The Digital Michelangelo Project

Marc Levoy

Computer Science Department
Stanford University

Executive summary

Atlas

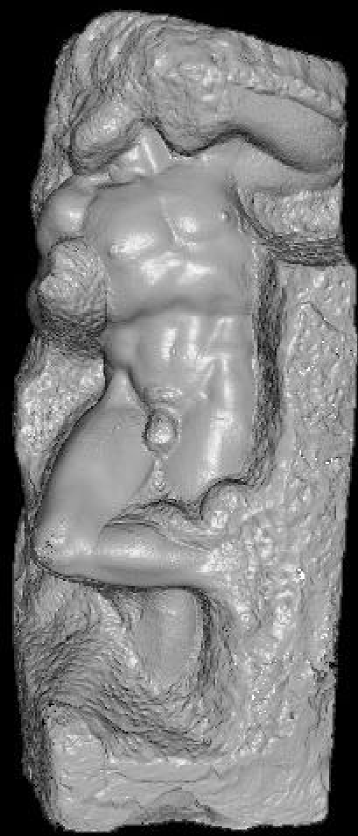

Awakening

Bearded

Youthful

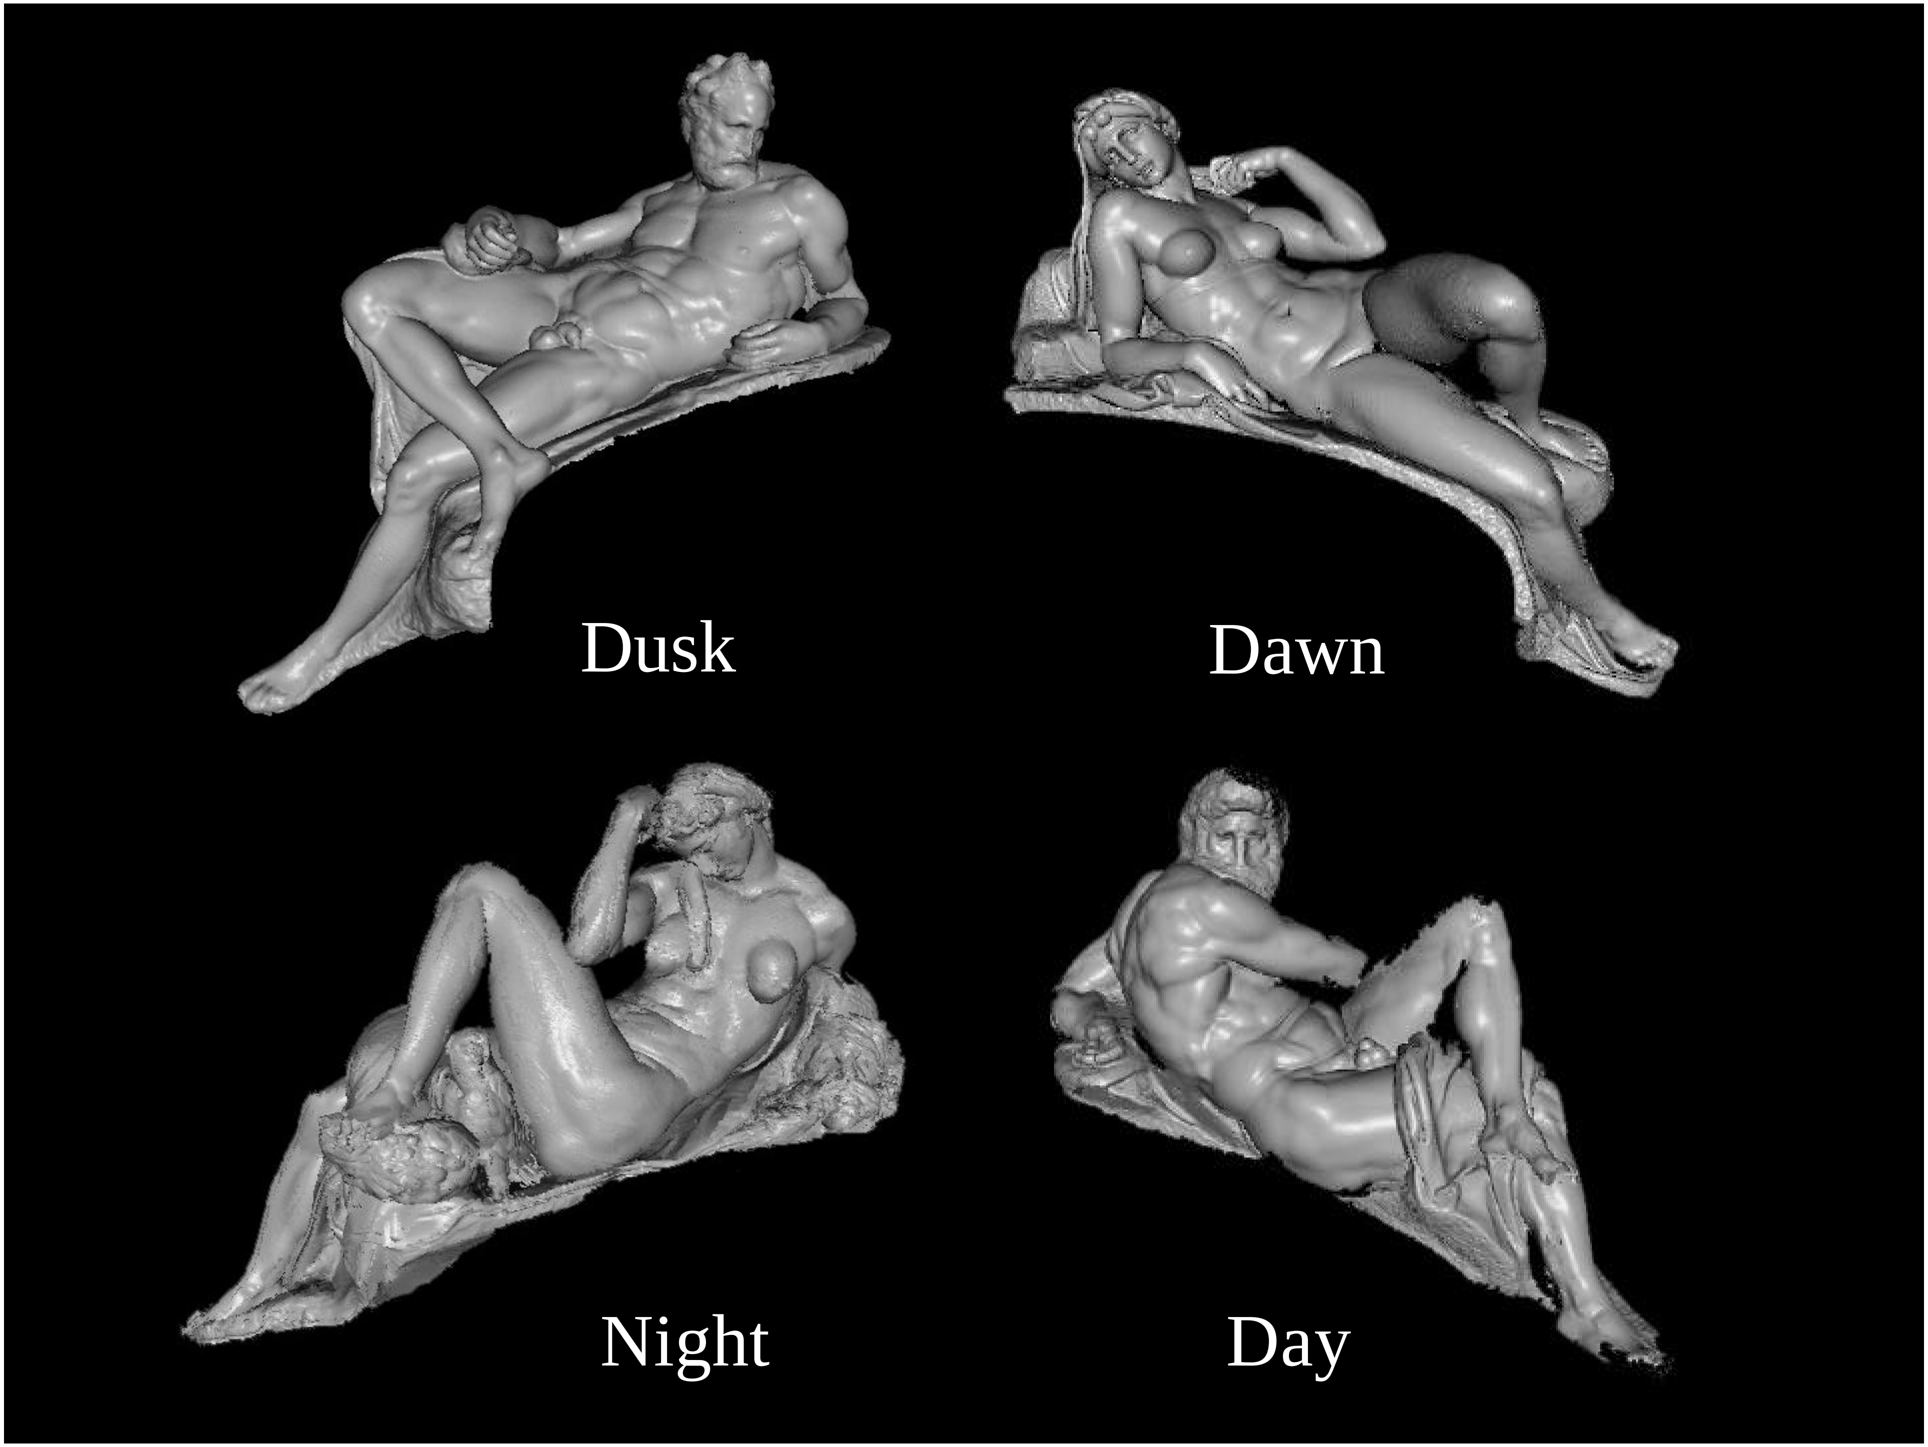
Dusk
Dawn
Night
Day

St. Matthew

David

Forma Urbis Romae

Executive summary

Motivations

- push 3D scanning technology
- tool for art historians
- lasting archive

Technical goals

- scan a big statue → 5 meters
- capture chisel marks → 1/4 mm
- capture reflectance → 1/4 mm

5 meters, 1/4 mm → 20,000:1 → $20{,}000^2$ 1 billion

Why capture chisel marks?

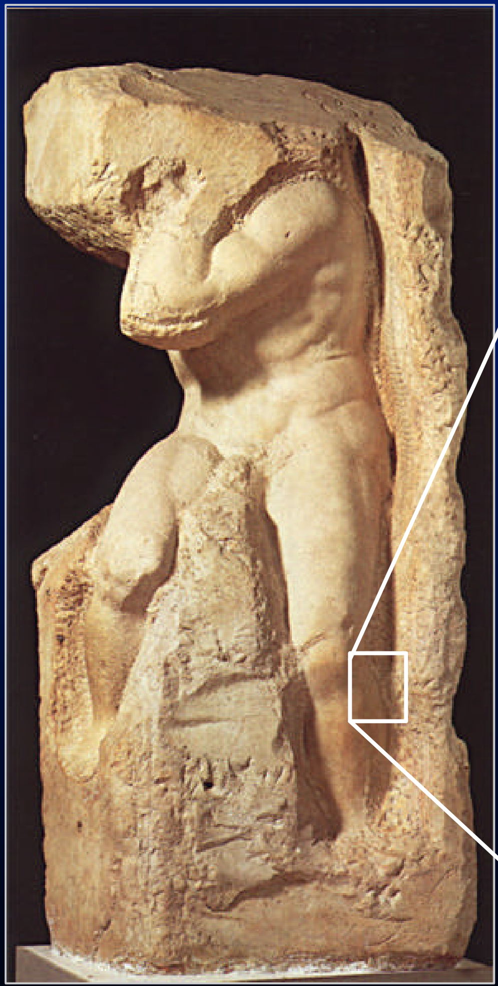

Atlas (Accademia)

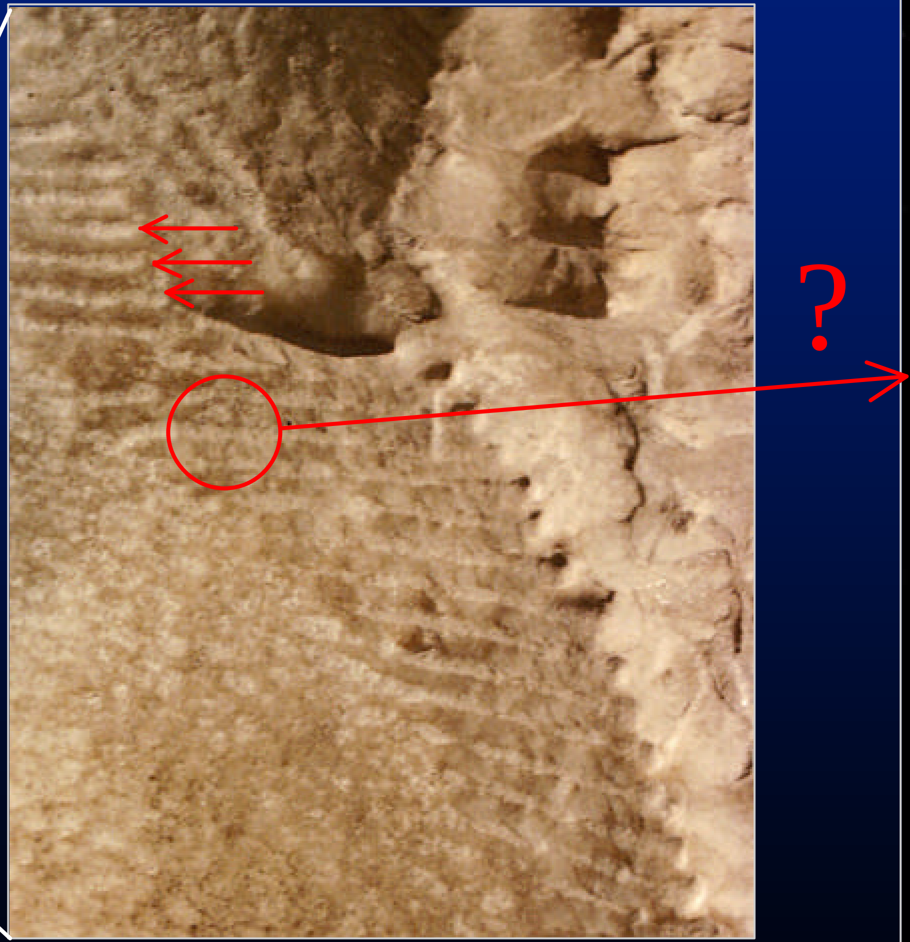

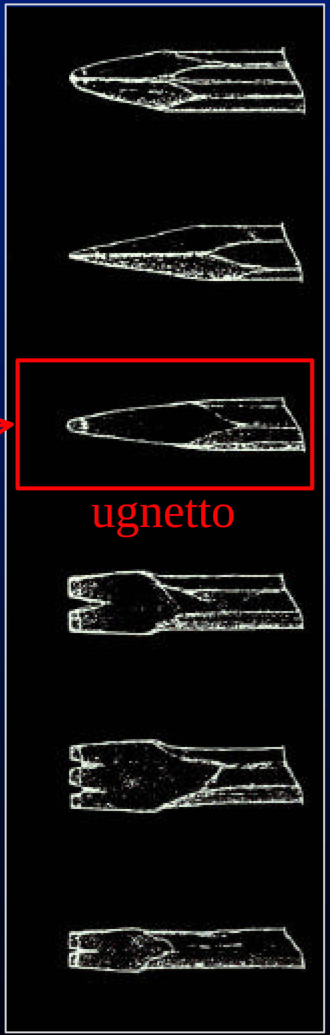

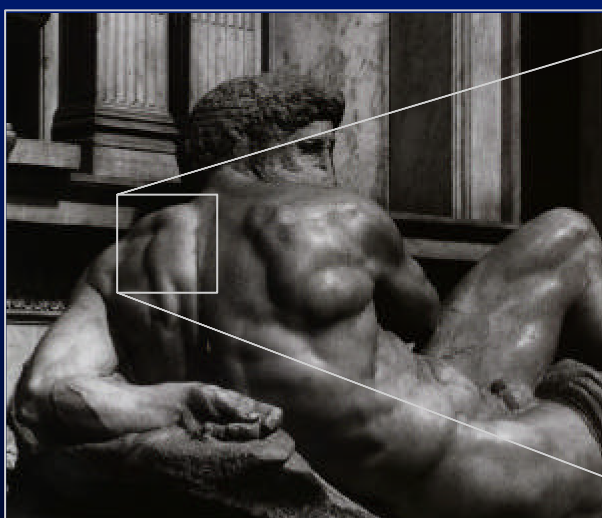

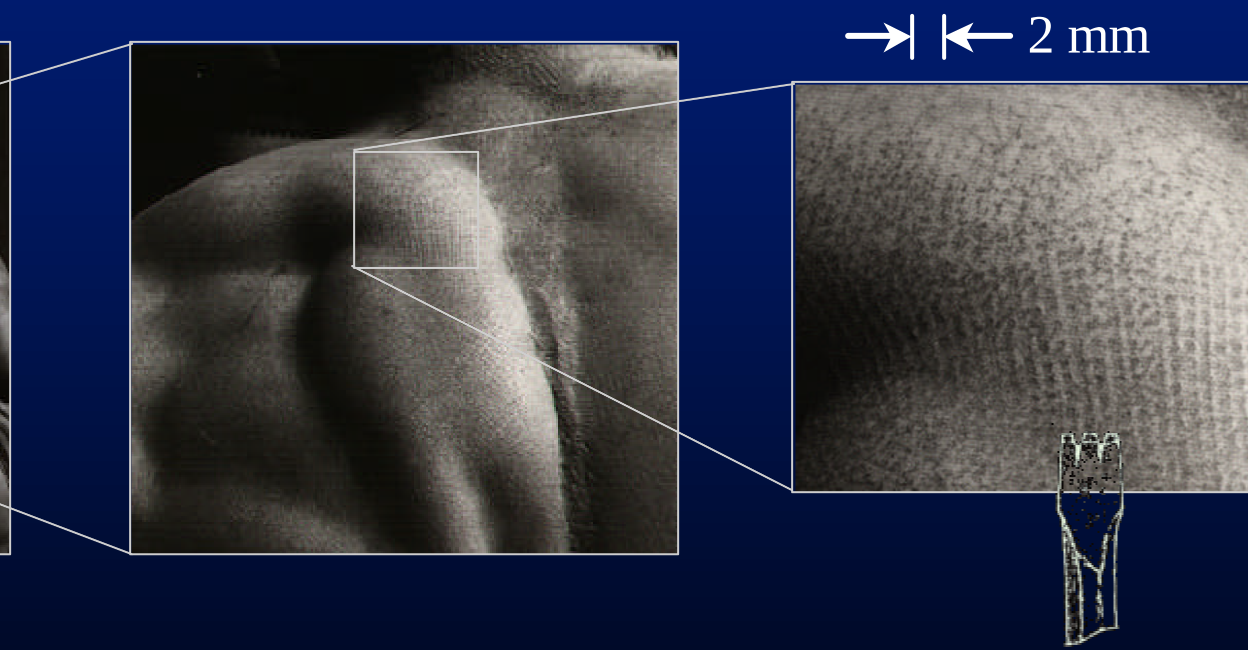

Day (Medici Chapel)

Outline of talk

- scanner design
- processing pipeline
- scanning the David
- problems faced and lessons learned
- some side projects
- uses for our models
- an archeological jigsaw puzzle

Scanner design

4 motorized axes

truss extensions
for tall statues

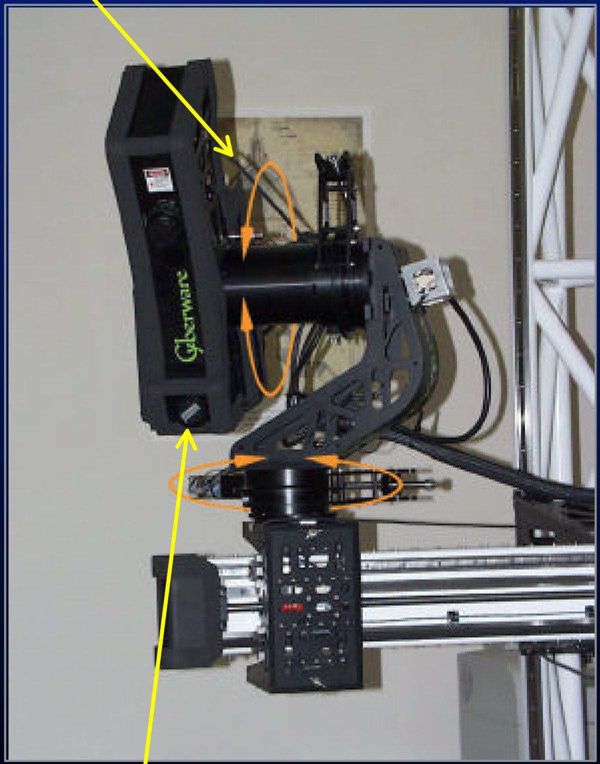

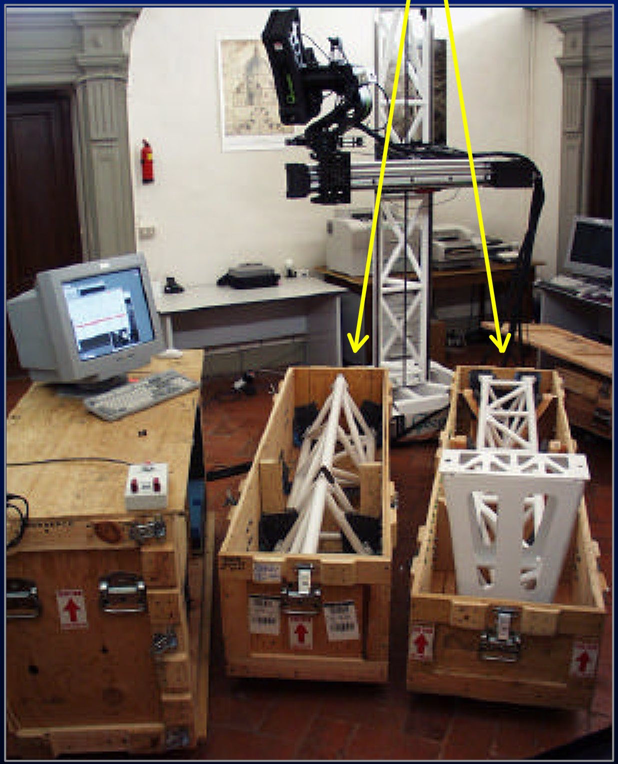

laser, range camera,
white light, and color camera

Scanning St. Matthew

working in
the museum

scanning
geometry

scanning
color

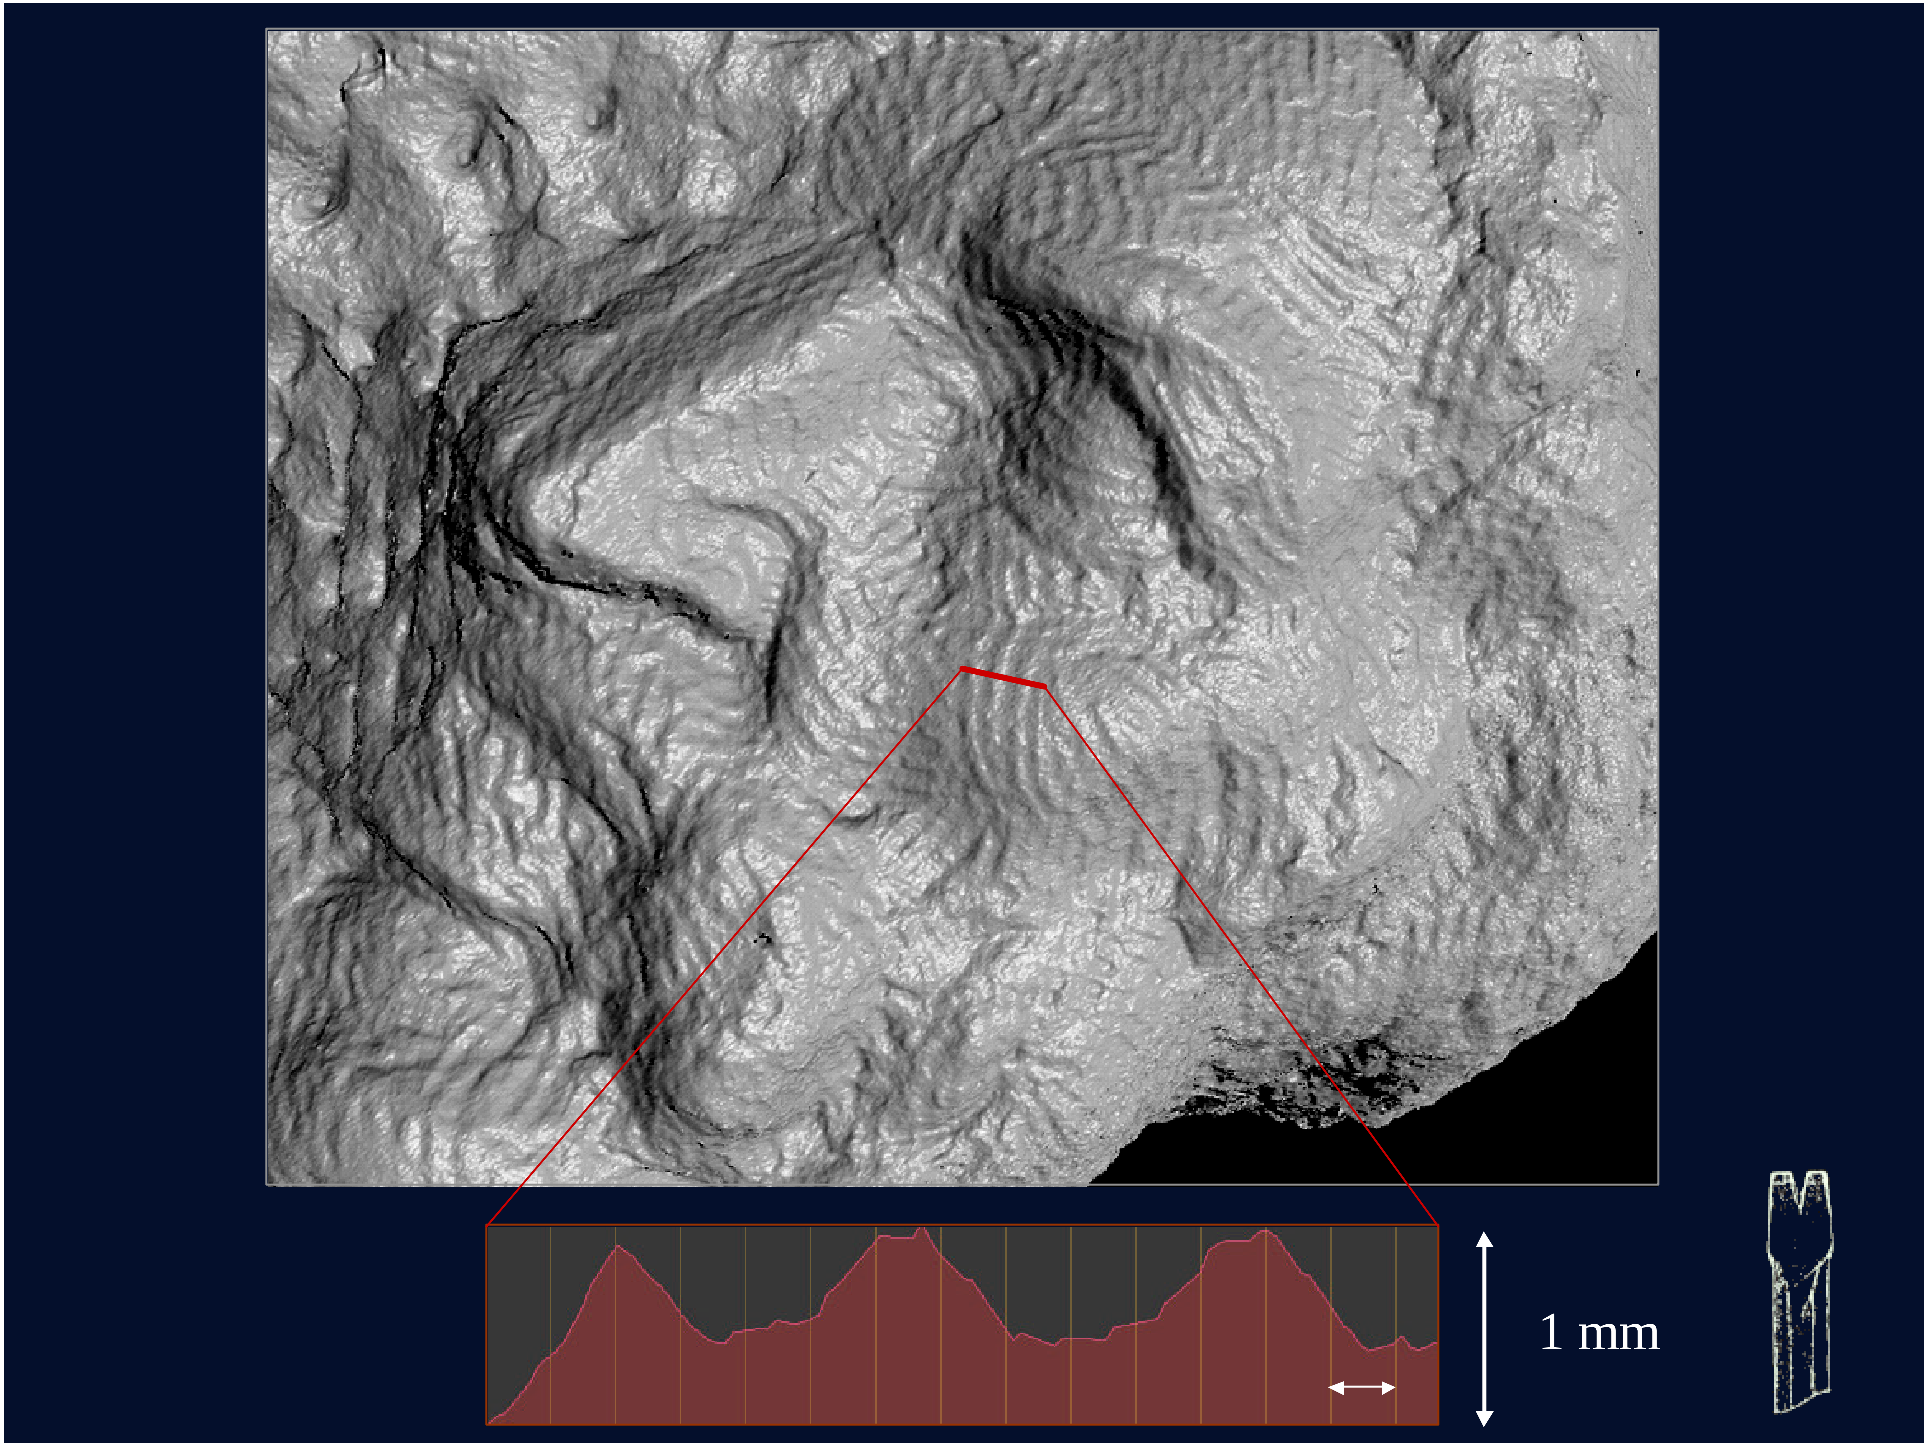
1 mm

How optically cooperative is marble?

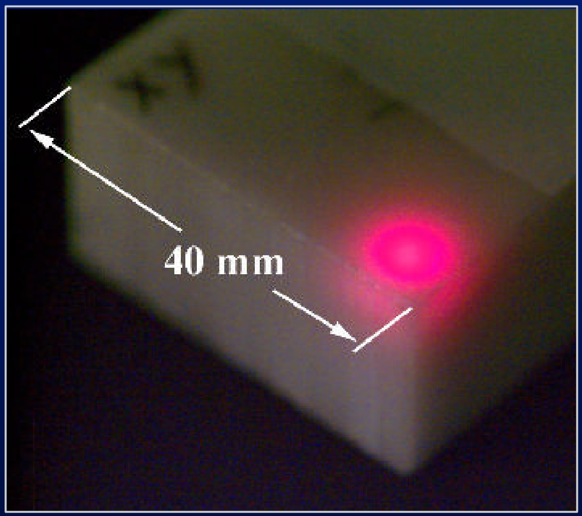

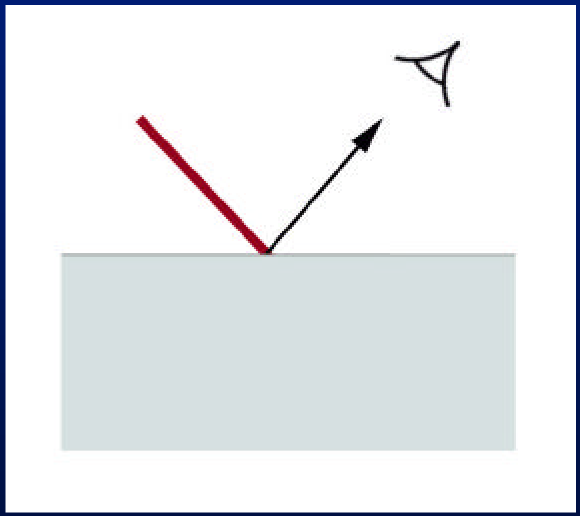

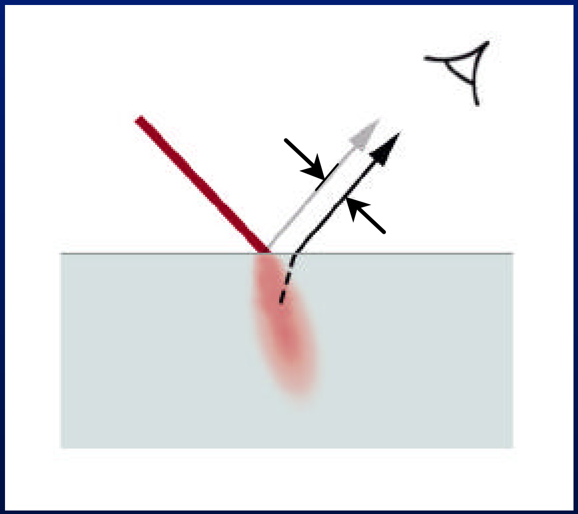

- systematic bias of 40 microns
- noise of 150 – 250 microns
  - worse at oblique angles of incidence
  - worse for polished statues

Scanning a large object

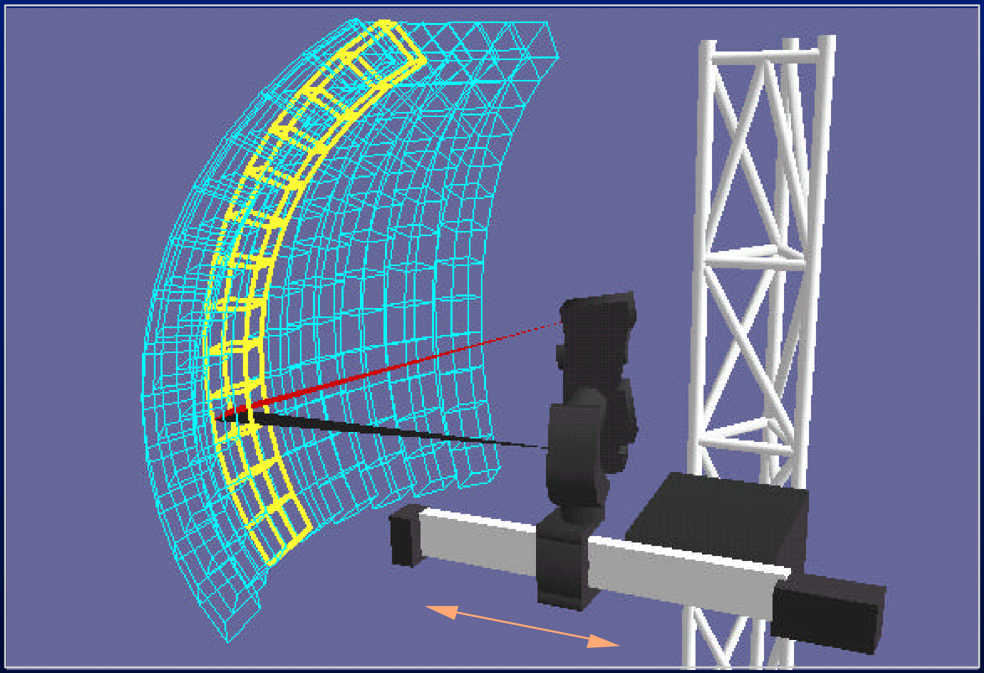

- calibrated motions
  - pitch (yellow)
  - pan (blue)
  - horizontal translation (orange)
- uncalibrated motions
  - vertical translation
  - remounting the scan head
  - moving the entire gantry

Our scan of St. Matthew

- 104 scans
- 800 million polygons
- 4,000 color images
- 15 gigabytes
- 1 week of scanning

Range processing pipeline

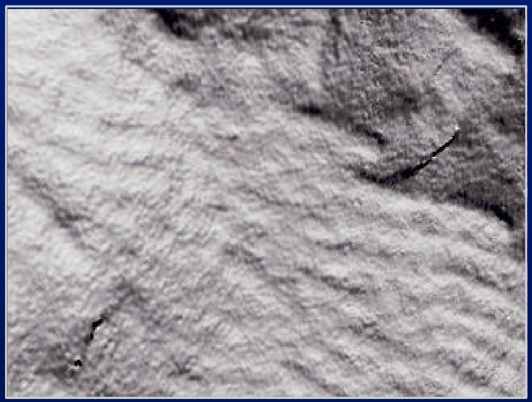

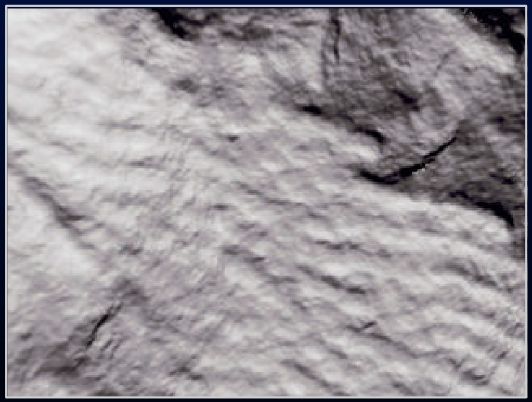

- steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method
- lessons learned
  – should have tracked the gantry location
  – ICP is unstable on smooth surfaces

Color processing pipeline

- steps
  1. compensate for ambient illumination
  2. discard shadowed or specular pixels
  3. map onto vertices – one color per vertex
  4. correct for irradiance → diffuse reflectance
- limitations
  – ignored interreflections
  – ignored subsurface scattering
  – treated diffuse as Lambertian
  – used aggregate surface normals

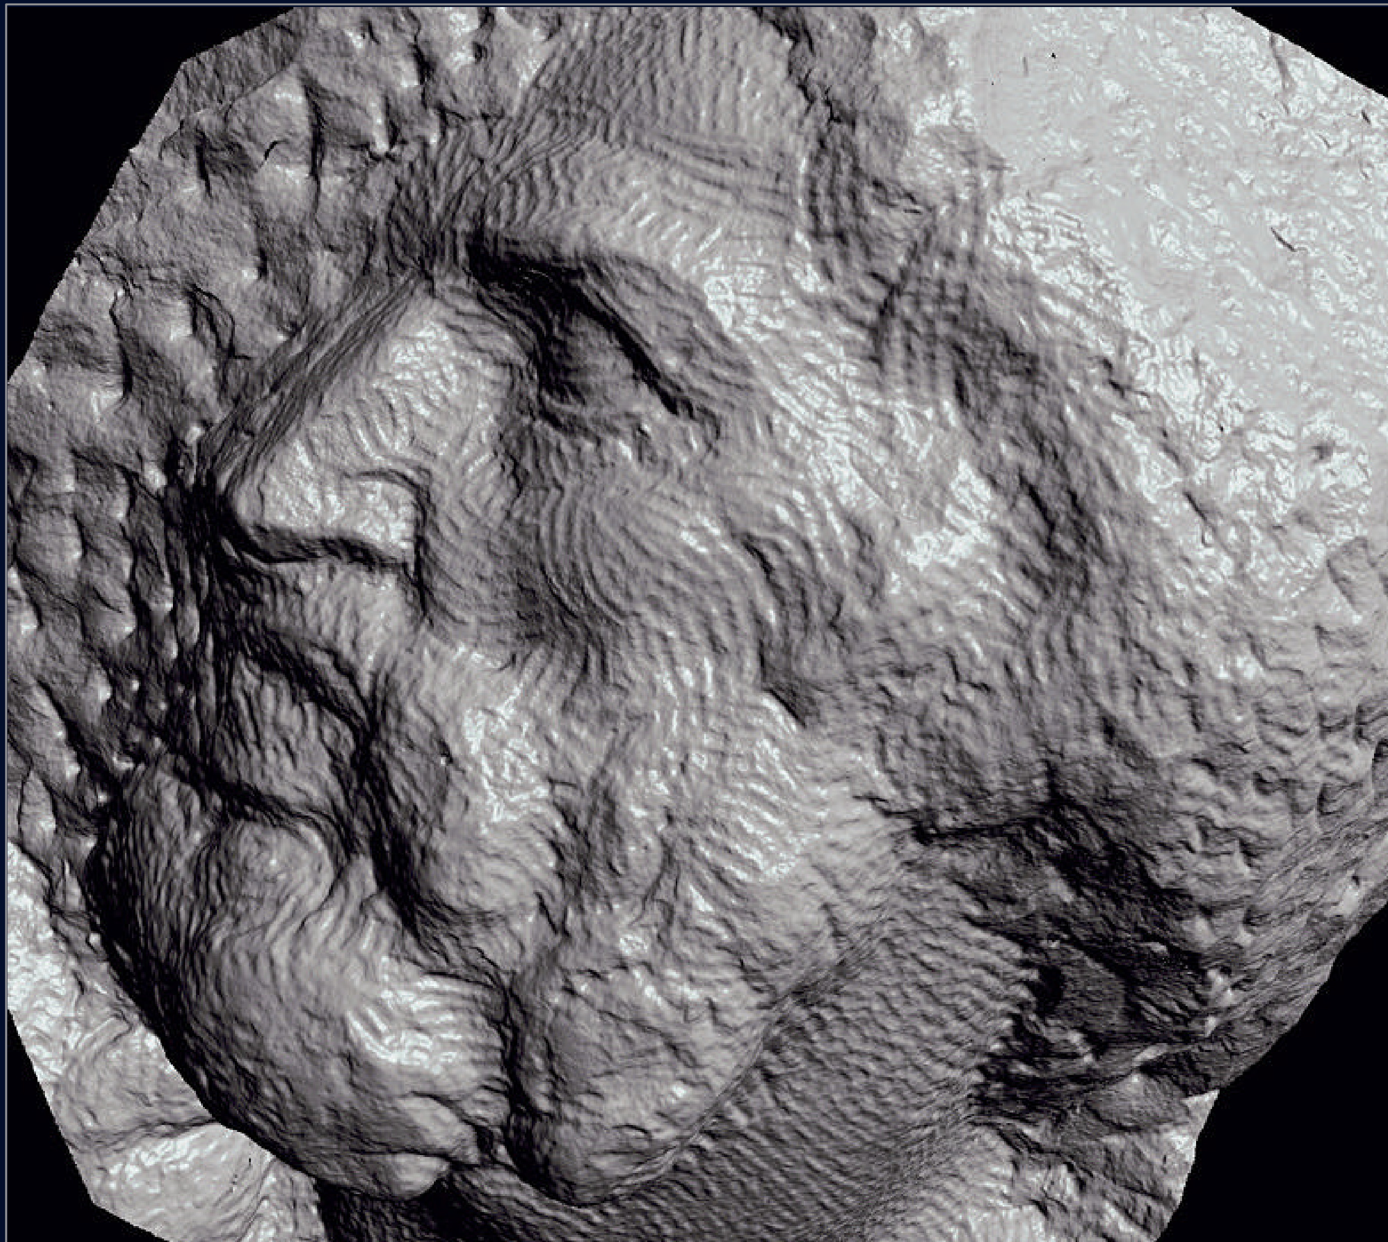

artificial surface reflectance

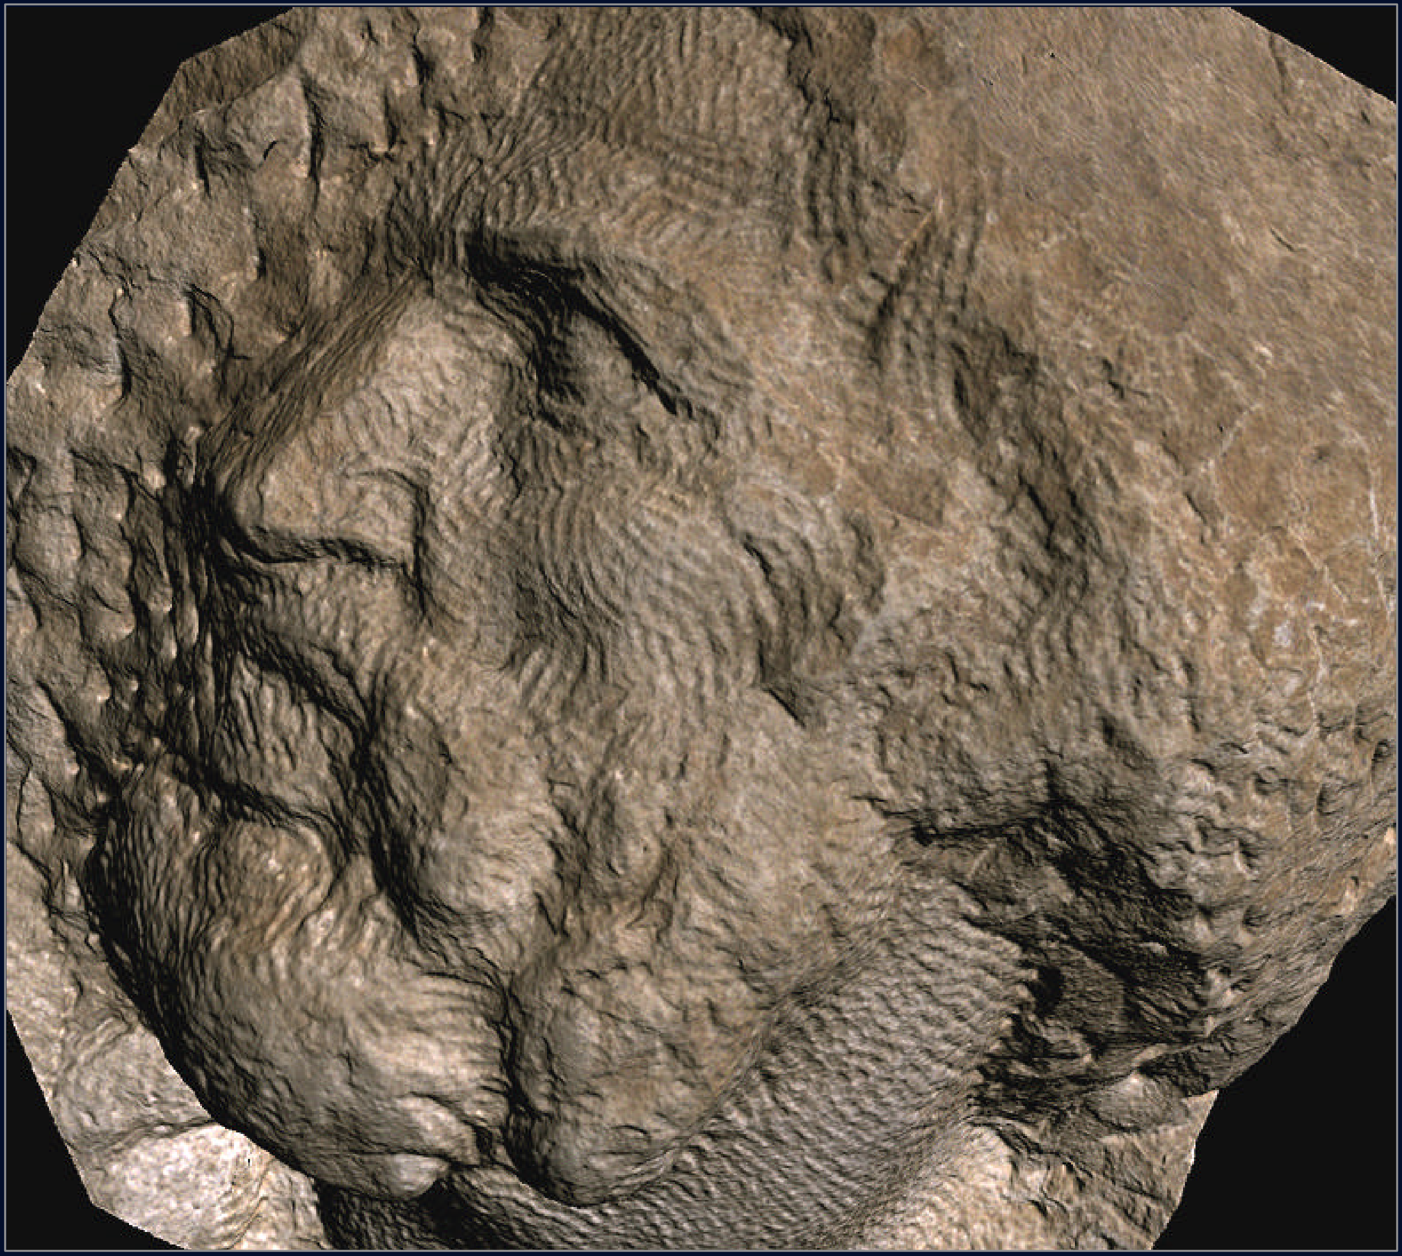

estimated diffuse reflectance

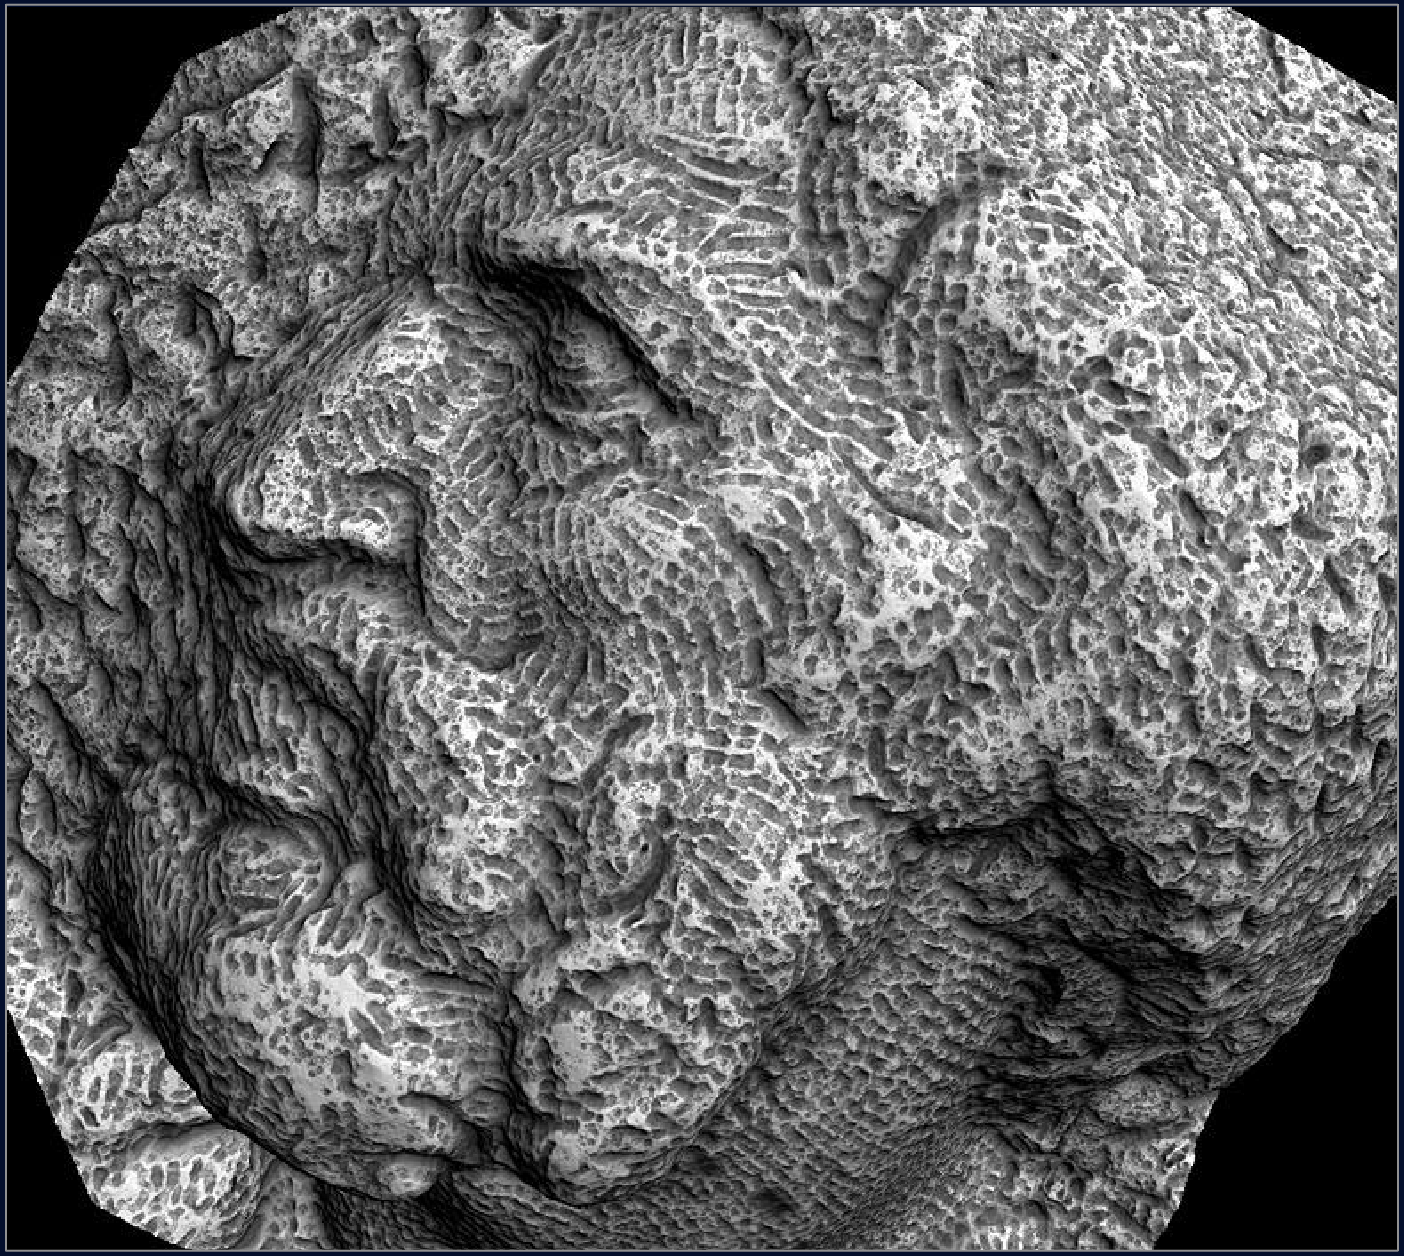

accessibility shading

Scanning the David

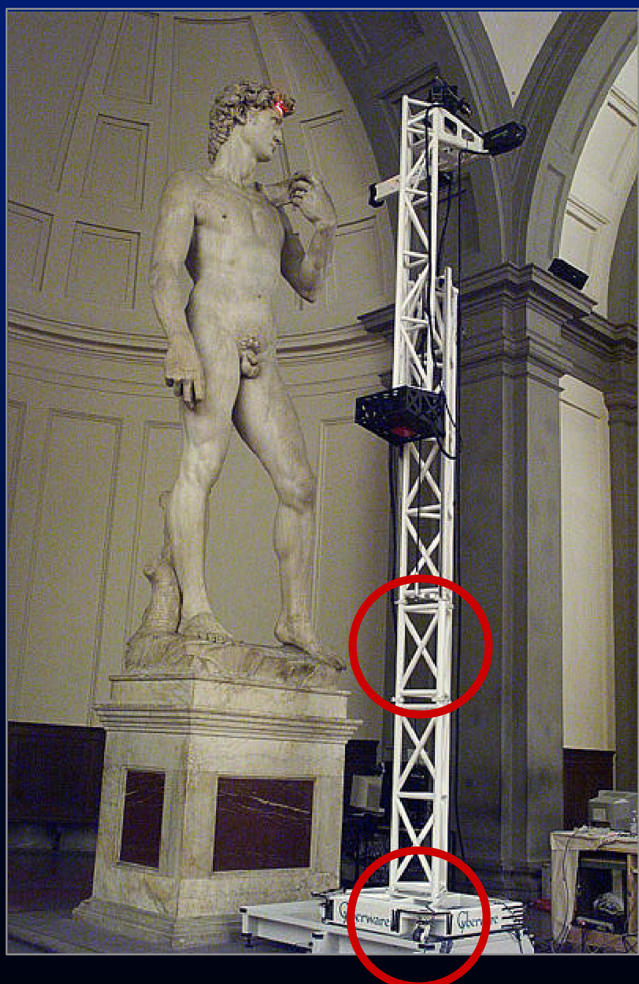

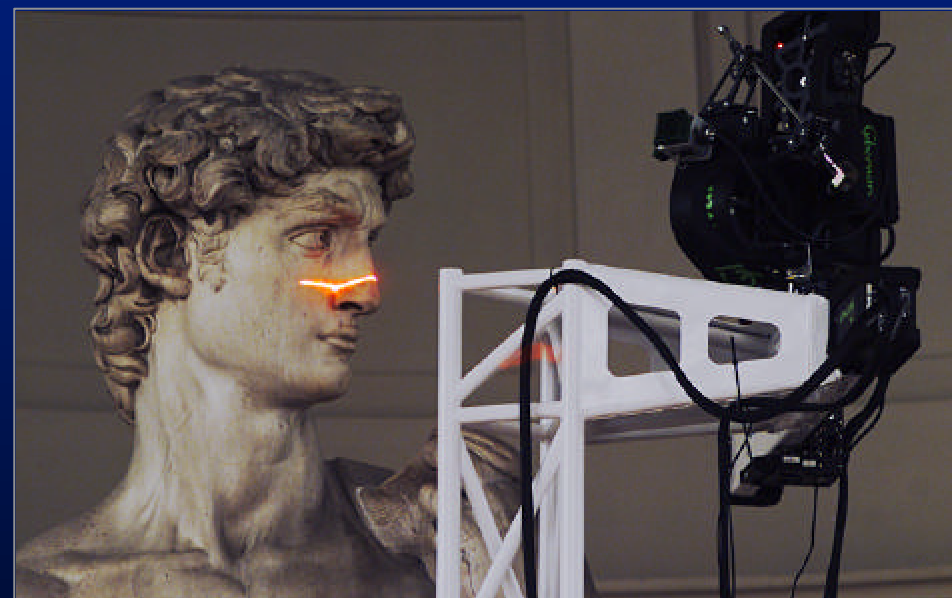

height of gantry: 7.5 meters

weight of gantry: 800 kilograms

Statistics about the scan

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

Hard problem #1: view planning

- procedure
  - manually set scanning limits
  - run scanning script

```
for horizontal = min to max by 12 cm
   for pan = min to max by 4.3 °
      for tilt = min to max continuously
         perform fast pre-scan (5 ° /sec)
         search pre-scan for range data
      for tilt = all occupied intervals
         perform slow scan (0.5 ° /sec)
   on every other horizontal position,
      for pan = min to max by 7 °
         for tilt = min to max by 7 °
            take photographs without spotlight
      warm up spotlight
      for pan = min to max by 7 °
         for tilt = min to max by 7 °
            take photographs with spotlight
```

- lessons learned
  - need automatic view planning – especially in the endgame
  - 50% of time on first 90%, 50% on next 9%, ignore last 1%

Hard problem #2: accurate scanning in the field

- error budget
  - 0.25mm of position, 0.013° of orientation
- design challenges
  - minimize deflection and vibration during motions
  - maximize repeatability when remounting
- lessons learned
  - motions were sufficiently accurate and repeatable
  - remounting was not sufficiently repeatable
  - used ICP to circumvent poor repeatability

Head of Michelangelo's David

photograph

1.0 mm computer model

The importance of viewpoint

classic 3/4 view

left profile

face-on view

The importance of lighting

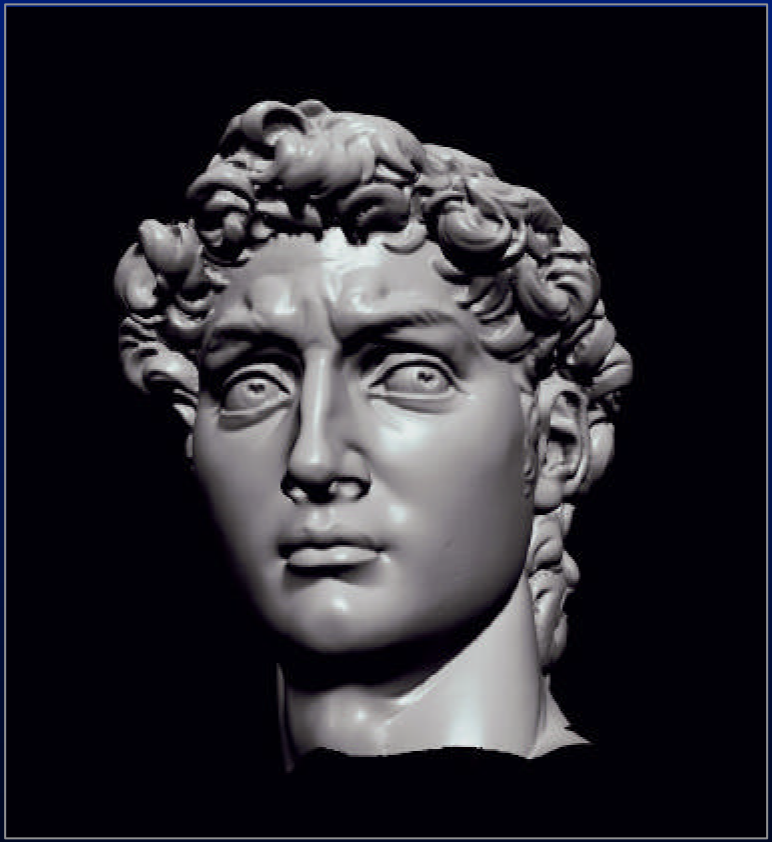

lit from above

lit from below

David’s left eye

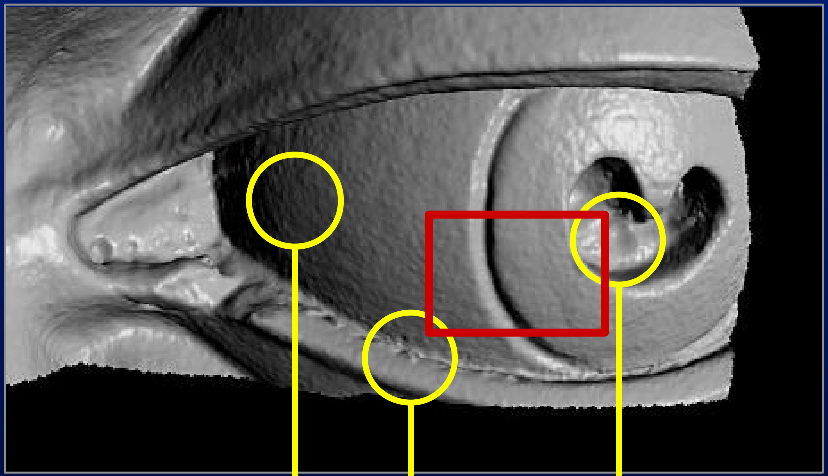

0.25 mm model

photograph

holes from Michelangelo’s drill

artifacts from space carving

noise from laser scatter

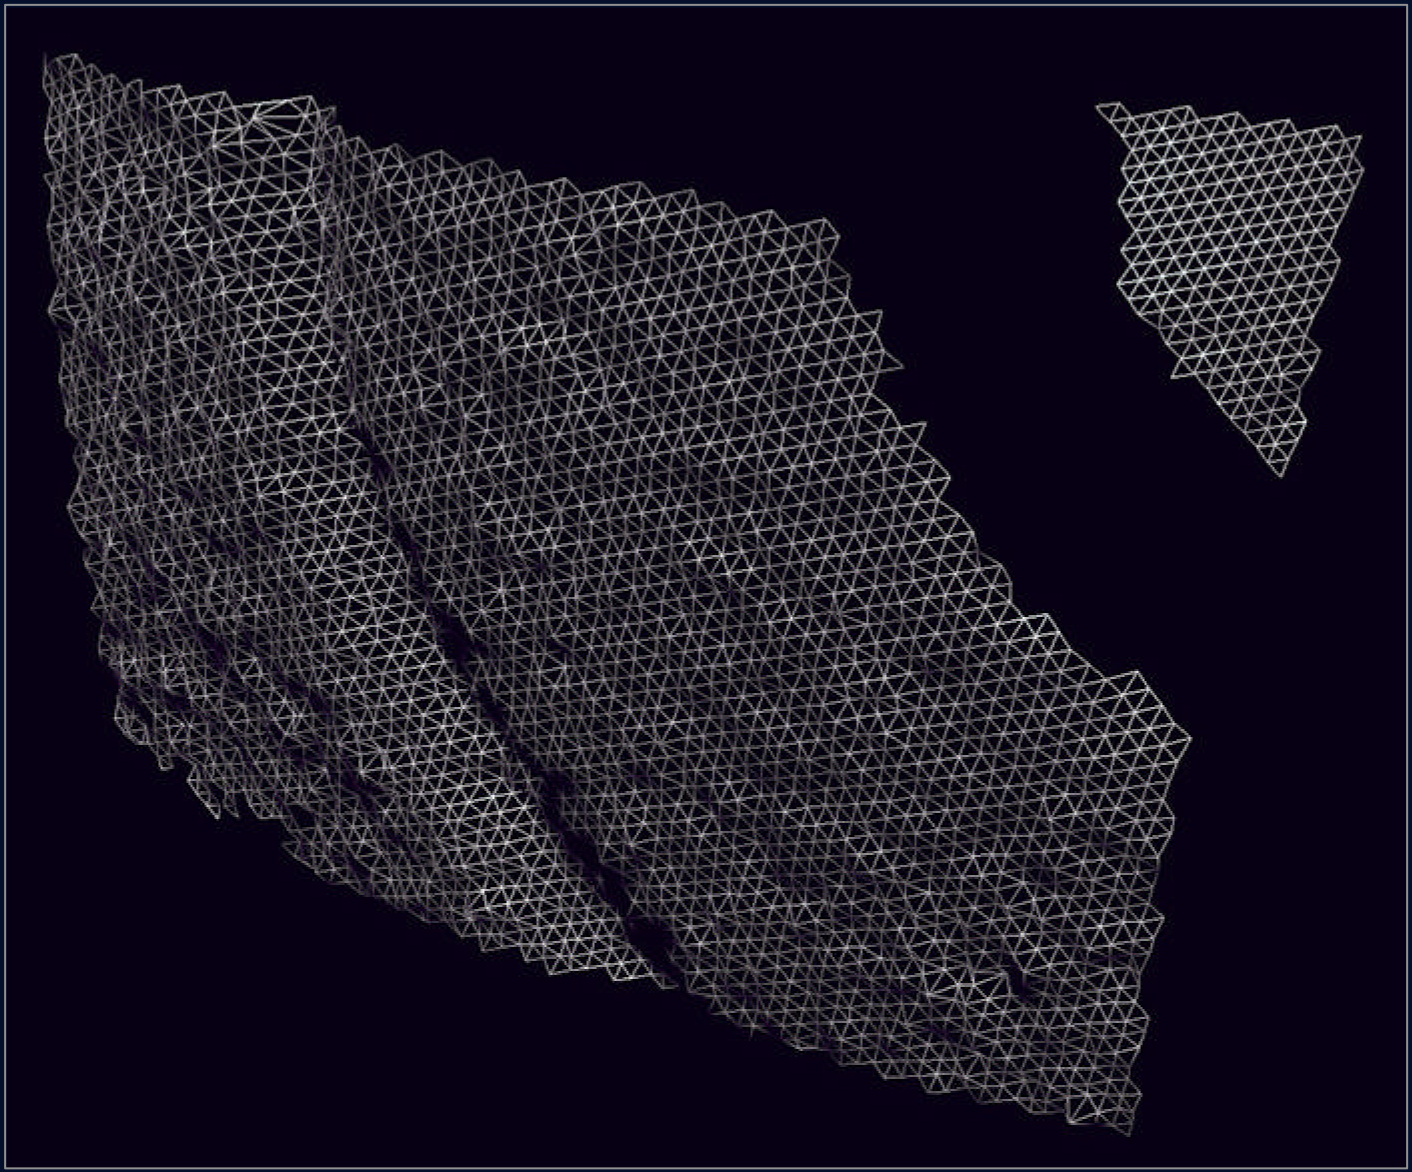

Single scan of David’s cornea

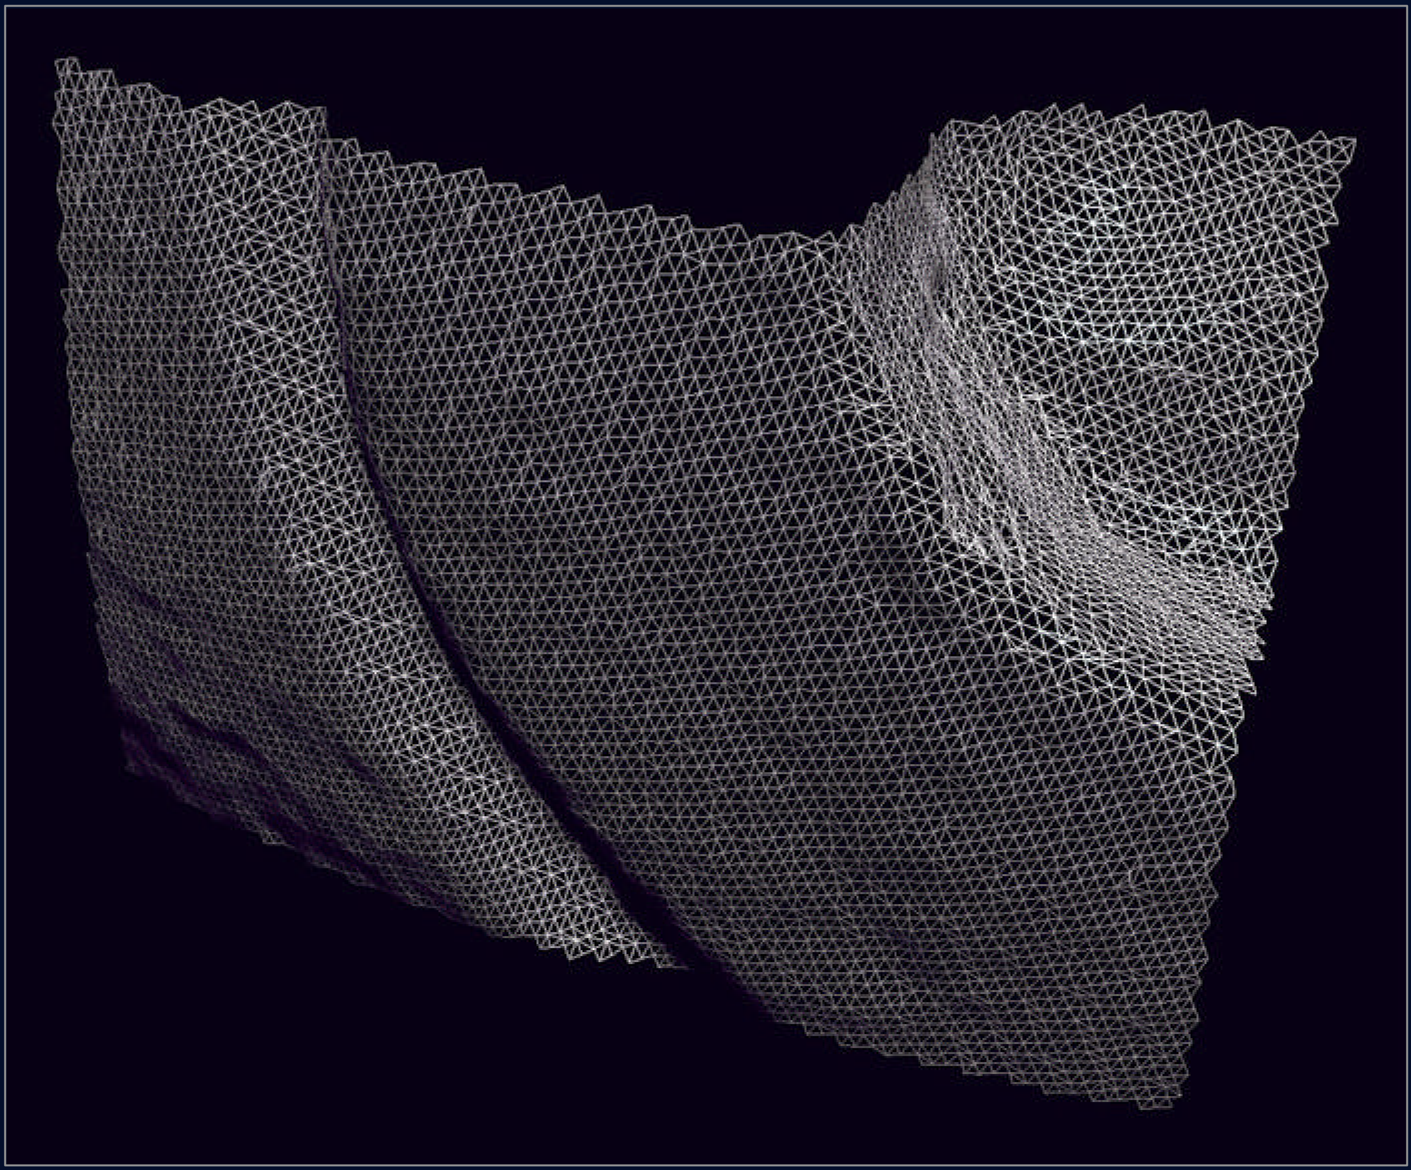

Mesh constructed from several scans

Hard problem #3: insuring safety for the statues

- energy deposition
  - not a problem in our case
- avoiding collisions
  - manual motion controls
  - automatic cutoff switches
  - one person serves as spotter
  - avoid time pressure
  - get enough sleep
- surviving collisions
  - pad the scan head

Hard problem #4: handling large datasets

- range images instead of polygon meshes
  - z(u,v)
  - yields 18:1 lossless compression
  - multiresolution using (range) image pyramid
- multiresolution viewer for polygon meshes
  - 2 billion polygons
  - immediate launching
  - real-time frame rate when moving
  - progressive refinement when idle
  - compact representation
  - fast pre-processing

The Qsplat viewer

- hierarchy of bounding spheres with position, radius, normal vector, normal cone, color

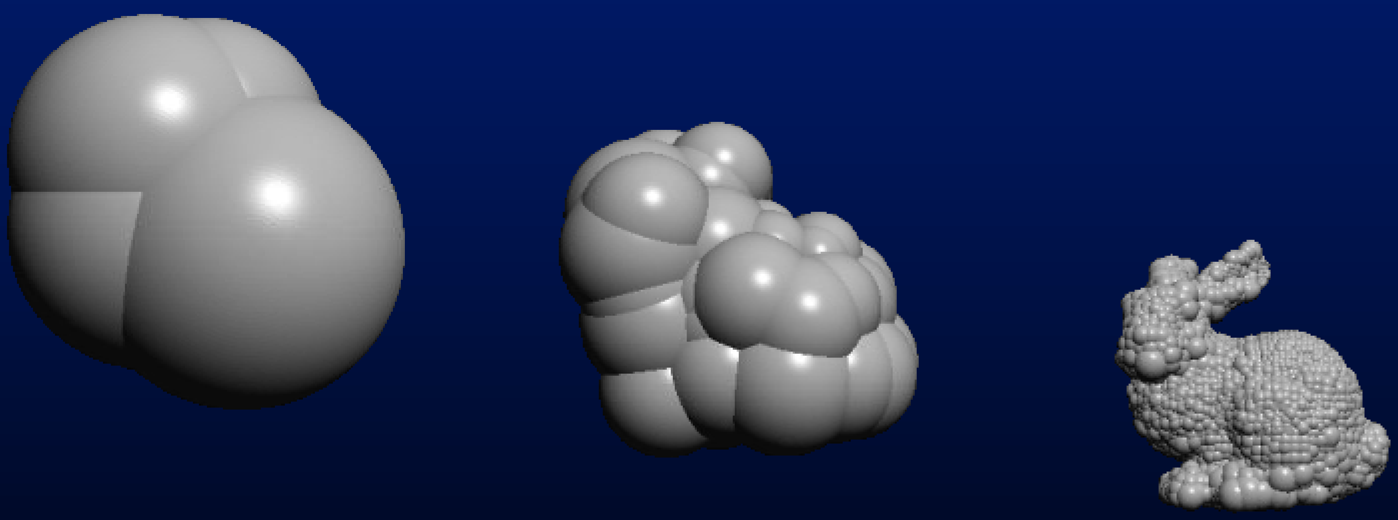

- traversed recursively subject to time limit
- spheres displayed as splats

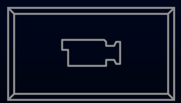

Side project #1: ultraviolet imaging

under white light

under ultraviolet light

Side project #2: architectural scanning

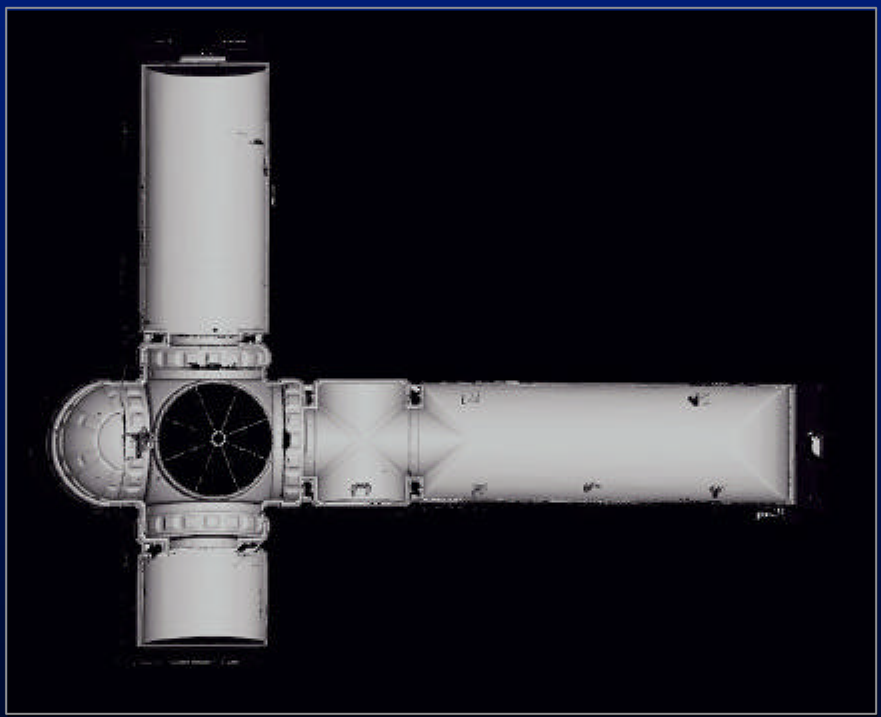

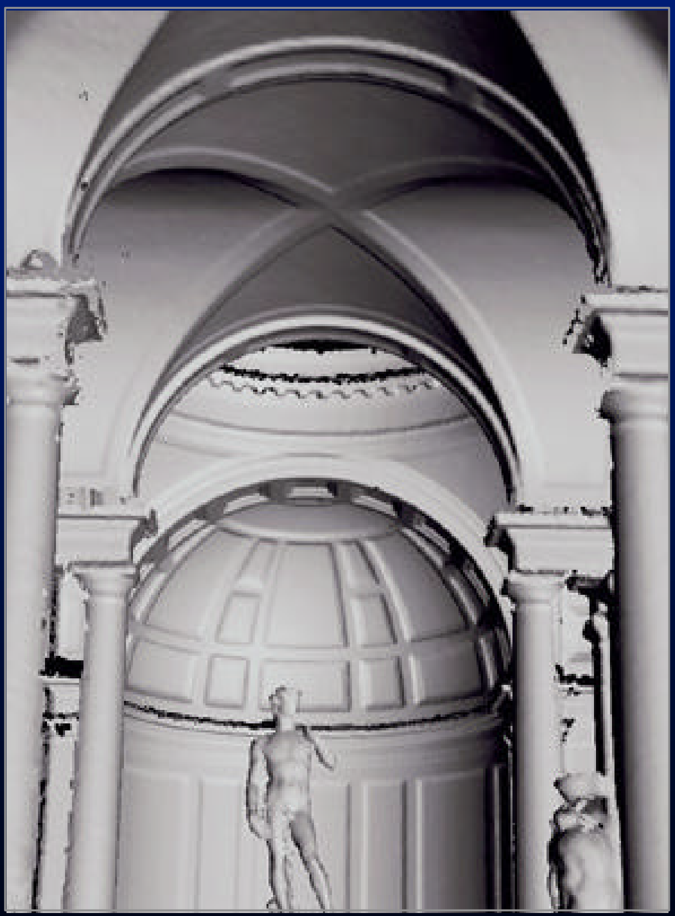

- Galleria dell'Accademia
- Cyra time-of-flight scanner
- 4mm model

Side project #3: light field acquisition

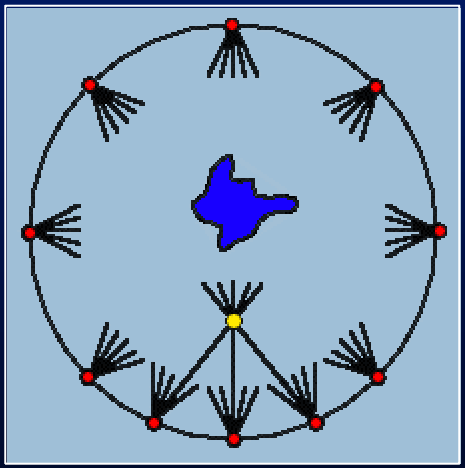

- a form of image-based rendering (IBR)
  - create new views by rebinning old views
- advantages
  - doesn't need a 3D model
  - less computation than rendering a model
  - rendering cost independent of scene complexity
- disadvantages
  - fixed lighting
  - static scene geometry
  - must stay outside convex hull of object

A light field is an array of images

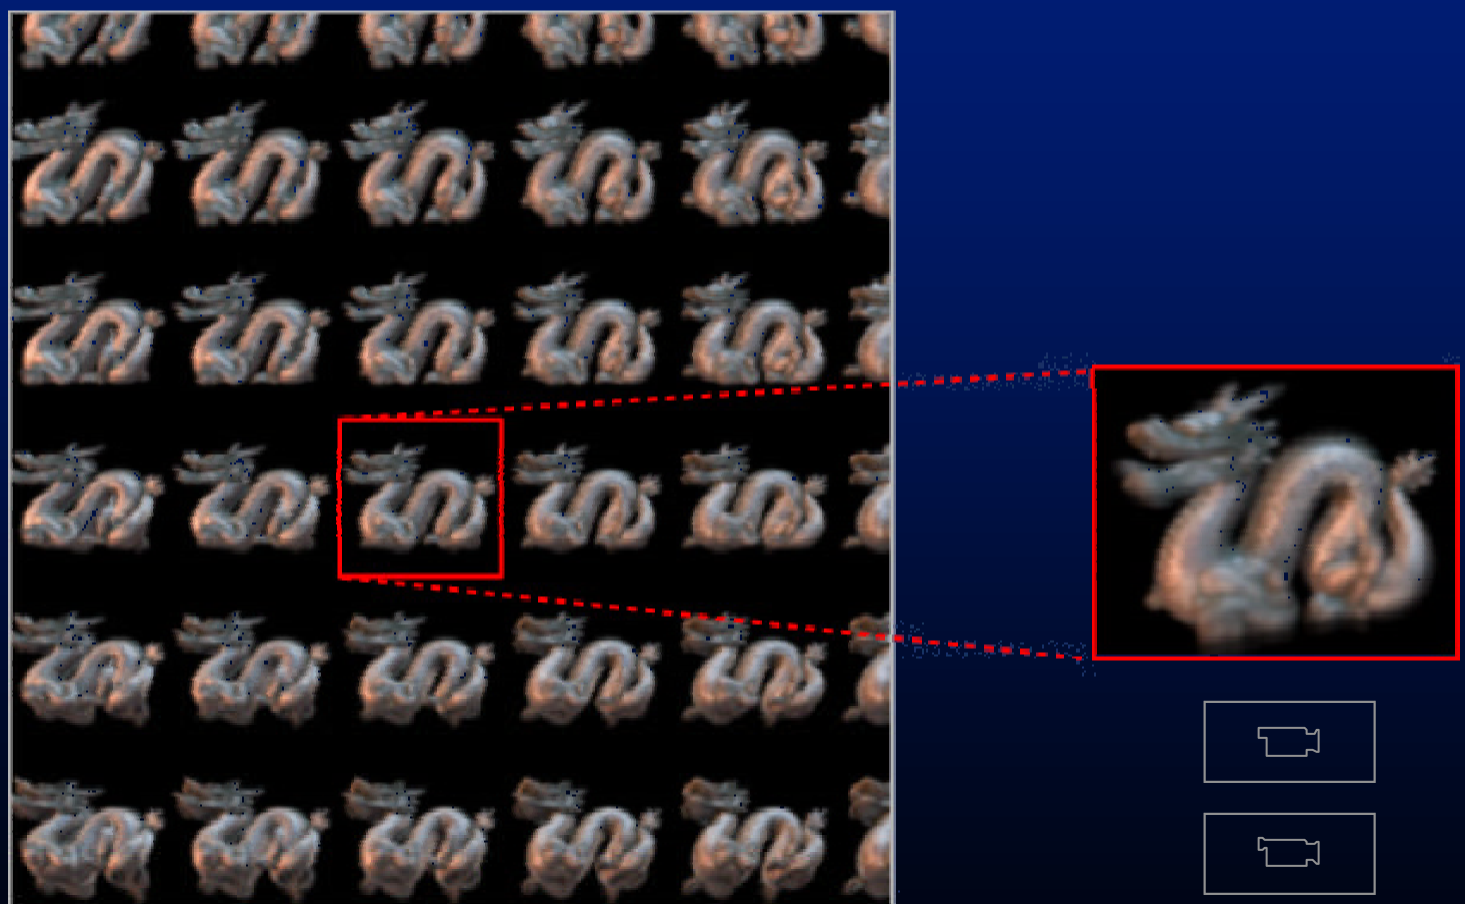

An optically complex statue

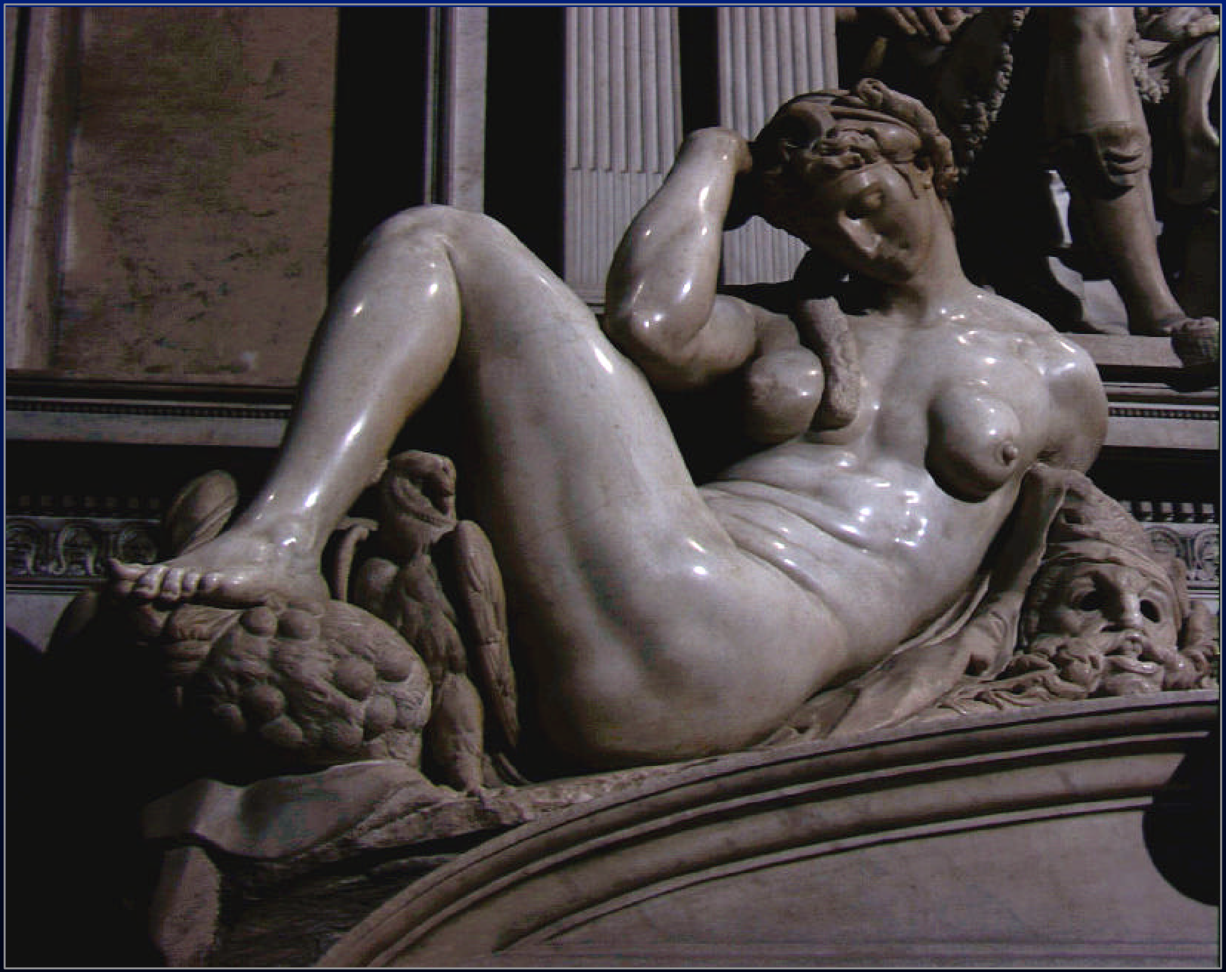

Night (Medici Chapel)

Acquiring the light field

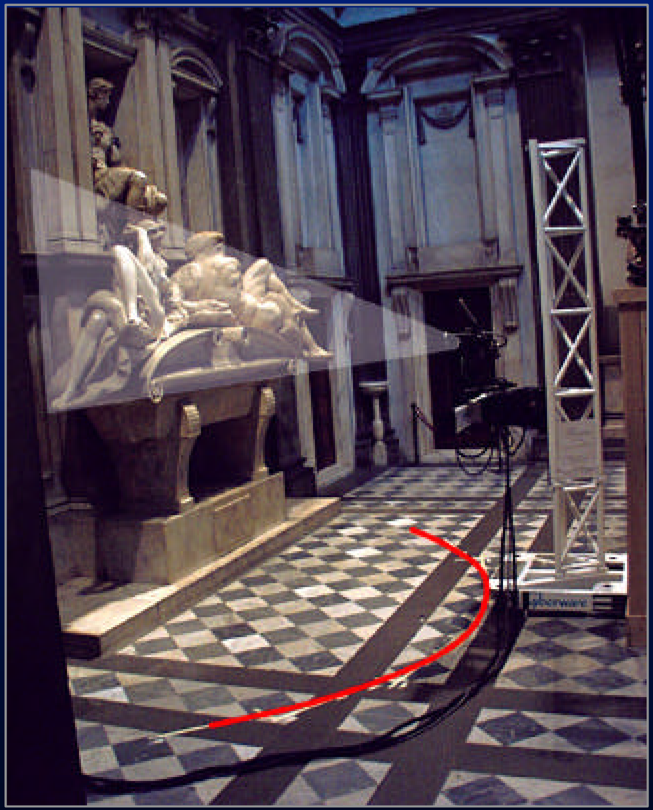

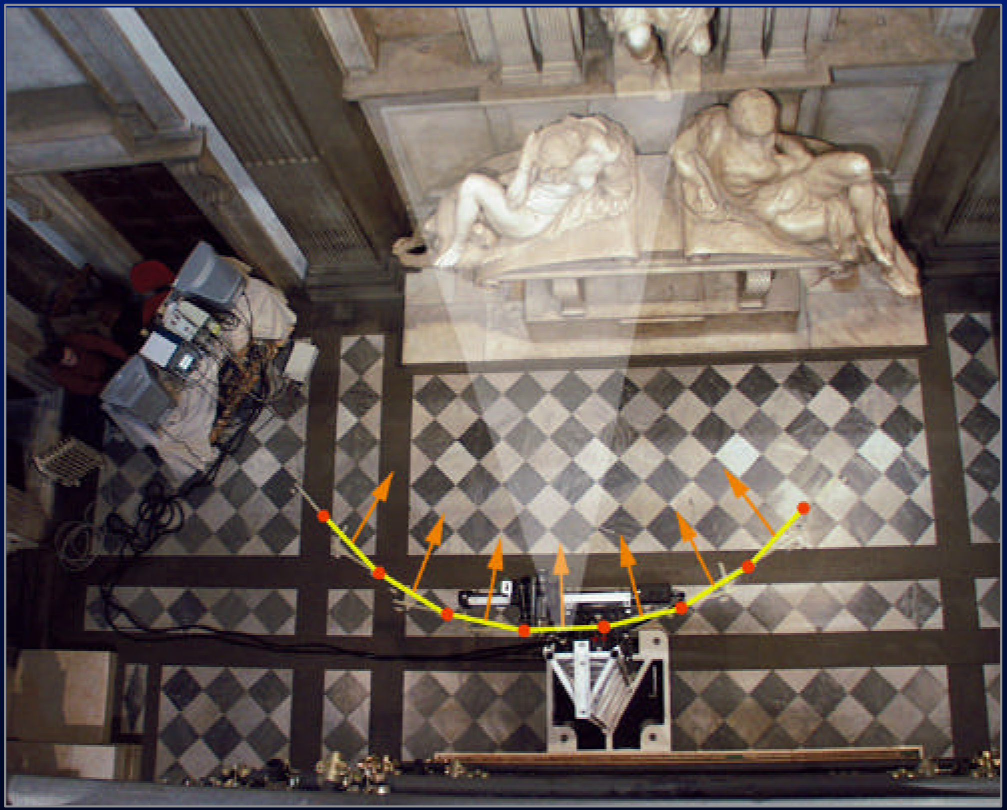

- natural eye level
- artificial illumination

7 light slabs,
each 70cm x 70cm

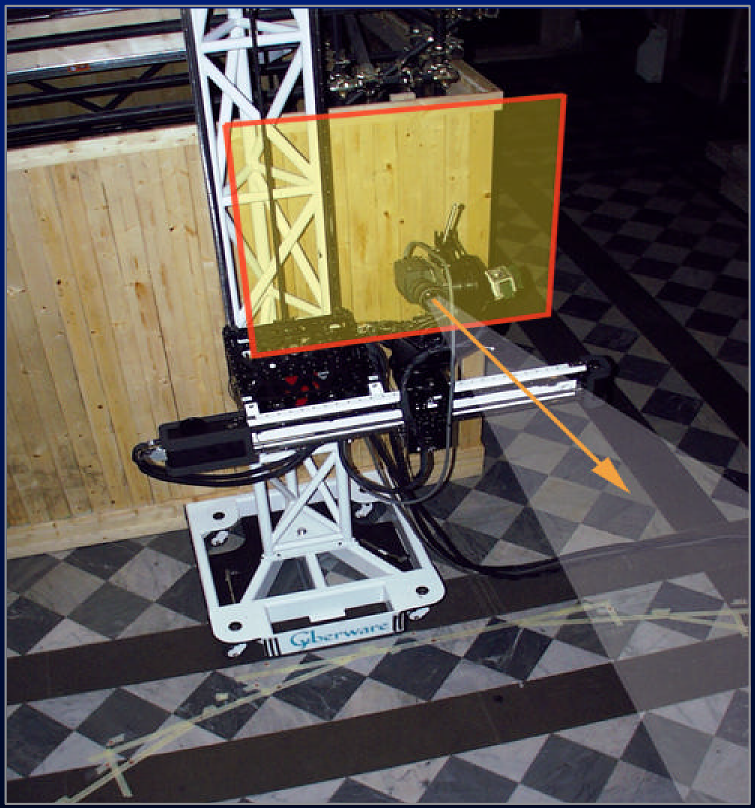

each slab contained 56 x 56 images spaced 12.5mm apart

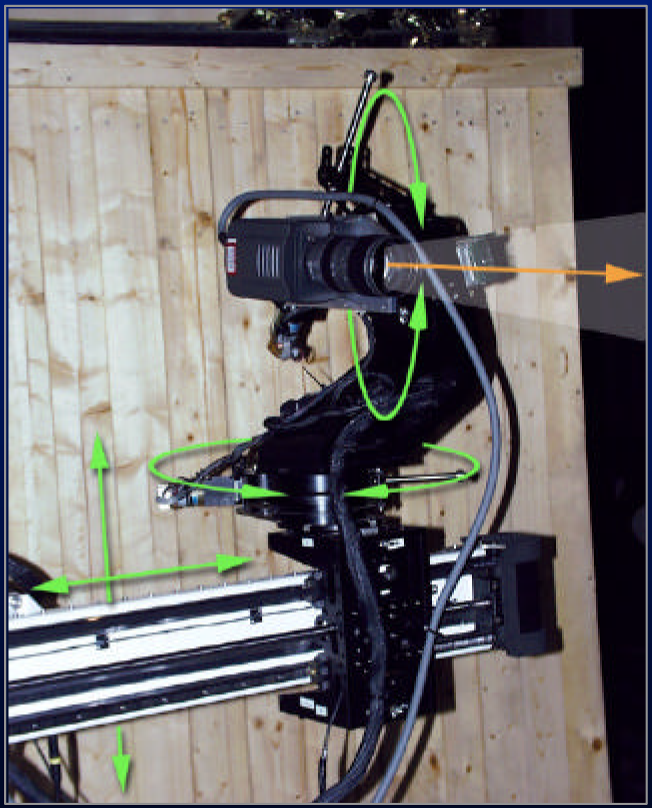

the camera was always aimed at the center of the statue

Statistics about the light field

- 392 x 56 images
- 1300 x 1000 pixels each
- 96 gigabytes (uncompressed)
- 35 hours of shooting (over 4 nights)
- also acquired a 0.29 mm 3D model of statue

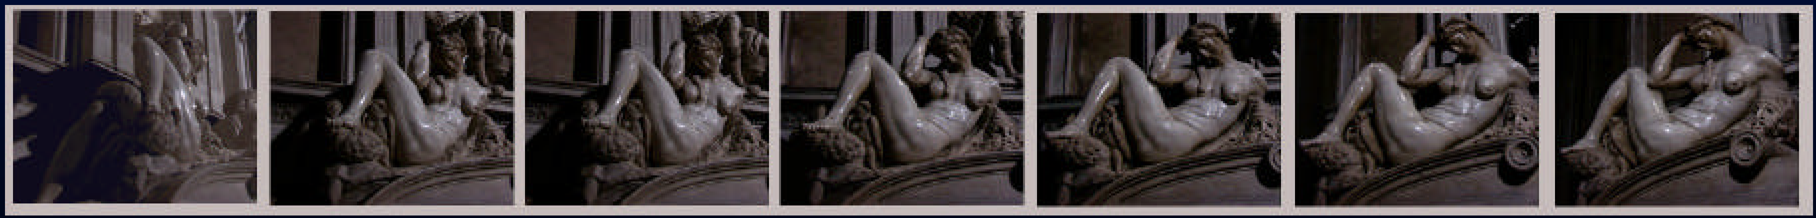

Some obvious uses for these models

- unique views of the statues
- permanent archive
- virtual museums
- physical replicas
- 3D stock photography

Michelangelo's Pieta

handmade replica

Some not-so-obvious uses

- restoration record
- geometric calculations
- projection of images onto statues

Side project #4: an archeological jigsaw puzzle

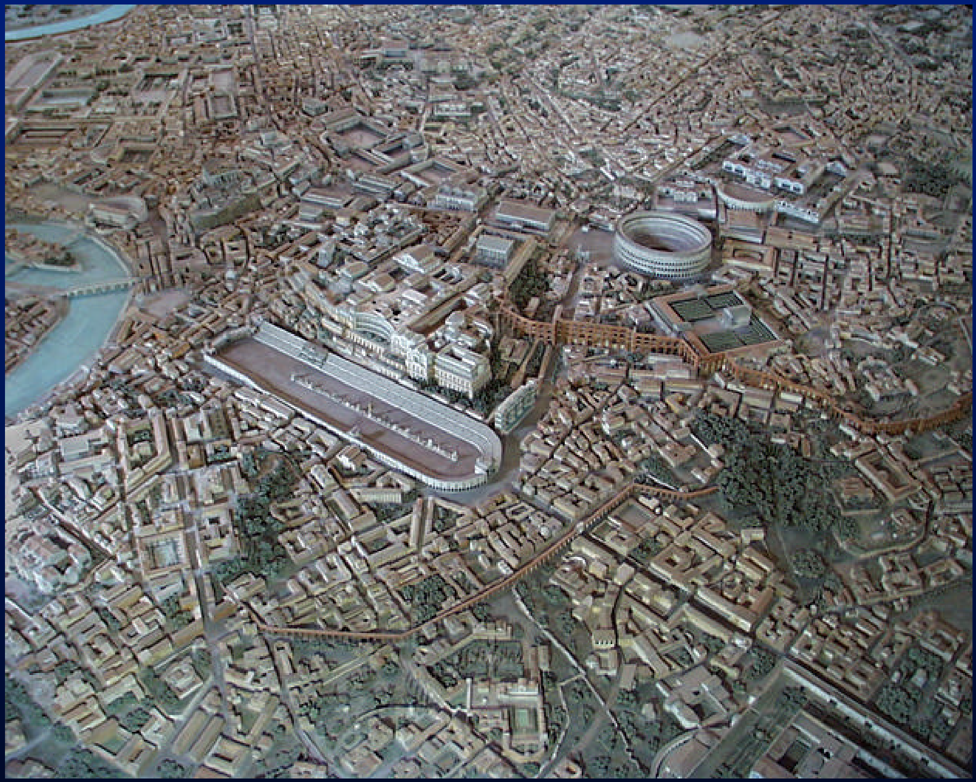

- Il Plastico – a model of ancient Rome
- made in the 1930's
- measures 60 feet on a side

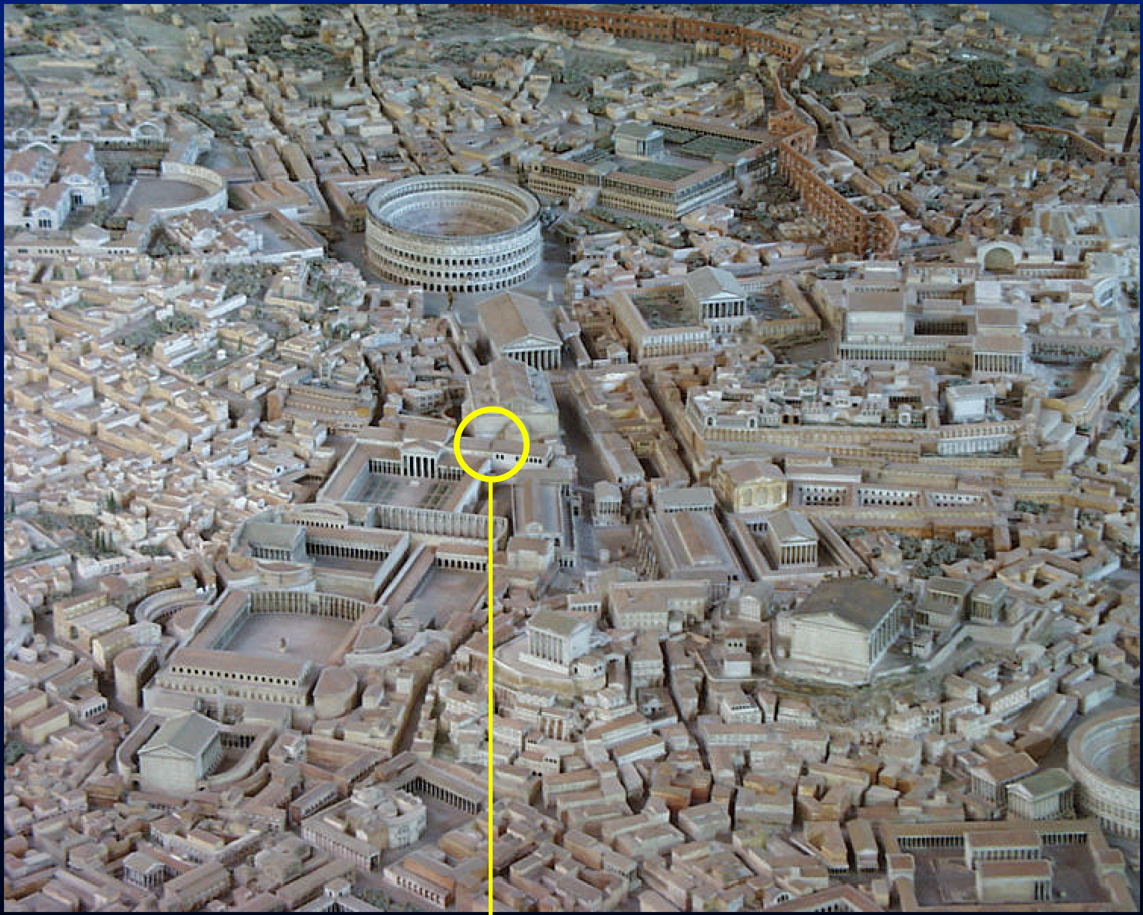

the Roman census bureau

The Forma Urbis Romae: a map of ancient Rome

- carved circa 200 A.D.
- 60 wide x 45 feet high
- marble, 4 inches thick
- showed the entire city at 1:240
- single most important document about ancient Roman topography

its back wall still exists, and on it was hung...

Fragment #10g

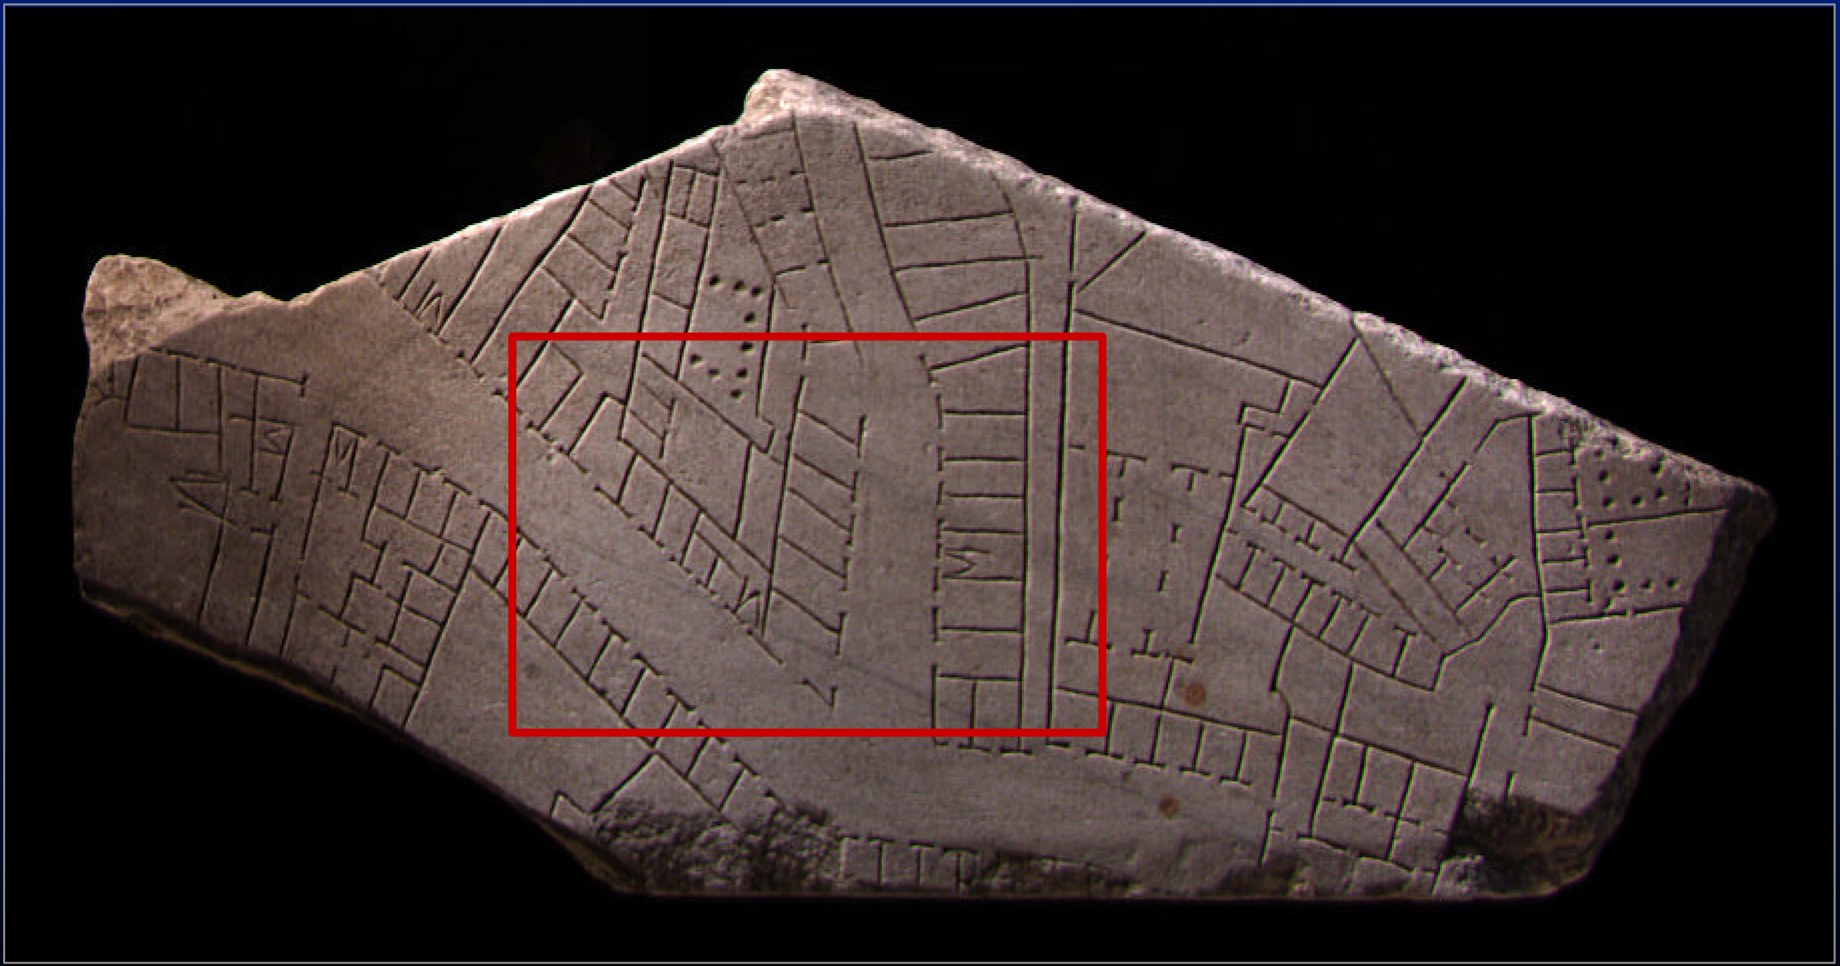

Fragment #10g

interior courtyard with
columned portico

18 cm on map
43 meters on the ground

staircase

room with door

Solving the jigsaw puzzle

- 1,163 fragments
  - 200 identified
  - 500 unidentified
  - 400 unincised
- 15% of map remains
  - but strongly clustered
- available clues
  - fragment shape (2D or 3D)
  - incised patterns
  - marble veining
  - matches to ruins

Scanning the fragments

uncrating...

Scanning the fragments

positioning...

Scanning the fragments

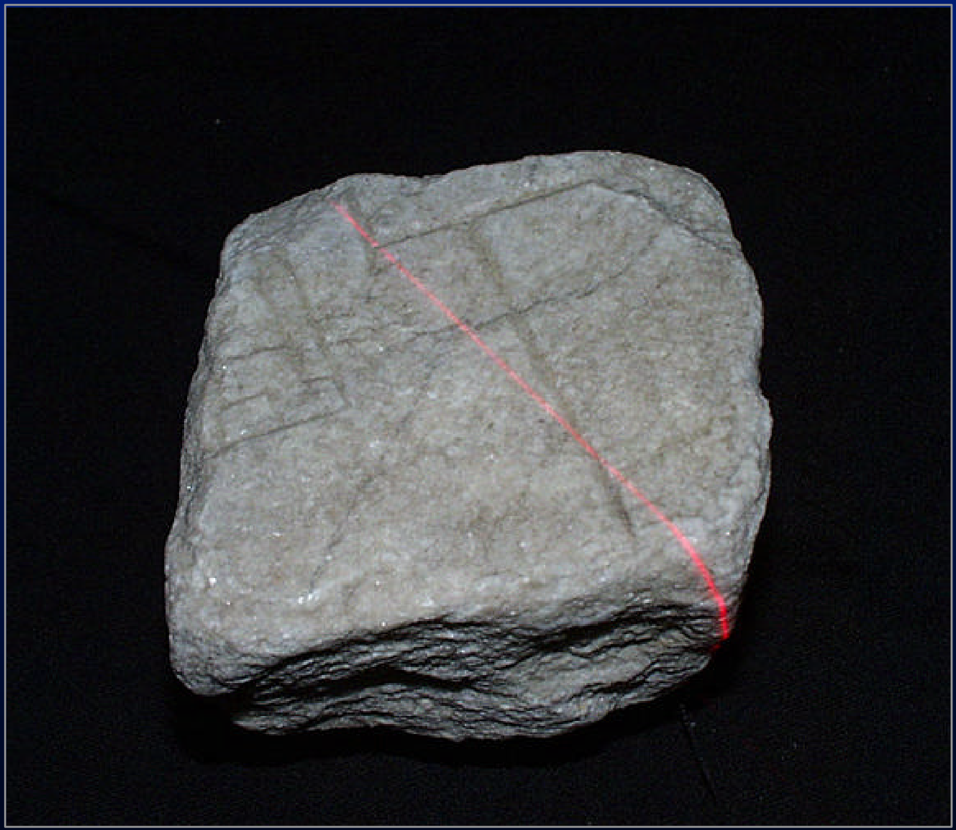

scanning...

Scanning the fragments

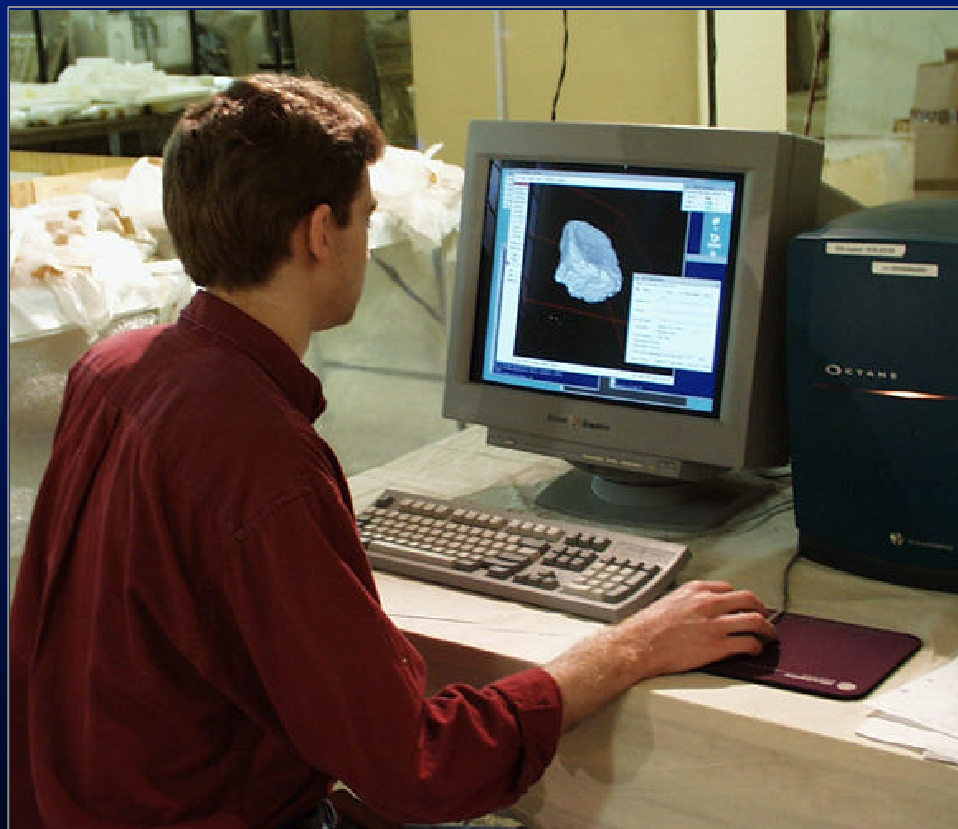

aligning...

Fragment #642

3D model

color photograph

forma urbis romae

Future work

1. hardware
   – scanner design
   – scanning in tight spots
   – tracking scanner position
   – better calibration methodologies
   – scanning uncooperative materials
   – insuring safety for the statues
2. software
   – automated view planning
   – accurate, robust global alignment
   – more sophisticated color processing
   – handling large datasets
   – filling holes

3. uses for these models
   - permanent archive
   - virtual museums
   - physical replicas
   - restoration record
   - geometric calculations
   - projection of images onto statues

4. digital archiving
   - central versus distributed archiving
   - insuring longevity for the archive
   - authenticity, versioning, variants
   - intellectual property rights
   - permissions, distribution, payments
   - robust 3D digital watermarking
   - detecting violations, enforcement
   - real-time viewing on low-cost PCs
   - indexing, cataloguing, searching
   - viewing, measuring, extracting data

Acknowledgements

Faculty and staff

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

Graduate students

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

Undergraduates

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

In Florence

Dott.ssa Cristina Acidini
Dott.ssa Licia Bertani
Matti Auvinen
Dott.ssa Franca Falletti
Alessandra Marino

In Rome

Prof. Eugenio La Rocca
Dott.ssa Anna Somella
Dott.ssa Susanna Le Pera
Dott.ssa Laura Ferrea

In Pisa

Roberto Scopigno

Sponsors

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

Equipment donors

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

Project: http://graphics.stanford.edu/projects/mich/
Software: /software/qsplat/
3D models: /data/mich/