3D photography (II)

Digital Visual Effects, Spring 2005

Yung-Yu Chuang

2005/5/25

with slides by Szymon Rusinkiewicz, Richard Szeliski, Steve Seitz and Brian Curless

Announcements

- Final project will be online tomorrow
- Proposal presentation on next Wednesday
- I will send out your current grades by next Wednesday
- Scribe (SIGGRAPH 2005, CVPR 2005, readings)
- Schedule for the next few weeks
  - 6/1 proposal
  - 6/8 making face/human
  - 6/15 random topics
  - 6/28 final project presentation

Outline

- Range acquisition techniques
- Full model reconstruction
  - ICP
  - Volumetric reconstruction
- Systems, projects and applications
- Final project

Range acquisition

Range acquisition

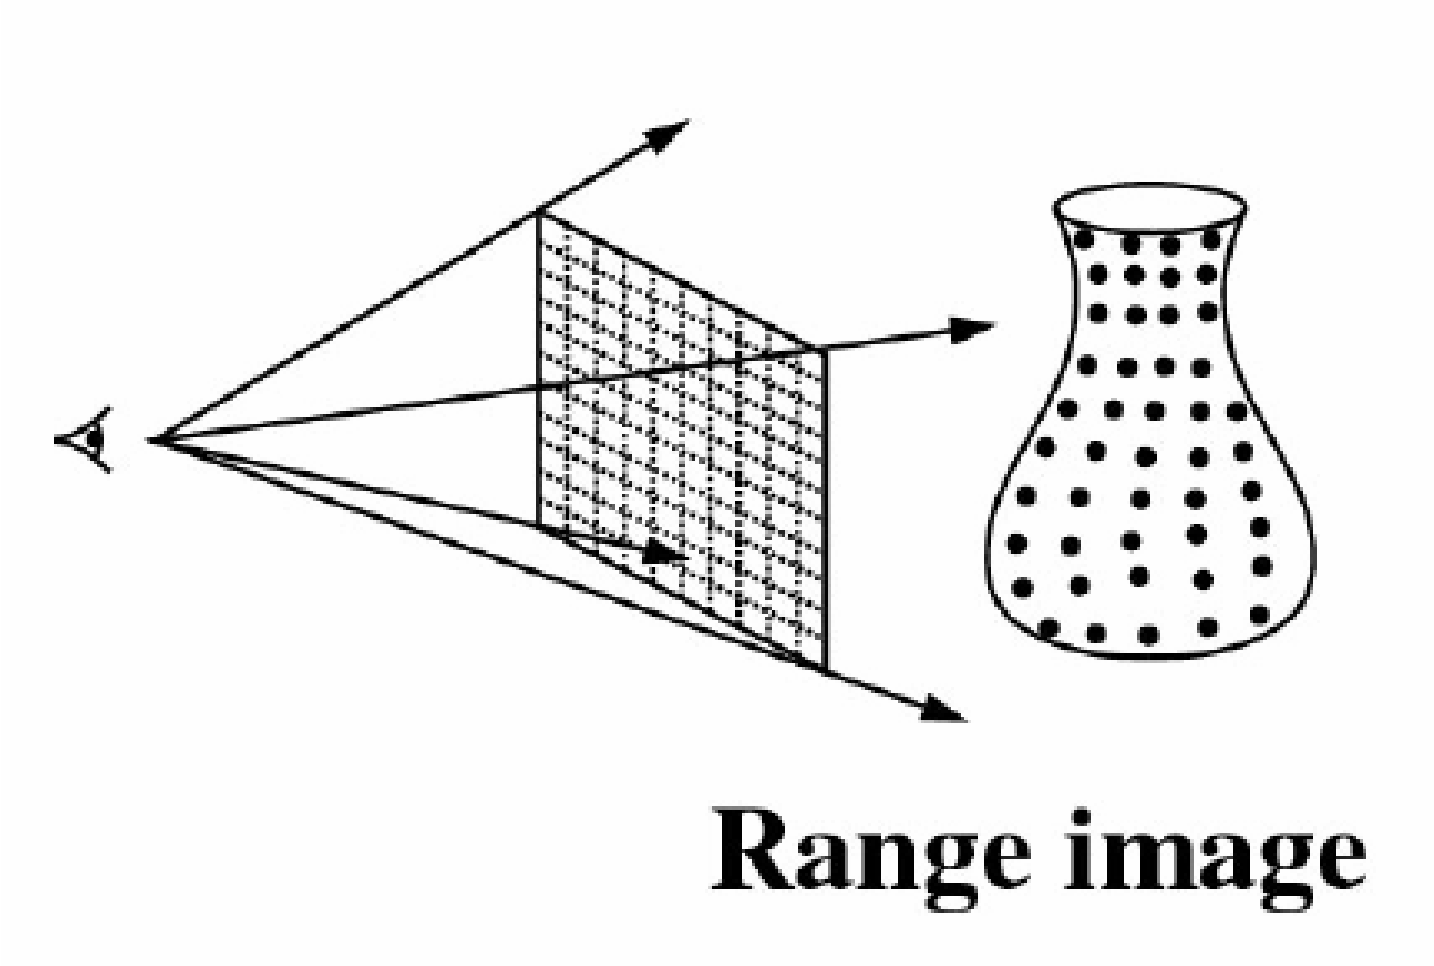

Range acquisition taxonomy

- range acquisition
  - contact
    - mechanical (CMM, jointed arm)
    - inertial (gyroscope, accelerometer)
    - ultrasonic trackers
    - magnetic trackers
  - transmissive
    - industrial CT
    - ultrasound
    - MRI
  - reflective
    - non-optical
      - radar
      - sonar
    - optical

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

Passive approaches

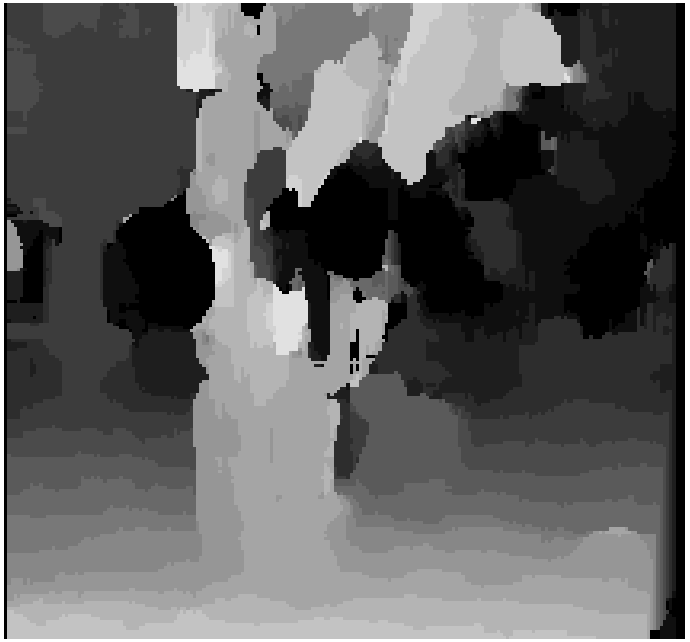

stereo

space carving

Active approaches

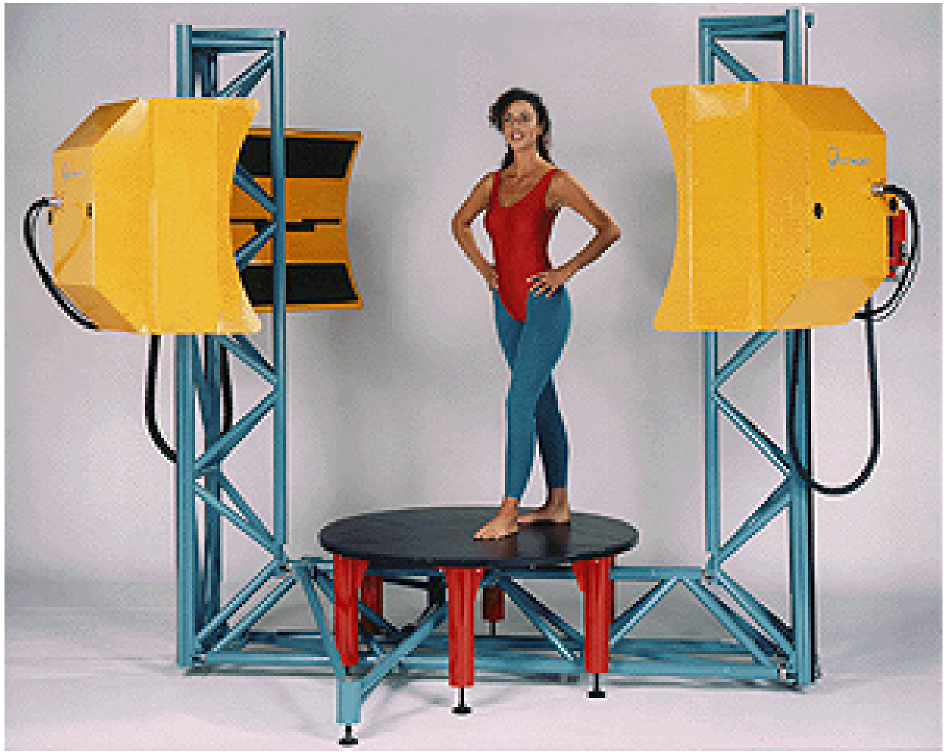

Cyberware whole body scanner

shadow scanning

Active variants

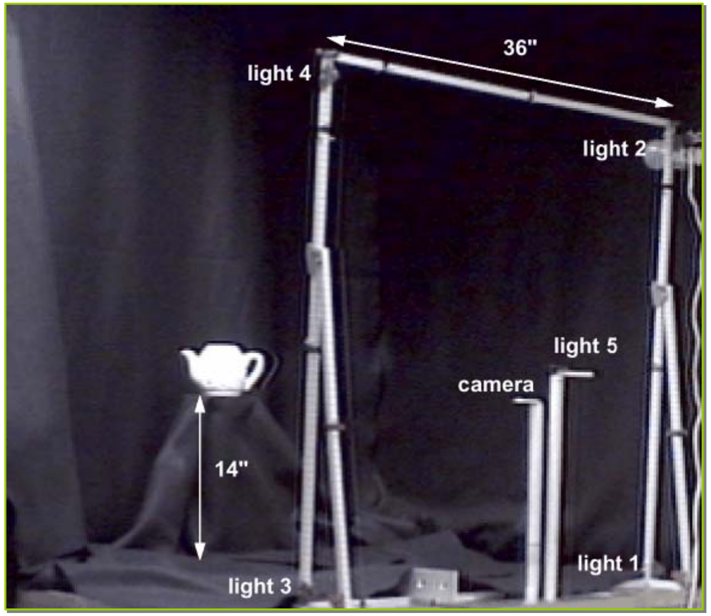

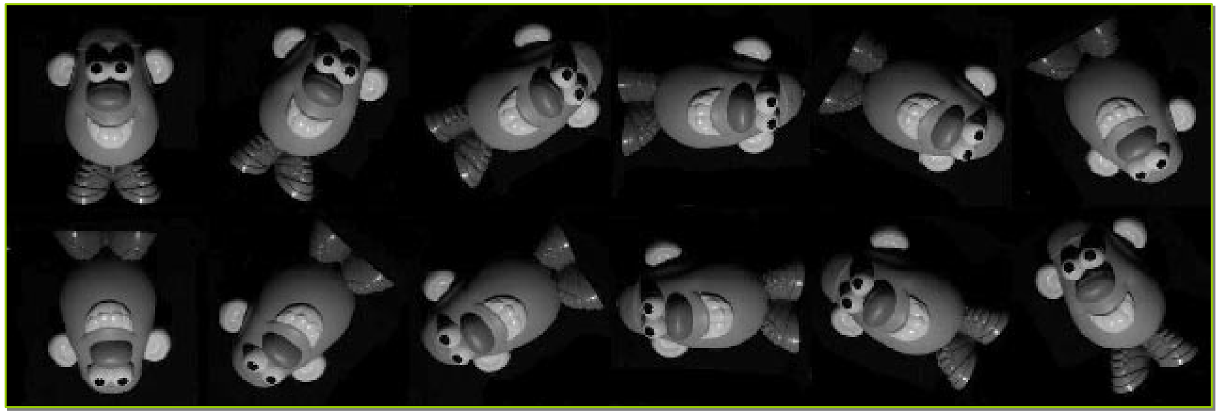

Full model reconstruction

3D Model Acquisition Pipeline

3D Scanner

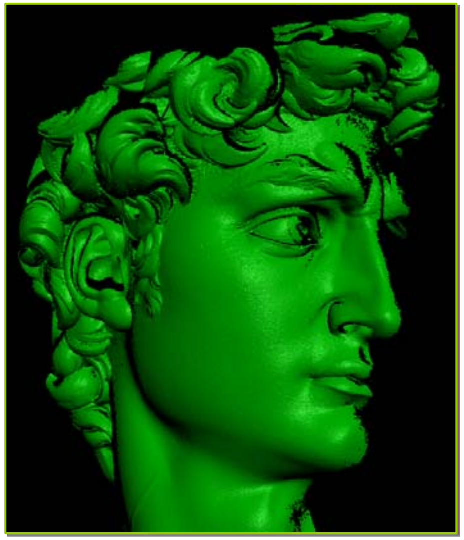

3D Model Acquisition Pipeline

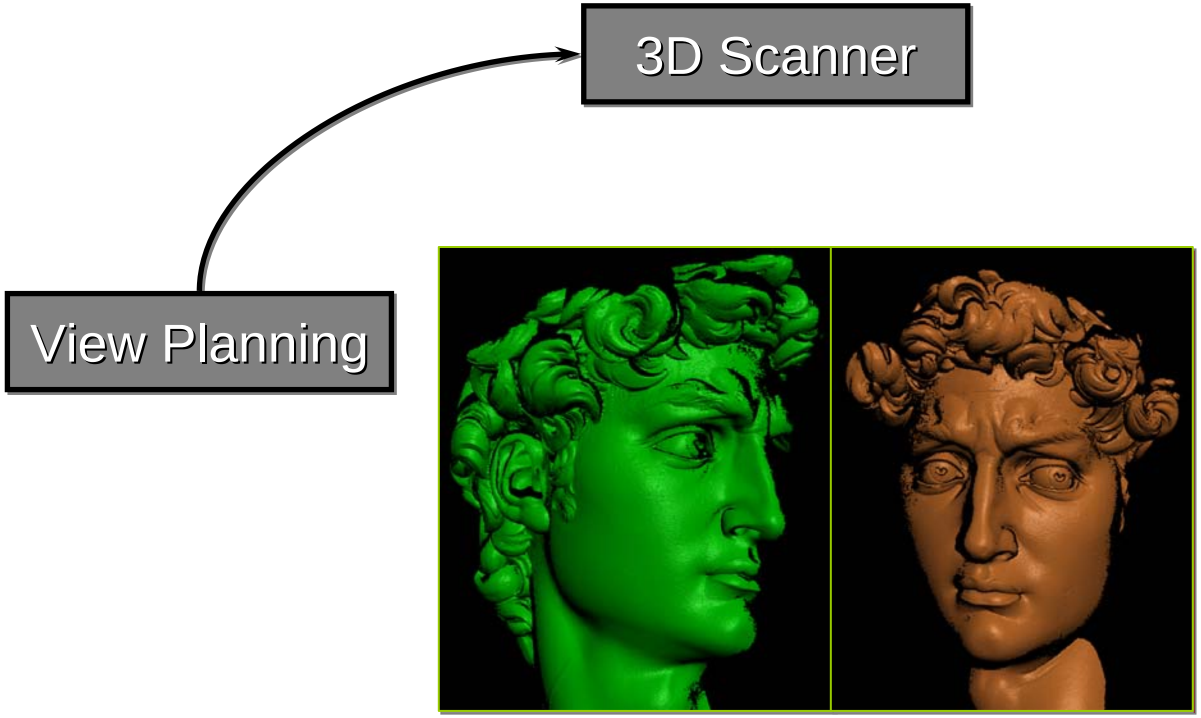

3D Model Acquisition Pipeline

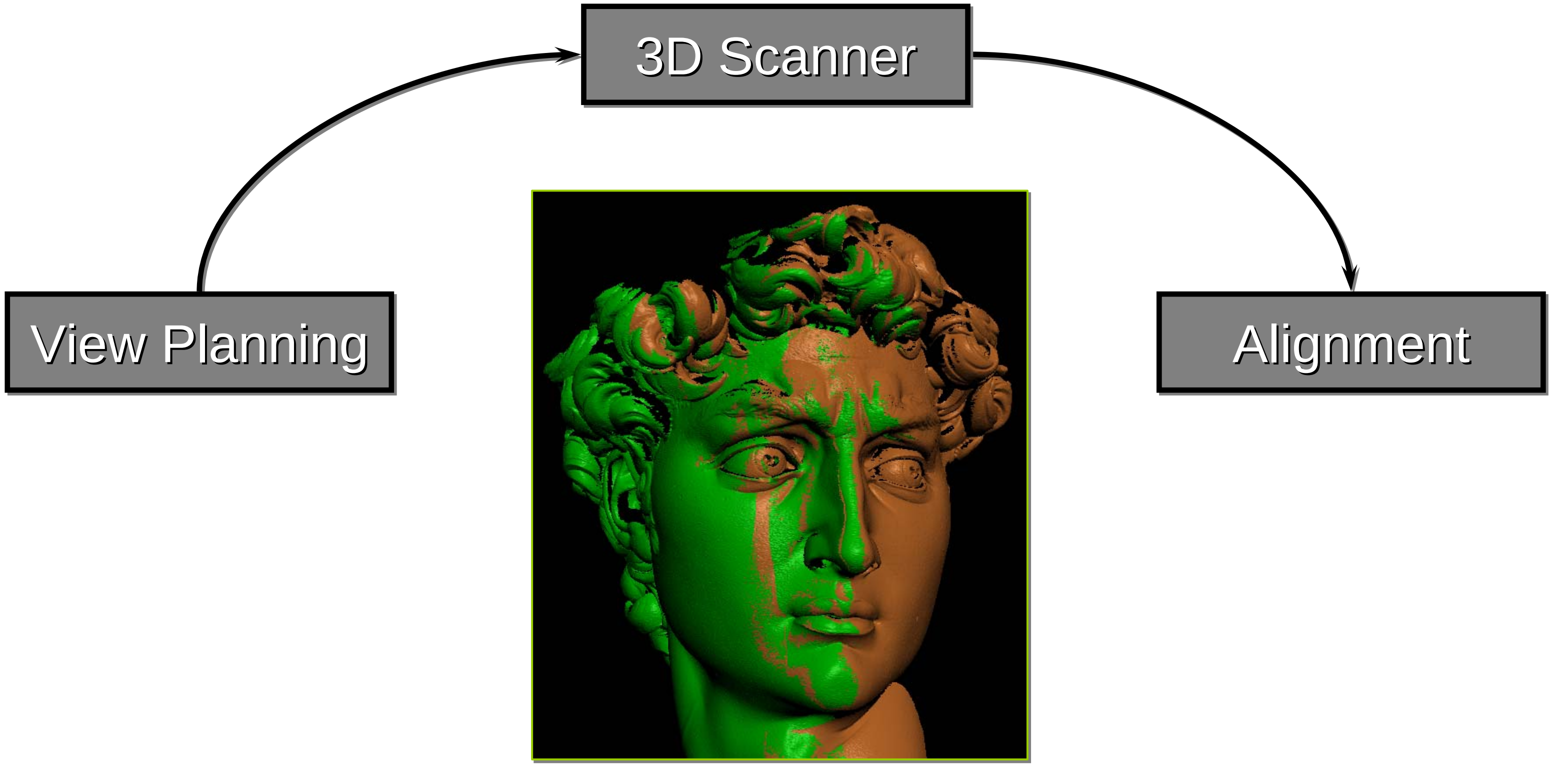

3D Model Acquisition Pipeline

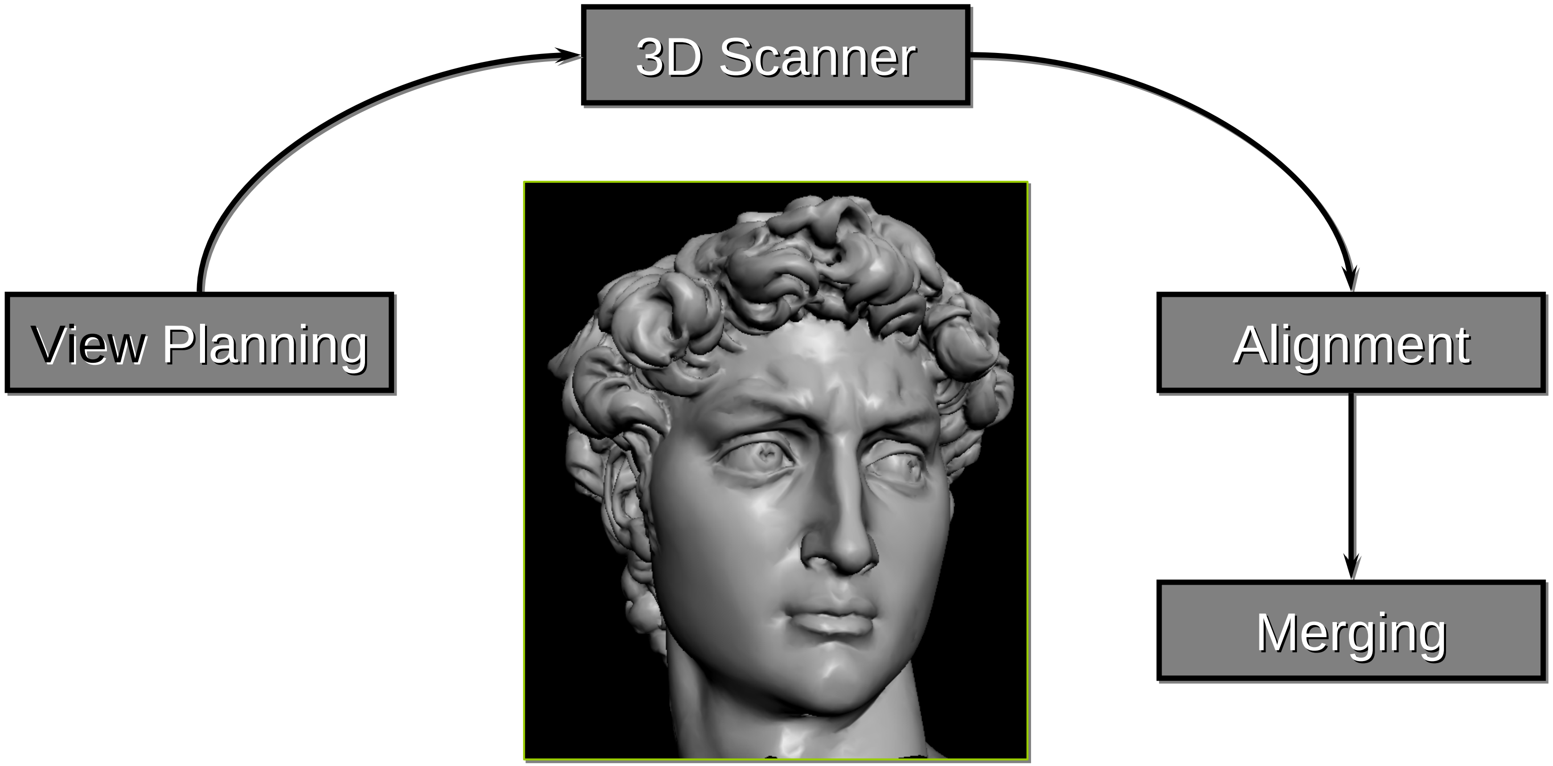

Problem

- Align two partially-overlapping meshes given initial guess for relative transform

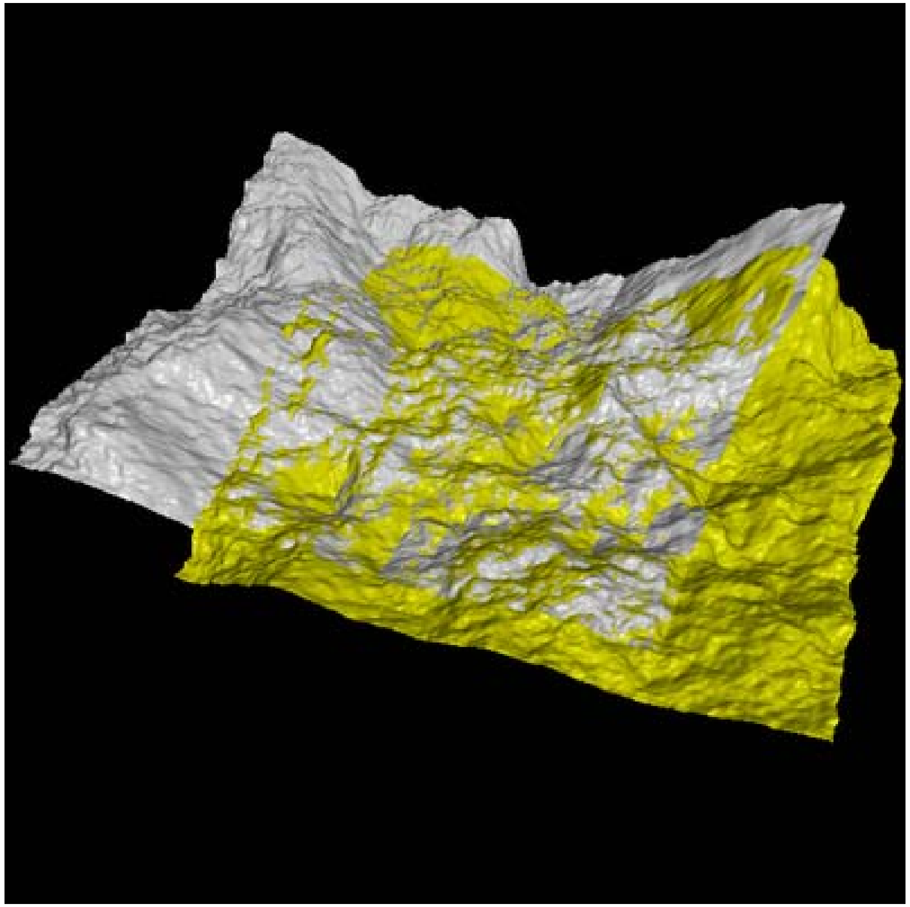

Aligning 3D data

- If correct correspondences are known, it is possible to find correct relative rotation/translation

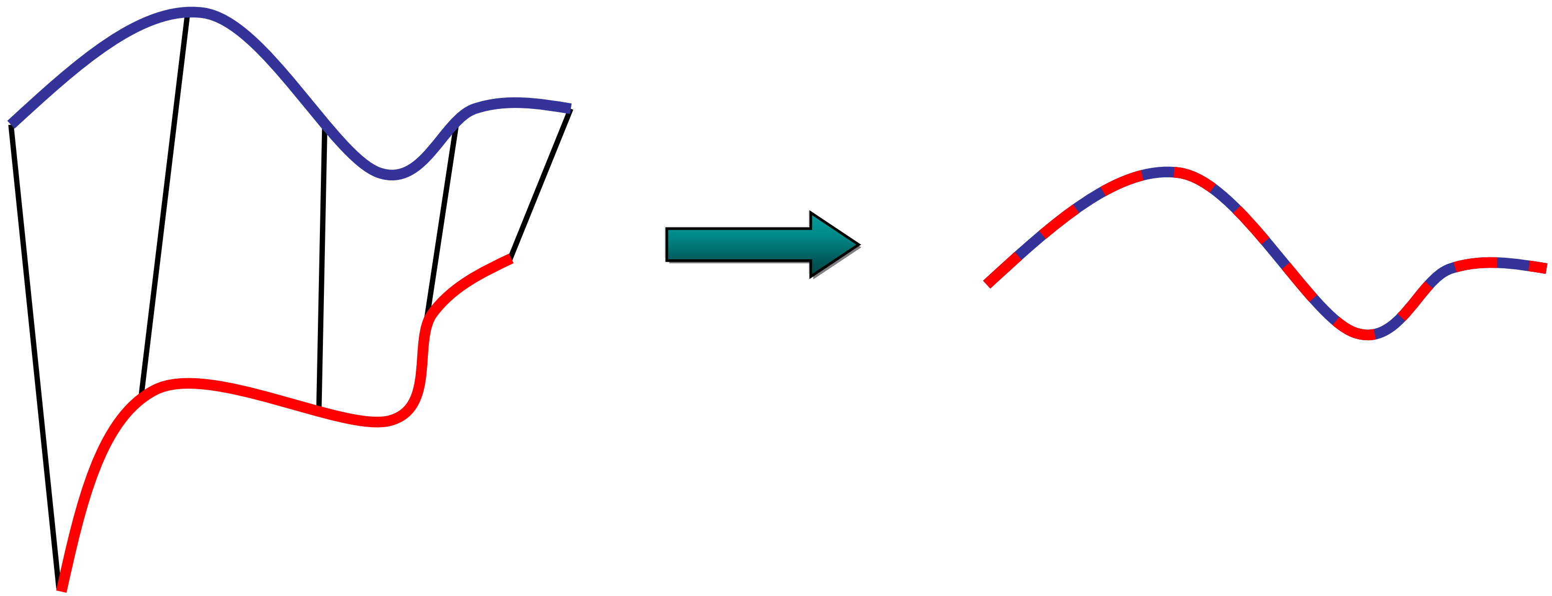

Find the optimal transformation

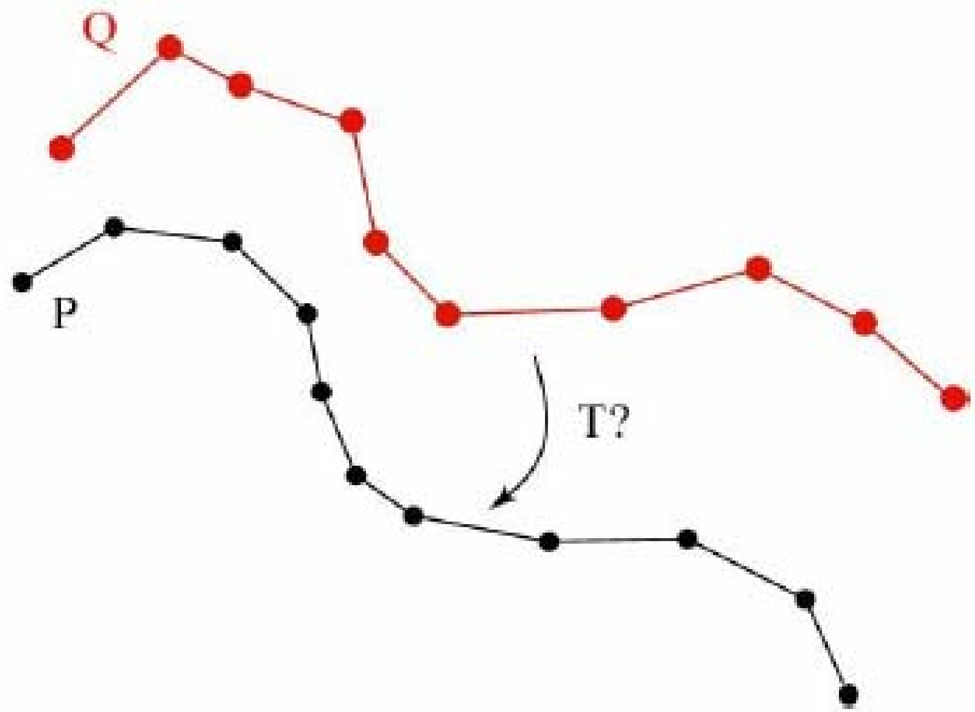

$$E = \sum_{i}^{N} \left\| \mathbf{p}_i - (\mathbf{R}\mathbf{q}_i + \mathbf{t}) \right\|^2$$

Find the optimal transformation

To solve, we first compute the centroid of each point set:

$$\overline{\mathbf{p}} = \sum_{i}^{N} \mathbf{p}_i \quad \overline{\mathbf{q}} = \sum_{i}^{N} \mathbf{q}_i$$

Horn showed that the best rotation satisfies:

$$\underset{\mathbf{R}}{\operatorname{argmax}} \sum_{i}^{N} (\mathbf{p}_i - \overline{\mathbf{p}})^T \mathbf{R} (\mathbf{q}_i - \overline{\mathbf{q}})$$

Find the optimal transformation

DigiVFX

In other words:

1. Convert the points into vectors relative to their centroids.
2. Find a rotation that makes corresponding vectors have dot products as close to 1 as possible.

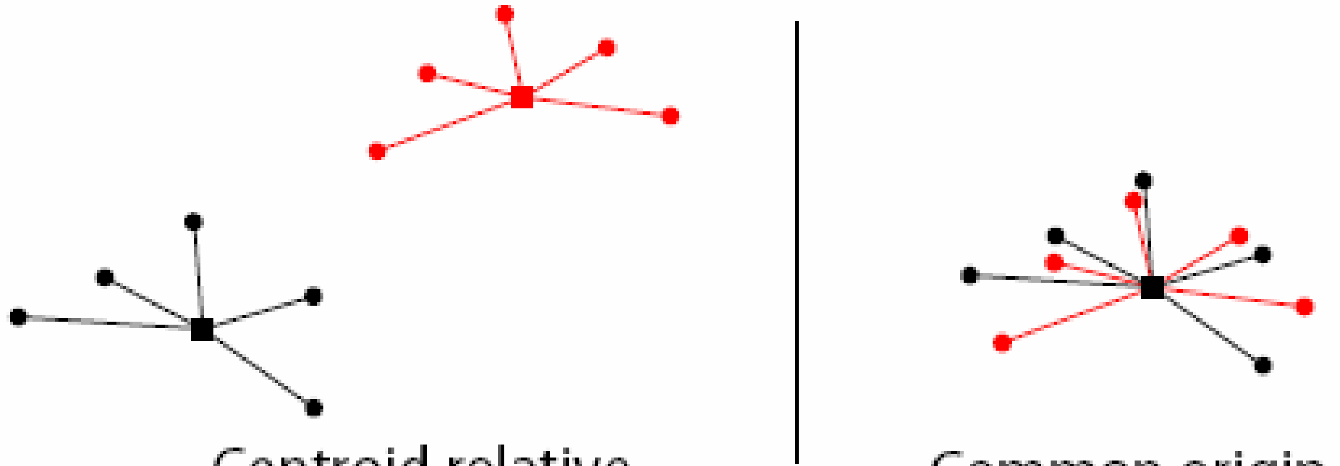

Centroid relative

Common origin

Aligning 3D data

- How to find corresponding points?
- Previous systems based on user input, feature matching, surface signatures, etc.

Aligning 3D data

- Alternative: assume closest points correspond to each other, compute the best transform…

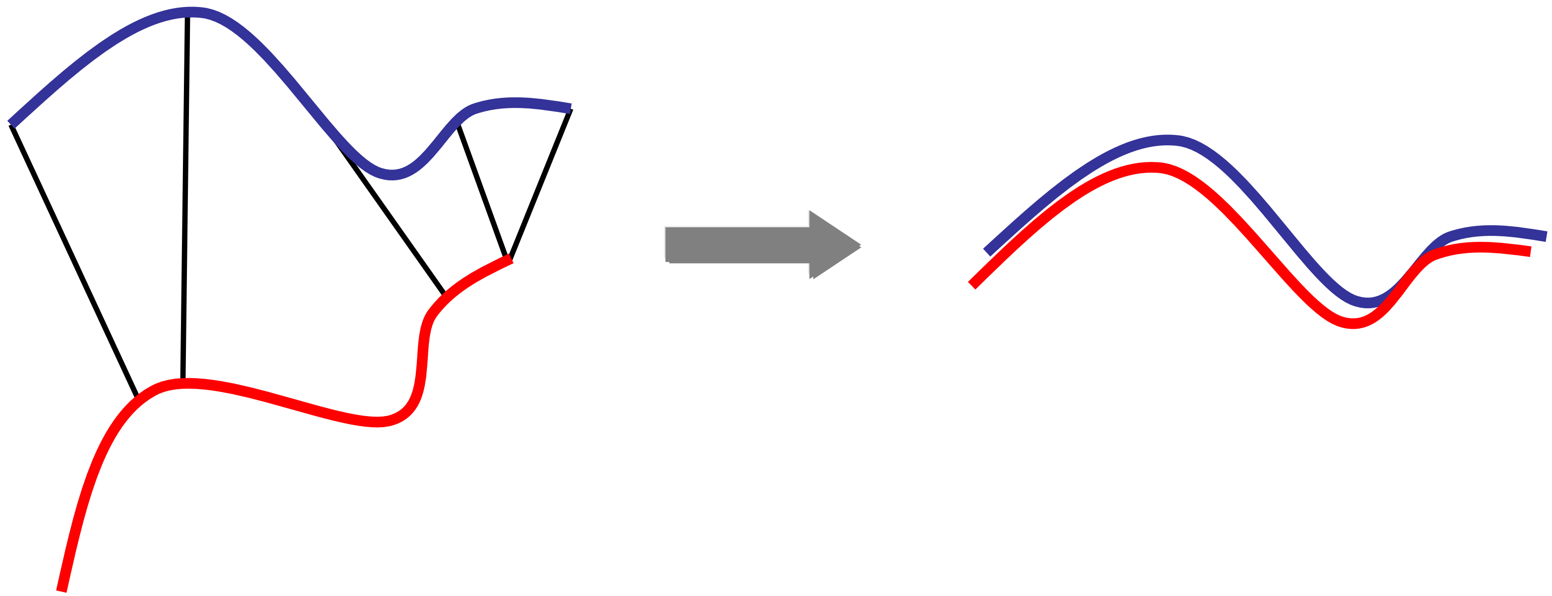

Aligning 3D Data

- … and iterate to find alignment
  - Iterated Closest Points (ICP) [Besl & McKay 92]
- Converges if starting position “close enough“

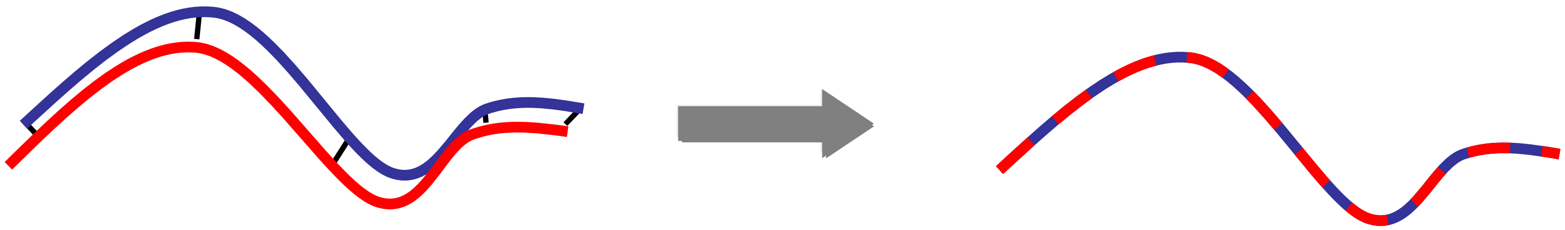

ICP variants

- Variants on the following stages of ICP have been proposed:

1. Selecting source points (from one or both meshes)
2. Matching to points in the other mesh
3. Weighting the correspondences
4. Rejecting certain (outlier) point pairs
5. Assigning an error metric to the current transform
6. Minimizing the error metric

ICP variants

1. Selecting source points (from one or both meshes)
2. Matching to points in the other mesh
3. Weighting the correspondences
4. Rejecting certain (outlier) point pairs
5. Assigning an error metric to the current transform
6. Minimizing the error metric

Selecting source points

- Use all points
- Uniform subsampling
- Random sampling
- Normal-space sampling
  - Ensure that samples have normals distributed as uniformly as possible

Normal-space sampling

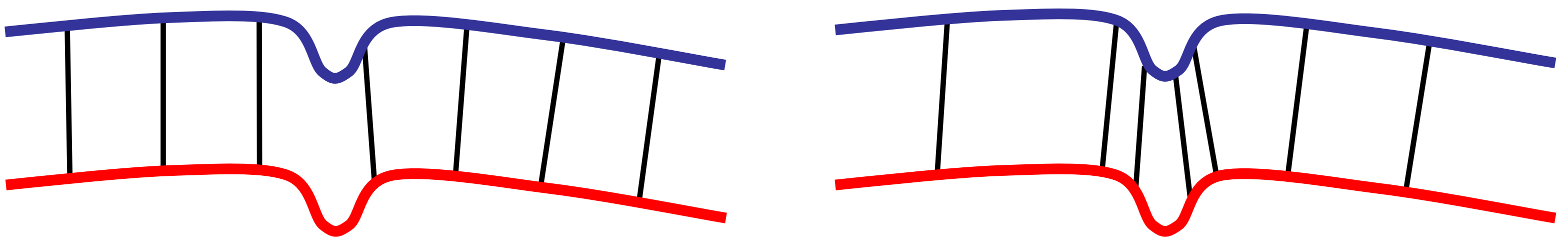

Uniform Sampling Normal-Space Sampling

Normal-space sampling

- Conclusion: normal-space sampling better for mostly-smooth areas with sparse features

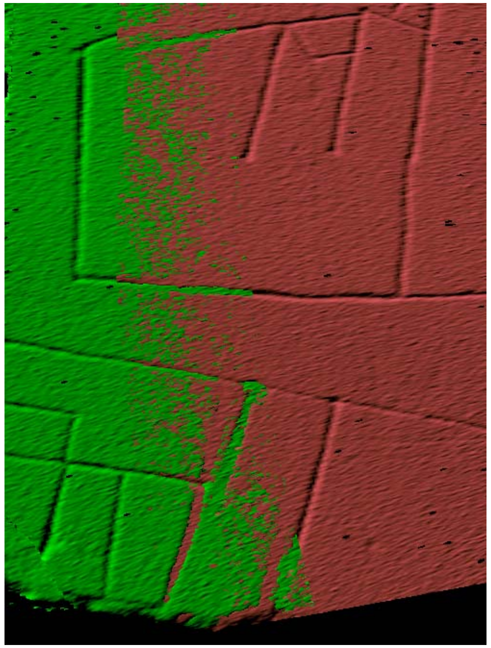

Random sampling

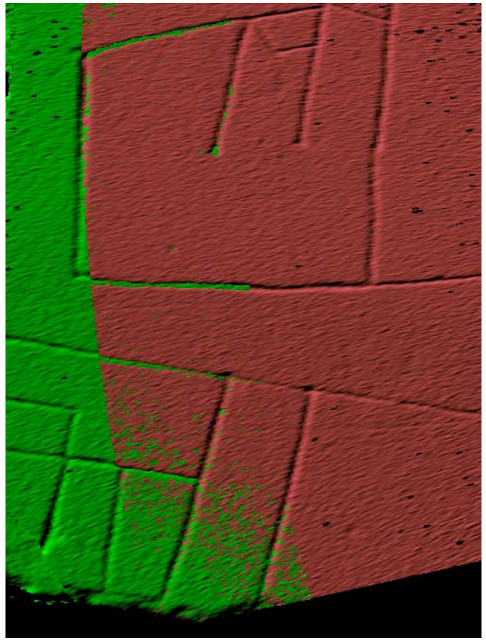

Normal-space sampling

ICP variants

1. Selecting source points (from one or both meshes)
2. Matching to points in the other mesh
3. Weighting the correspondences
4. Rejecting certain (outlier) point pairs
5. Assigning an error metric to the current transform
6. Minimizing the error metric

Matching

- Matching strategy has greatest effect on convergence and speed
- Closest point
- Normal shooting
- Closest compatible point
- Projection

Closest-point matching

- Find closest point in other mesh

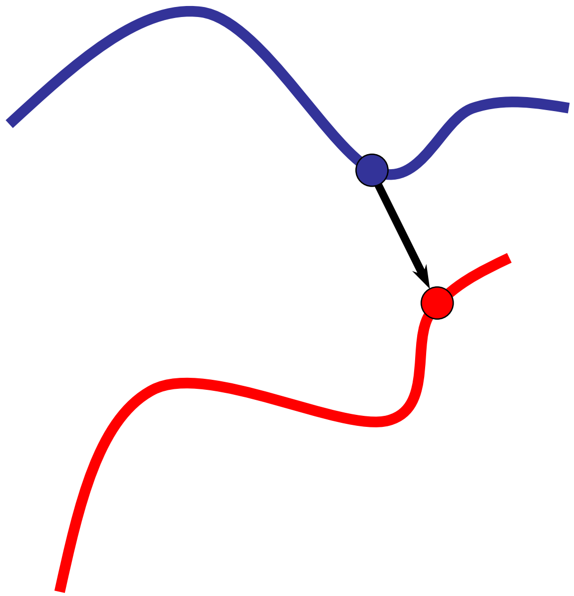

Closest-point matching generally stable, but slow and requires preprocessing

Normal shooting

- Project along normal, intersect other mesh

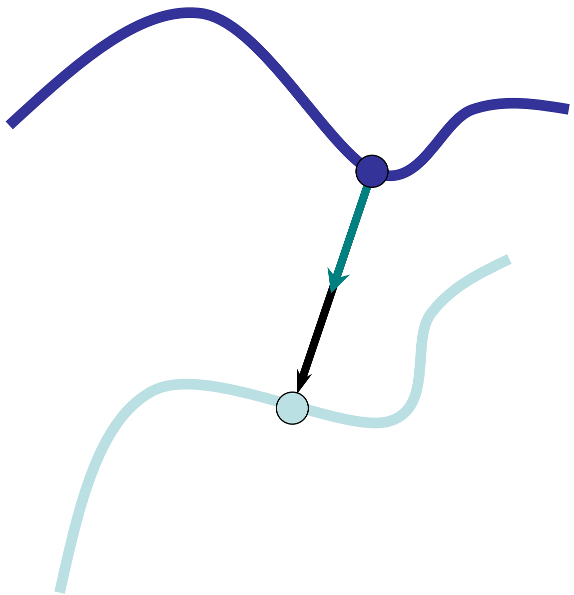

Slightly better than closest point for smooth meshes, worse for noisy or complex meshes

Closest compatible point

- Can improve effectiveness of both of the previous variants by only matching to compatible points
- Compatibility based on normals, colors, etc.
- At limit, degenerates to feature matching

Projection to find correspondences

- Finding closest point is most expensive stage of the ICP algorithm
- Idea: use a simpler algorithm to find correspondences
- For range images, can simply project point [Blais 95]

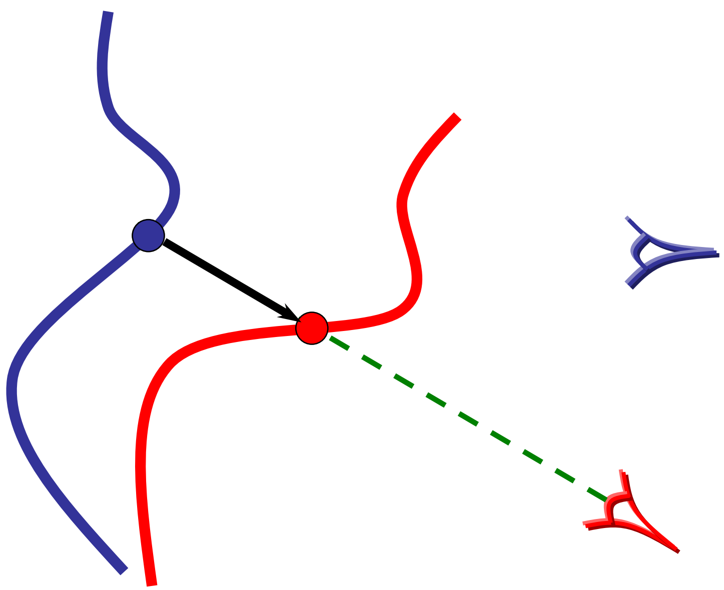

Projection-based matching

- Slightly worse performance per iteration
- Each iteration is one to two orders of magnitude faster than closest-point

ICP variants

1. Selecting source points (from one or both meshes)
2. Matching to points in the other mesh
3. Weighting the correspondences
4. Rejecting certain (outlier) point pairs
5. Assigning an error metric to the current transform
6. Minimizing the error metric

ICP variants

1. Selecting source points (from one or both meshes)
2. Matching to points in the other mesh
3. Weighting the correspondences
4. Rejecting certain (outlier) point pairs
5. Assigning an error metric to the current transform
6. Minimizing the error metric

Point-to-plane error metric

- Using point-to-plane distance instead of point-to-point lets flat regions slide along each other [Chen & Medioni 91]

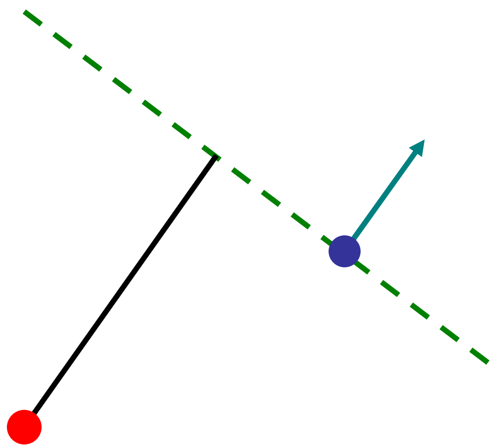

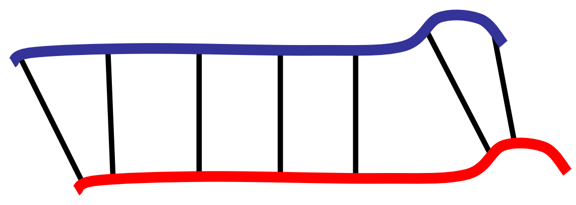

High-speed ICP algorithm

- ICP algorithm with projection-based correspondences, point-to-plane matching can align meshes in a few tens of ms. (cf. over 1 sec. with closest-point)

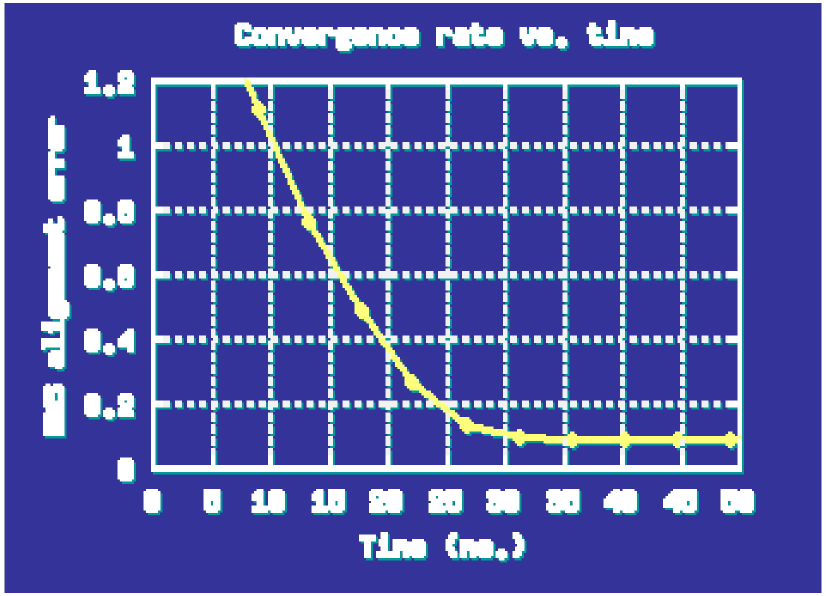

Range processing pipeline

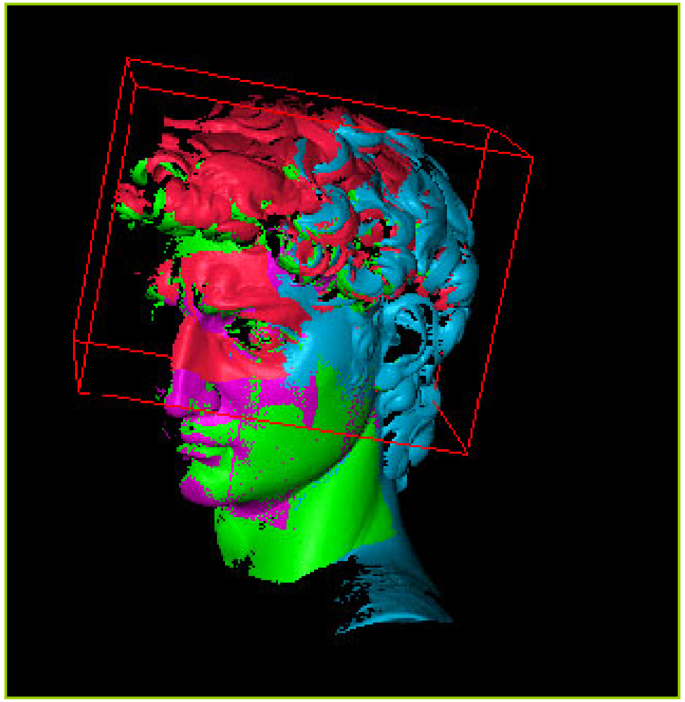

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Range processing pipeline

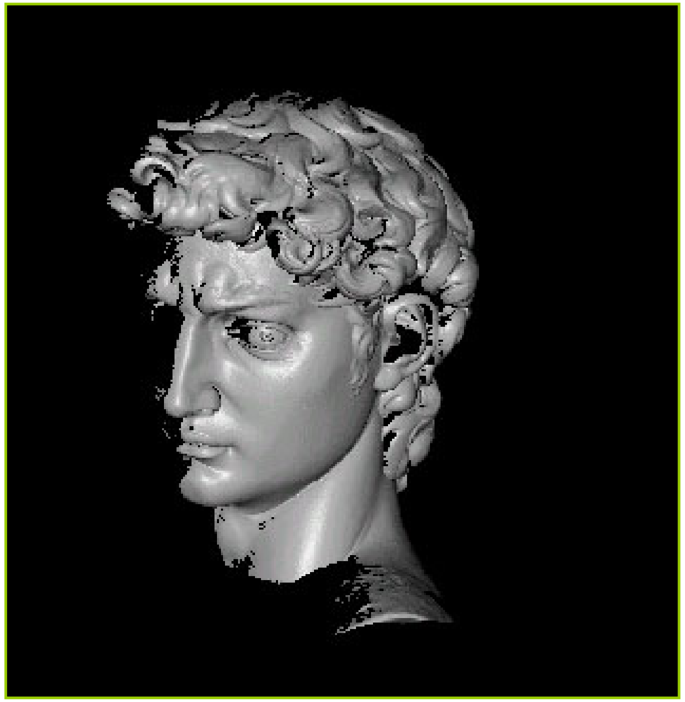

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Volumetric reconstruction

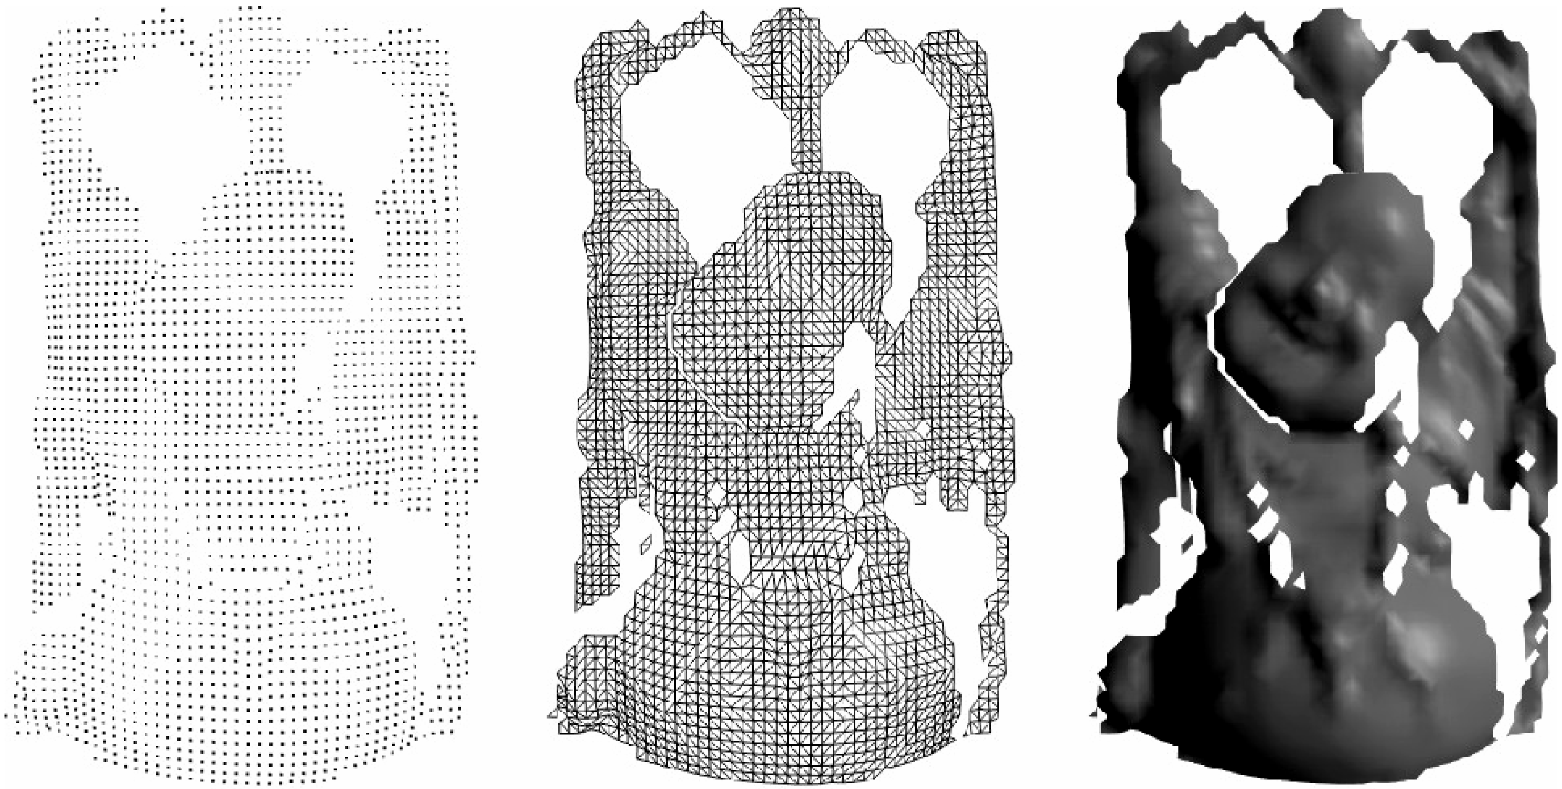

Range image Tesellation Range surface

Signed distance function

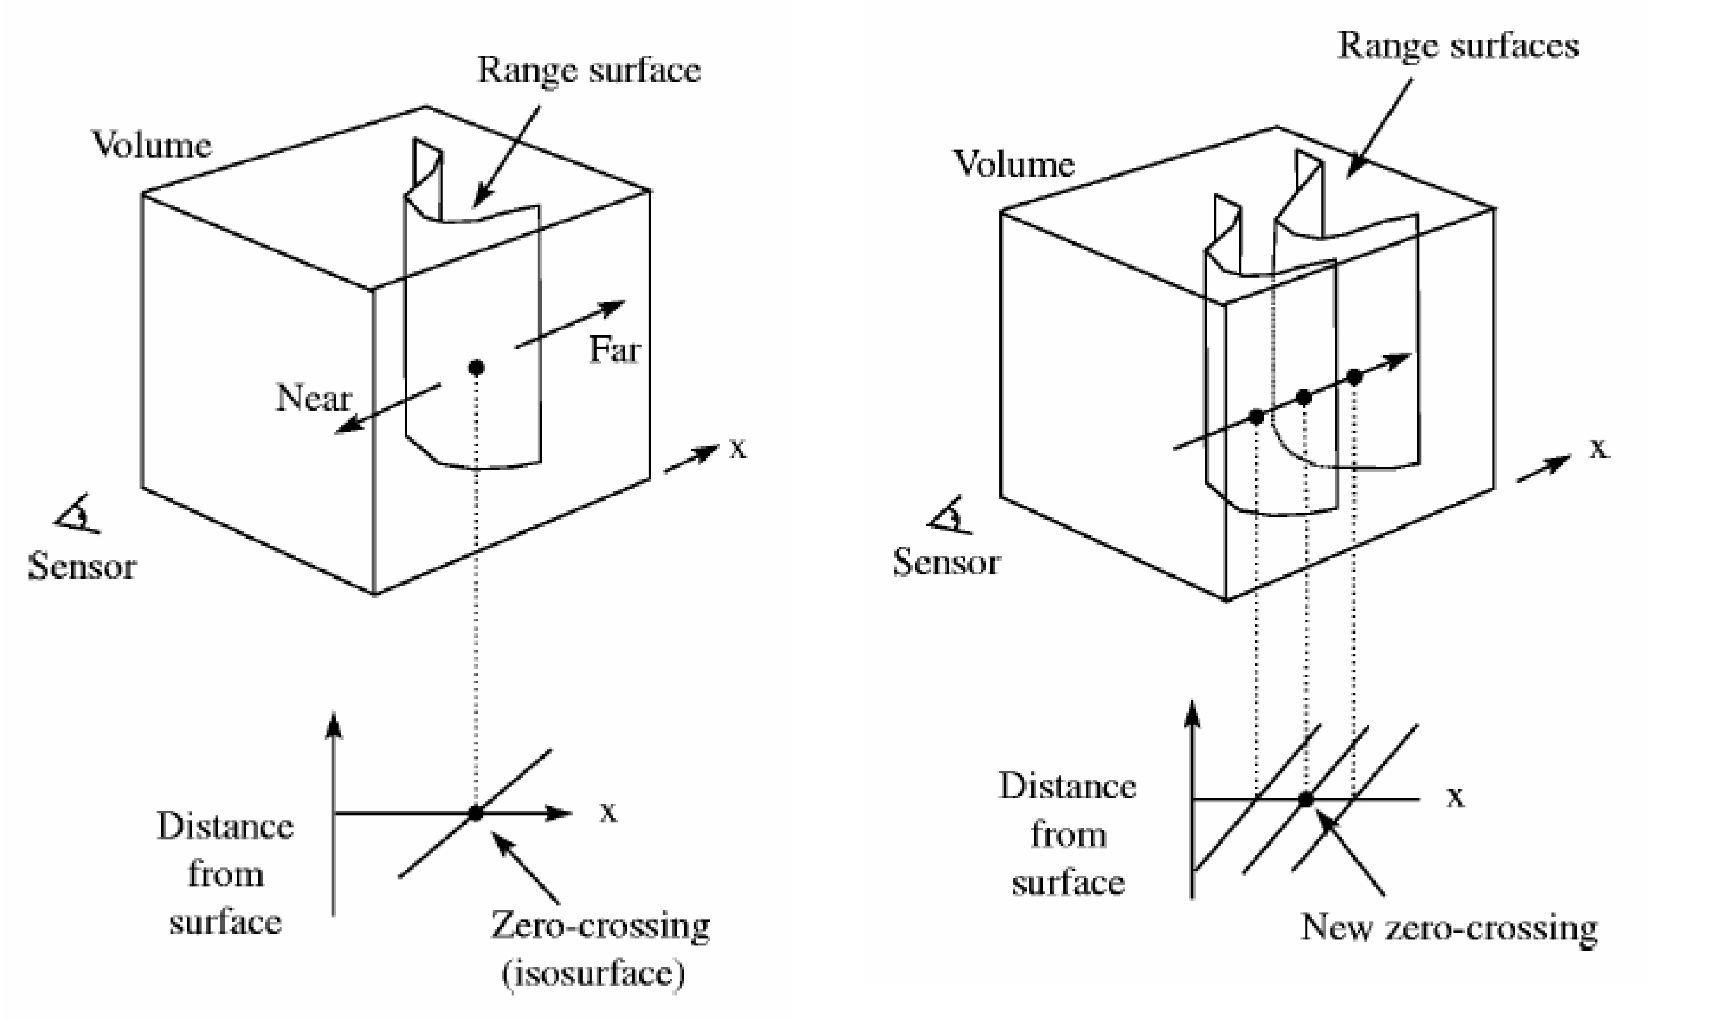

Overview

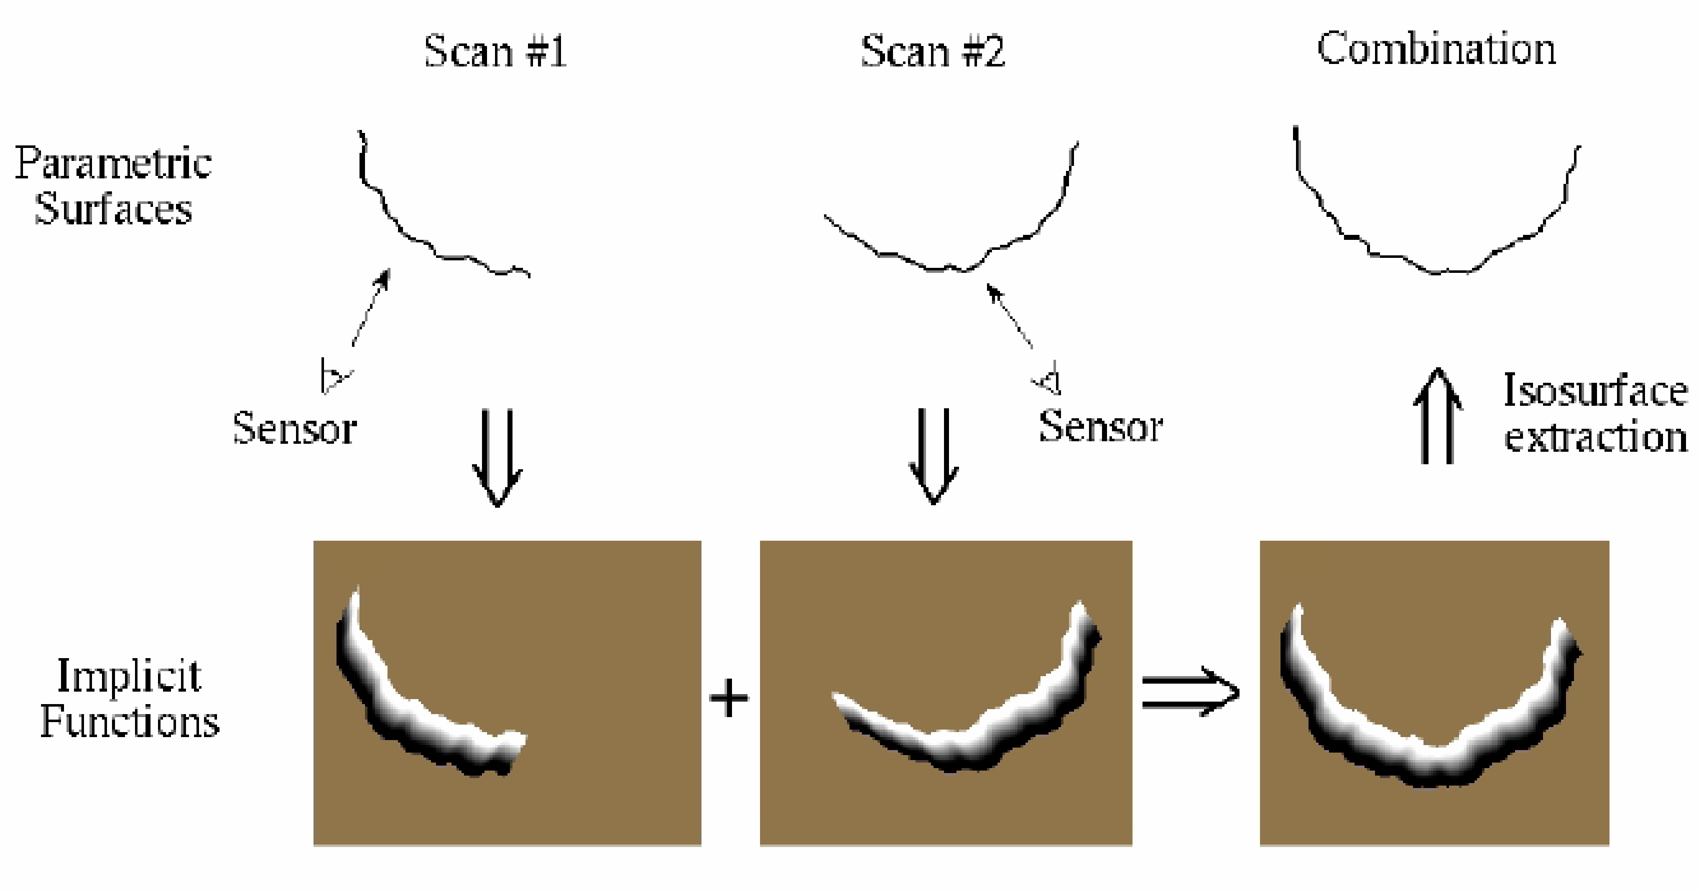
Scan #1
Scan #2
Combination
Parametric Surfaces
Sensor
Sensor
Isosurface extraction
Implicit Functions
+

Weighting

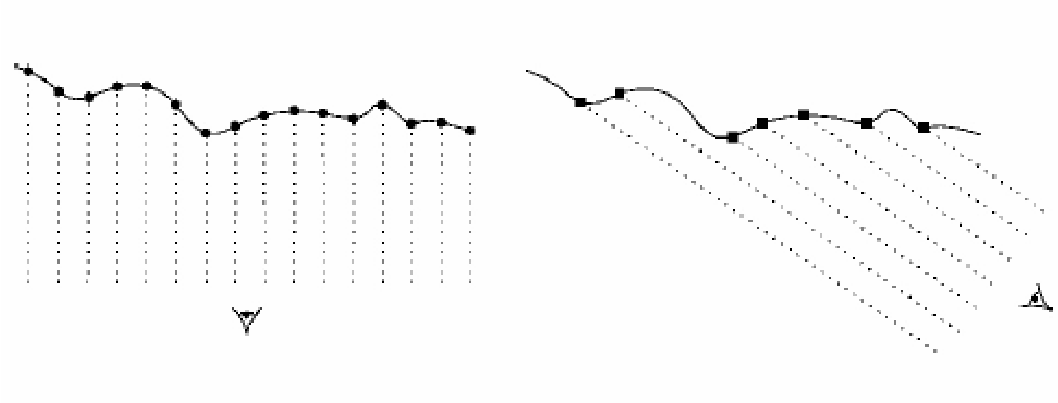

Range surface

Confidence rendering

Volumetric integration

(Curless and Levoy, Siggraph´96)

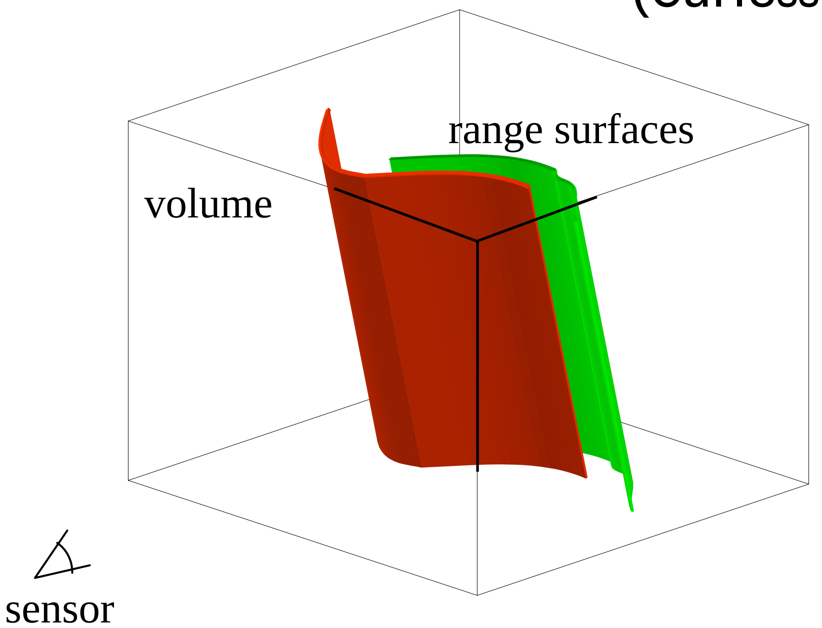

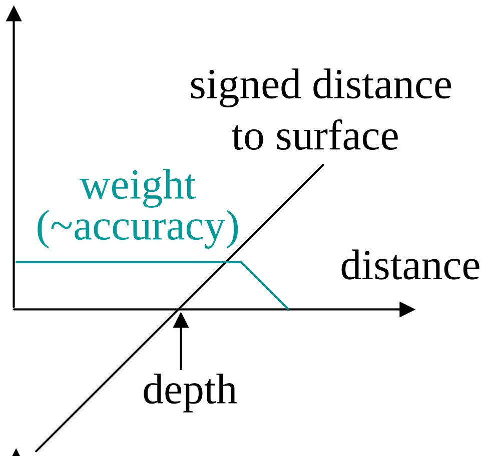

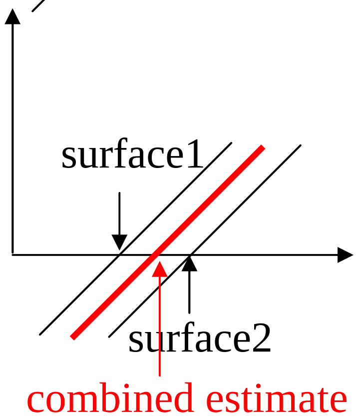

- use voxel space
- new surface as zero-crossing
  (find using marching cubes)
- least-squares estimate
  (zero derivative=minimum)

Isosurfacing: marching cubes

First 2D, marching squares

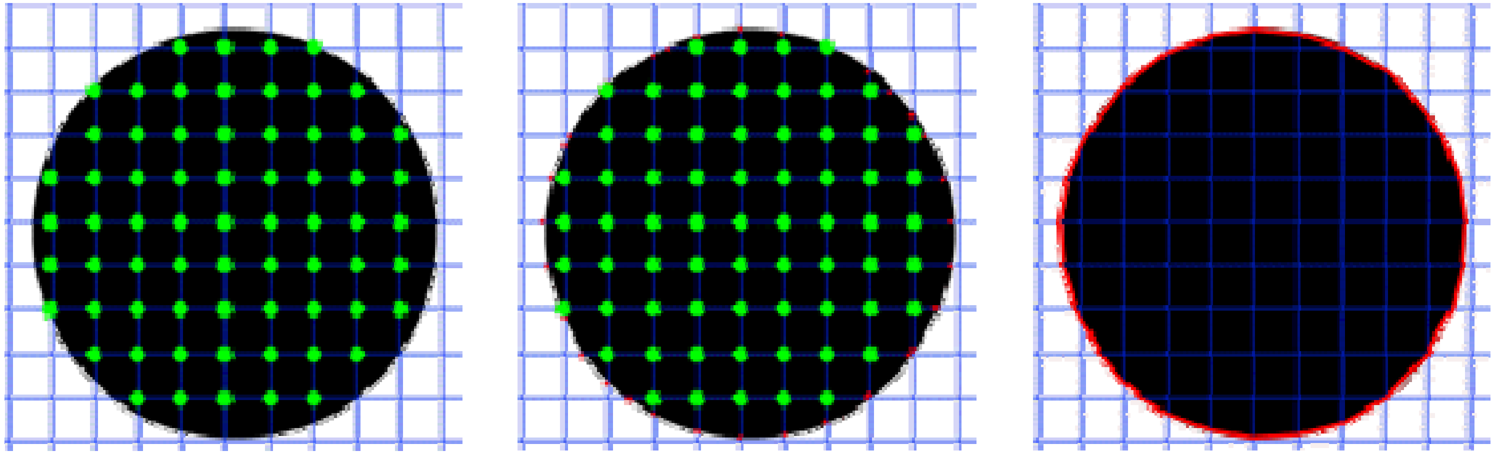

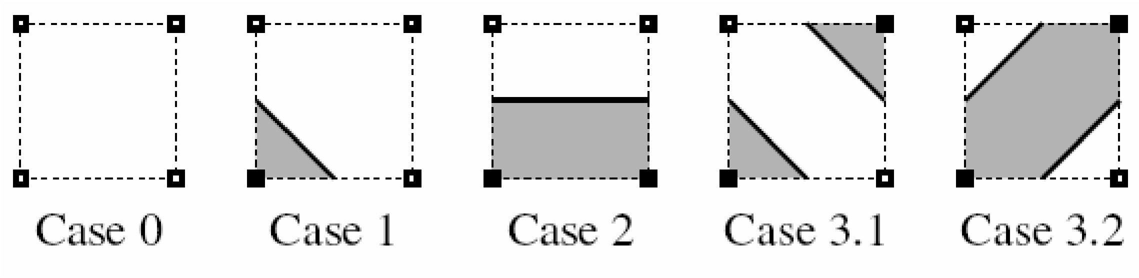

Marching cubes

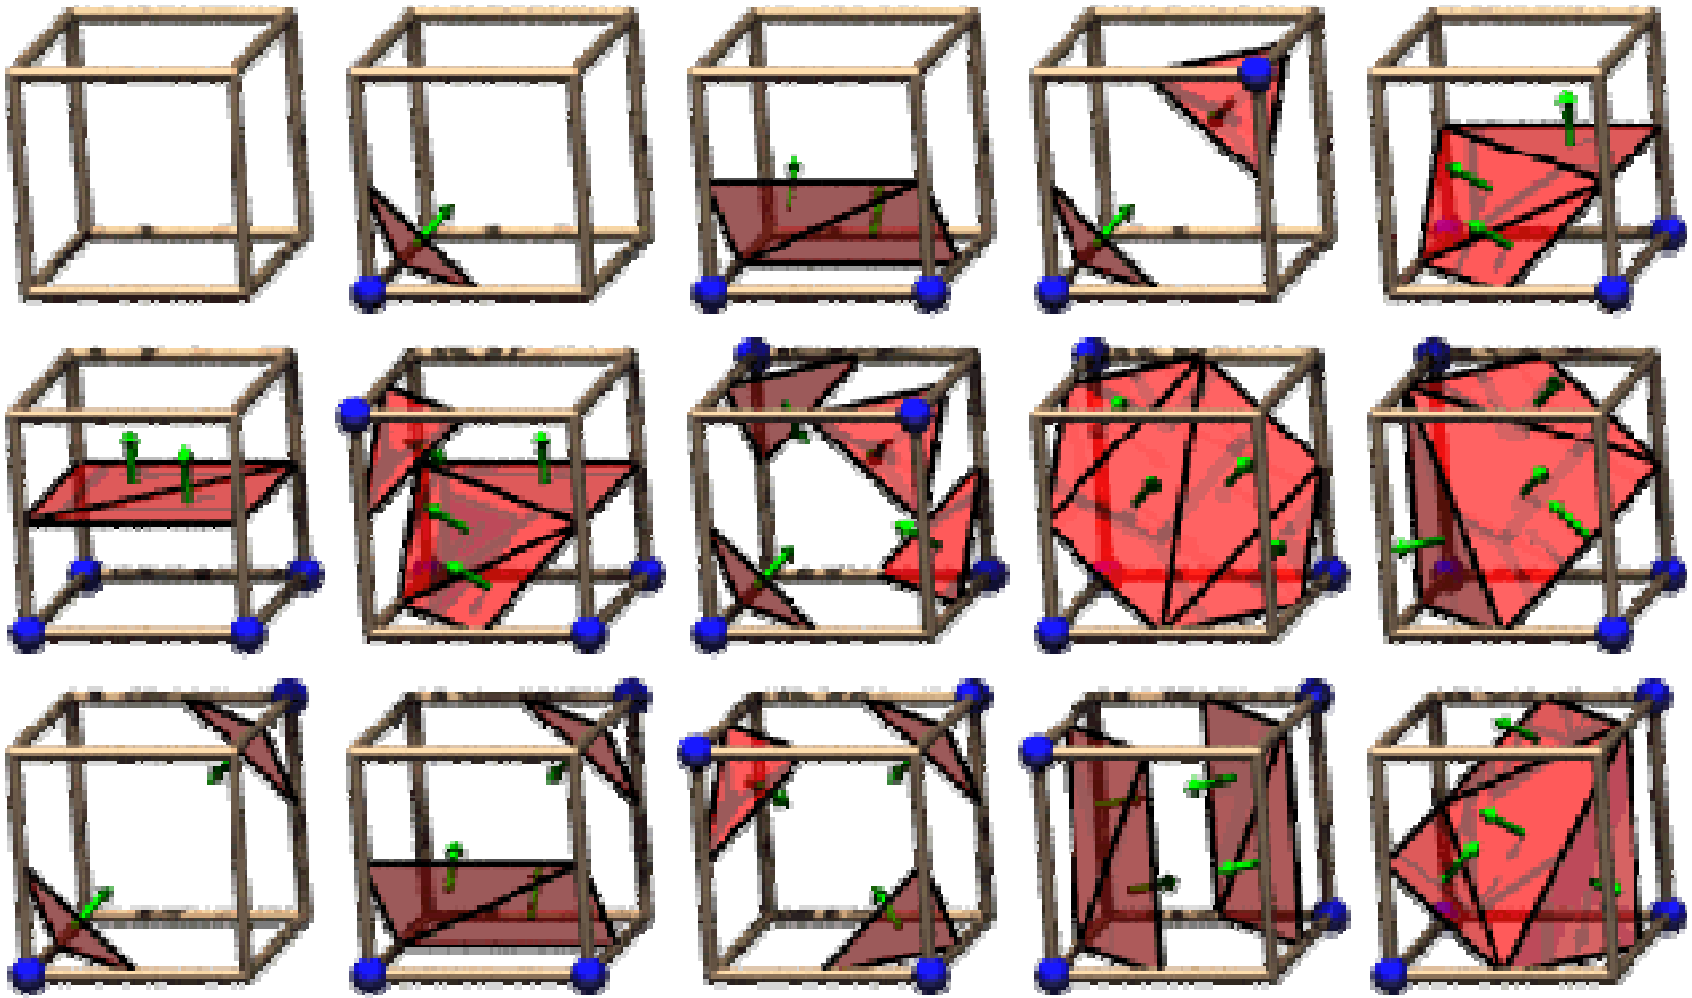

The 15 Cube Combinations

Marching cubes

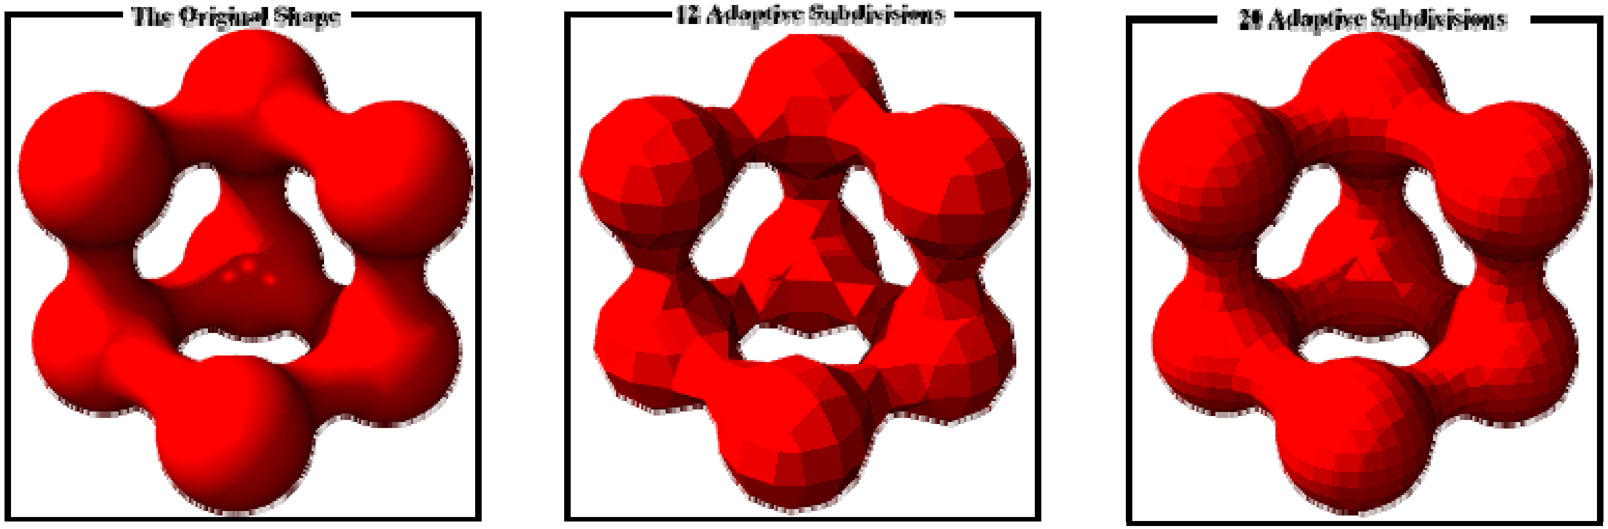

Results

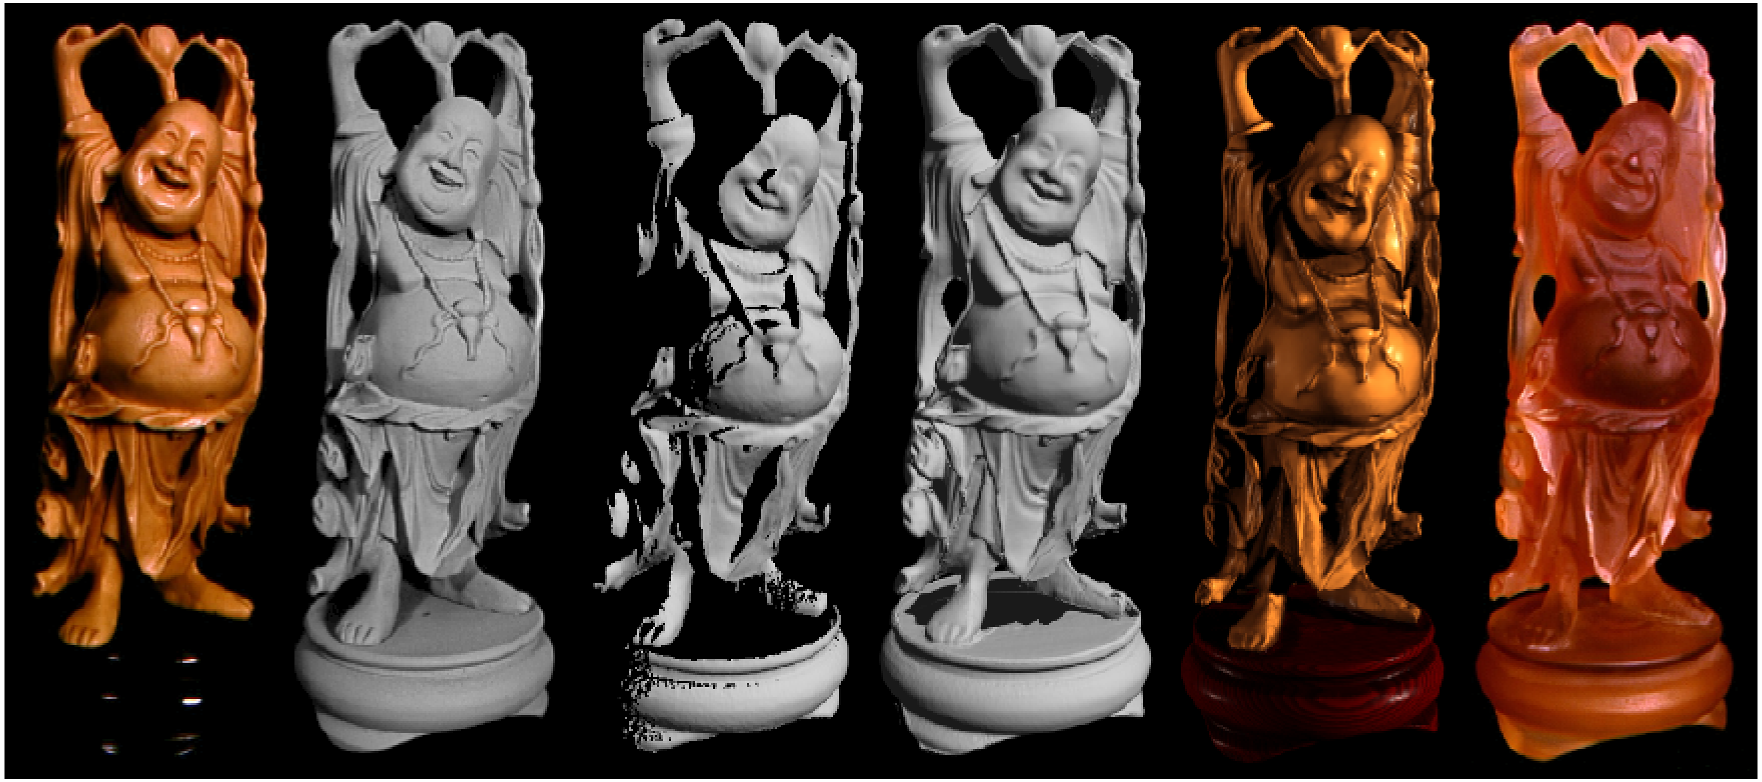

Photograph of original model	Photograph of painted original	Range surface from one scan	Reconstruction before hole–filling	Reconstruction after hole–filling	Hardcopy

Systems, projects and applications

3D model acquisition pipeline

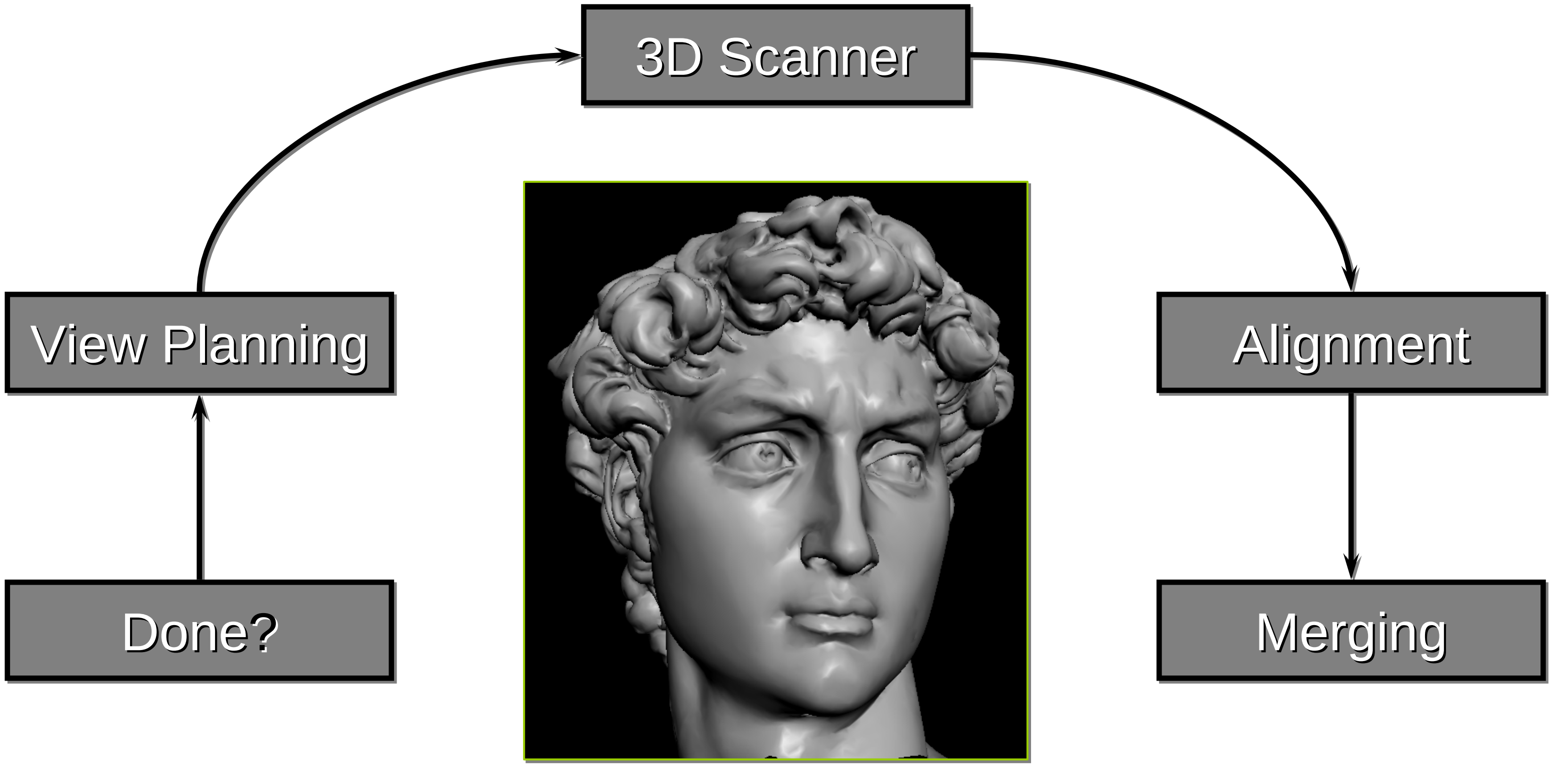

3D model acquisition pipeline

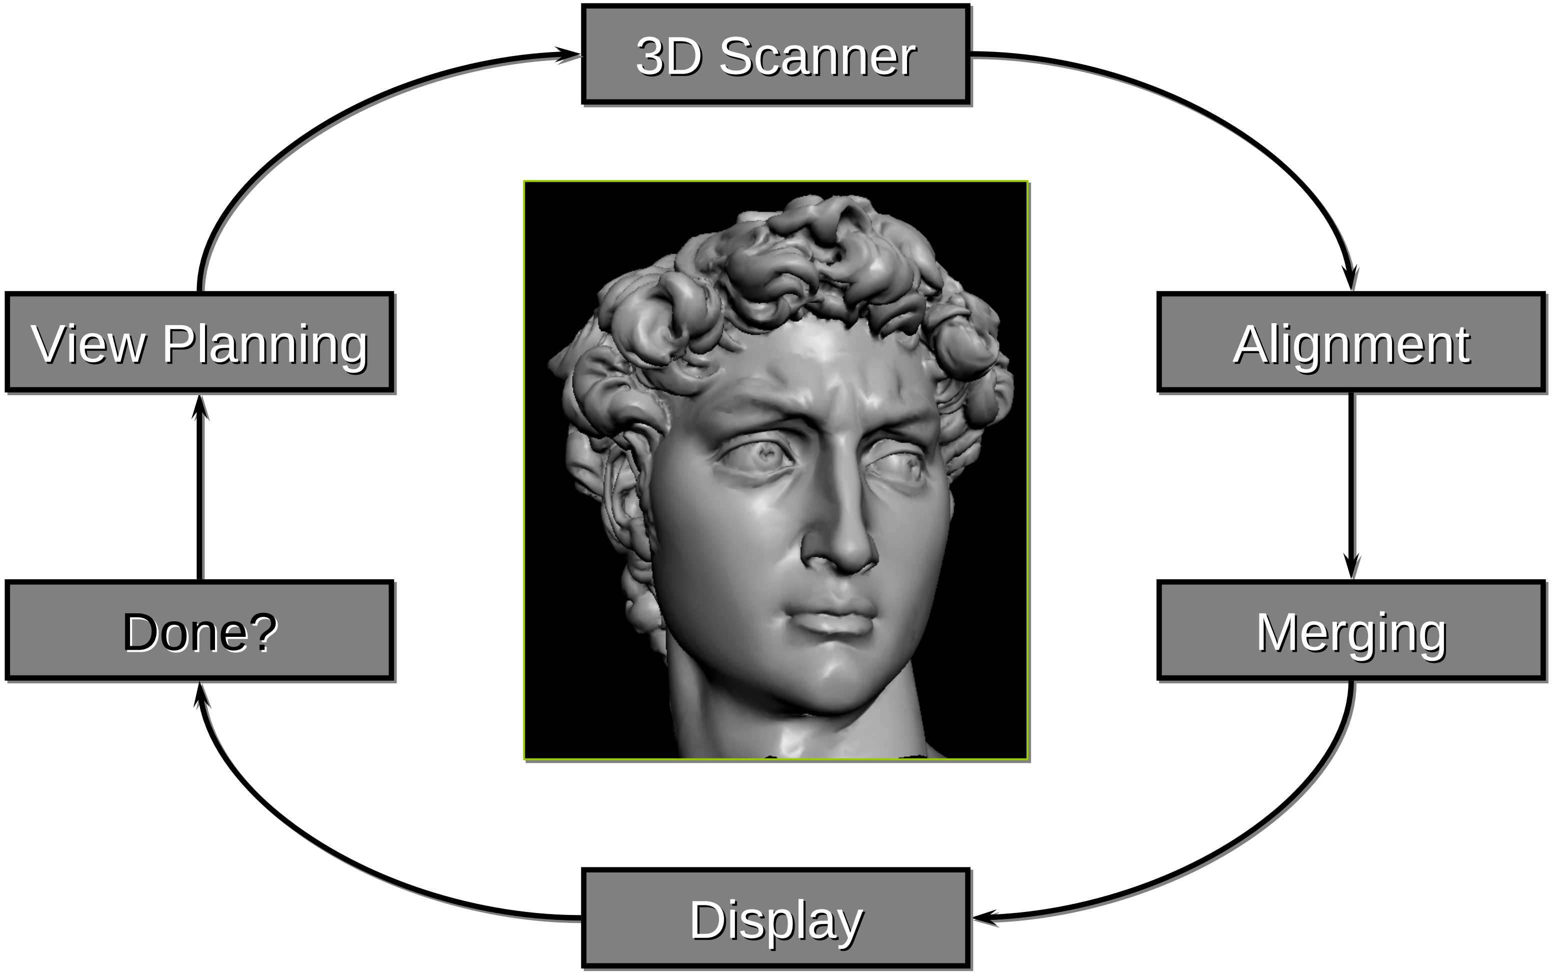

Real-time 3D model acquisition

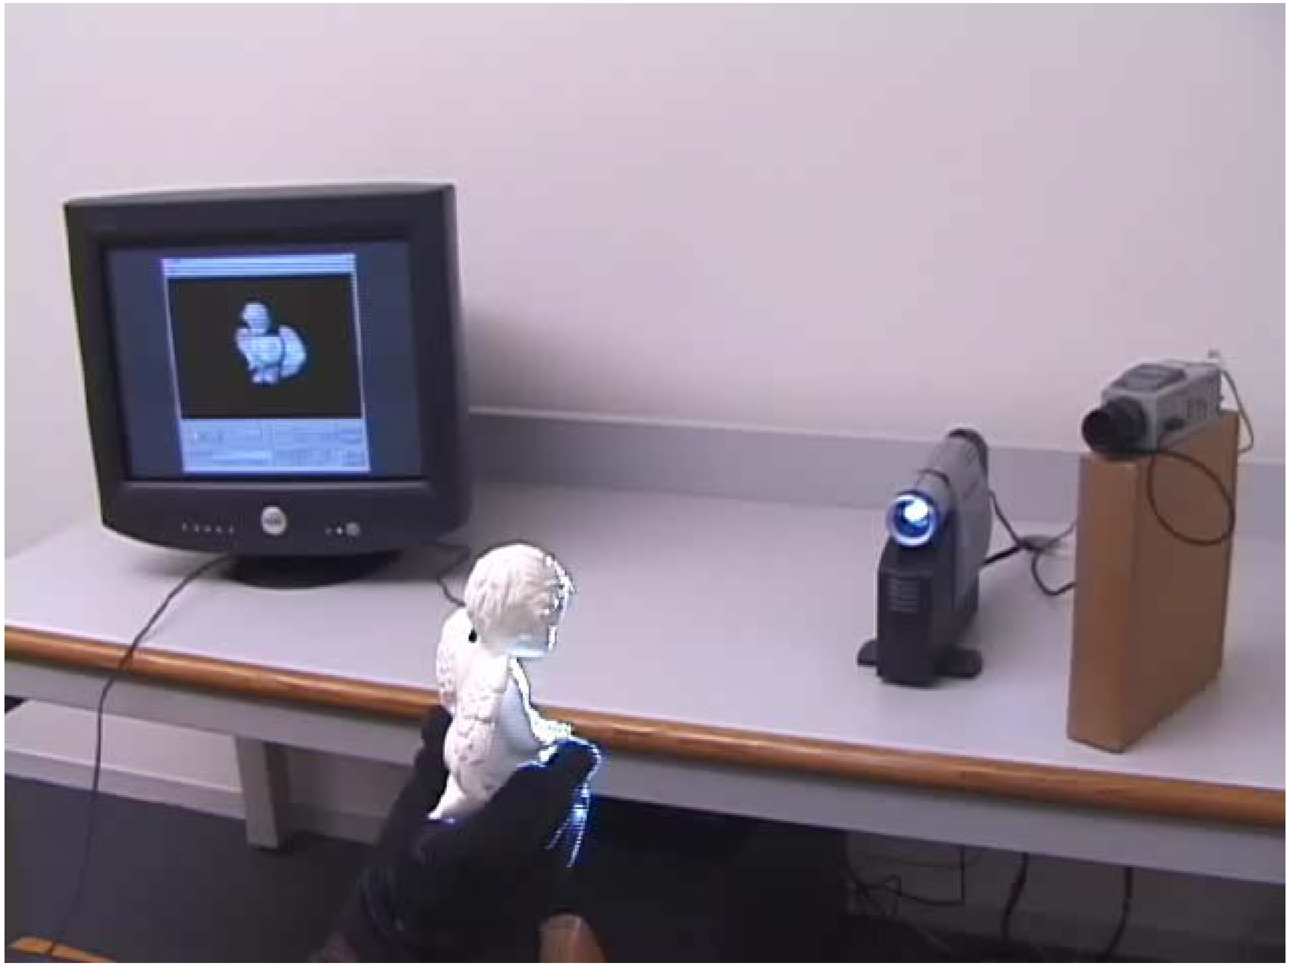

Real-time 3D model acquisition

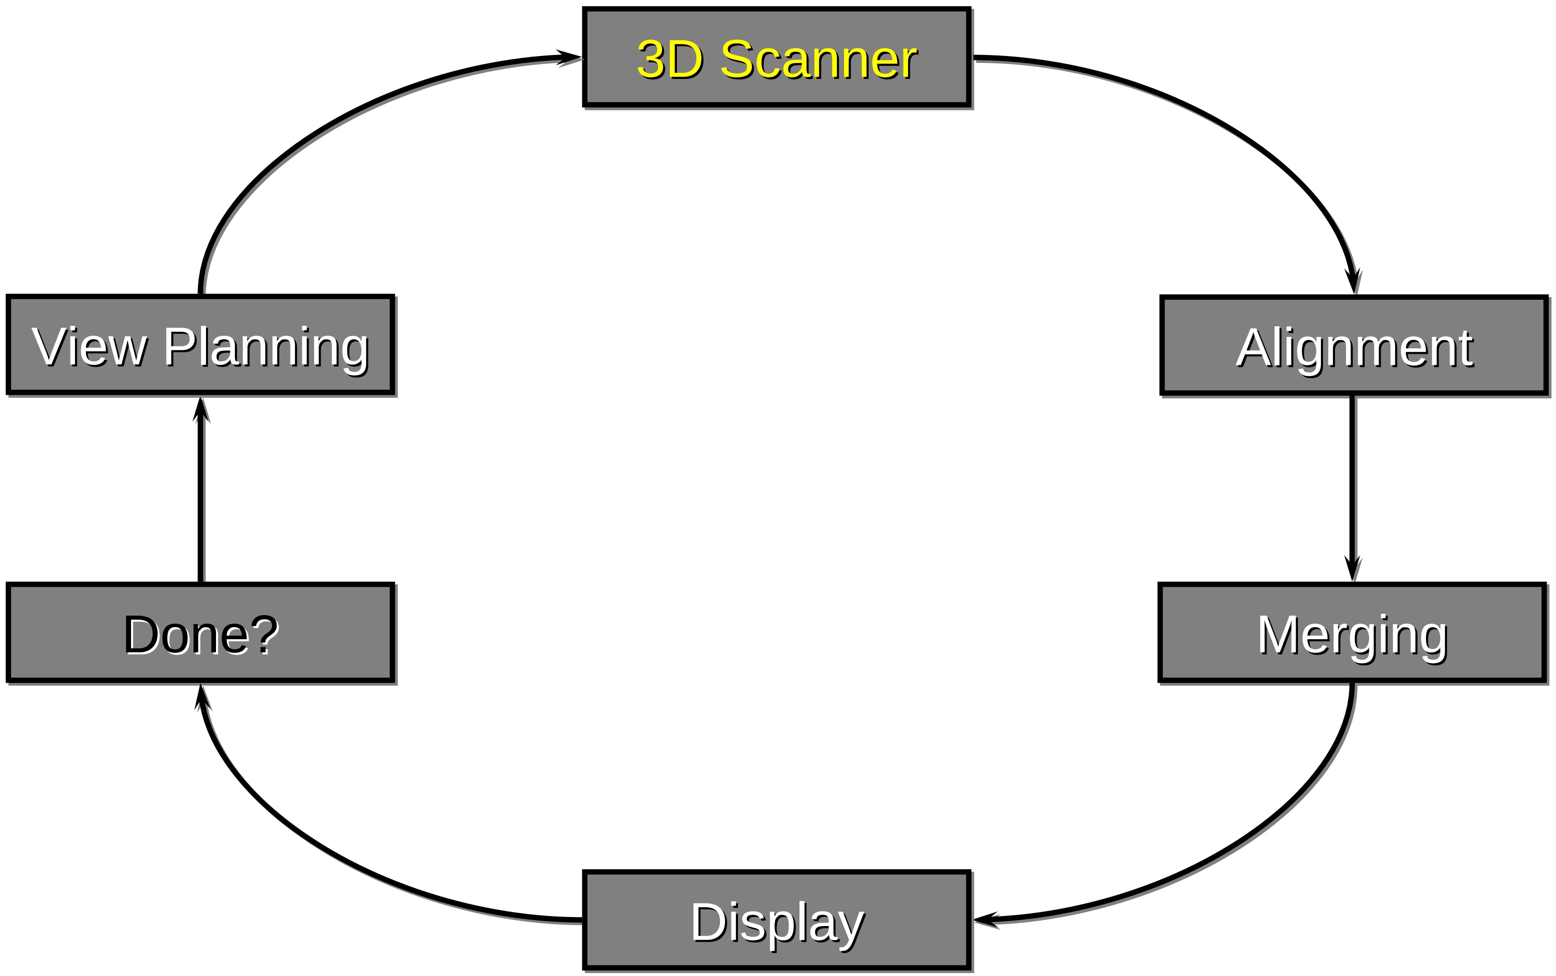

Triangulation

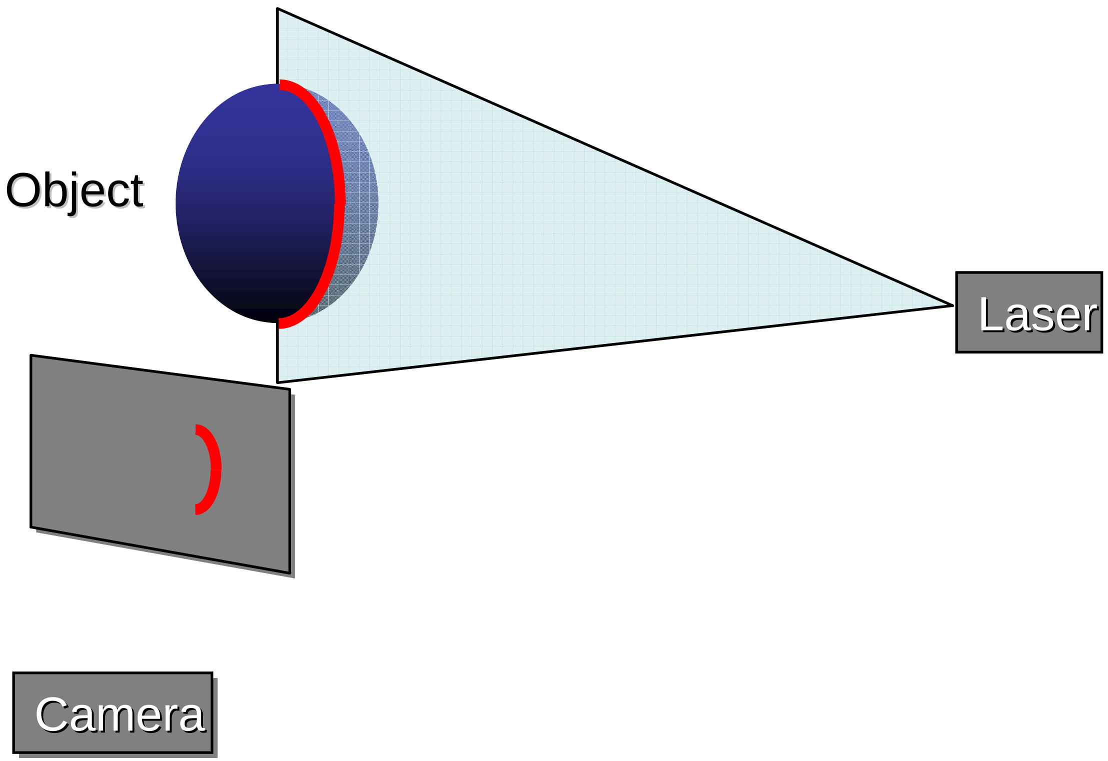

- Project laser stripe onto object

Triangulation

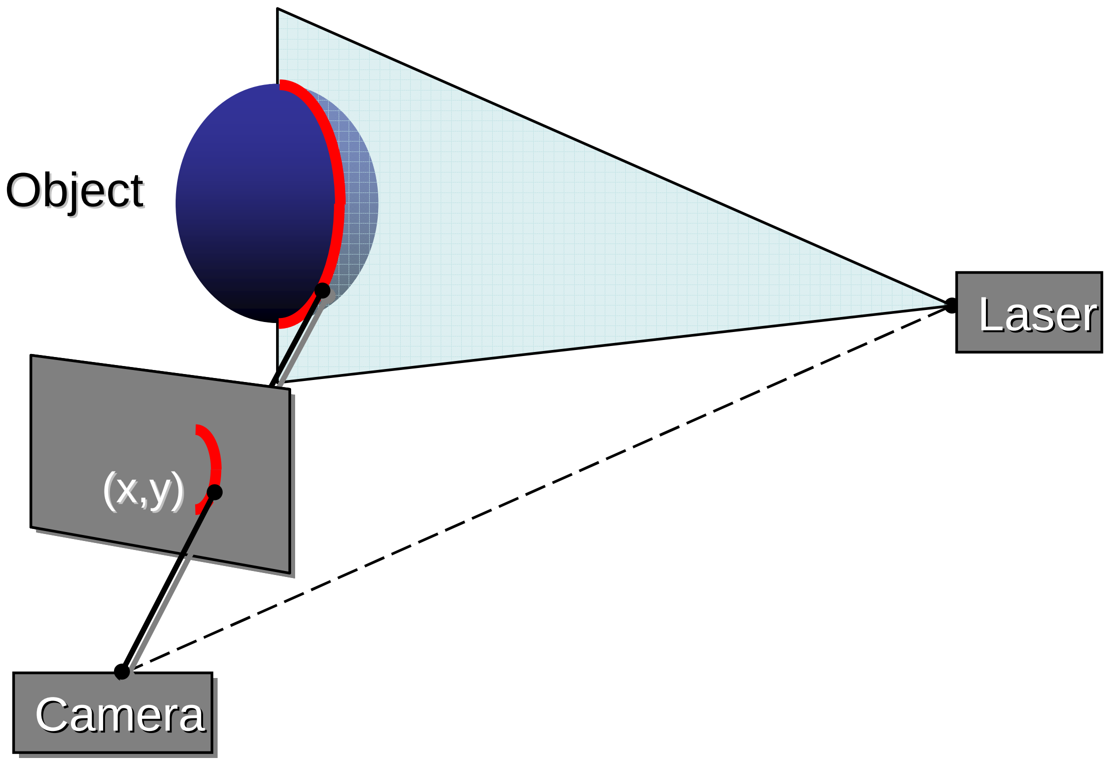

- Depth from ray-plane triangulation

Triangulation

- Faster acquisition: project multiple stripes
- Correspondence problem: which stripe is which?

Continuum of triangulation methods

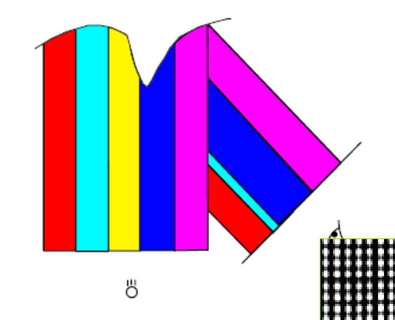

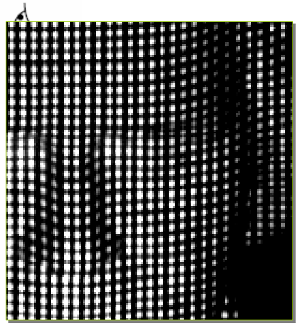

Single-stripe

Multi-stripe
Multi-frame

Single-frame

Slow, robust

Fast, fragile

Time-coded light patterns

- Assign each stripe a unique illumination code over time

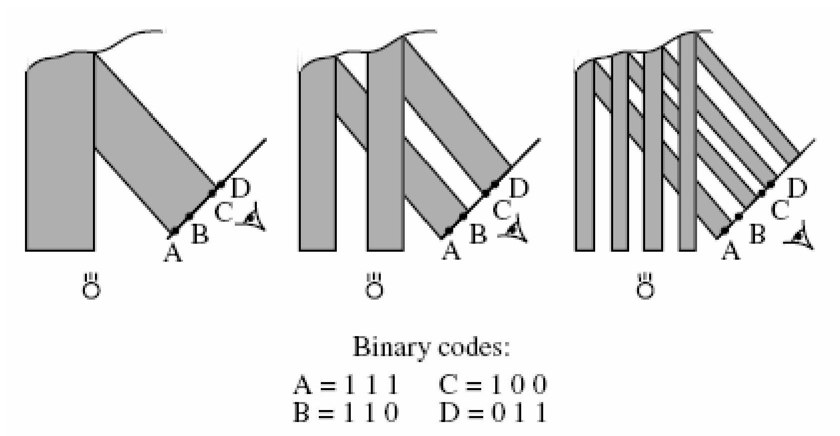

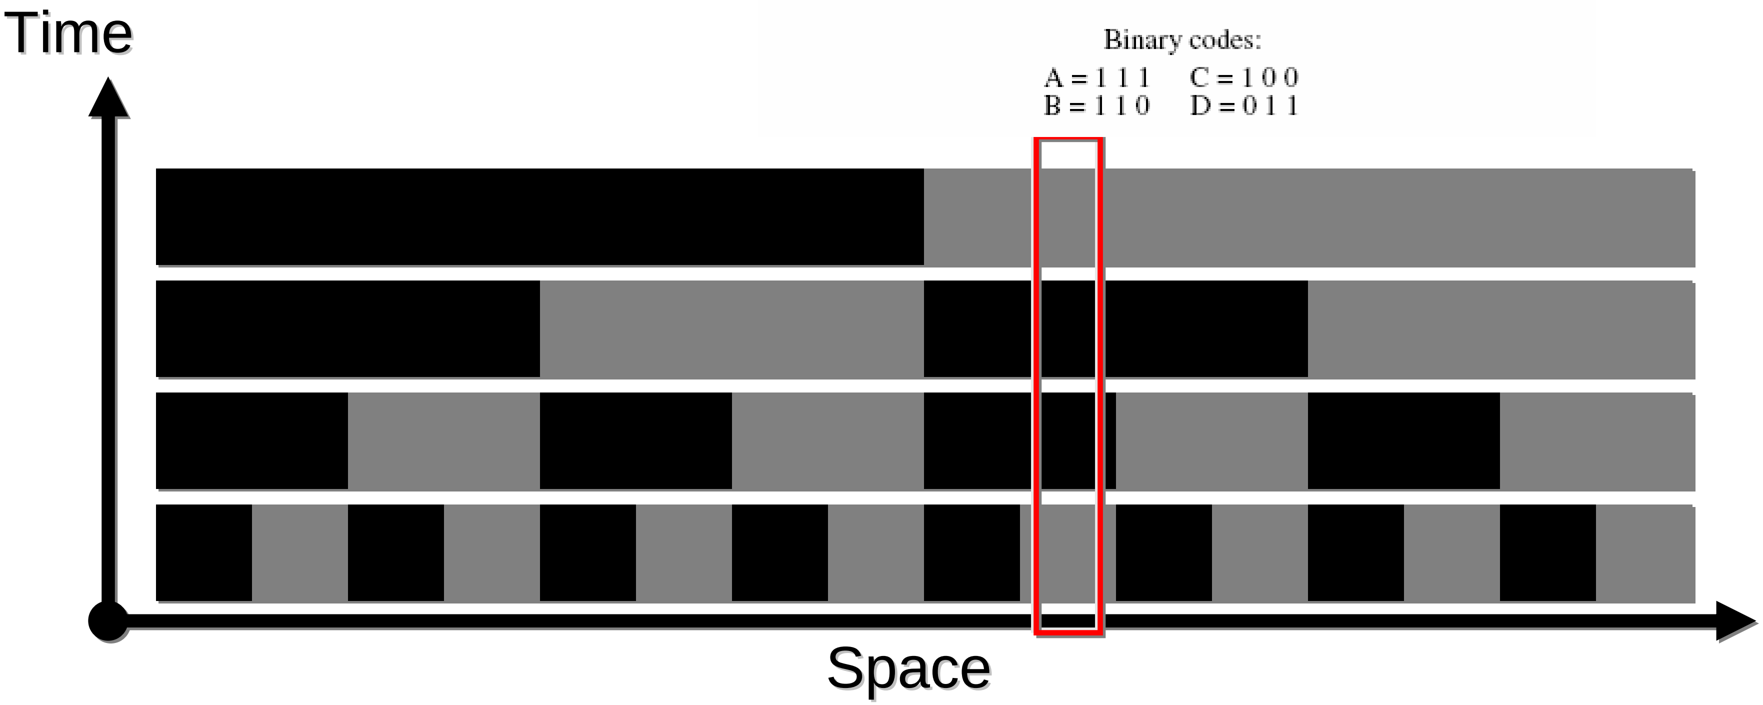

Codes for moving scenes

- Assign time codes to stripe boundaries
- Perform frame-to-frame tracking of corresponding boundaries
  - Propagate illumination history

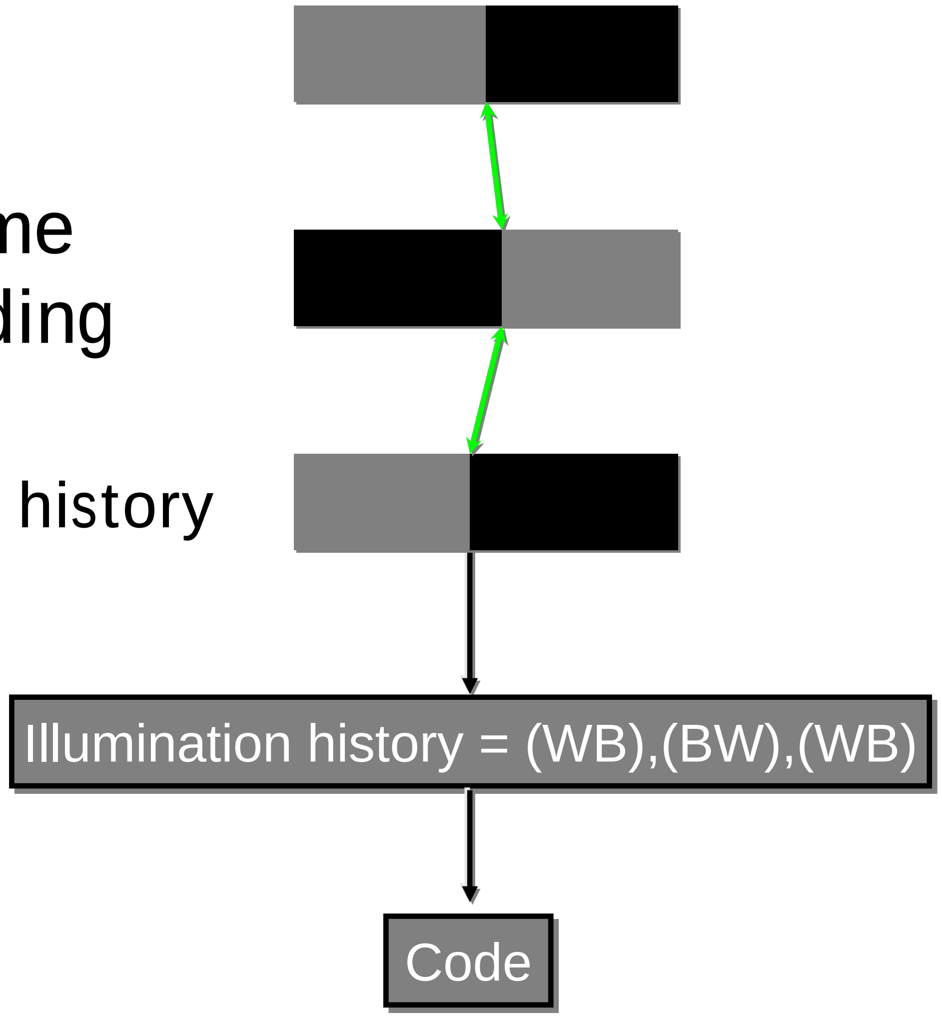

Codes for moving scenes

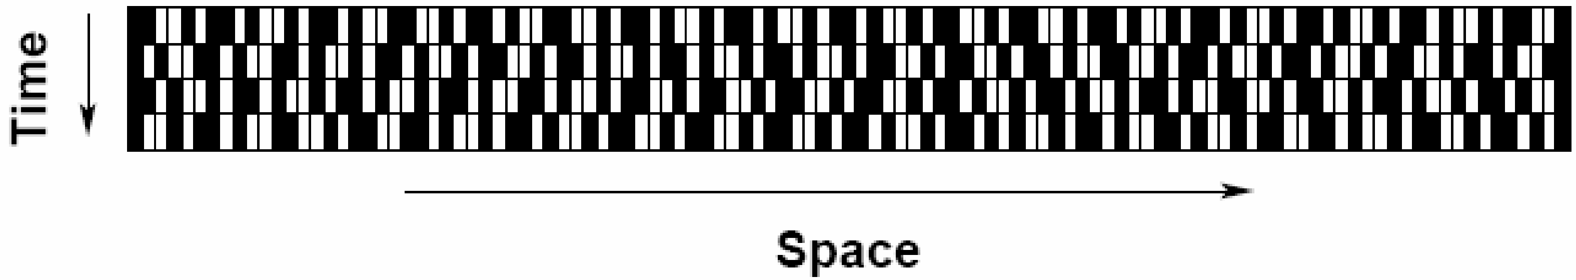

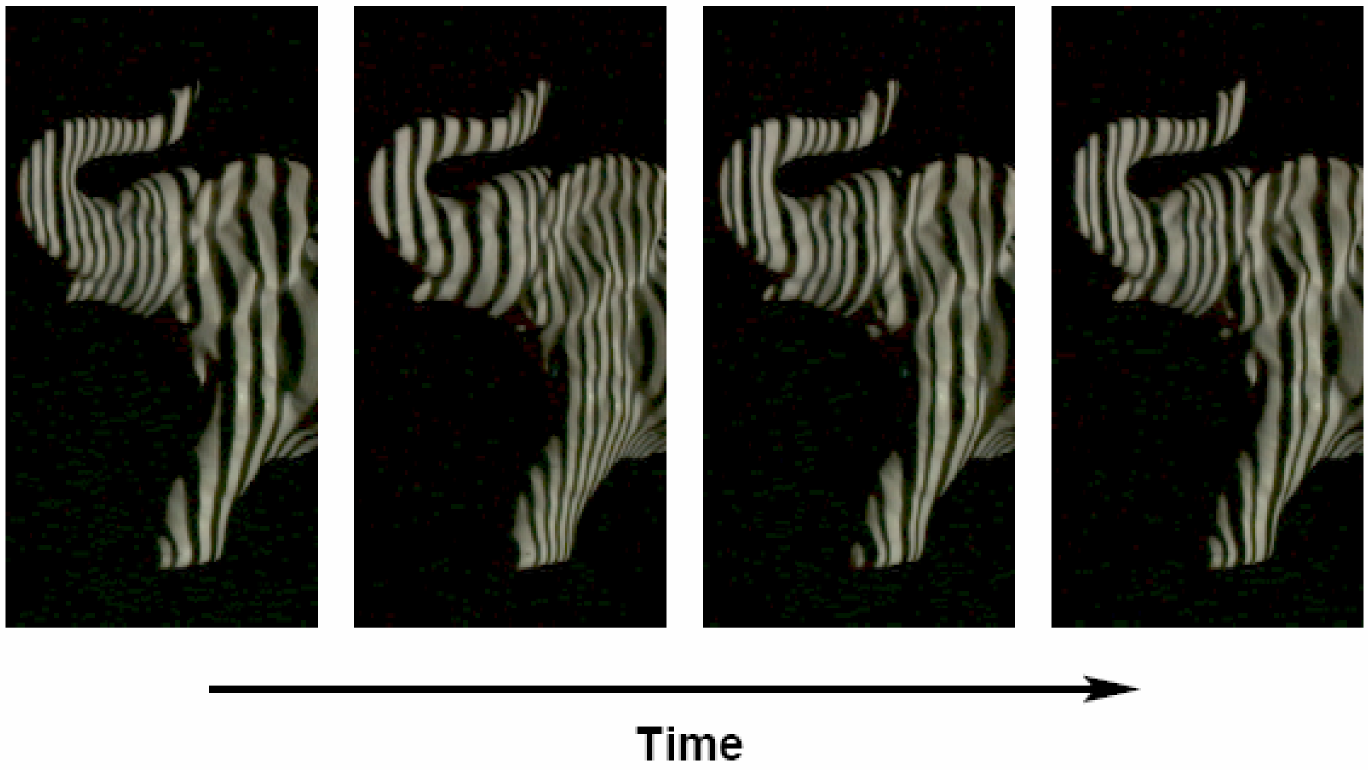

Real-time 3D model acquisition

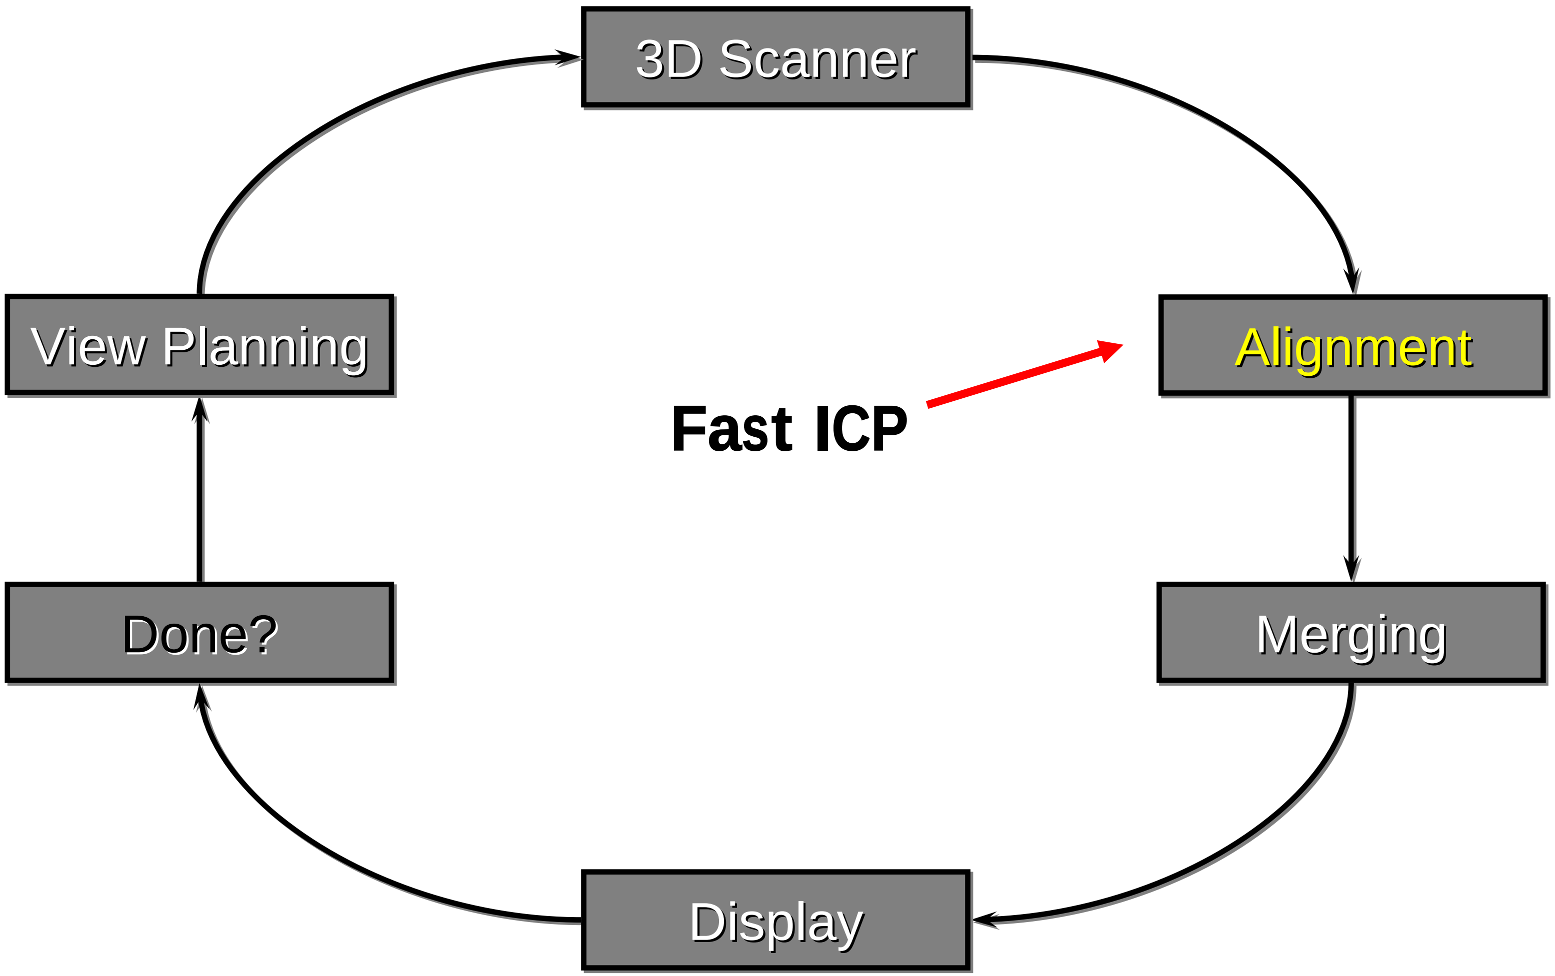

Real-time 3D model acquisition

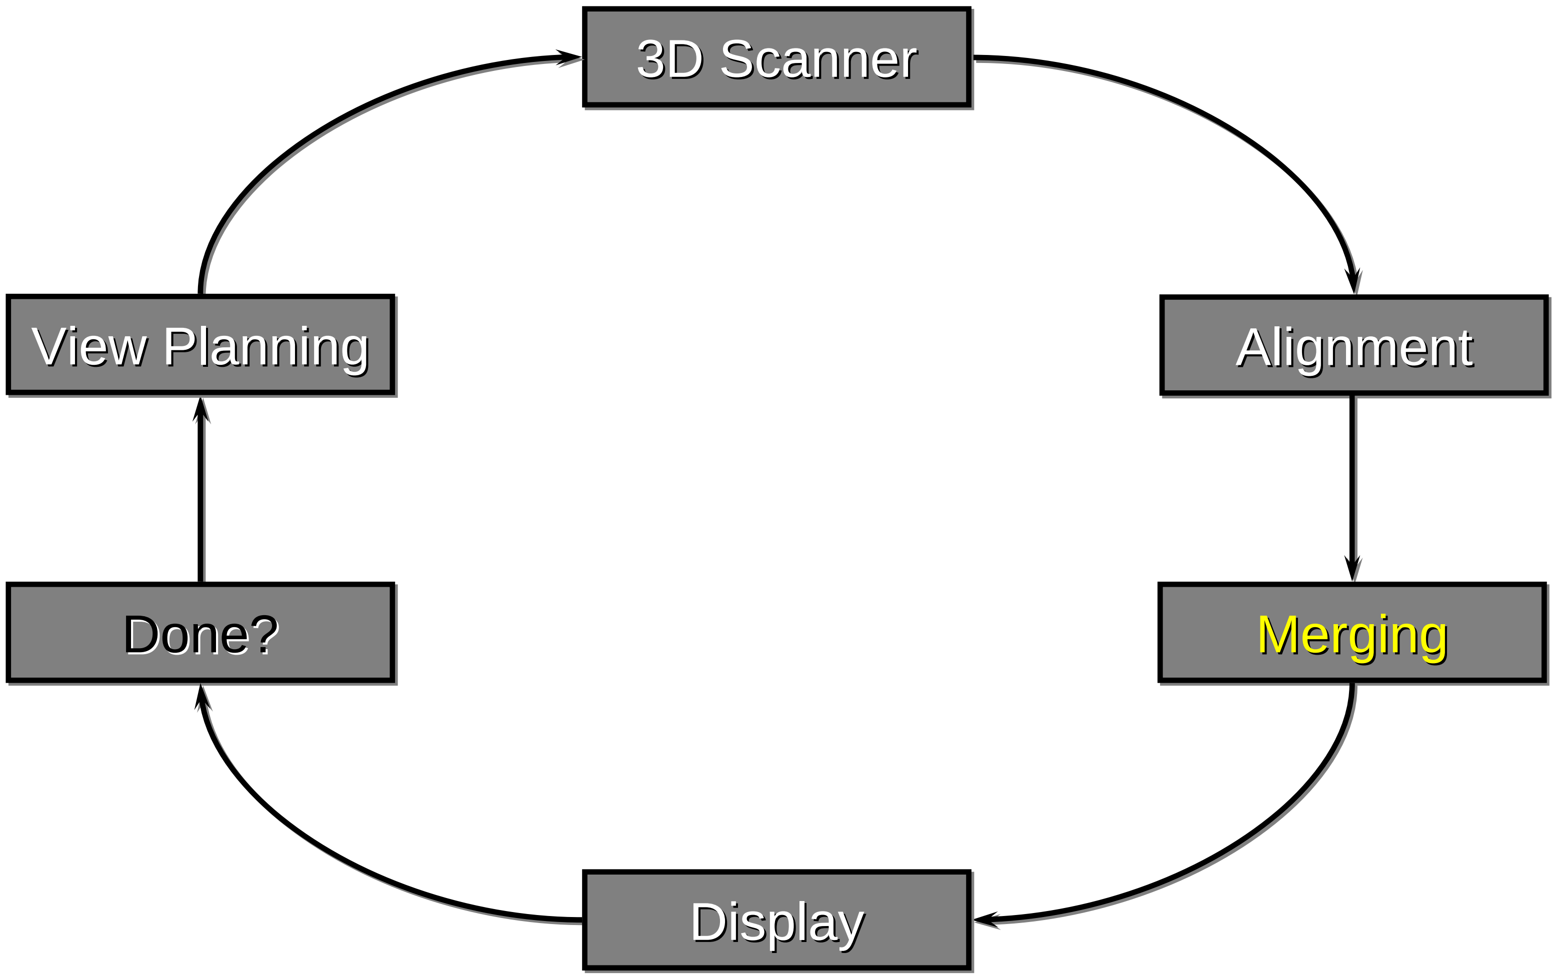

Merging and Rendering

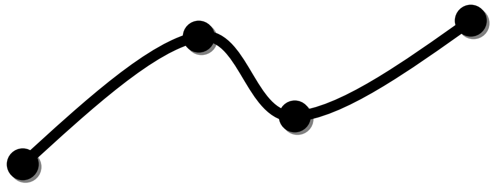

Merging and Rendering

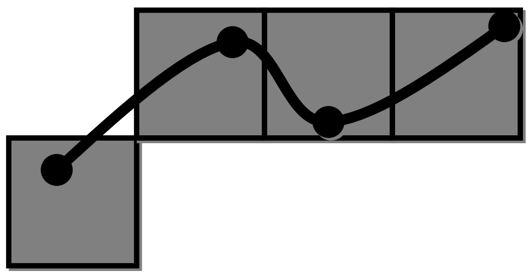

Merging and Rendering

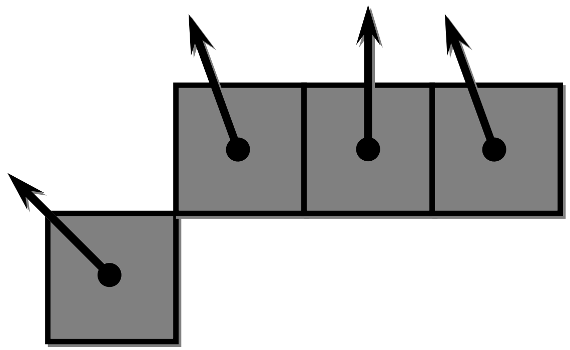

Merging and Rendering

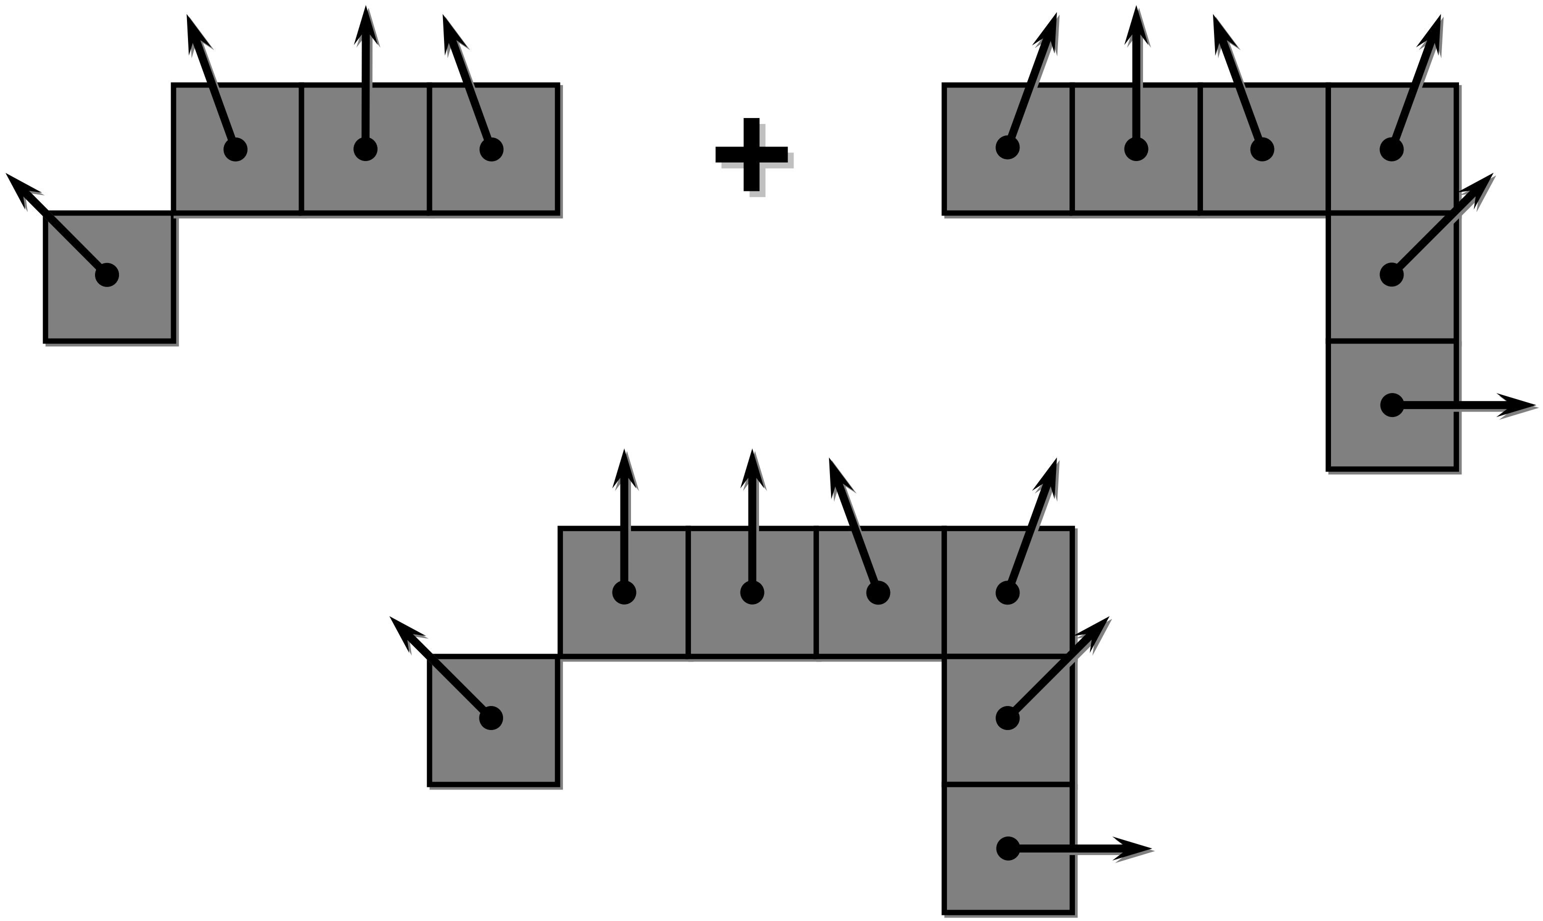

- Point rendering, using accumulated normals for lighting

Example: photograph

18 cm.

Result

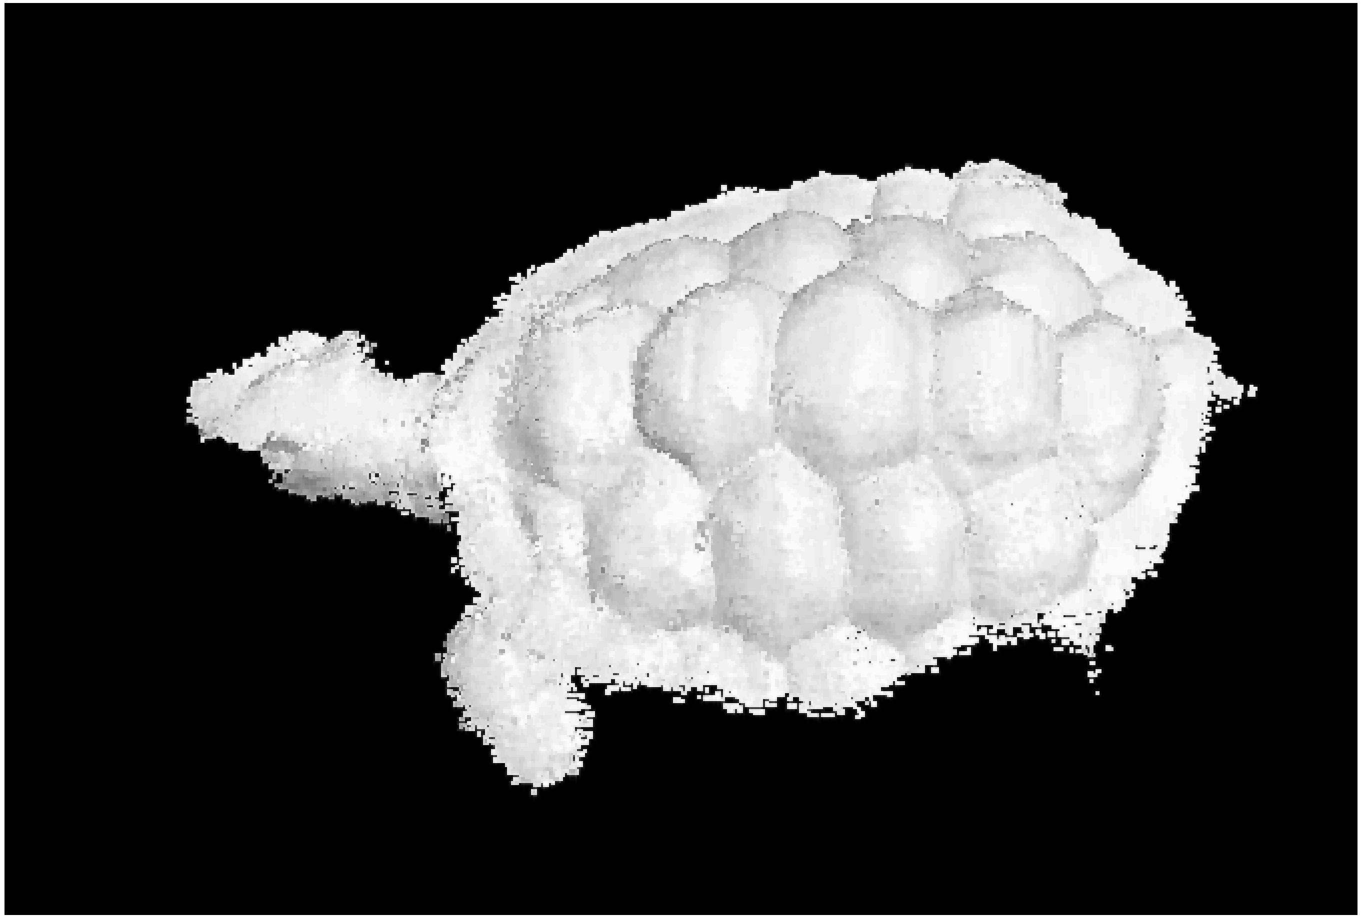

Postprocessed model

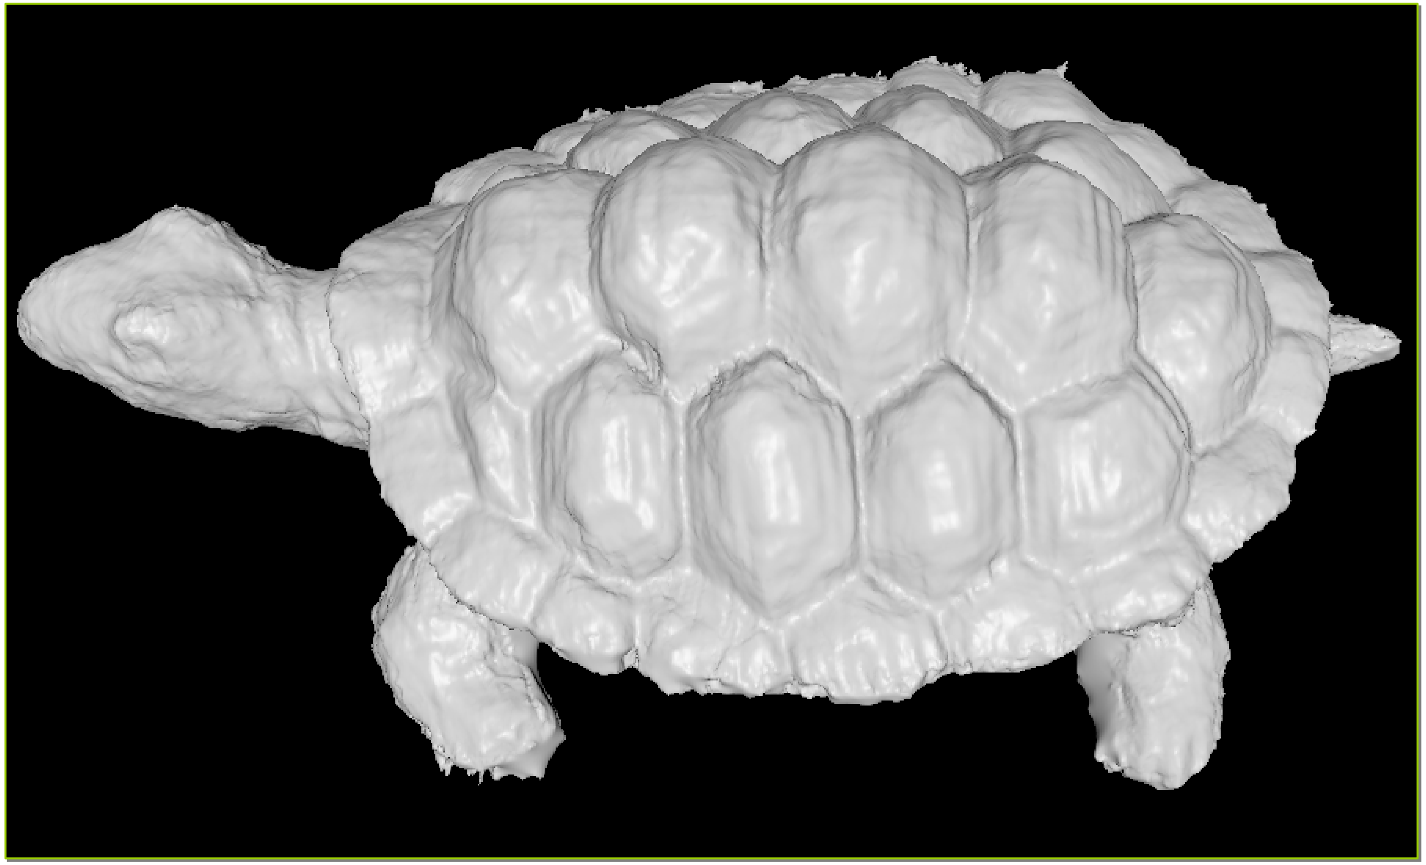

The Digital Michelangelo Project

- Goal: scan 10 sculptures by Michelangelo
- High-resolution ("quarter-millimeter") geometry
- Stanford University, led by Marc Levoy

Scanning the David

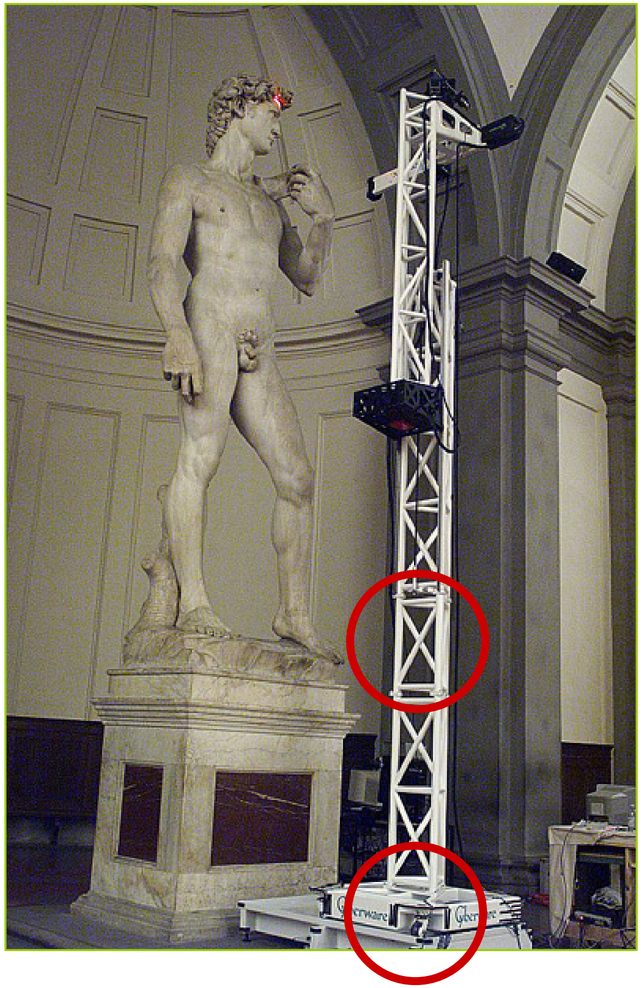

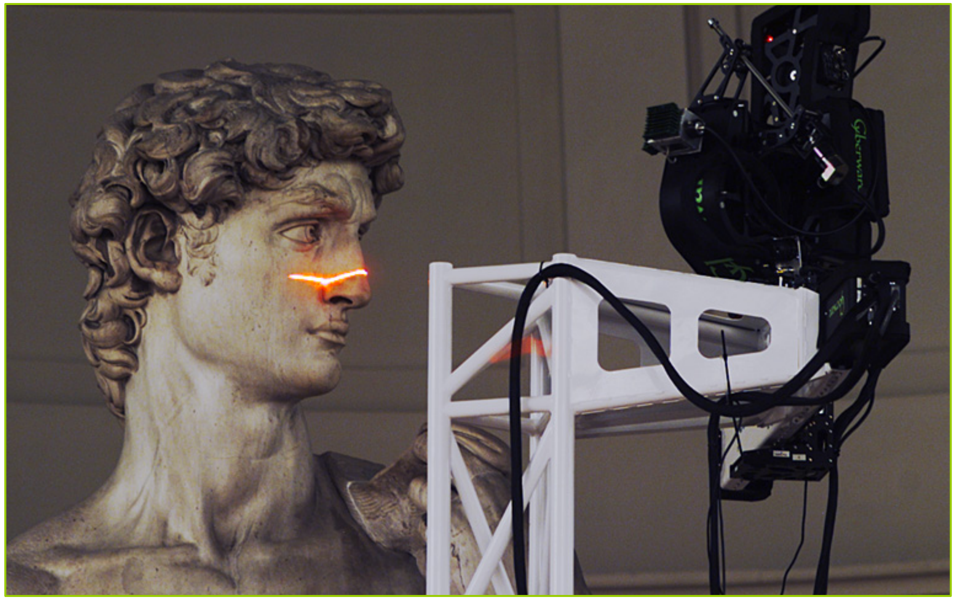

height of gantry: 7.5 meters
weight of gantry: 800 kilograms

Range processing pipeline

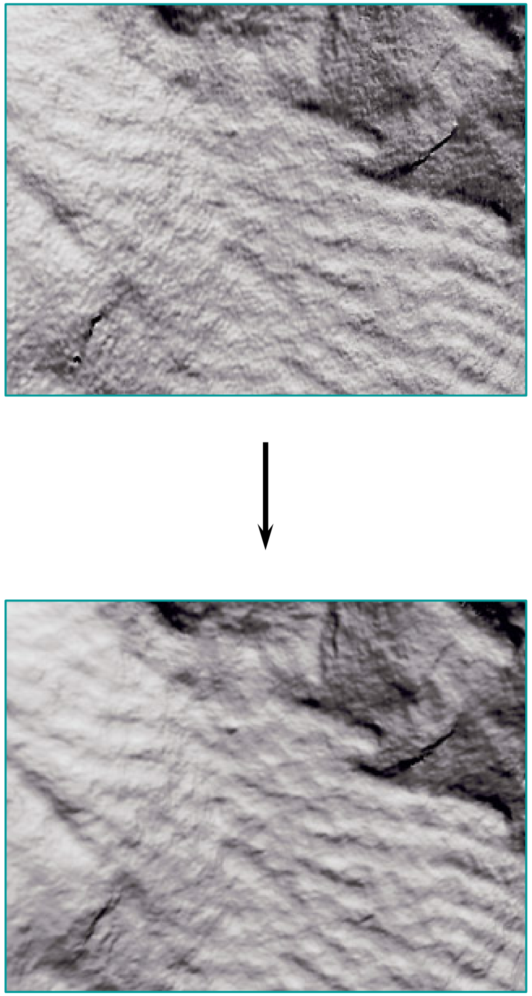

- steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Statistics about the scan

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

Comparison

photograph

1.0 mm computer model

Results

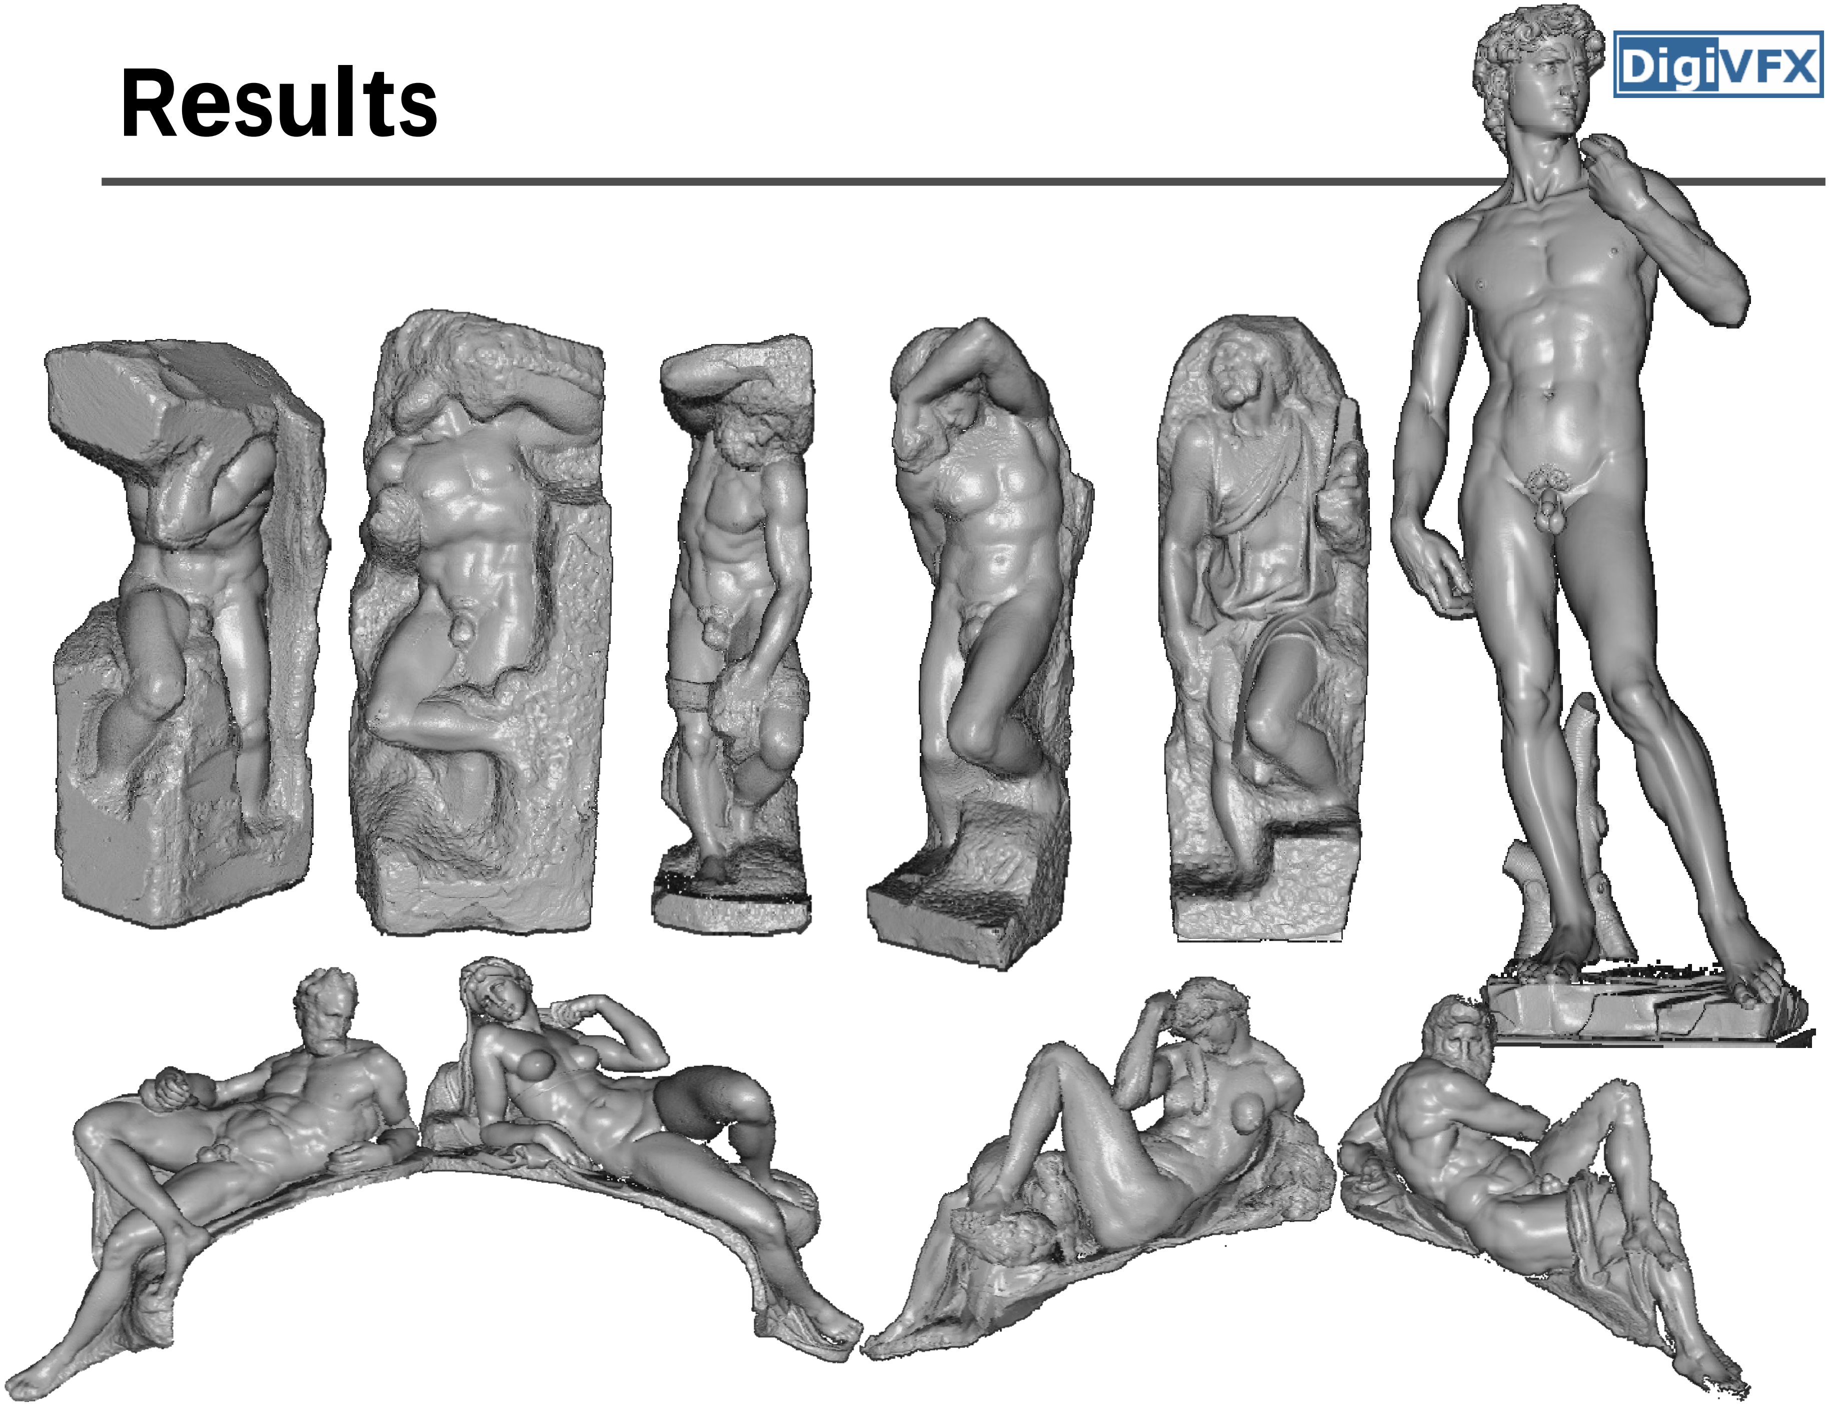

The Great Buddha Project

- 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
- Institute of Industrial Science, University of Tokyo, led by 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

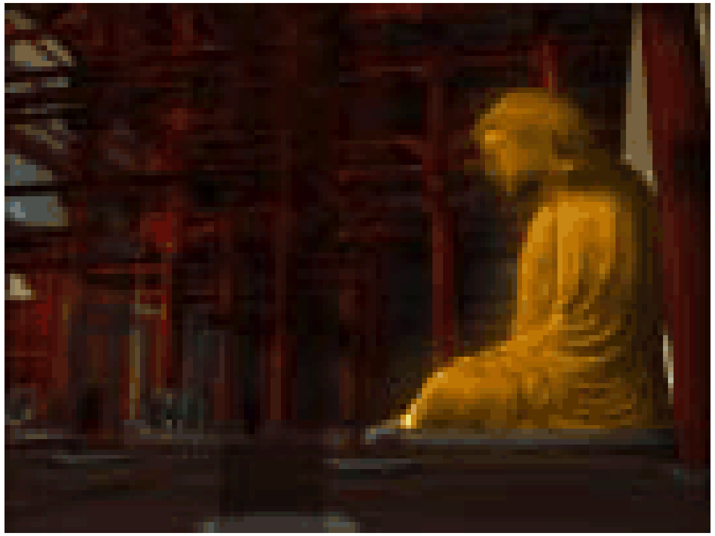
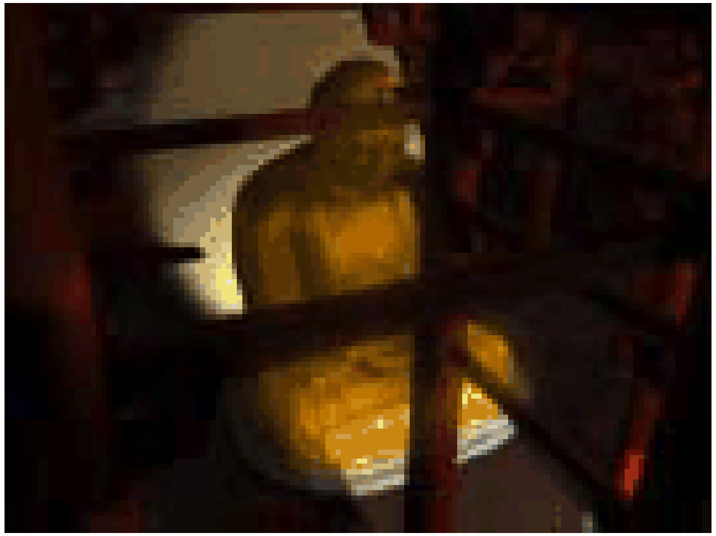
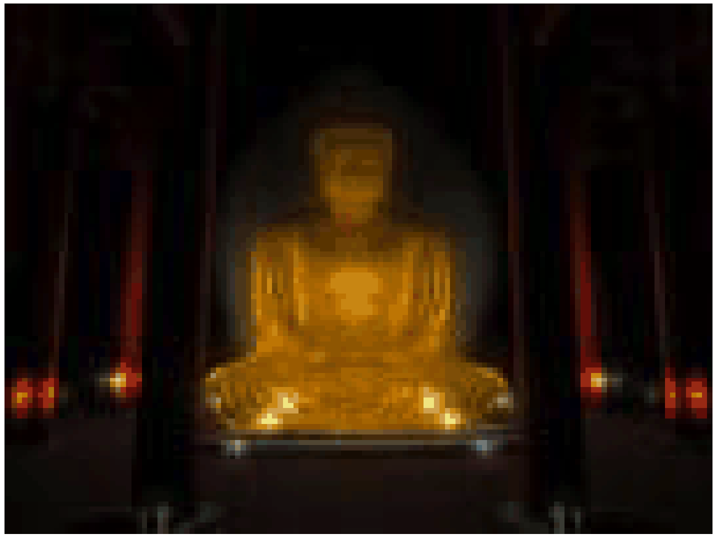
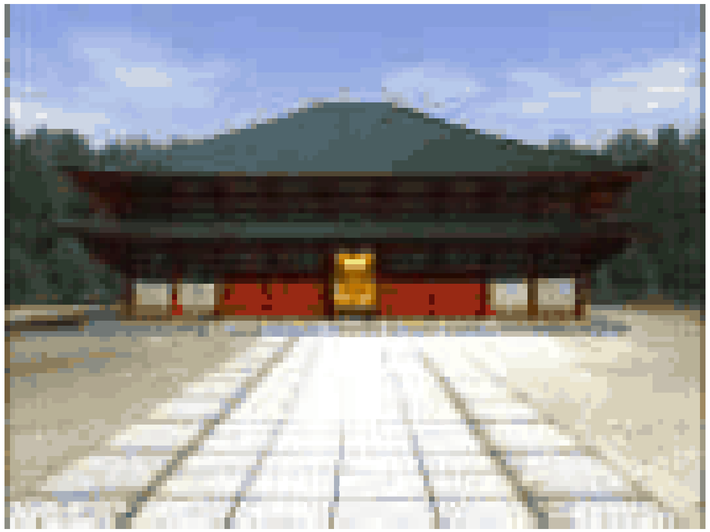

IMAX 3D

- 6K resolution, 42 linear bits per pixel
- For CG, it typically takes 6 hours for a frame
- 45-minute IMAX 3D CG film requires a 100-CPU rendering farm full-time for about a year just for rendering
- For live-action, camera is bulky (like a refrigerator)

Hybrid stereo camera

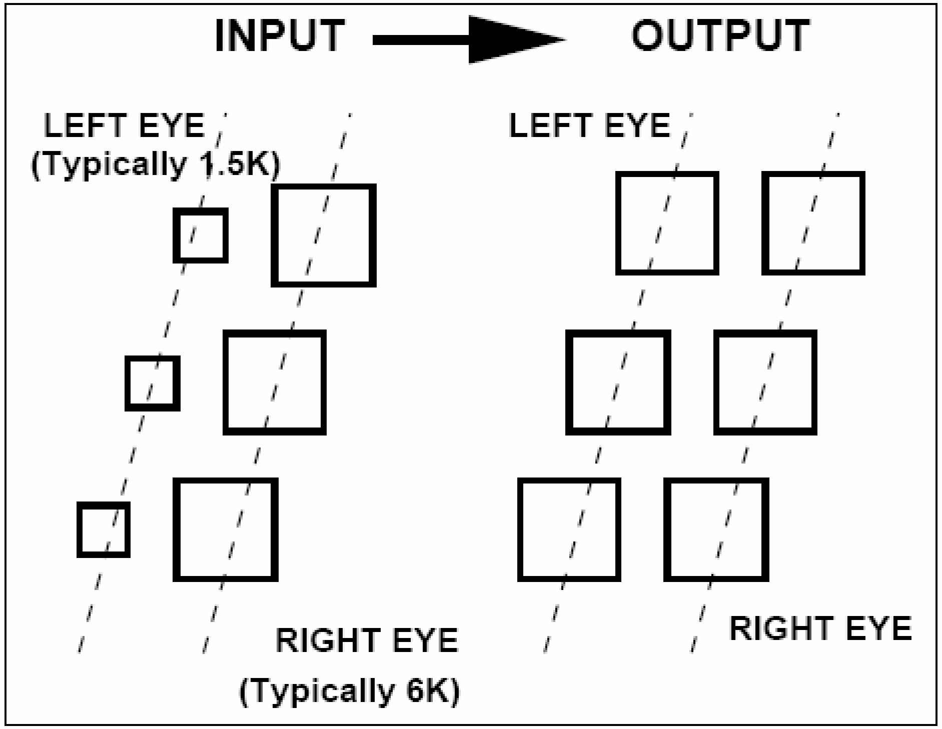

Live-action sequence

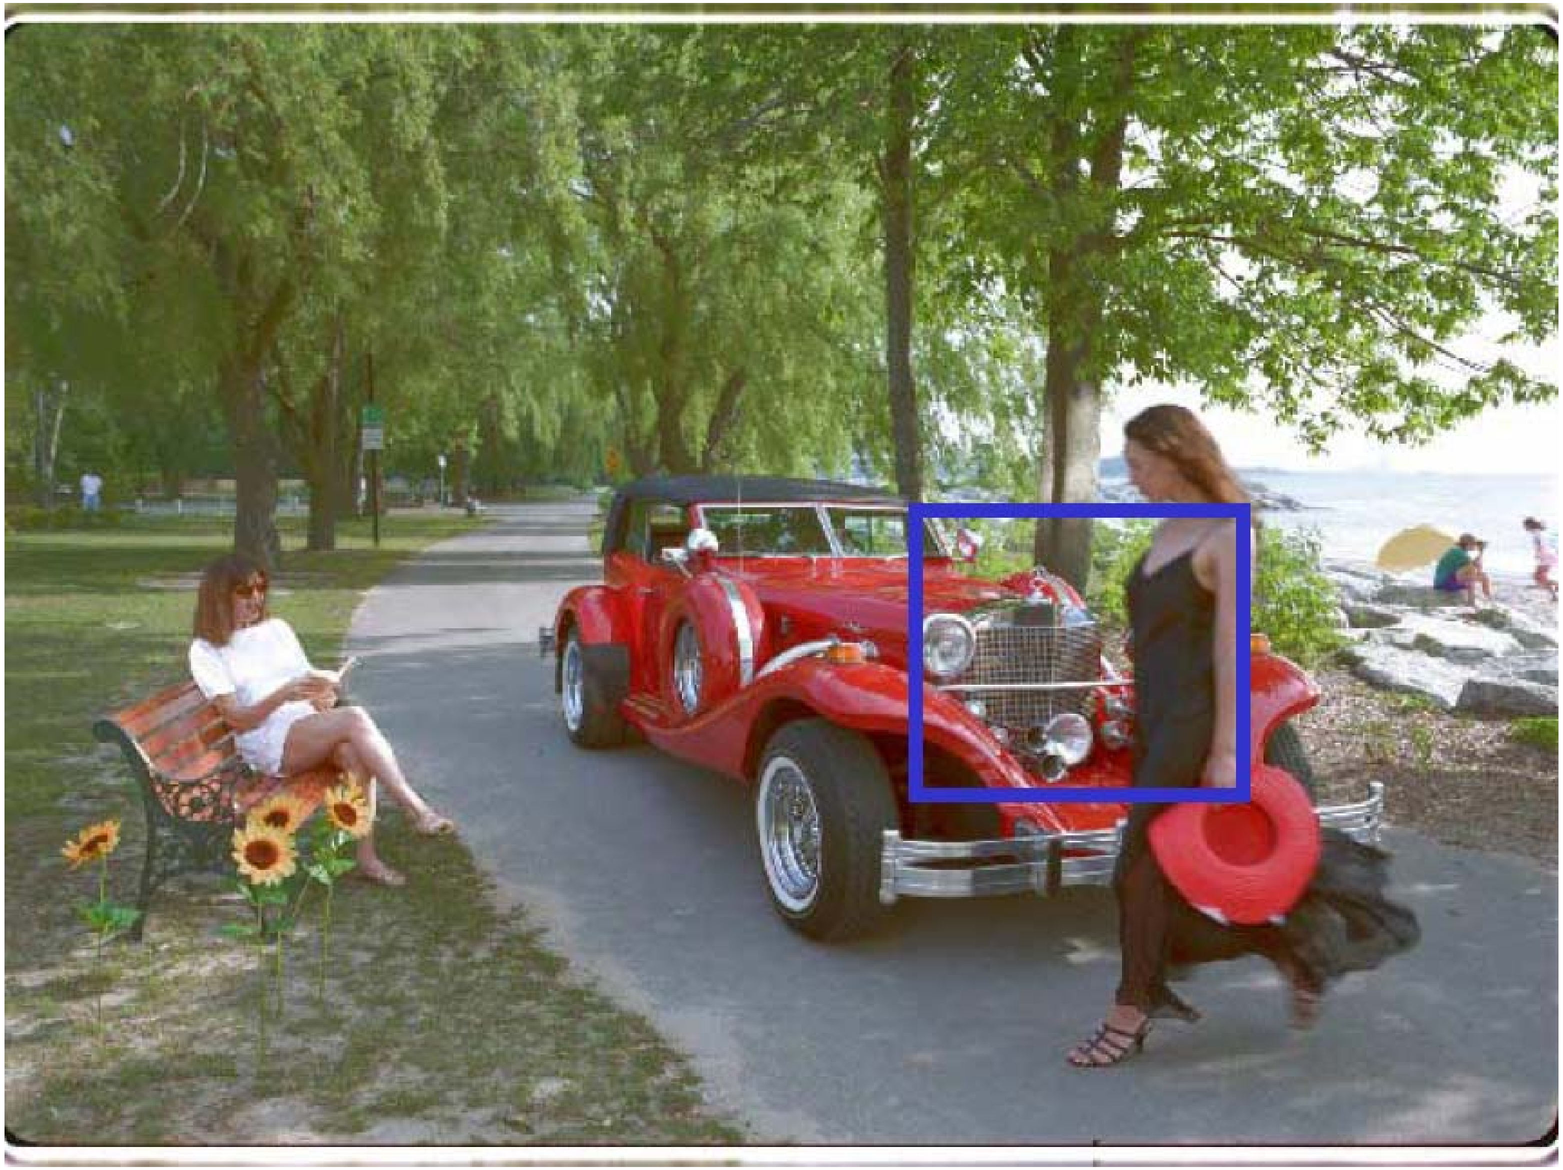

Hybrid input

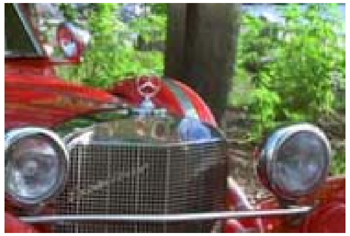

left

right

Hybrid input

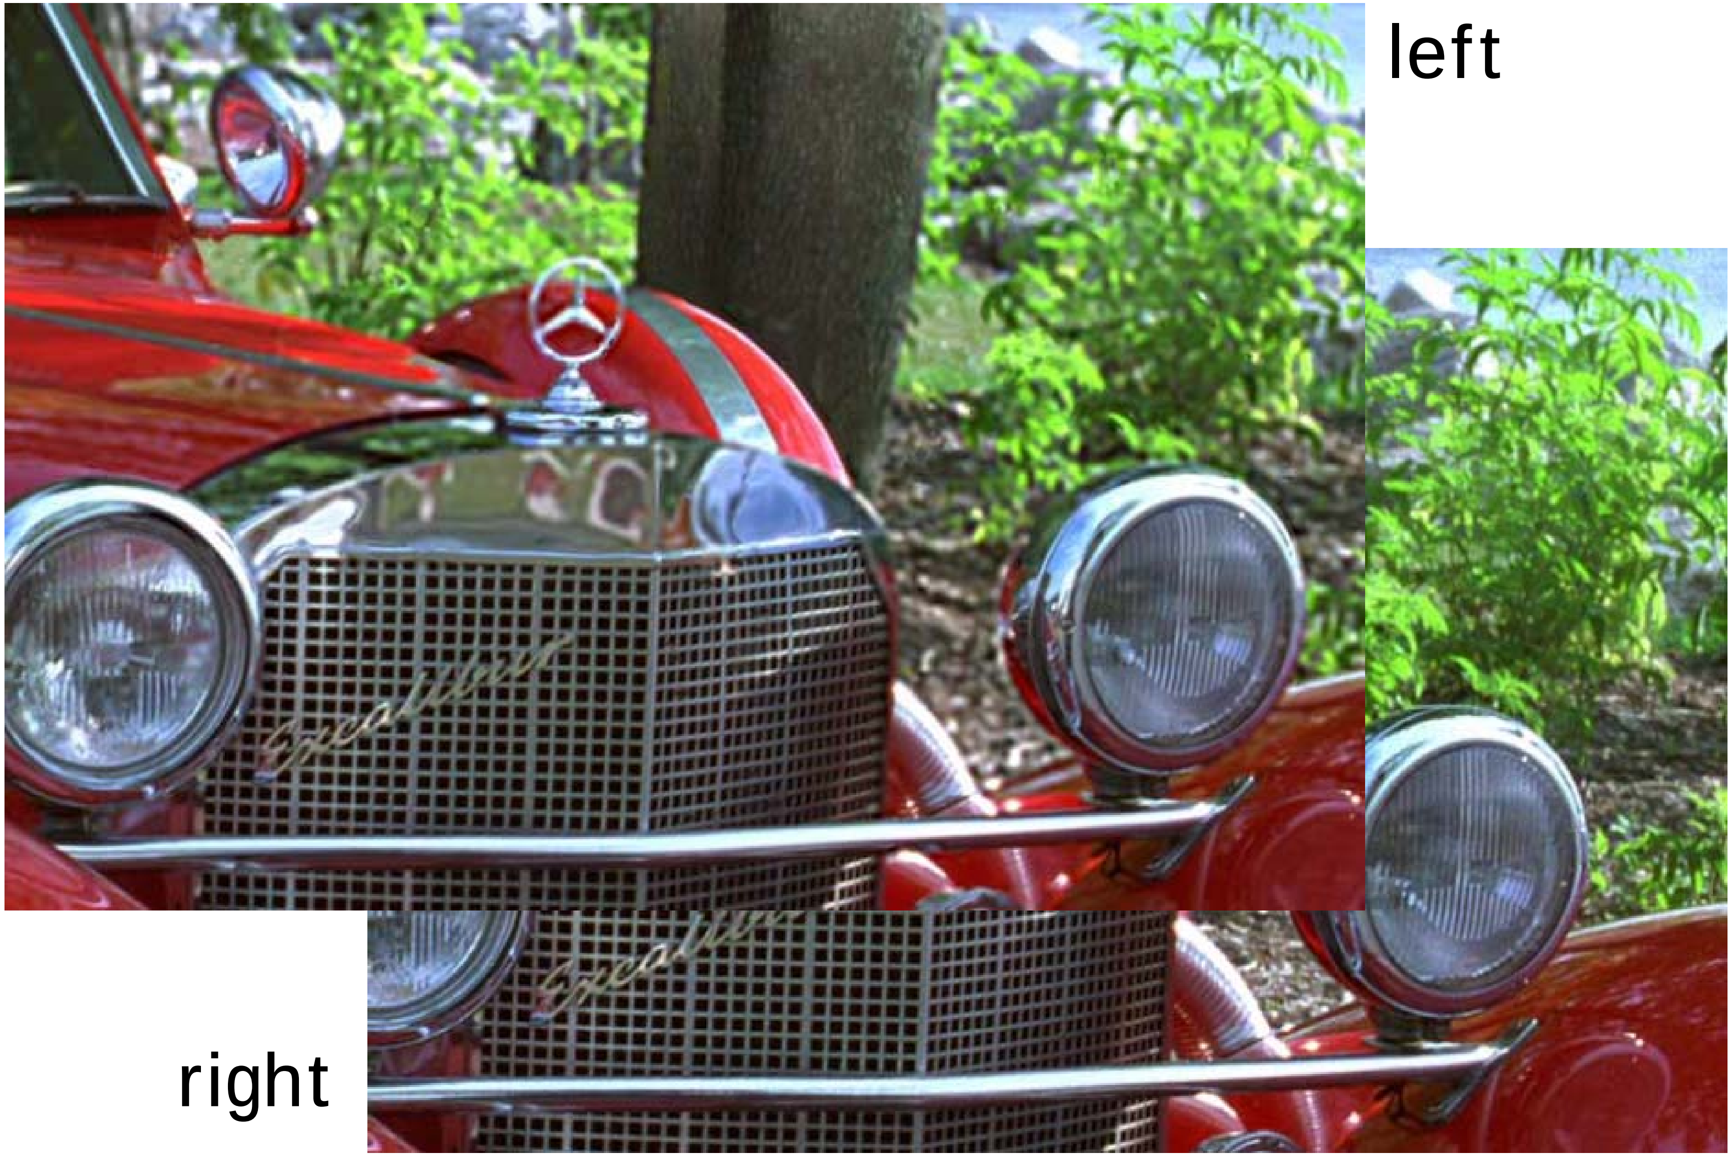

left

right

Combine multiple hires to lores

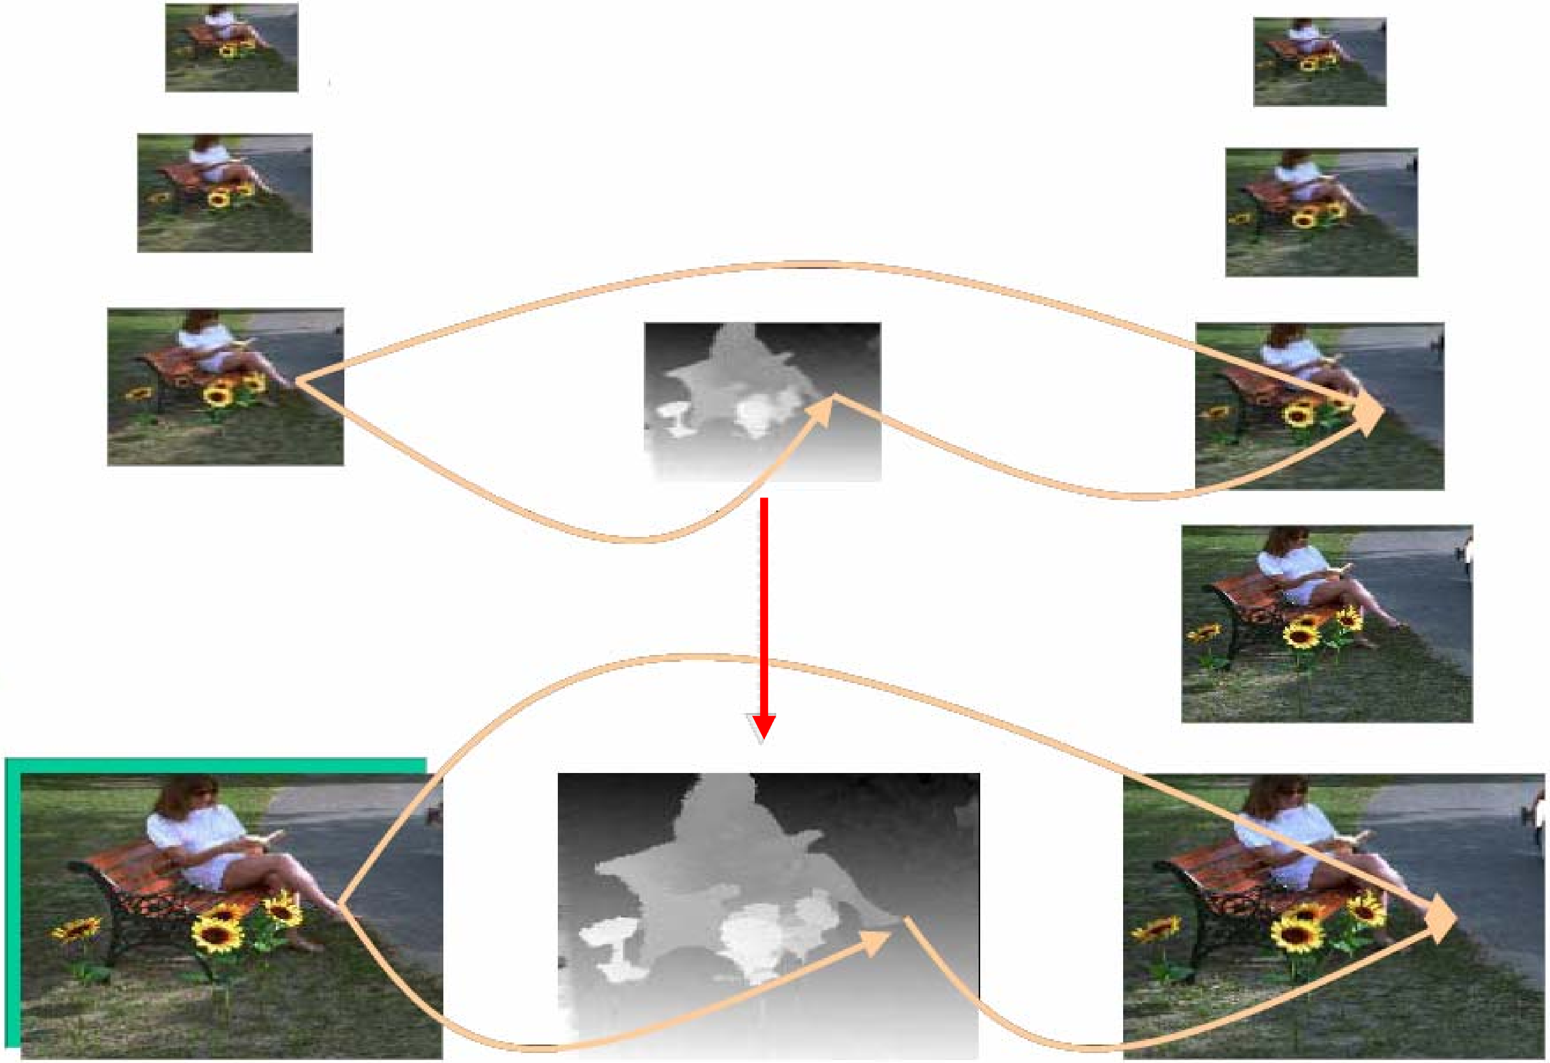

View interpolation

Bullet time video

View interpolation

High-quality video view interpolation

Final project

Final project

- Assigned: today
- Due: 6/28 Tuesday
- Final presentation: 6/28 Tuesday, 1:30pm?
- Proposal presentation next week, send me slides by 1:00pm next Wednesday if you want to use my laptop for your presentation
- Send me the topic and team members by next Tuesday

Final project

- Research (1-2 people)
- System (1-3 people)
- Film (3-4 people)

Research

- Define a problem and try to solve it
- You don't need to solve it all, but have to make a reasonable progress, for example, solve a simplified version.
- Find inspirations from SIGGRAPH/CVPR/ICCV papers

System

- Implement existing algorithm into a useful system

Film

- It must be an “effect” film.
- You can use any tools as you want. But, I guess that you have to write some on your own.
- Find inspirations from

  Gatech’s vfx course

  http://www.cc.gatech.edu/classes/AY2004/cs4480_spring/

  independent film makers

  http://www.peerlessproductions.com/
- Submit two videos, final and making-of.

Reference

- S. Rusinkiewicz and M. Levoy. Efficient Variants of the ICP Algorithm, 3DIM 2001.
- B. Curless and M. Levoy. A Volumetric Method for Building Complex Models from Range Images, SIGGRAPH 1996.
- S. Rusinkiewicz and O. Hall-Holt and M. Levoy. Real-Time 3D Model Acquisition, SIGGRAPH 2001.
- H. Sawhney, Y. Guo, K. Hanna, R. Kumar, S. Adkins and S. Zhou. Hybrid Stereo Camera: An IBR Approach for Synthesis of Very High Resolution Stereoscopic Image Sequences, SIGGRAPH 2001.
- C. L. Zitnick, S. B. Kang, M. Uyttendaele, S. Winder and R. Szeliski. High-quality video view interpolation using a layered represenation, SIGGRAPH 2004.