Active range scanning

“Traditional” pipeline for shape capture

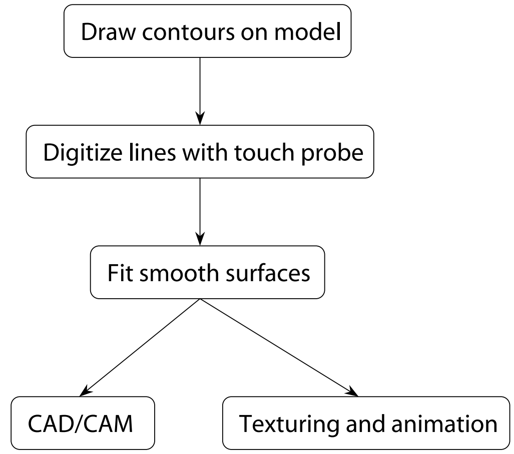

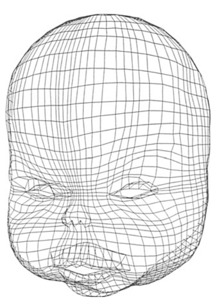

Hand digitized along parameter lines

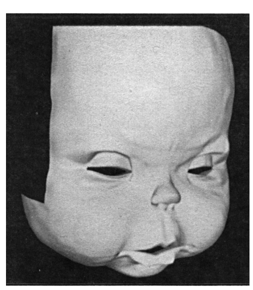

Fitted smooth surface

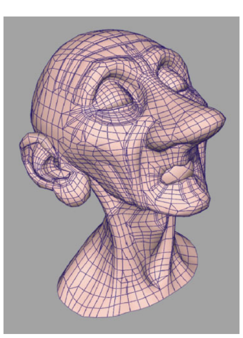

New pipeline for shape capture

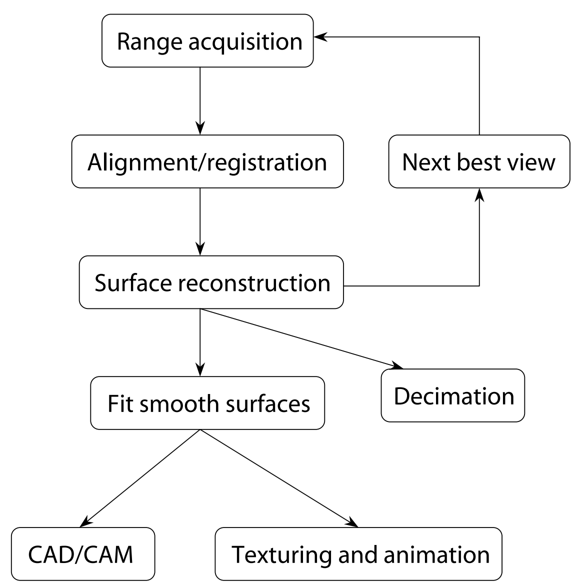

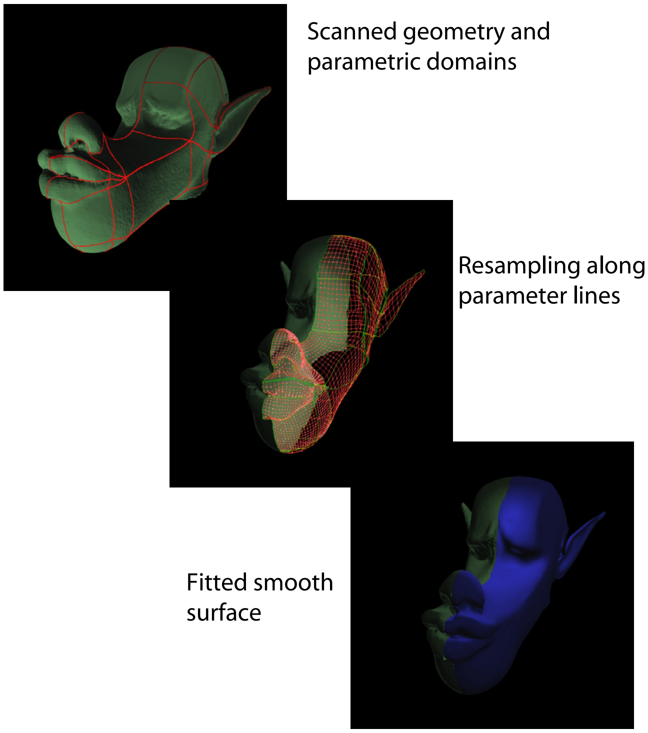

Outline

Overview

- Structure of data
- Quality measures

Non-optical methods

- Touch probes
- Industrial CT

Optical methods

- Lots of them…

Structure of the data

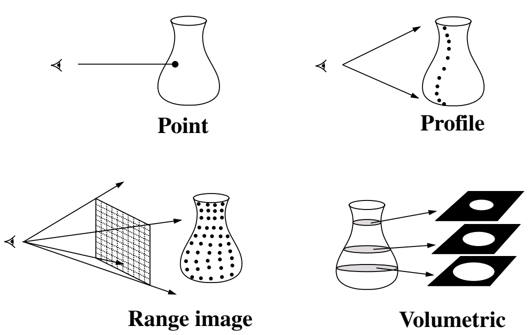

Quality measures

Resolution

- Smallest change in depth that sensor can report? Quantization? Spacing of samples?

Accuracy

- Statistical variations among repeated measurements of known value. Measurement of what?

Repeatability

- Do the measurements drift?

Environmental sensitivity

- Does temperature or wind speed influence measurements?

Speed

- Points per second? Seconds per point? Off-line or true realtime?

A taxonomy of shape capture

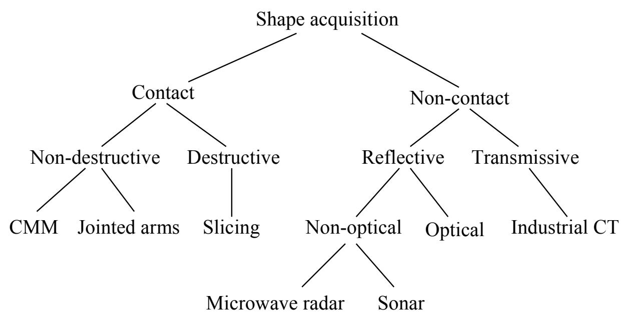

Coordinate measuring machine (CMM)

Principle:

Typically use X, Y, Z translation assembly with high precision displacement encoders. The probe may have a rotational degree of freedom to measure point of contact.

Working volume:

- 1-5 meters

Accuracy:

- 10 microns in 3 meter => 1 part in 300,000!

Notes:

- Expensive
- Slow

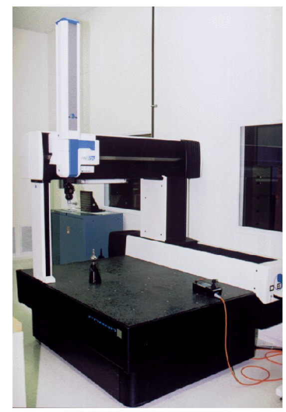

Overall view of coordinate measuring machine

Jointed arms

Principle:

A series of rotating joints provide freedom of movement. A sequence of joint-to-joint transformations provide position and orientation of stylus.

Working volume (all stats for FaroArm Silver series):

- 2.4 meters

Accuracy (2σ):

- 0.08mm or 1 part in 30,000

Notes:

- Much less expensive than CMM
- Slow

Magnetic trackers

Principle:

Magnetic fields are generated by transmitters. A sensor containing a set of orthogonal coils determines orientation and position based on relative and absolute strengths of induced currents.

Working volume (stats for Polhemus FastTrak):

- 3-10 meters

Accuracy:

- 0.75 mm and 0.15 degrees => ~1 part in 5000 positional accuracy

Notes:

- Wireless versions for motion tracking
- Sensitivity to metals and magnetic sources (e.g., monitors)

Industrial CT

Principle:

A set of full-surround, parallel-slice X-rays are taken as an object translates perpendicular to the slicing direction. After some fun math, the result is a volumetric density function, $\rho(\mathbf{x})$.

Surface can be extracted as a level set (a.k.a., isosurface) for some density value, ρ_o:

$$\rho(\mathbf{x}) = \rho_o \quad \text{or} \quad F(\mathbf{x}) = \rho(\mathbf{x}) - \rho_o = 0$$

Working volume:

- 0.2-2 meters

Accuracy:

- 0.25mm voxel spacing over 0.2 meters

Notes:

- Acquires internal cavities
- Expensive
- Hazardous radiation
- Accuracy varies with spatial variation of materials

Isocontour extraction

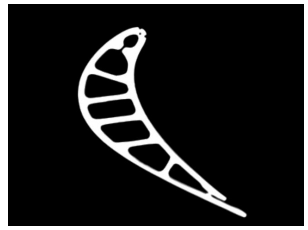

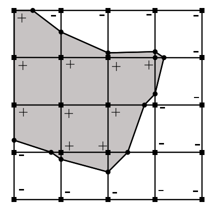

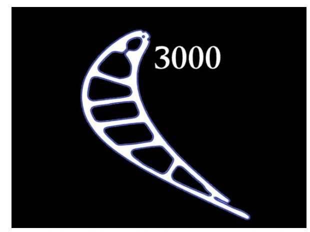

Isosurface extraction

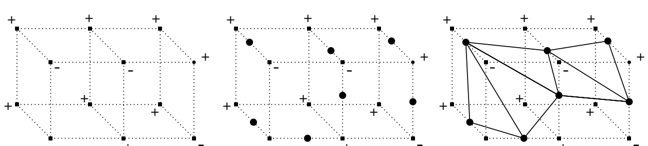

Isosurface extraction in practice

Q: Can you think of ambiguous cases that might cause problems?

Q: How could you overcome these problems?

Q: Is there a bound on the size or aspect ratio of triangles generated?

Q: If the surface were a single single connected component, how might you extract the surface efficiently?

Taxonomy of shape capture

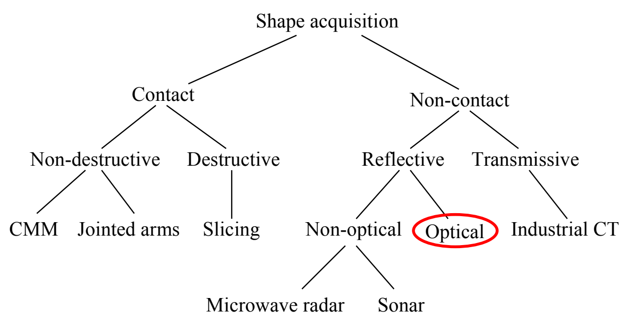

Taxonomy of optical shape capture

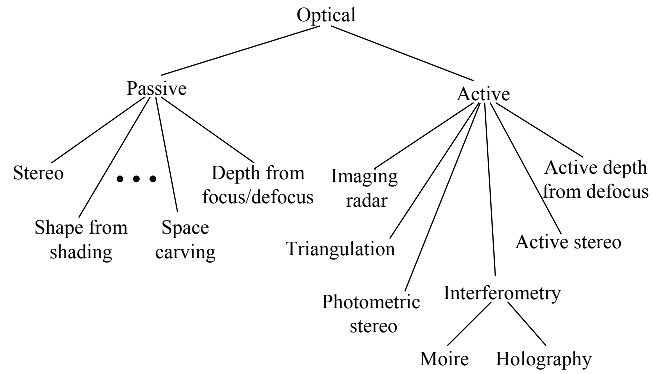

Optical shape capture

Optical shape capture only uses light and cameras to measure geometry.

Q: What would be some advantages of optical methods?

Q: What would be some **dis**advantages of optical methods?

Illumination

For active shape capture, need controllable illumination.

The basic choices are:

- Incoherent light (incandescent bulb, projector)
- Coherent light (lasers)

Q: What advantages might lasers offer?

Q: What **dis**advantages might they have?

Focusing a laser beam

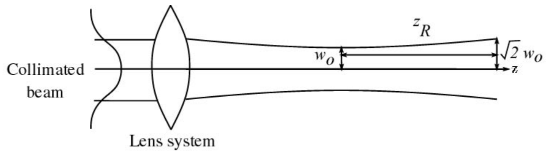

$$z_R = \frac{\pi w_o^2}{\lambda}$$

z_R = Rayleigh range
w_o = beam waist (narrowest laser width)
λ = wavelength of laser

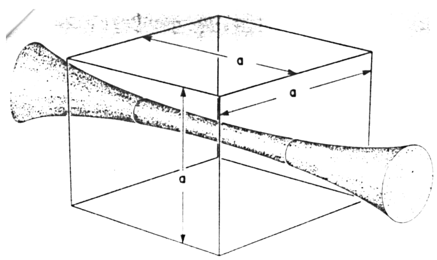

Laser speckle

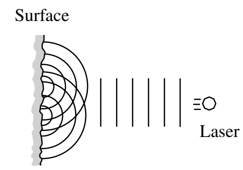

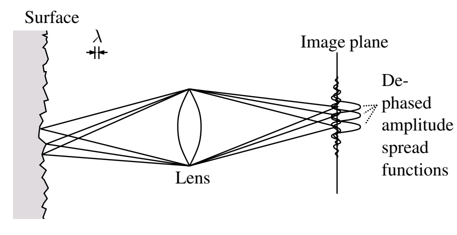

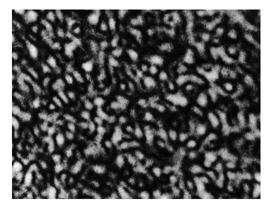

Imaging radar: time of flight

A pulse of light is emitted, and the time of the reflected pulse is recorded:

$c\,t = 2\,r$ = roundtrip distance

Typical scanning configuration:

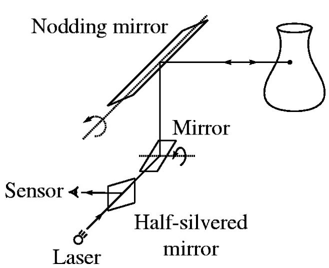

Imaging radar: time of flight

Working volume (for Cyra Cyrax):

- 50-100 meters, 60 degree fov

Accuracy:

- 1 cm => 1 part in 10,000

Notes:

- 1 cm accuracy requires 10 picosecond resolution
- Does not scale down readily

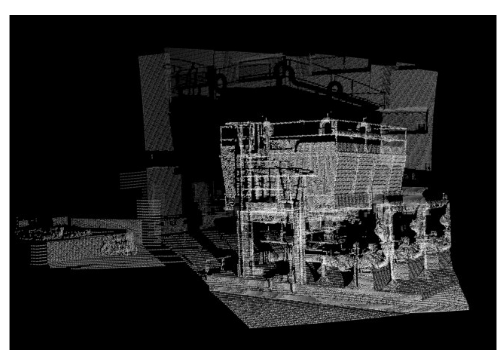

Imaging radar: Amplitude Modulation

The current to a laser diode is driven at frequency:

$$f_{AM} = \frac{c}{\lambda_{AM}}$$

The phase difference between incoming and outgoing signals, $\Delta\phi$, gives the range.

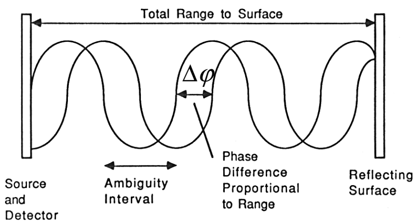

Imaging radar: Amplitude Modulation

Solving for the range:

$$2r = \frac{\Delta\varphi}{2\pi}\lambda_{AM} + n\lambda_{AM}$$

or:

$$r = \frac{\Delta\varphi}{4\pi}\lambda_{AM} + \frac{n\lambda_{AM}}{2}$$

Q: What kind of range ambiguity is there?

Q: How can it be overcome?

Optical triangulation

A beam of light strikes the surface, and some of the light bounces toward an off-axis sensor.

The center of the imaged reflection is triangulated against the laser line of sight.

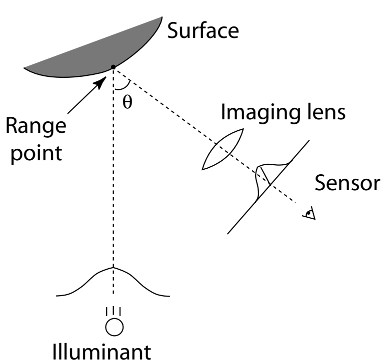

The Scheimpflug condition

A lens images a plane to a plane. If the object plane is tilted, then so is the image plane.

The image tilt is related to object tilt through the Scheimpflug condition:

$$\tan\alpha = \frac{\tan\theta}{M}$$

Where *M* is the on-axis magnification.

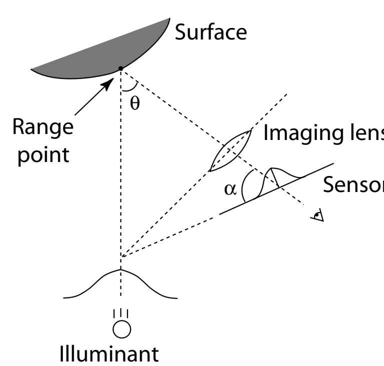

Triangulation angle

When designing an optical triangulation, we want:

- Small triangulation angle
- Uniform resolution

These requirements are at odds with each other.

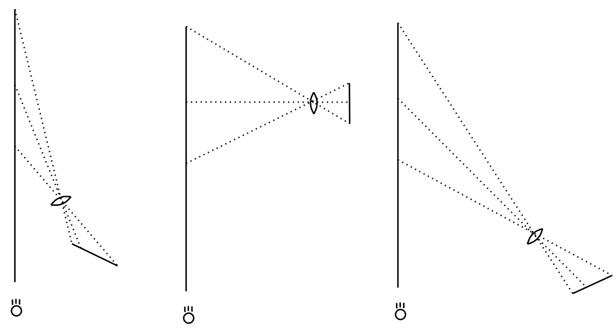

Scanning: move illuminant

A scene can be scanned by sweeping the illuminant.

Problems:

- Loss of resolution due to defocus
- Large variation in field of view
- Large variation in resolution

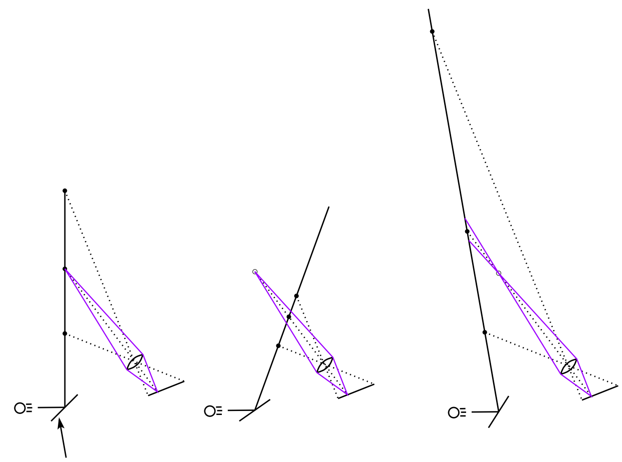

Scanning: move whole unit

Can instead move the laser and camera together, e.g., by translating or rotating a scanning unit.

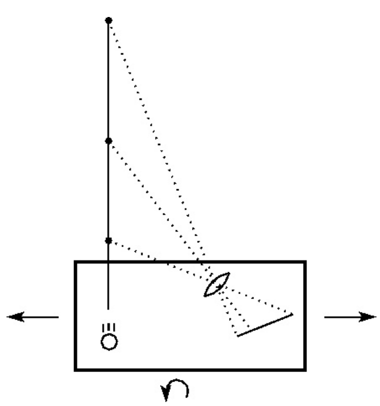

Synchronized scanning

A novel design was created and patented at the NRC of Canada [Rioux'87].

Basic idea: sweep the laser and sensor *simultaneously*.

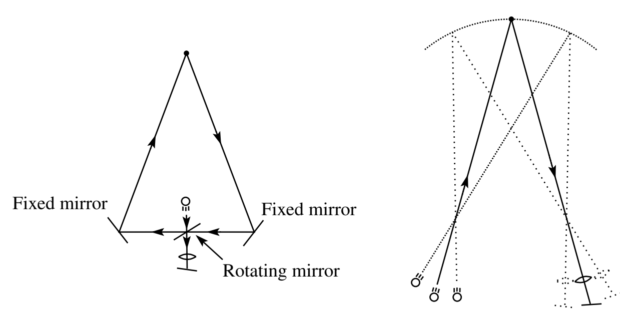

After unfolding the optics

Triangulation in 3D

Can extend into 3D by spreading the light beam into a plane:

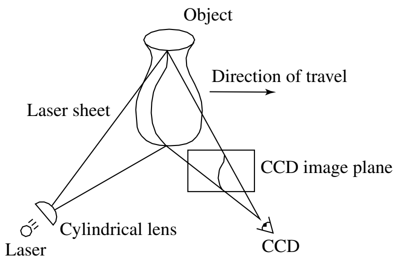

Working volume (Cyberware MS3030):

- 30 cm

Triangulation angle:

- 30 degrees

Accuracy:

- 0.3 mm => one part in 1000

Speed:

- 15,000 pts/sec

Triangulation in 3D

Can also extend into 3D with a nodding mirror:

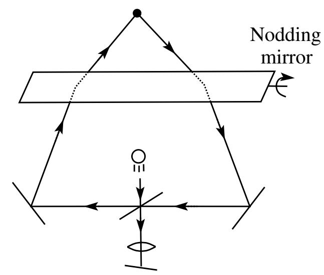

Working volume (Rioux scanner):

- 8 cm

Triangulation angle:

- ~15 degrees

Accuracy:

- 50 microns => one part in 1500

Speed:

- 10,000 pts/sec

Errors in optical triangulation

Finding the center of the imaged pulse is tricky.

If the surface exhibits variations in shape or reflectance, then laser width limits accuracy.

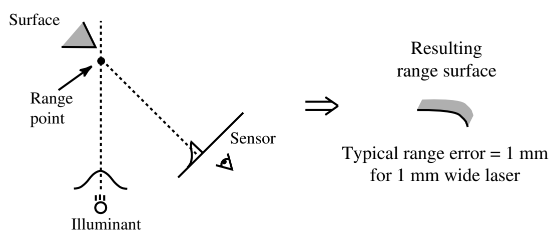

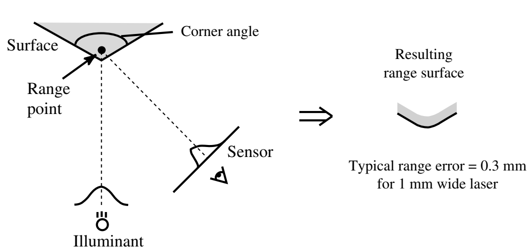

Errors in optical triangulation

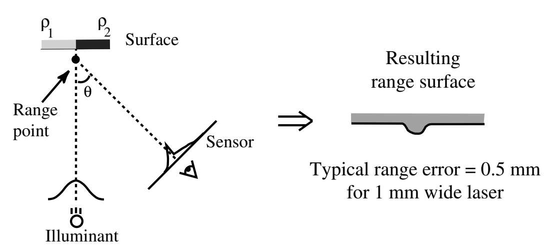

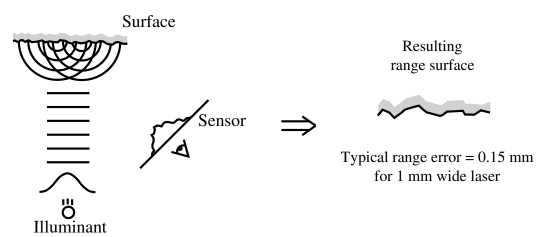

Spacetime analysis

A solution to this problem for scanning systems is spacetime analysis [Curless 95]:

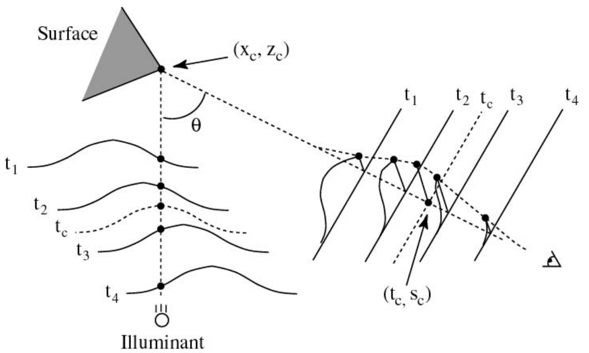

Spacetime analysis: results

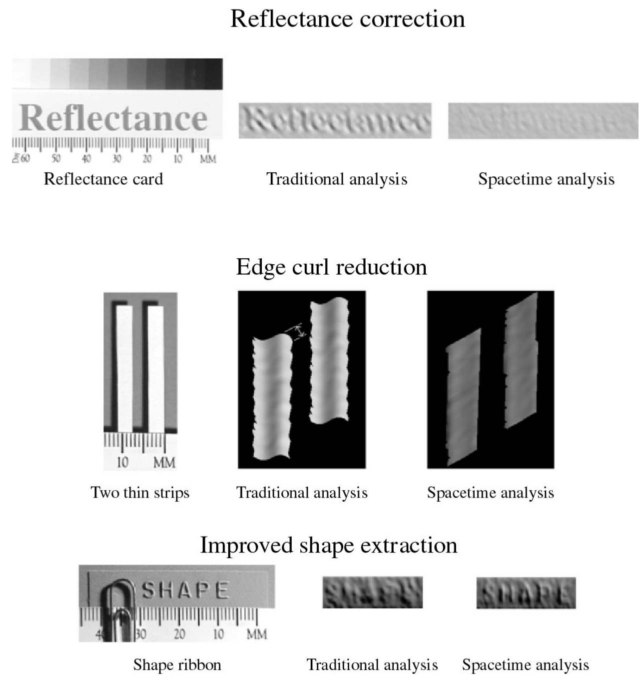

Multi-spot and multi-stripe triangulation

For faster acquisition, some scanners use multiple spots or stripes.

Trade off depth-of-field for speed.

Problem: ambiguity.

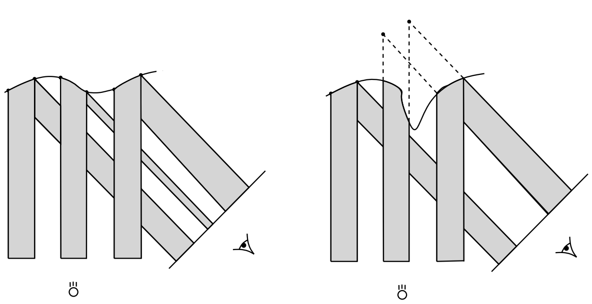

Q: How can we address this ambiguity?

"One shot" active triangulation

One approach is to look at the whole 2D image and borrow information from adjacent scanlines.

Proesmans96 developed such an approach, the basis for Eyetronics's product.

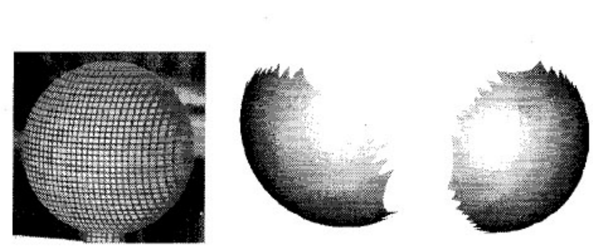

“One shot” active triangulation

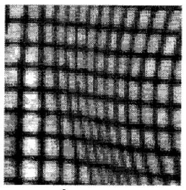
Image

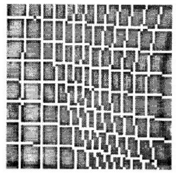
Pixel labeling

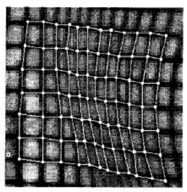
Line edges

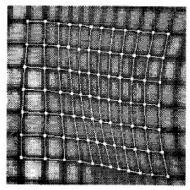
Snake-based curves

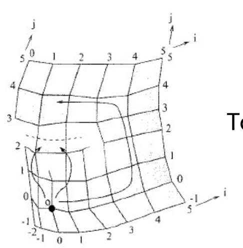
Topology checking

Q: What assumptions does this make about the surface?

Color-coded stripes

Another possibility is to project stripes of different colors.

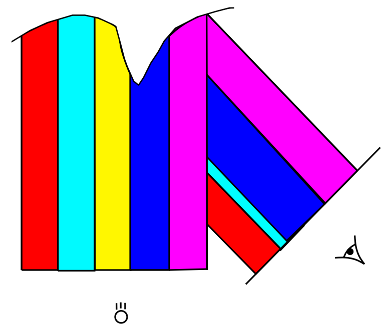

Q: What assumptions does this make about the surface?

Binary coded illumination

Alternative: resolve visibility hierarchically (logN).

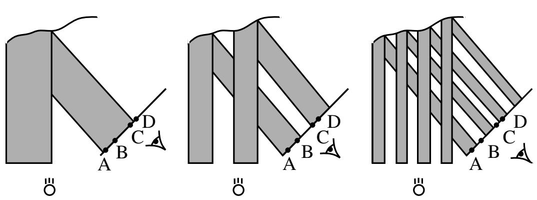

Binary codes:
A = 1 1 1 C = 1 0 0
B = 1 1 0 D = 0 1 1

Moire

Illuminate a surface with a sinusoidal grating and photograph from a different angle.

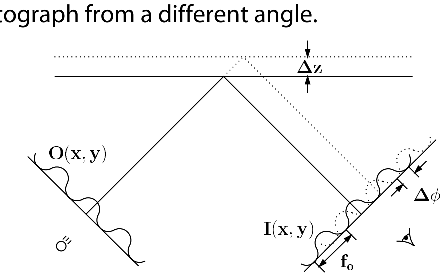

Reflected intensity is of the form:

$$I(x,y) = a(x,y)\cos[2\pi fx + \phi(x,y)]$$

where:

- $a(x,y)$ is reflected amplitude
- f is the period of the grating reflected from a flat surface
- $\phi(x,y)$ is the phase difference due to depth

Note that:

$$\Delta\phi(x,y) = \Delta z(x,y)$$

Quadrature moire

We need to "demodulate" to get phase. Multiply by a sinusoid of same frequency:

$$\begin{aligned} I(x,y)\cos[2\pi fx] &= a(x,y)\cos[2\pi fx+\phi(x,y)]\cos[2\pi fx] \\ &= \frac{1}{2}a(x,y)\{\cos[\phi(x,y)]+\cos[4\pi fx+\phi(x,y)]\} \end{aligned}$$

If *a*(*x,y*) and ϕ(*x,y*) are slowing varying, then we can filter out the high frequencies with a low-pass filter:

$$M_1(x,y) = LPF\{I(x,y)\cos[2\pi fx]\} = \frac{1}{2}a(x,y)\cos[\phi(x,y)]$$

Next, multiply by a phase shifted sinusoid:

$$\begin{aligned} I(x,y)\sin[2\pi fx] &= a(x,y)\cos[2\pi fx+\phi(x,y)]\cos[2\pi fx+\pi/2] \\ &= \frac{1}{2}a(x,y)\{\sin[\phi(x,y)]-\sin[4\pi fx+\phi(x,y)]\} \end{aligned}$$

Filtering again:

$$M_2(x,y) = LPF\{I(x,y)\sin[2\pi fx]\} = \frac{1}{2}a(x,y)\sin[\phi(x,y)]$$

Quadrature moire

Thus, we have:

$$\tan\phi(x,y) = M_2(x,y)/M_1(x,y)$$

and we can solve for the phase (and thus the depth).

Q: What assumptions does this technique make about the surface?

Shadow moire

Shadow moire:

- Place a grating (e.g., stripes on a transparency) near the surface.
- Illuminate with a lamp.
- Instant moire!

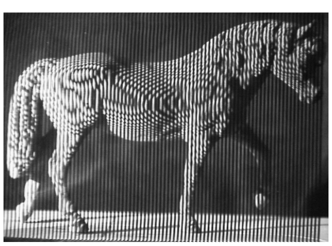

Shadow moire

Filtered image

Active stereo

Passive stereo methods match features observed by two cameras and triangulate.

Active stereo simplifies feature-finding with structured light.

Problem: ambiguity.

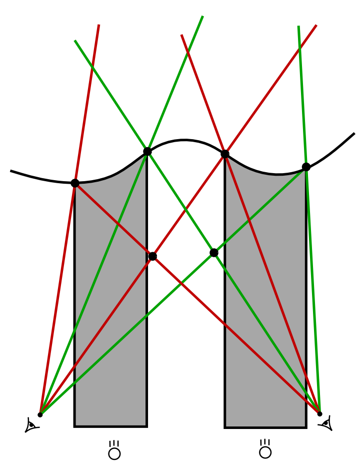

Active multi-baseline stereo

Using multiple cameras reduces likelihood of false matches.

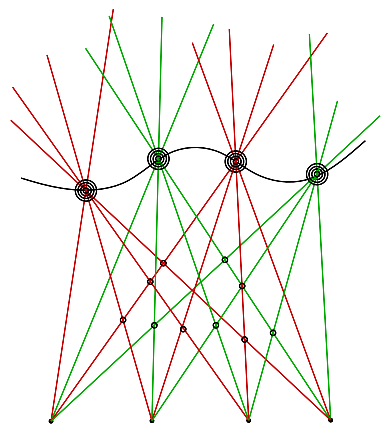

Depth from defocus

Depth of field for large apertures will cause the image of a point to blur.

The amount of blur indicates distance to the point.

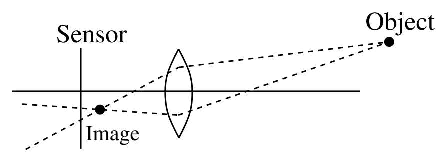

Depth from defocus

Problem: defocus ambiguity.

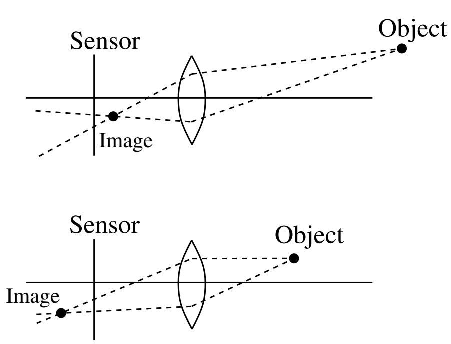

Depth from defocus

Solution: two sensor planes.

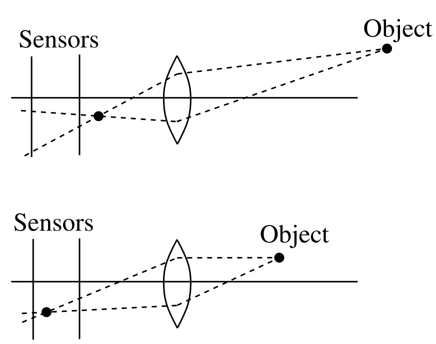

Problem: Does not work for objects without texture.

Active depth from defocus

Solution: project structured lighting onto surface.

[Nayar 95] demonstrates a real-time system utilizing telecentric optics.

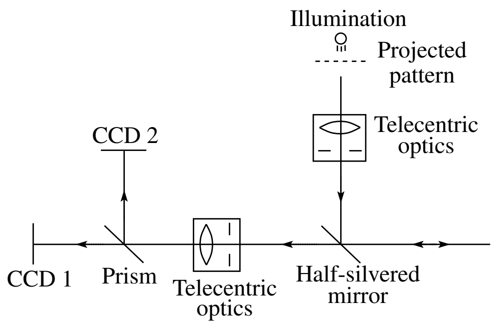

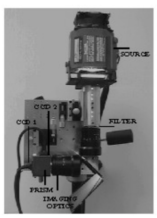

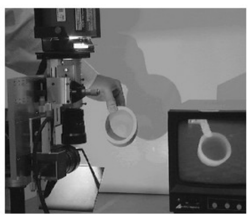

Bibliography

Besl, P. Advances in Machine Vision. "Chapter 1: Active optical range imaging sensors," pp. 1-63, Springer-Verlag, 1989.

Curless, B. and Levoy, M., "Better optical triangulation through spacetime analysis." In Proceedings of IEEE International Conference on Computer Vision, Cambridge, MA, USA, 20-23 June 1995, pp. 987-994.

Nayar, S.K., Watanabe, M., and Noguchi, M. "Real-time focus range sensor", Fifth International Conference on Computer Vision (1995), pp. 995-1001.

Proesmans, M. and Van Gool, L., "A sensor that extracts both 3D shape and surface texture," Proceedings of the 1996 IEEE/SICE/RSJ International Conference on Multisensor Fusion and Integration for Intelligent Systems, pp. 485-492.

Rioux, M., Bechthold, G., Taylor, D., and Duggan, M. "Design of a large depth of view three-dimensional camera for robot vision," Optical Engineering (1987), vol. 26, no. 12, pp. 1245-1250.