3D Scanning

COS 429: Computer Vision

PRINCETON UNIVERSITY

3D Scanning Applications

- Computer graphics
- Product inspection
- Robot navigation
- As-built floorplans
- Product design
- Archaeology
- Clothes fitting
- Art history

Industrial Inspection

- Determine whether manufactured parts are within tolerances

Medicine

- Plan surgery on computer model, visualize in real time

Medicine

- Plan surgery on computer model, visualize in real time

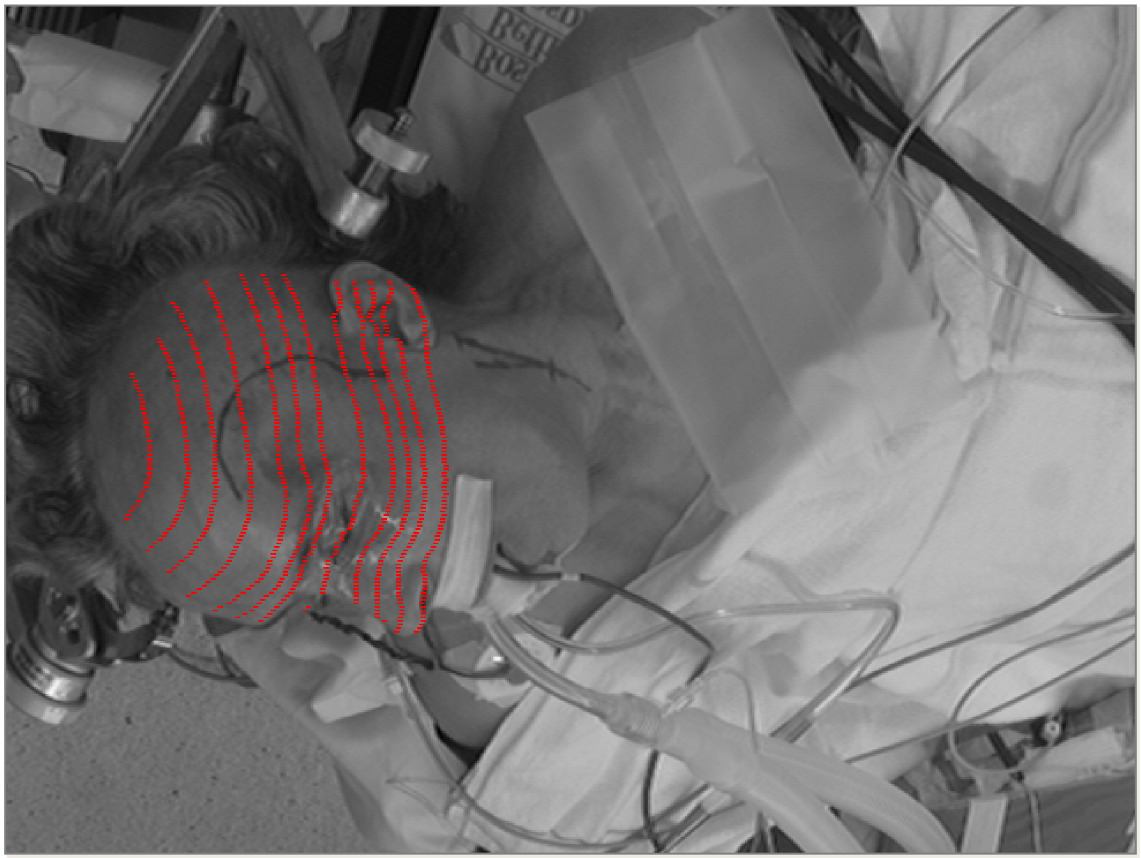

Medicine

- Plan surgery on computer model, visualize in real time

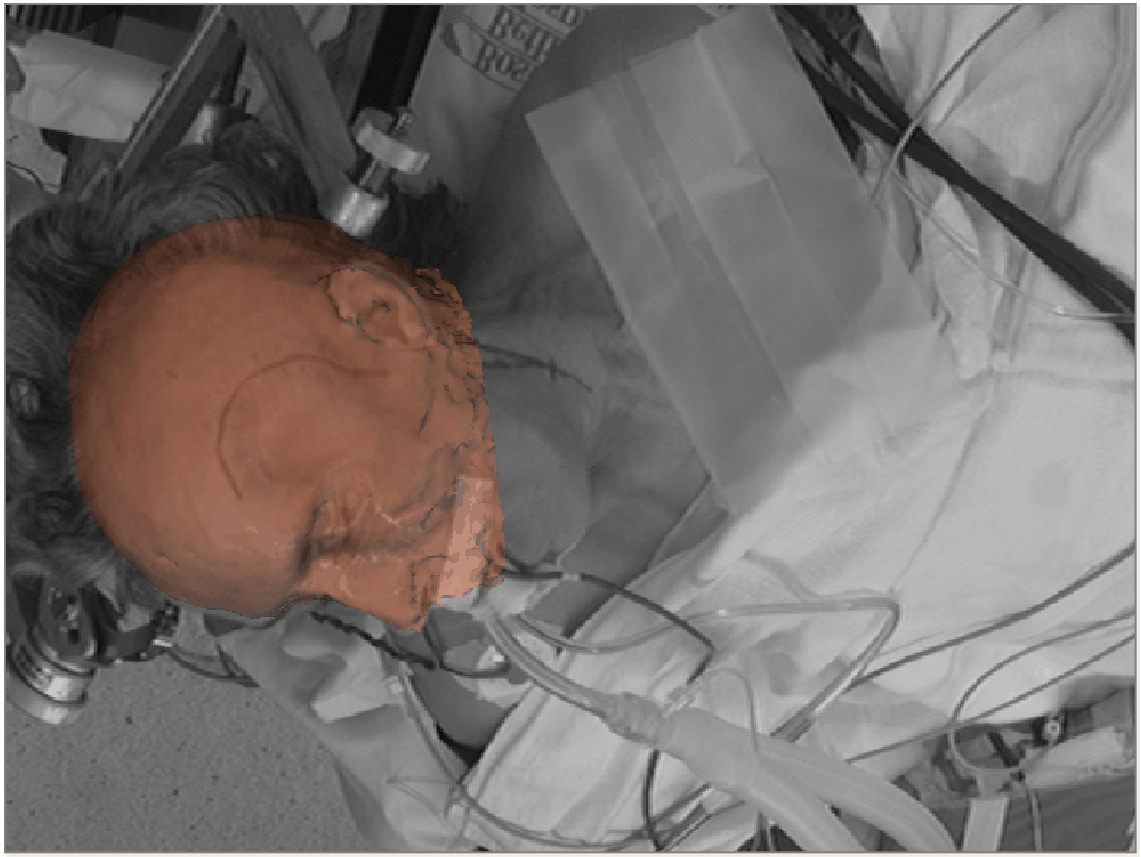

Medicine

- Plan surgery on computer model, visualize in real time

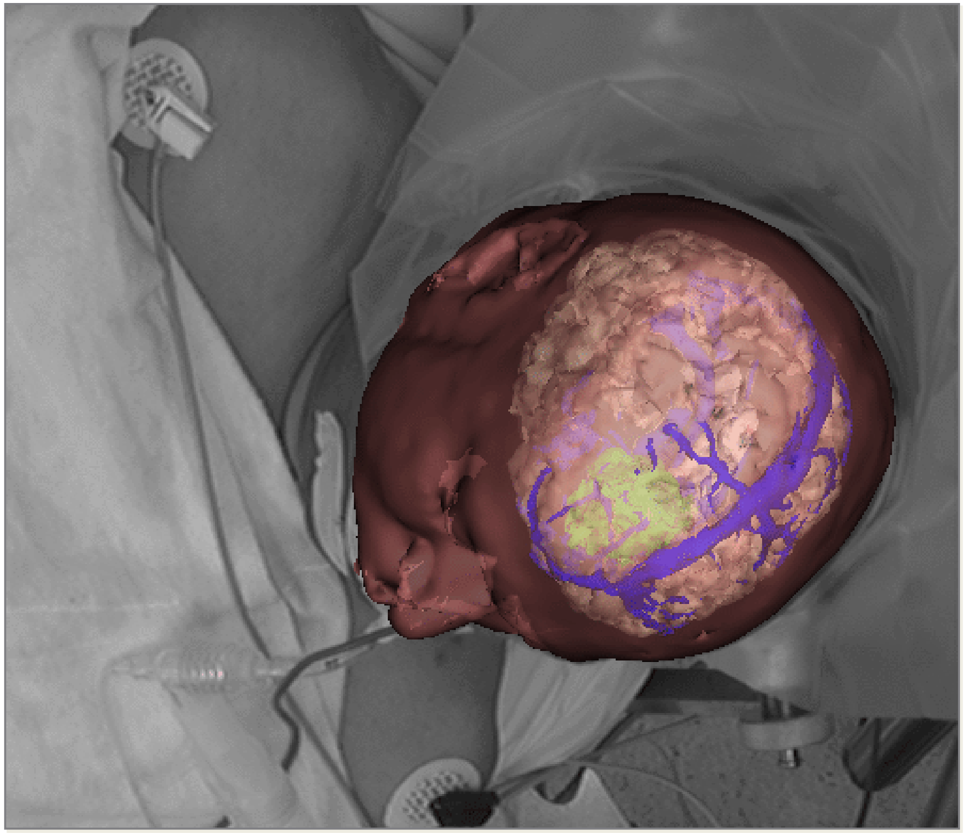

Scanning Buildings

- Quality control during construction
- As-built models

Scanning Buildings

- Quality control during construction
- As-built models

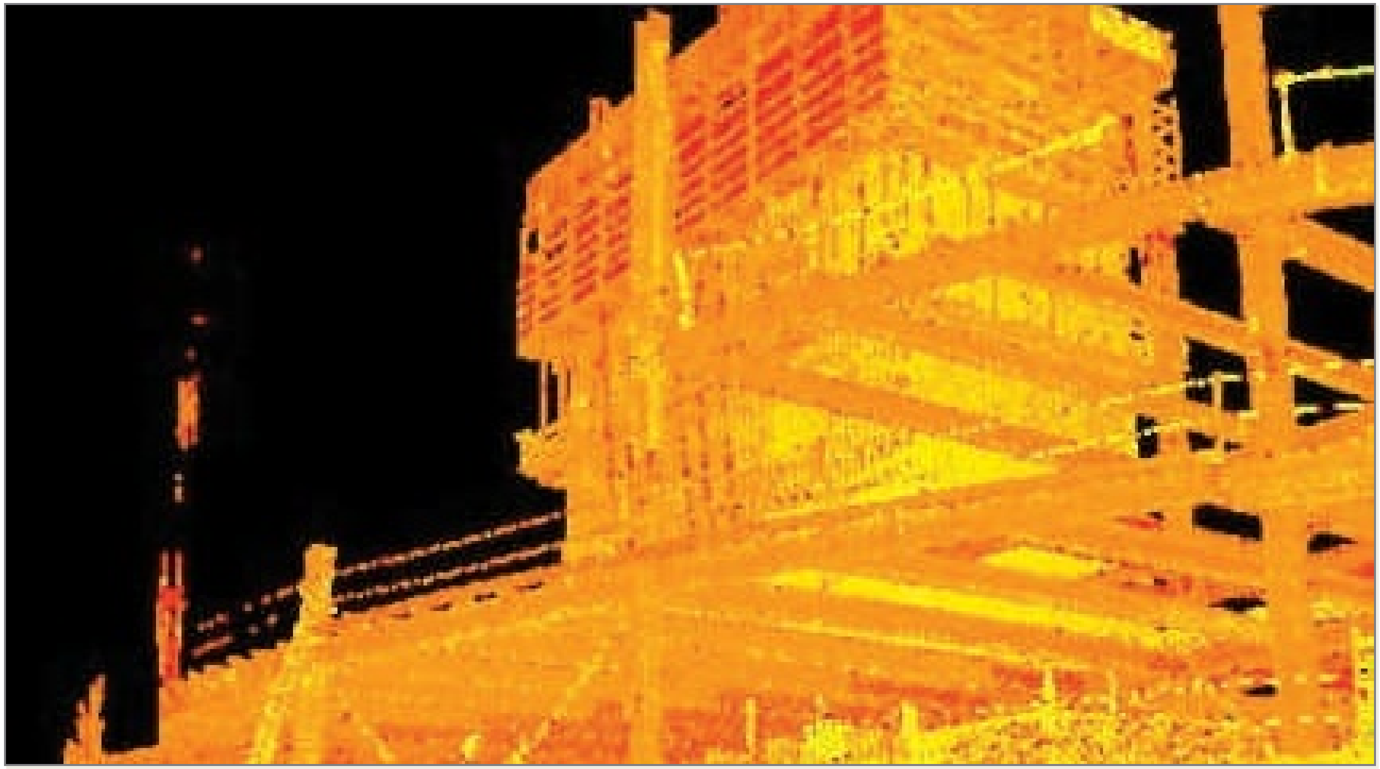

Clothing

- Scan a person, custom-fit clothing
- U.S. Army; booths in malls

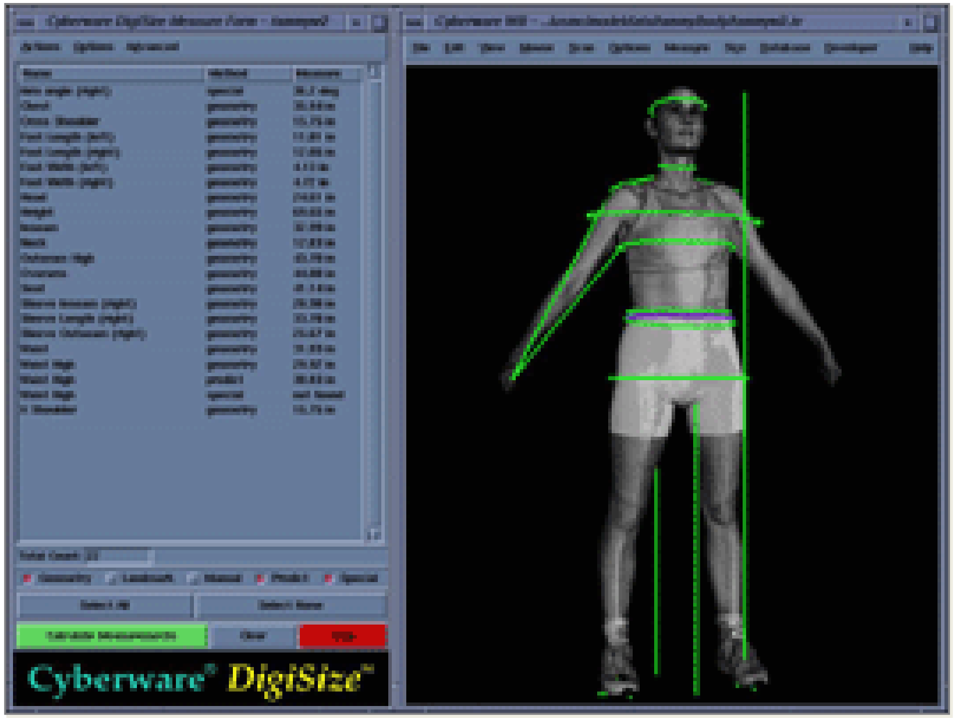

Range Acquisition Taxonomy

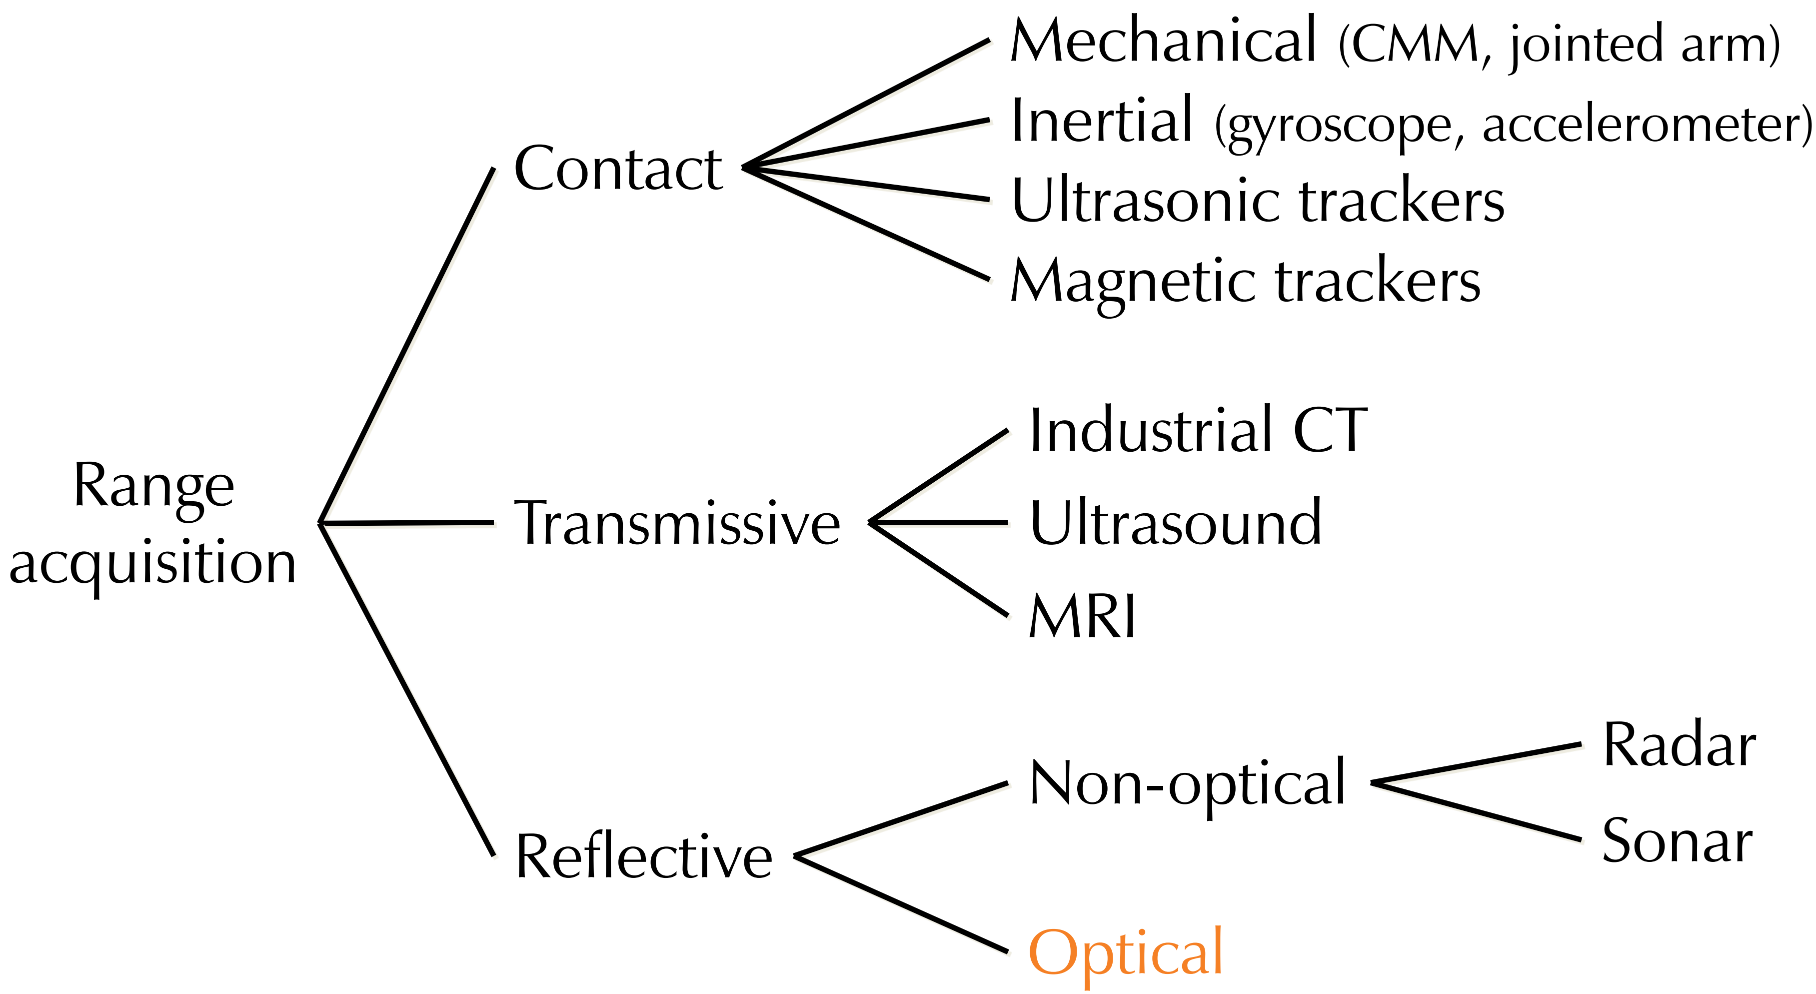

Range Acquisition Taxonomy

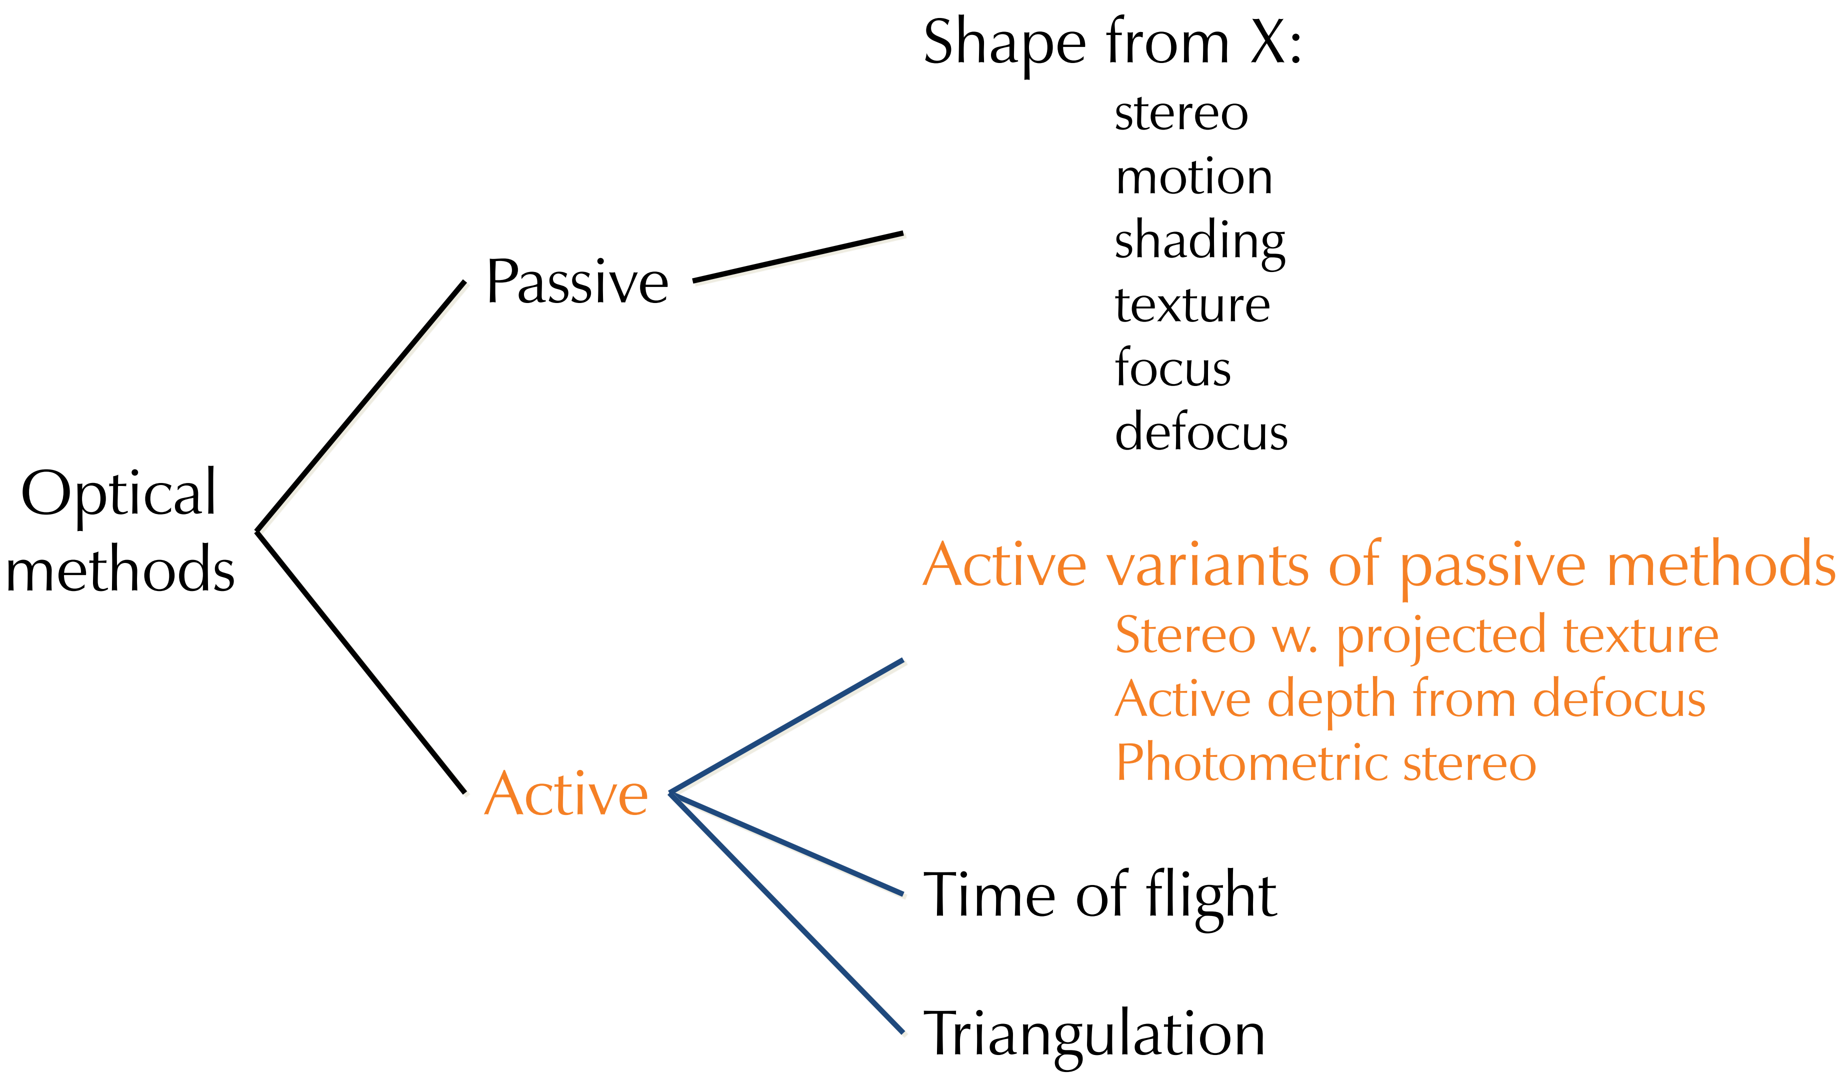

Active Optical Methods

- Advantages:
  - Usually can get dense data
  - Usually much more robust and accurate than passive techniques
- Disadvantages:
  - Introduces light into scene (distracting, etc.)
  - Not motivated by human vision

Active Variants of Passive Techniques

- Regular stereo with projected texture
  - Provides features for correspondence
- Active depth from defocus
  - Known pattern helps to estimate defocus
- Photometric stereo
  - Shape from shading with multiple known lights

Lambertian Reflectance Model

- Diffuse surfaces appear equally bright from all directions

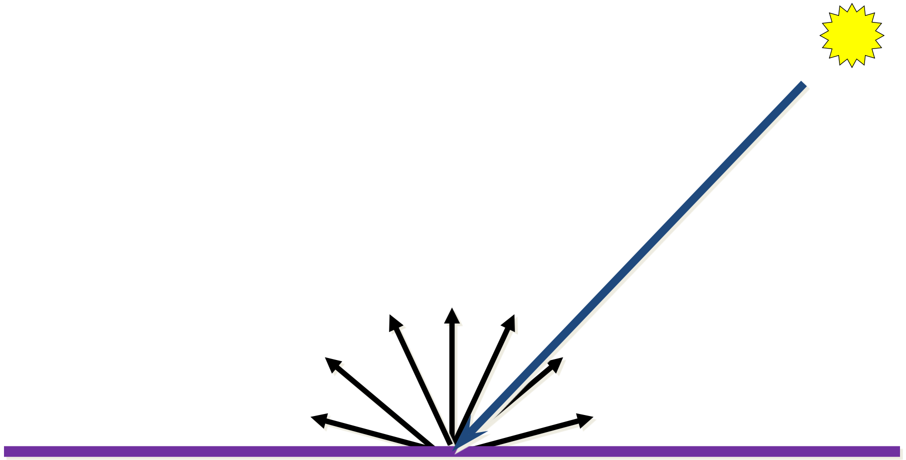

Lambertian Reflectance Model

- Diffuse surfaces appear equally bright from all directions

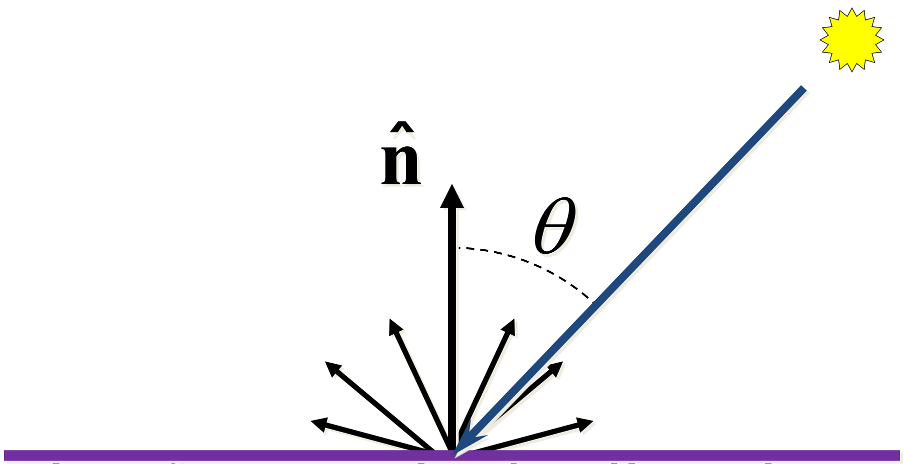

- For illumination coming from a single direction, brightness proportional to cos θ

Lambertian Reflectance Model

- Therefore, for a constant-colored object with distant illumination, we can write

$$E = L\,\rho\,l \cdot n$$

$E =$ observed brightness

$L =$ brightness of light source

$\rho =$ reflectance (albedo) of surface

$l =$ direction to light source

$n =$ surface normal

Shape from Shading

- The above equation contains some information about shape, and in some cases is enough to recover shape completely (in theory)
  if L, ρ, and l are known
- Similar to integration (surface normal is like a derivative), but only know a part of derivative
- Have to assume surface continuity

Active Shape from Shading

- Idea: several (user-controlled) light sources
- More data
  - Allows determining surface normal directly
  - Allows spatially-varying reflectance
  - Redundant measurements: discard shadows and specular highlights
- Often called "photometric stereo"

Photometric Stereo Setup

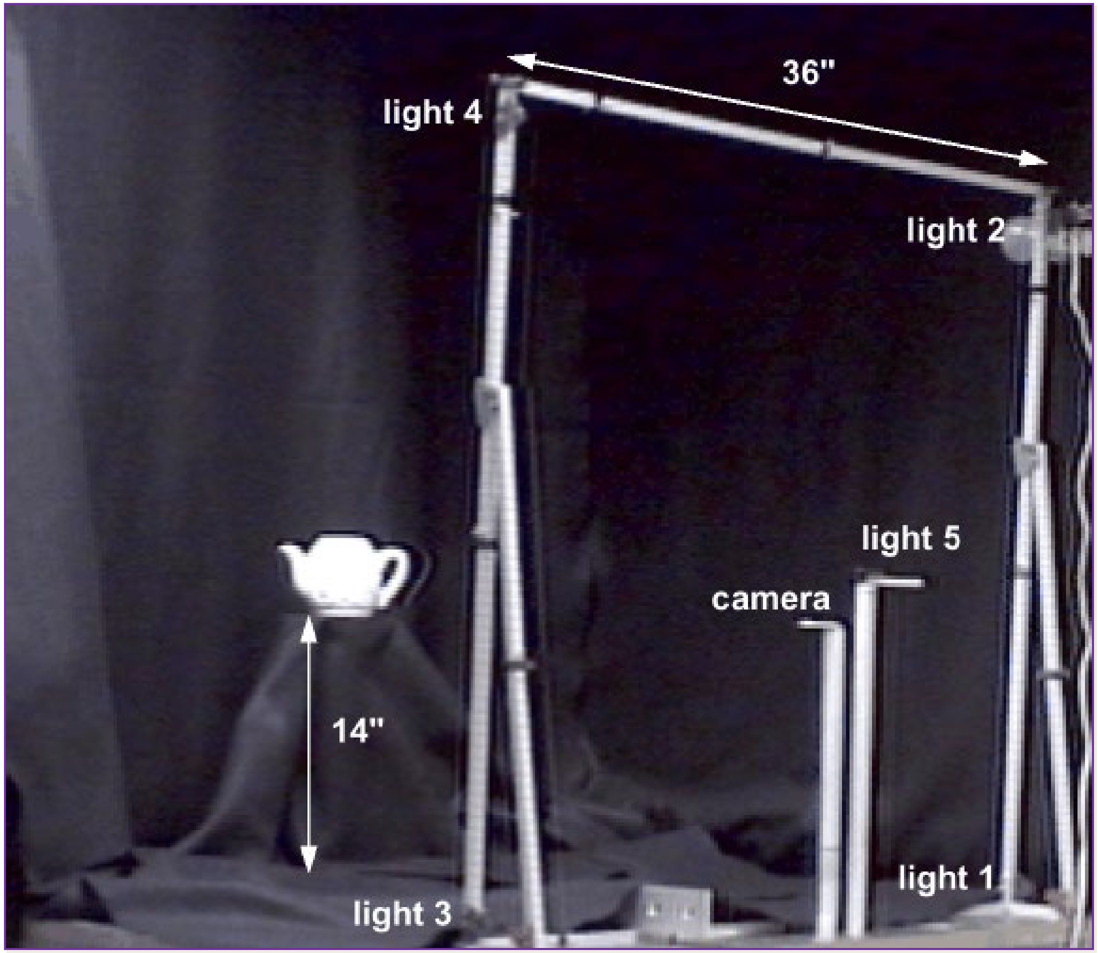

[Rushmeier et al., 1997]

Photometric Stereo Math

- For each point p, can write

$$\rho_p \begin{bmatrix} l_{1,x} & l_{1,y} & l_{1,z} \\ l_{2,x} & l_{2,y} & l_{2,z} \\ l_{3,x} & l_{3,y} & l_{3,z} \end{bmatrix} \begin{bmatrix} n_{p,x} \\ n_{p,y} \\ n_{p,z} \end{bmatrix} = \alpha \begin{bmatrix} E_{p,1} \\ E_{p,2} \\ E_{p,3} \end{bmatrix}$$

- Constant α incorporates light source brightness, camera sensitivity, etc.

Photometric Stereo Math

- Solving above equation gives $(\rho/\alpha)\,\mathrm{n}$
- n must be unit-length $\Rightarrow$ uniquely determined
- Determine ρ up to global constant
- With more than 3 light sources:
  - Discard highest and lowest measurements
  - If still more, solve by least squares

Photometric Stereo Results

Input images

Recovered normals (re-lit)

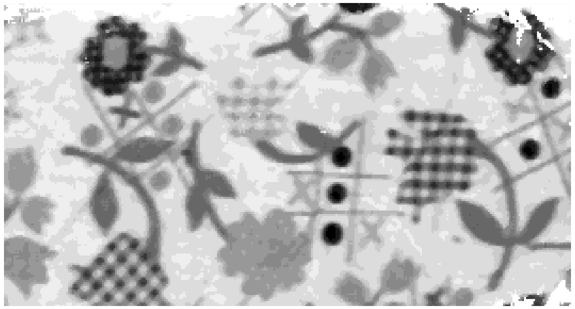

Recovered color

[Rushmeier et al., 1997]

Range Acquisition Taxonomy

- 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

Pulsed Time of Flight

- Basic idea: send out pulse of light (usually laser), time how long it takes to return

$$d = \frac{1}{2} c \Delta t$$

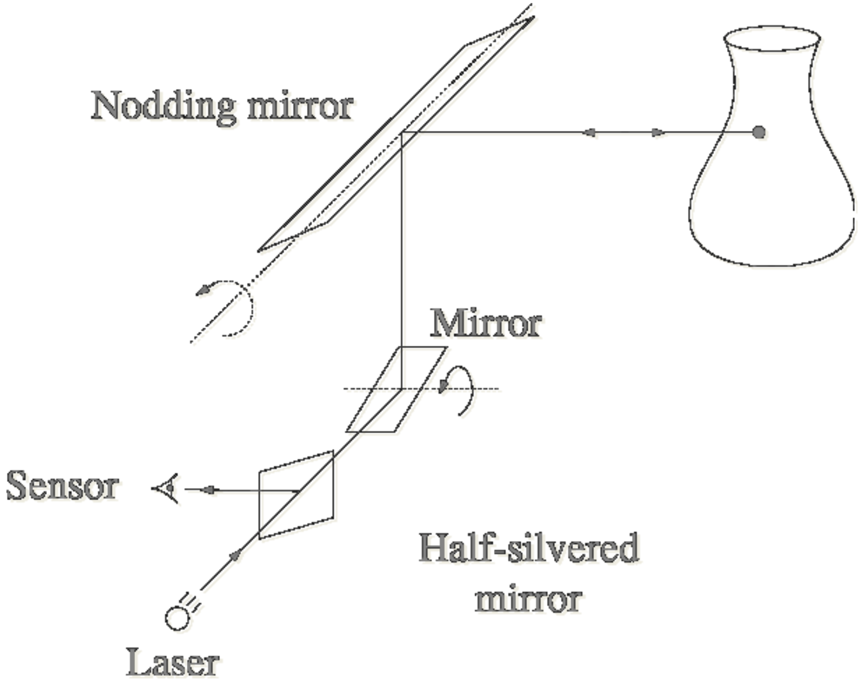

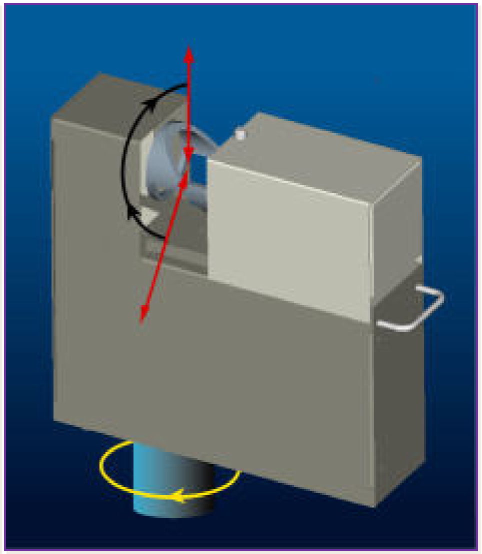

Pulsed Time of Flight

- Advantages:
  - Large working volume (up to 100 m.)
- Disadvantages:
  - Not-so-great accuracy (at best ~5 mm.)
    - Requires getting timing to ~30 picoseconds
    - Does not scale with working volume
- Often used for scanning buildings, rooms, archeological sites, etc.

AM Modulation Time of Flight

- Modulate a laser at frequency ν_m, it returns with a phase shift $\Delta\varphi$

$$d = \frac{1}{2}\left(\frac{c}{\nu_m}\right)\left(\frac{\Delta\varphi \pm 2\pi n}{2\pi}\right)$$

- Note the ambiguity in the measured phase!
  $\Rightarrow$ Range ambiguity of $^1/_2\lambda_m n$

AM Modulation Time of Flight

- Accuracy / working volume tradeoff
  (e.g., noise ~ $^{1}/_{500}$ working volume)
- In practice, often used for room-sized environments
  (cheaper, more accurate than pulsed time of flight)

Range Acquisition Taxonomy

- 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

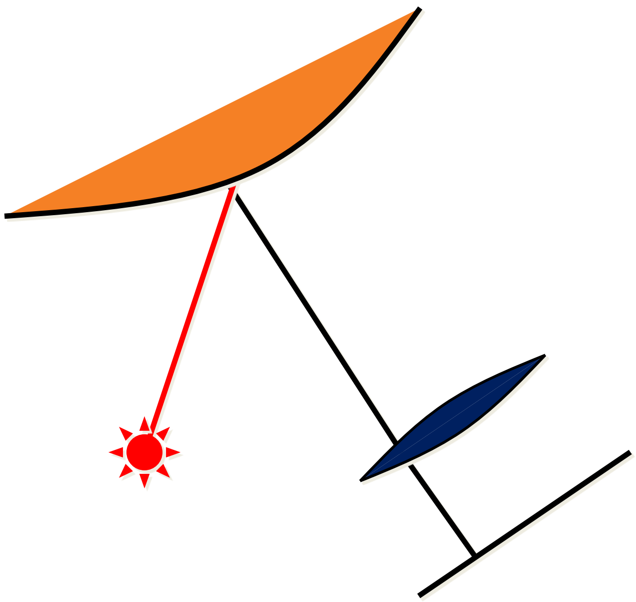

Triangulation: Moving the Camera and Illumination

- Moving independently leads to problems with focus, resolution
- Most scanners mount camera and light source rigidly, move them as a unit

Triangulation: Moving the Camera and Illumination

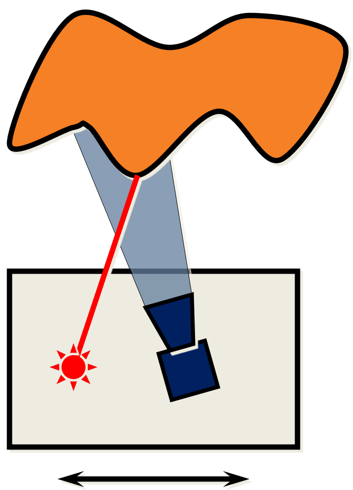

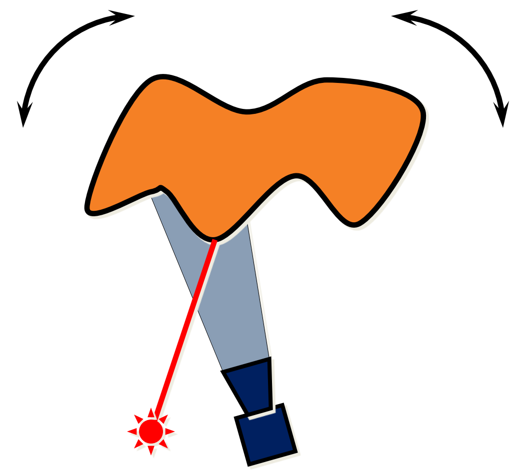

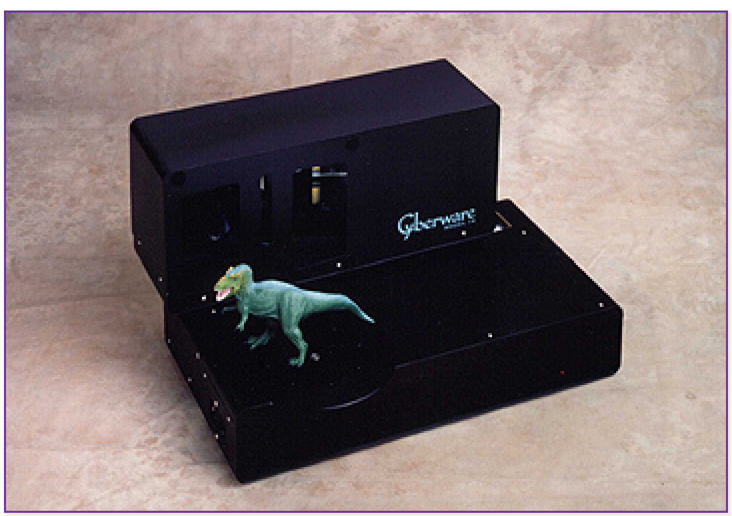

Triangulation: Moving the Camera and Illumination

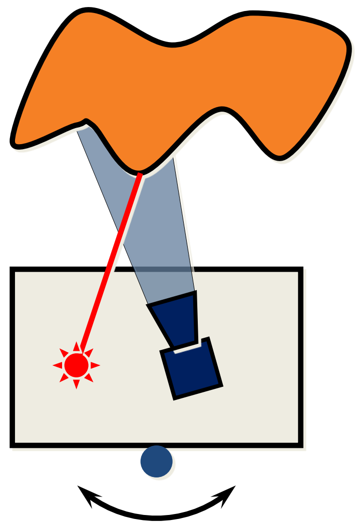

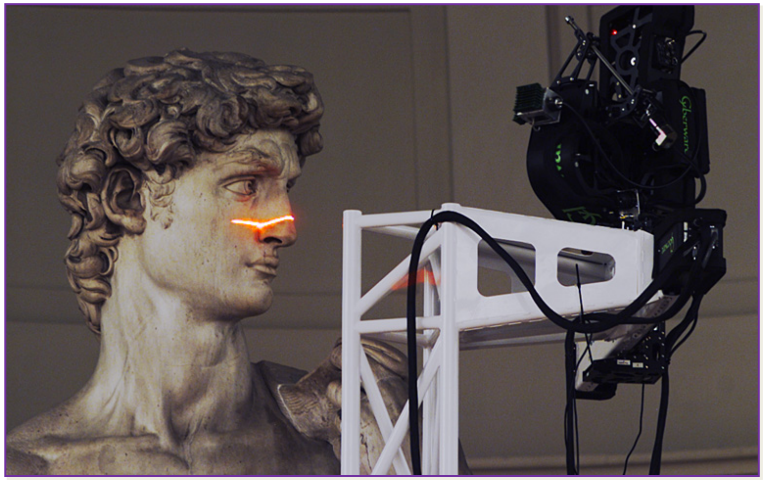

Triangulation: Extending to 3D

- Possibility #1: add another mirror (flying spot)
- Possibility #2: project a stripe, not a dot

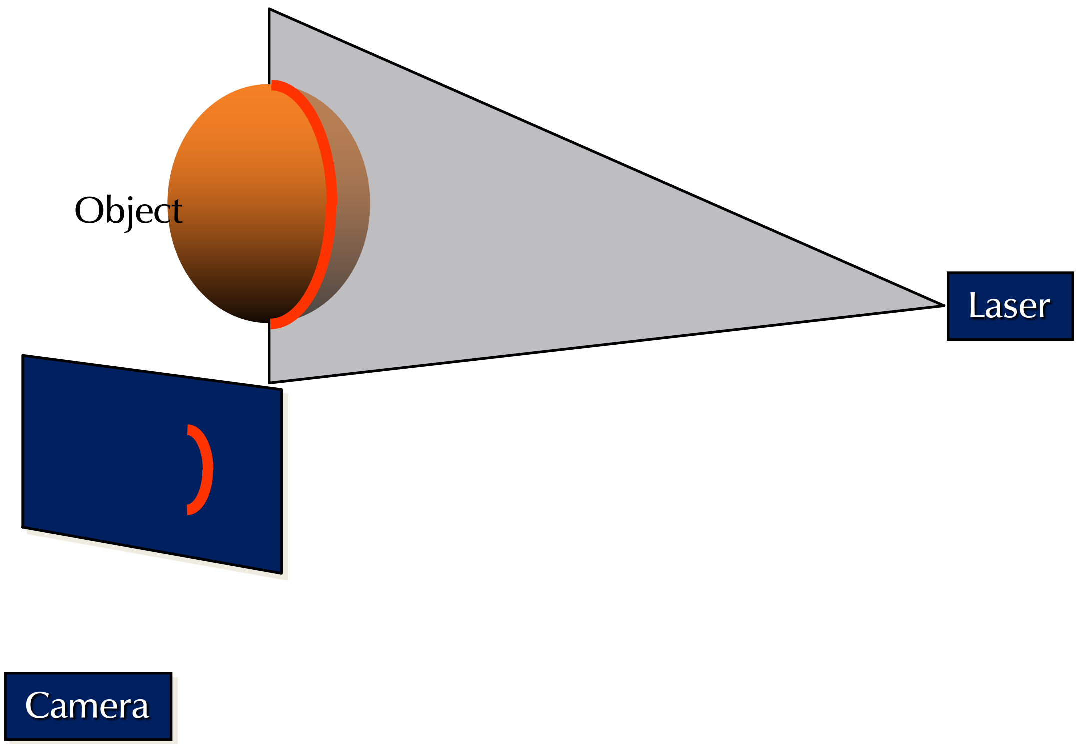

Triangulation Scanner Issues

- Accuracy proportional to working volume (typical is ~1000:1)
- Scales down to small working volume (e.g. 5 cm. working volume, 50 μm. accuracy)
- Does not scale up (baseline too large…)
- Two-line-of-sight problem (shadowing from either camera or laser)
- Triangulation angle: non-uniform resolution if too small, shadowing if too big (useful range: 15°-30°)

Triangulation Scanner Issues

- Material properties (dark, specular)
- Subsurface scattering
- Laser speckle
- Edge curl
- Texture embossing

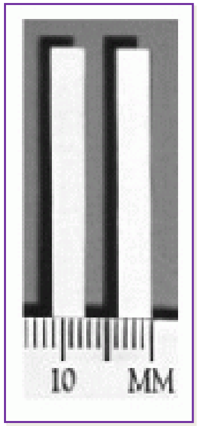

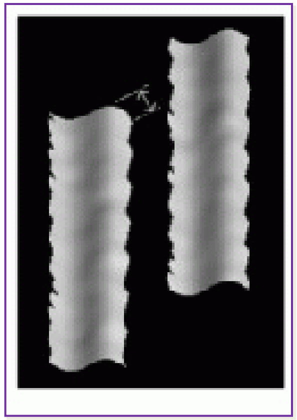

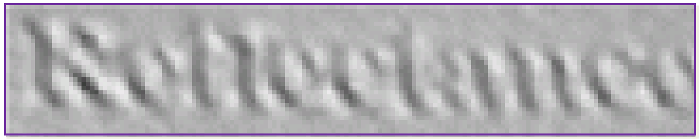

Multi-Stripe Triangulation

- To go faster, project multiple stripes
- But which stripe is which?
- Answer #1: assume surface continuity

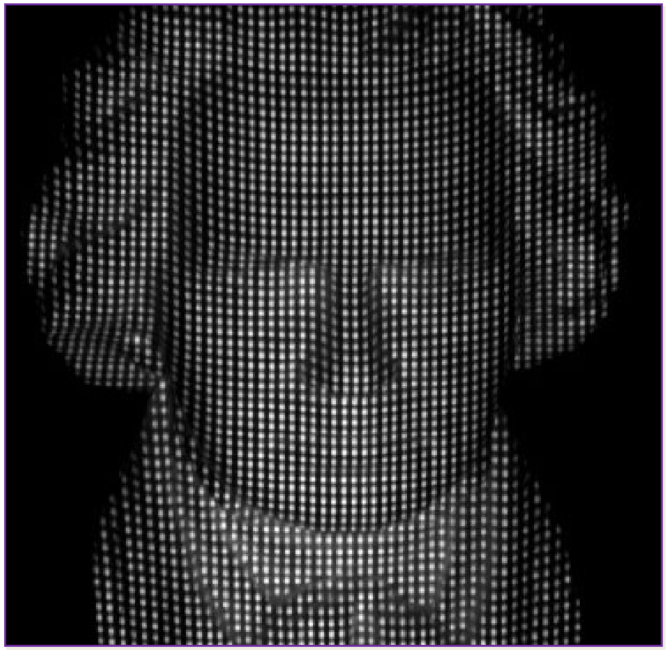

Multi-Stripe Triangulation

- To go faster, project multiple stripes
- But which stripe is which?
- Answer #2: colored stripes (or dots)

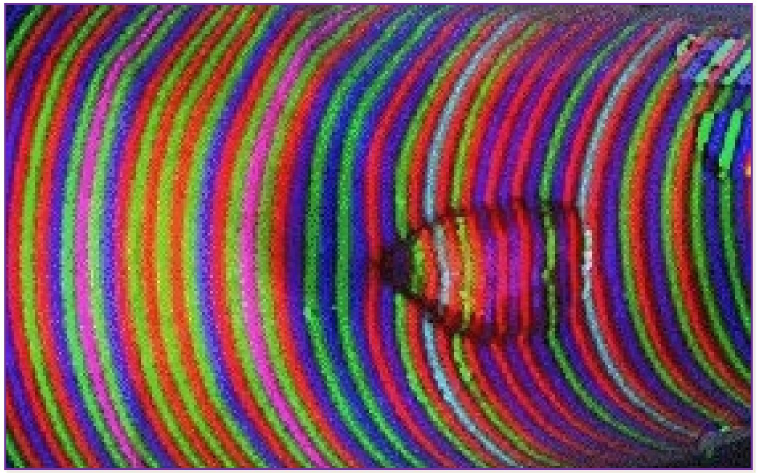

Multi-Stripe Triangulation

- To go faster, project multiple stripes
- But which stripe is which?
- Answer #3: time-coded stripes

Time-Coded Light Patterns

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

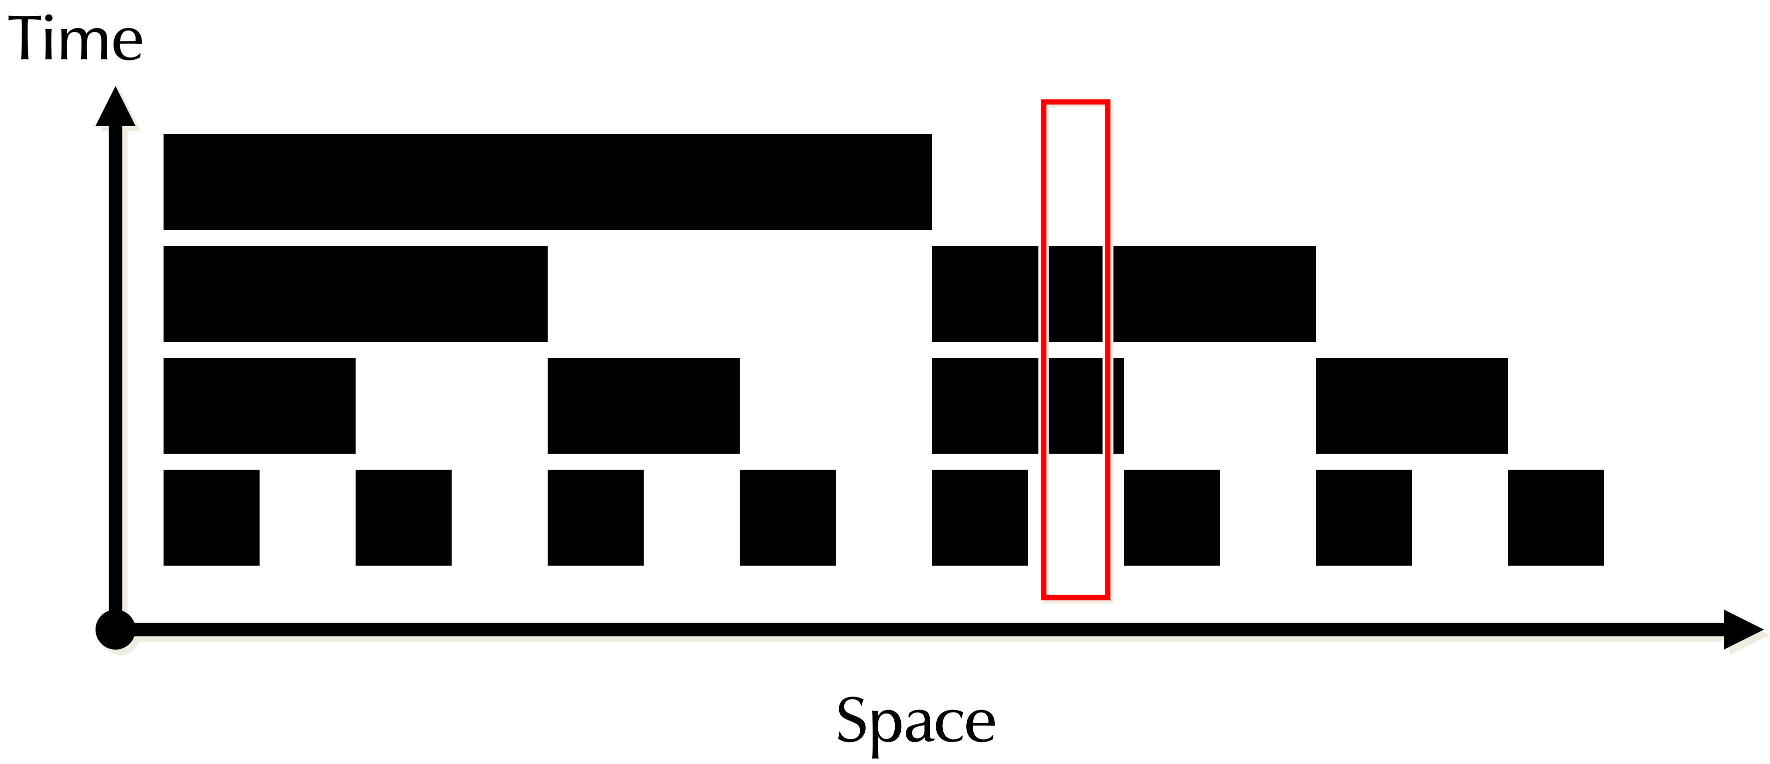

Multi-Stripe Triangulation

- To go faster, project multiple stripes
- But which stripe is which?
- Answer #4: space-coded stripes (or dots)

Microsoft Kinect (1st generation)

Range Processing Pipeline

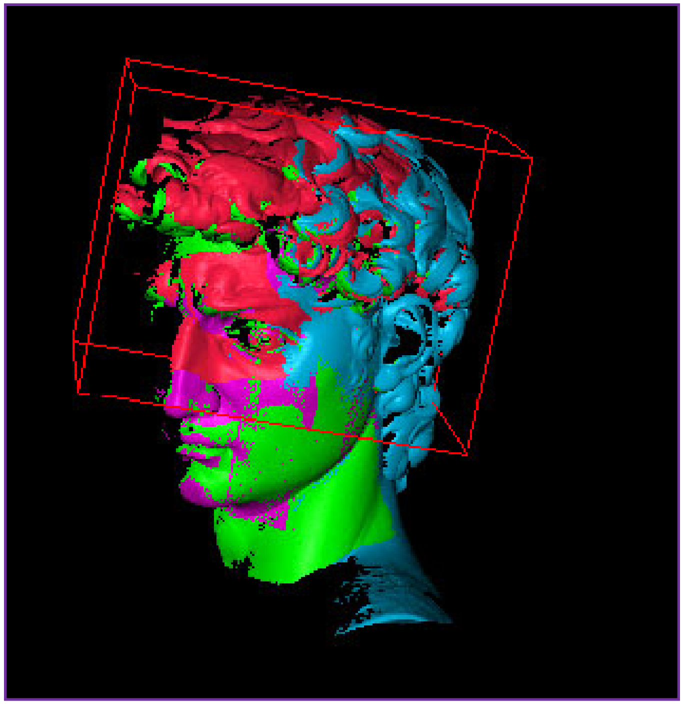

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Range Processing Pipeline

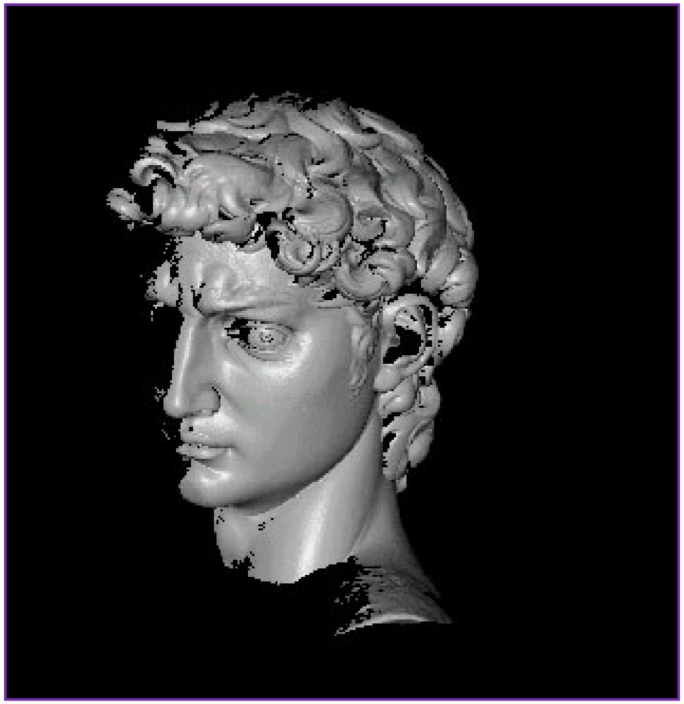

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Range Processing Pipeline

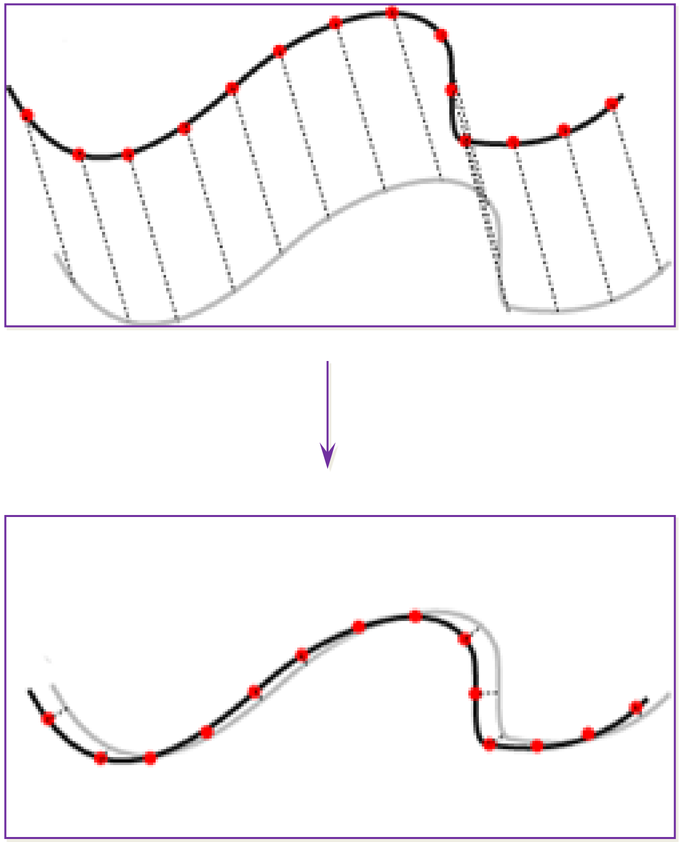

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Range Processing Pipeline

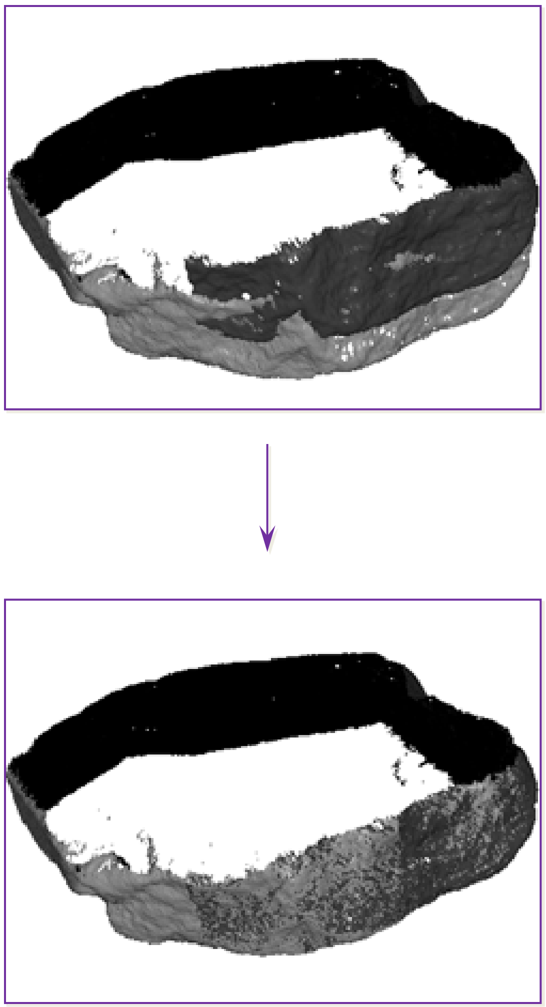

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Range Processing Pipeline

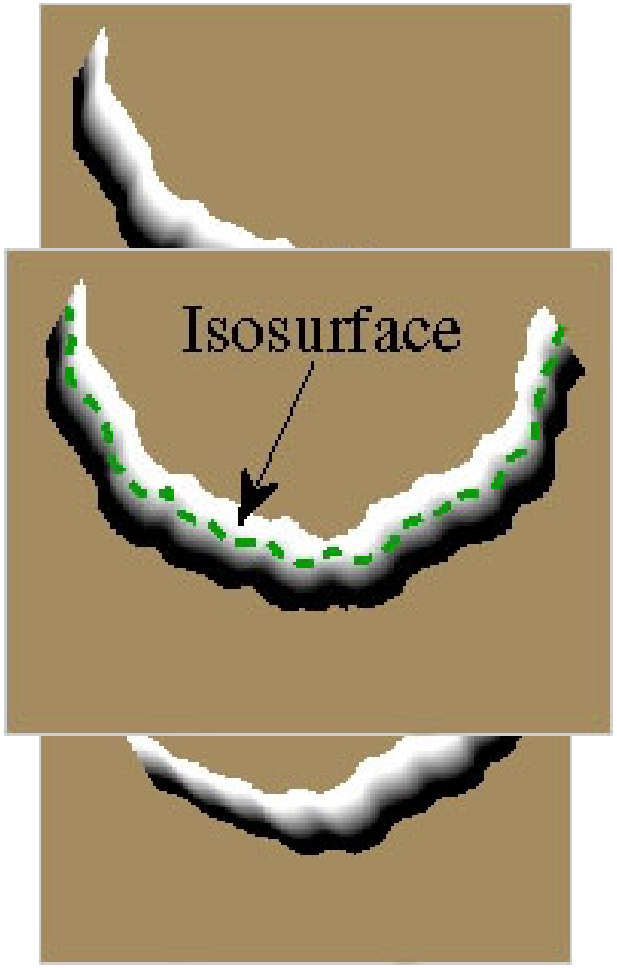

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Range Processing Pipeline

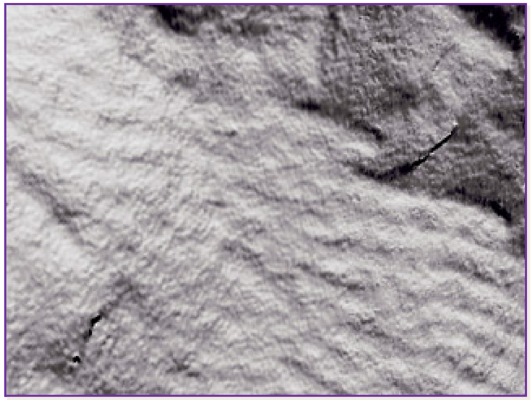

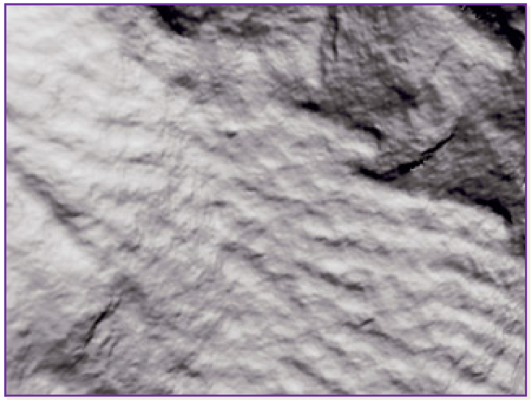

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

3D Scanning in Archaeology

Why 3D Scanning in Archaeology?

- Suggest matches based on shape

Visualization of Surface Markings

Even Less Realistic…

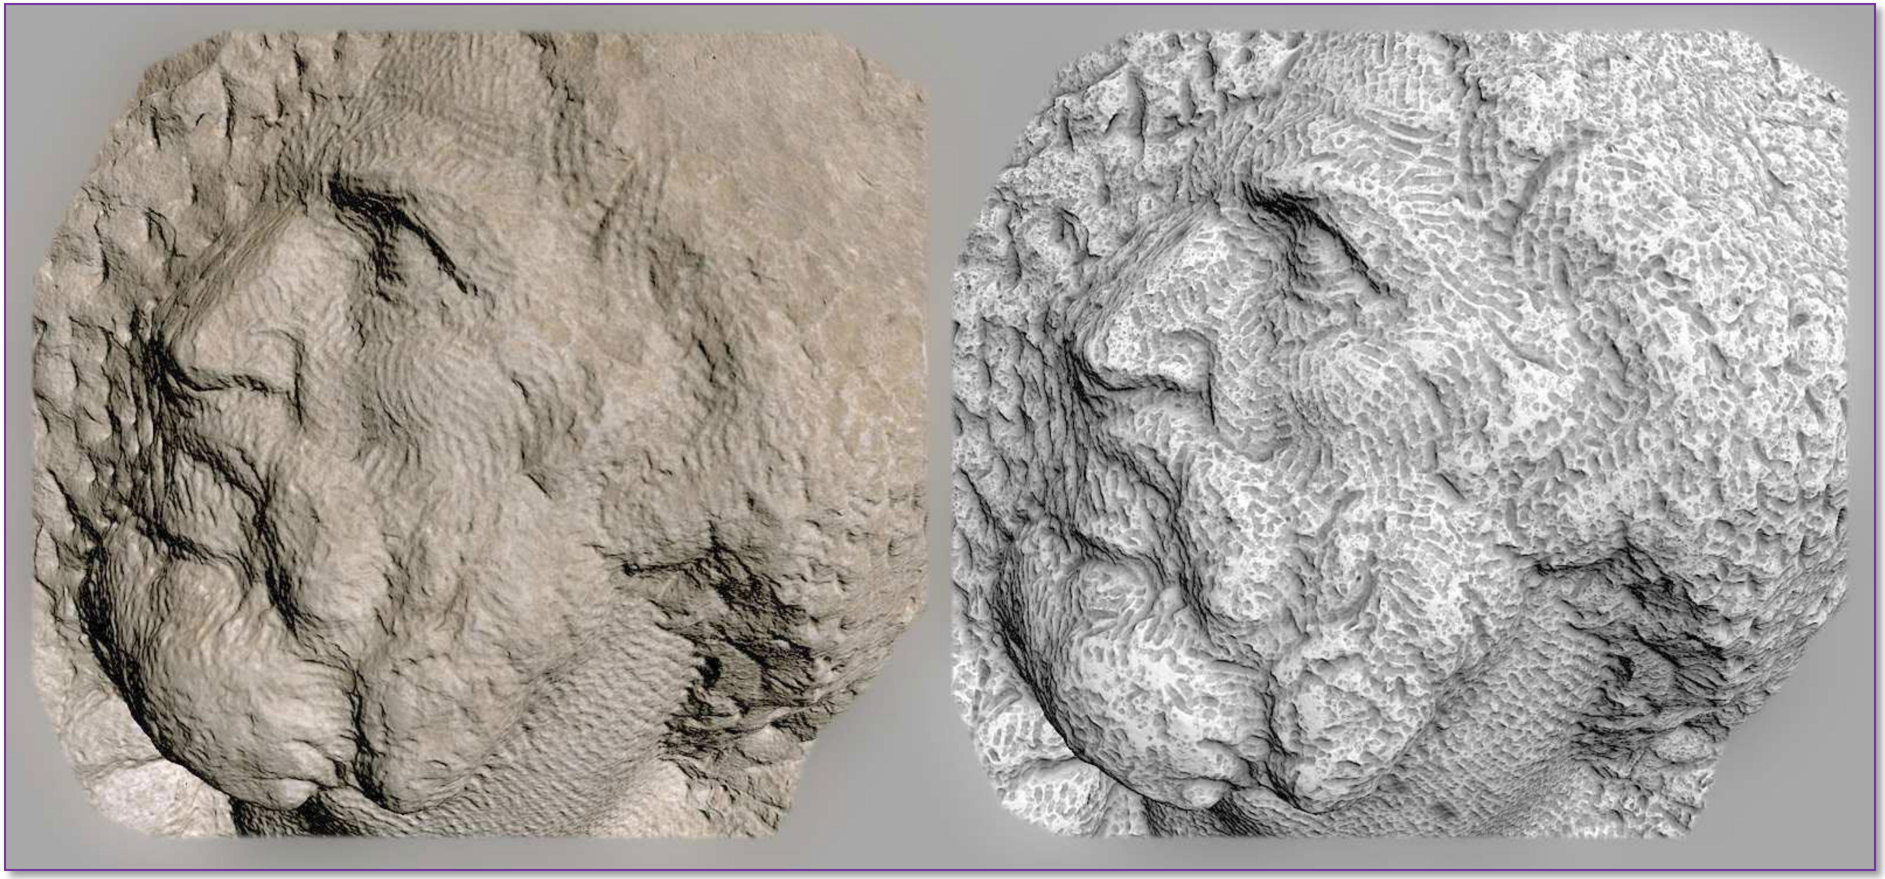

Changing Lighting

The Importance of Viewpoint

classic 3/4 view

left profile

Courtesy Marc Levoy

The Importance of Viewpoint

face-on view

Courtesy Marc Levoy

Conservation

Tensile stresses in the left leg with the statue tilted 3 degrees forward, as it was in 1871.

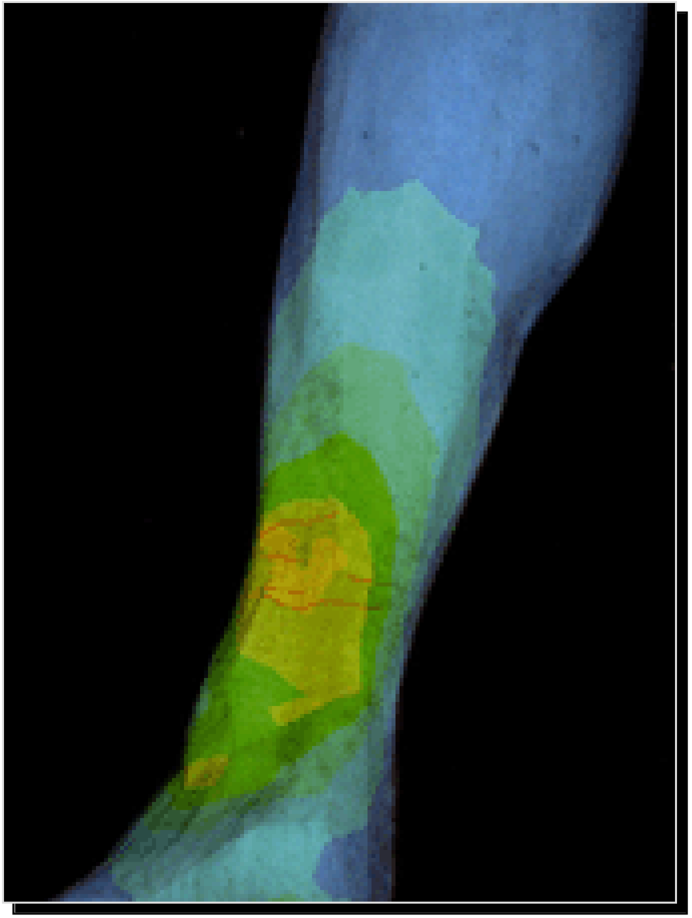

Courtesy Stanford University and Bracci et al.

Conservation

Exposure of the statue
to different contaminants

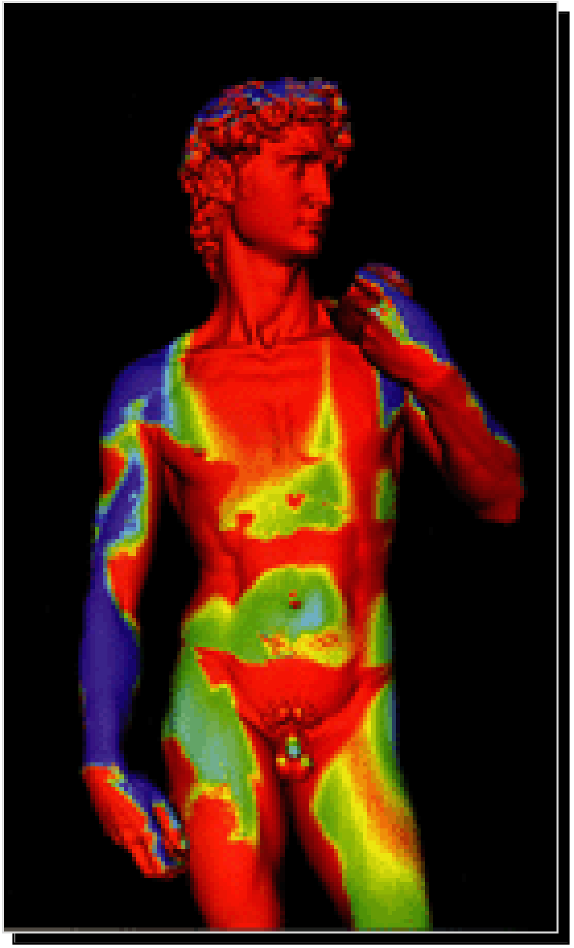

Courtesy Stanford University and Bracci et al.

Kiosk in Galleria dell'Accademia

Applications of 3D Scanning – Scanning Sculptures

- The Pietà Project

  IBM Research

- The Digital Michelangelo Project

  Stanford University

- The Great Buddha Project

  University of Tokyo

Why Scan Sculptures?

- Sculptures interesting objects to look at
- Introduce scanning to new disciplines
  - Art: studying working techniques
  - Art history
  - Cultural heritage preservation
  - Archeology
- High-visibility projects

Why Scan Sculptures?

- Challenging
  - High detail, large areas
  - Large data sets
  - Field conditions
  - Pushing hardware, software technology
- But not too challenging
  - Simple topology
  - Possible to scan most of surface

Issues Addressed

- Resolution
- Coverage
  - Theoretical: limits of scanning technologies
  - Practical: physical access, time
- Type of data
  - High-res 3D data vs. coarse 3D + normal maps
  - Influenced by eventual application
- Intellectual Property

IBM's Pietà Project

- Michelangelo's "Florentine Pietà"
- Late work (1550s)
- Partially destroyed by Michelangelo, recreated by his student
- Currently in the Museo dell'Opera del Duomo in Florence

Who?

- Dr. Jack Wasserman, professor emeritus of art history at Temple University
- Visual and Geometric Computing group @ IBM Research:

Fausto Bernardini	Gabriel Taubin
Holly Rushmeier	Andre Gueziec
Ioana Martin	Claudio Silva
Joshua Mittleman	

Scanner

- Visual Interface "Virtuoso"
- Active multibaseline stereo
- Projector (stripe pattern),
  6 B&W cameras, 1 color camera
- Augmented with 5 extra "point" light sources for photometric stereo (active shape from shading)

Data

- Range data has 2 mm spacing, 0.1mm noise
- Each range image: 10,000 points, 20×20 cm
- Color data: 5 images with controlled lighting, 1280×960, 0.5 mm resolution
- Total of 770 scans, 7.2 million points

Scanning

- Final scan June 1998, completed July 1999
- Total scanning time: 90 hours over 14 days (includes equipment setup time)

Postprocessing

- Use 11×11 grid of projected laser dots to help with pairwise alignment
- Align all scans to each other, then apply nonrigid "conformance smoothing"
- Reconstruct surface using BPA
- Compute normal and albedo maps, align to geometry

Results

The Digital Michelangelo Project

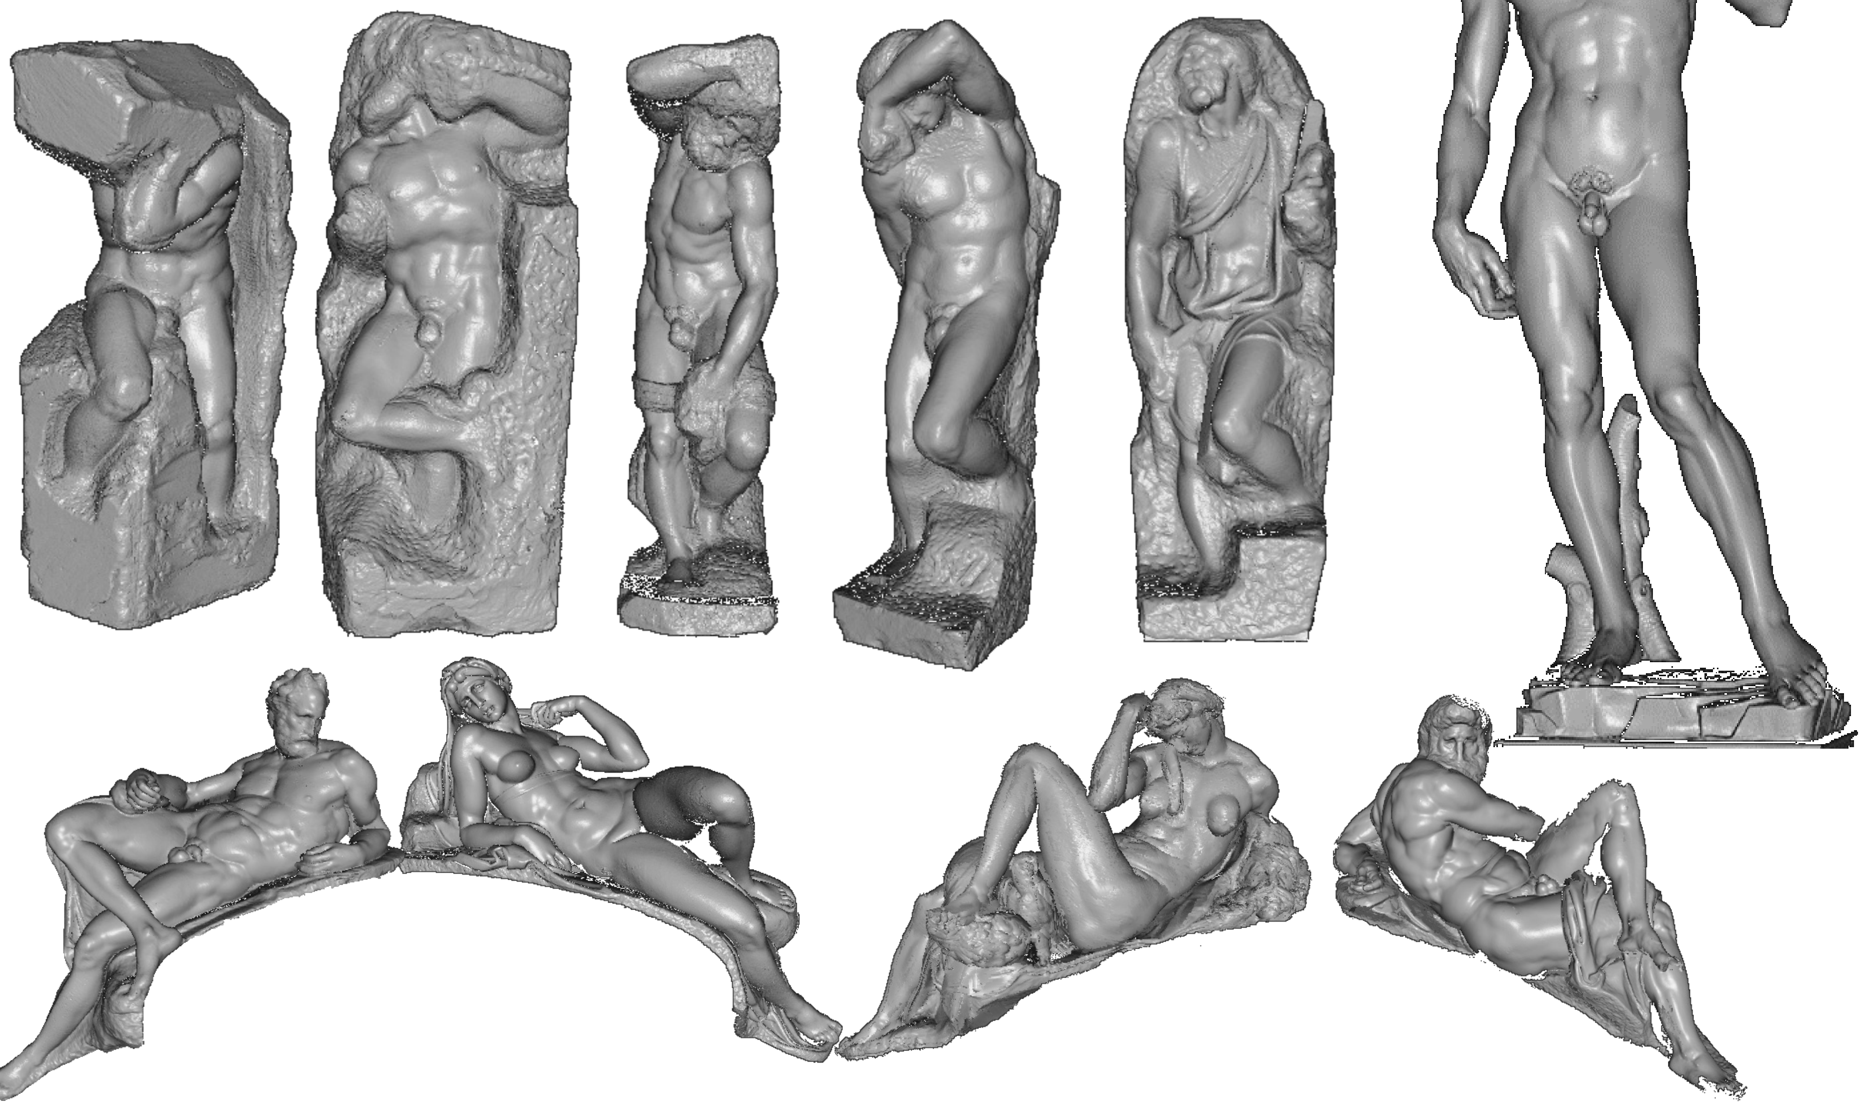

Goals

- Scan 10 sculptures by Michelangelo
- High-resolution (“quarter-millimeter”) geometry
- Side projects: architectural scanning (Accademia and Medici chapel), scanning fragments of Forma Urbis Romae

Why Capture Chisel Marks?

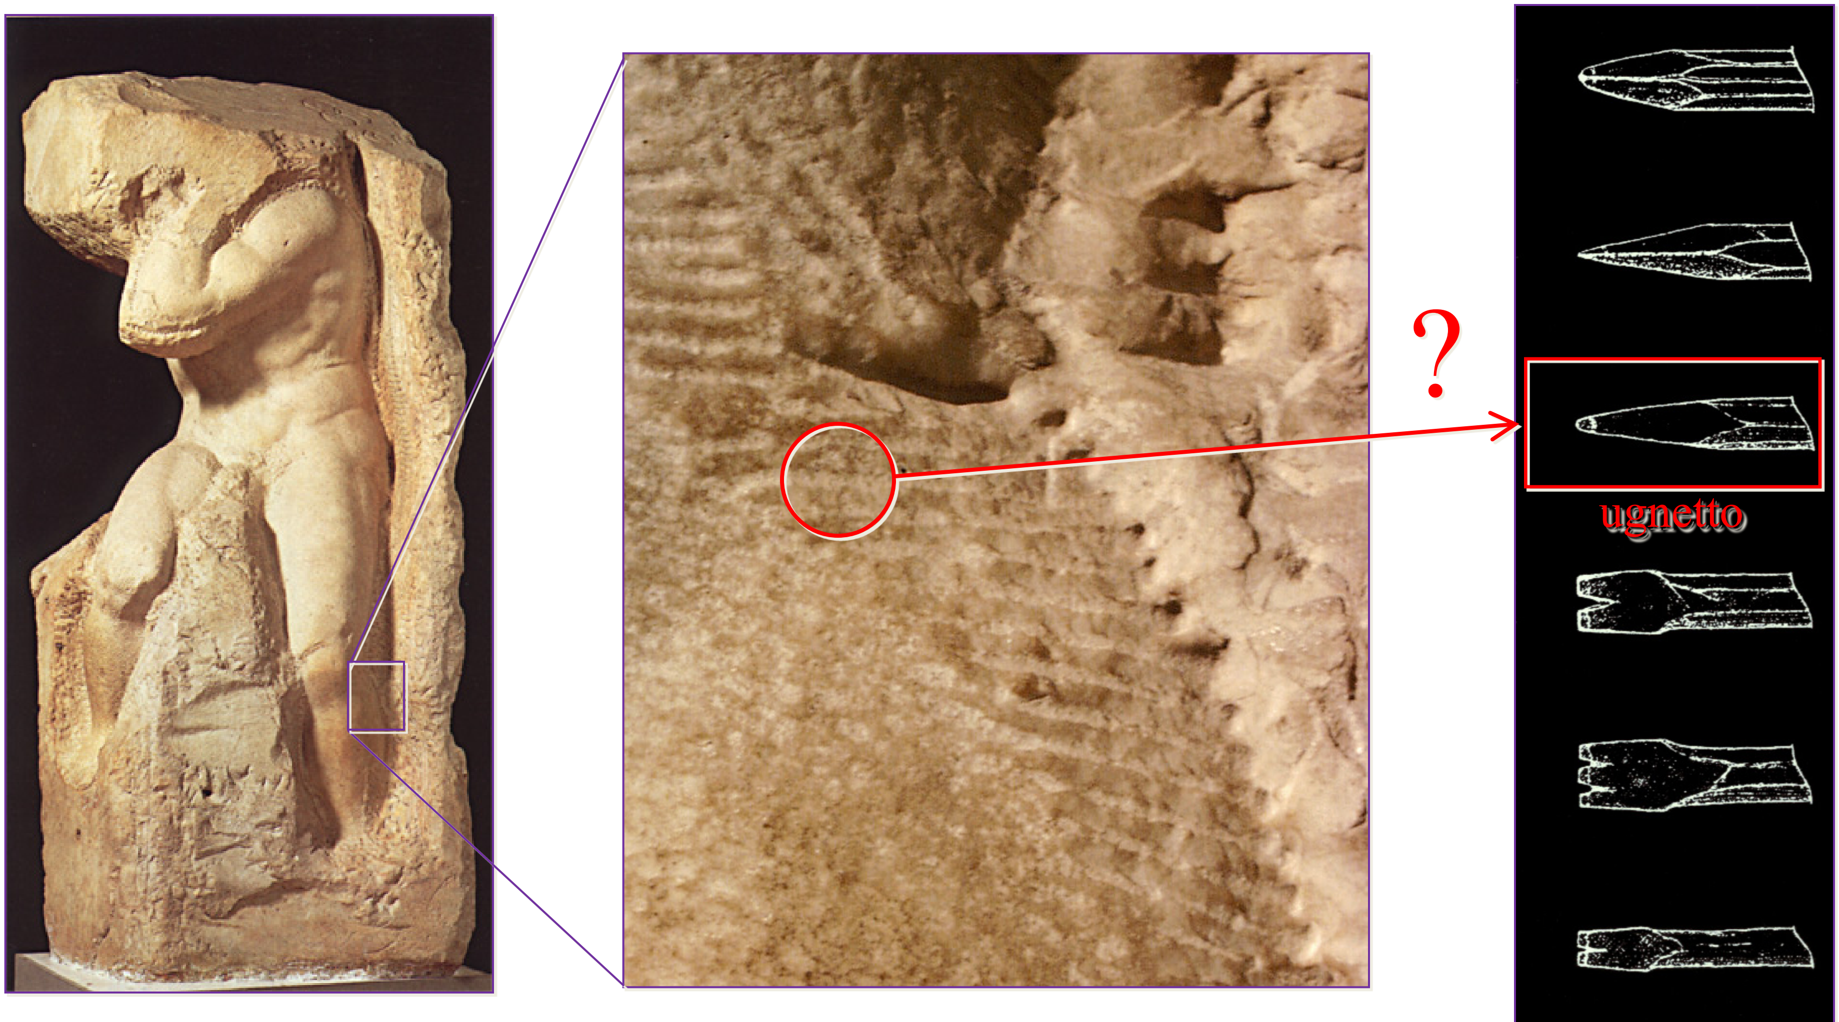

Atlas (Accademia)

Why Capture Chisel Marks as Geometry?

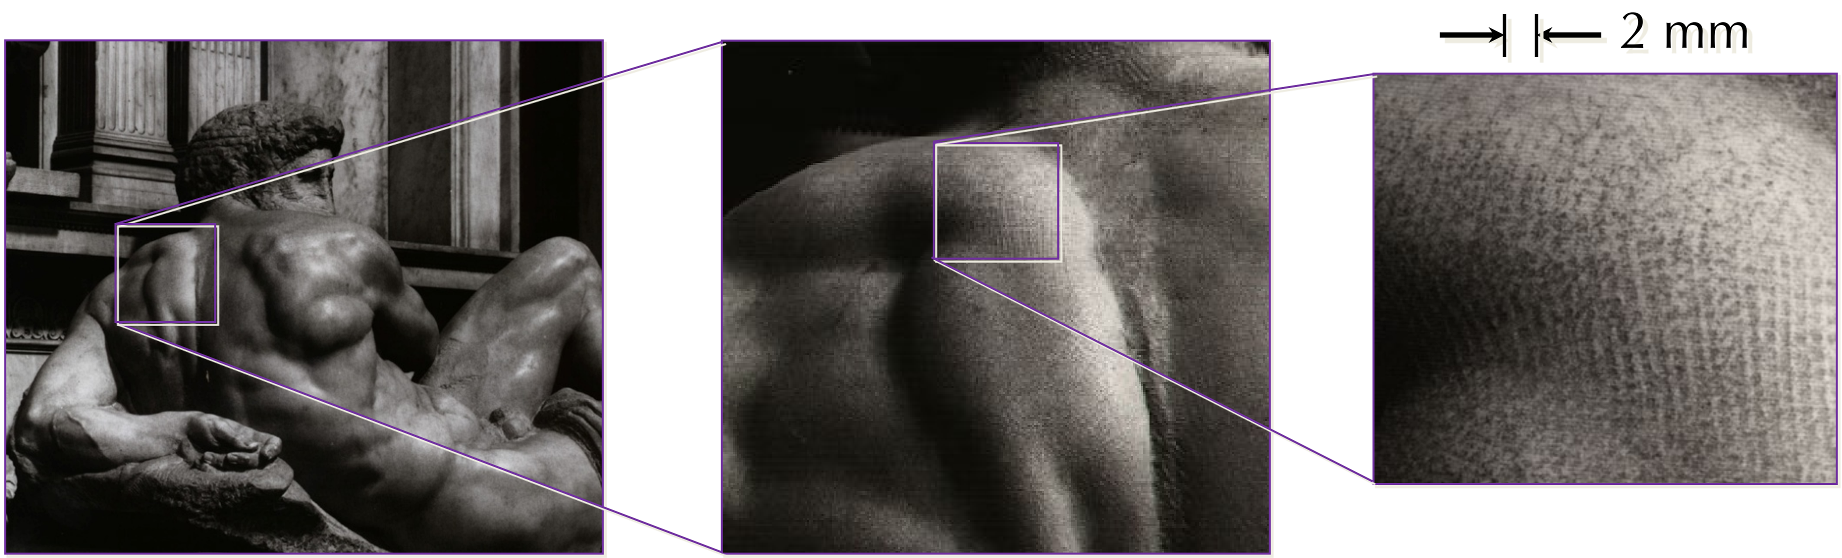

Day (Medici Chapel)

Who?

Faculty and staff

- Prof. Brian Curless
- John Gerth
- Jelena Jovanovic
- Prof. Marc Levoy
- Lisa Pacelle
- Domi Pitturo
- Dr. Kari Pulli

Graduate students

- Sean Anderson
- Barbara Caputo
- James Davis
- Dave Koller
- Lucas Pereira
- Szymon Rusinkiewicz
- Jonathan Shade
- Marco Tarini
- Daniel Wood

Undergraduates

- Alana Chan
- Kathryn Chinn
- Jeremy Ginsberg
- Matt Ginzton
- Unnur Gretarsdottir
- Rahul Gupta
- Wallace Huang
- Dana Katter
- Ephraim Luft
- Dan Perkel
- Semira Rahemtulla
- Alex Roetter
- Joshua Schroeder
- Maisie Tsui
- David Weekly

In Florence

- Dottssa Cristina Acidini
- Dottssa Franca Falletti
- Dottssa Licia Bertani
- Alessandra Marino
- Matti Auvinen

In Rome

- Prof. Eugenio La Rocca
- Dottssa Susanna Le Pera
- Dottssa Anna Somella
- Dottssa Laura Ferrea

In Pisa

- Roberto Scopigno

Sponsors

- Interval Research
- Paul G. Allen Foundation for the Arts
- Stanford University

Equipment donors

- Cyberware
- Cyra Technologies
- Faro Technologies
- Intel
- Silicon Graphics
- Sony
- 3D Scanners

Scanner Design

4 motorized axes

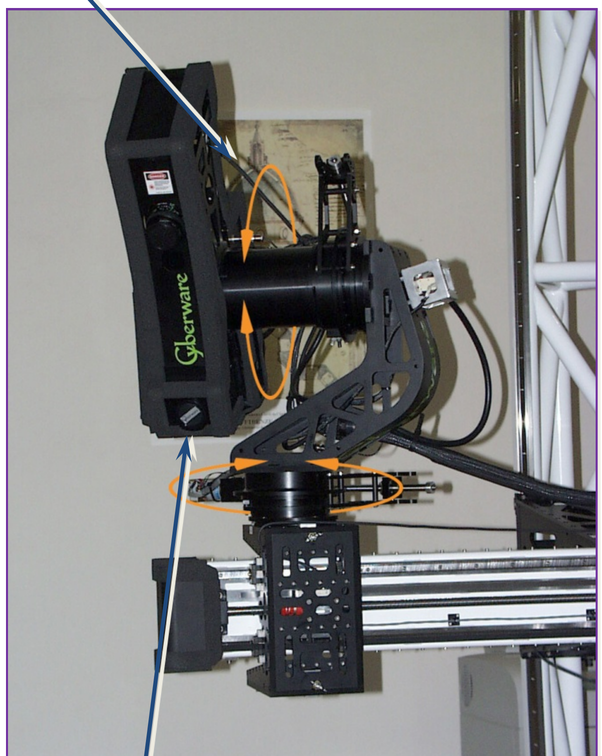

laser, range camera,
white light, and color camera

- Flexibility
  - outward-looking rotational scanning
  - 16 ways to mount scan head on arm
- Accuracy
  - center of gravity kept stationary during motions
  - precision drives, vernier homing, stiff trusses

Scanning a Large Object

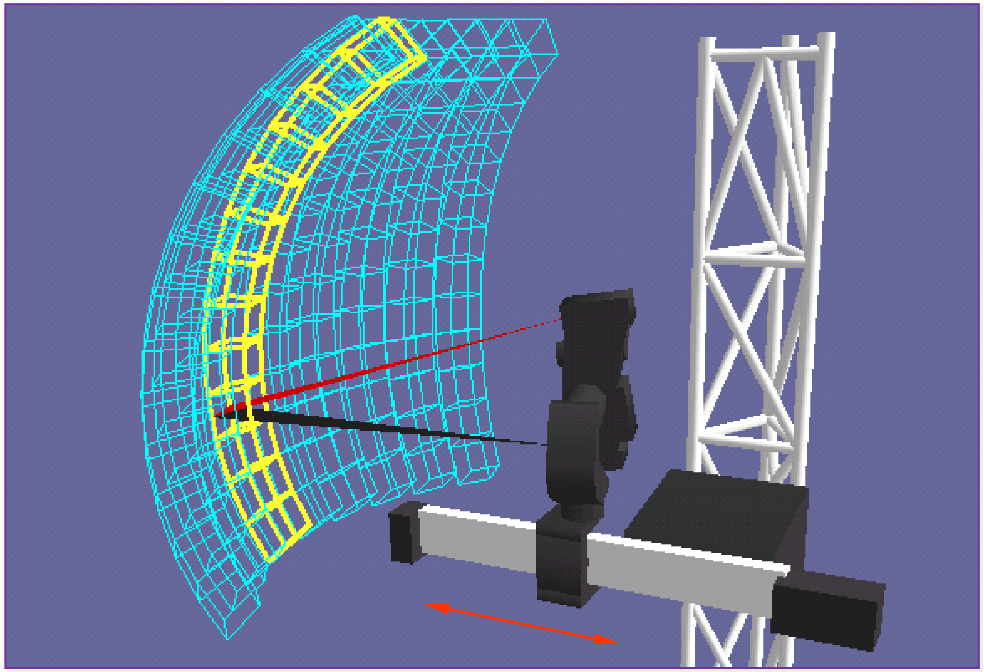

- Calibrated motions
  - pitch (yellow)
  - pan (blue)
  - horizontal translation (orange)
- Uncalibrated motions
  - vertical translation
  - rolling the gantry
  - remounting the scan head

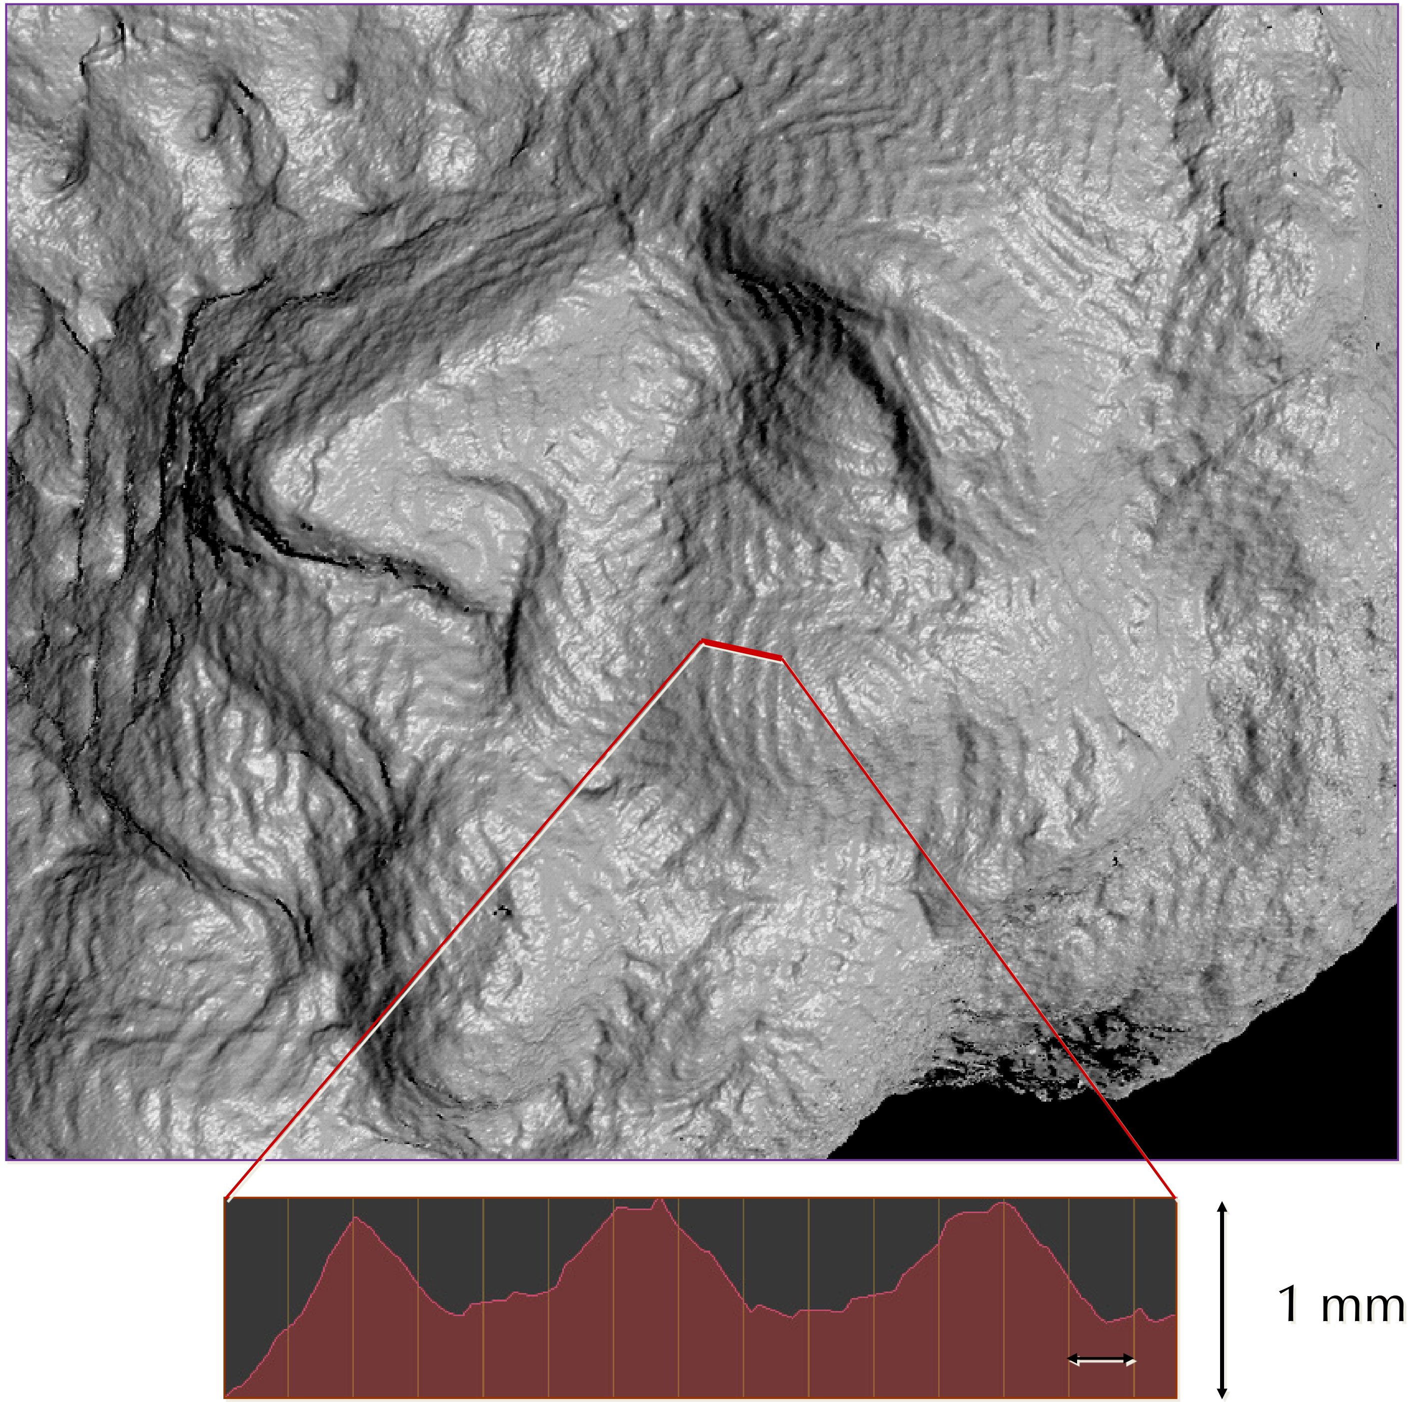
1 mm

Postprocessing

- Manual initial alignment
- Pairwise ICP, then global registration
- VRIP (parallelized across subvolumes)
- Use high-res geometry to discard bad color data, perform inverse lighting calculations

Statistics About the Scan of David

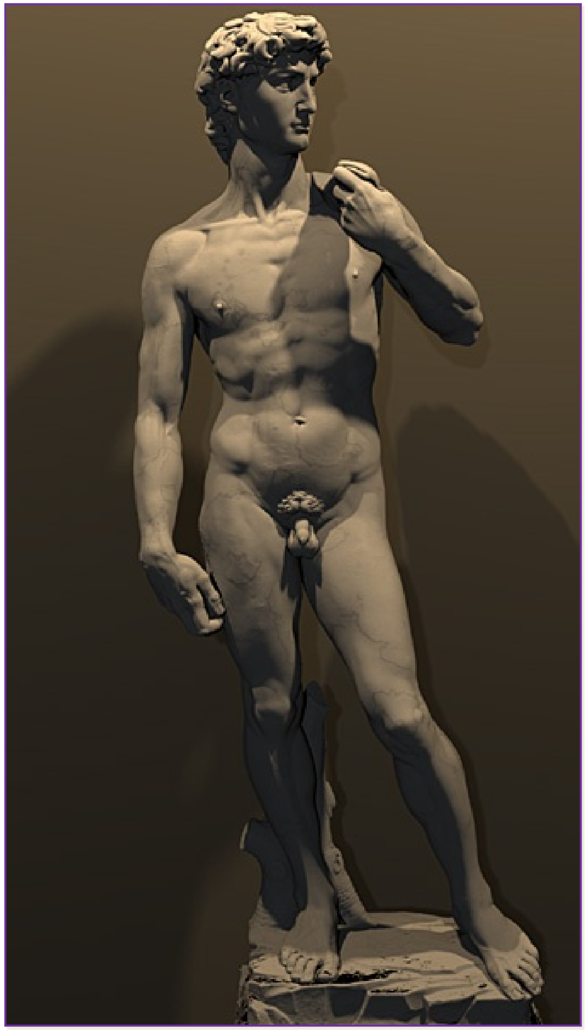

- 480 individually aimed scans
- 0.3 mm sample spacing
- 2 billion polygons
- 7,000 color images
- 32 gigabytes
- 30 nights of scanning
- 22 people

Head of Michelangelo's David

Photograph

1.0 mm computer model

Side project:
The Forma Urbis Romae

Forma Urbis Romae Fragment

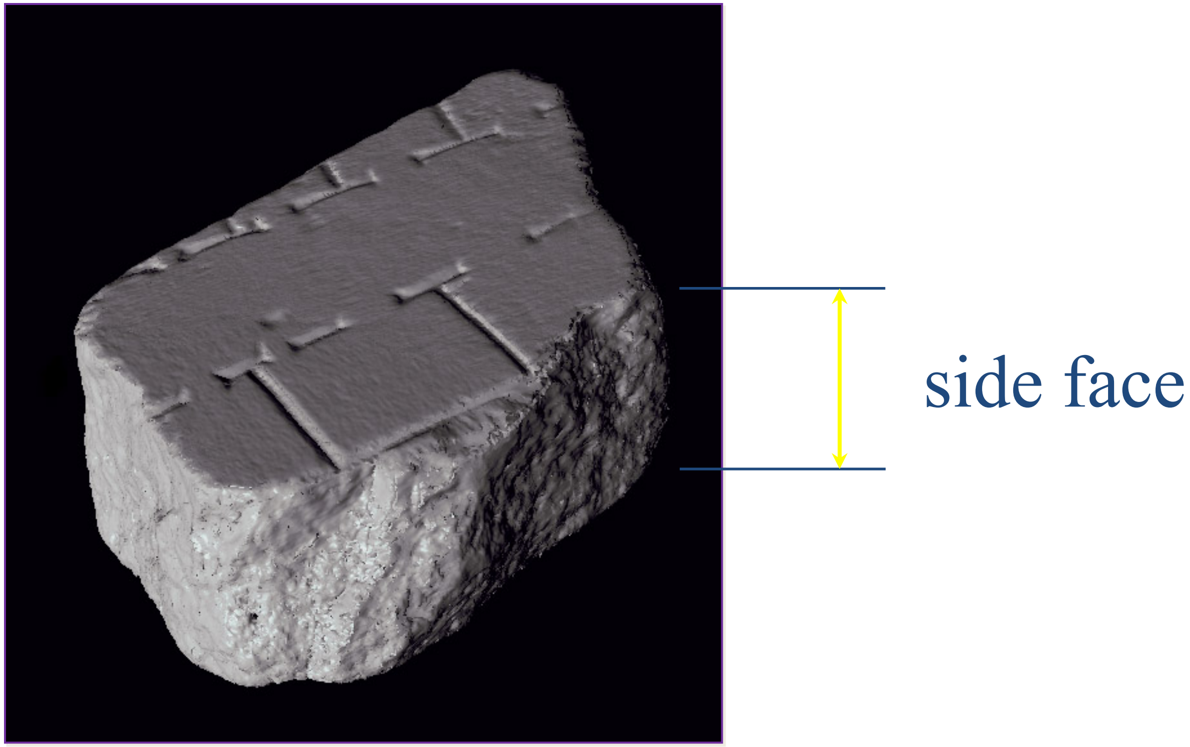

Forma Urbis Romae

Hard Problems

- Keeping scanner calibrated is hard in the lab, really hard in the museum
- Dealing with large data sets is painful
- Filling all the holes converges only asymptotically (if it converges at all…)

The Great Buddha Project

- Great Buddha of Kamakura
- Original made of wood, completed 1243
- Covered in bronze and gold leaf, 1267
- Approx. 15 m tall
- Goal: preservation of cultural heritage

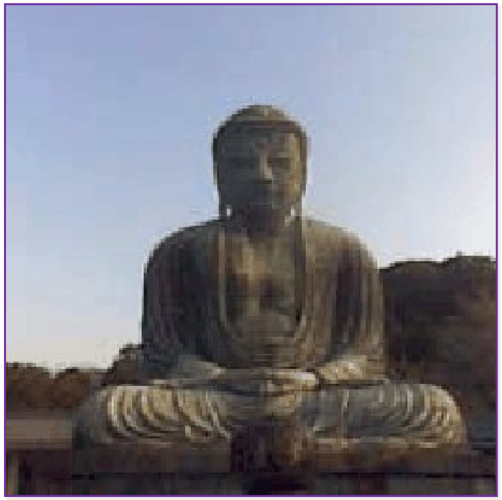

Who?

- Institute of Industrial Science, University of Tokyo

Daisuke Miyazaki
Takeshi Ooishi
Taku Nishikawa
Ryusuke Sagawa
Ko Nishino
Takashi Tomomatsu
Yutaka Takase
Katsushi Ikeuchi

Scanner

- Cyrax range scanner
  by Cyra Technologies
- Laser pulse time-of-flight
- Accuracy: 4 mm
- Range: 100 m

Processing

- 20 range images (a few million points)
- Simultaneous all-to-all ICP
- Variant of volumetric merging (parallelized)

Results

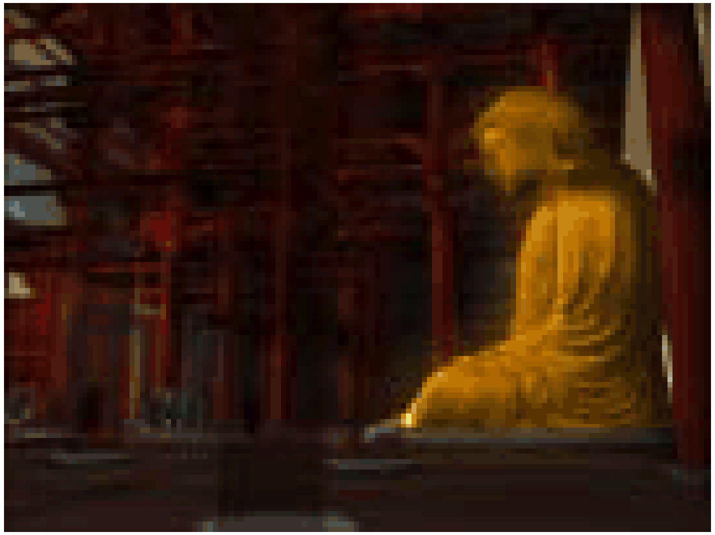

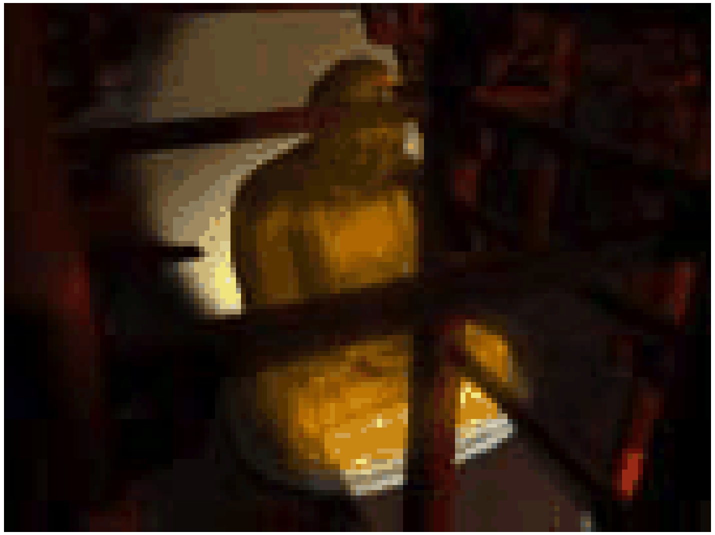

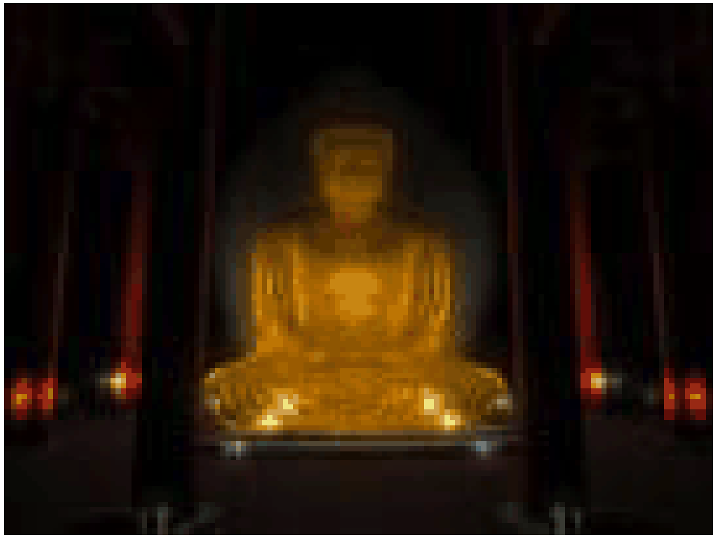

[redacted]