altschuler

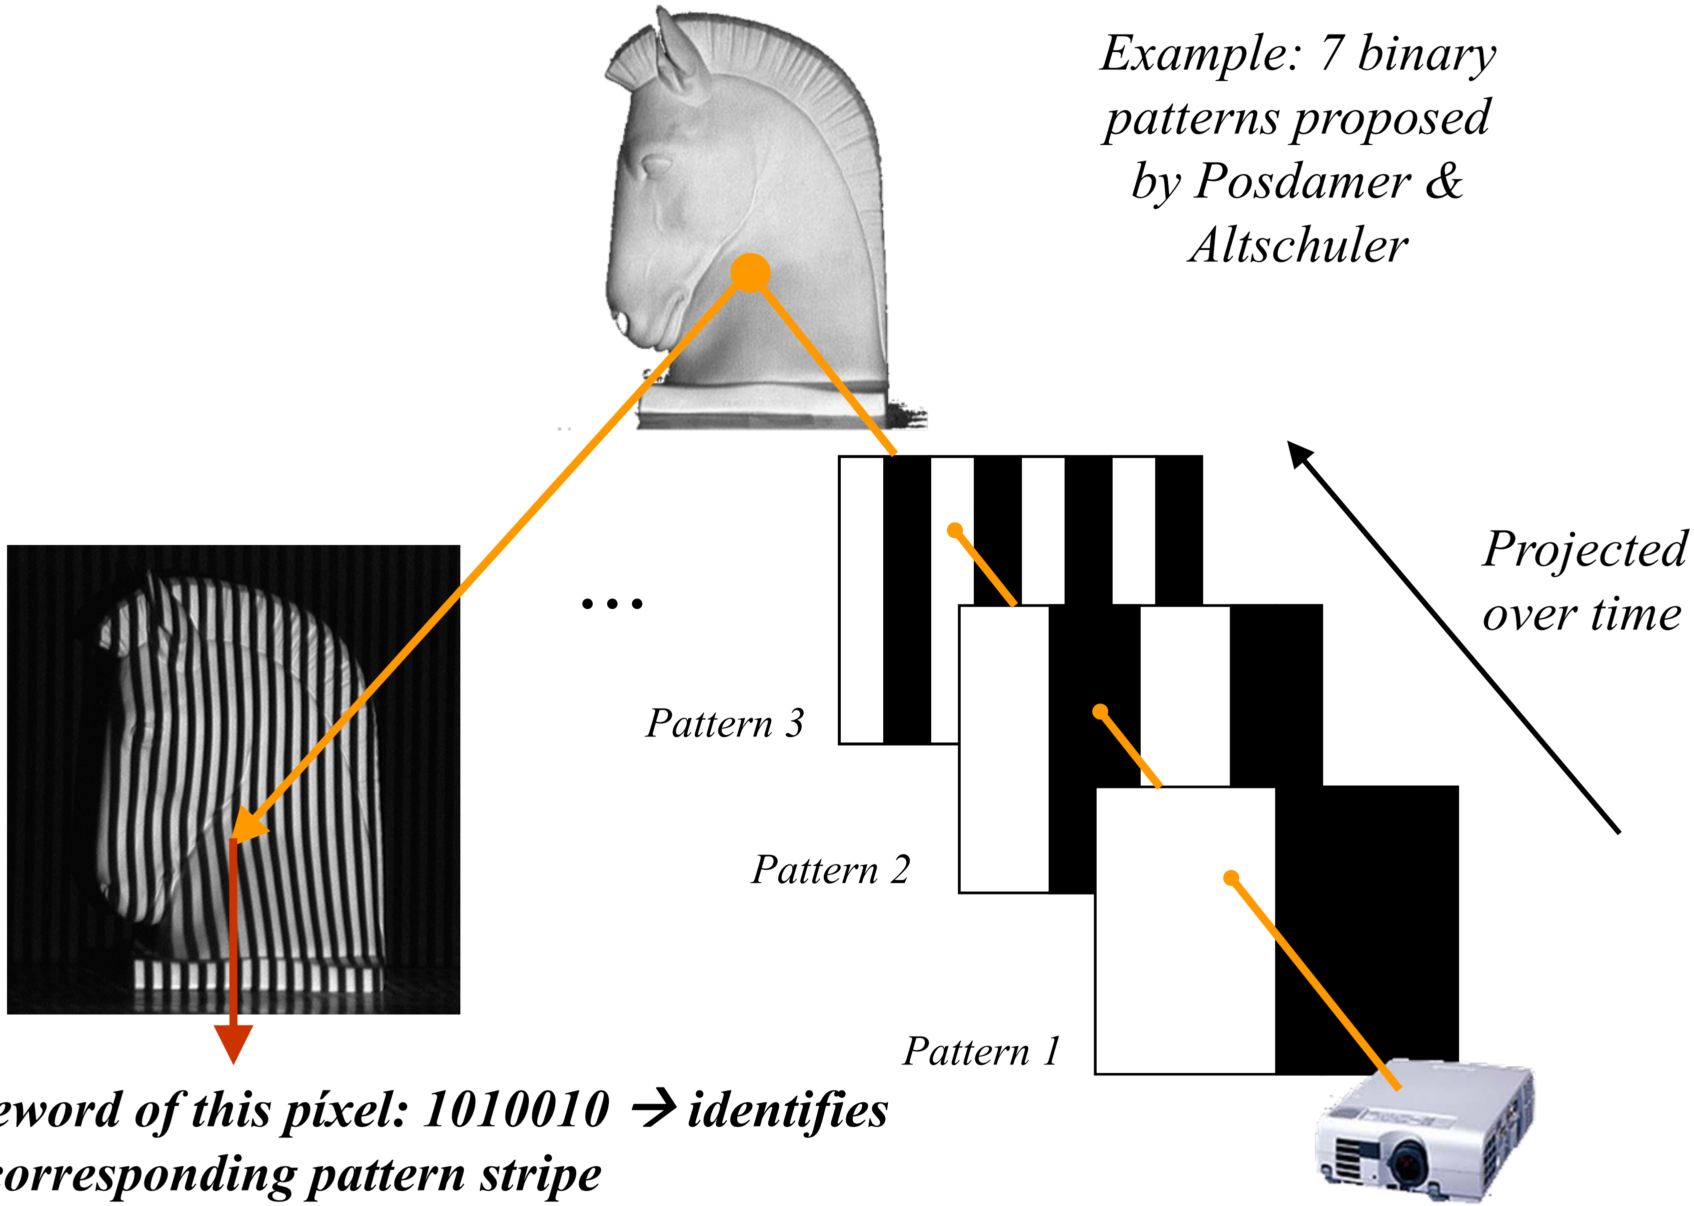

Codeword of this pixel: 1010010 → identifies the corresponding pattern stripe

More complex patterns

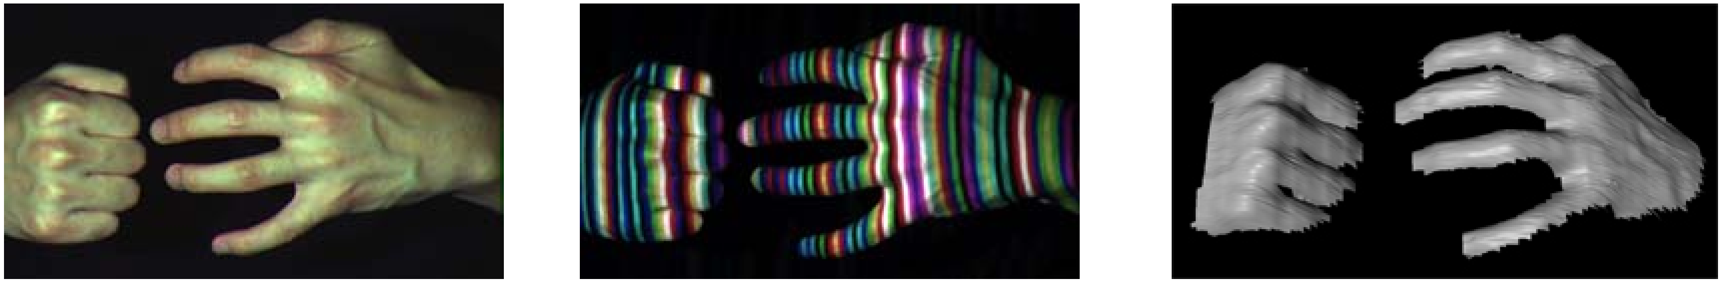

Works despite complex appearances

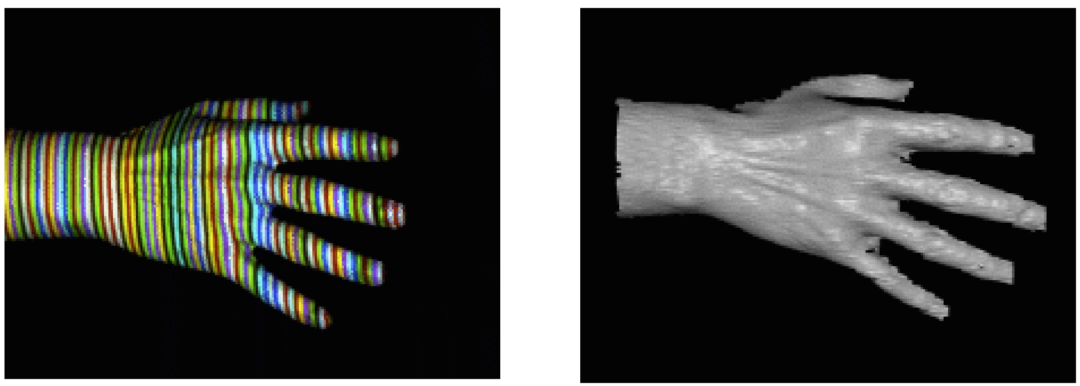

Works in real-time and on dynamic scenes

- *Need very few images (one or two).*
- *But needs a more complex correspondence algorithm*

Zhang et a

Continuum of Triangulation Methods

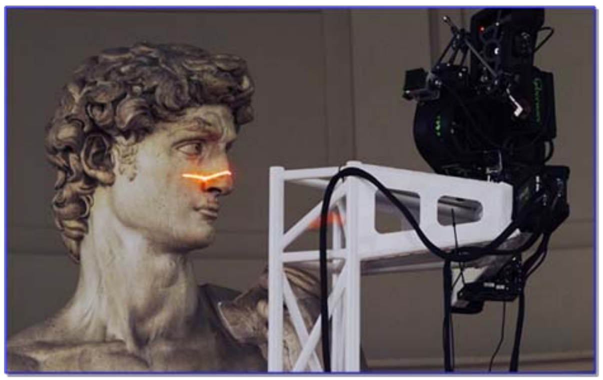

Multi-stripe
Multi-frame

Single-stripe

Single-frame

Slow, robust

Fast, fragile

Time-of-flight

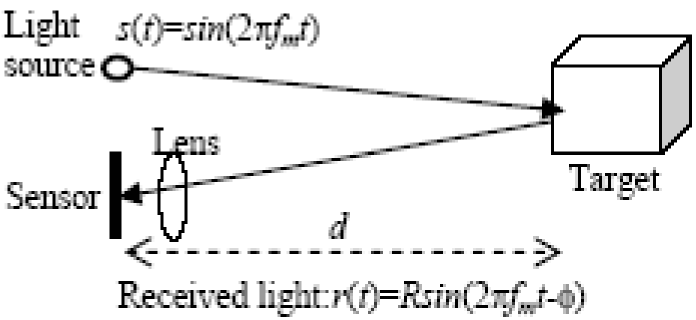

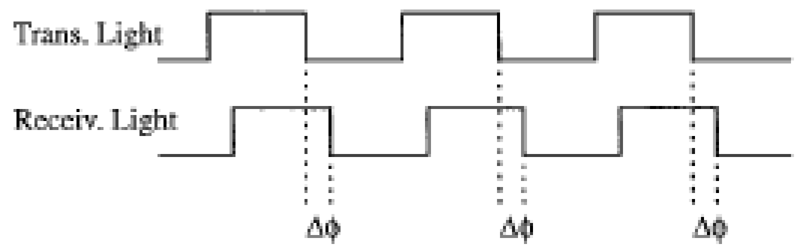

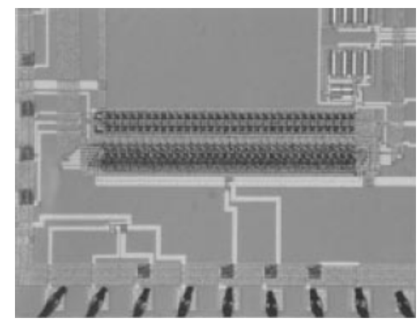

+ *No baseline, no parallax shadows*
+ *Mechanical alignment is not as critical*
− *Low depth accuracy*
− *Single viewpoint capture*

Miyagawa, R., Kanade, T., "CCD-Based Range Finding Sensor", IEEE Transactions on Electron Devices, 1997

Working Volume: 1500mm - Accuracy: 7%

Spatial Resolution: 1x32- Speed: ??

Comercial products

Canesta

64x64@30hz
Accuracy 1-2cm

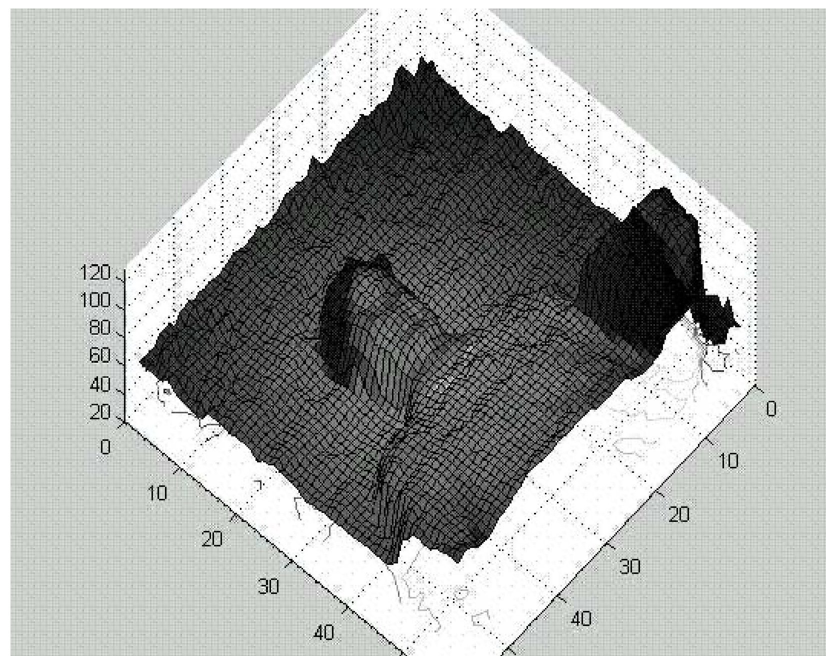

Not accurate enough for face modeling, but good enough for layer extraction.

Depth from Defocus

Depth from Defocus

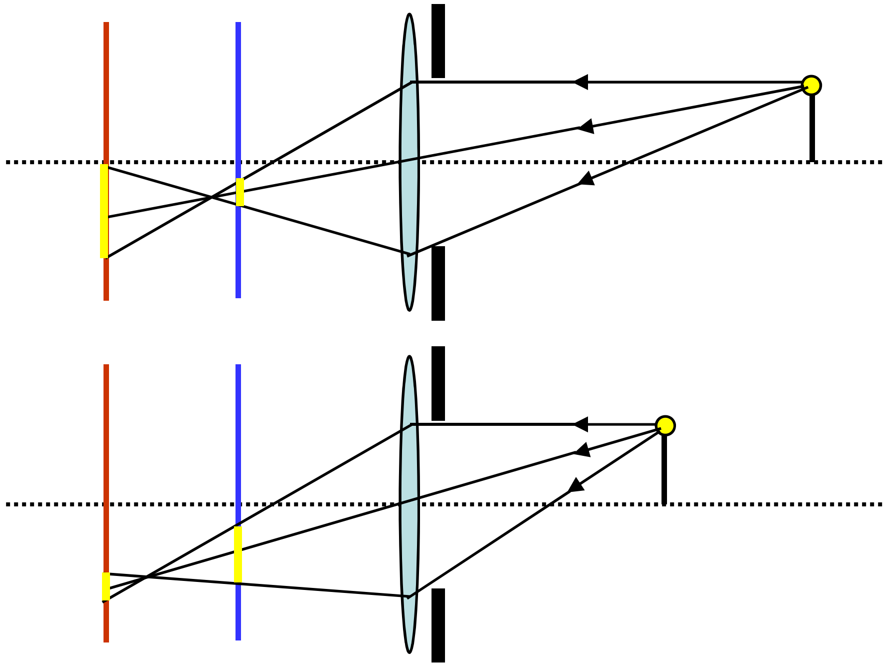

Depth from Defocus

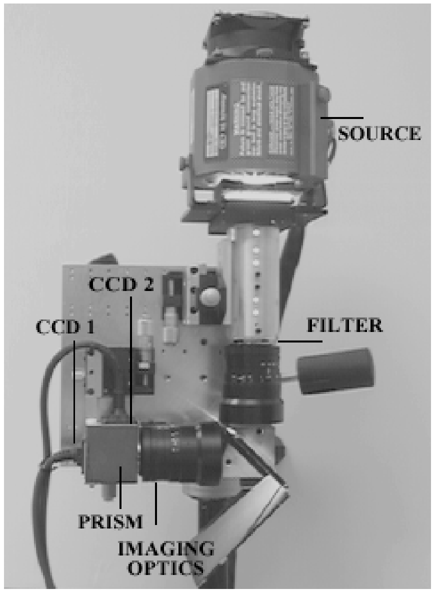

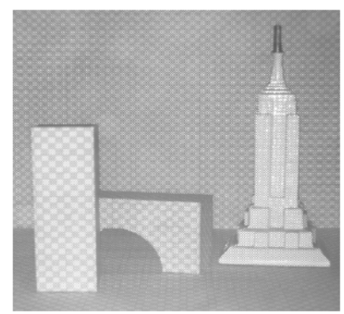

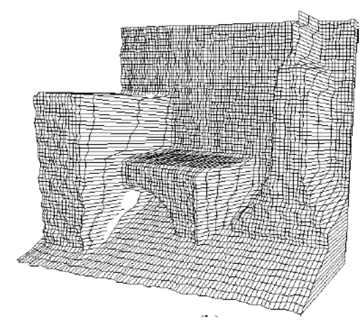

+ *Hi resolution and accuracy, real-time*
− *Customized hardware*
− *Single view capture?*

Nayar, S.K., Watanabe, M., Noguchi, M., "Real-Time Focus Range Sensor", ICCV 1995

Working Volume: 300mm - Accuracy: 0.2%

Spatial Resolution: 512x480 - Speed: 30Hz