SIGGRAPH 99 Course on
3D Photography

Overview of Active Vision Techniques

Brian Curless
University of Washington

Overview

Introduction

Active vision techniques

- Imaging radar
- Triangulation
- Moire
- Active Stereo
- Active depth-from-defocus

Capturing appearance

A taxonomy

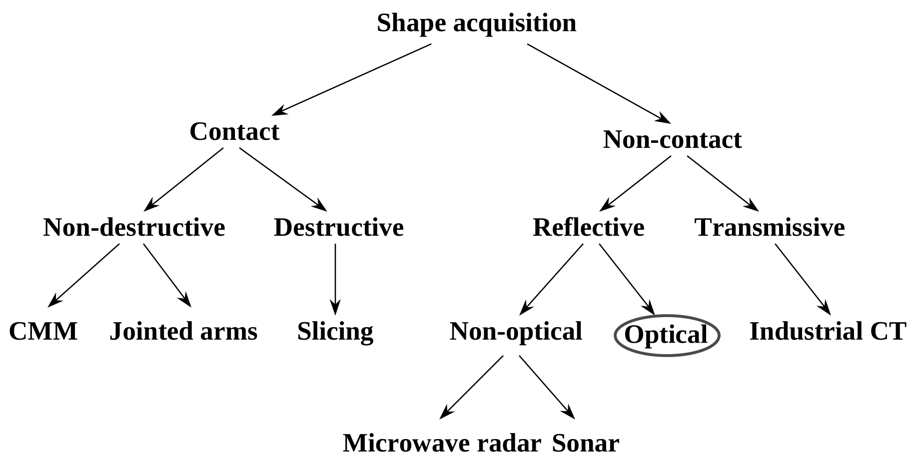

A taxonomy

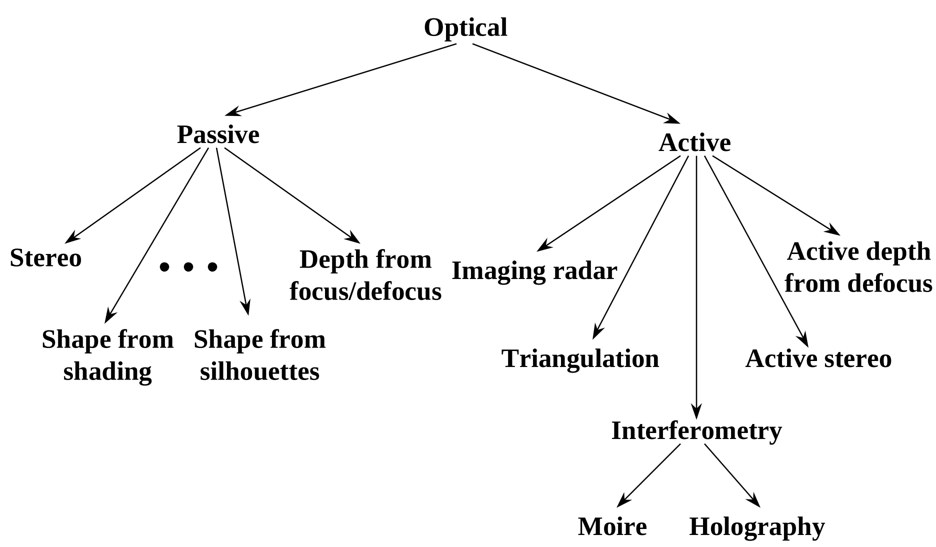

Structure of the data

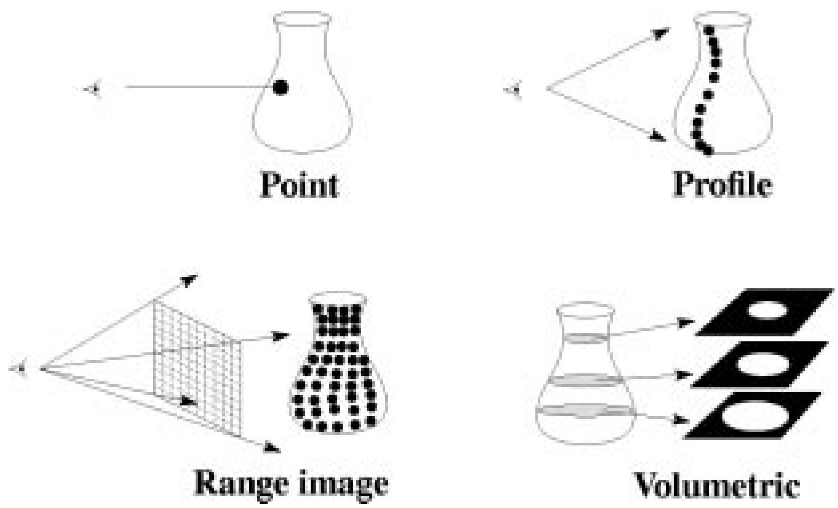

Quality measures

Resolution

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

Accuracy

Statistical variations among repeated measurements of known value.

Repeatability

Do the measurements drift?

Environmental sensitivity

Does temperature or wind speed influence measurements?

Speed

Optical range acquisition

Strengths

- ***Non-contact***
- ***Safe***
- ***Inexpensive (?)***
- ***Fast***

Limitations

- ***Can only acquire visible portions of the surface***
- ***Sensitivity to surface properties***
  - > *transparency, shininess, rapid color variations, darkness (no reflected light), subsurface scatter*
- ***Confused by interreflections***

Illumination

Why are lasers a good idea?

- ***Compact***
- ***Low power***
- ***Single wavelength is easy to isolate***
- ***No chromatic aberration***
- ***Tight focus over long distances***

Illumination

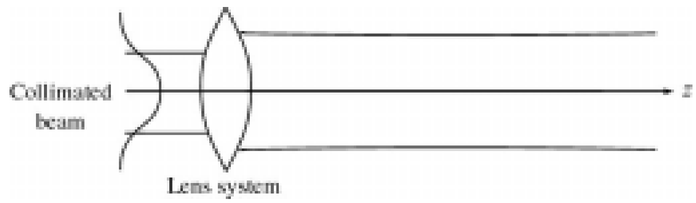

Illumination

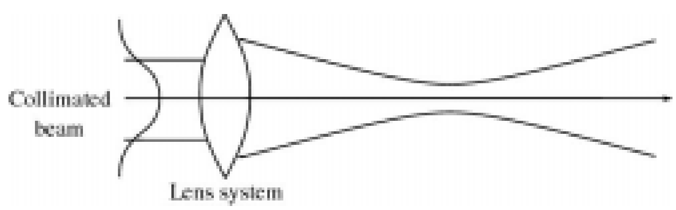

Illumination

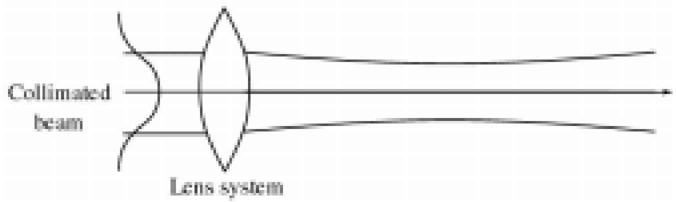

Illumination

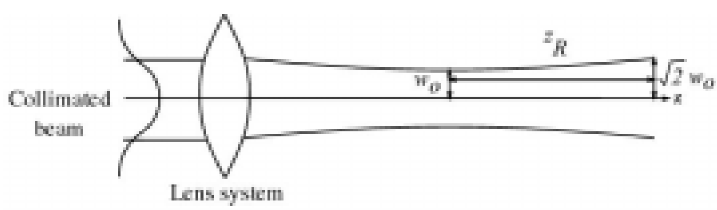

$$z_R = \frac{\pi w_o^2}{\lambda}$$

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

Illumination

Limitations of lasers

- ***Eye safety concerns***
- ***Laser speckle adds noise***
  - *> Narrowing the aperture increases the noise*

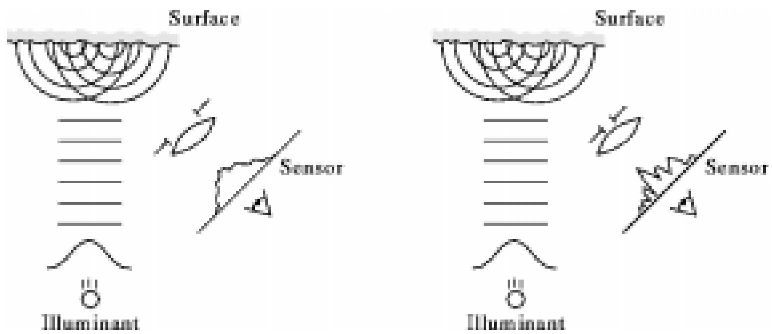

Imaging radar: time of flight

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

c t = 2 r = roundtrip distance

Typical scanning configuration:

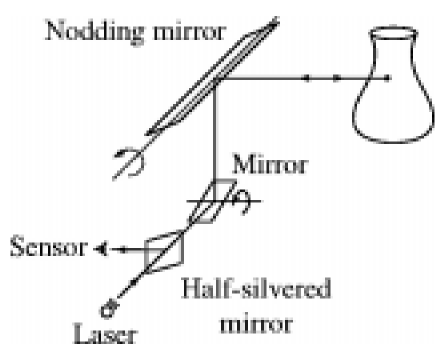

Imaging radar: Amplitude Modulation

The current to a laser diode is driven at frequency:

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

The phase difference between incoming and outgoing signals gives the range:

$$2r(\Delta\varphi) = n\lambda_{AM} + \frac{\Delta\varphi}{2\pi}\lambda_{AM} \Rightarrow r(\Delta\varphi) = \frac{1}{2}\lambda_{AM}\frac{(\Delta\varphi + 2\pi n)}{2\pi}$$

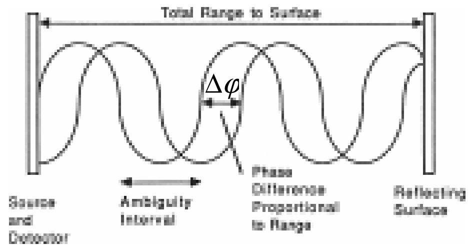

Figure from [Besl89]

Imaging radar: Amplitude Modulation

Note the ambiguity due to the + 2*πn*. This translates into range ambiguity:

$$r_{ambig} = \frac{n\lambda_{AM}}{2}$$

The ambiguity can be overcome with sweeps of increasingly finer wavelengths.

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.

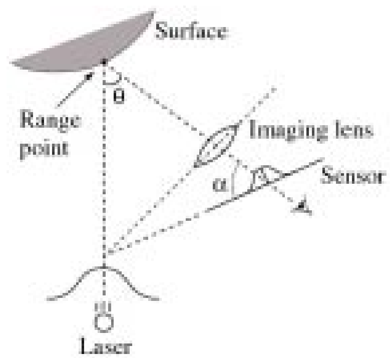

Optical triangulation

Lenses map planes to planes. If the object plane is tilted, then so should the image plane.

The image plane tilt is described by 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.

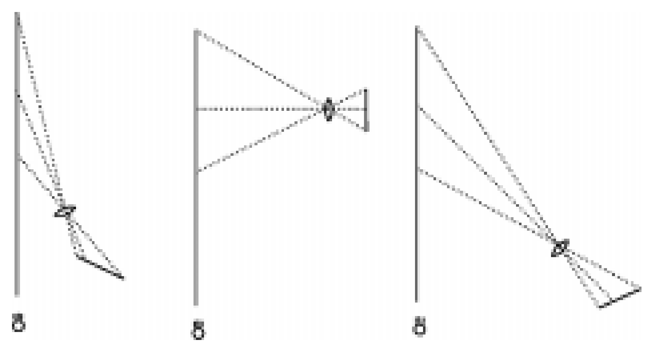

Triangulation scanning configurations

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***

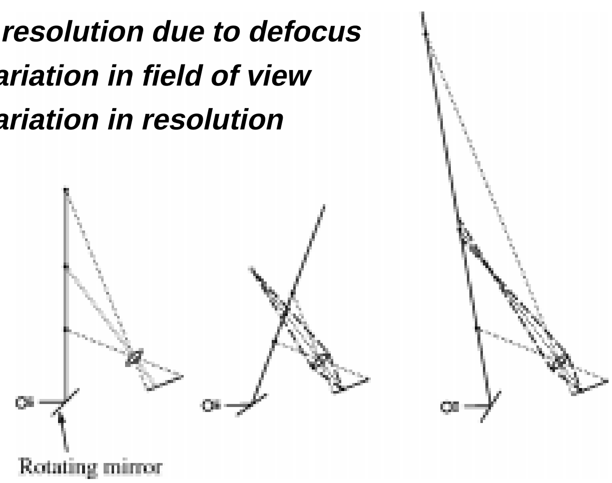

Triangulation scanning configurations

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

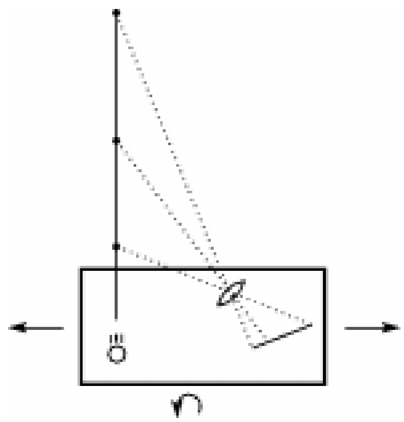

Triangulation scanning configurations

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

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

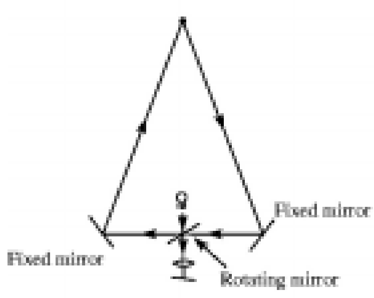

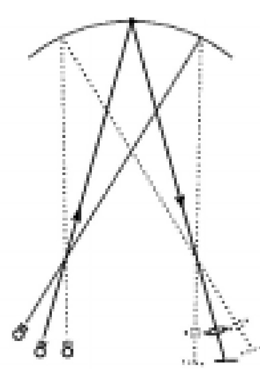

After unfolding the optics

Triangulation scanning configurations

Extension to 3D achievable as:

- ***flying spot***
- ***sweeping light stripe***
- ***hand-held light stripe on jointed arm***

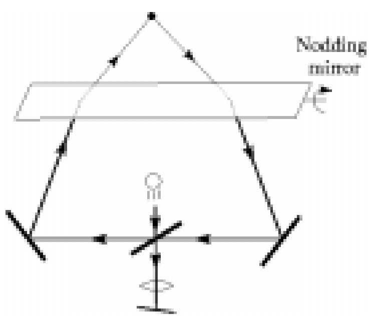

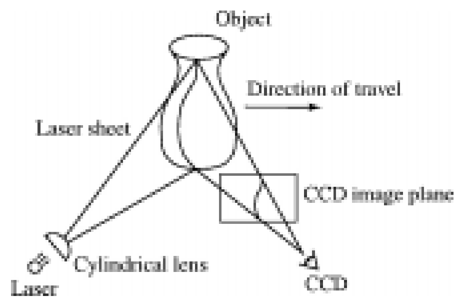

Errors in optical triangulation

Finding the center of the imaged pulse is tricky.

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

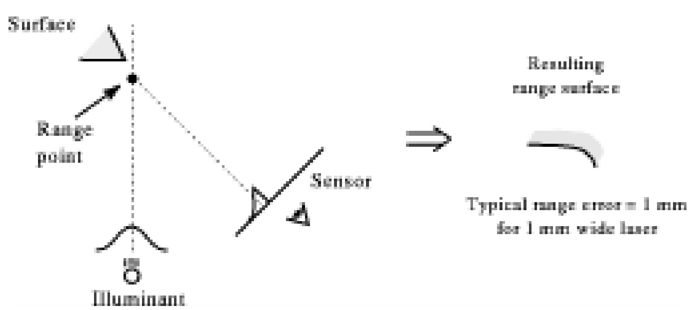

Errors in optical triangulation

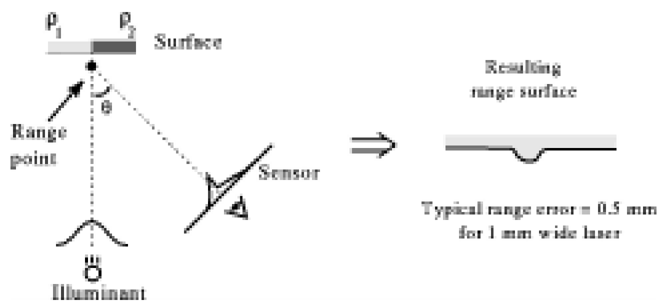

Spacetime analysis

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

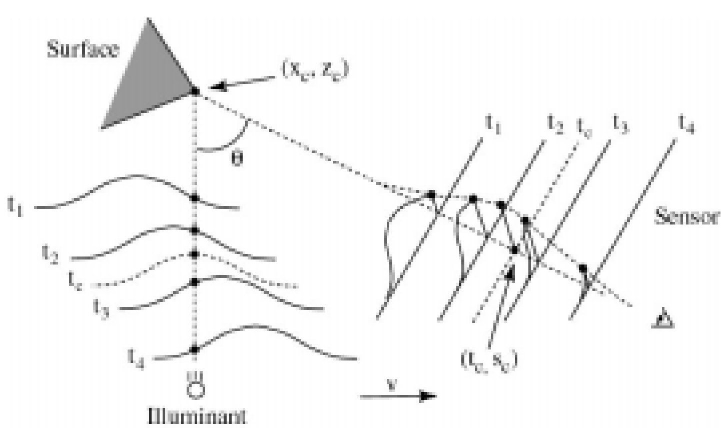

Spacetime analysis

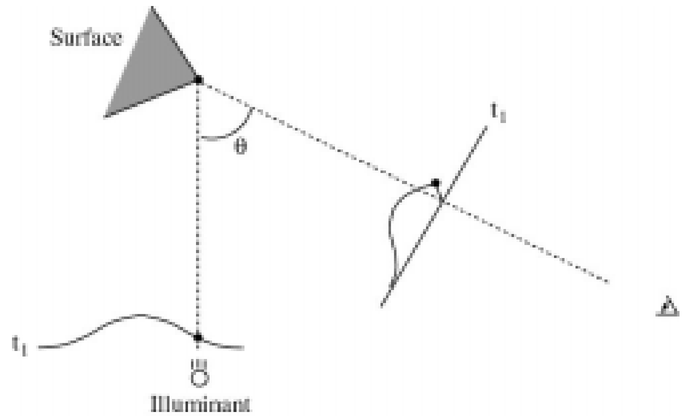

Spacetime analysis

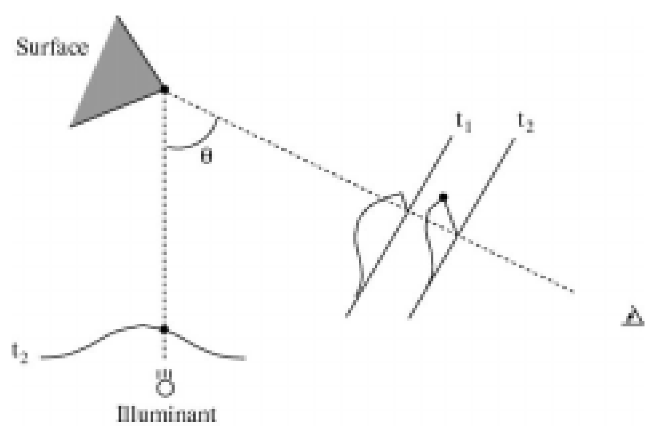

Spacetime analysis

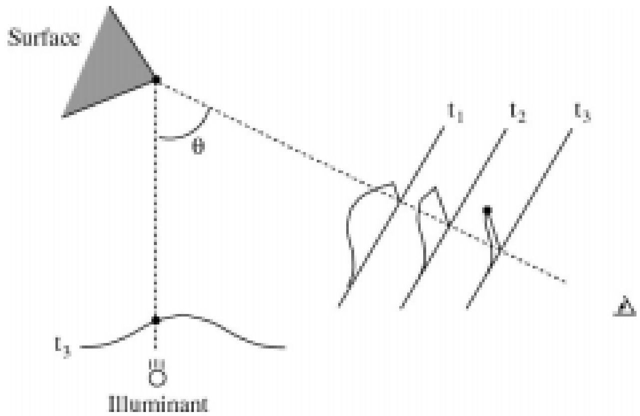

Spacetime analysis

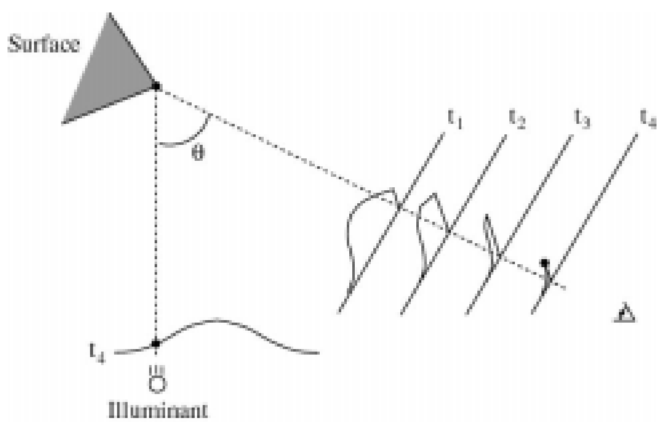

Spacetime analysis

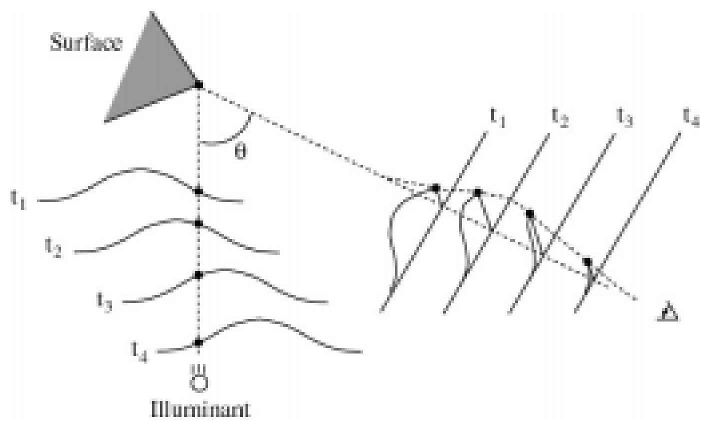

Spacetime analysis

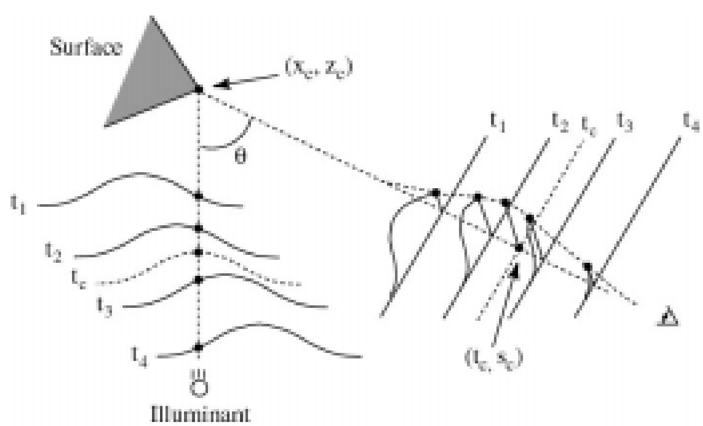

Spacetime analysis: results

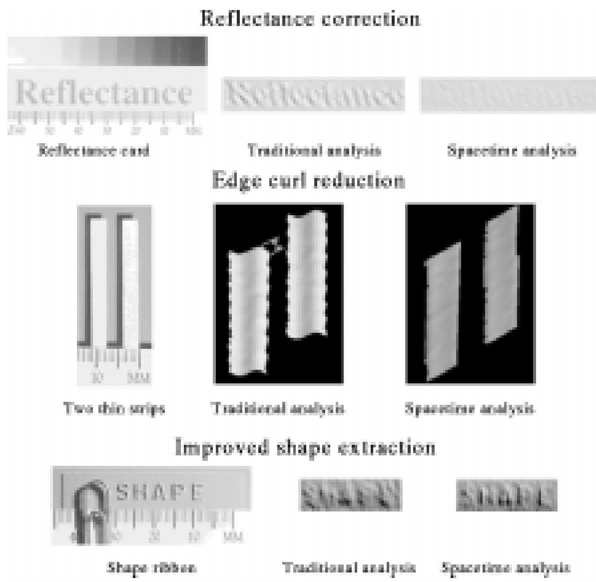

Multi-spot and multi-stripe triangulation

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

Trade off depth-of-field for speed.

Problem: ambiguity.

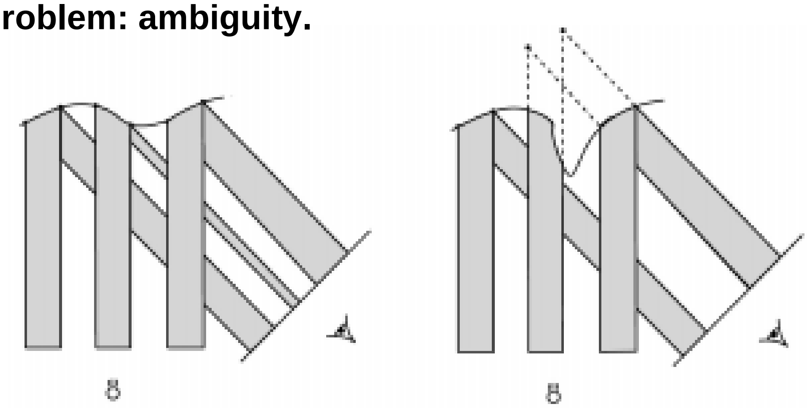

Binary coded illumination

Alternative: resolve visibility hierarchically (logN).

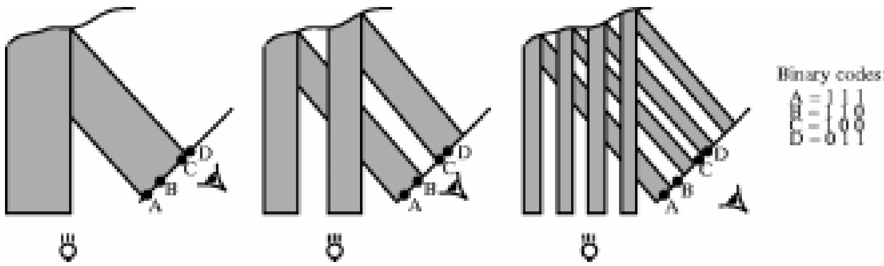

Moire

Moire methods extract shape from interference patterns:

- ***Illuminate a surface with a periodic grating.***
- ***Capture image as seen at an angle through another grating.***

  => interference pattern, phase encodes shape
- ***Low pass filter the image to extract the phase signal.***

Requires that the shape varies slowly so that phase is low frequency, much lower than grating frequency.

Example: shadow moire

Shadow moire:

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

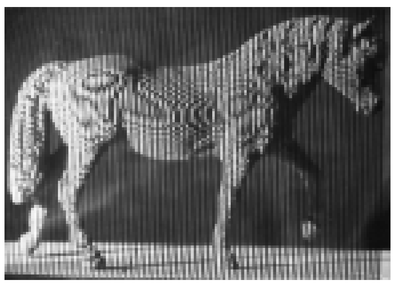

Shadow moire

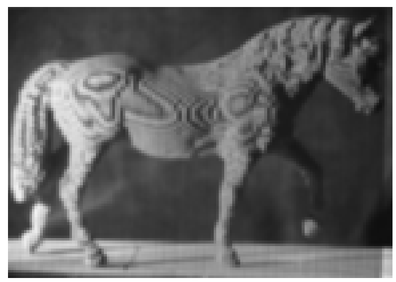

Filtered image

Active stereo

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

Active stereo simplifies feature finding with structured light. Problem: ambiguity.

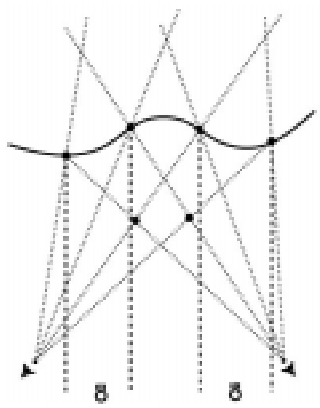

Active multi-baseline stereo

Using multiple cameras reduces likelihood of false matches.

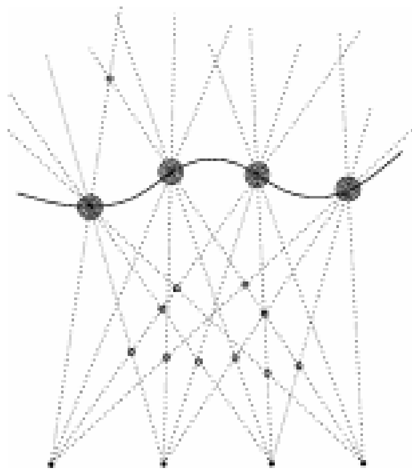

Active depth from defocus

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

The amount of blur indicates distance to the point.

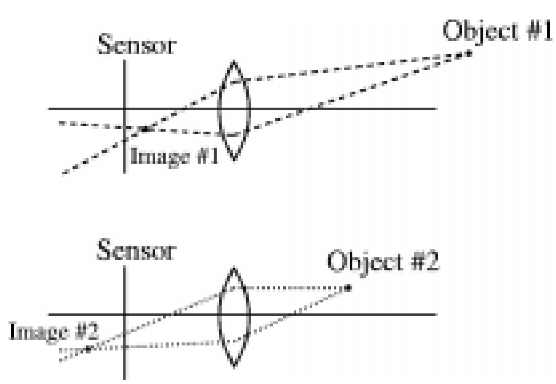

Active depth from defocus

Depth ambiguity can be resolved with two sensor planes.

Amount of defocus depends on presence of texture. Solution: project structured lighting onto surface.

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

Active depth from defocus

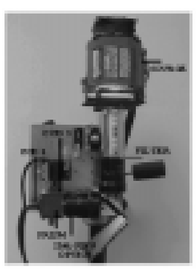

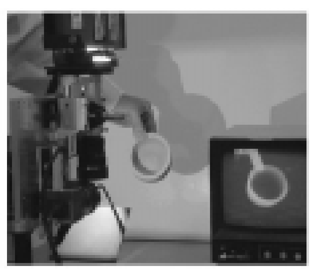

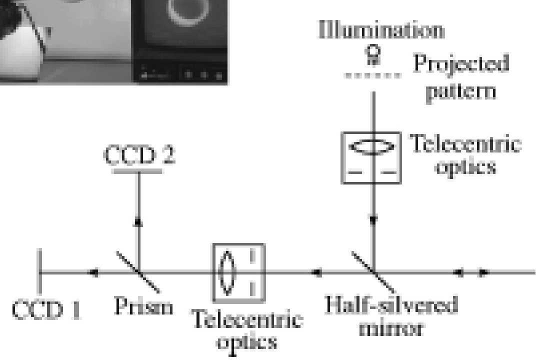

Capturing appearance

“Appearance” refers to the way an object reflects light to a viewer.

We can think of appearance under:

- ***fixed lighting***
- ***variable lighting***

Appearance under fixed lighting

Under fixed lighting, a static radiance field forms.

Each point on the object reflects a 2D (directional) radiance function.

We can acquire samples of these radiance functions with photographs registered to the geometry.

Appearance under variable lighting

To re-render the surface under novel lighting, we must capture the BRDF -- the bi-directional reflectance distribution function.

In the general case, this problem is *hard*:

- ***The BRDF is a 4D function -- may need many samples.***
- ***Interreflections imply the need to perform difficult inverse rendering calculations.***

Here, we mention ways of capturing the data needed to estimate the BRDF.

BRDF capture

To capture the BRDF, we must acquire images of the surface under known lighting conditions.

[Sato'97] captures color images with point source illumination. The camera and light are calibrated, and pose is determined by a robot arm.

[Baribeau'92] uses a white laser that is also used for optical triangulation. Reflectance samples are registered to range samples.

Key advantage: minimizes interreflection.

Bibliography

Baribeau, R., Rioux, M., and Godin, G., "Color reflectance modeling using a polychromatic laser range scanner," IEEE Transactions on PAMI, vol. 14, no. 2, Feb., 1992, pp. 263-269.

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.

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.

Sato, Y., Wheeler, M.D., Ikeuchi, K., "Object shape and reflectance modeling from observation." SIGGRAPH '97, p.379-387.