3D Scanning

Szymon Rusinkiewicz
Princeton University
COS 426 Guest Lecture
Spring 2003

Computer Graphics Pipeline

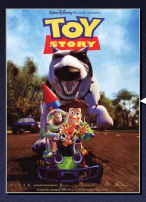

Rendering ← Shape ← 3D Scanning
Rendering ← Motion
Rendering ← Lighting and Reflectance

- Human time = expensive
- Sensors = cheap
  - Computer graphics increasingly relies on measurements of the real world

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

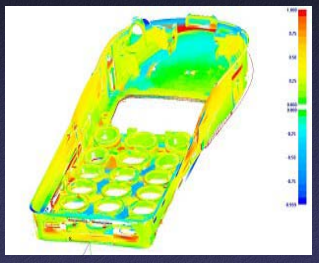

Medicine

- Plan surgery on computer model, visualize in real time

Medicine

- Plan surgery on computer model, visualize in real time

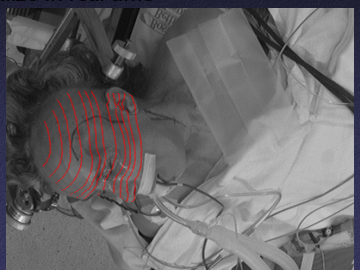

Medicine

- Plan surgery on computer model, visualize in real time

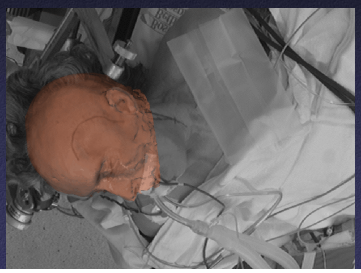

Medicine

- Plan surgery on computer model, visualize in real time

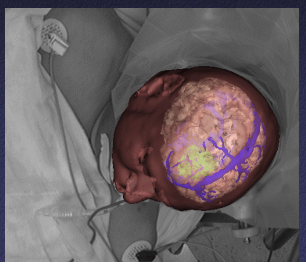

Scanning Buildings

- Quality control during construction
- As-built models

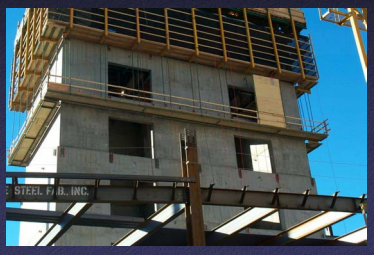

Scanning Buildings

- Quality control during construction
- As-built models

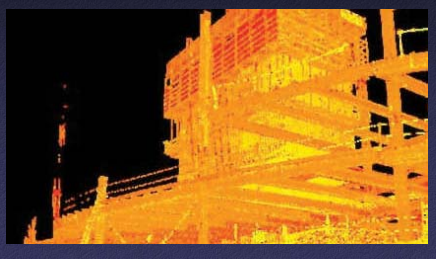

Clothing

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

The Digital Michelangelo Project

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 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?

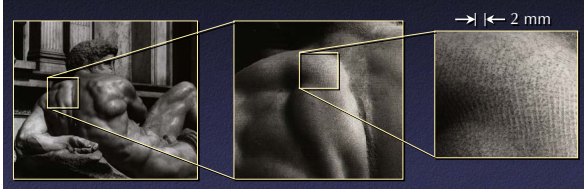

Day (Medici Chapel)

Side project: The Forma Urbis Romae

Forma Urbis Romae Fragment

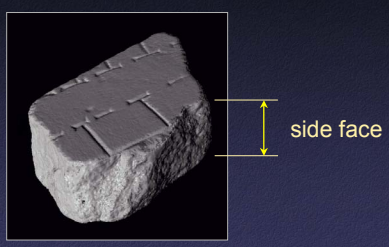

Range Acquisition Taxonomy

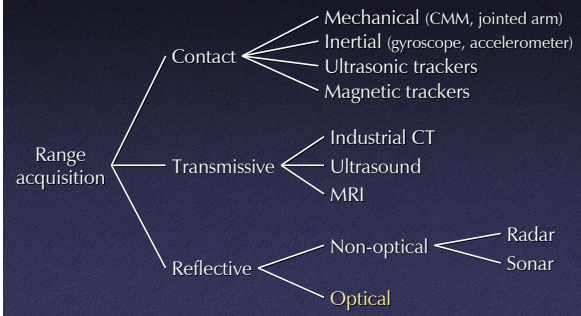

Range Acquisition Taxonomy

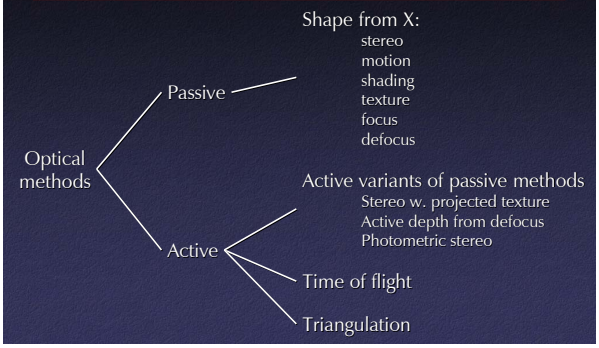

Touch Probes

- Jointed arms with angular encoders
- Return position, orientation of tip

Faro Arm – Faro Technologies, Inc.

Stereo

- Find feature in one image, search along epipolar line in other image for correspondence

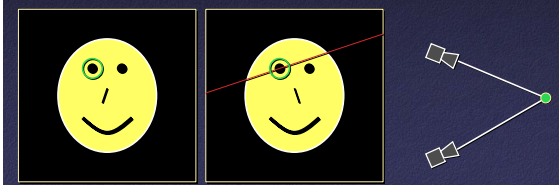

Why More Than 2 Views?

- Baseline
  - Too short – low accuracy
  - Too long – matching becomes hard

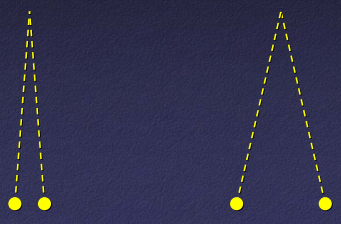

Why More Than 2 Views?

- Ambiguity with 2 views

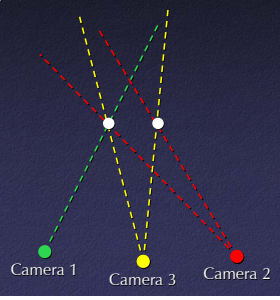

Multibaseline Stereo

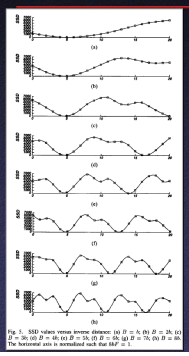

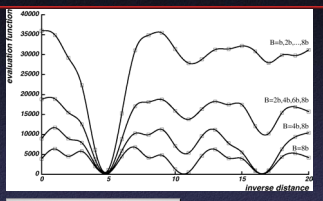

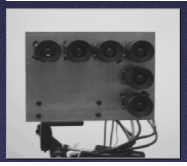

[Okutomi & Kanade]

Shape from Motion

- "Limiting case" of multibaseline stereo
- Track a feature in a video sequence
- For n frames and f features, have $2 \cdot n \cdot f$ knowns, $6 \cdot n + 3 \cdot f$ unknowns

Shape from Shading

- Given: image of surface with known, constant reflectance under known point light
- Estimate normals, integrate to find surface
- Problem: ambiguity

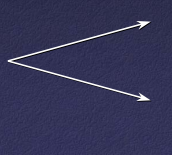

Shape from Shading

- Advantages:
  - Single image
  - No correspondences
  - Analogue in human vision
- Disadvantages:
  - Mathematically unstable
  - Can't have texture
- "Photometric stereo" (active method) more practical than passive version

Shape from Texture

- Mathematically similar to shape from shading, but uses stretch and shrink of a (regular) texture

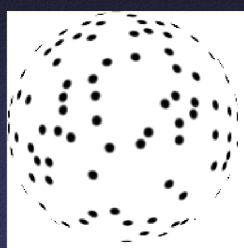

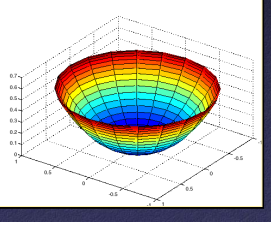

Shape from Focus and Defocus

- Shape from focus: at which focus setting is a given image region sharpest?
- Shape from defocus: how out-of-focus is each image region?
- Passive versions rarely used
- Active depth from defocus can be made practical

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

Pulsed Time of Flight

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

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

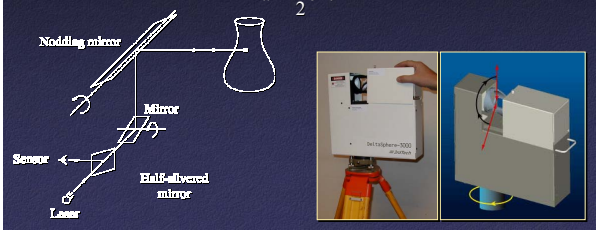

Pulsed Time of Flight

- 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.

Triangulation

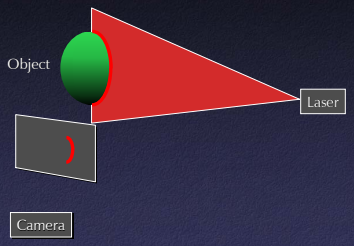

- Project laser stripe onto object

Triangulation

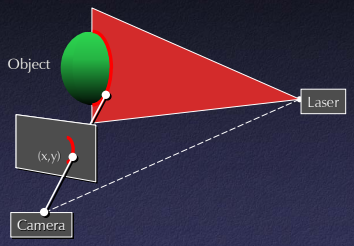

- Depth from ray-plane triangulation

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

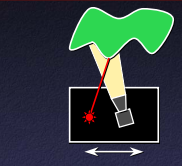


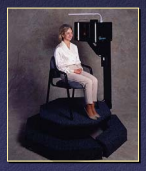

Triangulation: Moving the Camera and Illumination


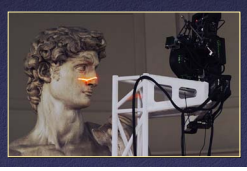
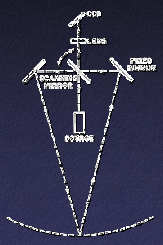

Scanning a Large Object

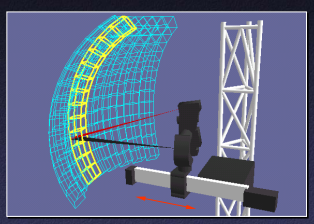

- Calibrated motions
  - pitch (yellow)
  - pan (blue)
  - horizontal translation (orange)
- Uncalibrated motions
  - vertical translation
  - rolling the gantry
  - remounting the scan head

Range Processing Pipeline

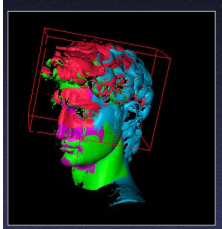

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Range Processing Pipeline

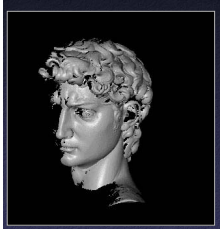

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Range Processing Pipeline

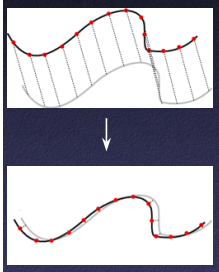

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Range Processing Pipeline

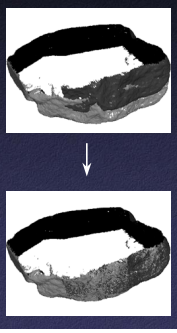

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Range Processing Pipeline

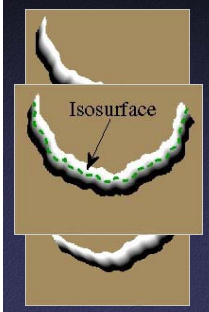

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

Range Processing Pipeline

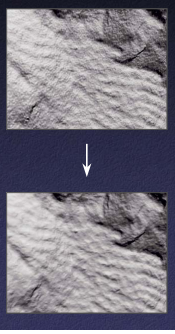

- Steps
  1. manual initial alignment
  2. ICP to one existing scan
  3. automatic ICP of all overlapping pairs
  4. global relaxation to spread out error
  5. merging using volumetric method

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