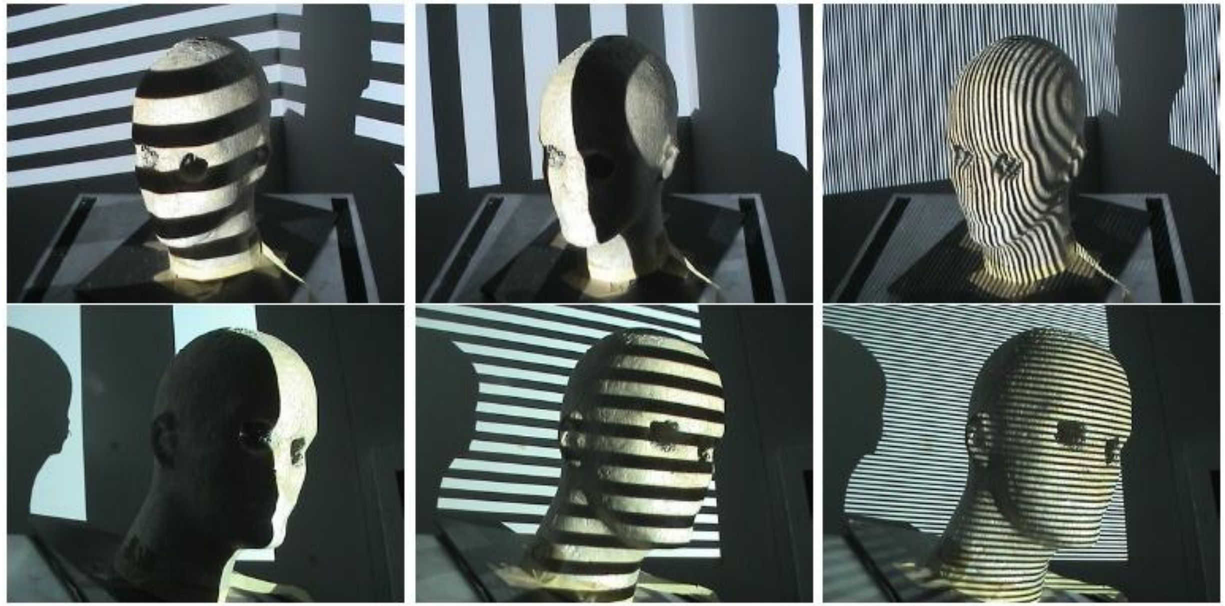

Structured Lighting

Guido Gerig

CS 6320, 3D Computer Vision

Spring 2012

(thanks: some slides S. Narasimhan CMU, Marc Pollefeys UNC)

http://www.cs.cmu.edu/afs/cs/academic/class/15385-s06/lectures/ppts/lec-17.ppt

Real-Time 3D Model Acquisition

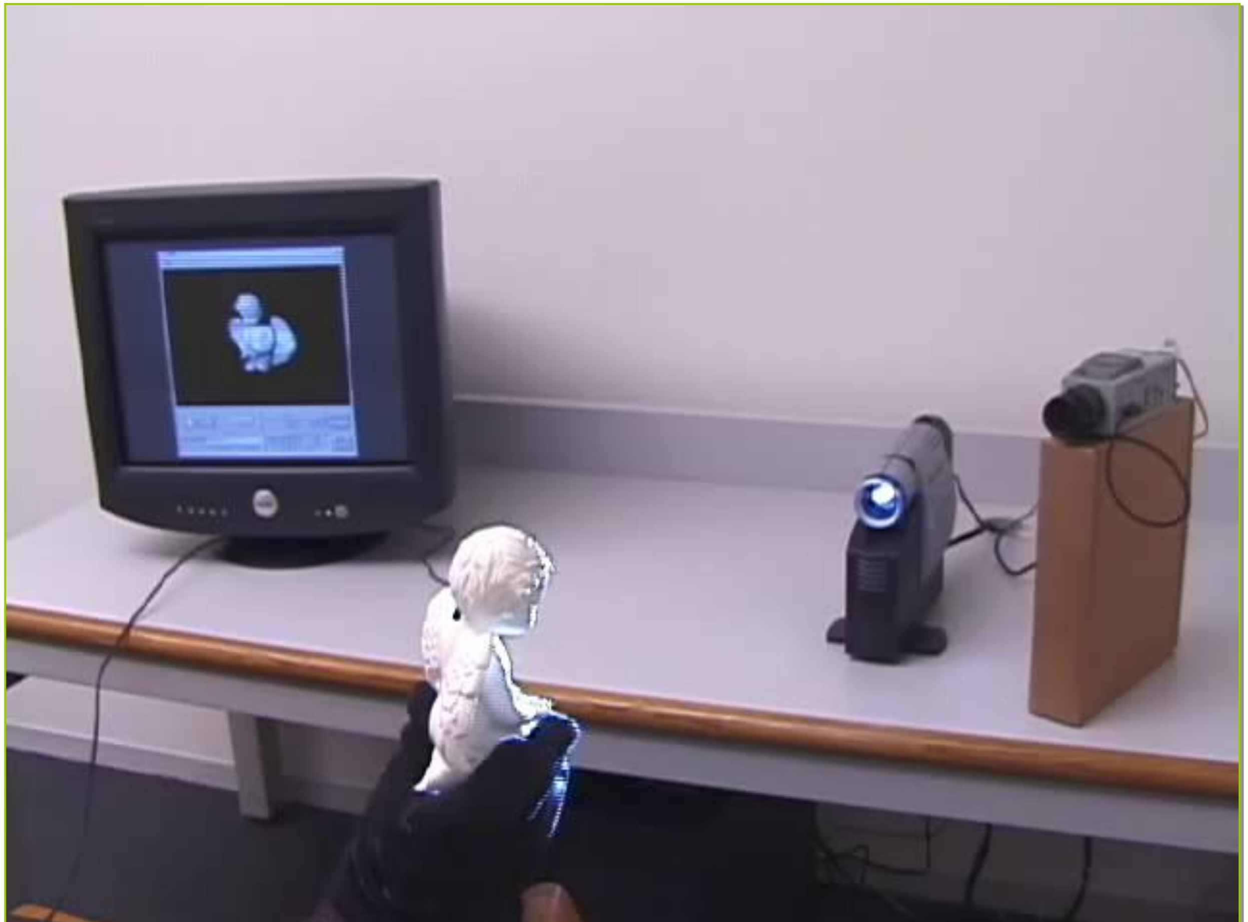

Link:
http://graphics.stanford.edu/papers/rt_model/

http://graphics.stanford.edu/papers/rt_model/

The SIGGRAPH Paper:
Full paper as PDF.
One-page abstract and Figure 1 as PDF.
Two-page abstract and Figure 1 as PDF.
A 5-minute video describing the system:
AVI file, 640 x 480 pixels (19MB)
RealVideo stream, 640 x 480 pixels, 1536 kbs
RealVideo stream, 320 x 240, 56 - 904 kbs
SIGGRAPH 2002 talk:
Talk as PPT
Embedded video clip: sig02_begin_m.avi
Embedded video clip: sig02_recap.avi
Embedded video clip: turtle2.avi

A Taxonomy

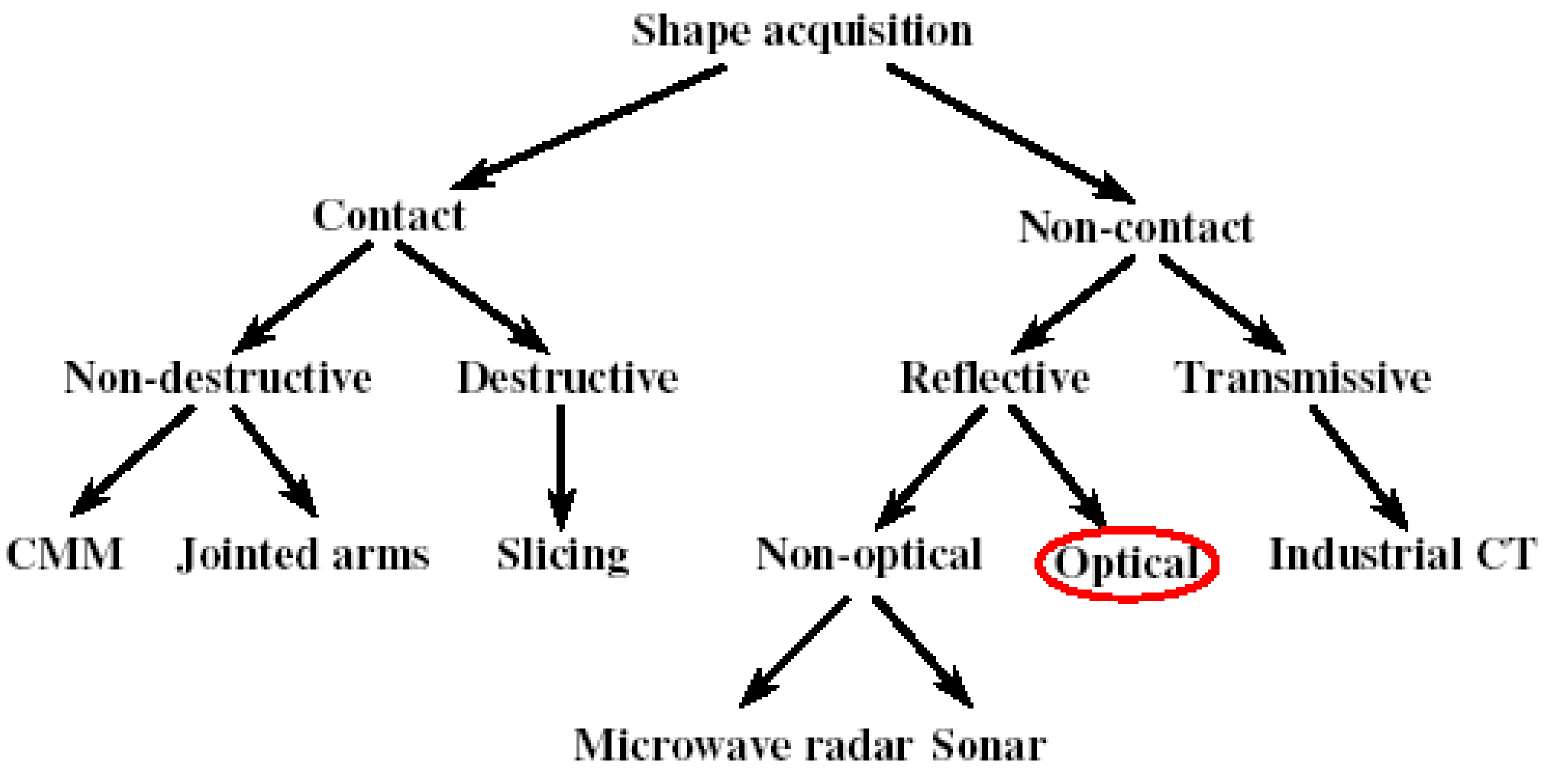
Shape acquisition
Contact
Non-contact
Non-destructive
Destructive
Reflective
Transmissive
CMM
Jointed arms
Slicing
Non-optical
Optical
Industrial CT
Microwave radar
Sonar

Excellent Additional Materials

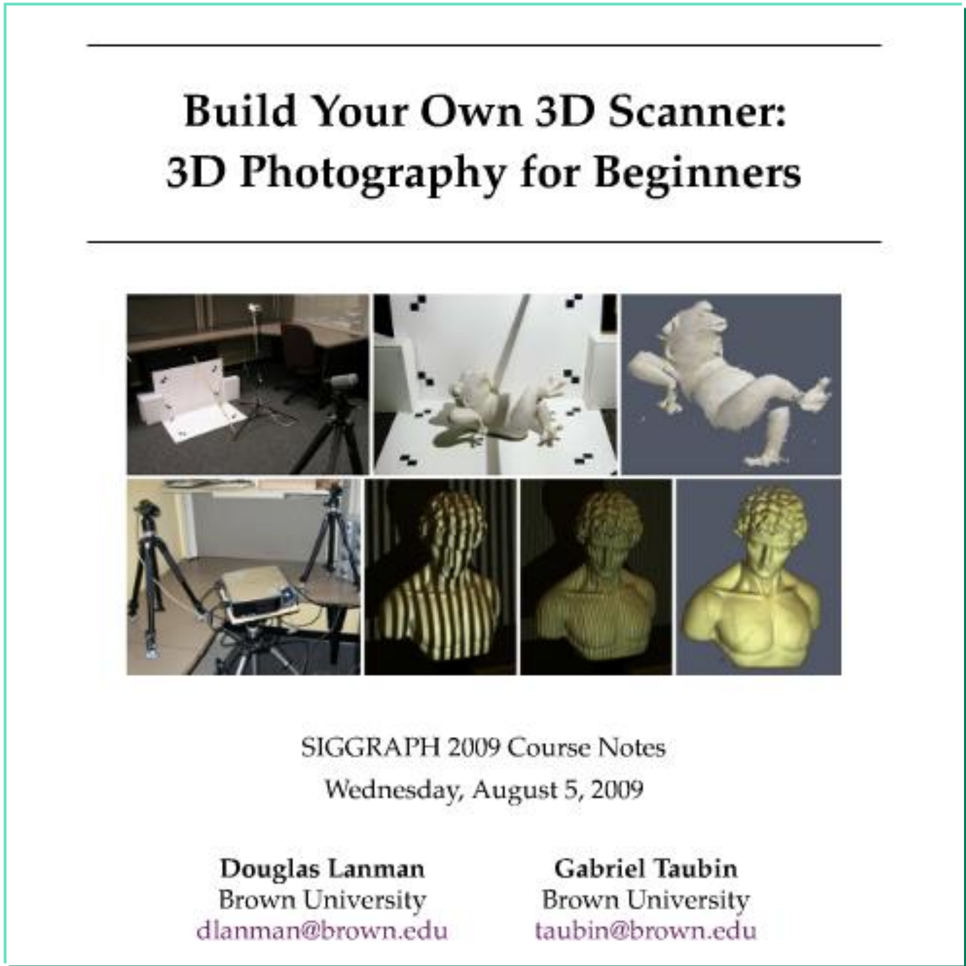

Build Your Own 3D Scanner:
3D Photography for Beginners

SIGGRAPH 2009 Course Notes
Wednesday, August 5, 2009

Douglas Lanman
Brown University
dlanman@brown.edu

Gabriel Taubin
Brown University
taubin@brown.edu

- Course notes: http://mesh.brown.edu/byo3d/notes/byo3D.pdf
- Slides: http://mesh.brown.edu/byo3d/slides.html
- Source code: http://mesh.brown.edu/byo3d/source.html

3D Scanning

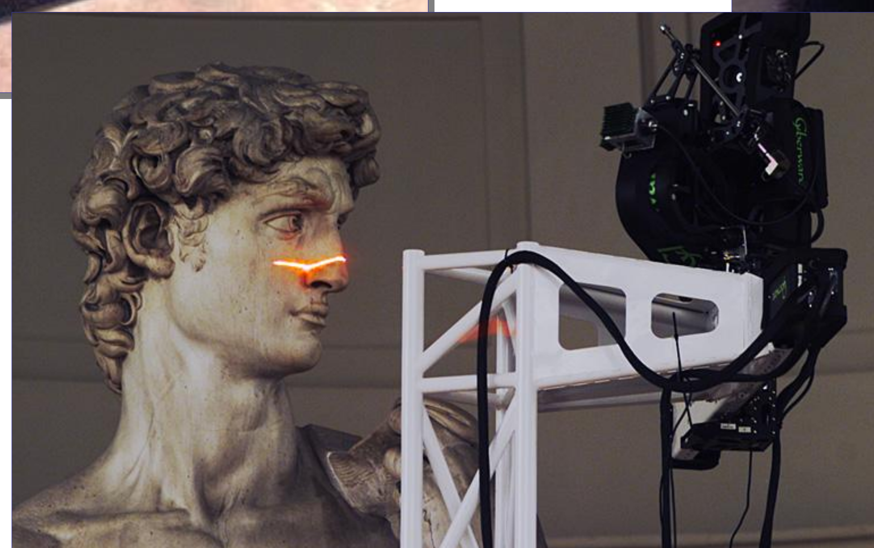

Courtesy S. Narasimhan, CMU

Typical Application

Figure 10: *The Minerva statue; note the corrosion of the bronze surface (pre-restoration status).*

Figure 11: *A digital model of the Minerva (simplified down to 1,034,029 faces).*

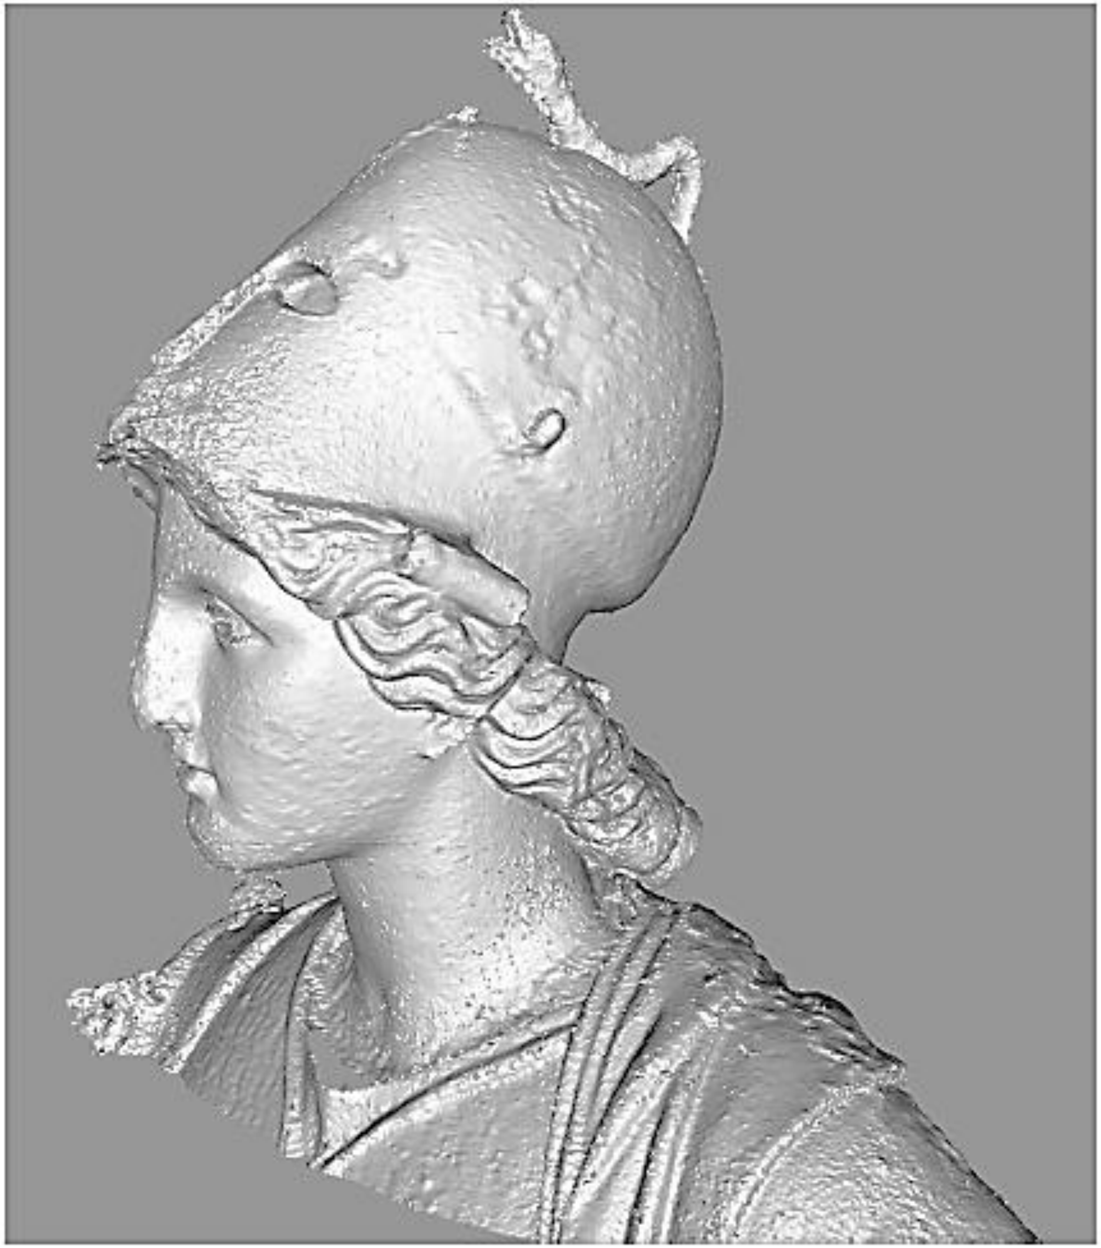

Figure 12: *The Minerva head section (simplified down to 1,094,368 faces).*

Microsoft Kinect

The Kinect combines structured light with two classic computer vision techniques: depth from focus, and depth from stereo.

Stage 1: The depth map is constructed by analyzing a speckle pattern of infrared laser light

The Kinect uses infrared laser light, with a speckle pattern

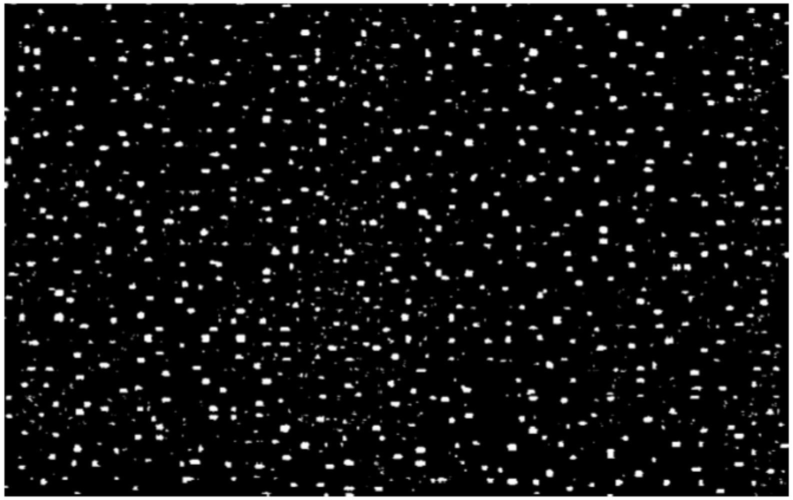

Shpunt et al, PrimeSense patent application US 2008/0106746

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

Microsoft Kinect

Inferring body position is a two-stage process: first compute a depth map, then infer body position

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

Space-time stereo

Zhang, Curless and Seitz, CVPR'03

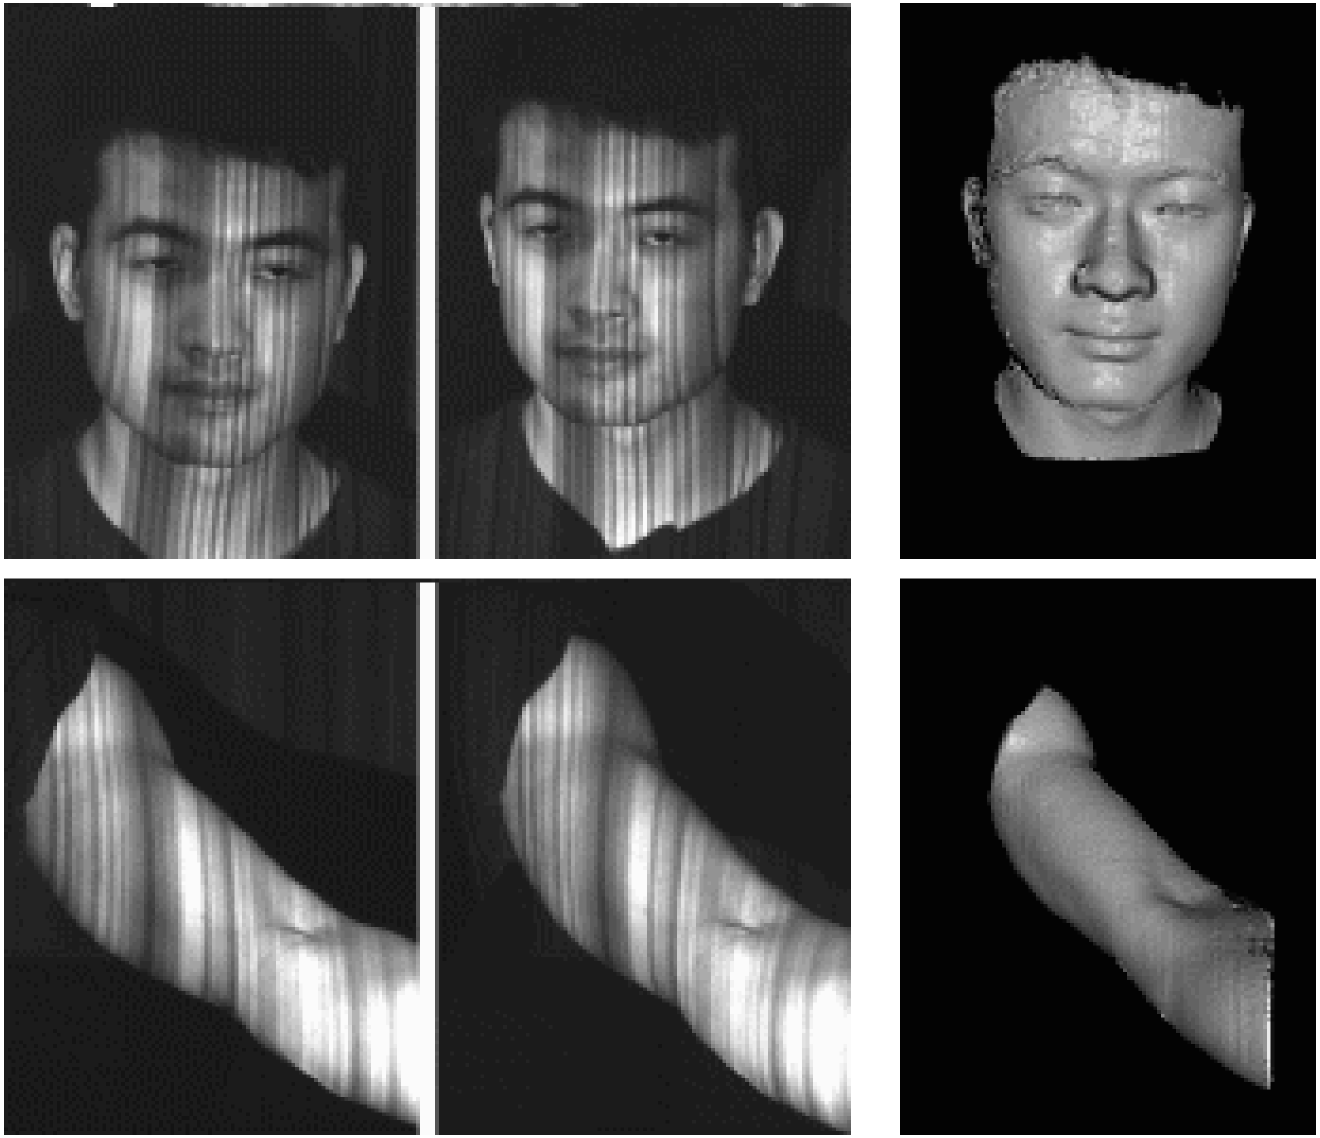

Real Time by Color Coding

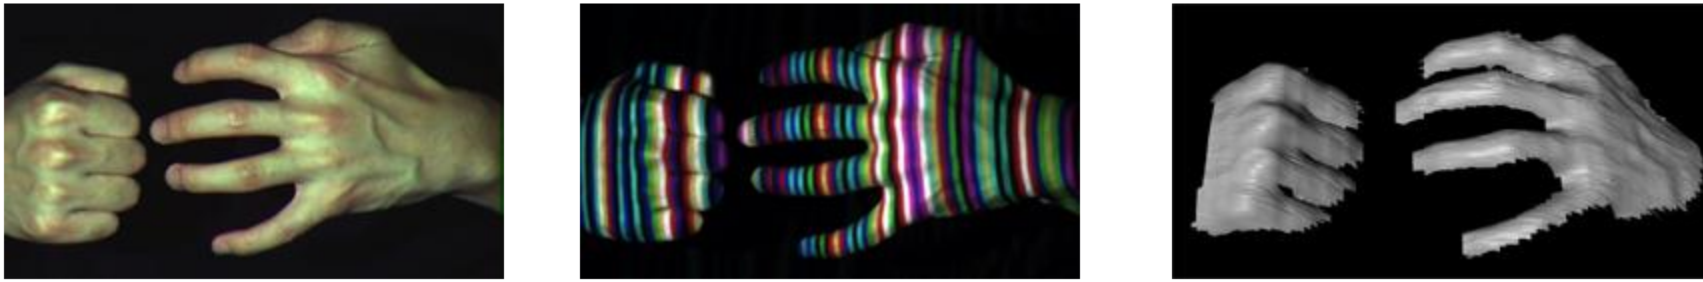

Works despite complex appearances

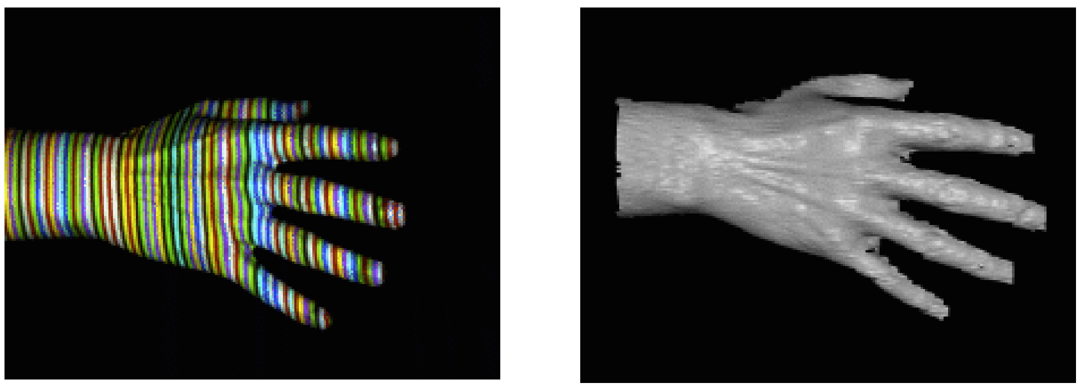

Works in real-time and on dynamic scenes

- Need very few images (one or two).
- But needs a more complex correspondence algorithm

Zhang et al, 3DPVT 2002

Concept: Active Vision

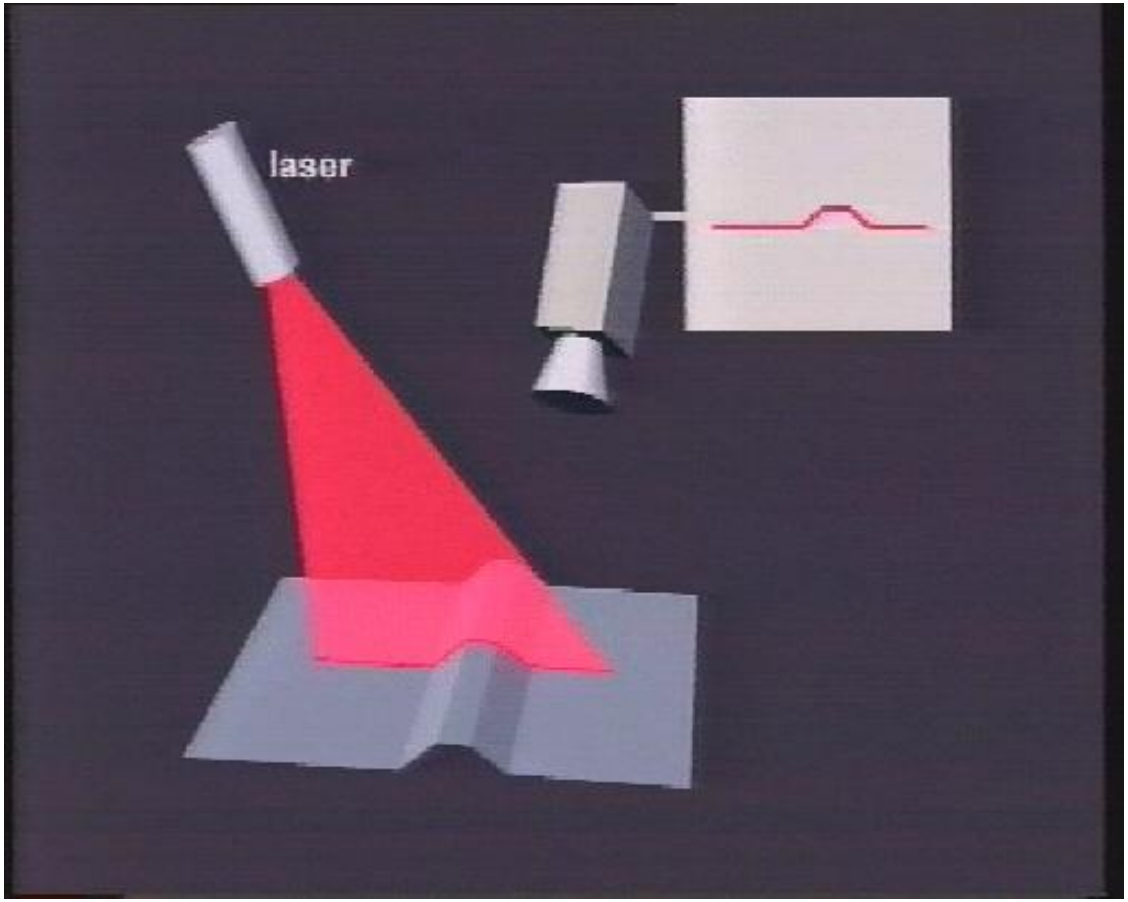

Active manipulation of scene: Project light pattern on object. Observe geometry of pattern via camera → 3D geometry

Passive triangulation: Stereo vision

- Correspondence problem
- Geometric constraints ⇒

  search along epipolar lines
- 3D reconstruction of matched pairs by triangulation

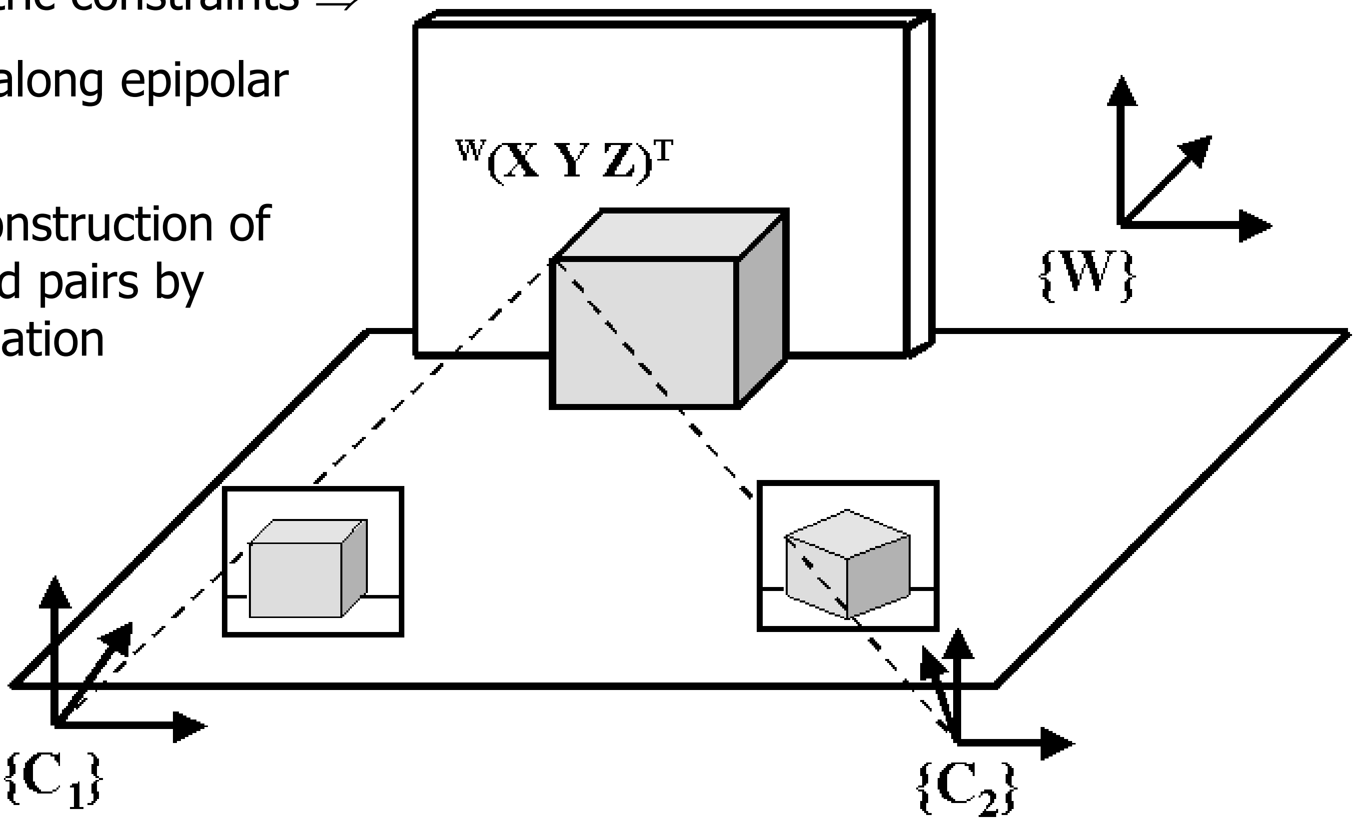

Active triangulation: Structured light

- One of the cameras is replaced by a light emitter
- Correspondence problem is solved by searching the pattern in the camera image (pattern decoding)
- No geometric constraints

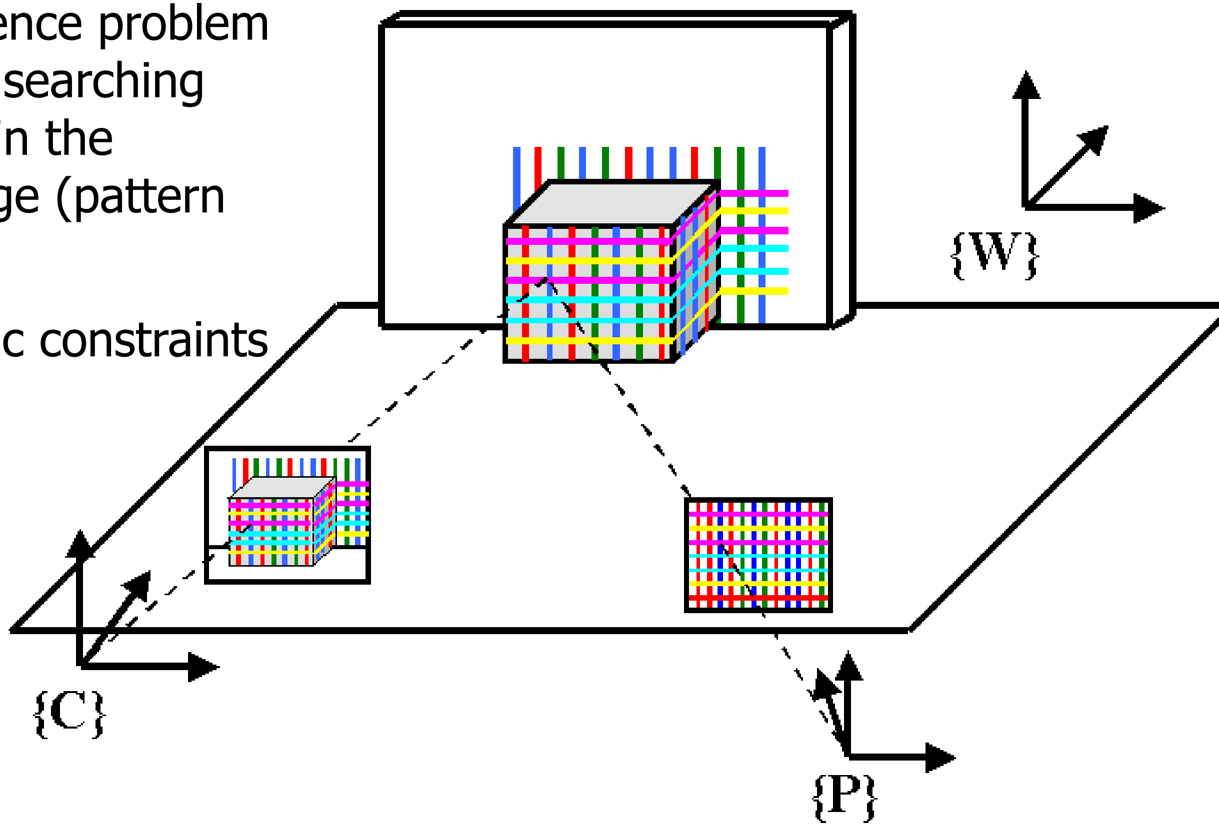

Overview

- Background
- General Setup
- Light Point Projection 2D and 3D
- Light Stripe Projection
- Static Light Pattern Projection
  - Binary Encoded Light Stripes
  - Segmenting Stripes
- 3D Photography on Your Desk

General Setup

- one camera
- one light source
  - types
    - slide projector
    - laser
  - projection
    - spot
    - stripe
    - pattern

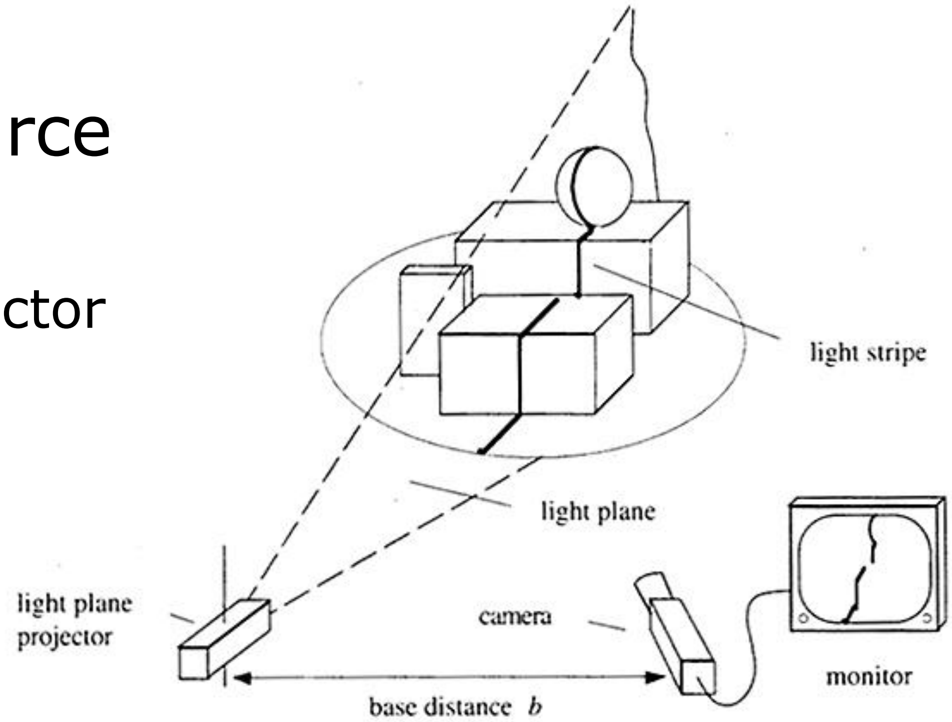

Figure 9.6: Image acquisition set-up for the light stripe projection technique.

Light Spot Projection 2D

Assume point-wise illumination by laser beam, only 2D

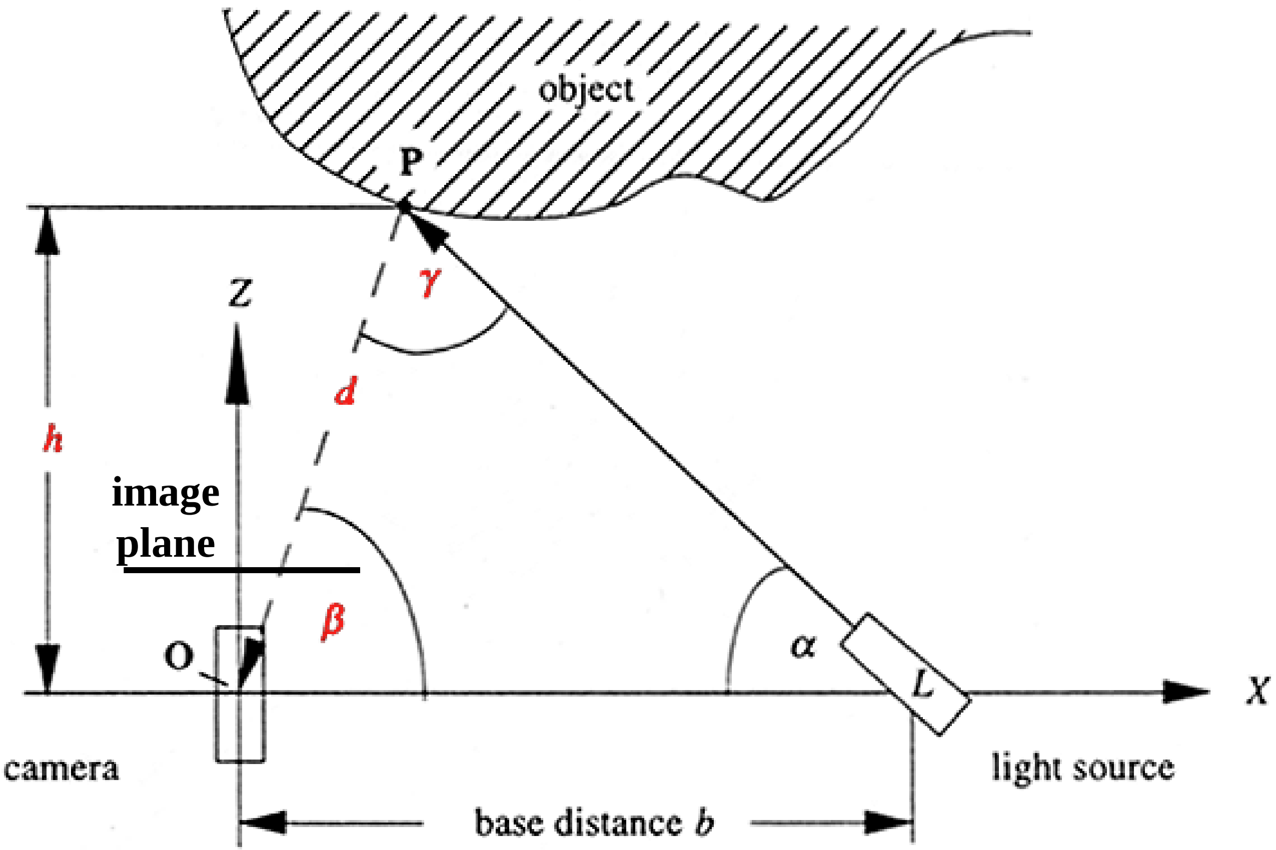

Light Spot Projection 2D

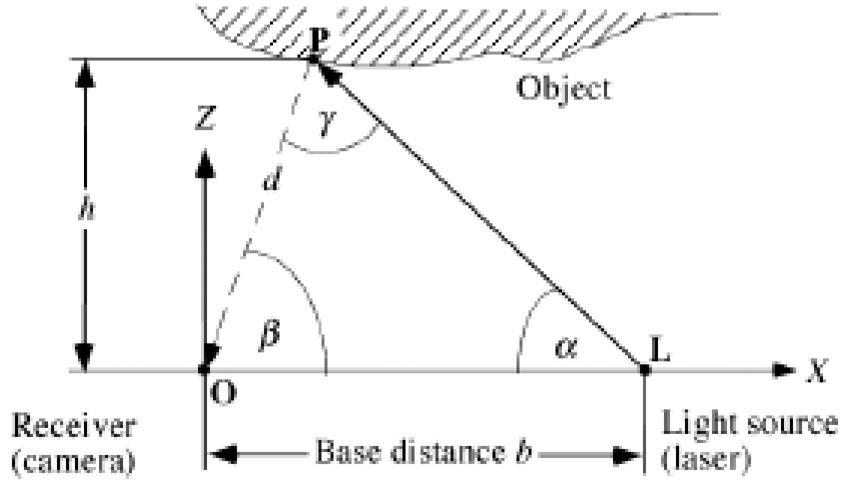

O, **L**, and **P** define a triangle, and we determine the position of P by *triangulation*, using basic formulas about triangles such as the law of sines:

$$\frac{d}{\sin\alpha} = \frac{b}{\sin\gamma}$$

It follows that

$$d = \frac{b \cdot \sin\alpha}{\sin\gamma} = \frac{b \cdot \alpha}{\sin(\pi - \alpha - \beta)} = \frac{b \cdot \alpha}{\sin(\alpha + \beta)}$$

and, finally, $\mathbf{P} = (d \cdot \cos\beta, d \cdot \sin\beta)^T$. Note that β is determined by the position of the projected (illuminated) point **P** in the 1D image.

Light Spot Projection 2D

- Coordinates found by triangulation
  - β can be found by projection geometry
  - $d = b*\sin(\alpha)/\sin(\alpha + \beta)$

  - $X_0 = d*\cos(\beta)$
  - $Z_0 = h = d*\sin(\beta)$

- Concept:
  - known b and α
  - β defined by projection geometry
  - Given image coordinate u and focal length f -> calculate β
  - Given b, α, β -> <u>calculate d</u>

Light Spot Projection 3D

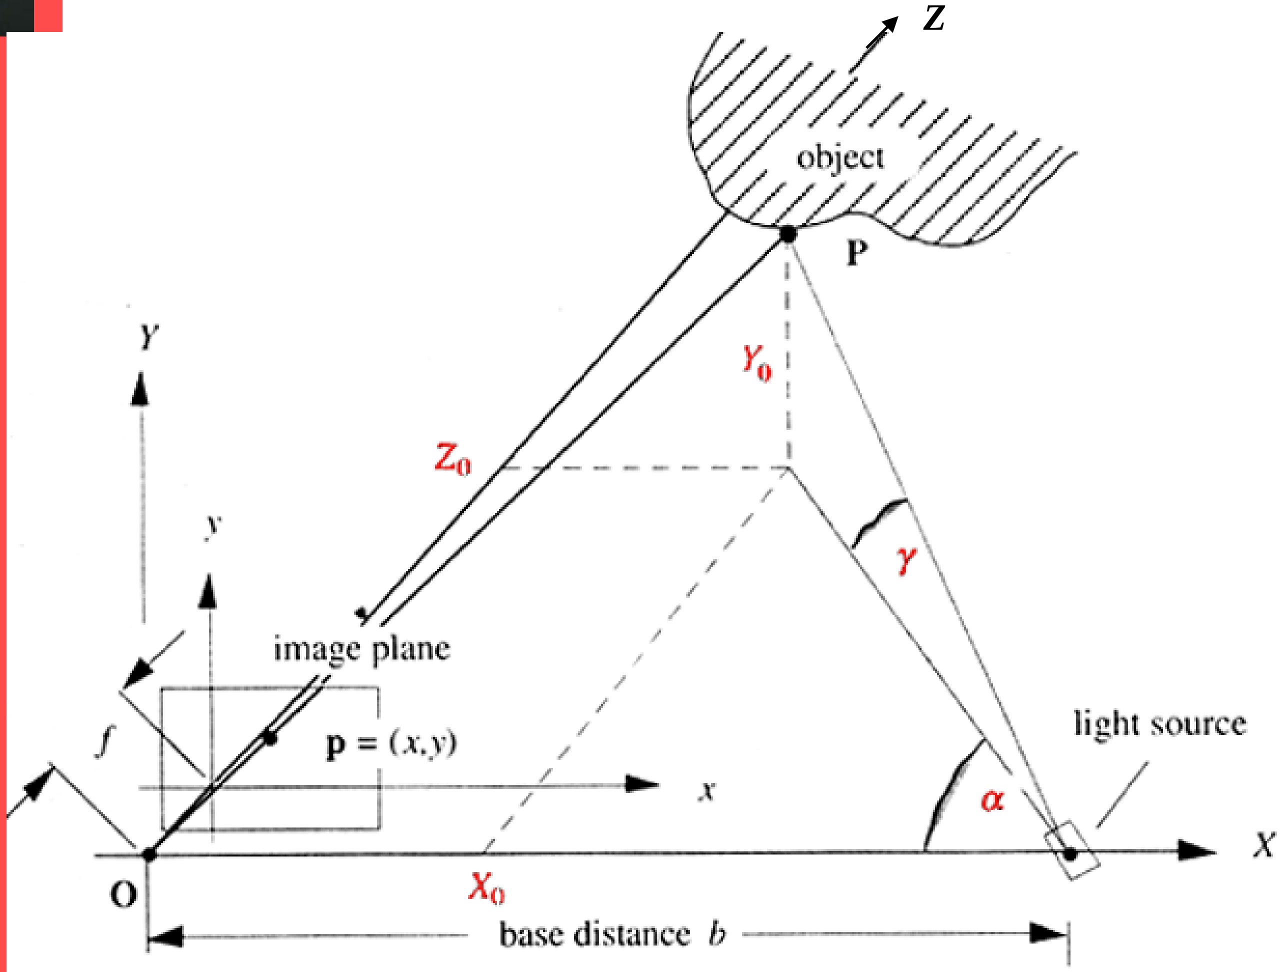

Light Spot Projection 3D

Now we consider the 3D case. Assume that α and b are given.

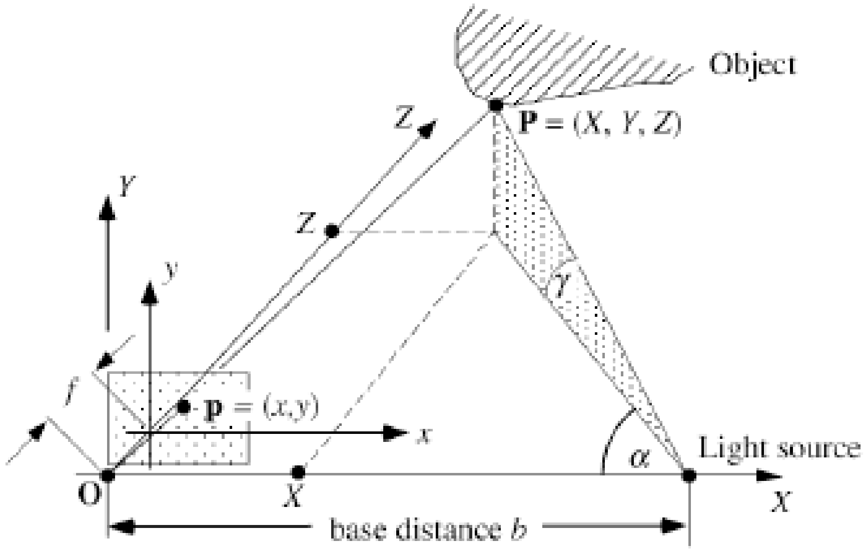

The ray theorem (of central projection) tells us that $\frac{X}{x} = \frac{Z}{f} = \frac{Y}{y}$, and from the trigonometry of right triangles we know that $\tan\alpha = \frac{Z}{b-X}$. It follows that

$$Z = \frac{X}{x} \cdot f = \tan\alpha \cdot (b - X) \quad \text{and} \quad X \cdot \left(\frac{f}{x} + \tan\alpha\right) = \tan\alpha \cdot b$$

The solution is

$$X = \frac{\tan\alpha \cdot b \cdot x}{f + x \cdot \tan\alpha}, \; Y = \frac{\tan\alpha \cdot b \cdot y}{f + x \cdot \tan\alpha}, \; Z = \frac{\tan\alpha \cdot b \cdot f}{f + x \cdot \tan\alpha}$$

Why does γ not appear in these equations?

Special Case: Light Spot Stereo

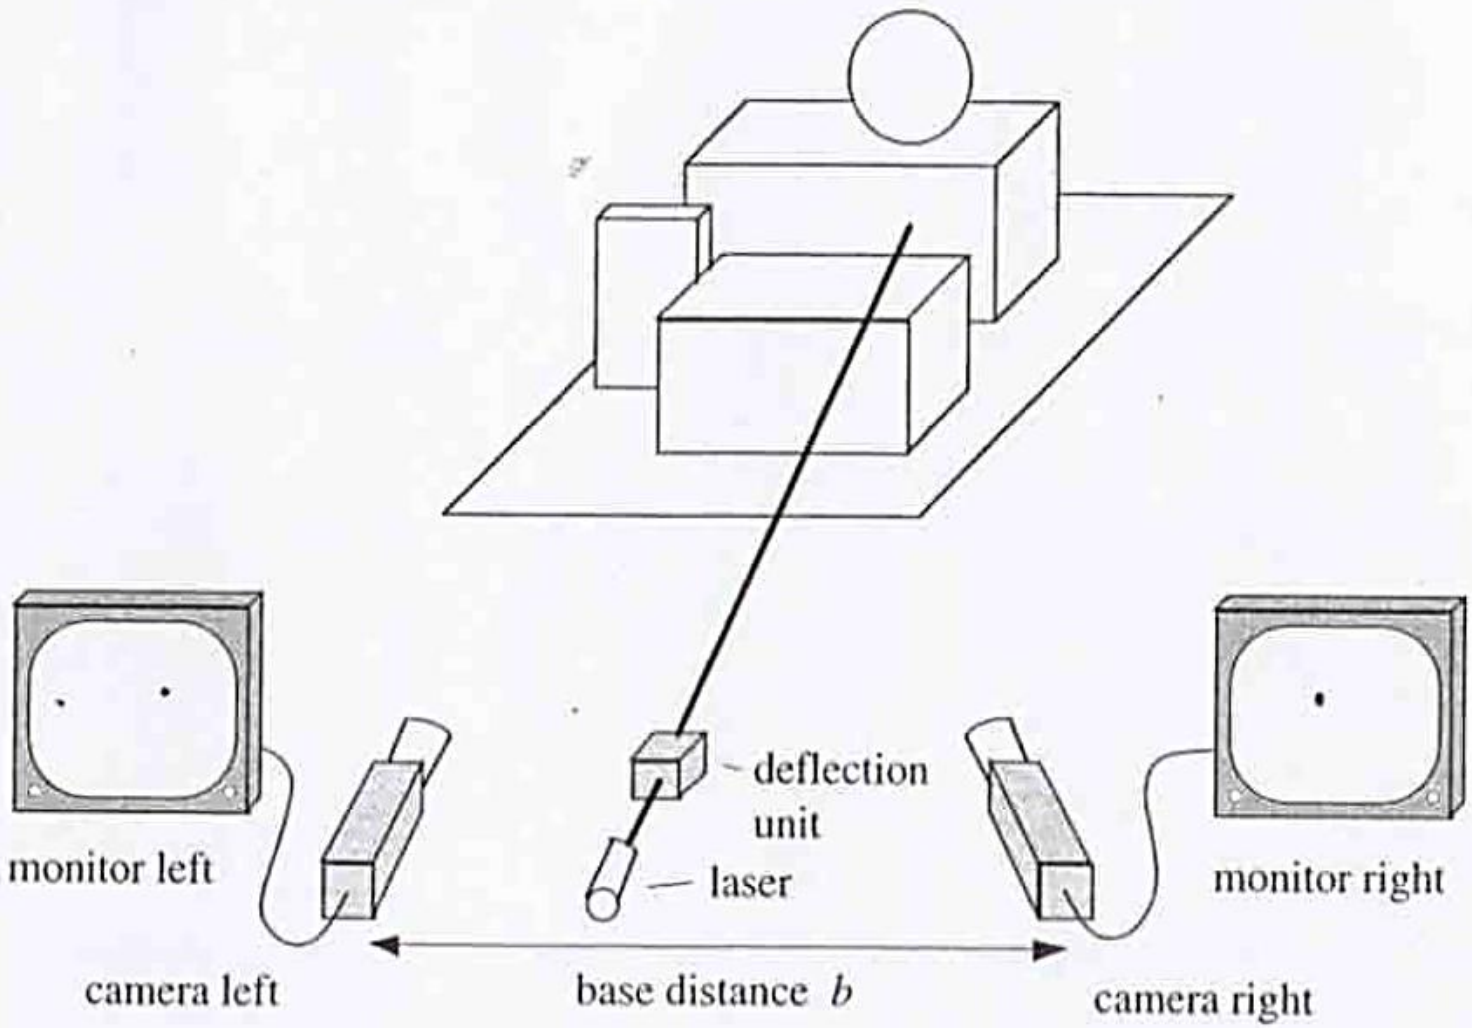

Figure 9.5: General arrangement for a method based on light spot stereo analysis.

Light Stripe Scanning – Single Stripe

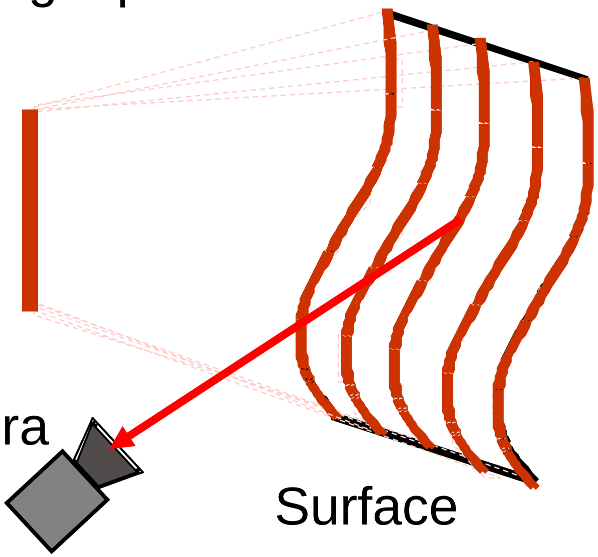

- Optical triangulation
  - Project a single stripe of laser light
  - Scan it across the surface of the object
  - This is a very precise version of structured light scanning
  - Good for high resolution 3D, but needs many images and takes time

Courtesy S. Narasimhan, CMU

Light Stripe Projection

Structured light is the projection of a light pattern (ray, plane, grid, encoded light, and so forth) under calibrated geometric conditions onto an object whose shape needs to be recovered.

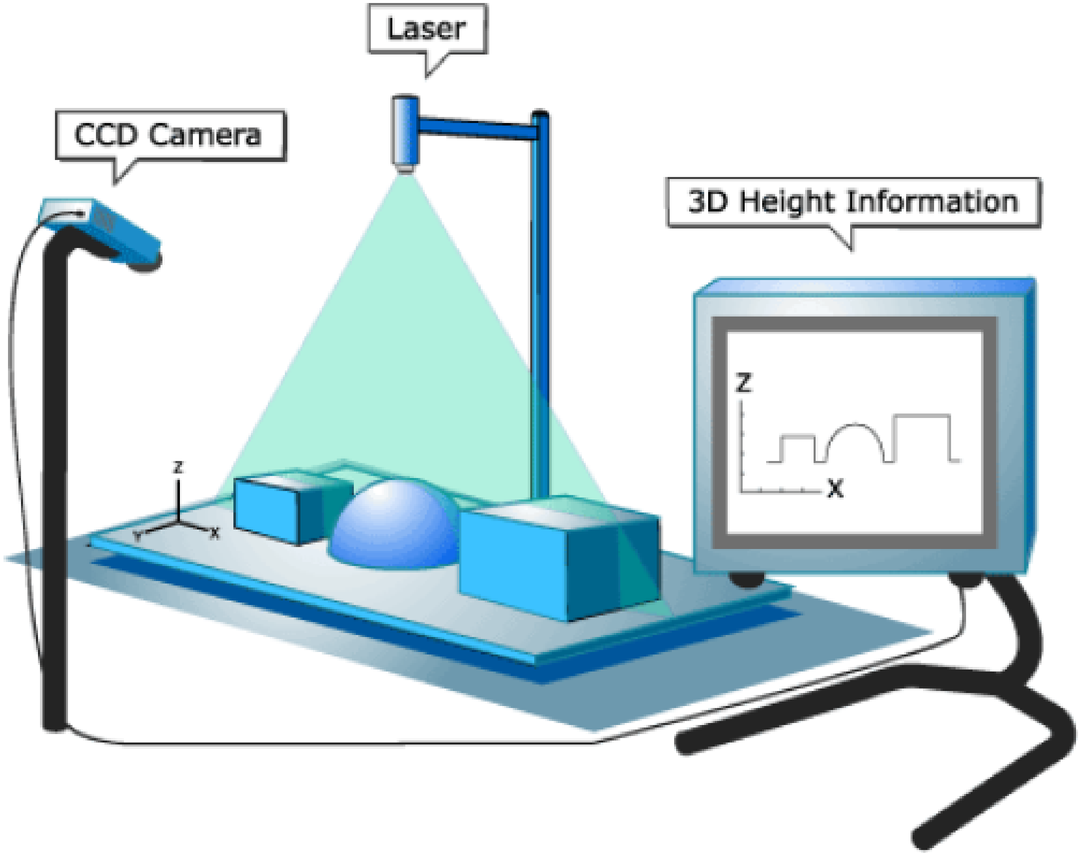

Triangulation

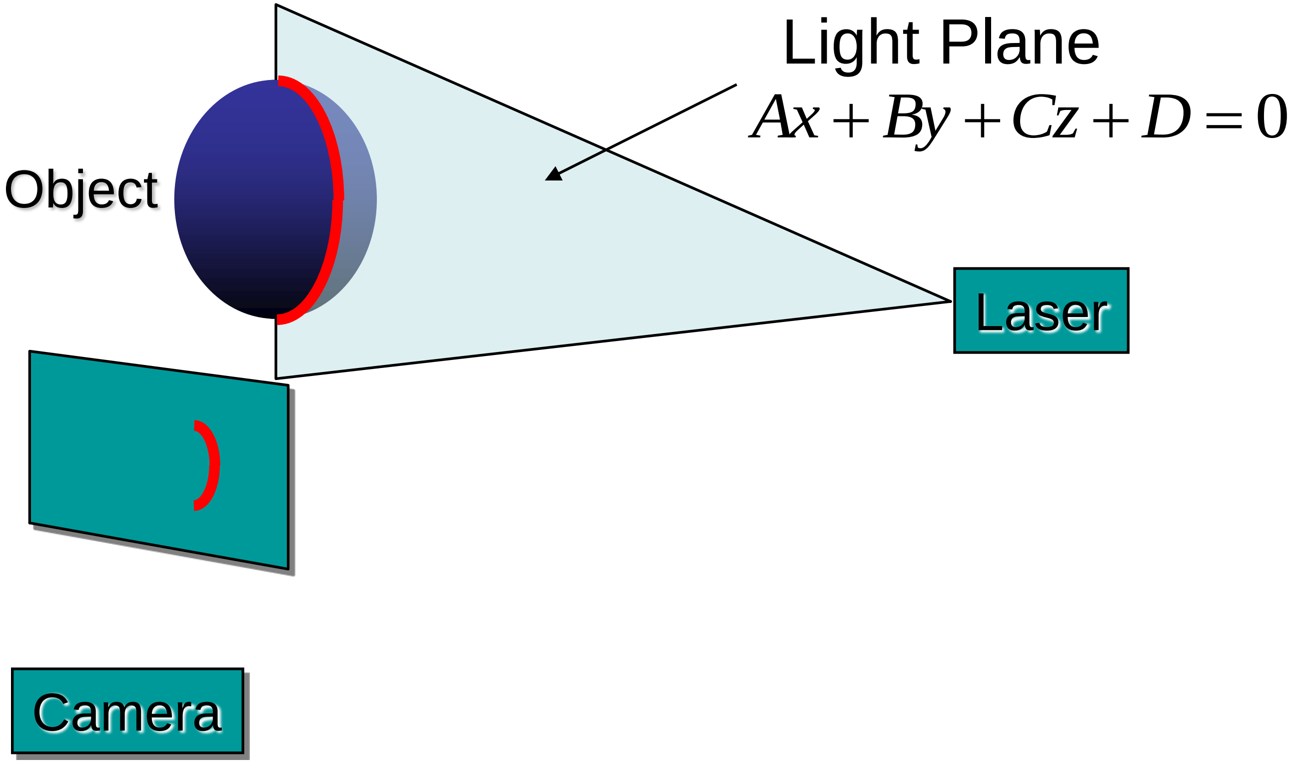

- Project laser stripe onto object

Courtesy S. Narasimhan, CMU

Triangulation

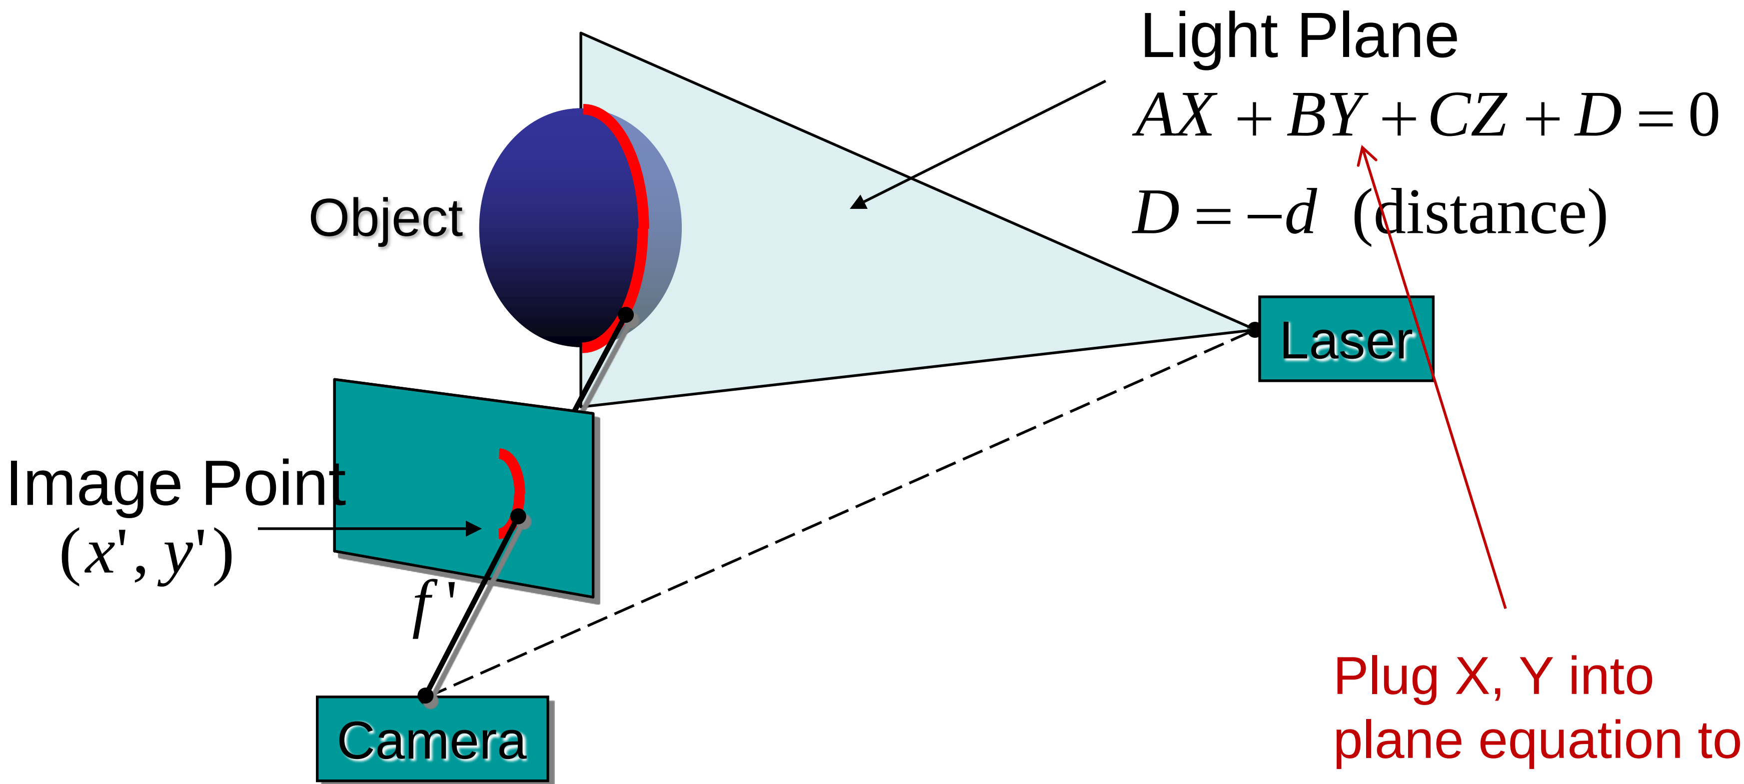

Plug X, Y into plane equation to get Z

- Depth from ray-plane triangulation:
  - Intersect camera ray with light plane

$$X = x'Z / f' \qquad Y = y'Z / f'$$

$$Z = \frac{-Df'}{Ax' + By' + Cf'}$$

Courtesy S. Narasimhan, CMU

Light Stripe Projection: Calibration

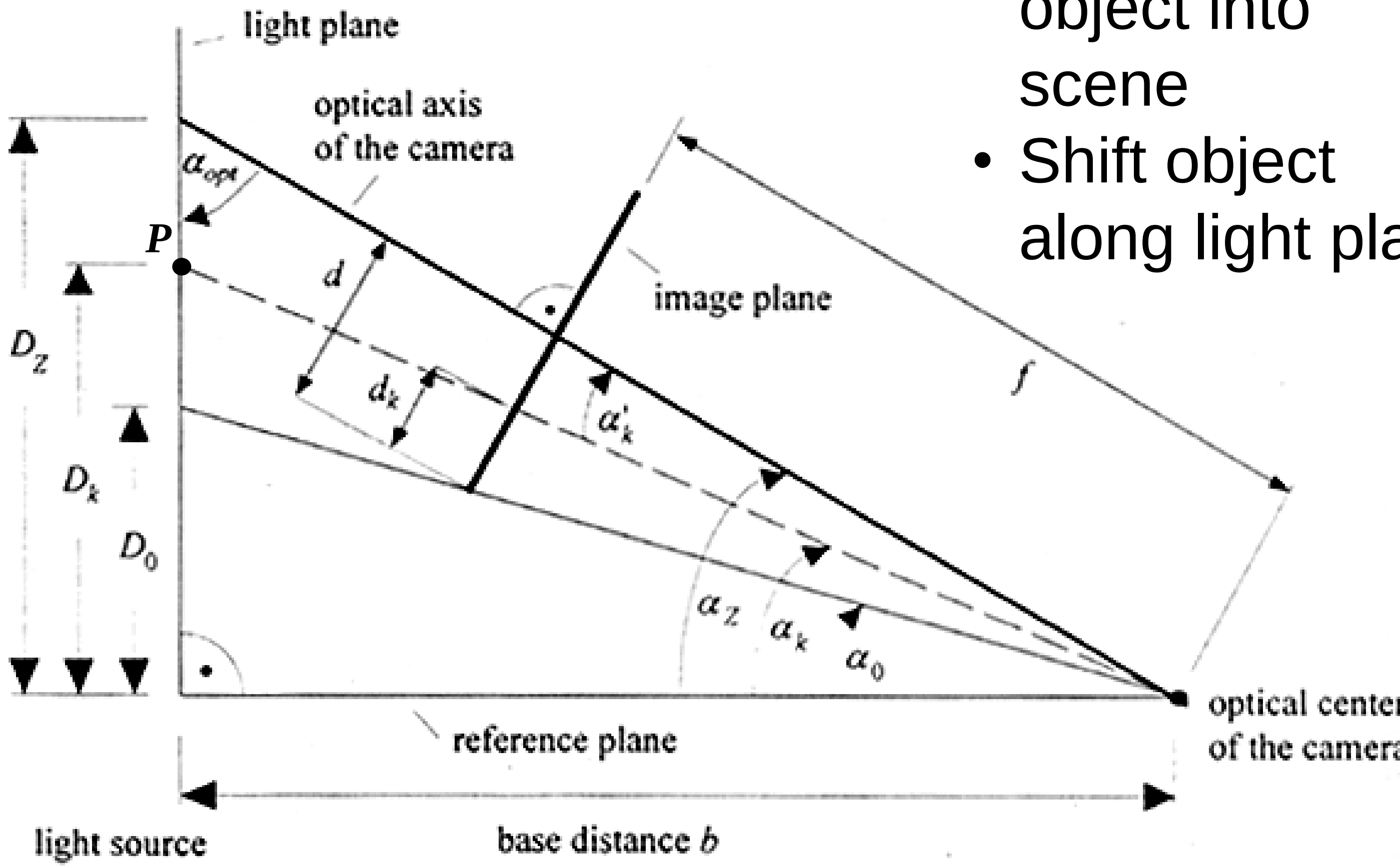

- Put calibration object into scene
- Shift object along light plane

Light Stripe Projection: Calibration

A simple method for calibrating distance for a system with a fixed light-plane is as follows: objects are assumed on a turntable, which has its rotation center at the origin of the $X_{light}Y_{light}$ plane.

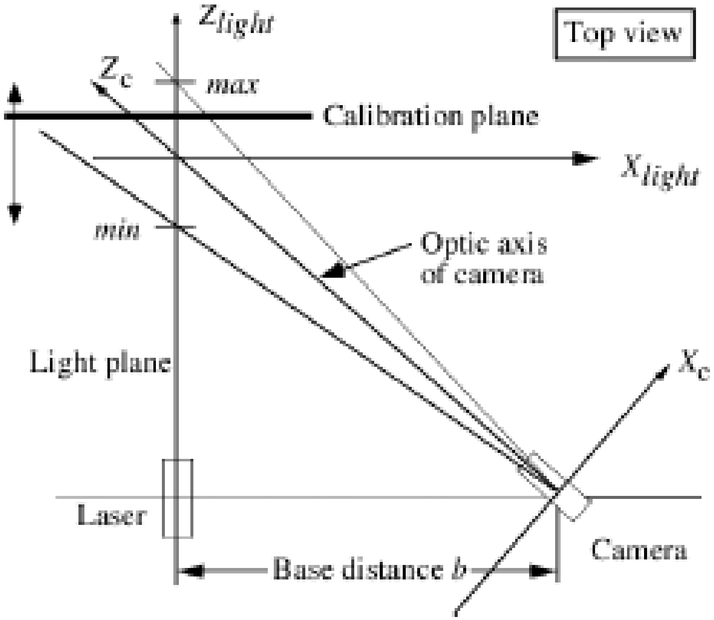

The calibration plane is now moving between two positions max and min, parallel to the X_{light}, Y_{light} plane. At max, the rightmost column in the image (of the camera) is illuminated, at min the leftmost column. For each column y we measure the Z_{light} coordinate of the calibration plane such that column y is illuminated at this position.

Straightforward: Single light stripe and rotating Object

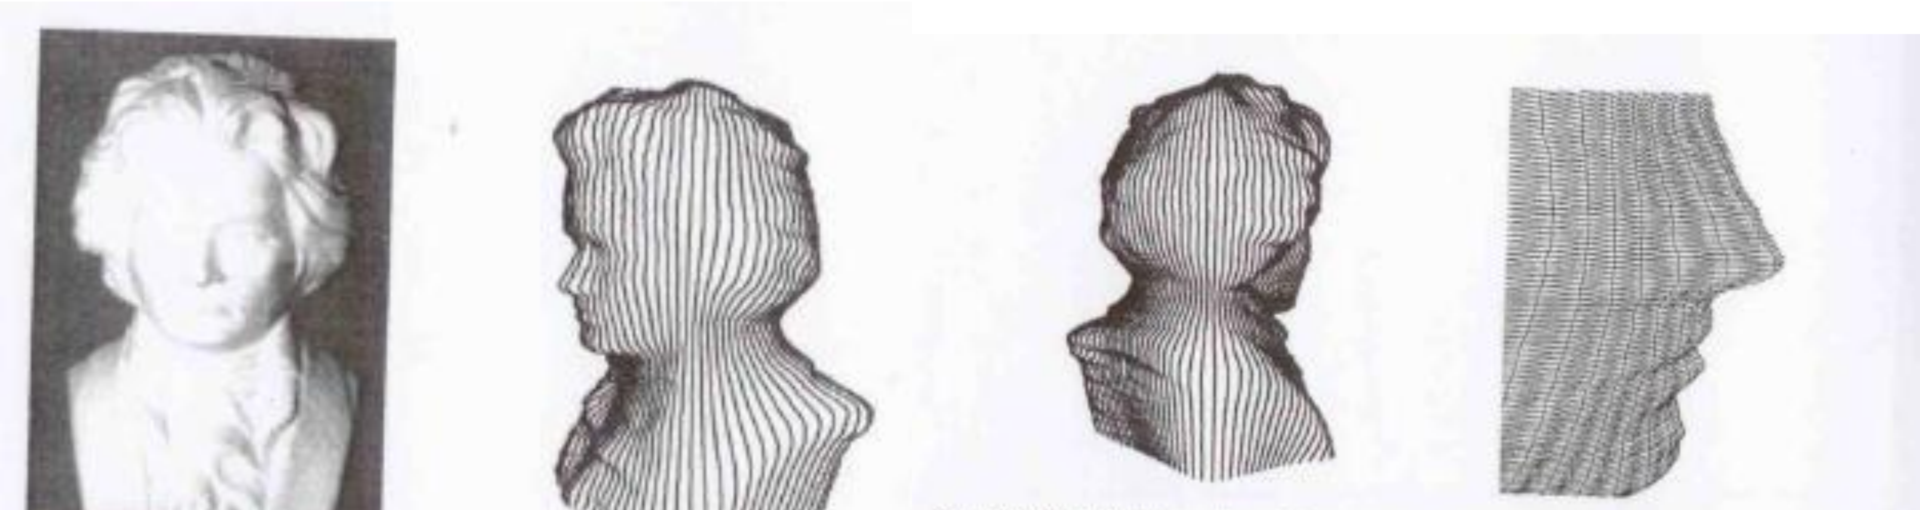

Figure 9.10: Left: Rear view of the recovered 3D profiles of the Beethoven statue shown on the left of Fig. 9.9. Right: Recovered triangular mesh around nose and mouth of the statue.

Object on turntable:

- Create P(X,Y,Z) profile for each rotation and fixed light slit
- Rotate object in discrete intervals and repeat
- Reconstruct 3D object by cylindric assembly of profiles → 3D mesh

Example: Laser scanner

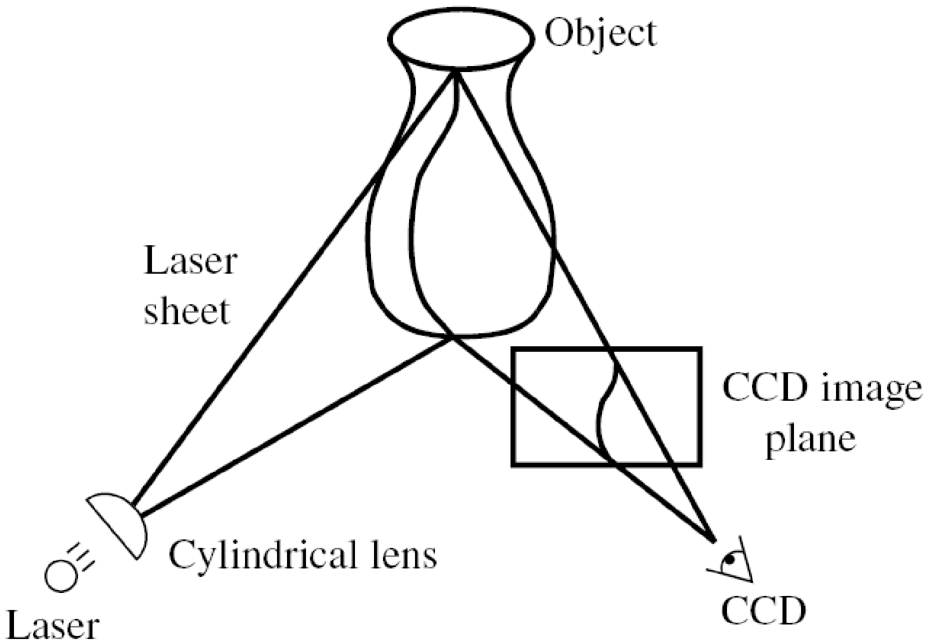

Cyberware® face and head scanner

- \+ very accurate < 0.01 mm
- − more than 10sec per scan

Portable 3D laser scanner (this one by Minolta)

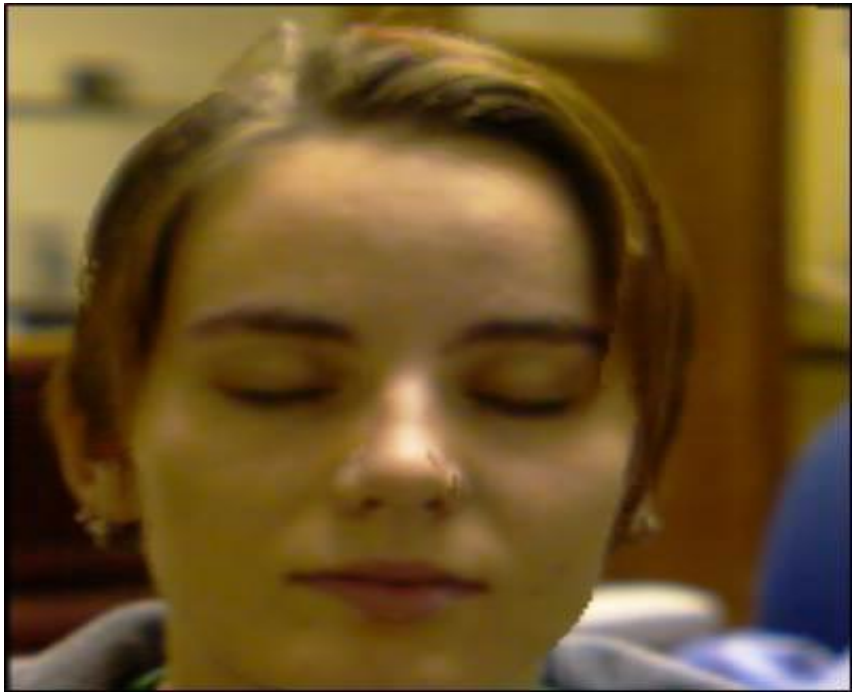

Example: Laser scanner

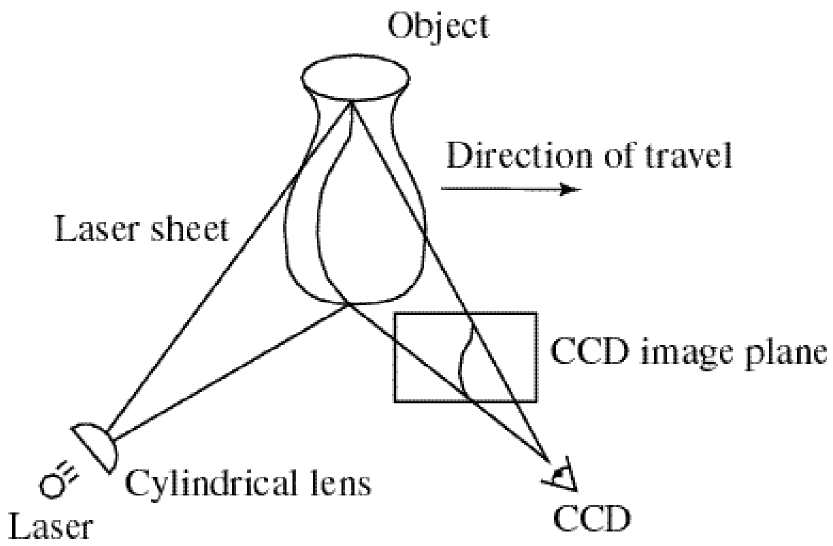

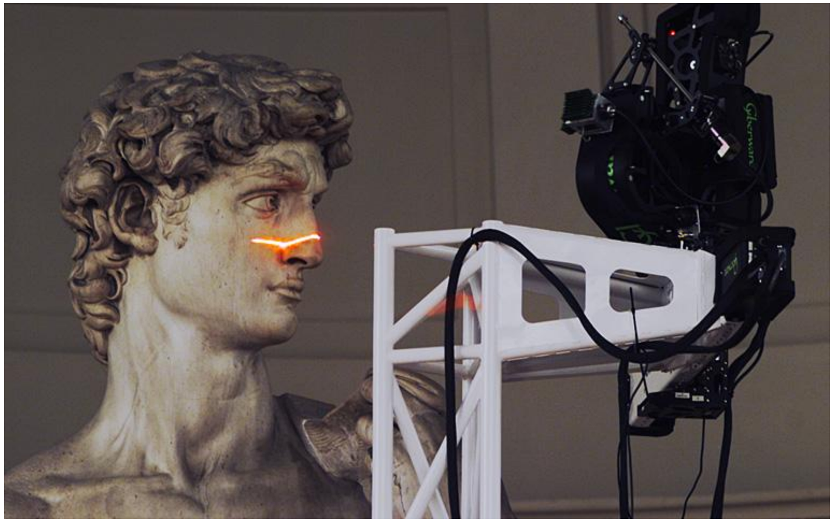

Digital Michelangelo Project

http://graphics.stanford.edu/projects/mich/

Can we do it without expensive equipment?

3D Acquisition from Shadows

3D photography on your desk

Jean-Yves Bouguet† and Pietro Perona†‡

† California Institute of Technology, 136-93, Pasadena, CA 91125, USA

‡ Università di Padova, Italy

{bouguetj,perona}@vision.caltech.edu

Abstract

A simple and inexpensive approach for extracting the three-dimensional shape of objects is presented. It is based on 'weak structured lighting'; it differs from other conventional structured lighting approaches in that it requires very little hardware besides the camera: a desk-lamp, a pencil and a checker-board. The camera faces the object, which is illuminated by the desk-lamp. The user moves a pencil in front of the light source casting a moving shadow on the object. The 3D shape of the object is extracted from the spatial and temporal location of the observed shadow. Experimental results are presented on three different scenes demonstrating that the error in reconstructing the surface is less than 1%.

Figure 1: **The general setup of the proposed method:** The camera is facing the scene illuminated by a halogen desk

Bouguet-Perona, ICCV 98

3D Photography on Your Desk

- “Cheap” method that uses very common tools to do 3D photography
- Requirements: PC, camera, stick, lamp, and a checker board
- Uses “weak structured light” approach

The idea

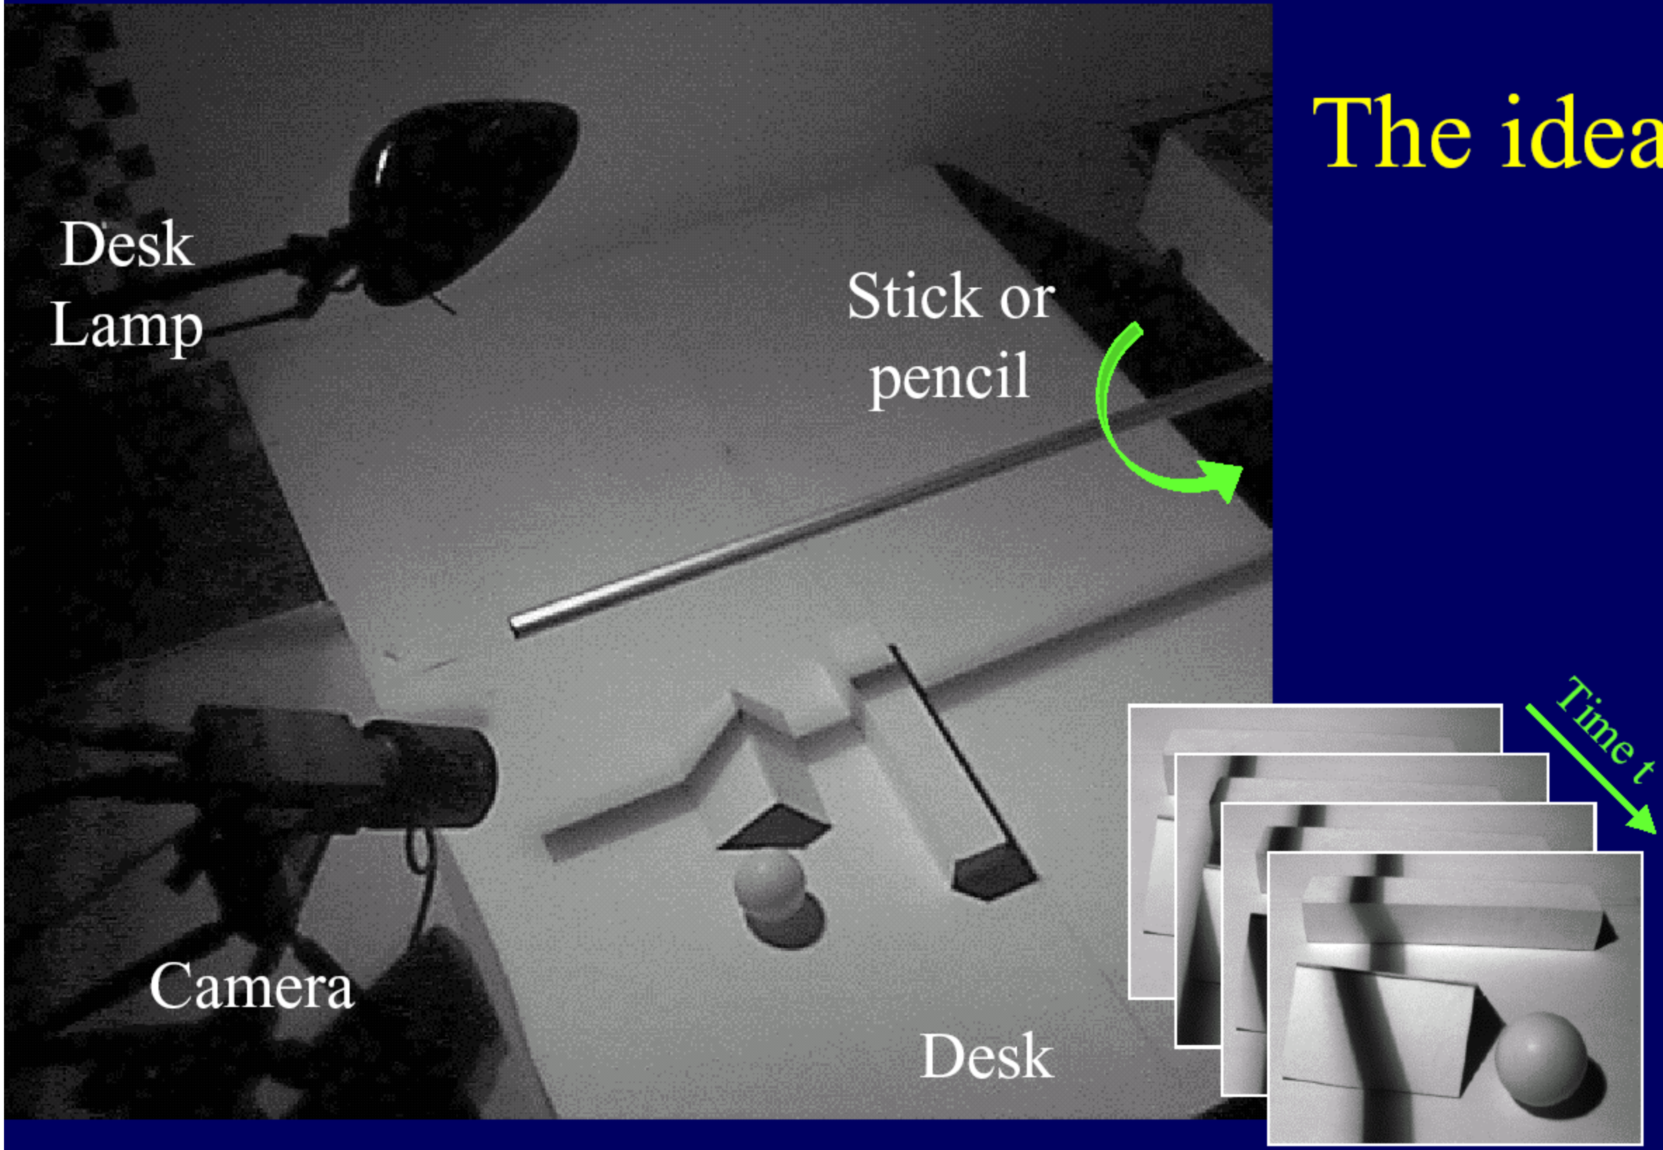

The geometry

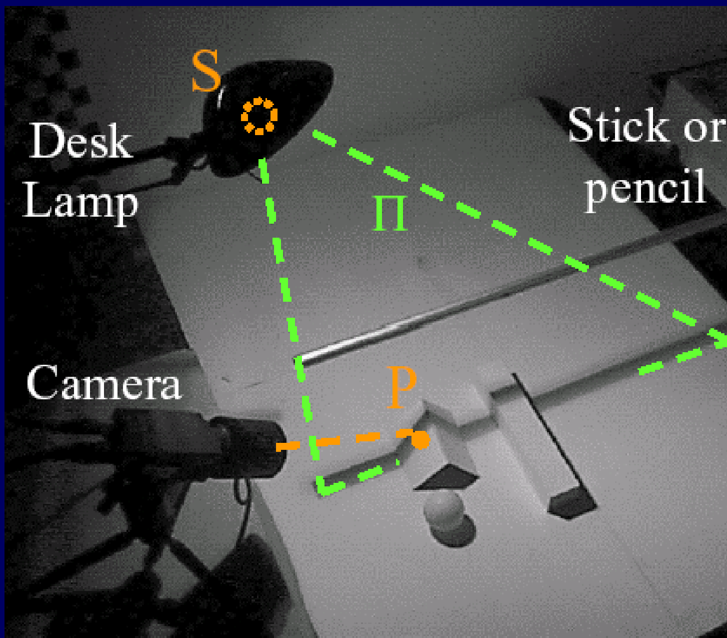

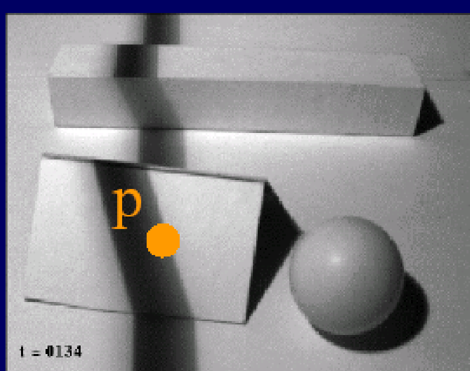

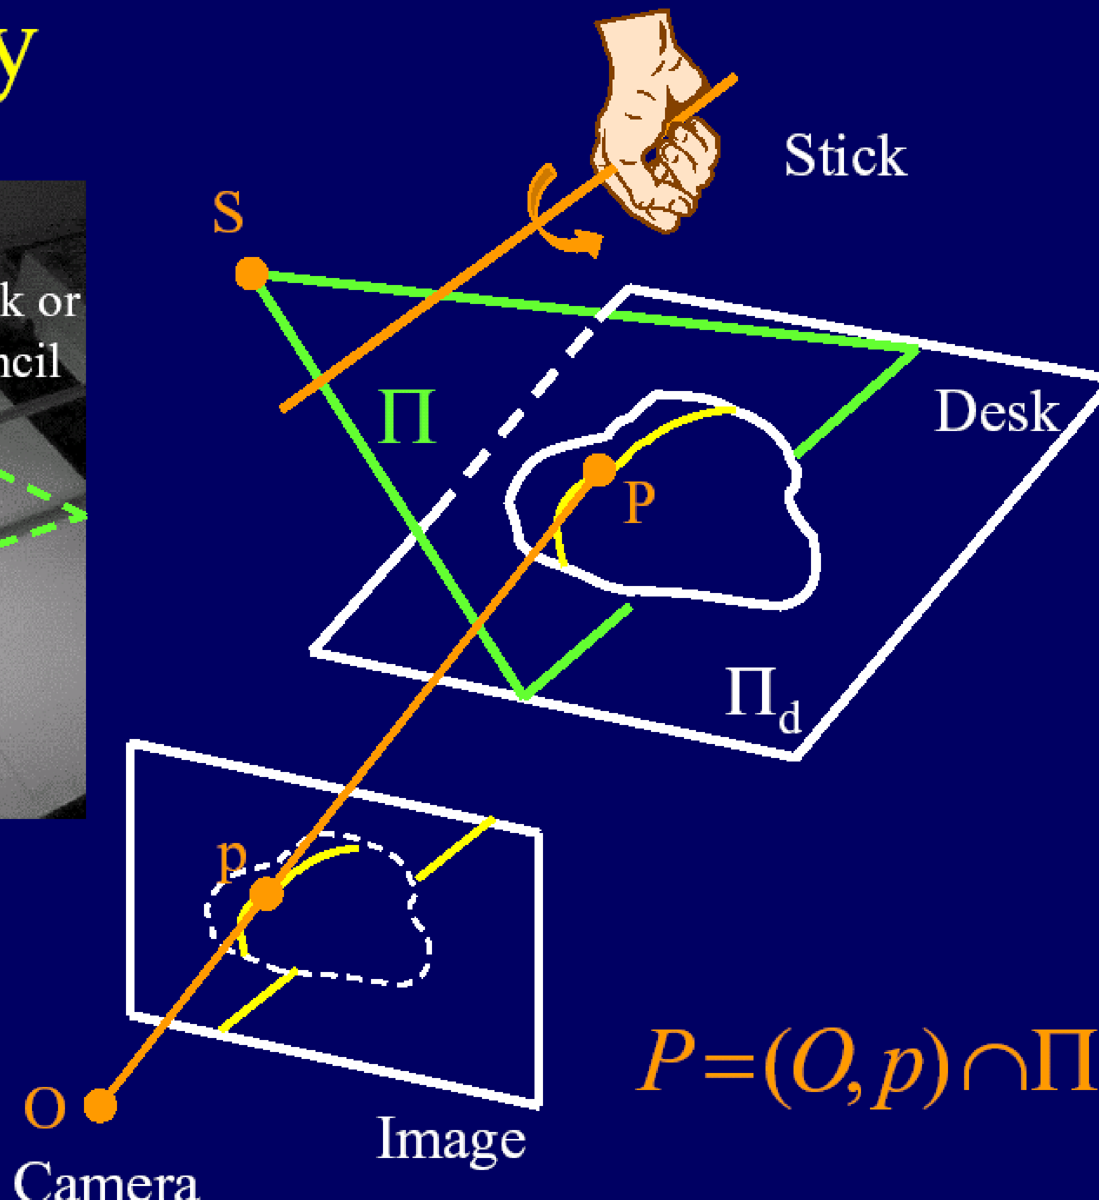

$$P = (O, p) \cap \Pi$$

Lamp Calibration

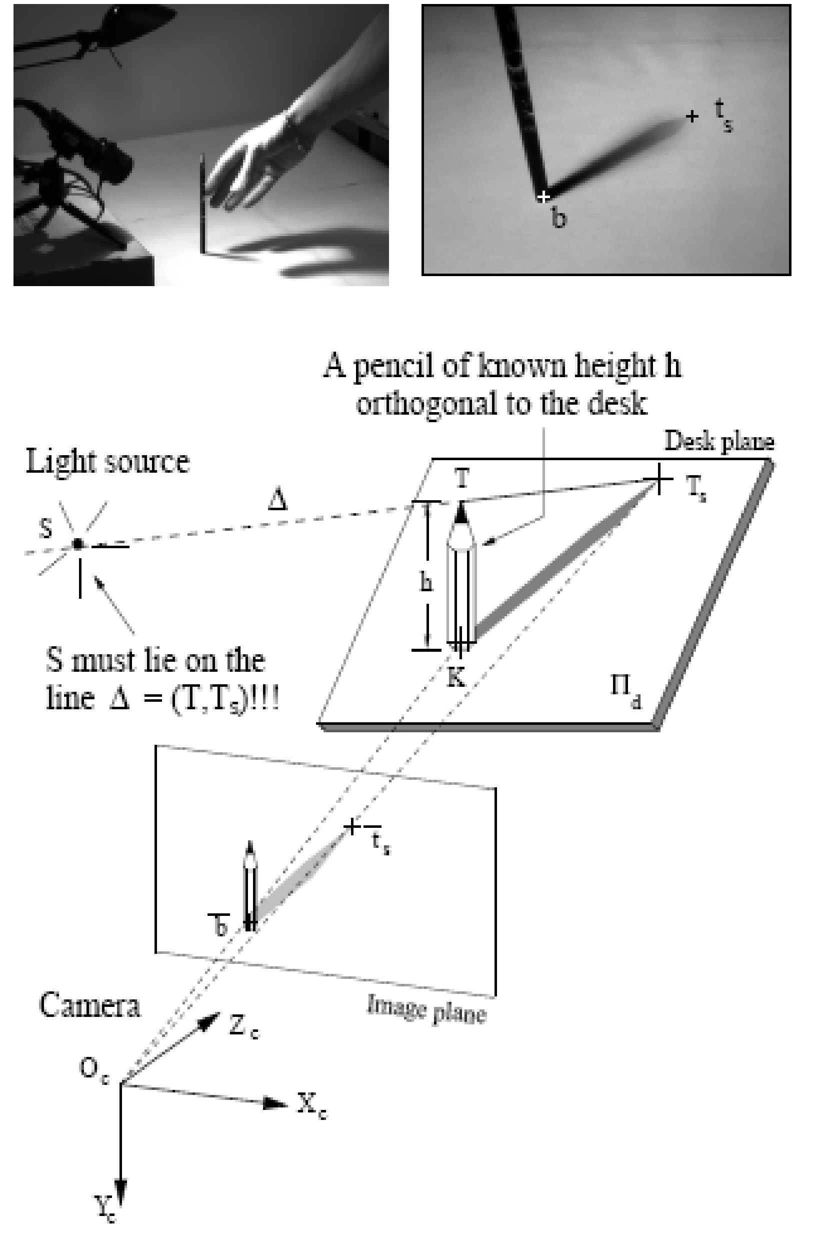

Figure 3: **Lamp calibration:** The operator places a pencil on the desk plane Π_d, orthogonal to it (top-left). The camera observes the shadow of the pencil projected on the tabletop.

The geometry

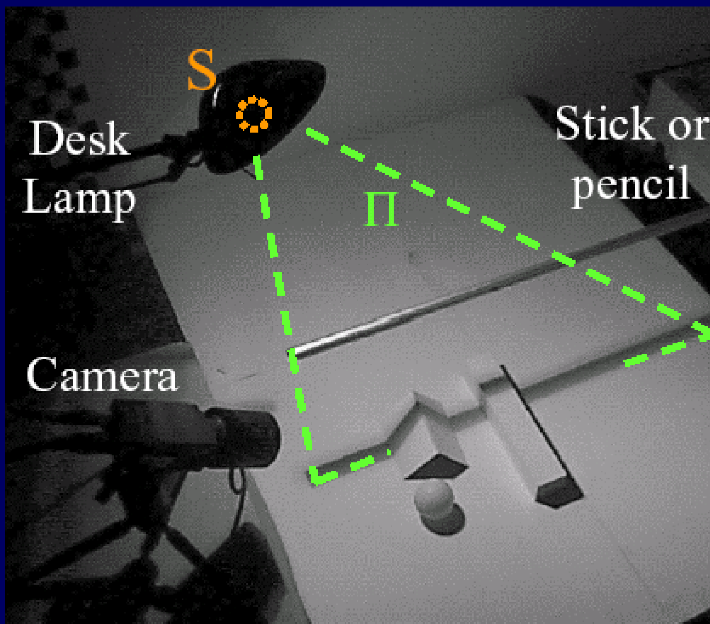

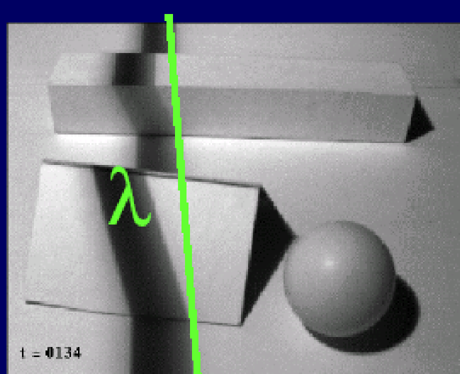

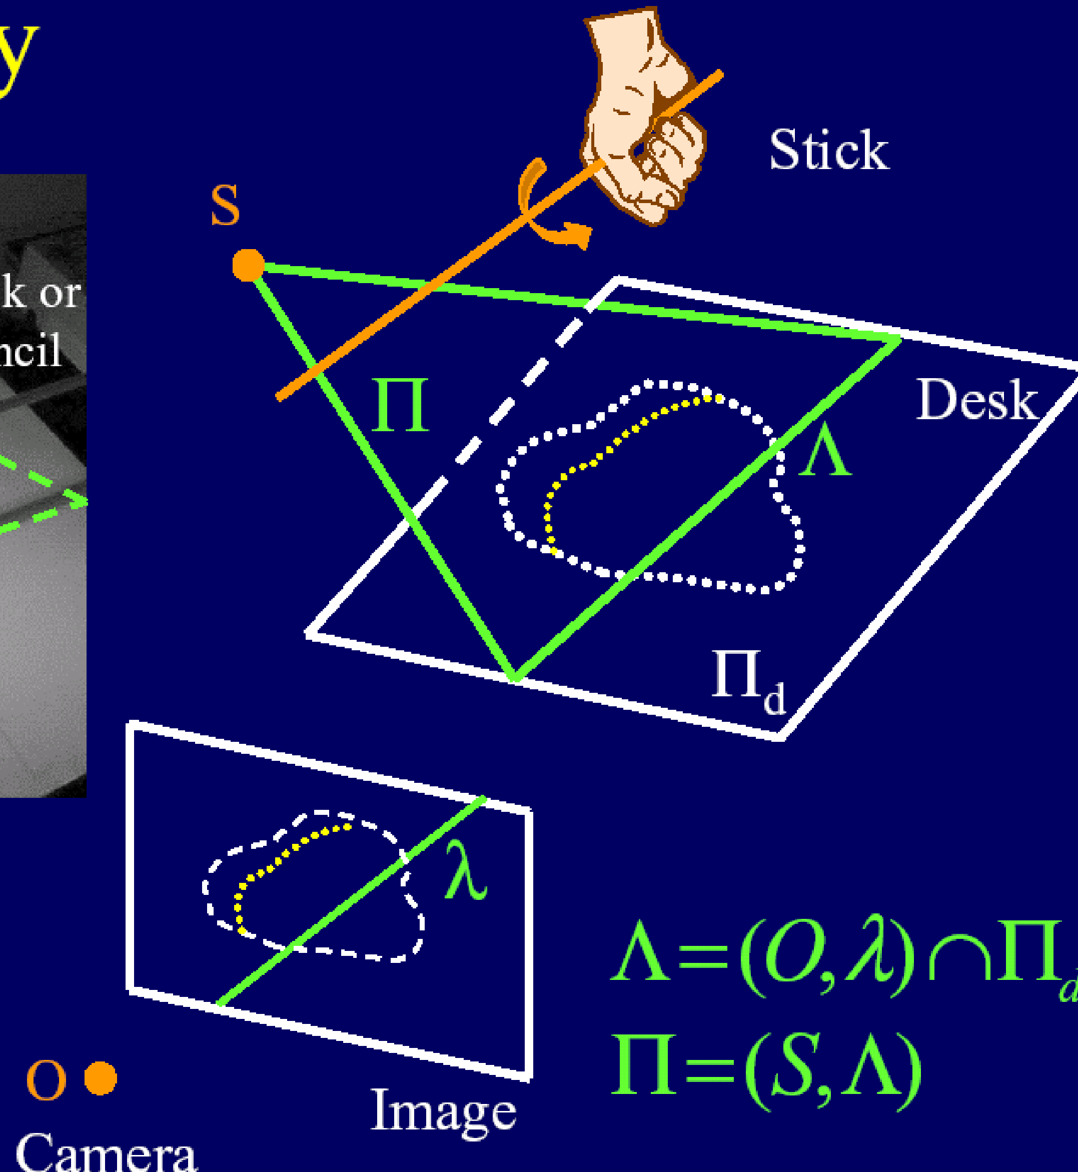

$$\Lambda = (O, \lambda) \cap \Pi_d$$

$$\Pi = (S, \Lambda)$$

The geometry

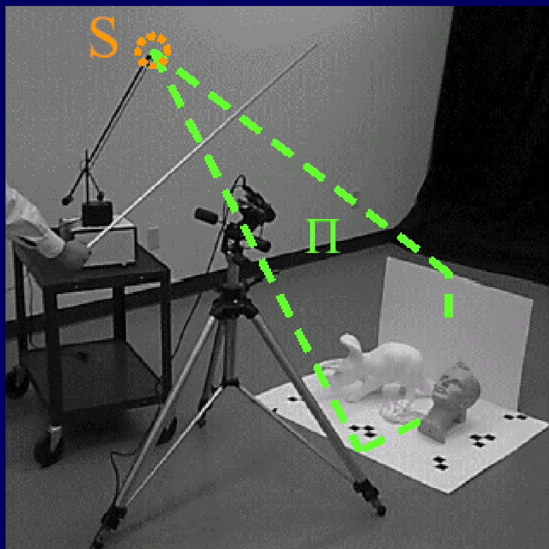

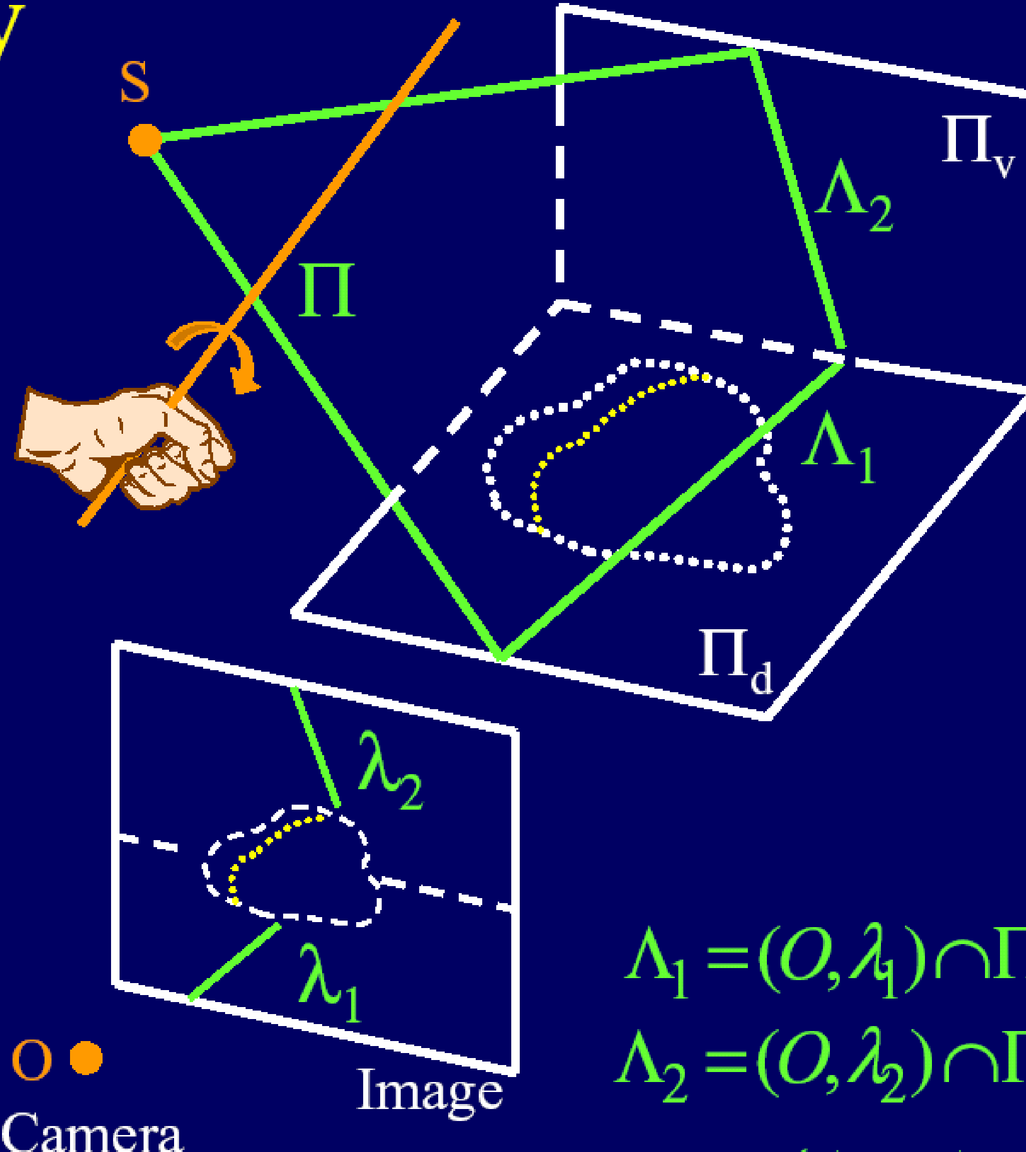

$$\Lambda_1 = (O, \lambda_1) \cap \Pi_d$$

$$\Lambda_2 = (O, \lambda_2) \cap \Pi_v$$

$$\Pi = (\Lambda_1, \Lambda_2)$$

Angel experiment

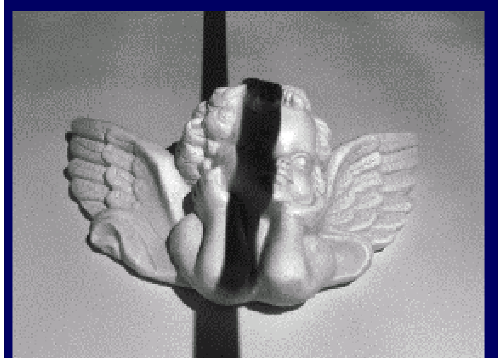

Accuracy: 0.1mm over 10cm → ~ 0.1% error

Scanning with the sun

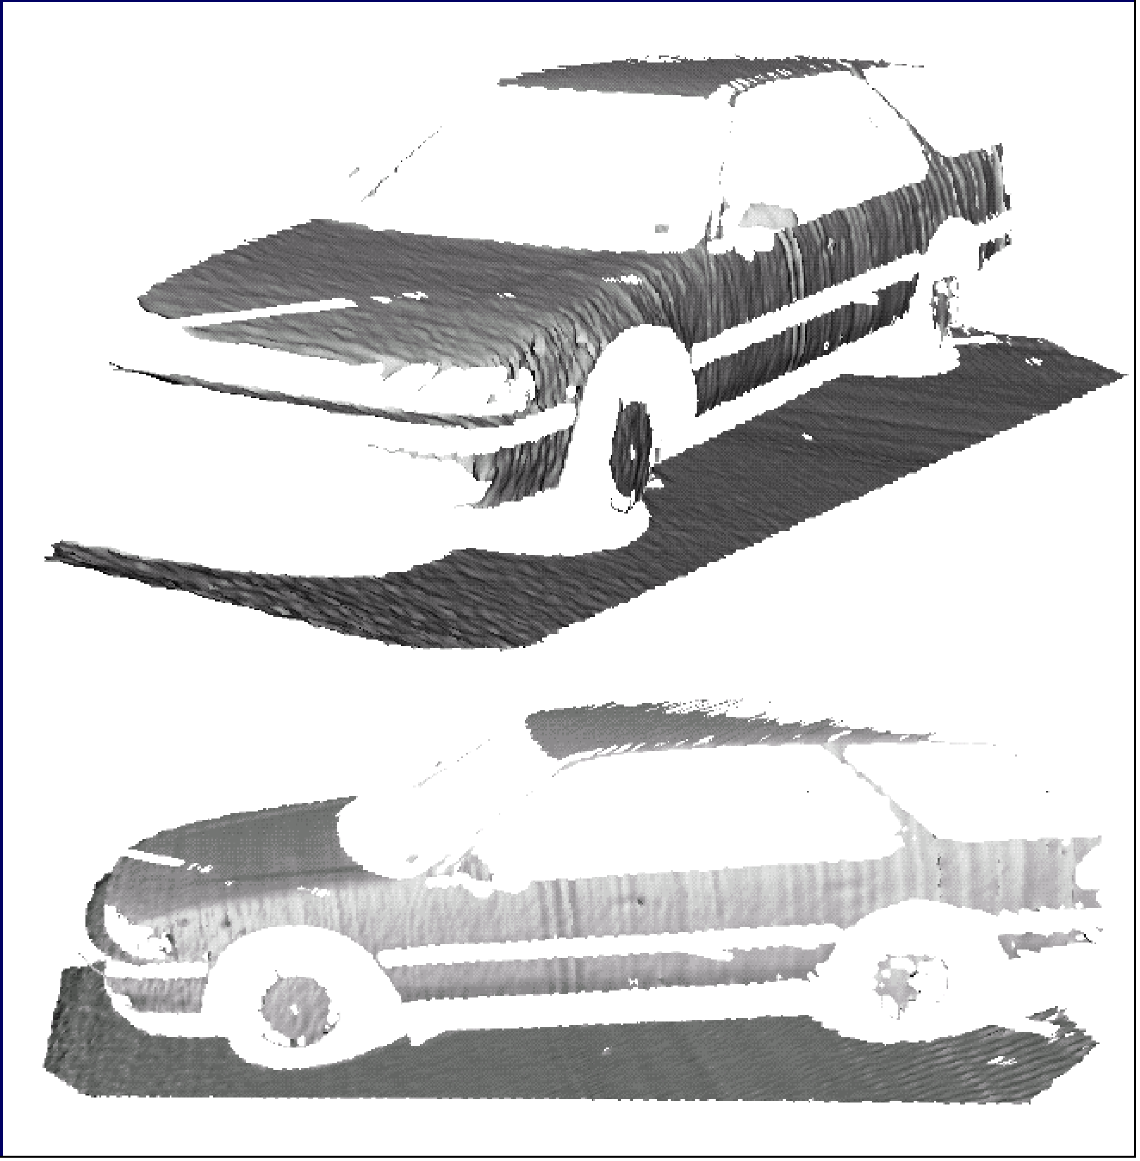

Accuracy: 1cm over 2m

→ ~ 0.5% error

Low-Cost 3D Scanner for Everyone

http://www.david-laserscanner.com/

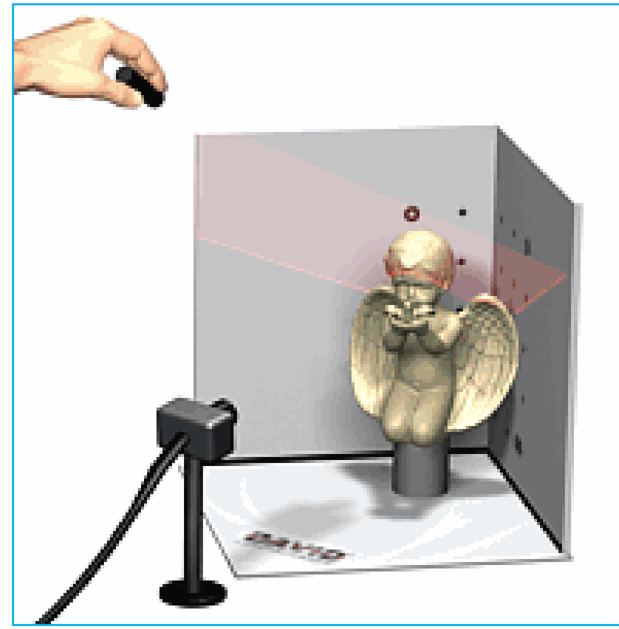

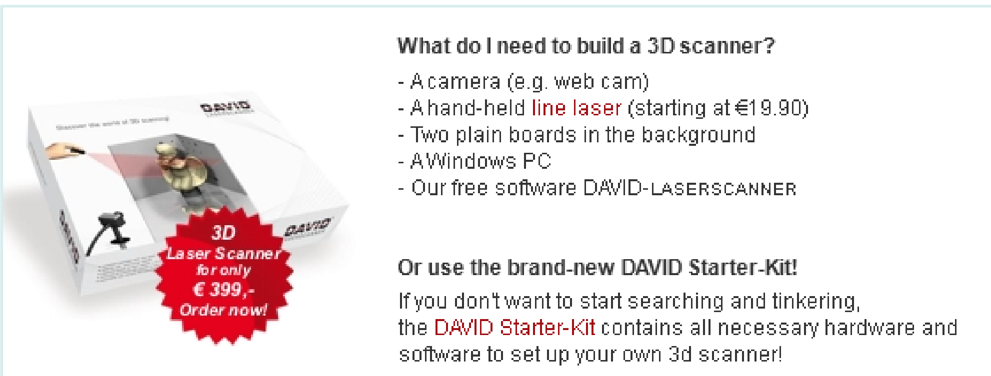

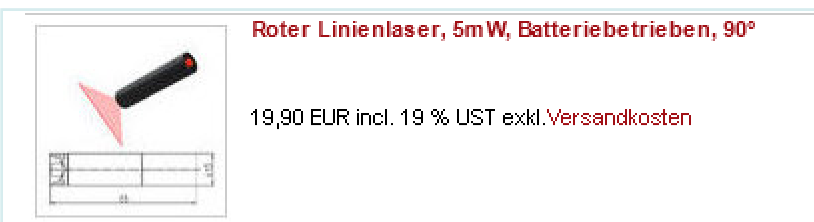

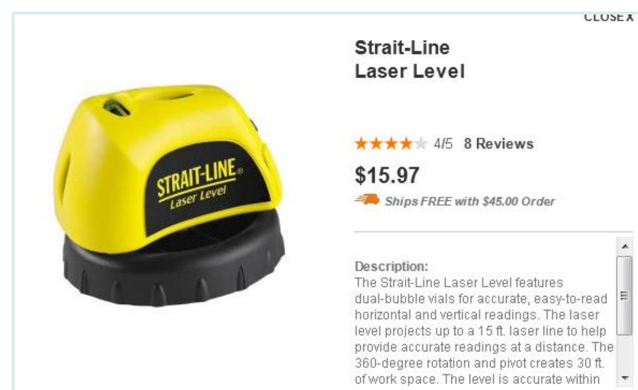

Low-Cost 3D Scanner for Everyone

http://www.david-laserscanner.com/wiki/user_manual/3d_laser_scanning

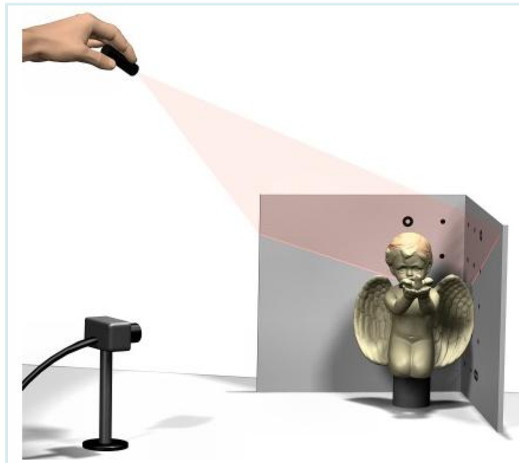

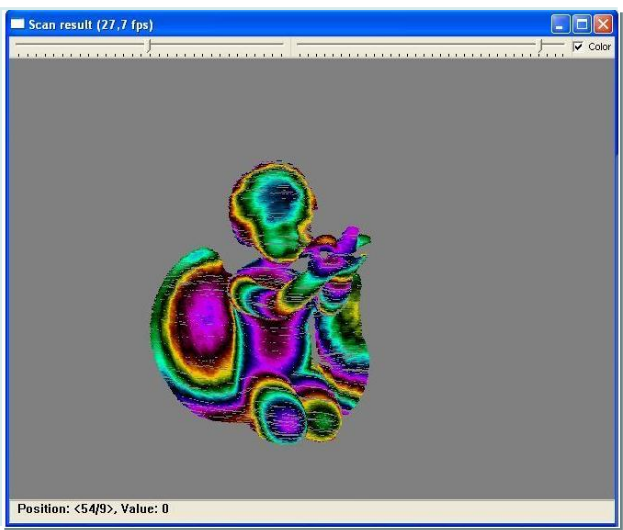

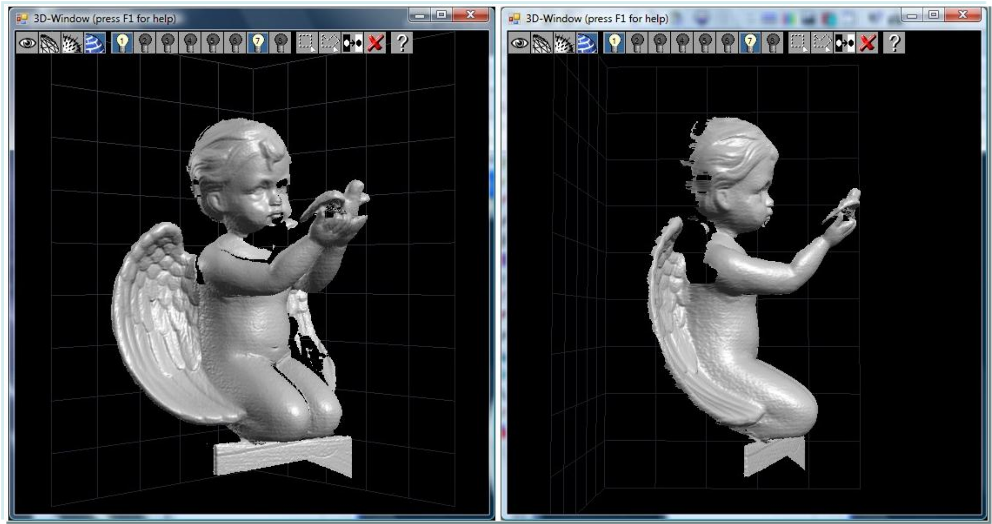

Image Processing Problem: Segmenting Stripes

- New Problem: How can we find the stripes in the images?
- Image thresholding is dependent on the contrast

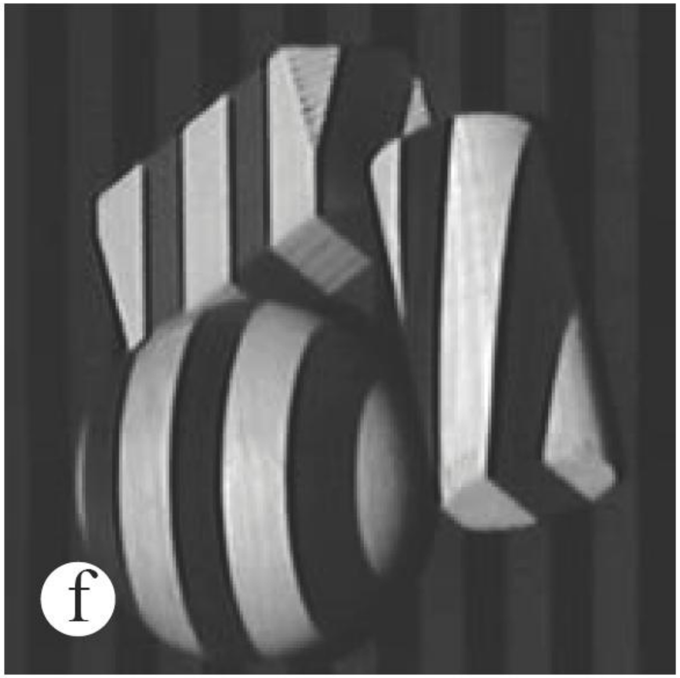

Image Processing Problem: How to detect stripes in images?

- Edge detection: Thresholding difficult
- Line detection: Lines of different width

- Solution: Project positive and negative strip pattern, detect intersections

Subpixel accuracy

1. **Zero crossings of 2nd derivative**
   - Gradient filter width must be chosen
   - Depends on stripe width
   - Problem: Width changes with orientation of surface

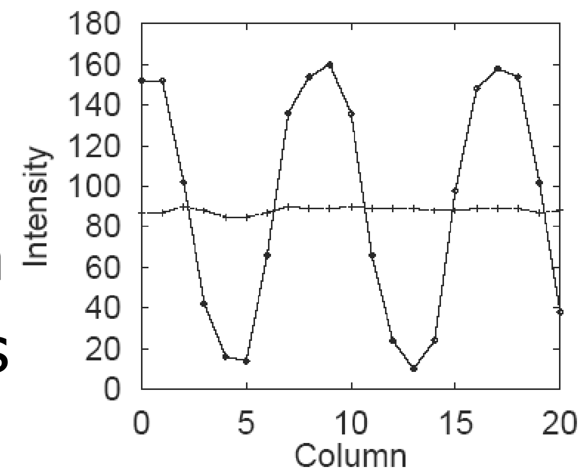

Subpixel accuracy

2. Linear interpolation

- With fully lit and completely dark images determine dynamic threshold T
- P determined by intersecting threshold and image profile
- Robust against changes in contrast

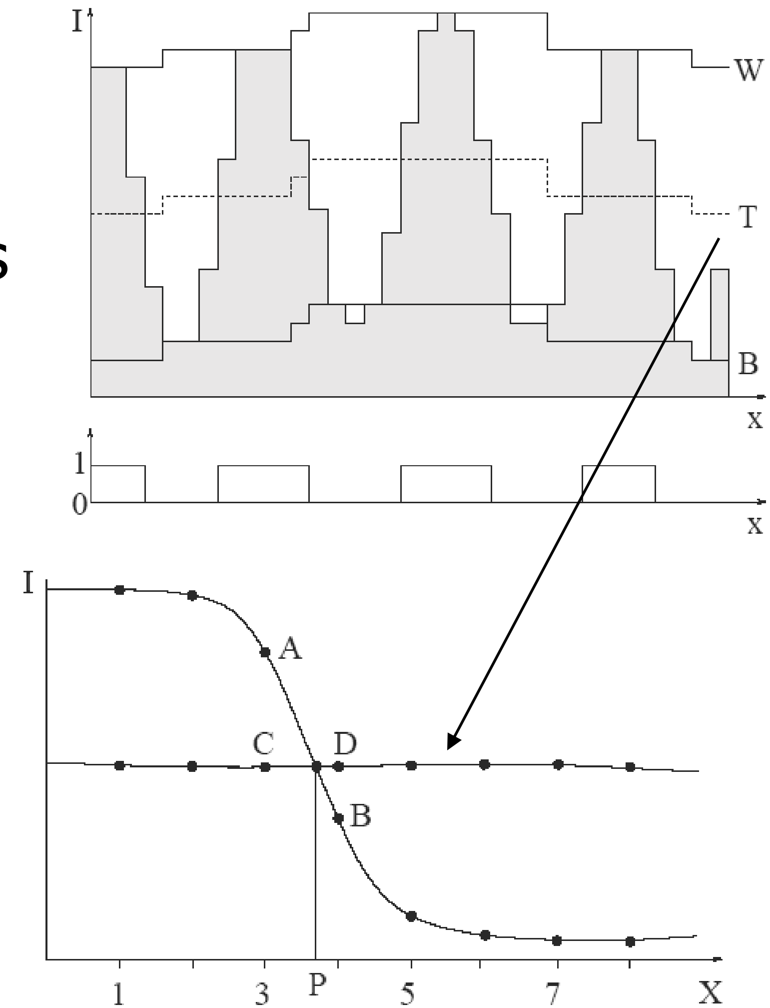

Subpixel accuracy

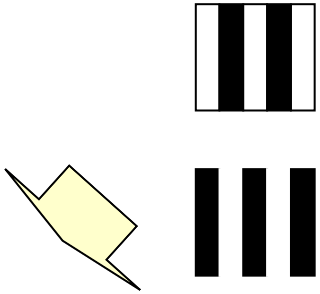

3. Inverse stripe pattern intersection

- Also robust against slightly different width of black and white stripes
- No bias from isolating gap between adjacent stripes in LCD array

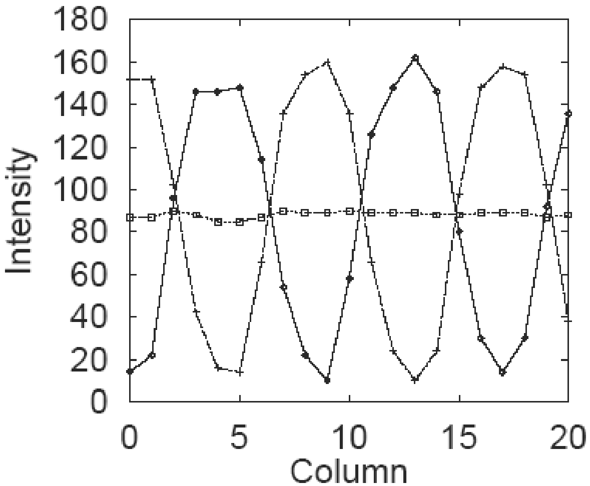

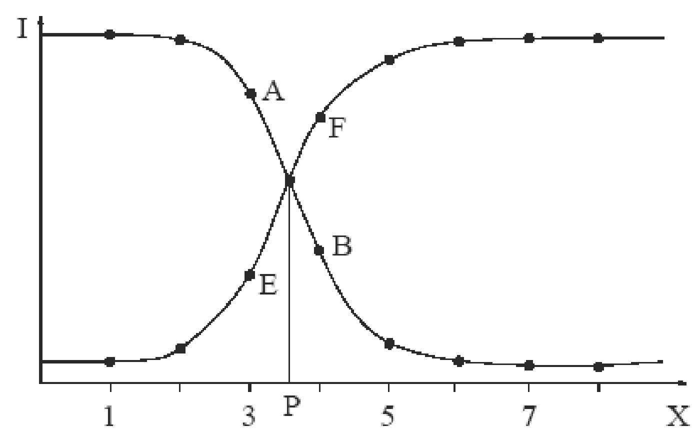

Image Processing Problem: How to detect stripes in images?

- Edge detection
- Line detection
- Solution: Project positive and negative strip pattern, detect intersections

- But: set of lines, uniqueness?, which part of the line corresponds to which light plane?

Next Lecture: Encoded Patterns

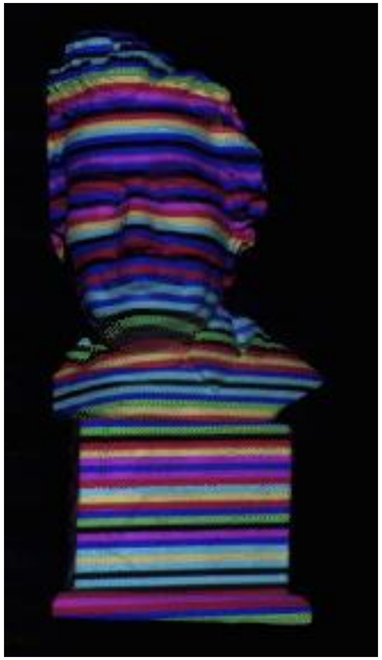

- Any spatio-temporal pattern of light projected on a surface (or volume).
- Cleverly illuminate the scene to extract scene properties (eg., 3D).
- Avoids problems of 3D estimation in scenes with complex texture/BRDFs.
- Very popular in vision and successful in industrial applications (parts assembly, inspection, etc).