3D Photography

Sensors

<Some slides by Szymon Rusinkiewicz, Princeton University>

Ioannis Stamos

Perspective projection

Pinhole & the Perspective Projection

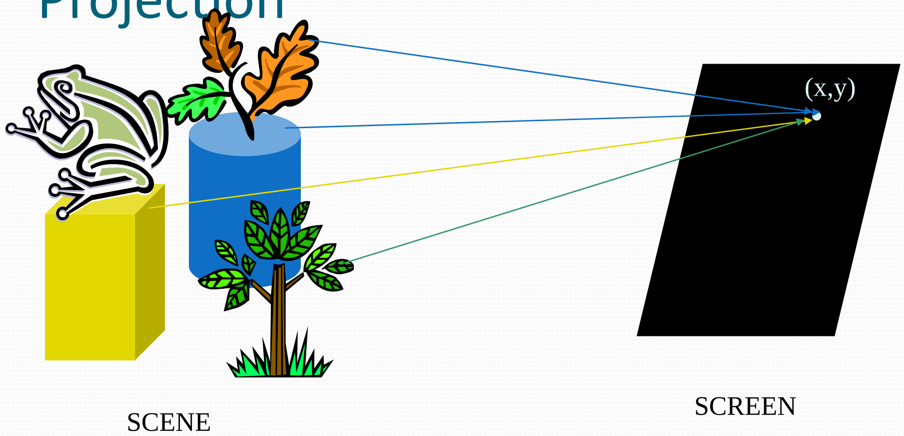

Is there an image being formed on the screen?

Pinhole Camera

- "Camera obscura" – known since antiquity

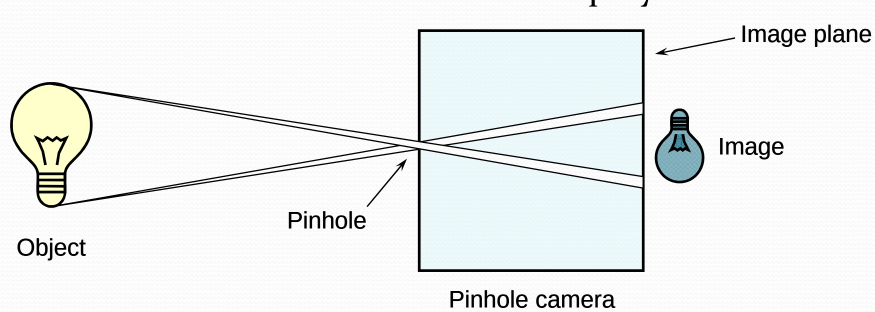

Perspective Camera

From Trucco & Verri

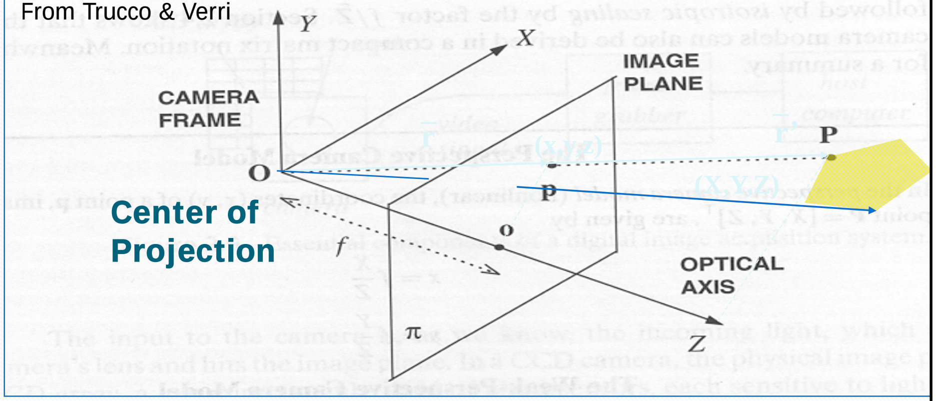

$\mathbf{r} = [x,y,z]^T$

$\mathbf{r}' = [X,Y,Z]^T$

$\mathbf{r}/f = \mathbf{r}'/Z$

f: effective focal length:
distance of image plane from O.

$x = f * X/Z$
$y = f * Y/Z$
$z = f$

Magnification

From Trucco & Verri

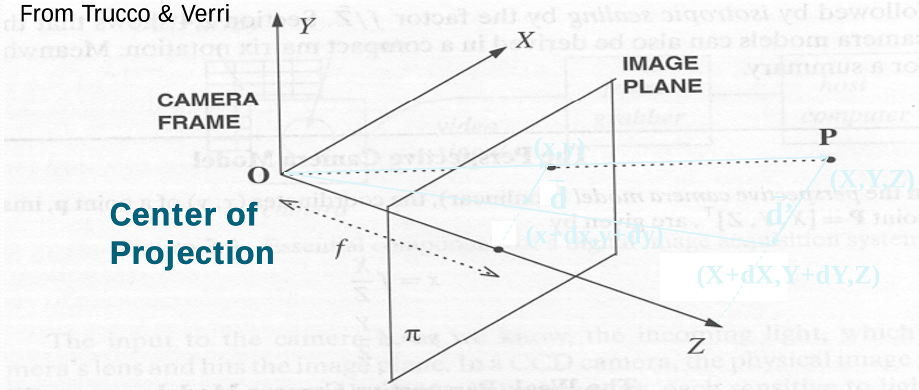

$x/f = X/Z$ $\quad (x+dx)/f = (X+dX)/Z \quad \Rightarrow \quad dx/f = dX/Z$

$y/f = Y/Z$ $\quad (y+dy)/z = (Y+dY)/Z \quad \quad dy/f = dY/Z$

Magnification

From Trucco & Verri

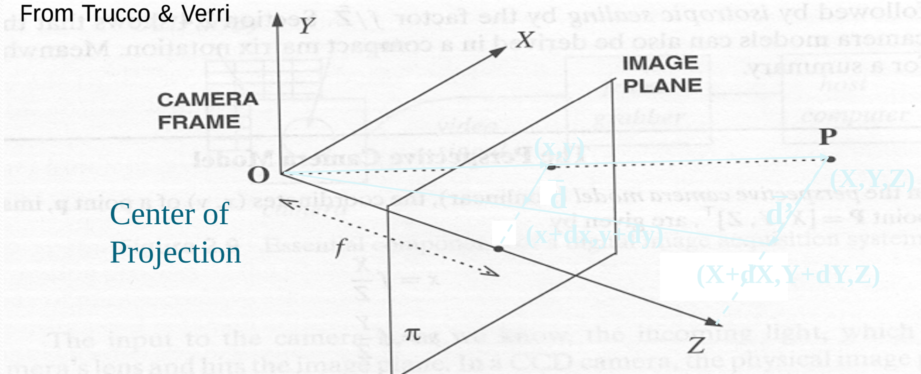

Magnification: |m|=||**d'**||/||**d**||=|f/Z|

or m=f/Z

m is negative when image is inverted…

Implications For Perception*

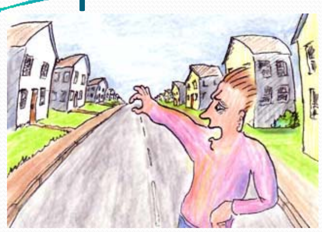

Same size things get smaller, we hardly notice…

Parallel lines meet at a point…

* A Cartoon Epistemology: http://cns-alumni.bu.edu/~slehar/cartoonepist/cartoonepist.html

Vanishing Points

Figure 2.4 Photograph illustrating a vanishing point. Parallel straight lines converge at a single point under perspective projection. This point is called the *vanishing point* of the straight lines. (Photograph by Herbert Gehr, from the magazine *Life*, July 1947, © Time Warner, Inc.)

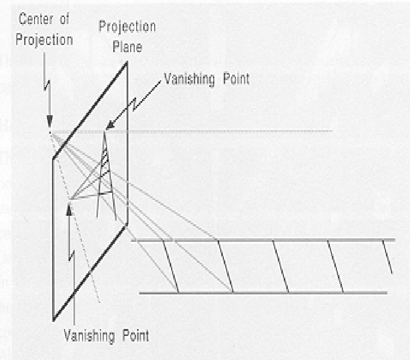

Figure 2.5 The vanishing point. The *vanishing point* of a straight line under perspective projection is that point on the projection surface at which the line would appear to "vanish" if the line were infinitely long in space. The location of the vanishing point of a straight line depends only on the orientation of the straight line in space, and not on the line's position: For any given spatial orientation, the vanishing point is located at that point on the projection surface where a straight line passing through the center of projection with the given orientation would intersect the projection surface.

(from NALWA)

Vanishi

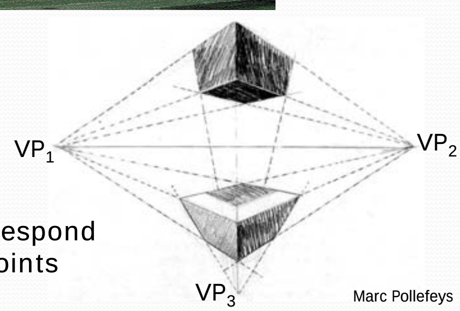

Different directions correspond to different vanishing points

Marc Pollefeys

3D is different…

3D Data Types: Volumetric Data

- Regularly-spaced grid in (x,y,z): "voxels"
- For each grid cell, store
  - Occupancy (binary: occupied / empty)
  - Density
  - Other properties
- Popular in medical imaging
  - CAT scans
  - MRI

3D Data Types: Surface Data

- Polyhedral
  - Piecewise planar
  - Polygons connected together
  - Most popular: "triangle meshes"
- Smooth
  - Higher-order (quadratic, cubic, etc.) curves
  - Bézier patches, splines, NURBS, subdivision surfaces, etc.

2½-D Data

- Image: stores an intensity / color along each of a set of regularly-spaced rays in space
- Range image: stores a depth along each of a set of regularly-spaced rays in space
- Not a complete 3D description: does not store objects occluded (from some viewpoint)
- View-dependent scene description

2½-D Data

- This is what most sensors / algorithms really return
- Advantages
  - Uniform parameterization
  - Adjacency / connectivity information
- Disadvantages
  - Does not represent entire object
  - View dependent

2½-D Data

- Range images
- Range surfaces
- Depth images
- Depth maps
- Height fields
- 2½-D images
- Surface profiles
- *xyz* maps
- …

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

Optical Range Acquisition Methods

- Advantages:
  - Non-contact
  - Safe
  - Usually inexpensive
  - Usually fast
- Disadvantages:
  - Sensitive to transparency
  - Confused by specularity and interreflection
  - Texture (helps some methods, hurts others)

Stereo

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

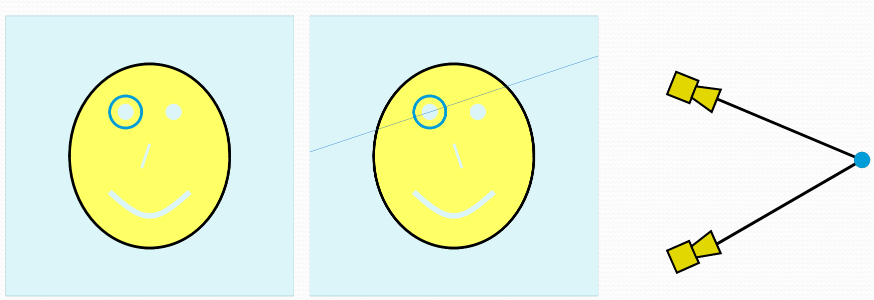

Stereo Vision

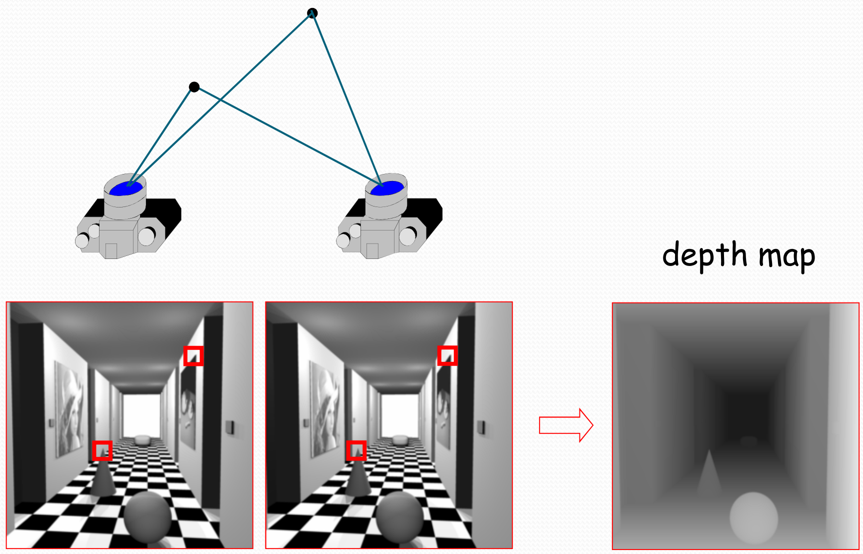

Triangulation

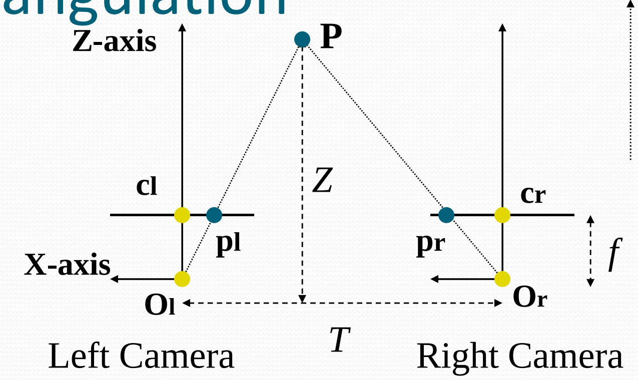

Fixation Point:Infinity. Parallel optical axes.

Calibrated Cameras

Similar triangles: $\frac{T + x_l - x_r}{Z - f} = \frac{T}{Z} \Rightarrow Z = f\frac{T}{d}, d = x_l - x_r$

d:disparity (difference in retinal positions).

T:baseline.

Depth (Z) is inversely proportional to d (fixation at infinity)

Traditional Stereo

Inherent problems of stereo:

- Need textured surfaces
- Matching problem
- Baseline trade-off
- Unstructured point set

→ Sparse estimates

→ Unreliable results

However

- Cheap (price, weight, size)
- Mobile
- Depth plus Color

Why More Than 2 Views?

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

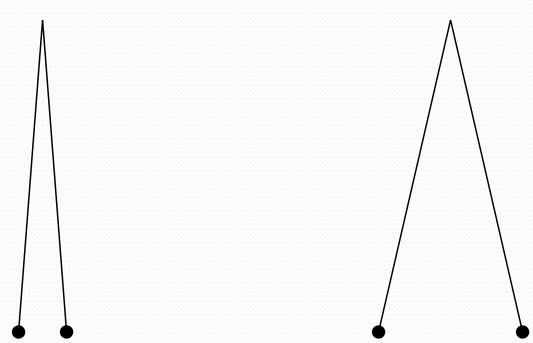

Multibaseline Stereo

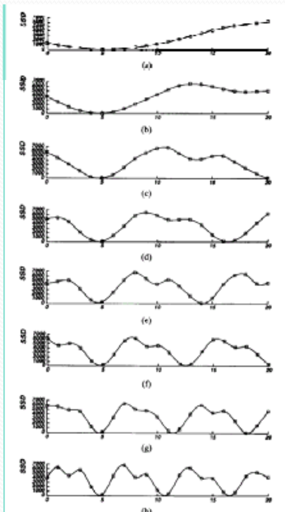

Fig. 5. SSD values versus inverse distance: (a) $B = b$; (b) $B = 2b$; (c) $B = 3b$; (d) $B = 4b$; (e) $B = 5b$; (f) $B = 6b$; (g) $B = 7b$; (h) $B = 8b$. The horizontal axis is normalized such that $8bF = 1$.

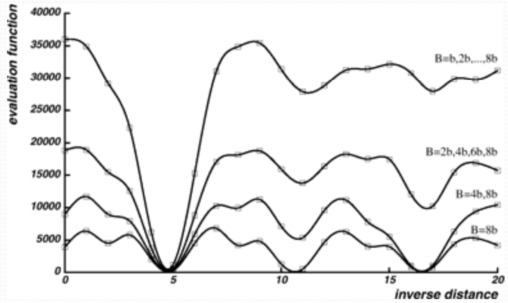

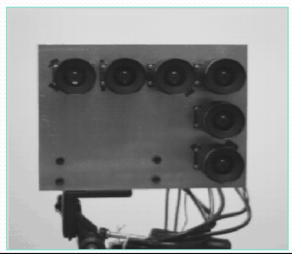

[Okutami & Kanade]

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

Photometric Stereo Setup

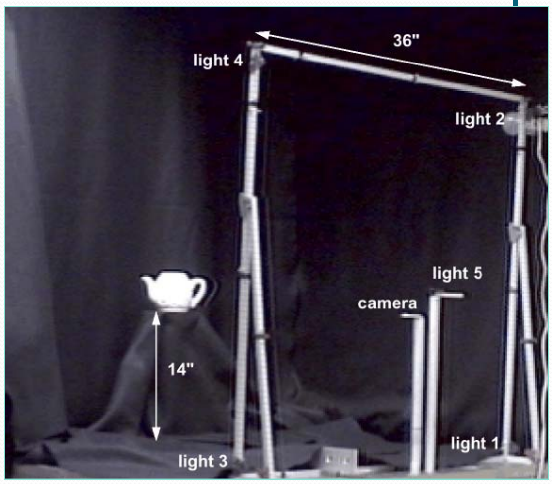

[Rushmeier et al., 1997

Photometric Stereo

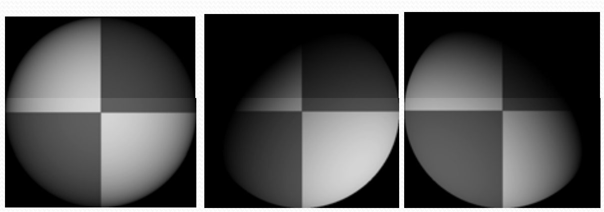

Use multiple light sources to resolve ambiguity
In surface orientation.

Note: Scene does not move – Correspondence
between points in different images is easy!

Notation: Direction of source i: s_i or (p_{s_i}, q_{s_i})
Image intensity produced by source i: $I_i(x, y)$

Lambertian Surfaces (special case)

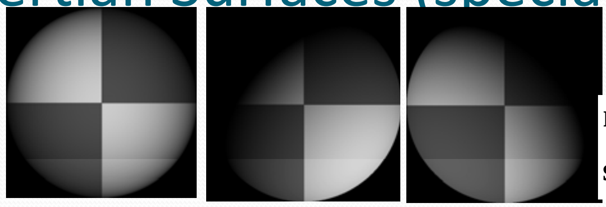

$\mathbf{n} = (n_x, n_y, n_z)$

$\mathbf{s_i} = (s_{x_i}, s_{y_i}, s_{z_i})$

Use THREE sources in directions $\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3$

Image Intensities measured at point (x,y):

$$I_1 = \rho(\mathbf{s}_1 \cdot \mathbf{n})$$
$$I_2 = \rho(\mathbf{s}_2 \cdot \mathbf{n})$$
$$I_3 = \rho(\mathbf{s}_3 \cdot \mathbf{n})$$

$$\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix} = \rho \begin{bmatrix} \mathbf{s}_1^T \\ \mathbf{s}_2^T \\ \mathbf{s}_3^T \end{bmatrix} \mathbf{n}$$

↑ $\mathbf{S}$

$$\rho\mathbf{n} = \mathbf{S}^{-1} \begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix} = \mathbf{N}$$

$\mathbf{n} = \dfrac{\mathbf{N}}{|\mathbf{N}|}$ ⟵ ← orientation

$\rho = |\mathbf{N}|$ ⟵ albedo

Photometric Stereo: RESULT

INPUT

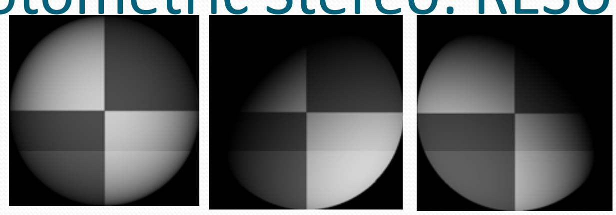

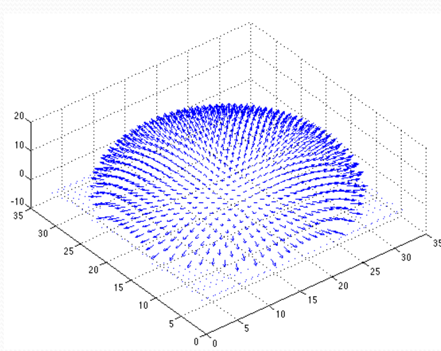

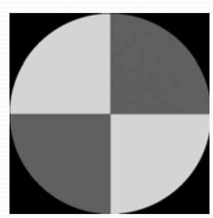

orientation

albedo

Data Acquisition Example

Color Image of Thomas Hunter Building, New York City.

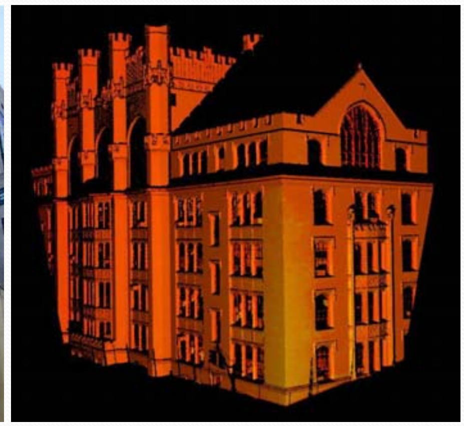

Range Image of same building. One million 3D points. Pseudocolor corresponds to laser intensity.

Time-of-flight scanners

Data Acquisition

- Spot laser scanner.
- Time of flight.
- Max Range: 100 meters.
- Scanning time: 16 minutes for one million points.
- Accuracy: ~6mm per range point

Data Acquisition

- Leica ScanStation 2
- Spherical field of view
- Registered color camera.
- Max Range: 300 meters.
- Scanning time: 2-3 times faster
- Accuracy: ~5mm per range point

Data Acquisition, Leica Scan Station 2, Park Avenue and 70th Street, NY

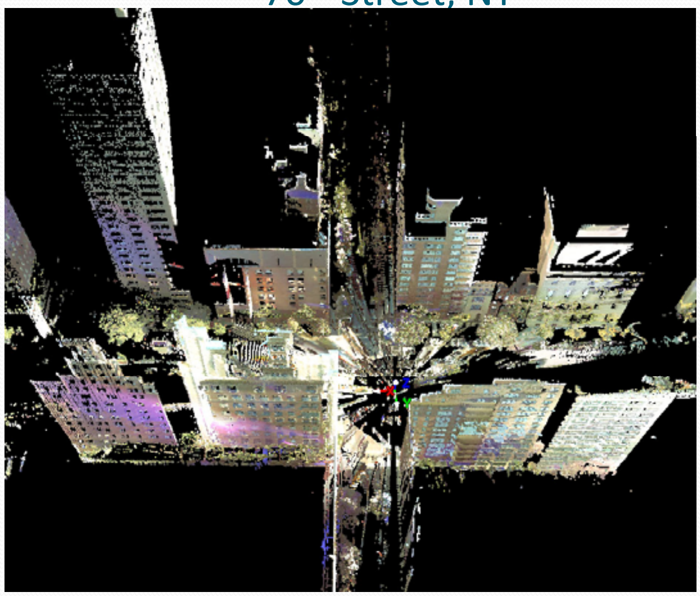

Cyclone view and Cyrax Video

Pulsed Time of Flight

- Send out pulse of light (usually laser), time how long it takes to return

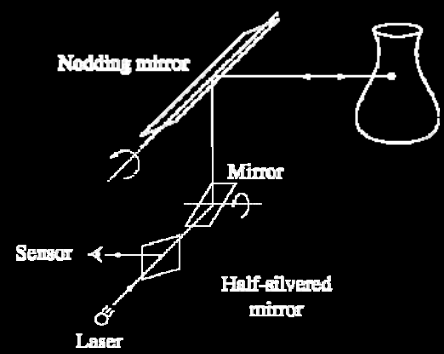

$$d = \frac{1}{2} c \Delta t$$

Pulsed Time of Flight

- Advantages:
  - Large working volume (more than 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.

Optical Triangulation

Sources of error:
1) grazing angle,
2) object boundaries.

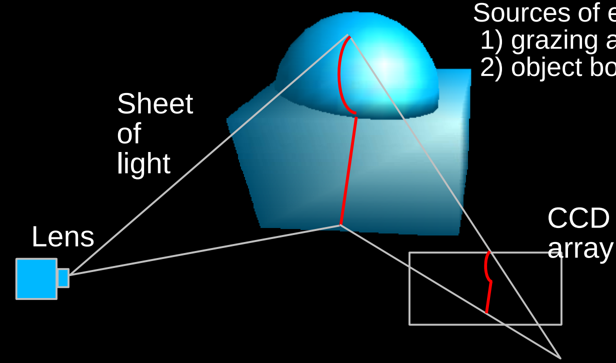

Optical Triangulation

Sources of error:
1) grazing angle,
2) object boundaries.

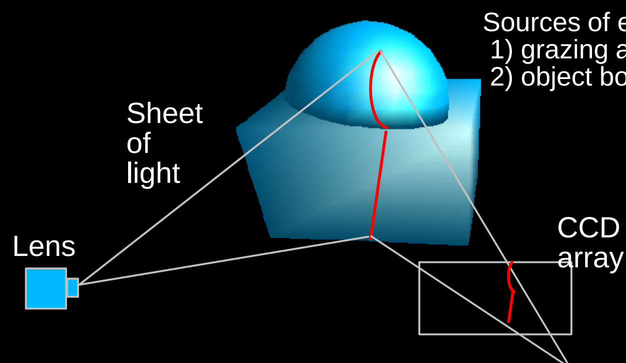

Active Optical Triangulation

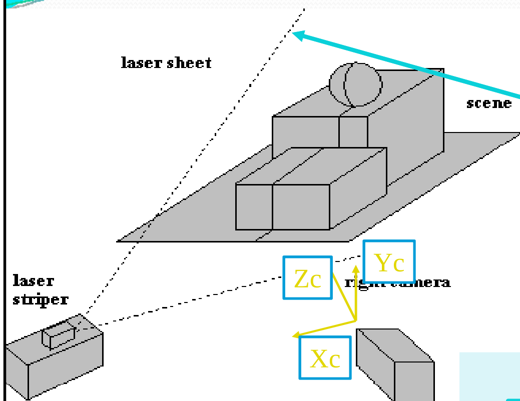

Light Stripe System.

Light Plane:
A**X**+B**Y**+C**Z**+D=0
(in camera frame)
Image Point:
x=f **X**/**Z**, y=f **Y**/**Z**
(perspective)

Image 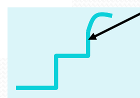

Triangulation: **Z**=-D f/(A x + B y + C f)

Move light stripe or object.

Triangulation

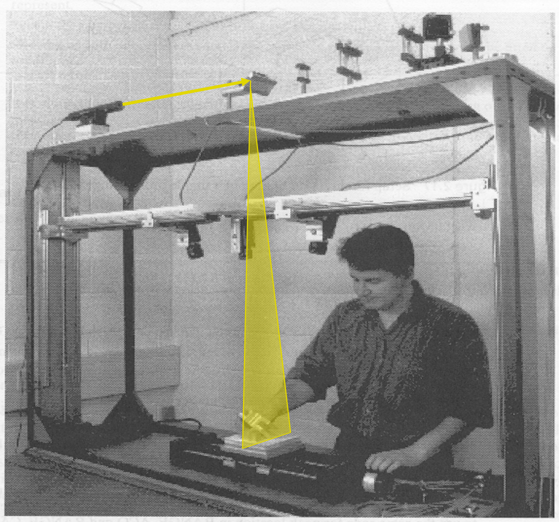

Figure 2.16 A real 3-D triangulation system, developed at Heriot-Watt University by A. M. Wallace and coworkers. Notice the laser source (top left), which generates a laser beam; the optical components forming the plane of laser light (top middle and left); the cameras; and the motorized platform (bottom middle) supporting the object and sweeping it through the stationary plane of light.

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

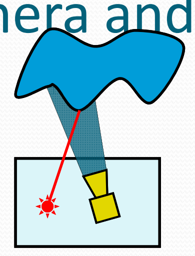

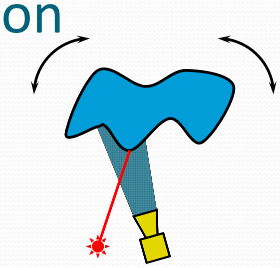

Triangulation: Moving the Camera and Illumination

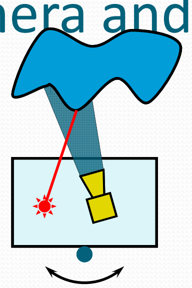

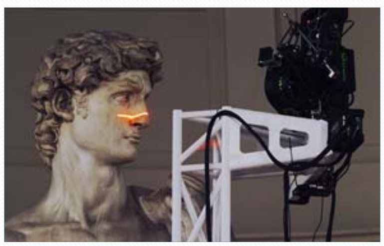

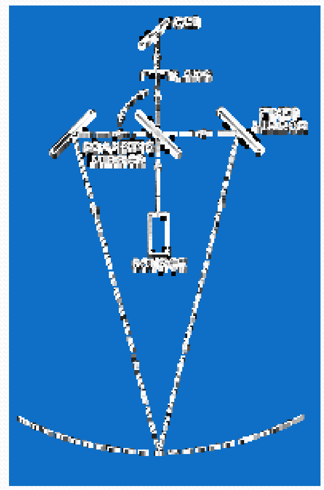

The Digital Michelangelo Project

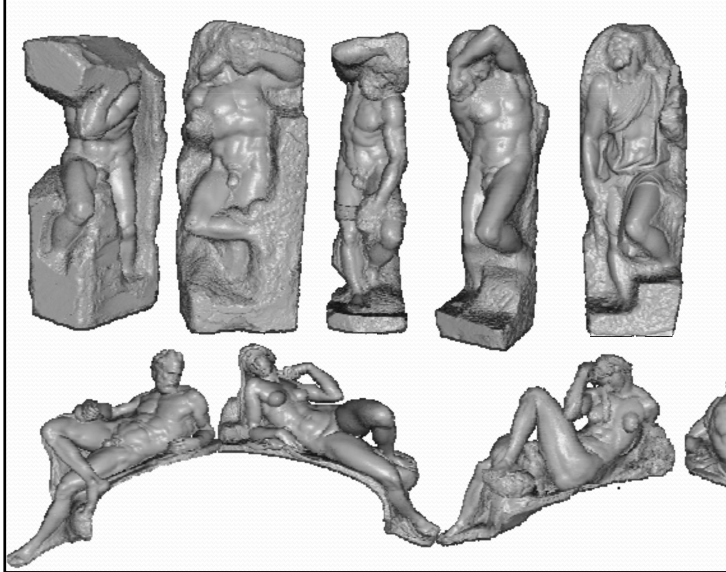

Marc Levoy (Stanford)

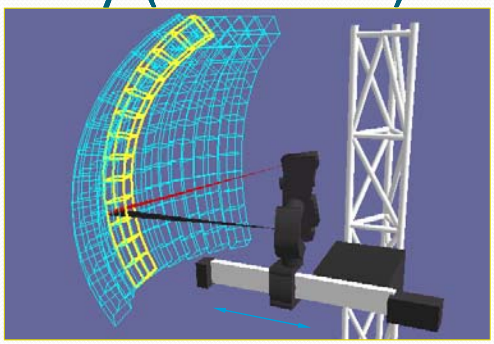

- calibrated motions
  - pitch (yellow)
  - pan (blue)
  - horizontal translation (orange)
- uncalibrated motions
  - vertical translation
  - remounting the scan head
  - moving the entire gantry

Scanner design

4 motorized axes

truss extensions
for tall statues

laser, range camera,
white light, and color camera

Scanning the David

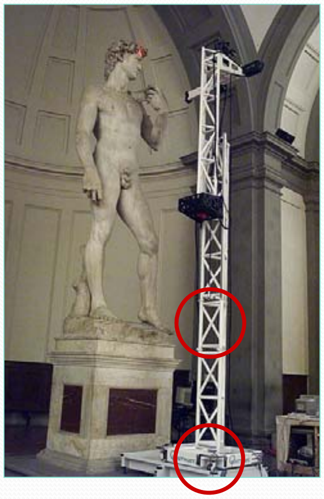

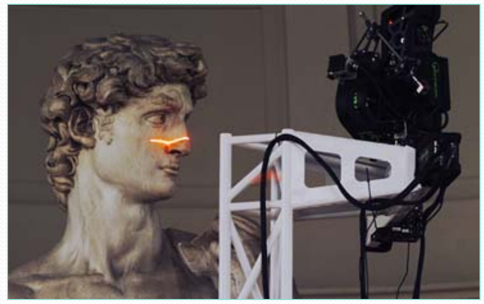

height of gantry: 7.5 meters
weight of gantry: 800 kilograms

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

- 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°)

Multi-Stripe Triangulation

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

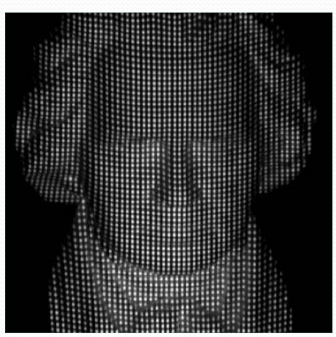

Multi-Stripe Triangulation

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

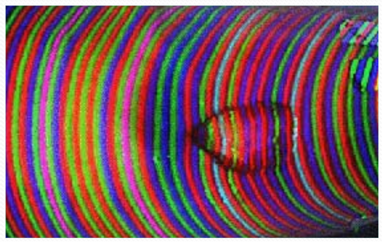

Multi-Stripe Triangulation

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

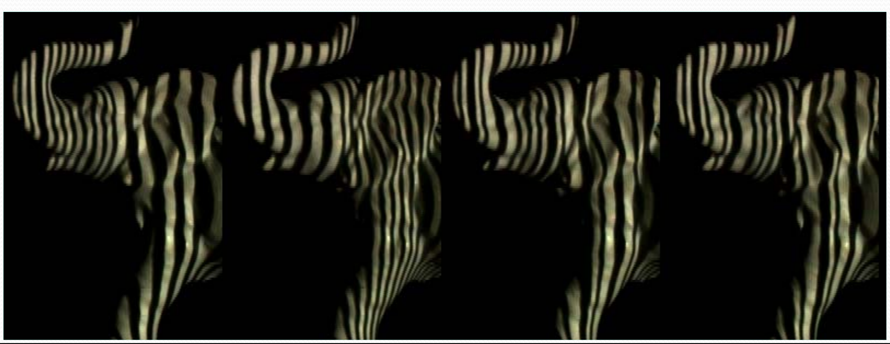

Time-Coded Light Patterns

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

Laser-based Sensing

Provides solutions

- Dense & reliable
- Regular grid → Adjacency info → Mesh

But

- Complex & large point sets
- Redundant point sets
- No color information (explain this)
- Expensive, non-mobile

Major Issues

- Registration of point sets
- Global and coherent geometry
  - remove redundancy
  - handle holes
  - handle all types of geometries
- Handle complexity
  - Fast Rendering

Representations of 3D Scenes

Geometry and Material	Complete description	Global Geometry
Geometry and Images	Traditional texture mapping	
Images with Depth		Local Geometry
Panorama with Depth		
Light-Field/Lumigraph	Scene as a light source	No/Approx. Geometry
Facade		
Panorama		
Colored Voxels		False Geometry

Head of Michelangelo's David

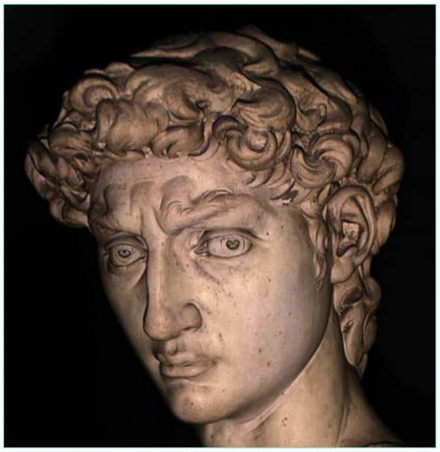

photograph

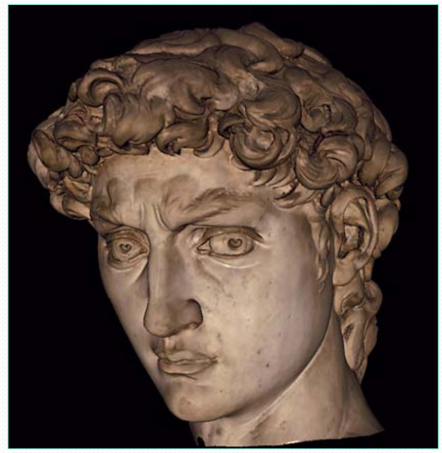

1.0 mm computer model