3D acquisition

è

Acknowledgement

Courtesy of Prof. Luc Van Gool

3D acquisition taxonomy

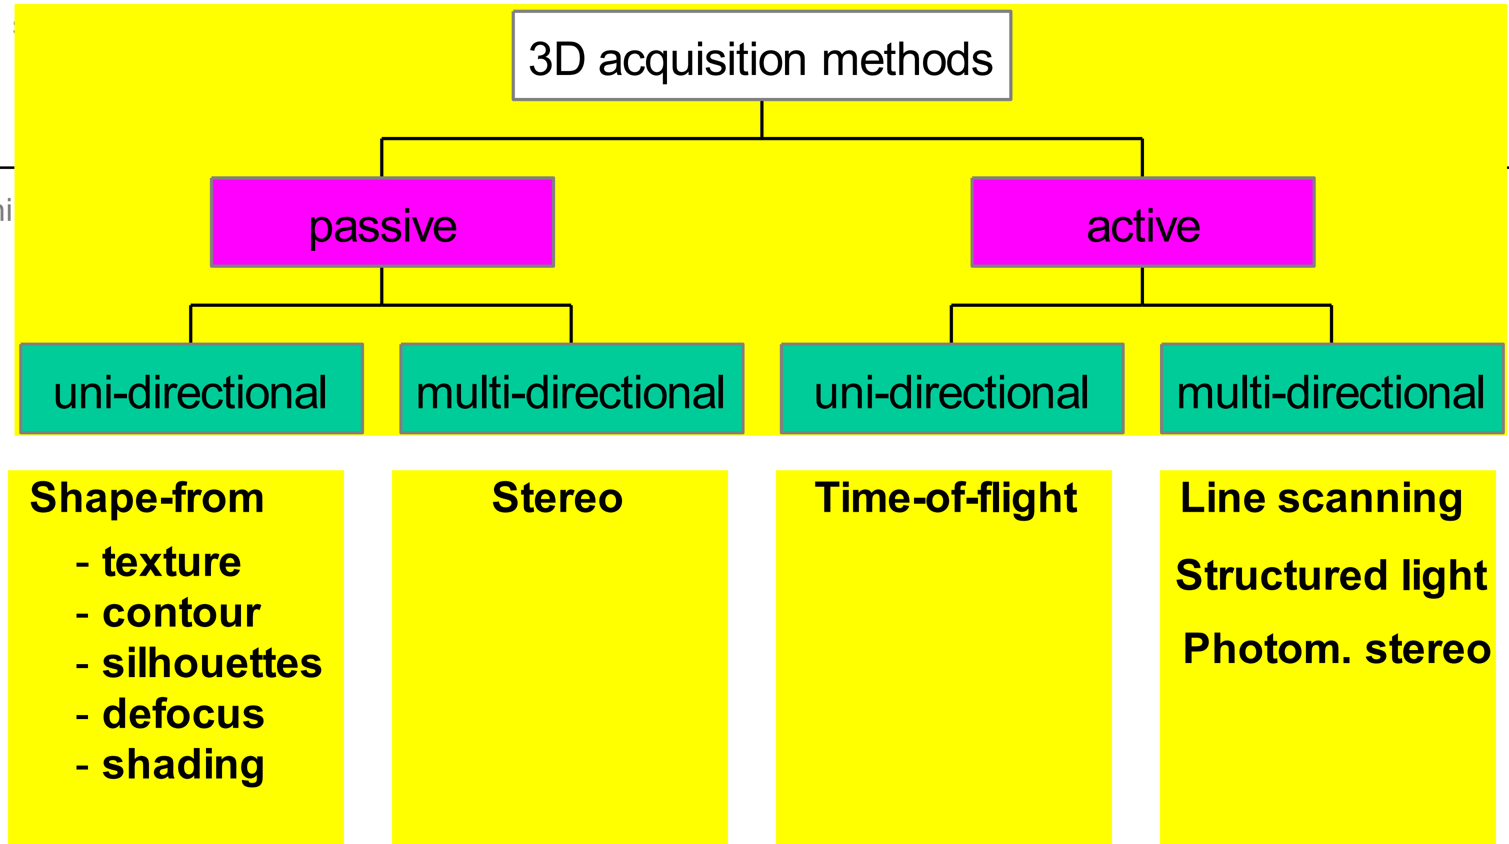

3D acquisition taxonomy

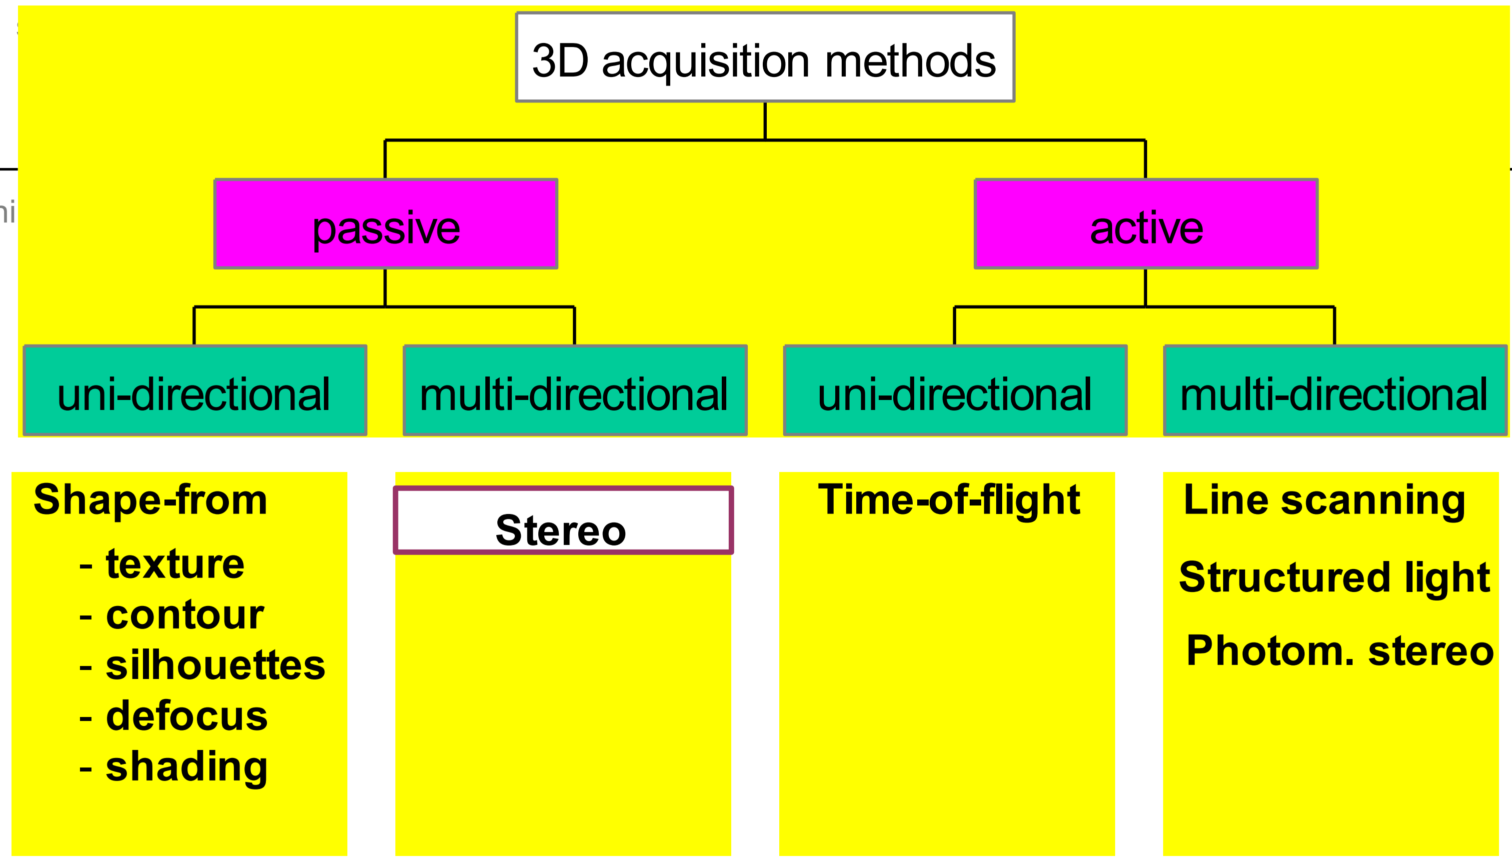

Stereo

The underlying principle is "triangulation" :

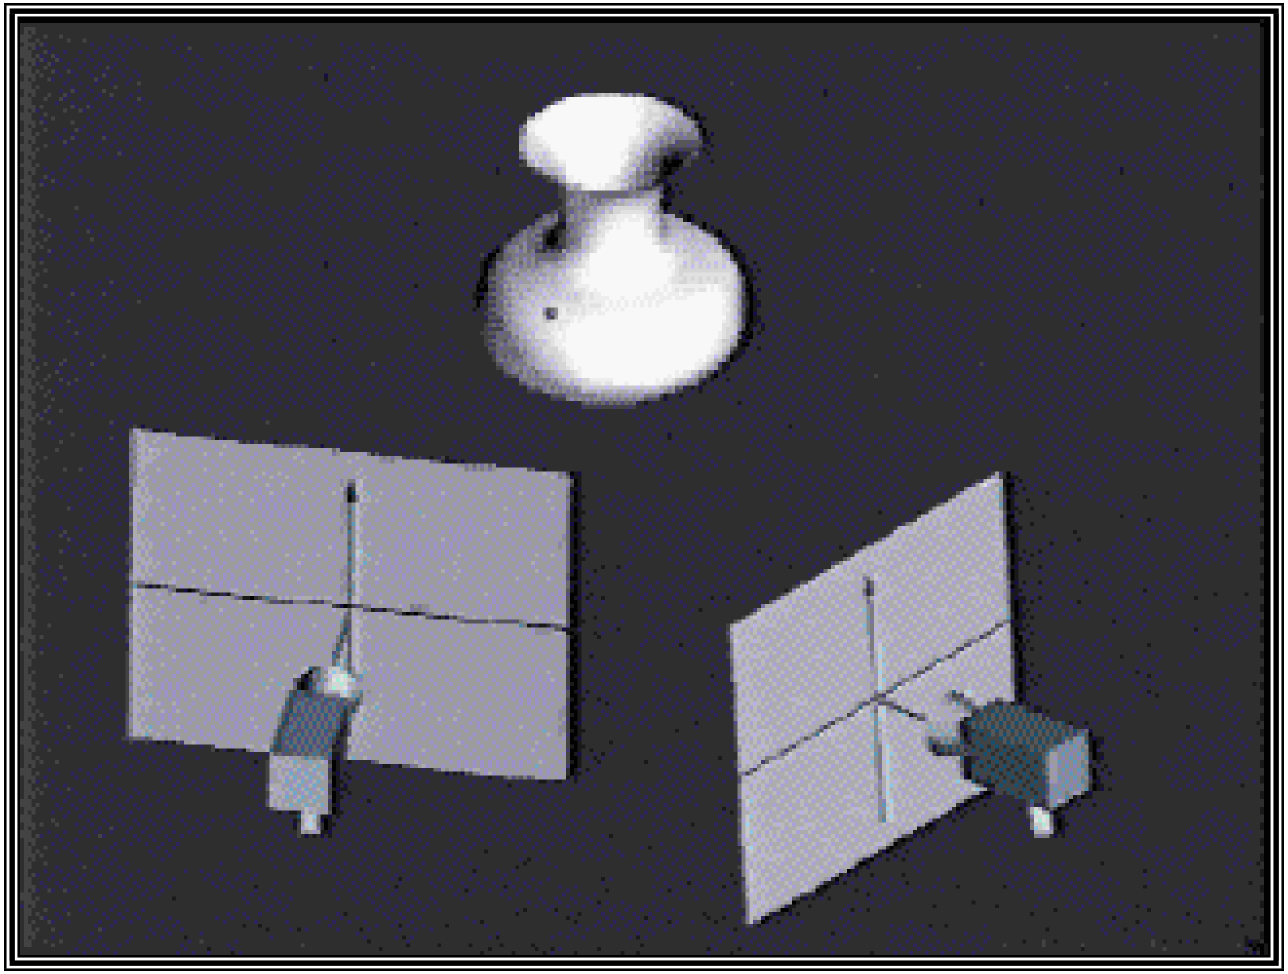

è

(Passive) stereo

Simple configuration :

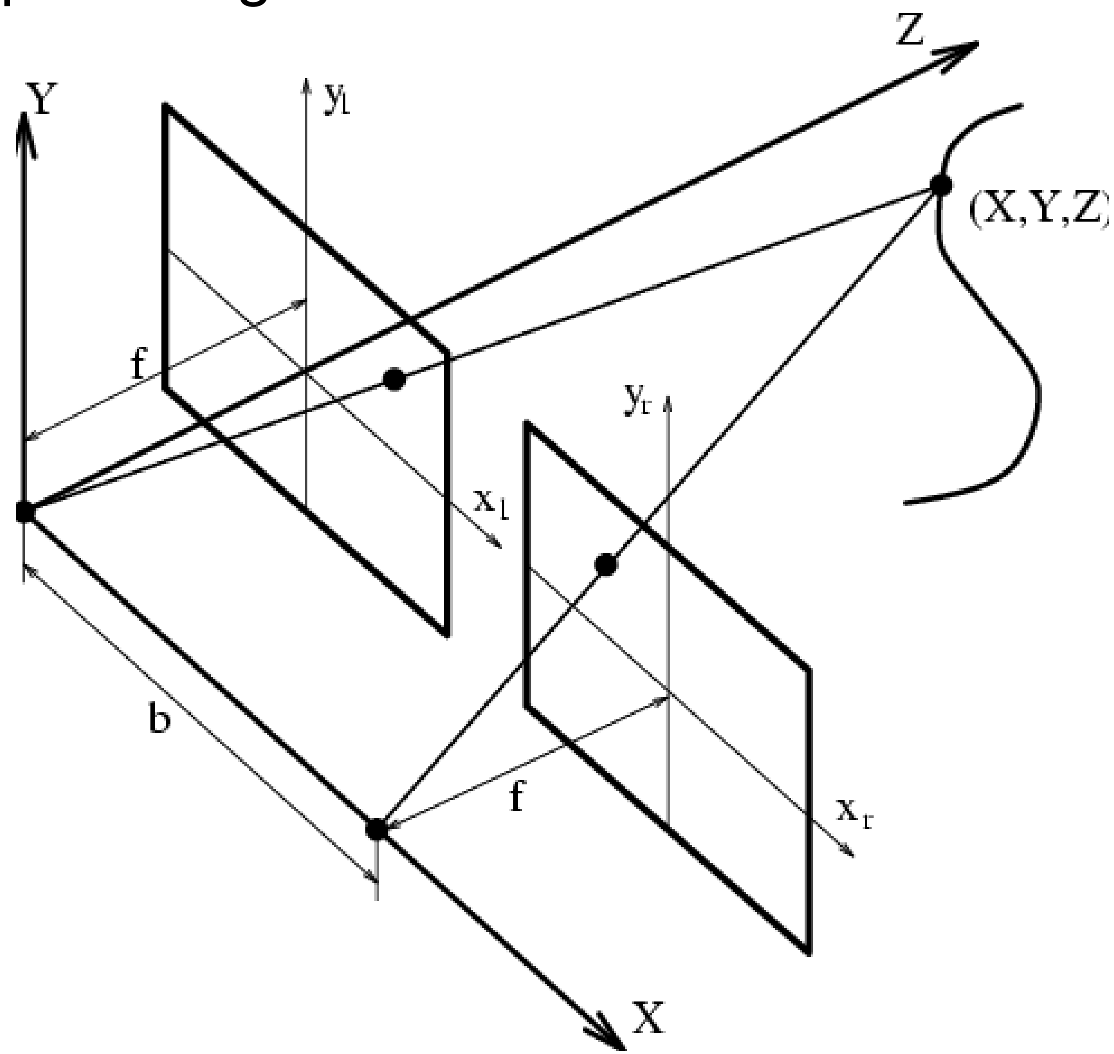

è

A simple stereo setup

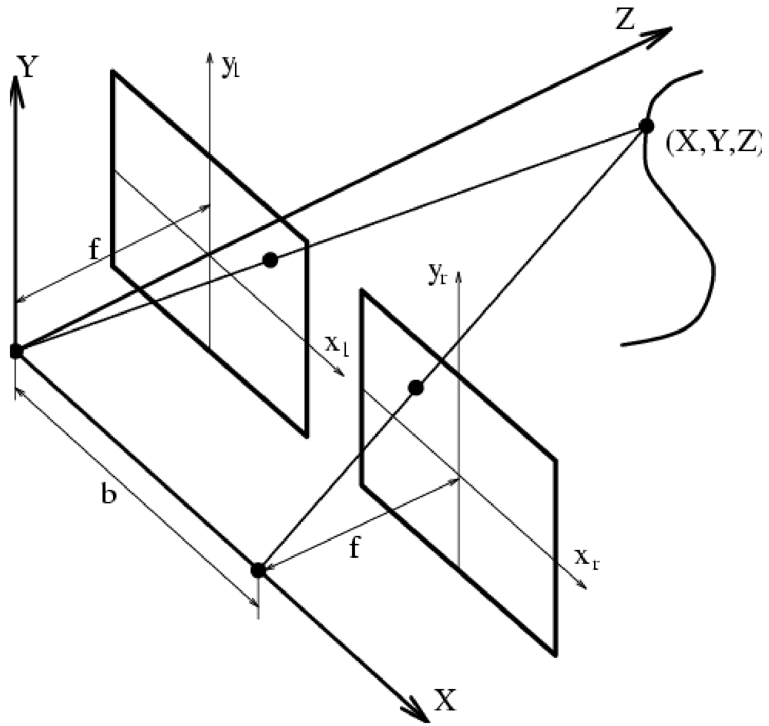

- identical cameras
- coplanar image planes
- aligned *x*-axes

è

A simple stereo setup

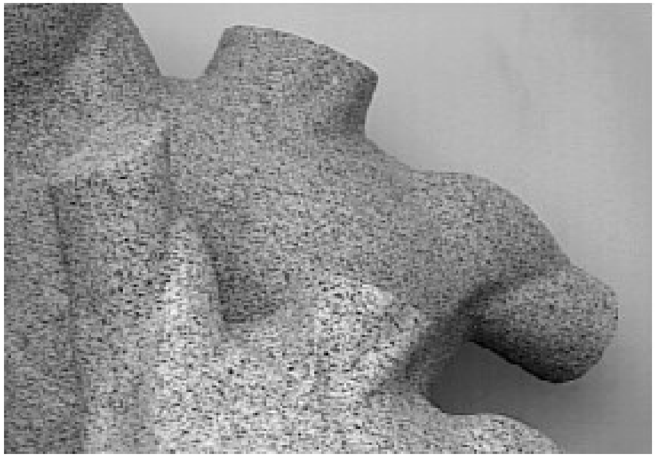

The HARD problem is finding the *correspondences*

Notice : no reconstruction for the untextured back wall...

è

Correspondence problem : methods

1. correlation
   - deformations…
   - small window ⇒ noise!
   - large window ⇒ bad localisation
2. feature-based
   - mainly edges and corners
   - sparse depth image
3. regularisation methods

è

3D city models

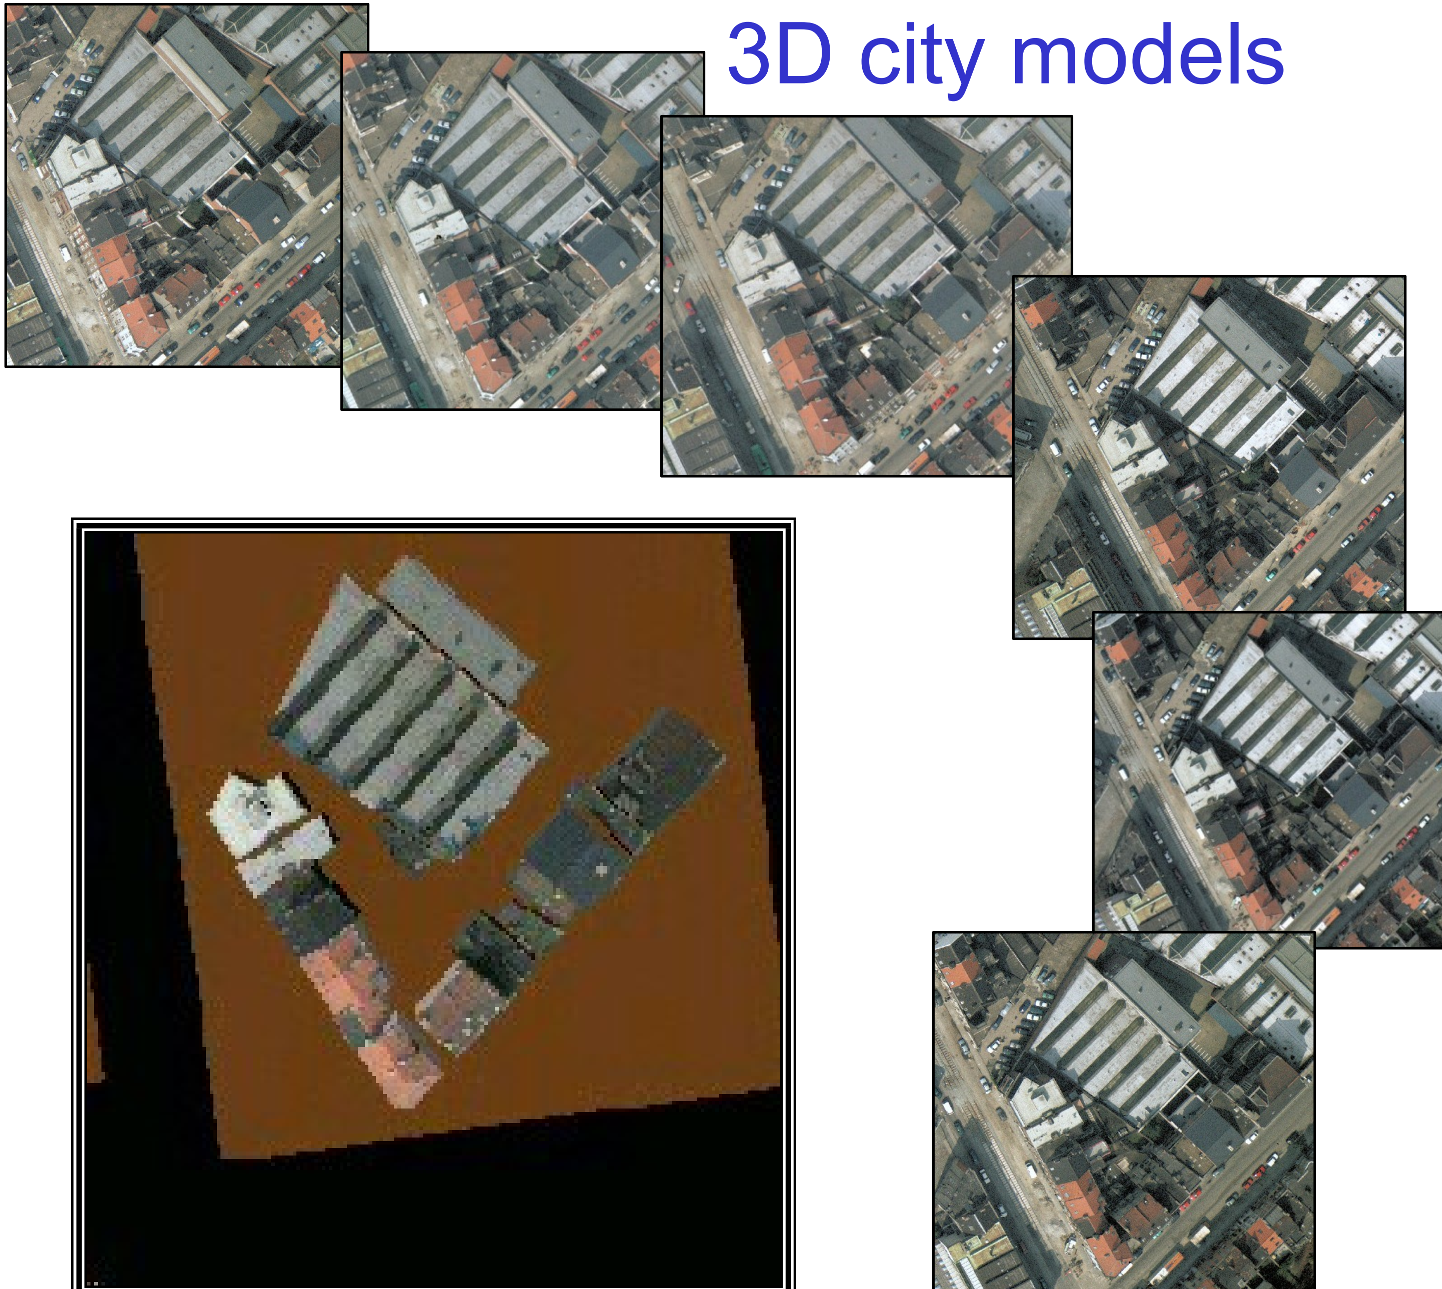

3D city models – ground level

Mobile mapping example – for measuring

Uncalibrated reconstruction

Tracking and Calibration → **Dense depth estimation** → **3D surface modeling**

Points and cameras

Depth map

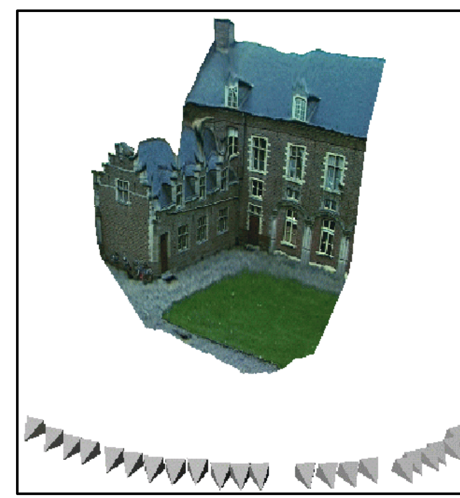

3D models

Uncalibrated reconstruction

Uncalibrated reconstruction - example

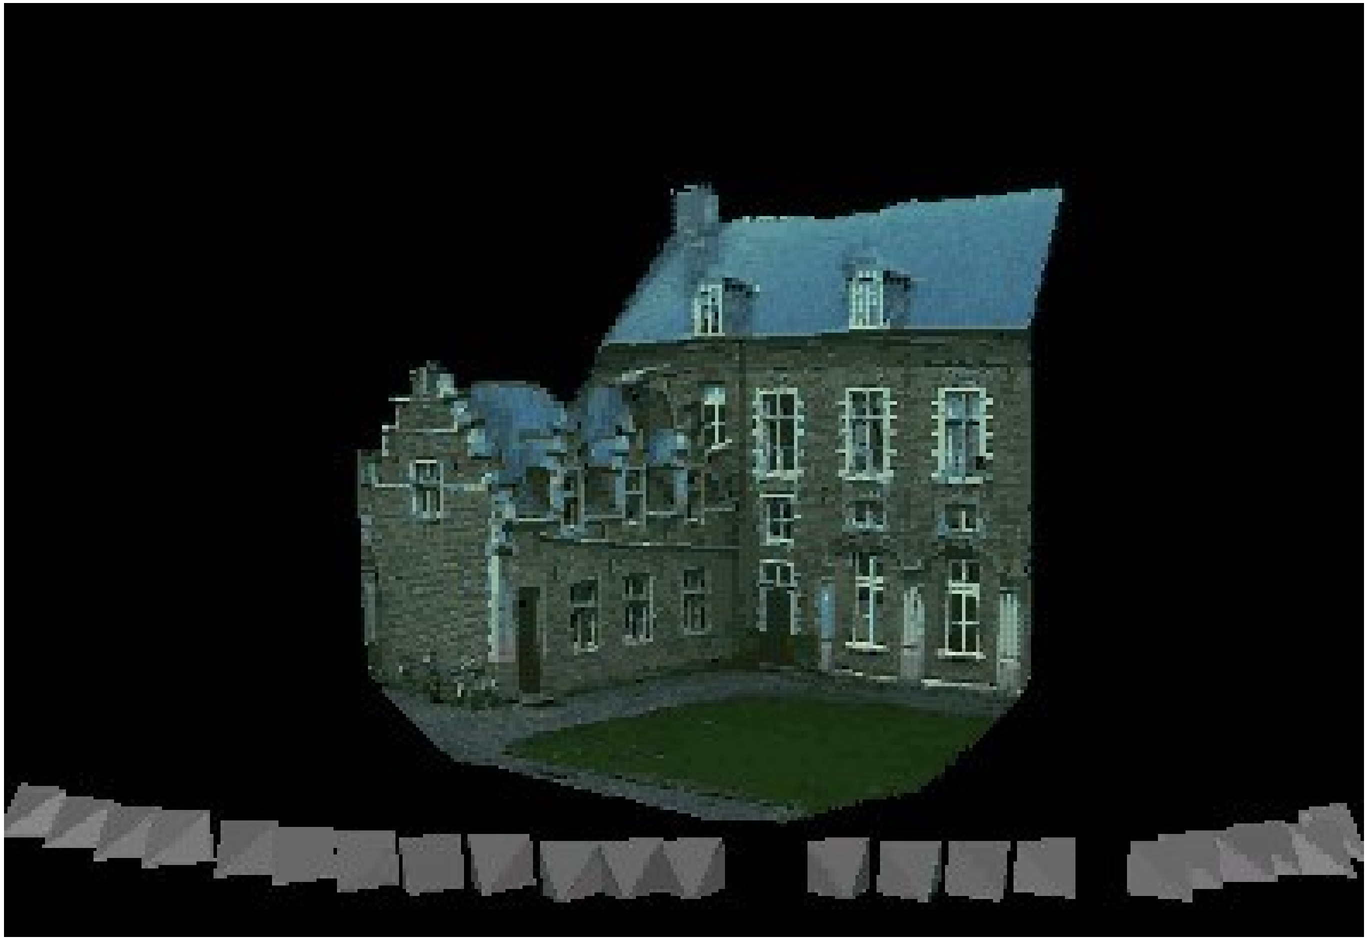

Univ. of Leuven

è

Shape-from-stills

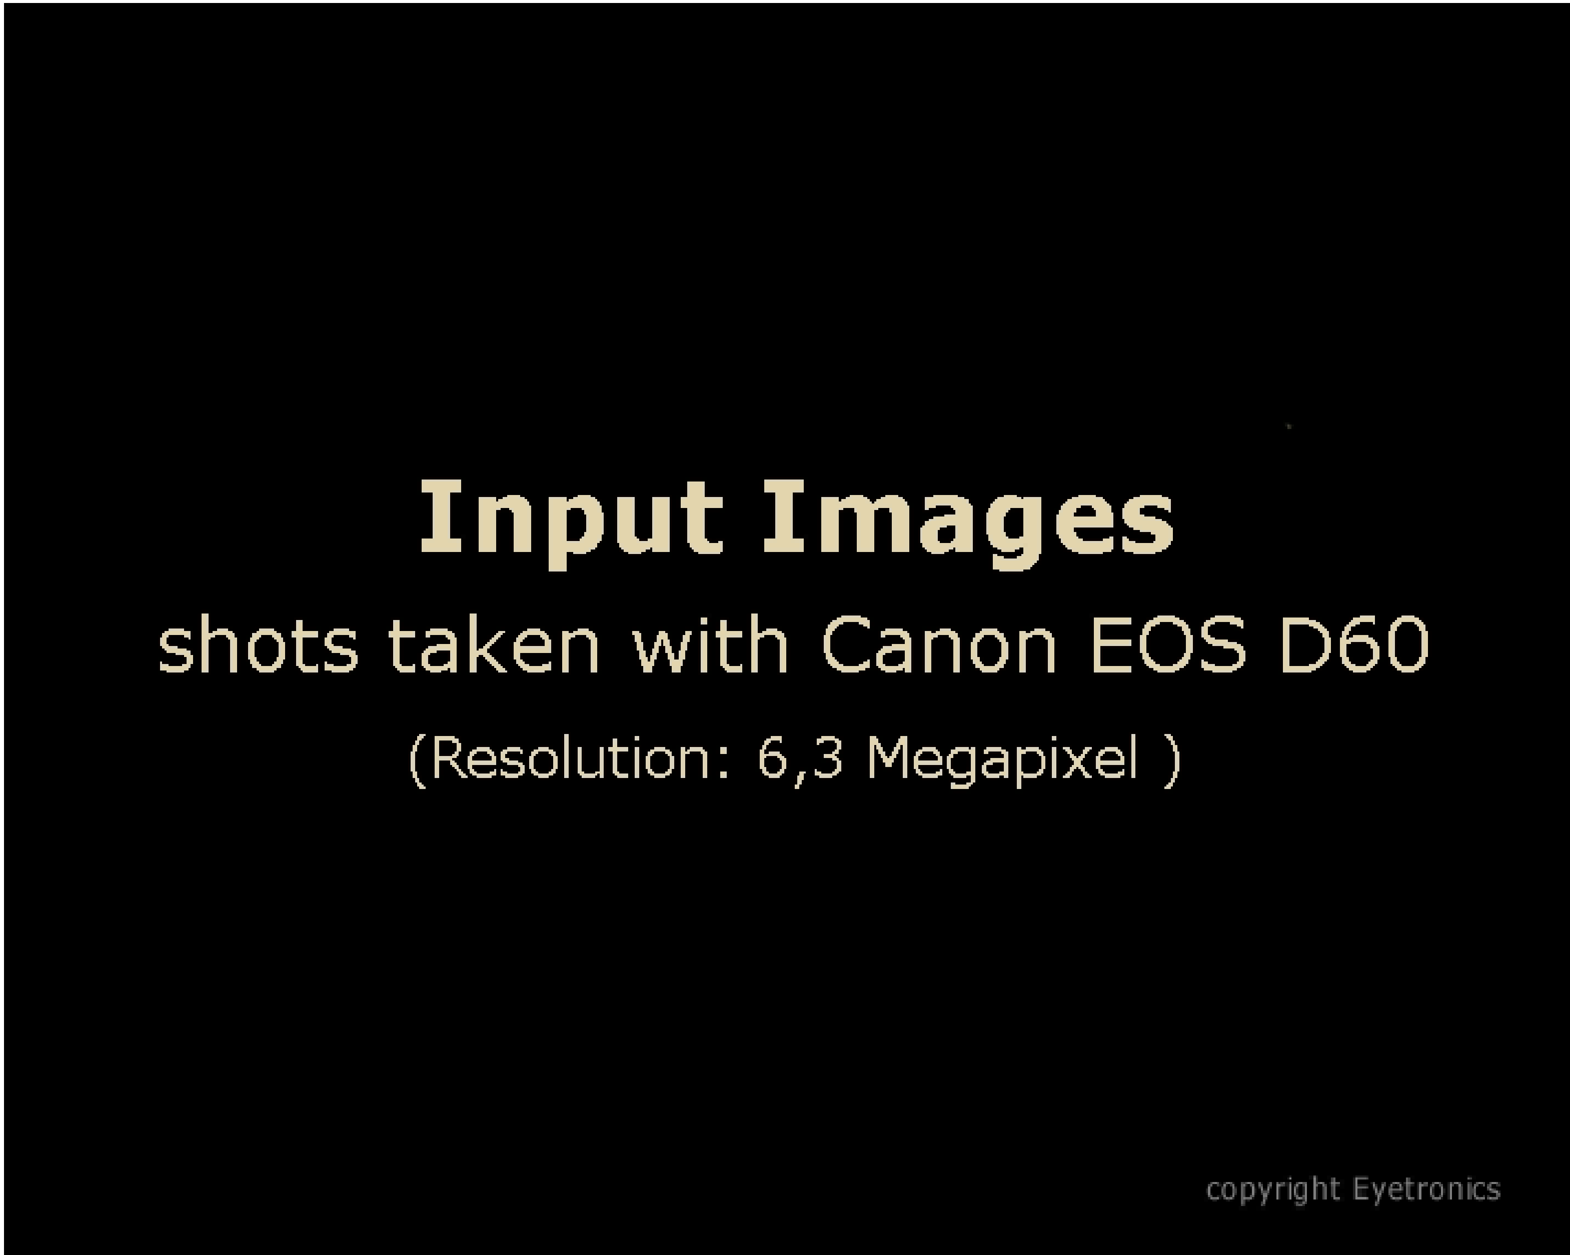

Shape-from-stills

www.arc3d.be

Webservice,
free for non-commercial use

3D acquisition taxonomy

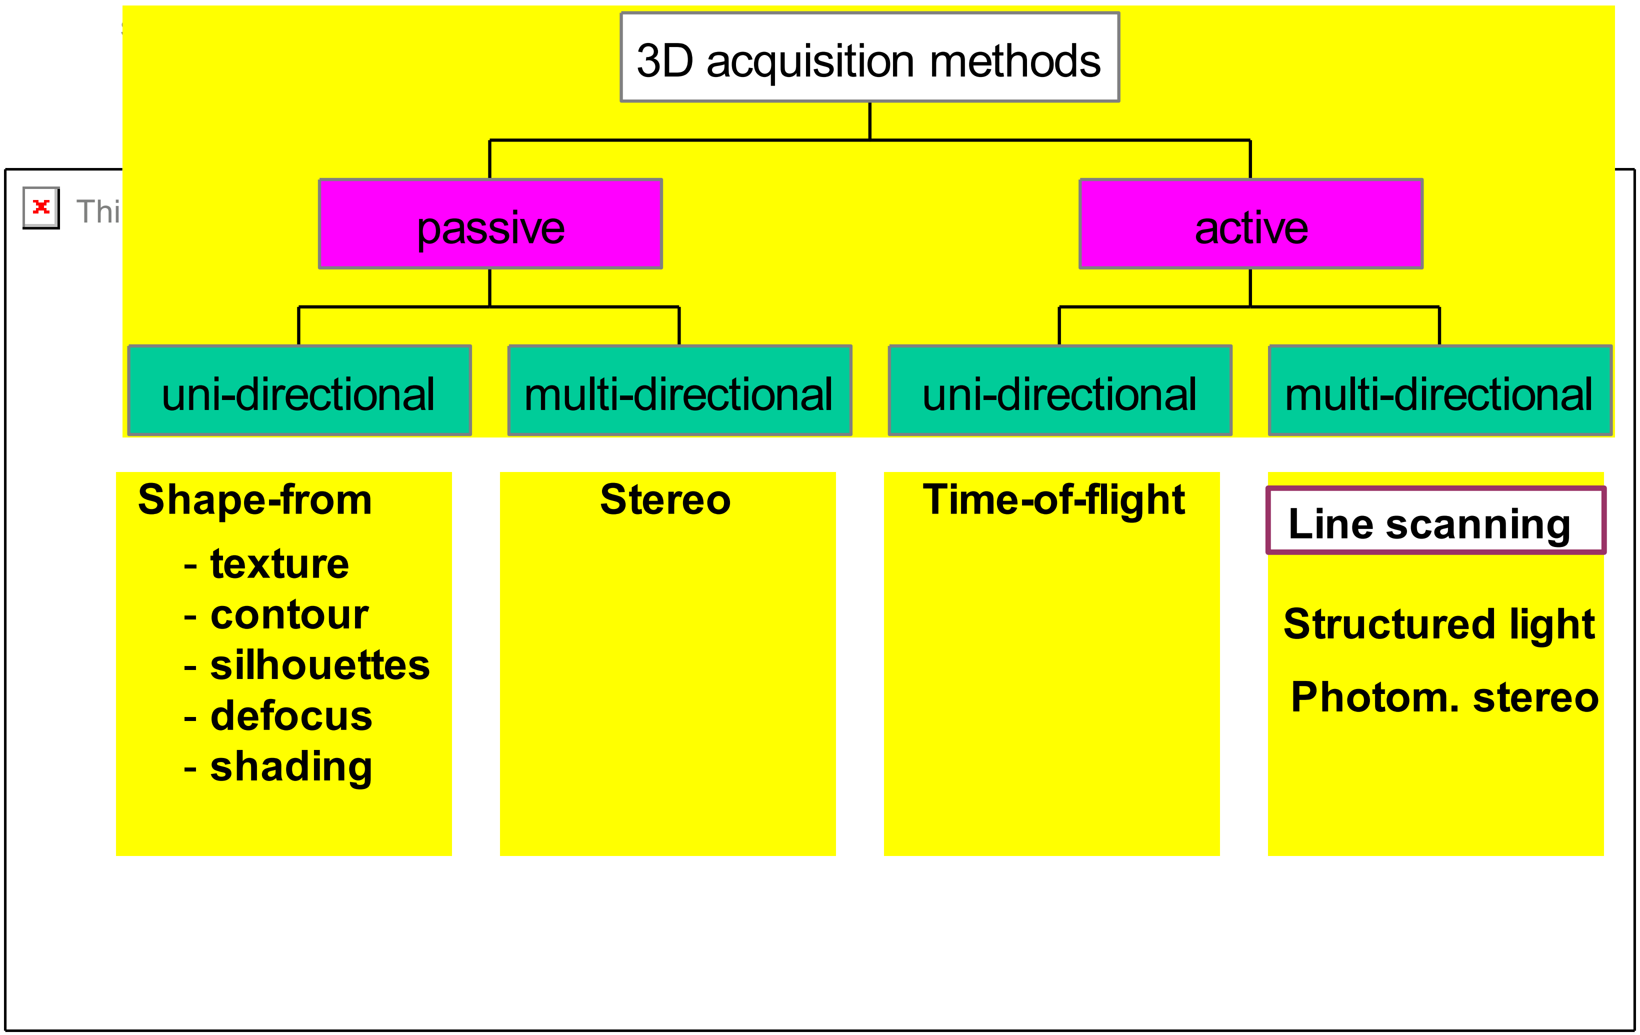

Active triangulation

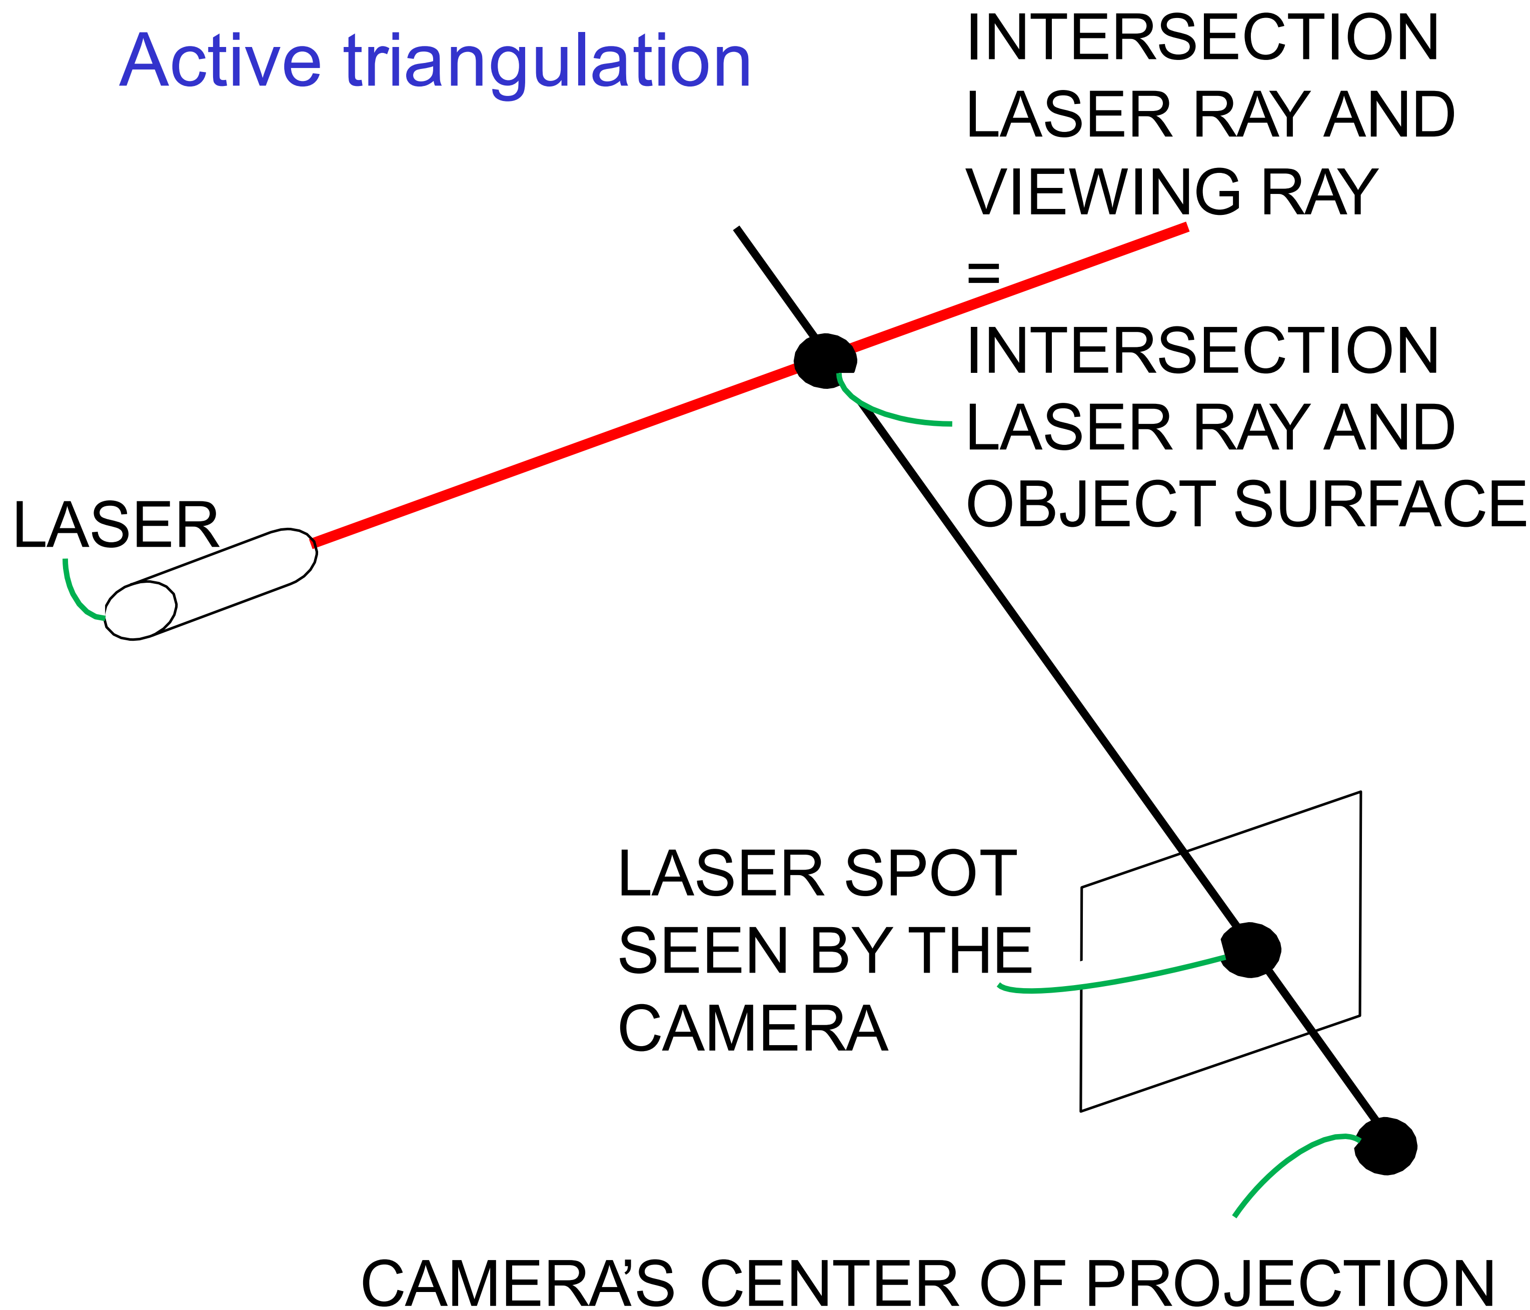

Active triangulation

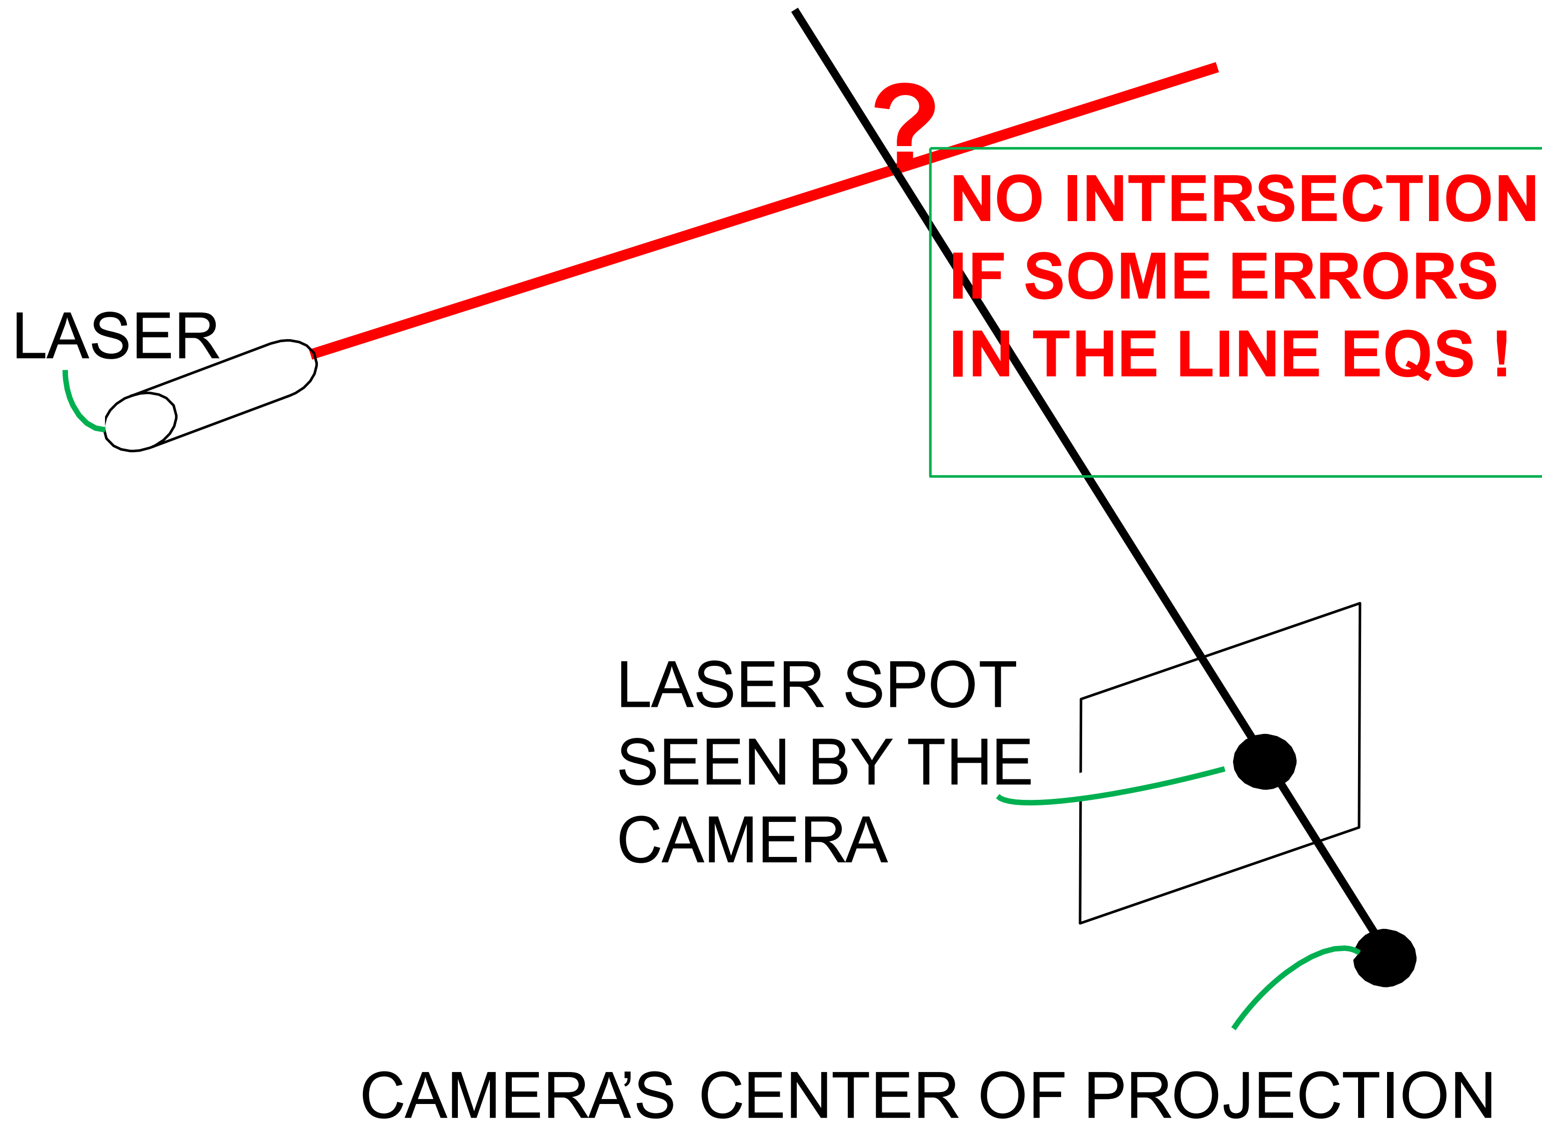

Active triangulation

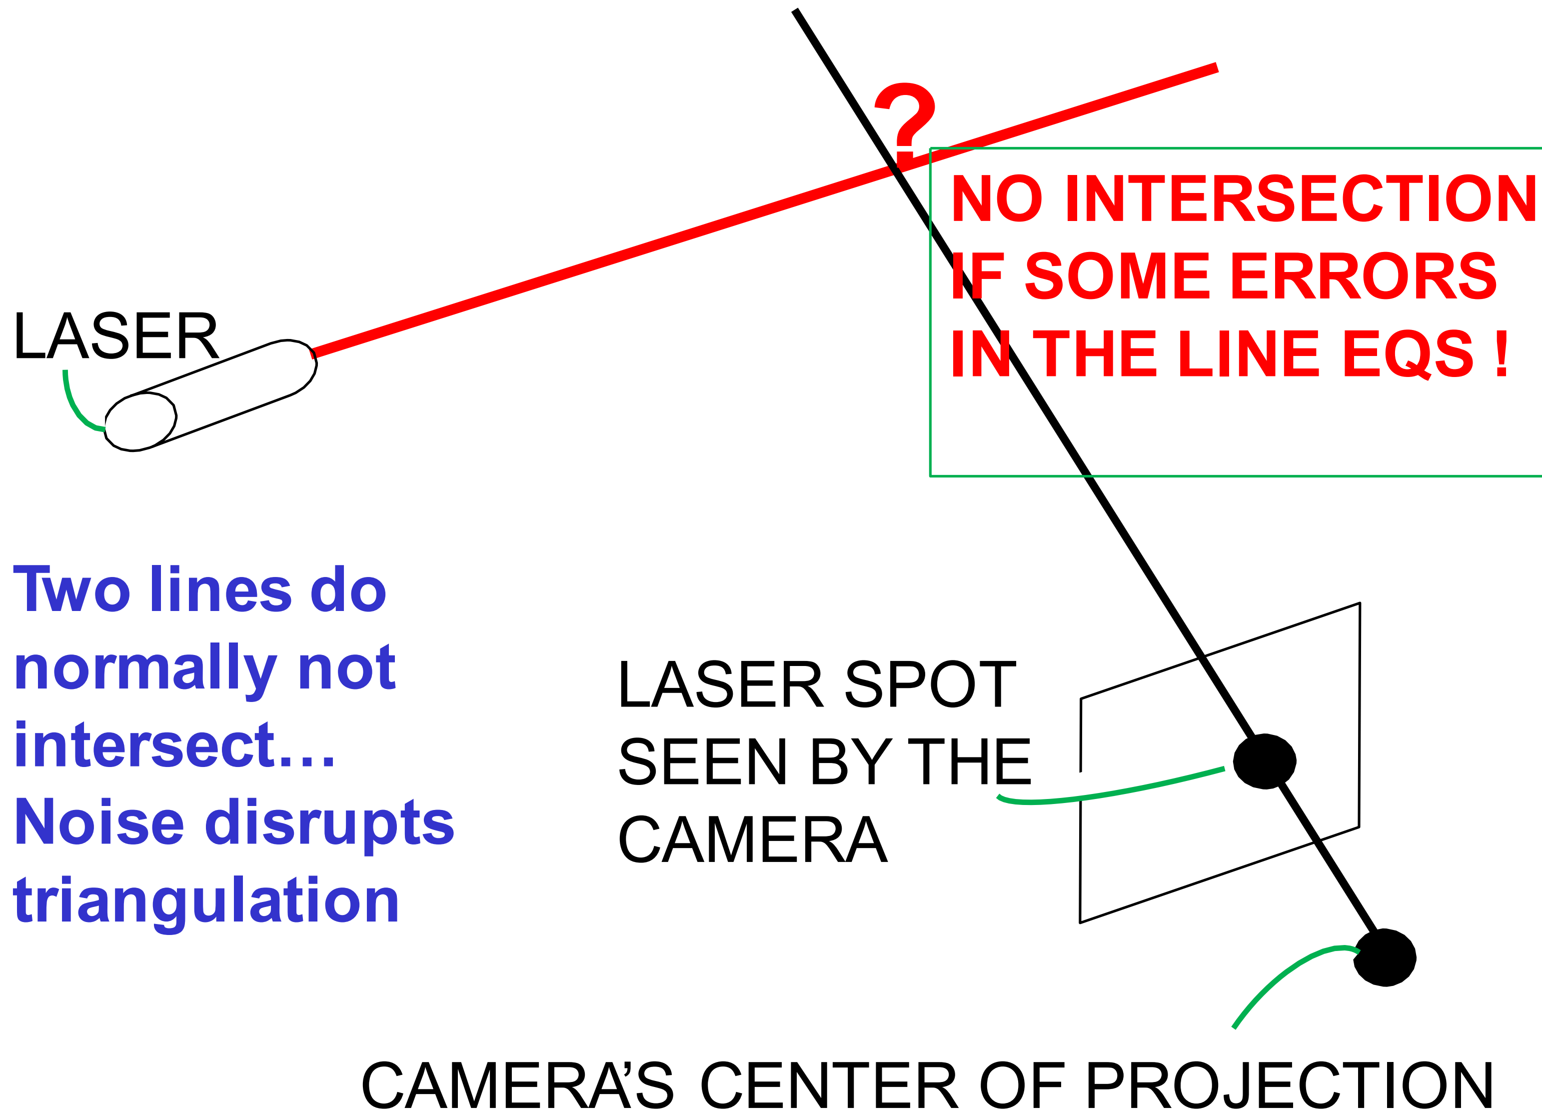

Two lines do normally not intersect… Noise disrupts triangulation

Active triangulation

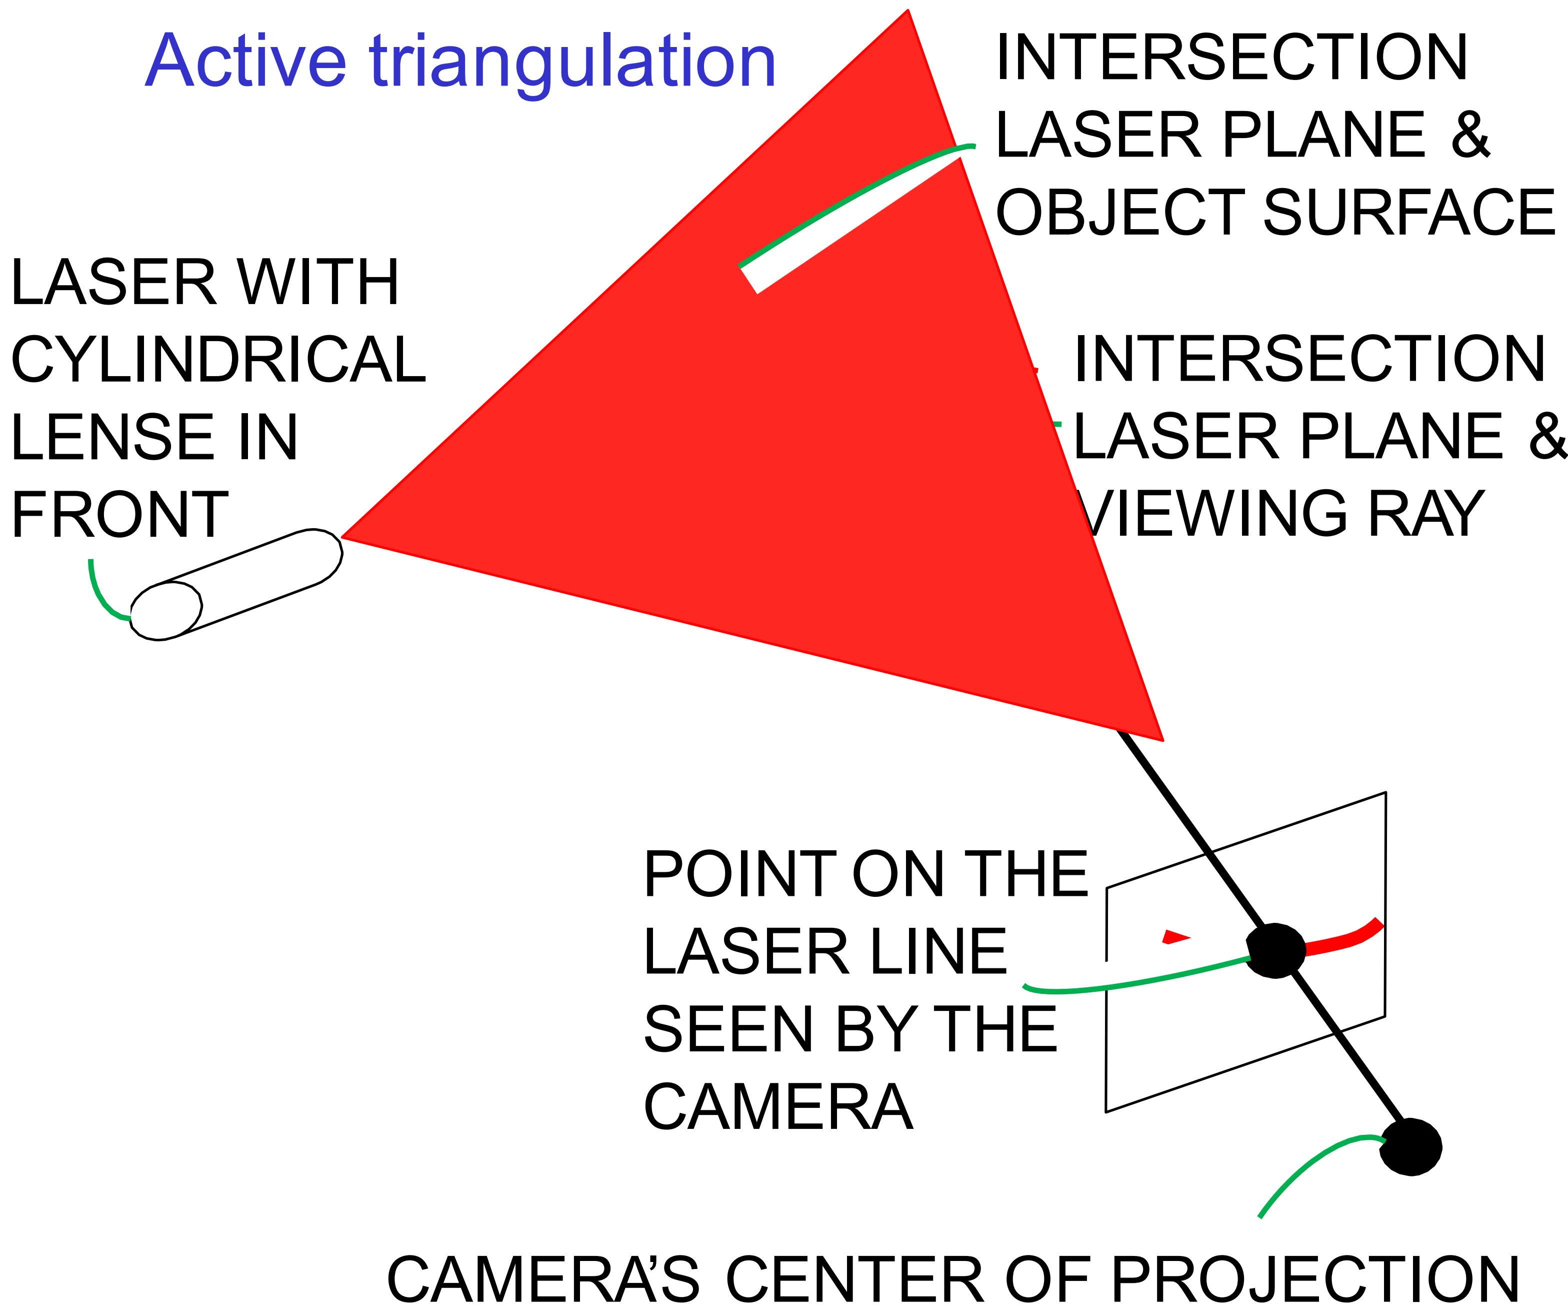

Active triangulation

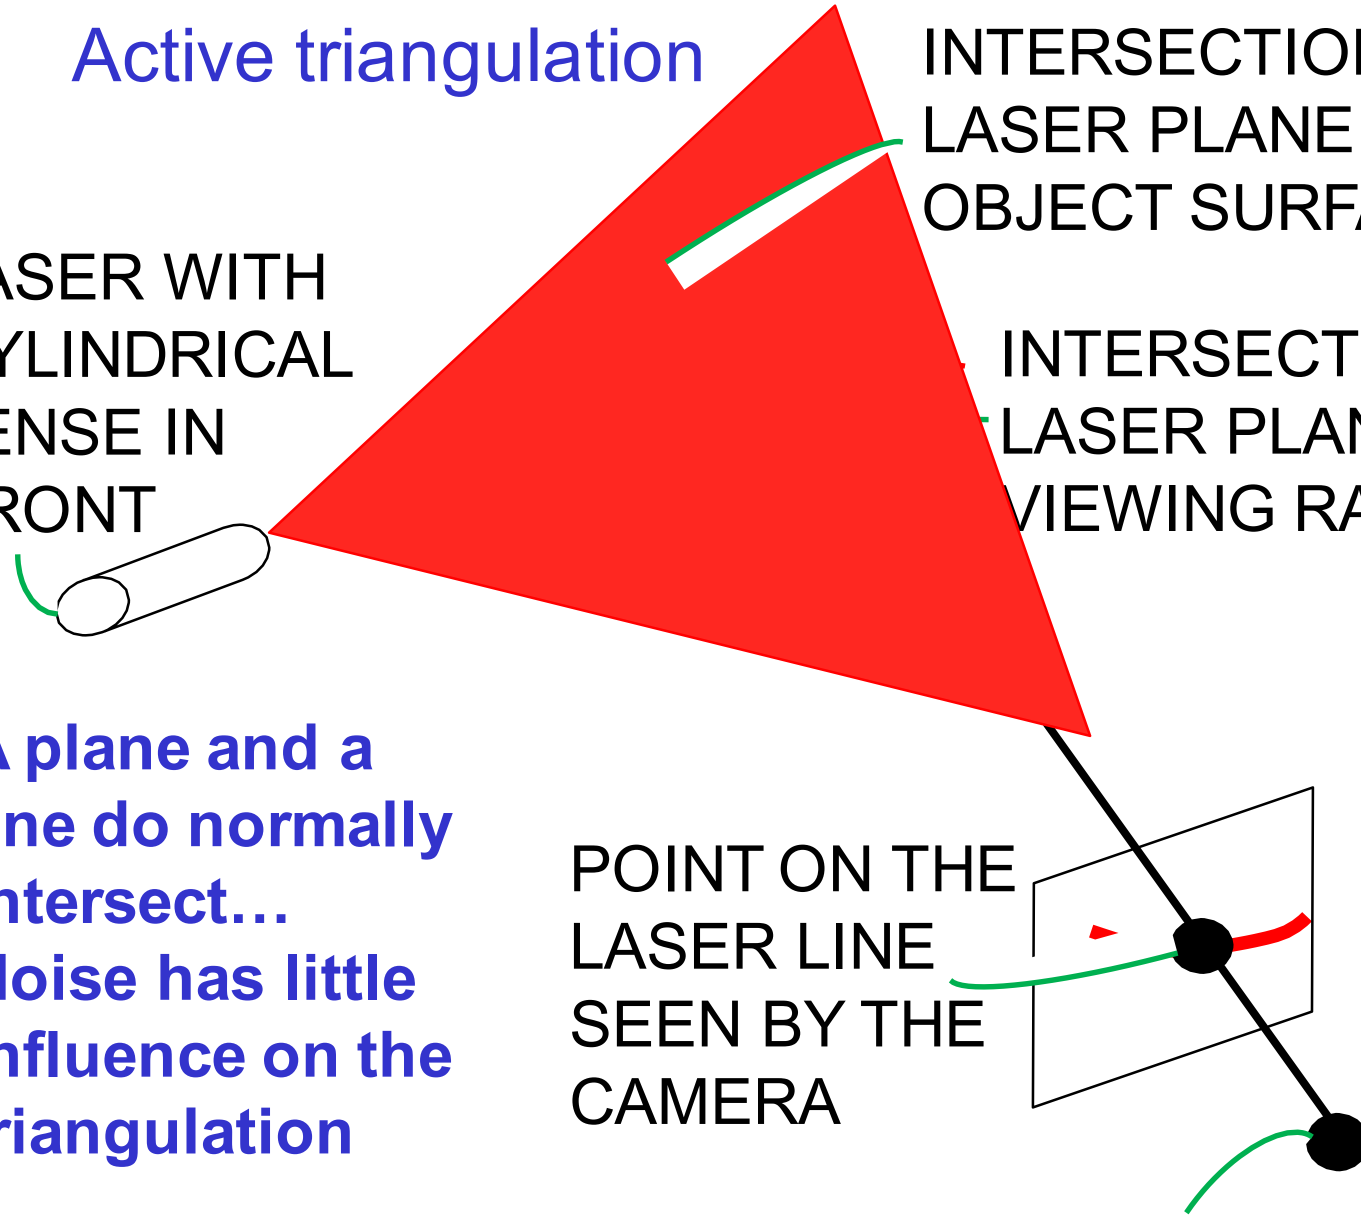

A plane and a line do normally intersect…
Noise has little Influence on the triangulation

Active triangulation

Active triangulation

Triangulation à 3D measurements

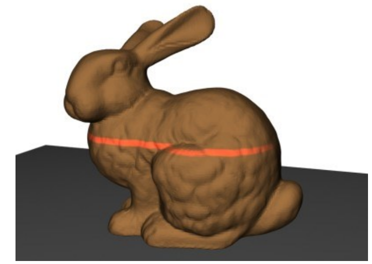

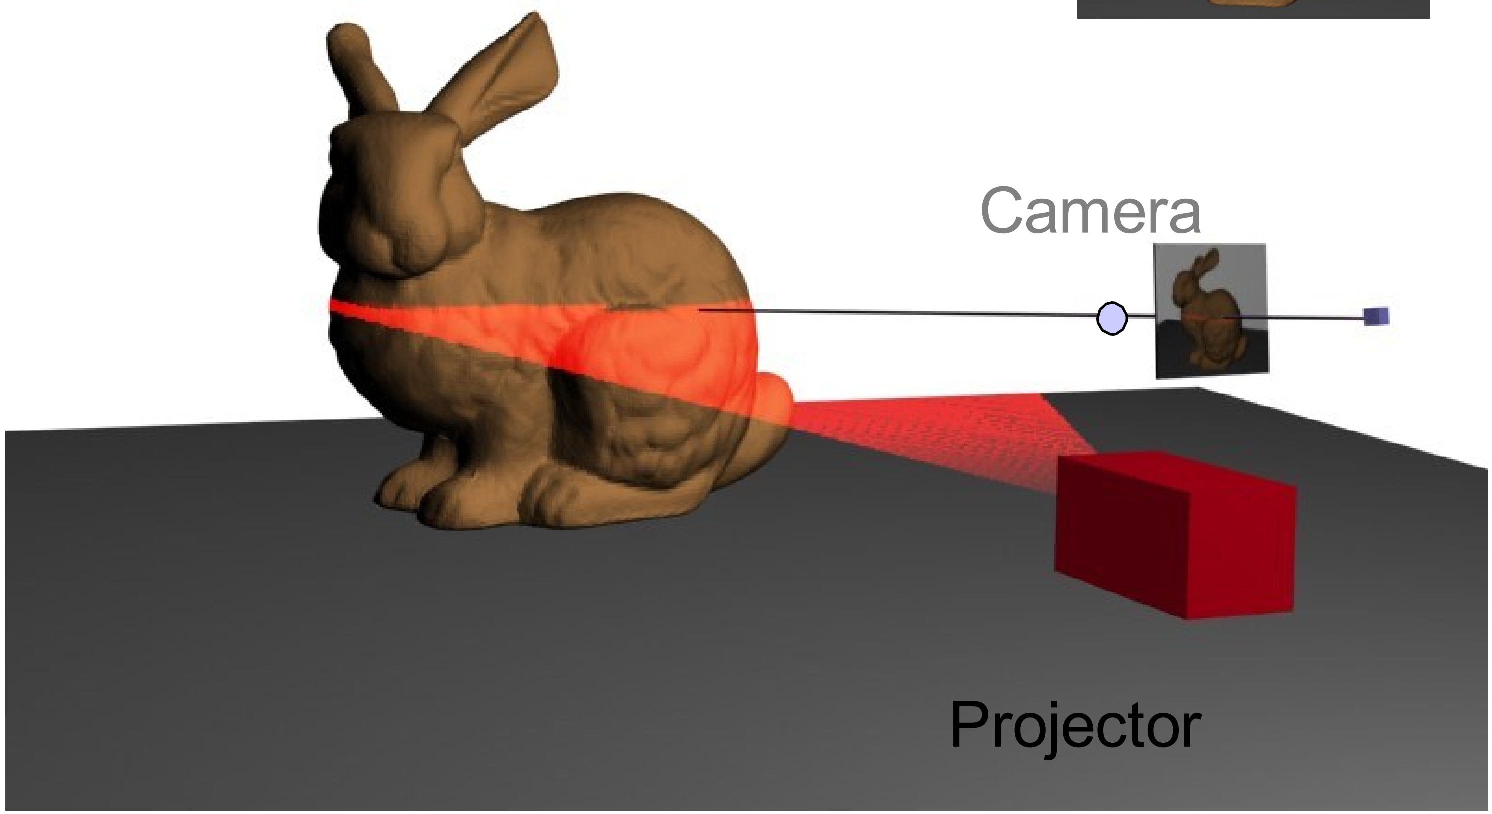

Active triangulation

Camera image

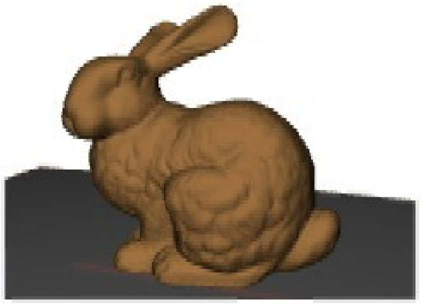

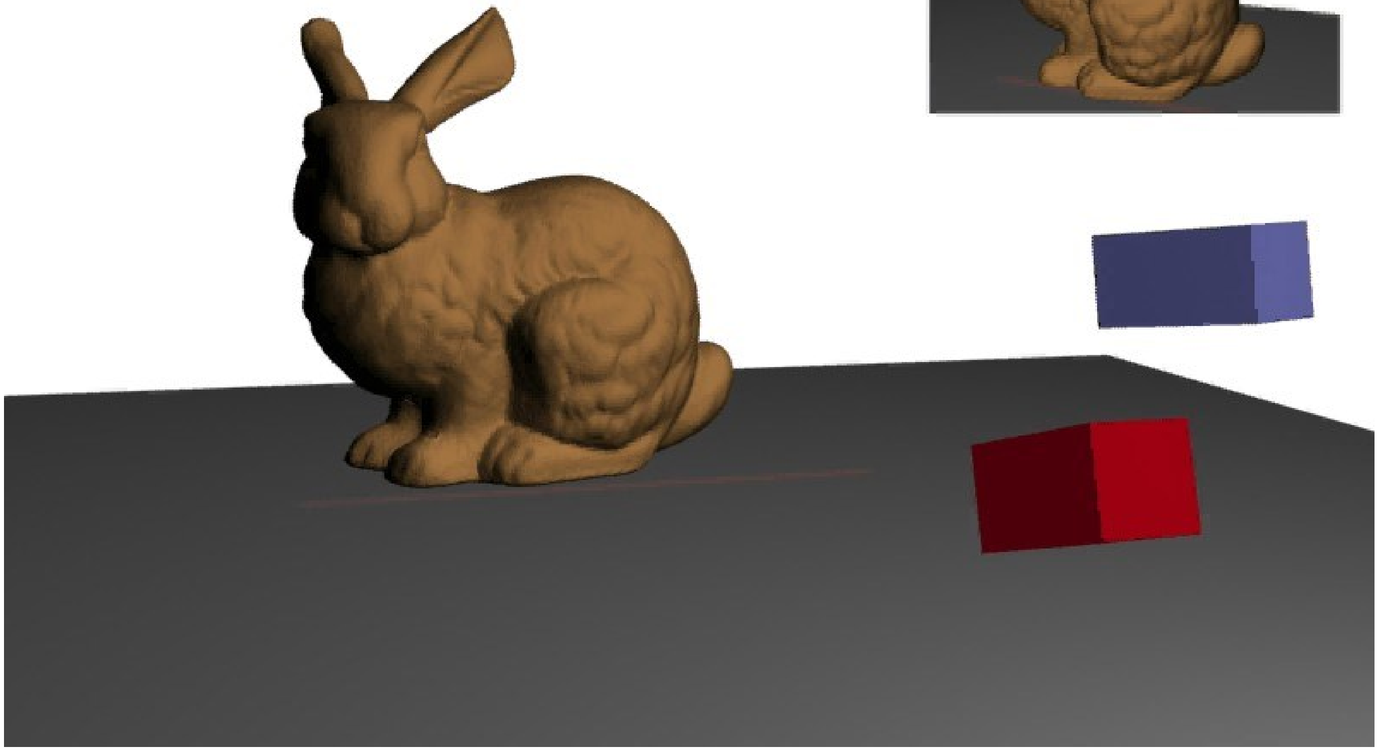

Active triangulation

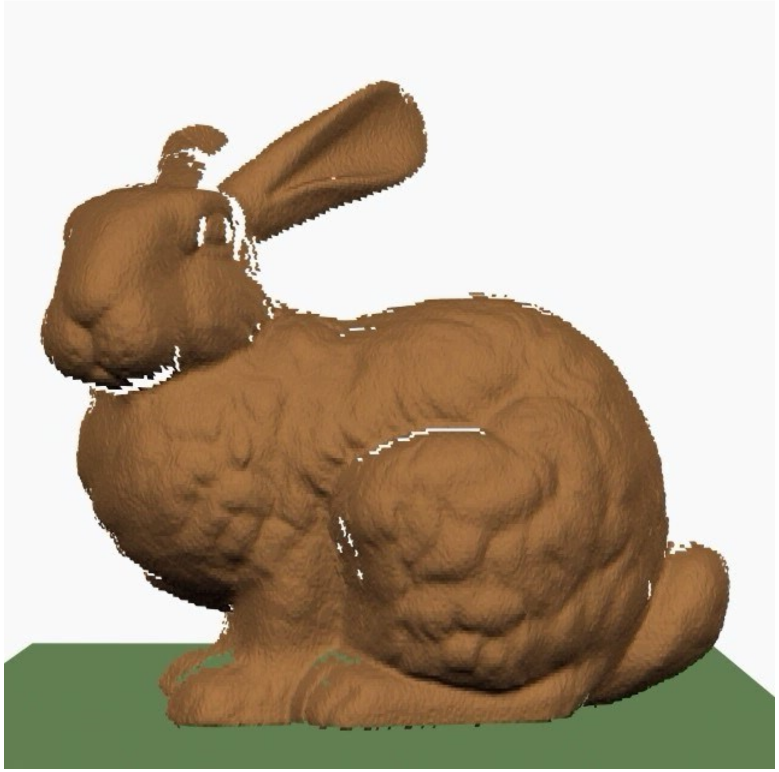

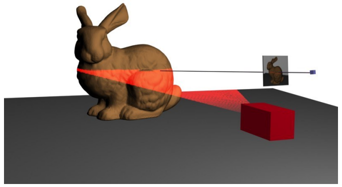

Active triangulation

Example 1 Cyberware laser scanners

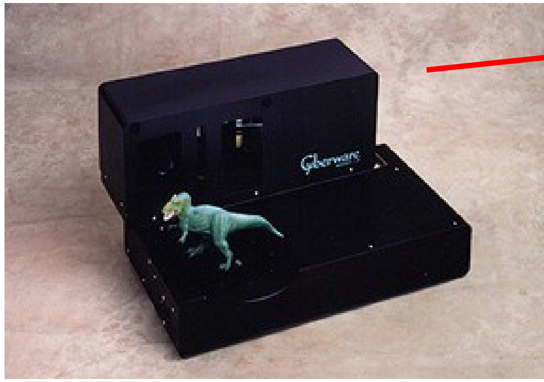

Desktop model
for small objects

Medium-sized objects

Body scanner

Head scanner

è

Active triangulation

Example 2 Minolta

Portable desktop model

https://www.youtube.com/watch?v=R_F66gwXSik

è

3D acquisition taxonomy

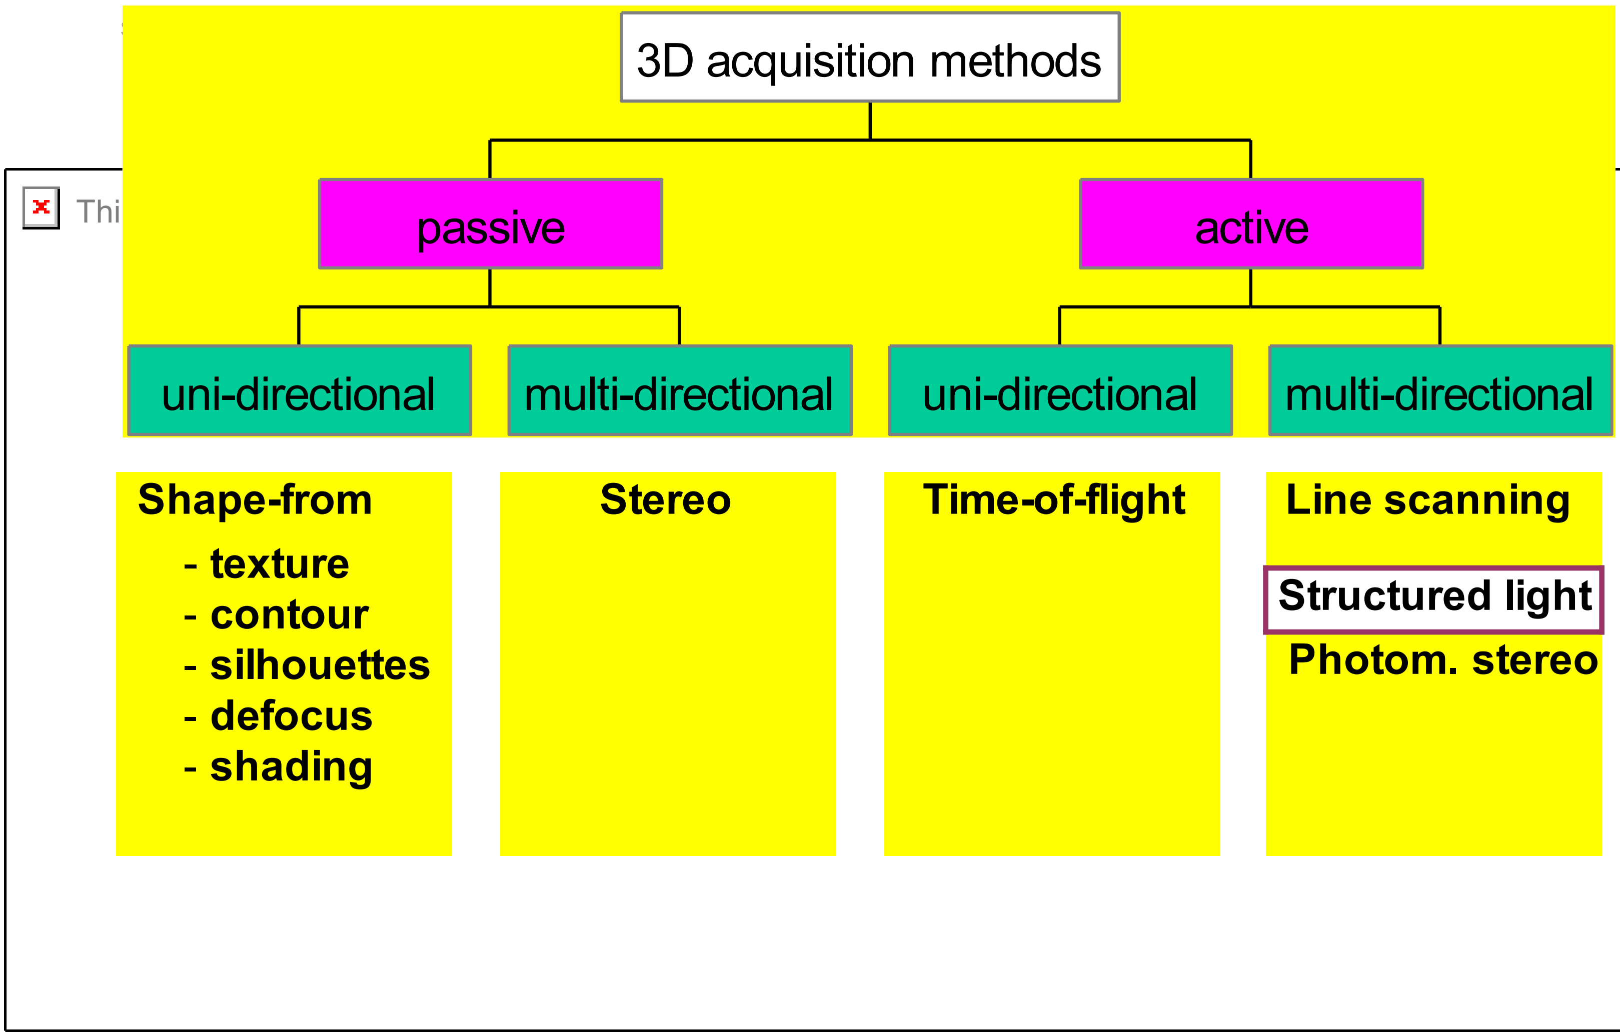

Structured Light Reconstruction

Avoid problems due to correspondence
Avoid problems due to surface appearance
Much more accurate
Very popular in industrial settings

Reading:

Marc Levoy's webpages (Stanford)
Katsu Ikeuchi's webpages (U Tokyo)
Peter Allen's webpages (Columbia)

https://www.youtube.com/watch?v=mSsnf5tqXnA

Stereo Triangulation

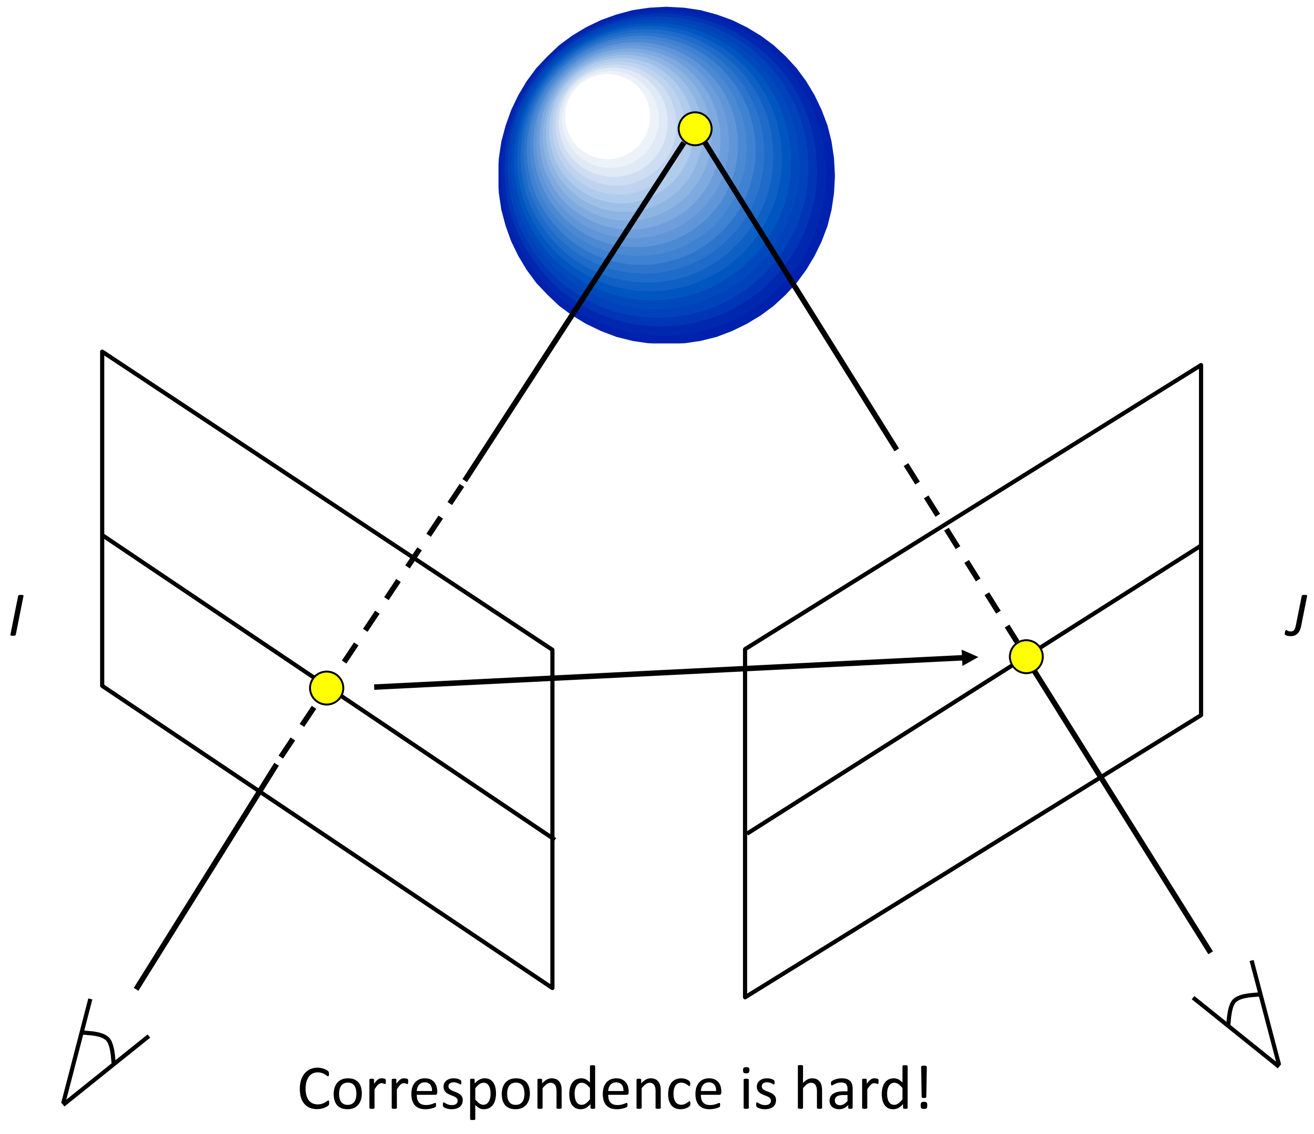

Correspondence is hard!

Structured Light Triangulation

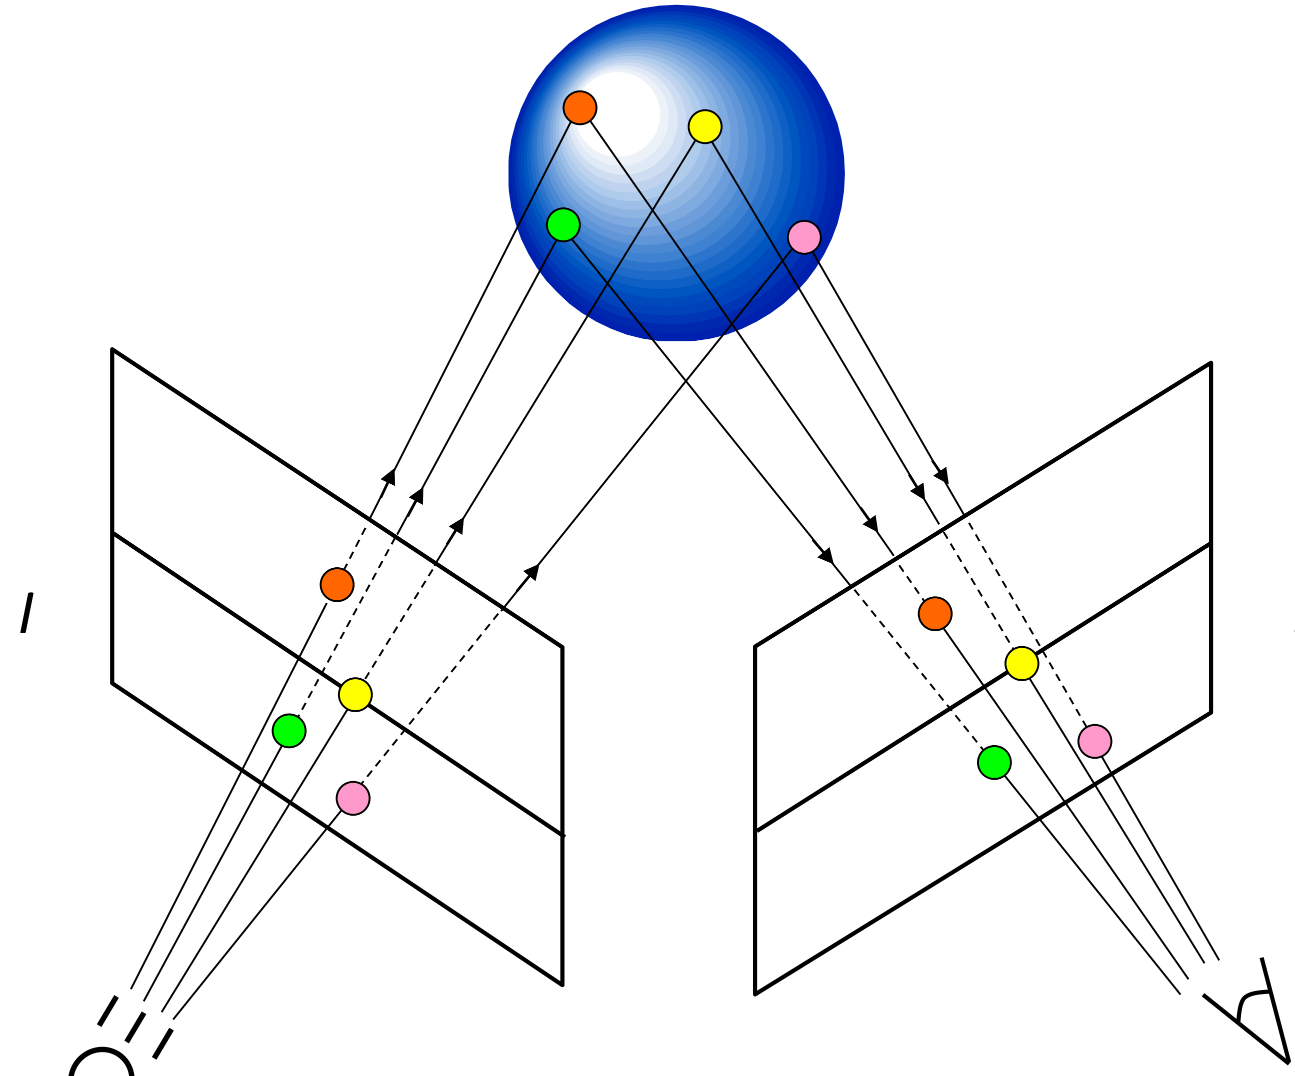

Correspondence becomes easier!

Structured Light

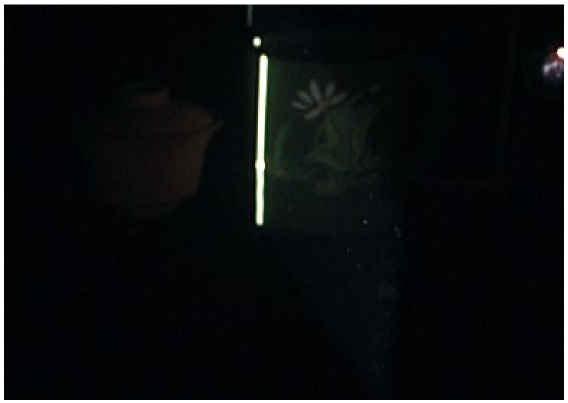 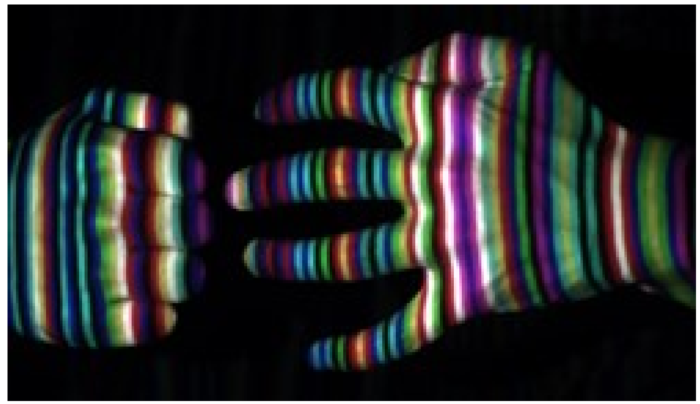 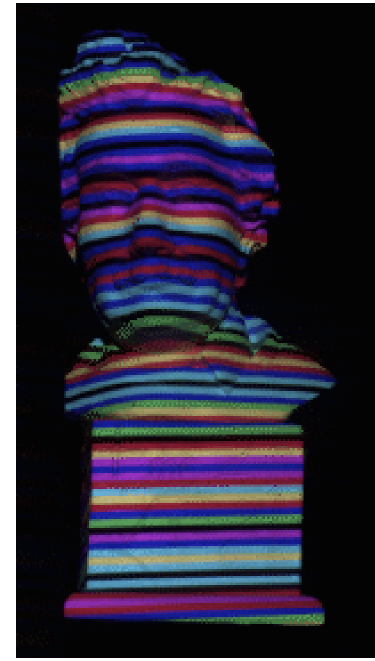

- Any spatio-temporal pattern of light projected on a surface (or volume).
- Cleverly illuminate the scene to extract scene properties (eg., 3D).
- Avoids problems of 3D estimation in scenes with complex texture/BRDFs.
- Very popular in vision and successful in industrial applications (parts assembly,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

Light Stripe Scanning – Single Stripe

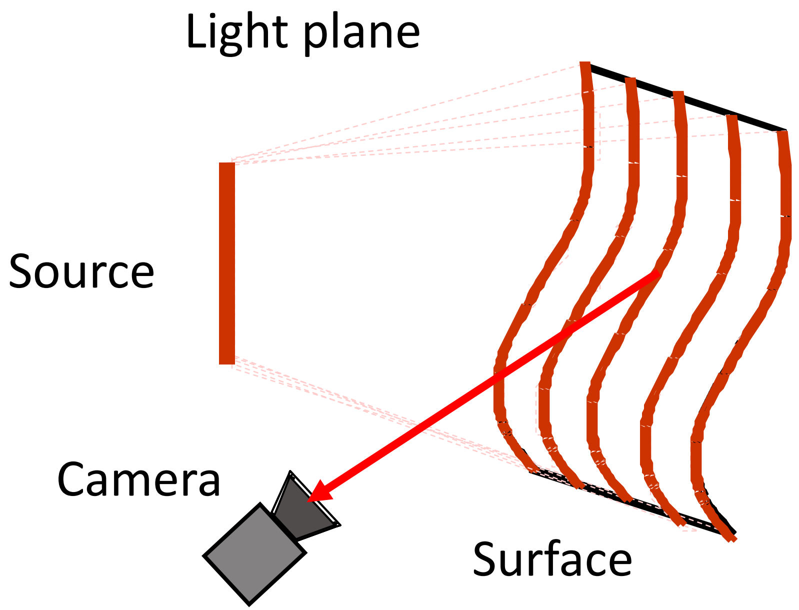

- Optical triangulation
  - Project a single stripe of laser light
  - Scan it across the surface of the object
  - This is a very precise version of structured light scanning
  - Good for high resolution 3D, but needs many images and takes time

Triangulation

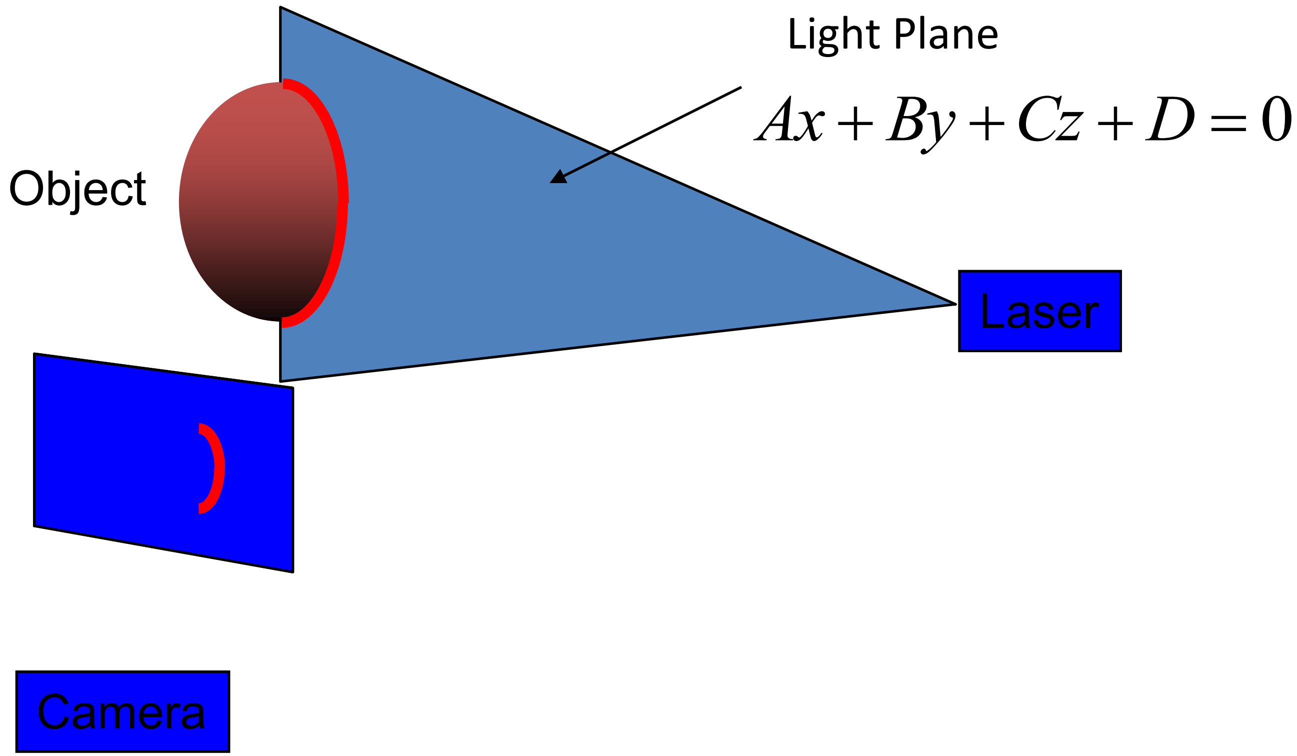

Project laser stripe onto object

Triangulation

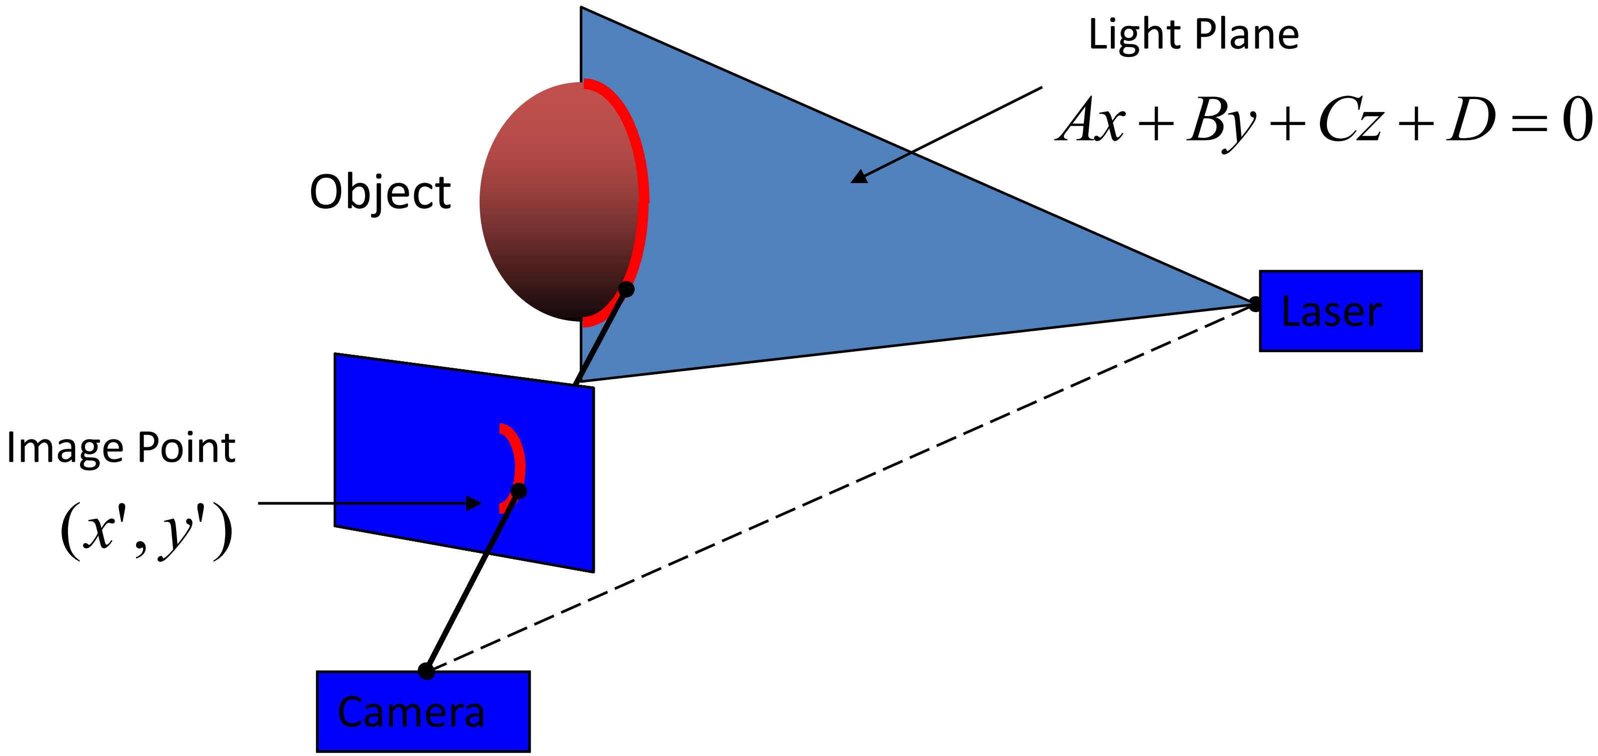

Depth from ray-plane triangulation:

Intersect camera ray with light plane

$$x = x' z / f \qquad z = \frac{-Df}{Ax' + By' + Cf}$$

$$y = y' z / f$$

Example: Laser scanner

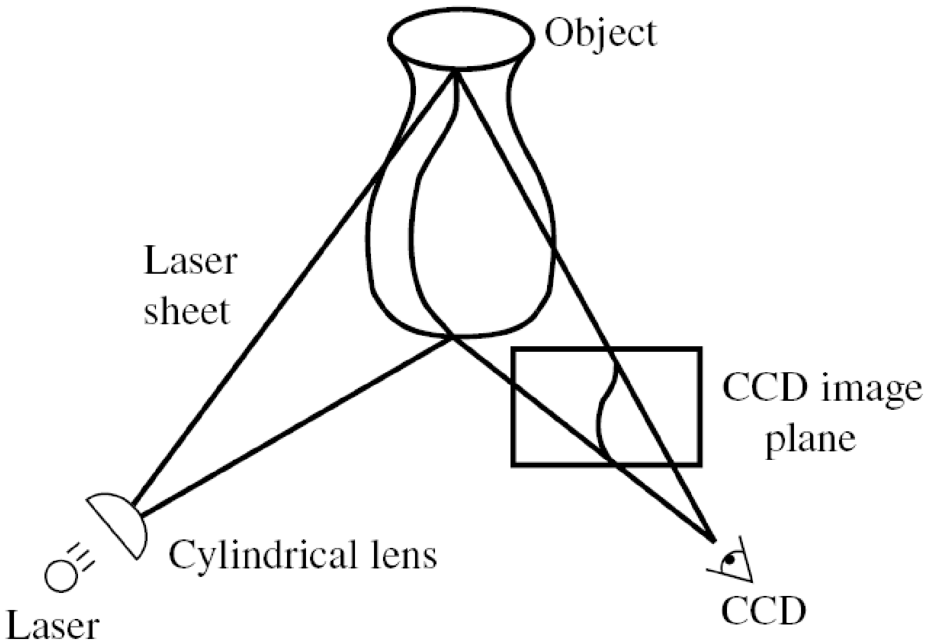

Cyberware® face and head scanner

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

Example: Laser scanner

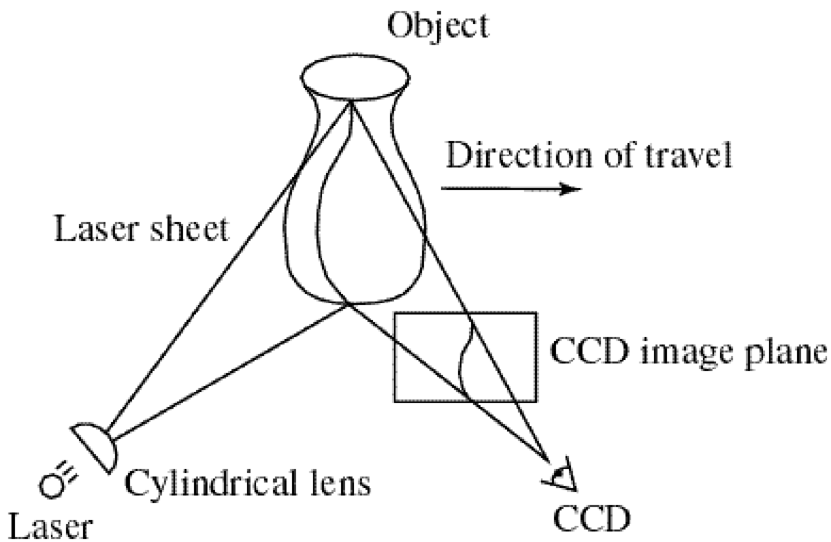

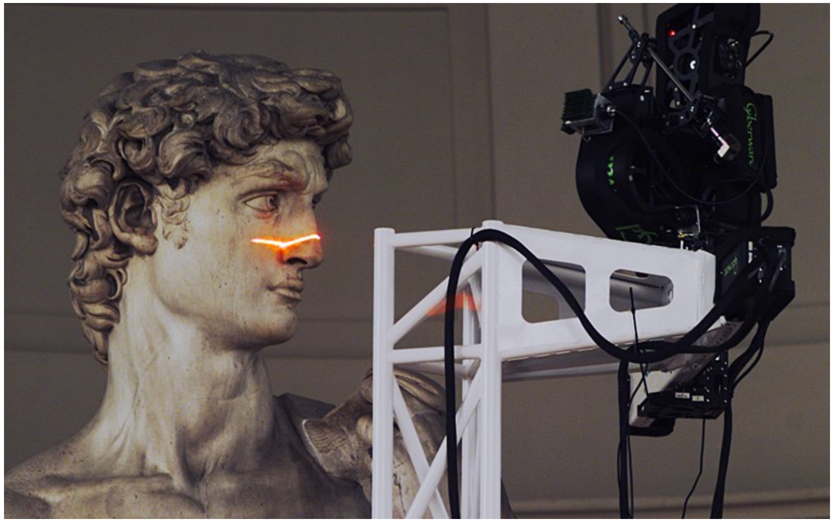

Digital Michelangelo Project

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

Portable 3D laser scanner (this one by Minolta)

Faster Acquisition?

Project multiple stripes simultaneously
Correspondence problem: which stripe is which?

Common types of patterns:

- Binary coded light striping
- Gray/color coded light striping

Binary Coding

Faster:

$2^n - 1$ stripes in n images.

Example:

3 binary-encoded patterns which allows the measuring surface to be divided in 8 sub-regions

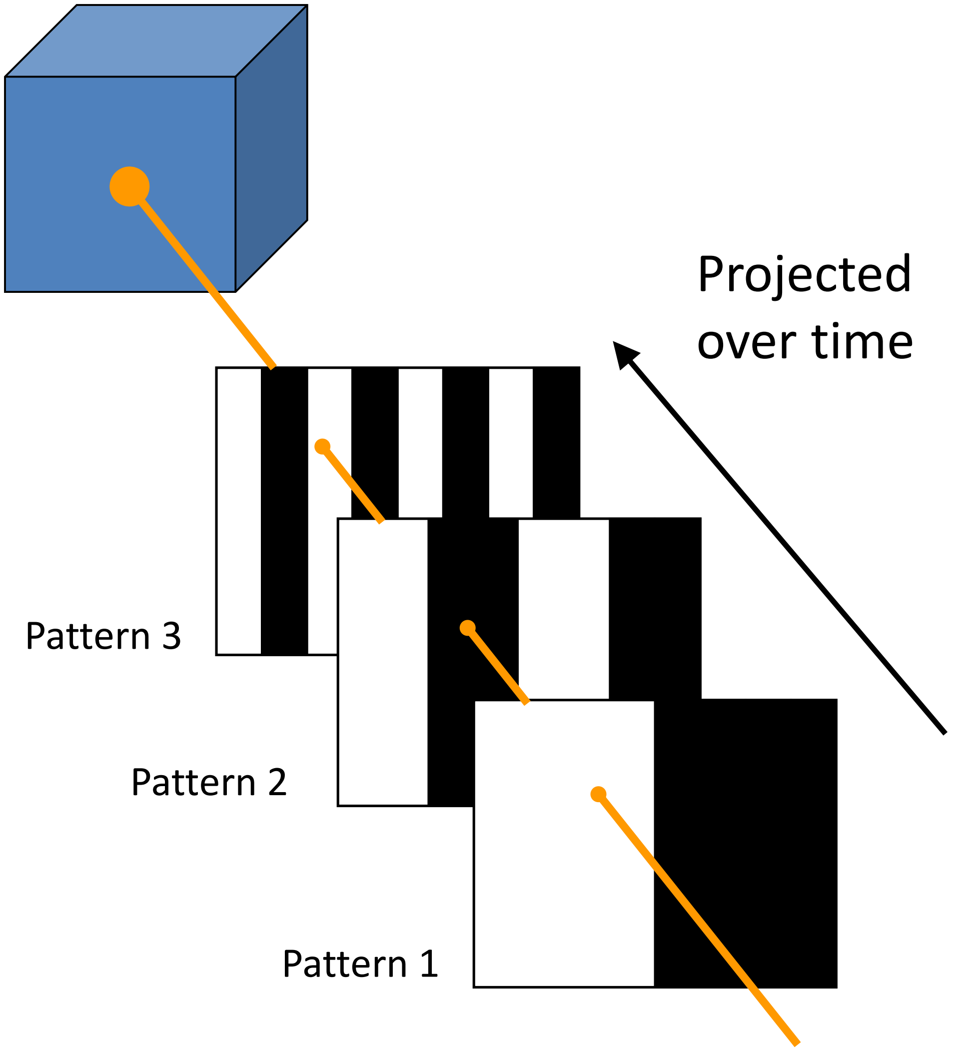

Binary Coding

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

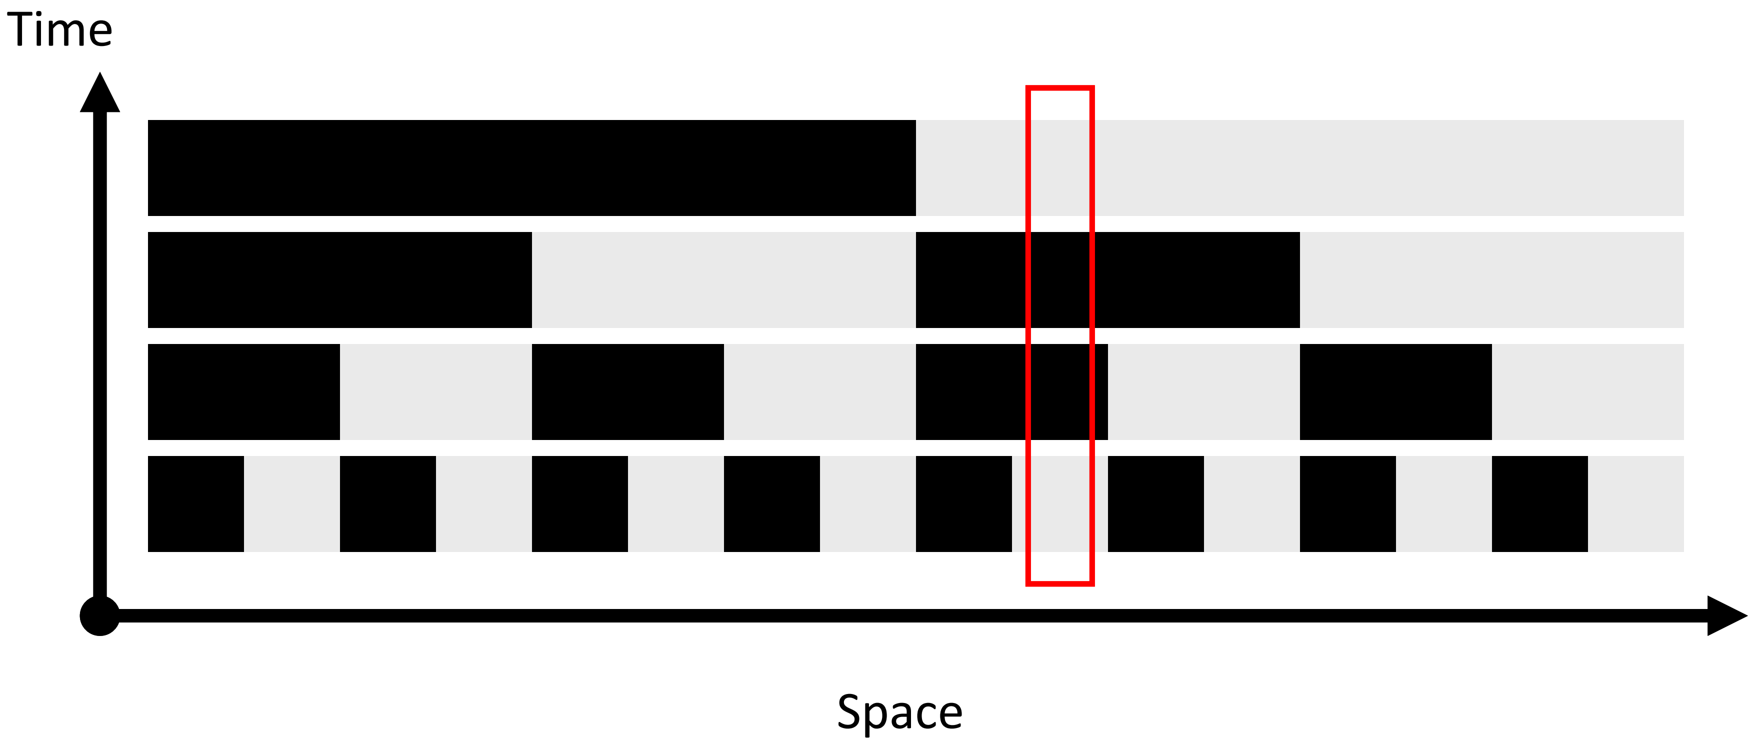

Binary Coding

Example: 7 binary patterns proposed by Posdamer & Altschuler

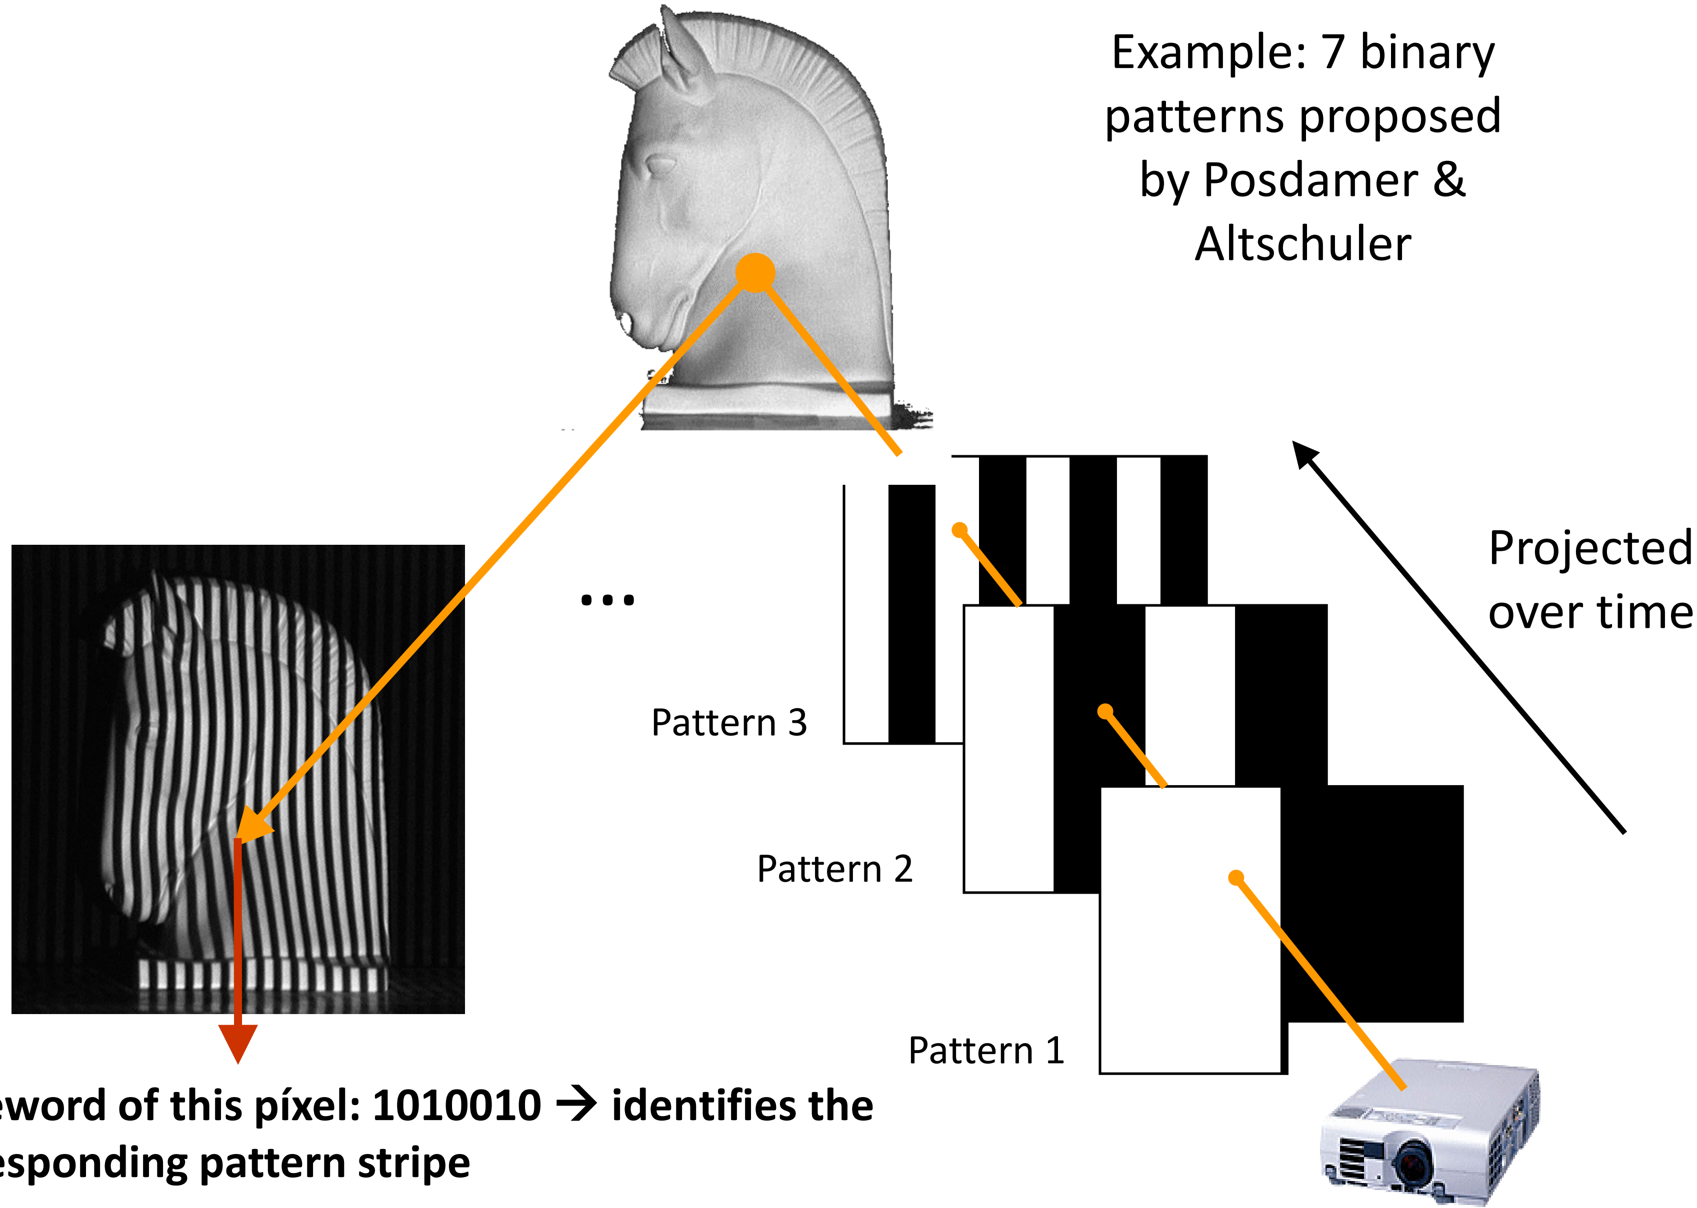

Codeword of this píxel: 1010010 → identifies the corresponding pattern stripe

More complex patterns

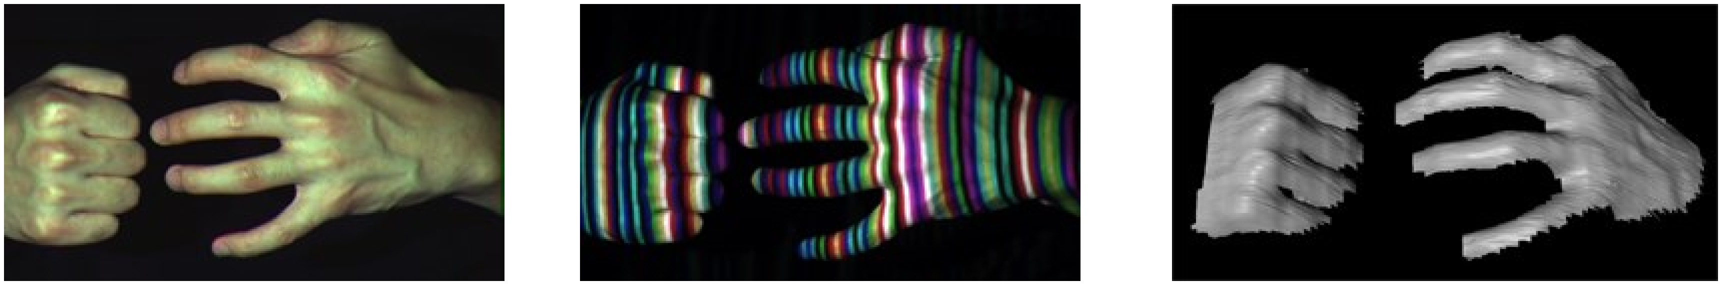

Works despite complex appearances

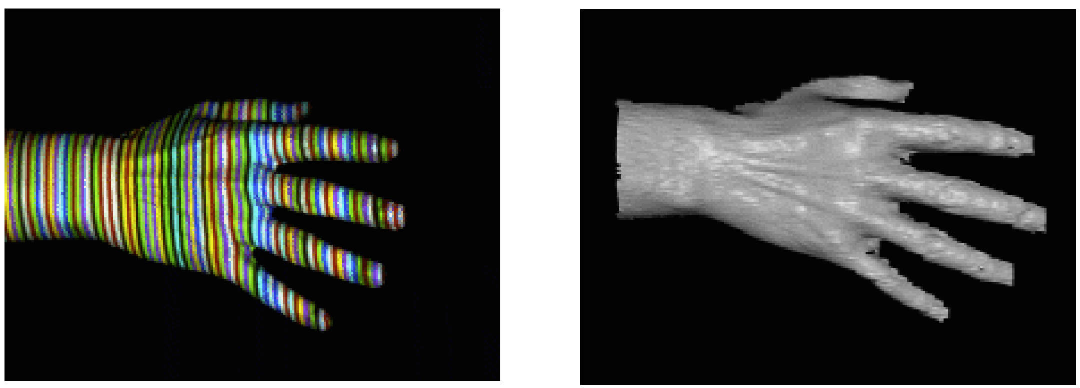

Works in real-time and on dynamic scenes

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

Zhang et al

Continuum of Triangulation Methods

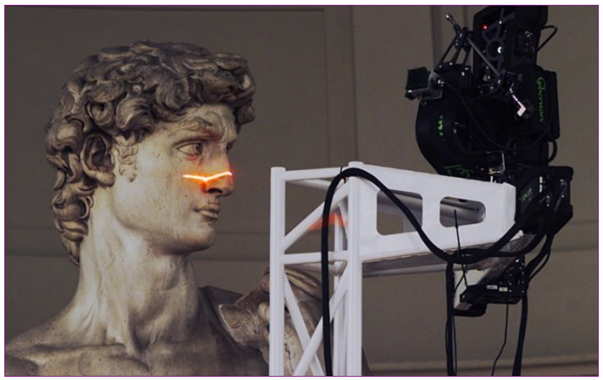

Multi-stripe
Multi-frame

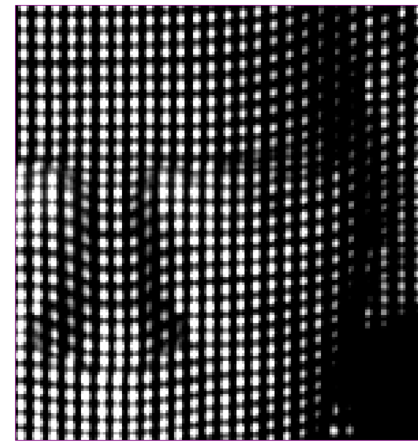

Single-stripe

Single-frame

Slow, robust

Fast, fragile

One-shot implementation

KULeuven '81: checkerboard pattern with column code

example :

è

3D reconstruction for the example

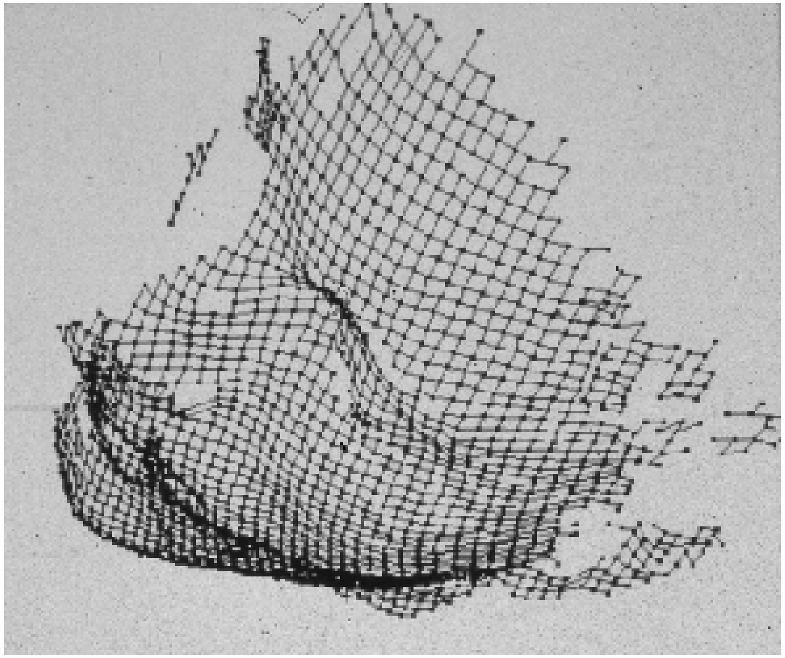 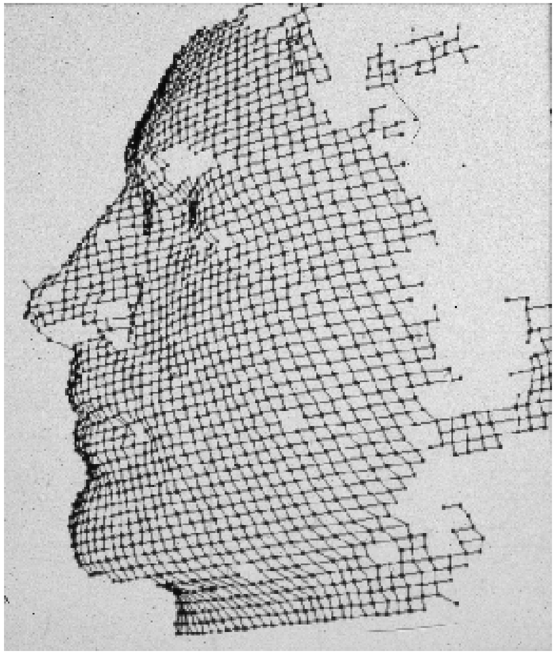

è

An application in agriculture

One-shot 3D acquisition

Leuven ShapeCam

Shape + texture often needed

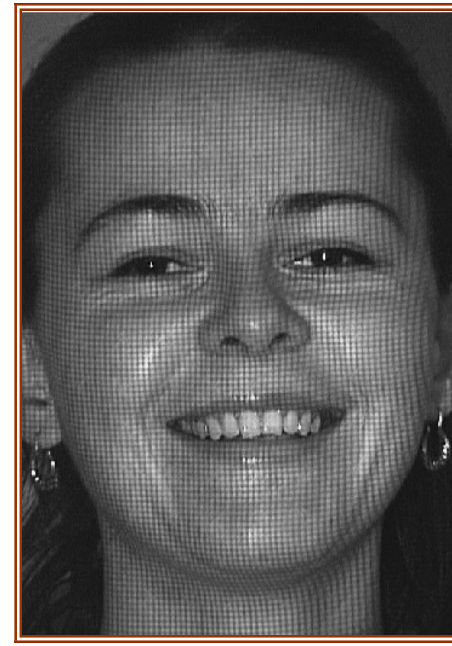

Higher resolution

Texture is also extracted

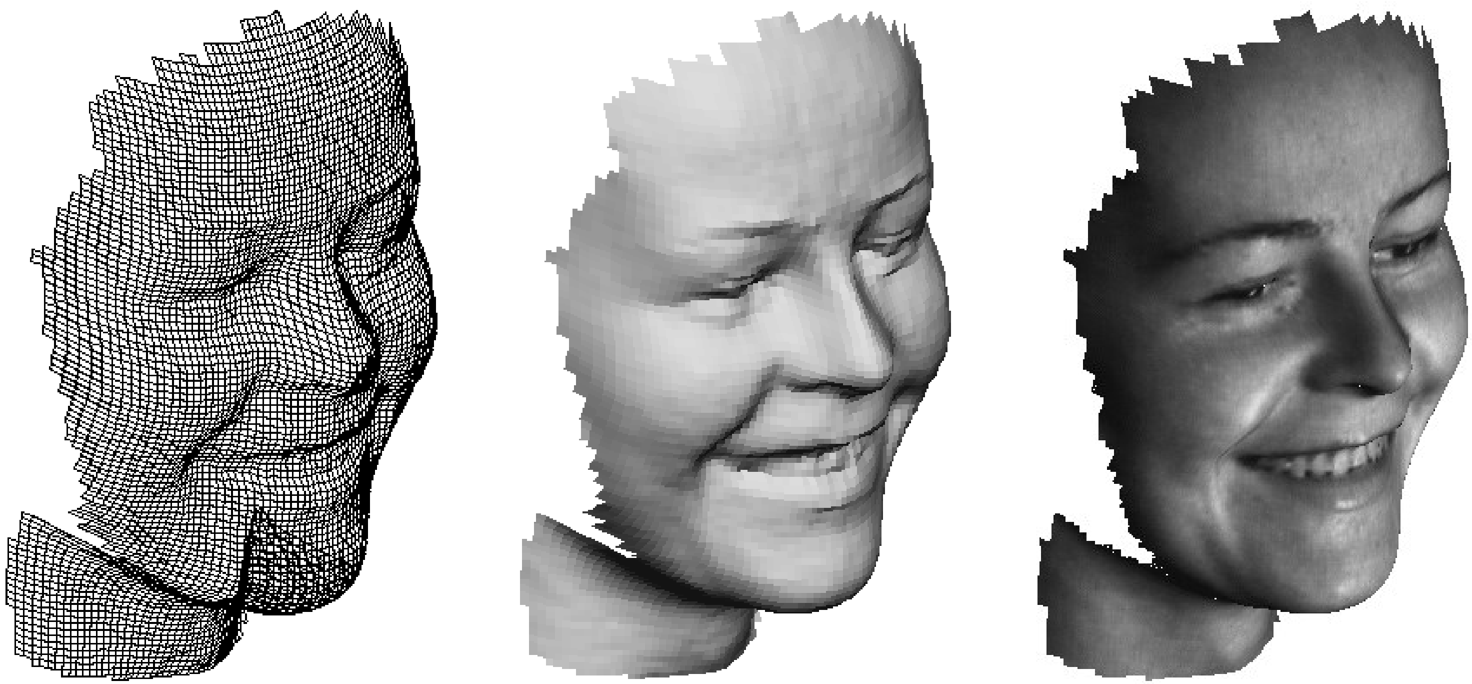

Active triangulation

Recent, commercial example

Kinect 3D camera, affordable
and compact solution by Microsoft.

Projects a 2D point pattern in the NIR,
to make it invisible to the human eye

è

Kinect: 9x9 patches with locally unique code

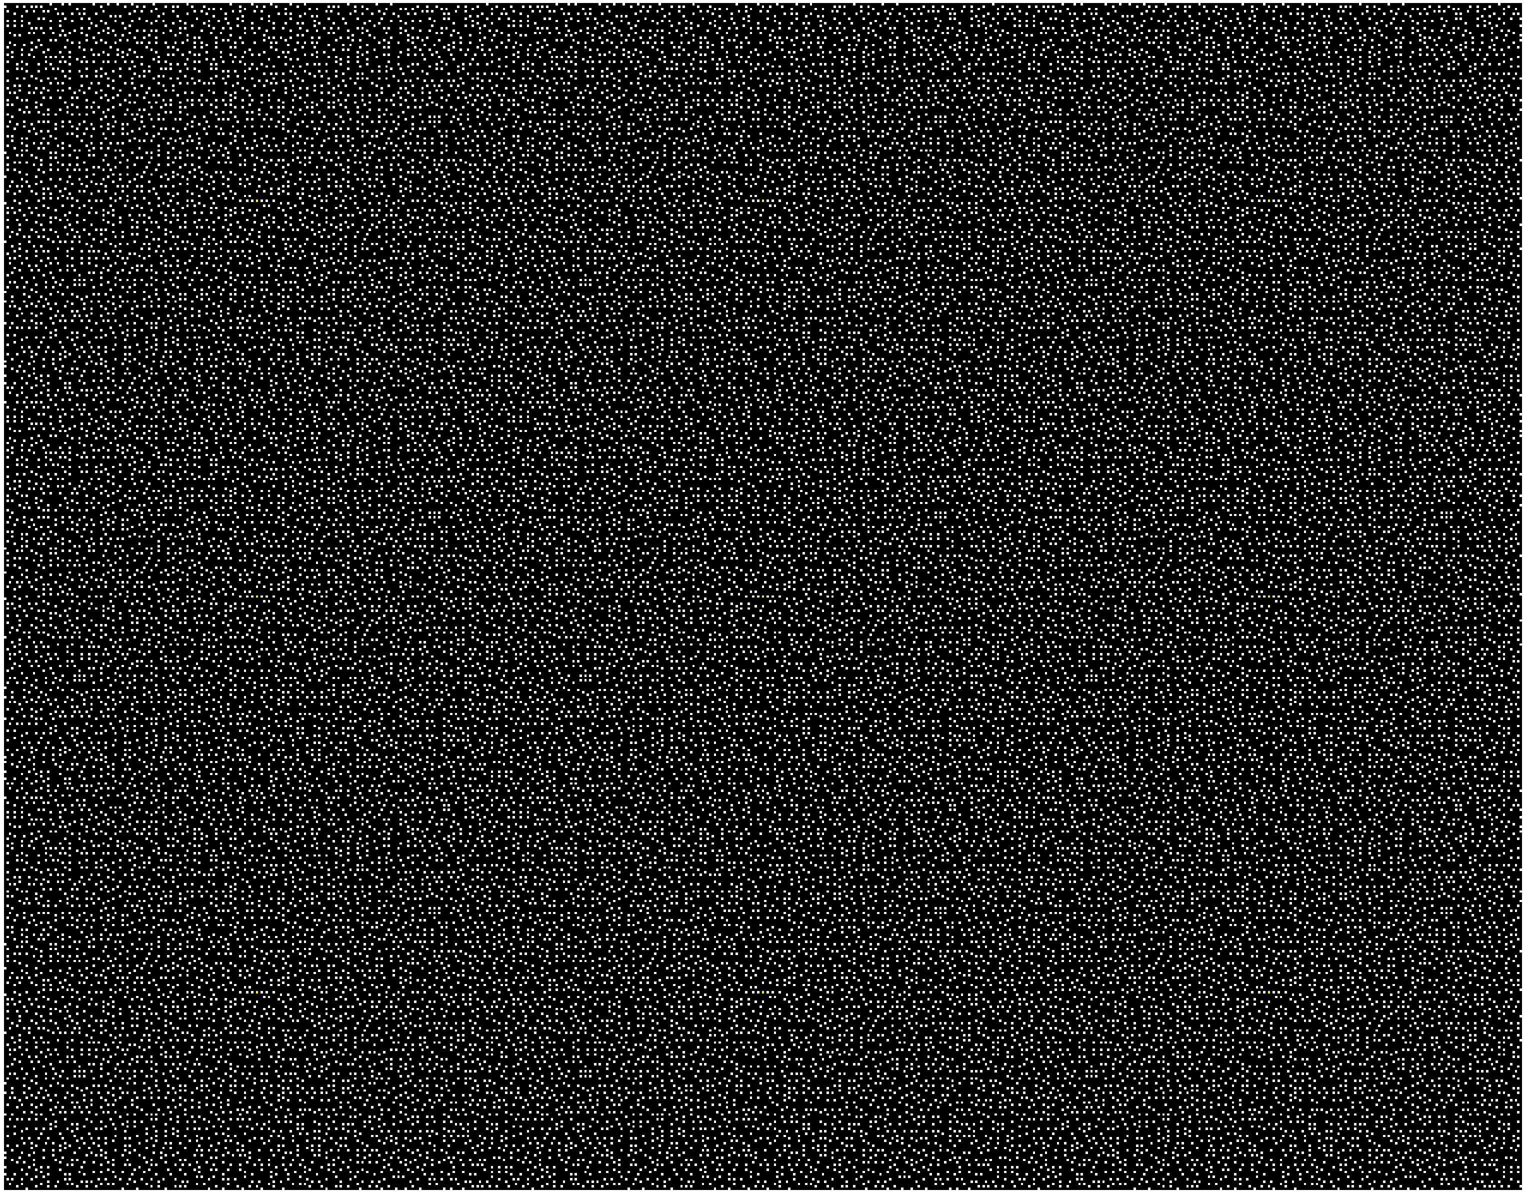

è

What’s in the sensor?

What’s in the sensor?

Color Sensor

IR Emitter
IR Depth Sensor

Inside the Kinect Depth Camera

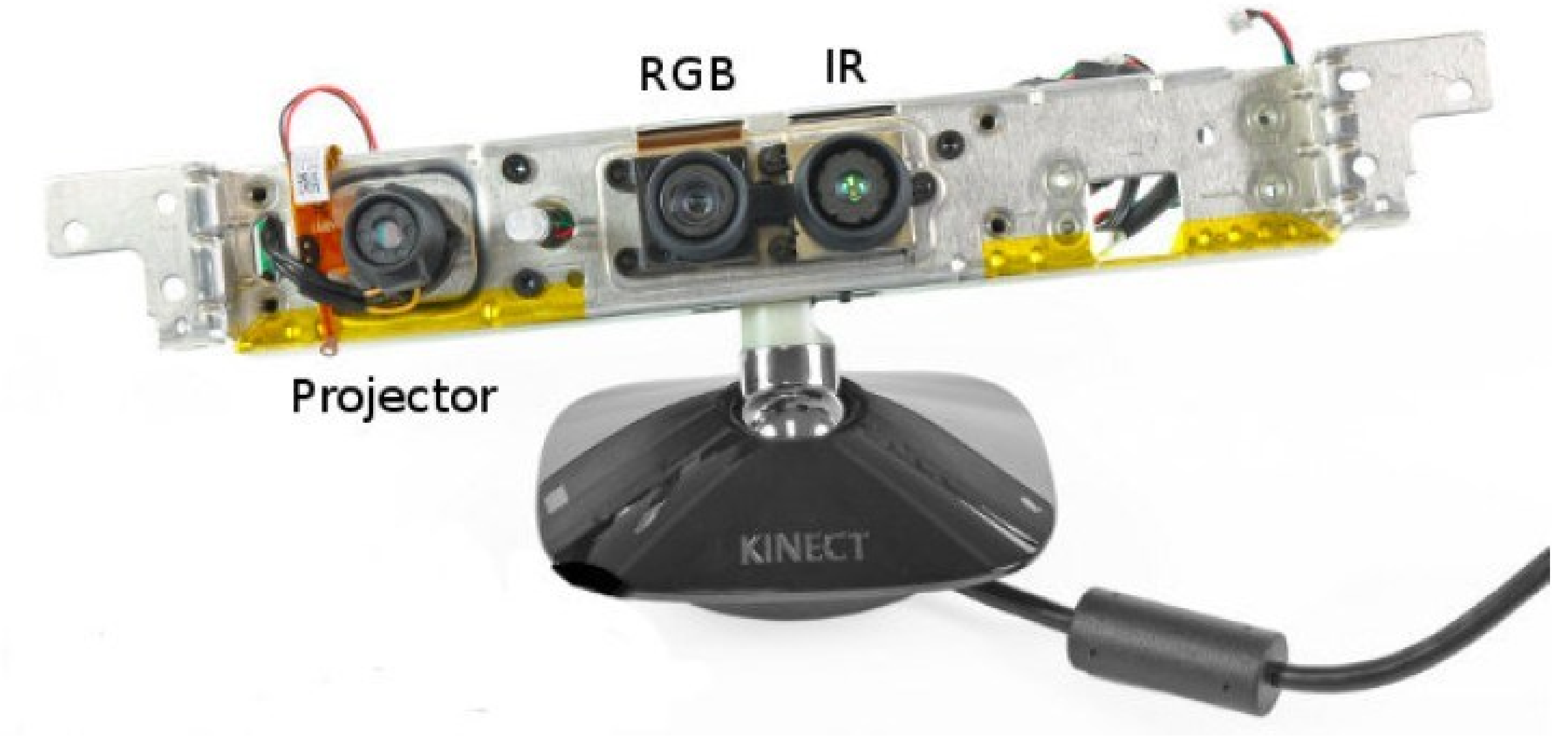

Kinect as one-shot, low-cost scanner

Excerpt from the dense NIR dot pattern:

Face animation - input

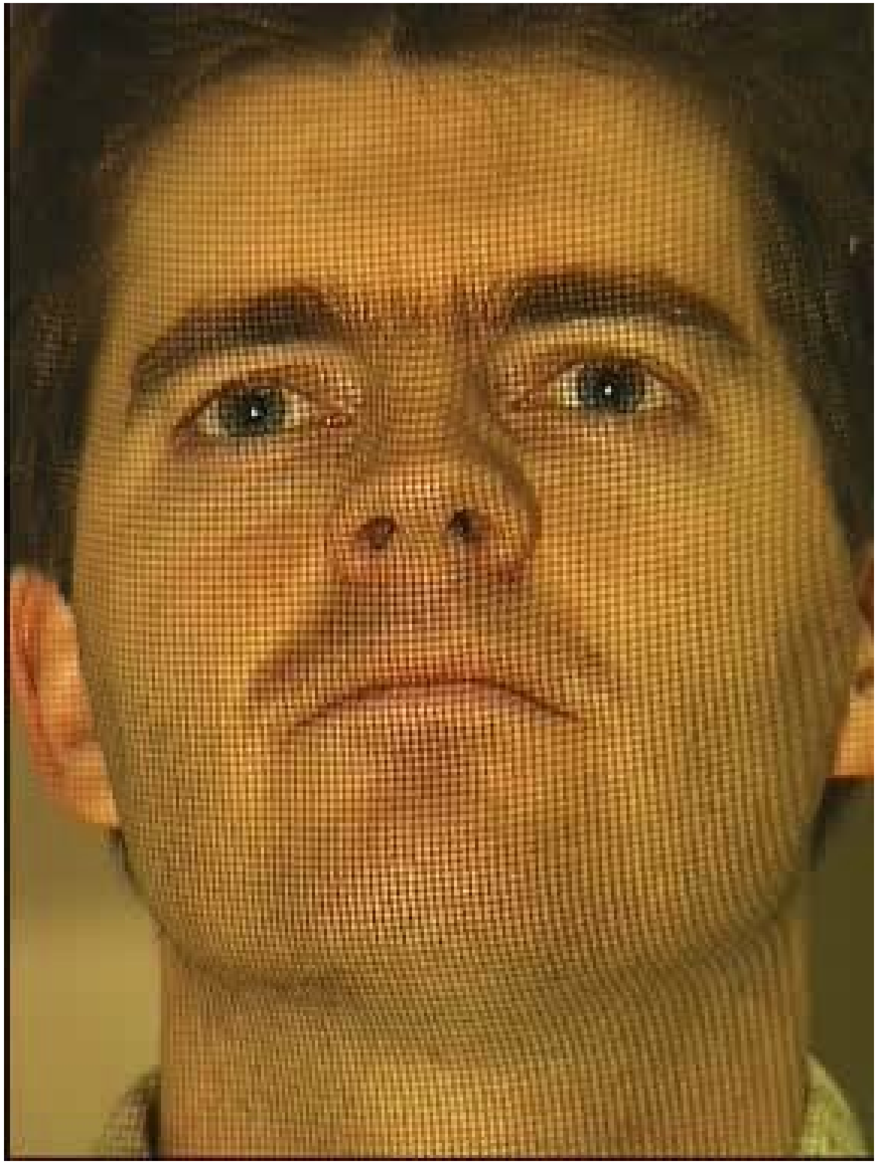

Face animation – replay + effects

3D acquisition taxonomy

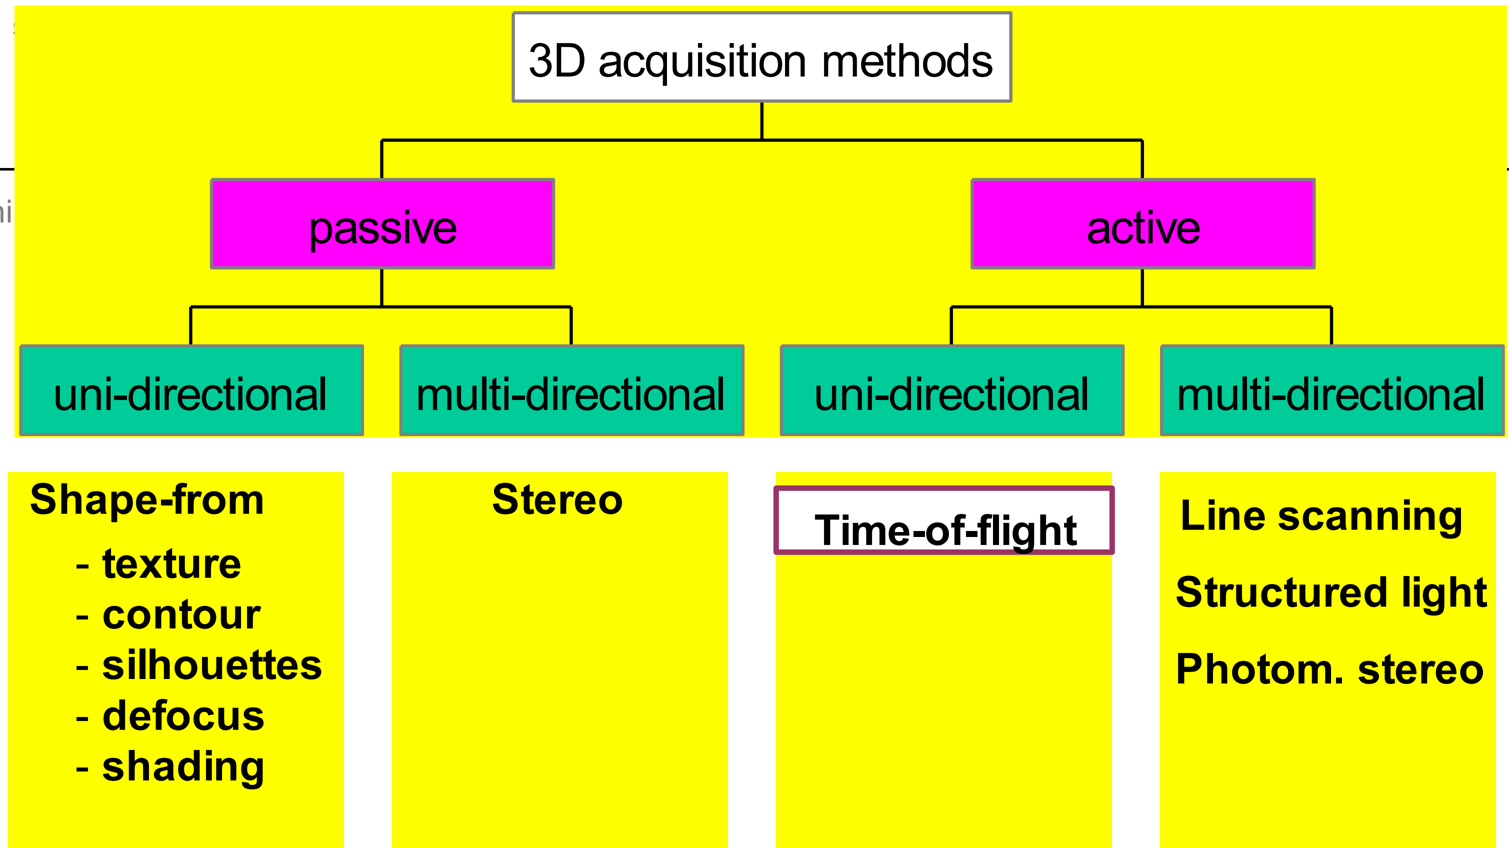

Time-of-flight

measurement of the time a modulated light signal needs to travel before returning to the sensor

this time is proportional to the distance

waves :

1. radar	low freq. electromagnetic
2. sonar	acoustic waves
3. optical radar	optical waves

working principles :

1. pulsed
2. phase shifts

è

Time-of-flight

Example 1: Cyrax

Example 2: Riegl

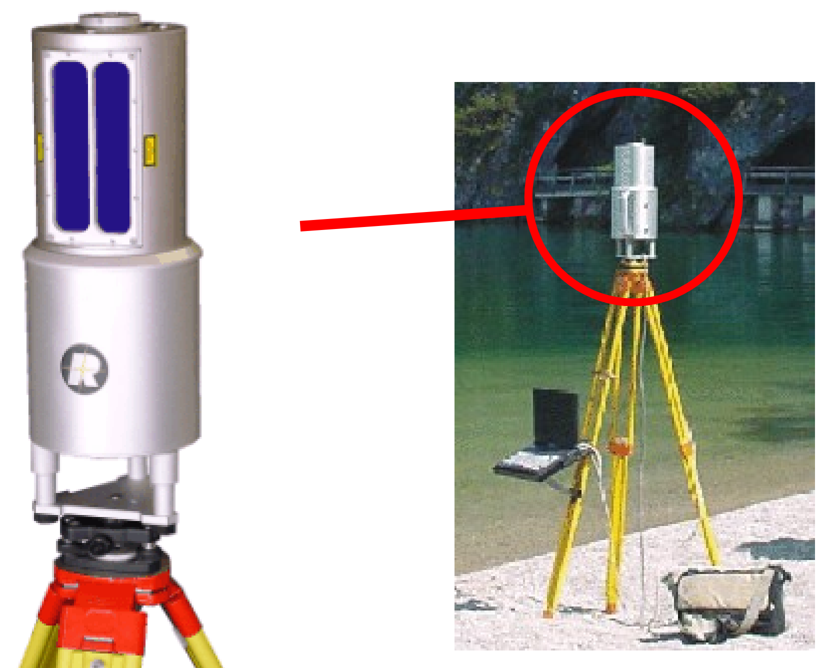

è

Time-of-flight: example

Cyrax **™**

3D Laser Mapping

System

Cyrax

Accurate, detailed, fast measuring

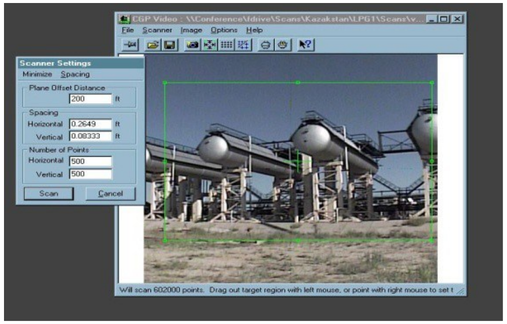

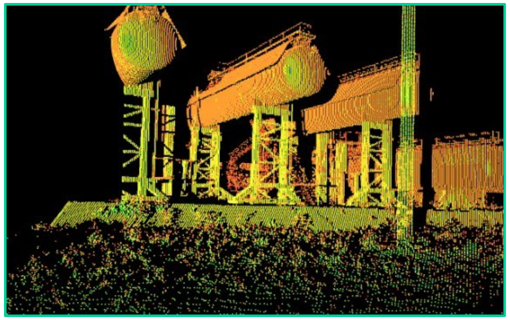

2D / 3D CAD

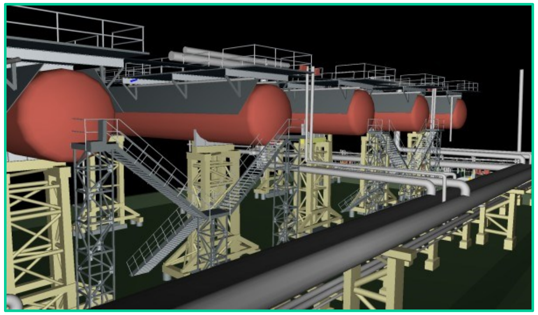

Integrated modeling

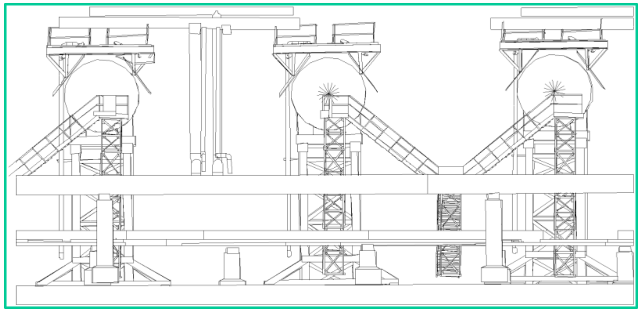

Cyrax (now Leica)

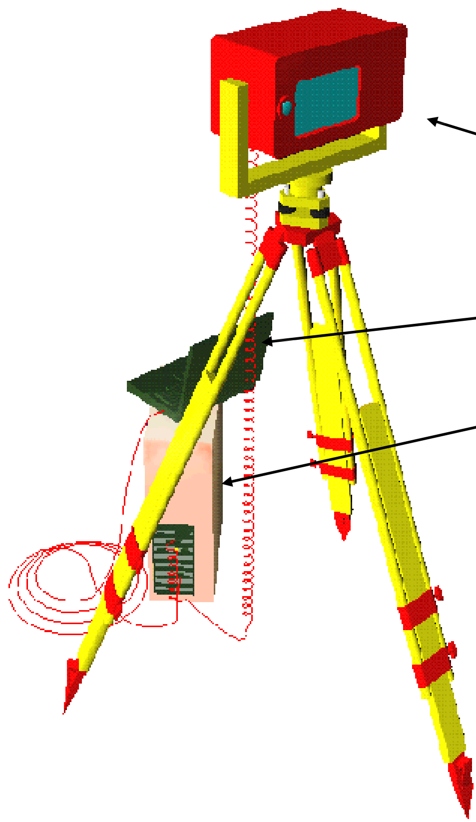

- 3 Components
  - Laser Scanner
  - Laptop Computer & Software
  - Electronics & Power Supply
- Field Portable

Pulsed laser (time-of-flight)

No reflectors needed

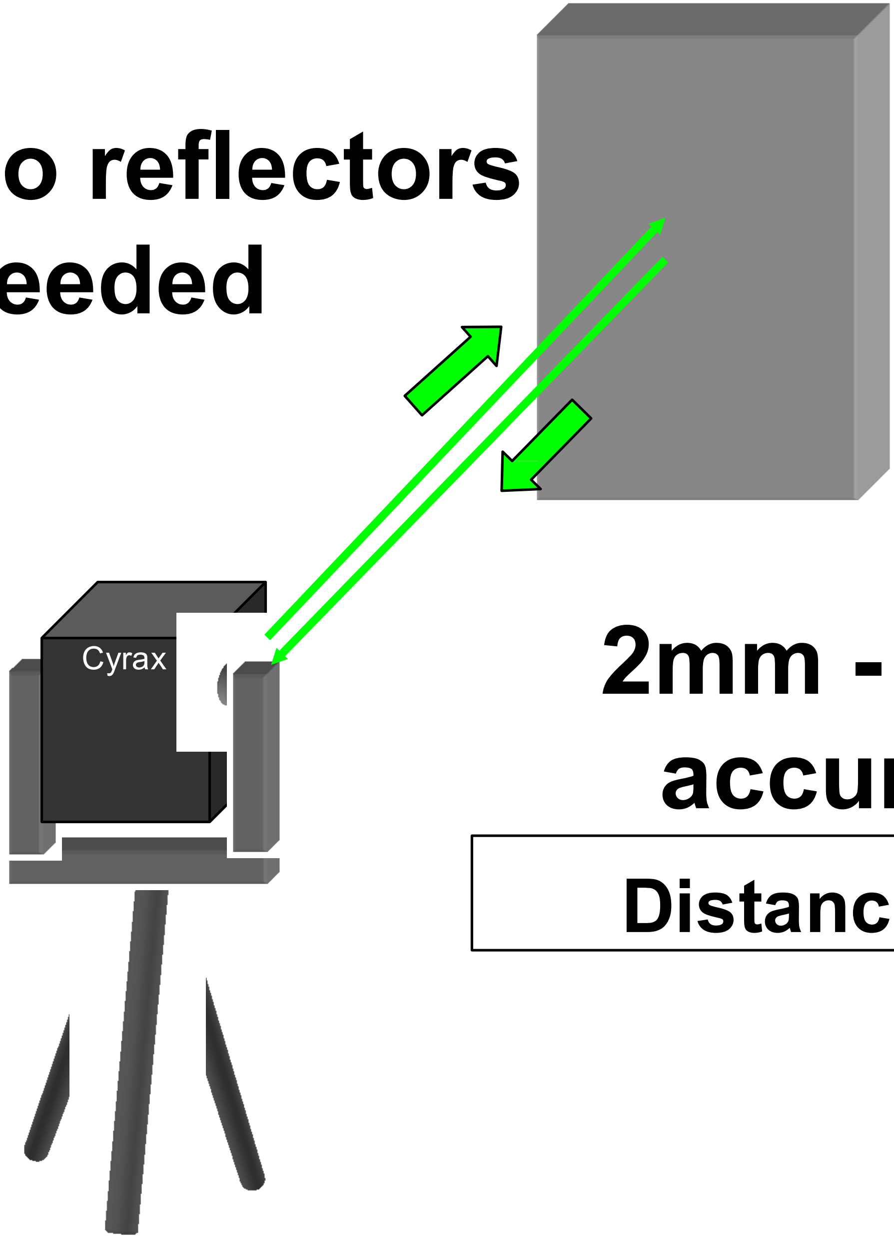

2mm - 6mm accuracy

Distance = C x ΔT÷2

Laser sweeps over surface

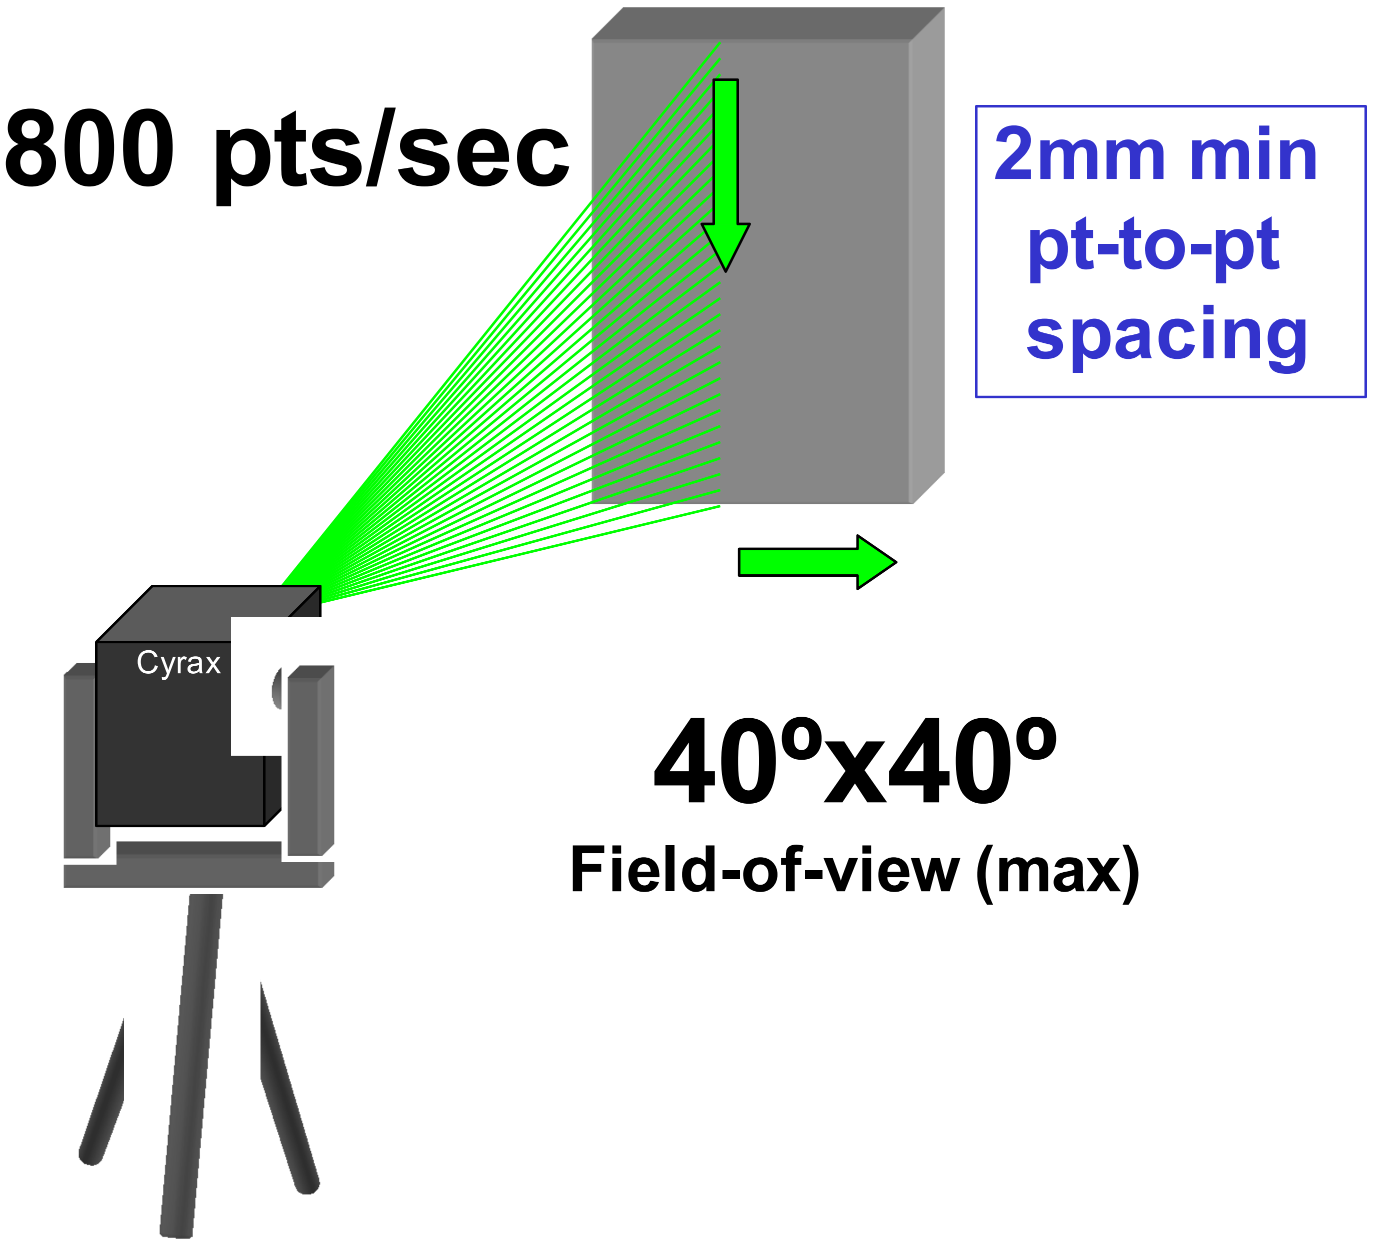

Up to 100m range (50m rec)

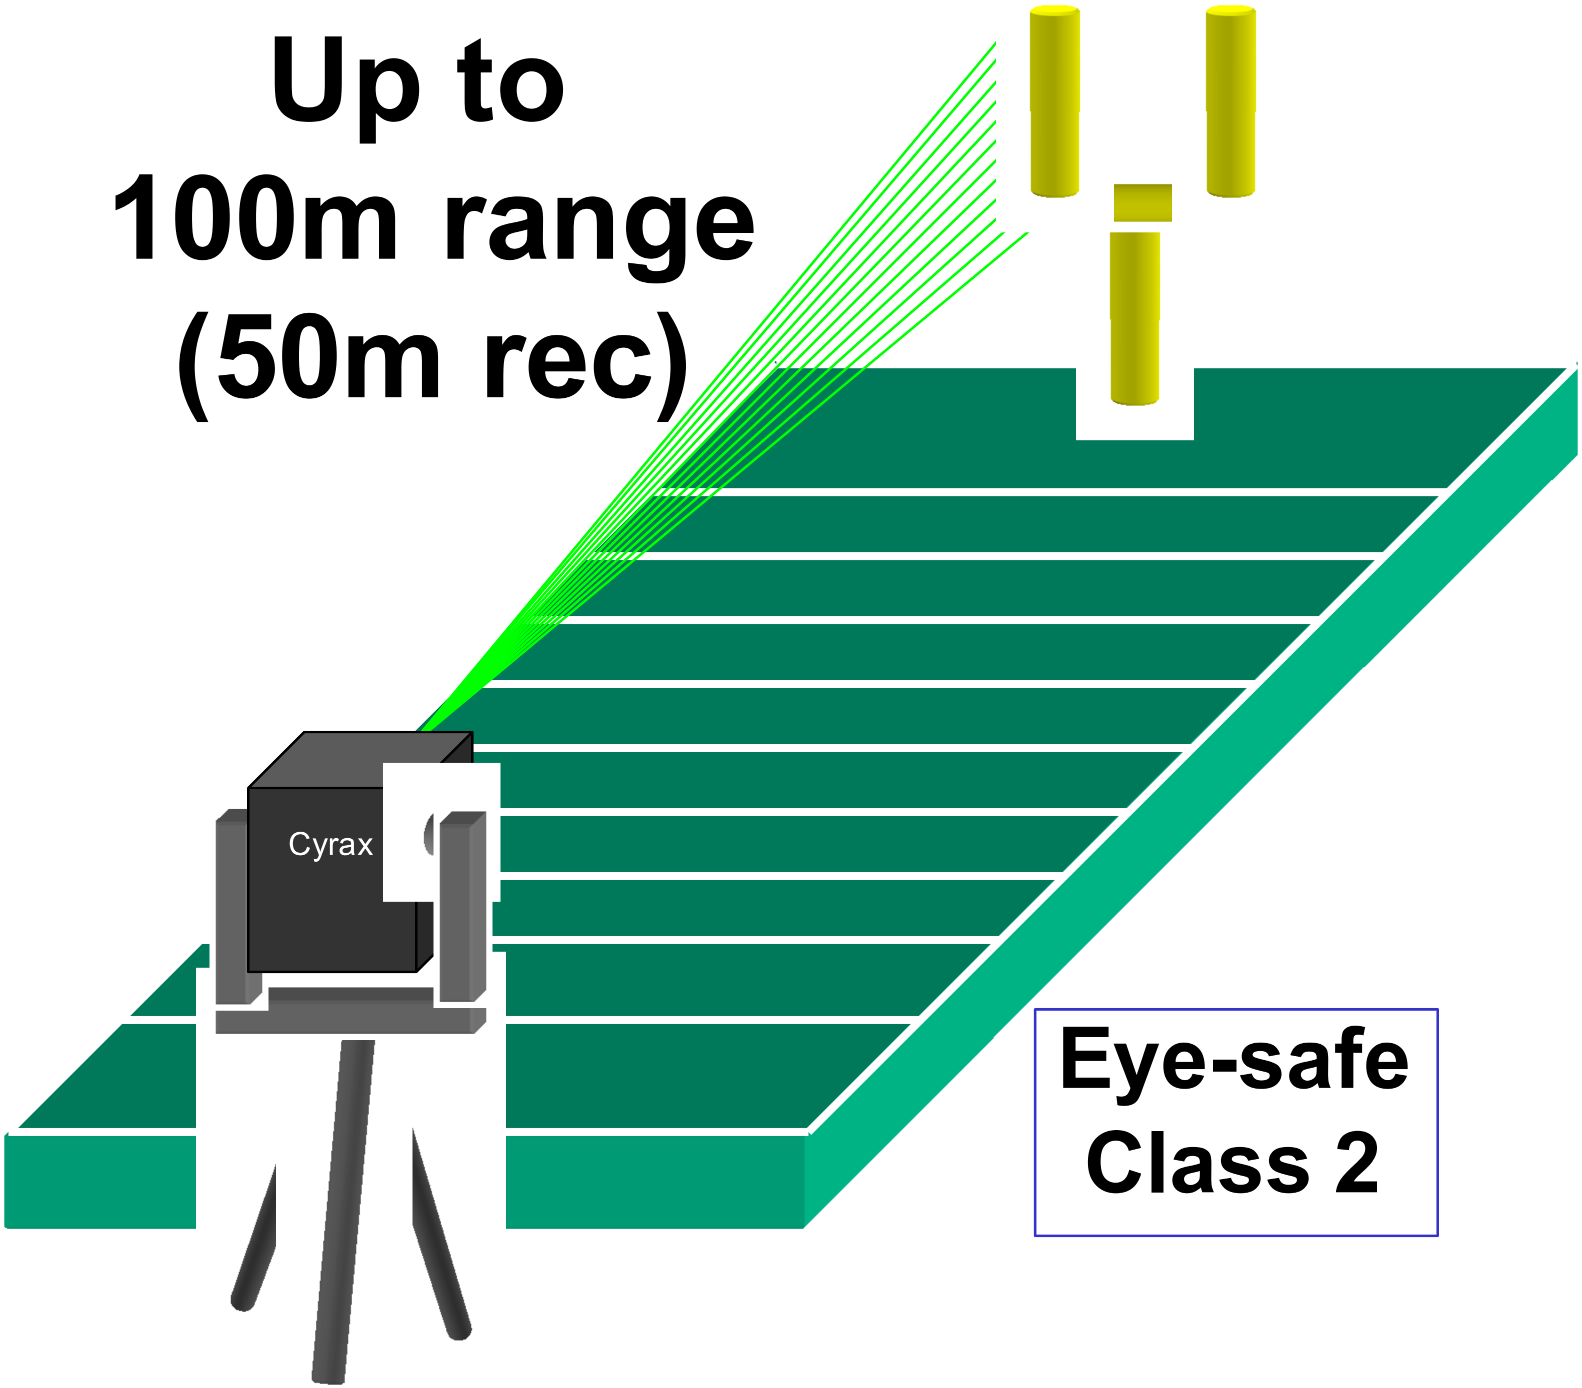

Eye-safe Class 2

Step 1:
Target the structure

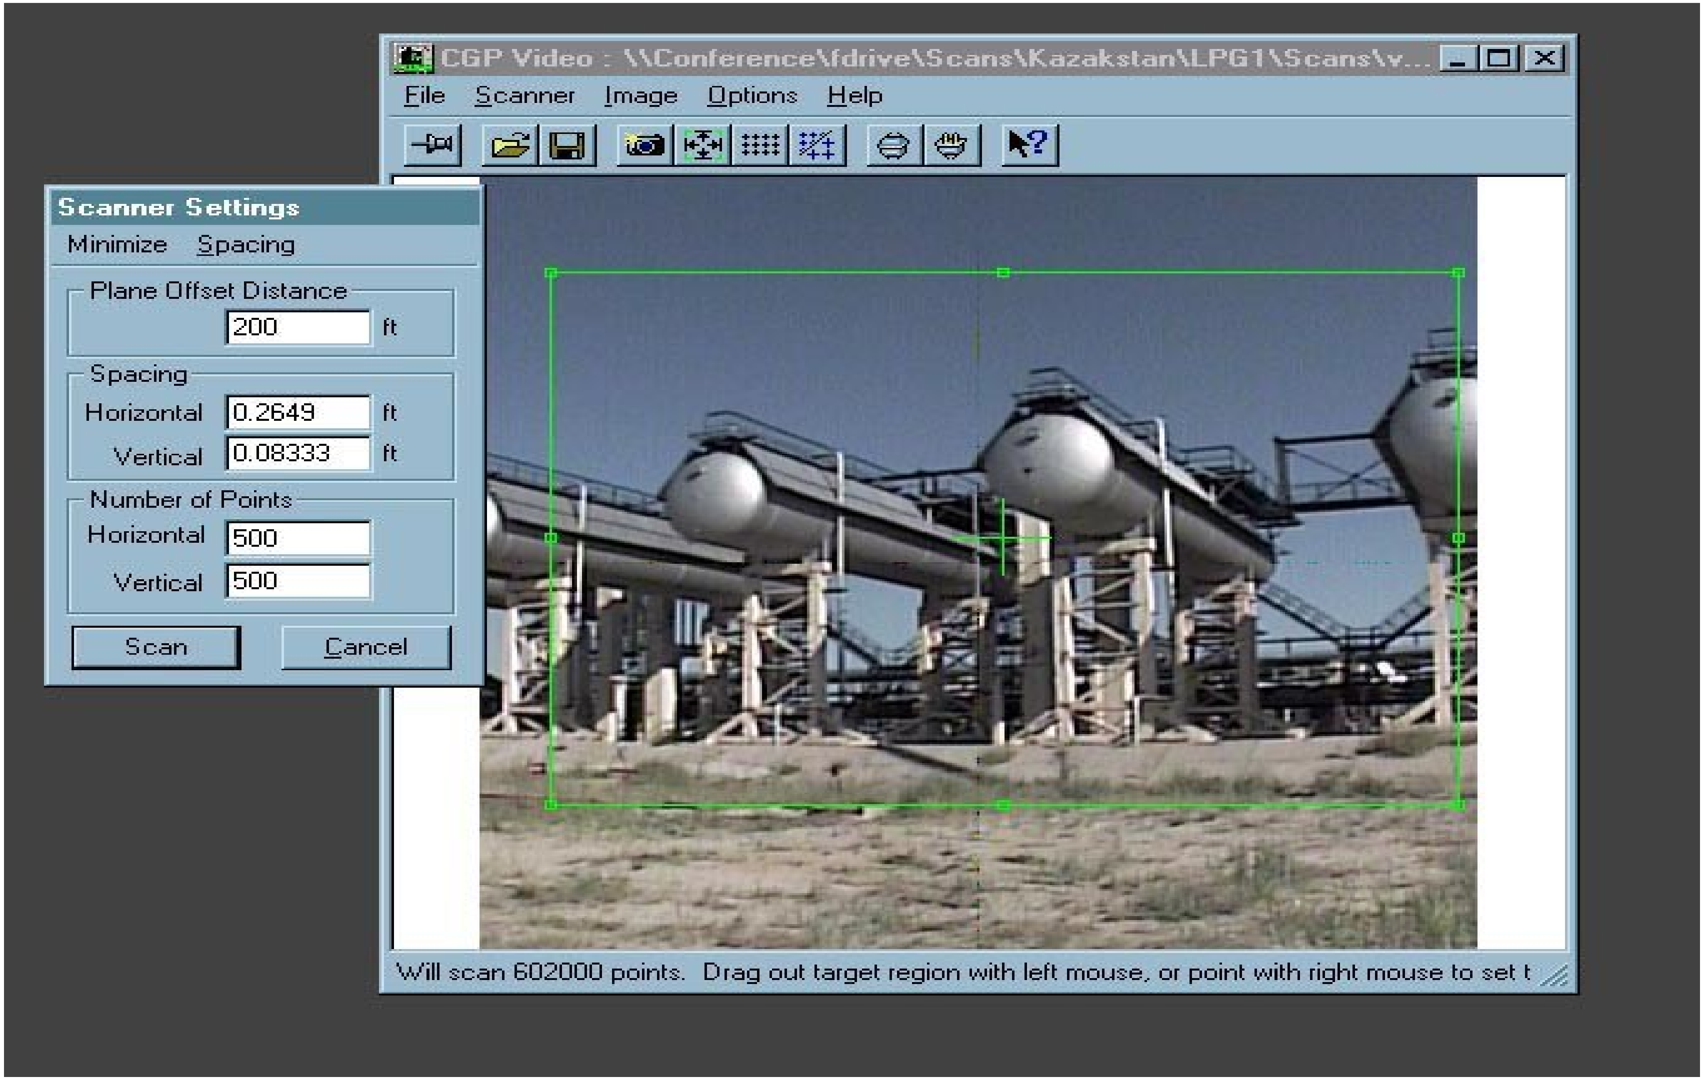

Step 2: Scan the structure

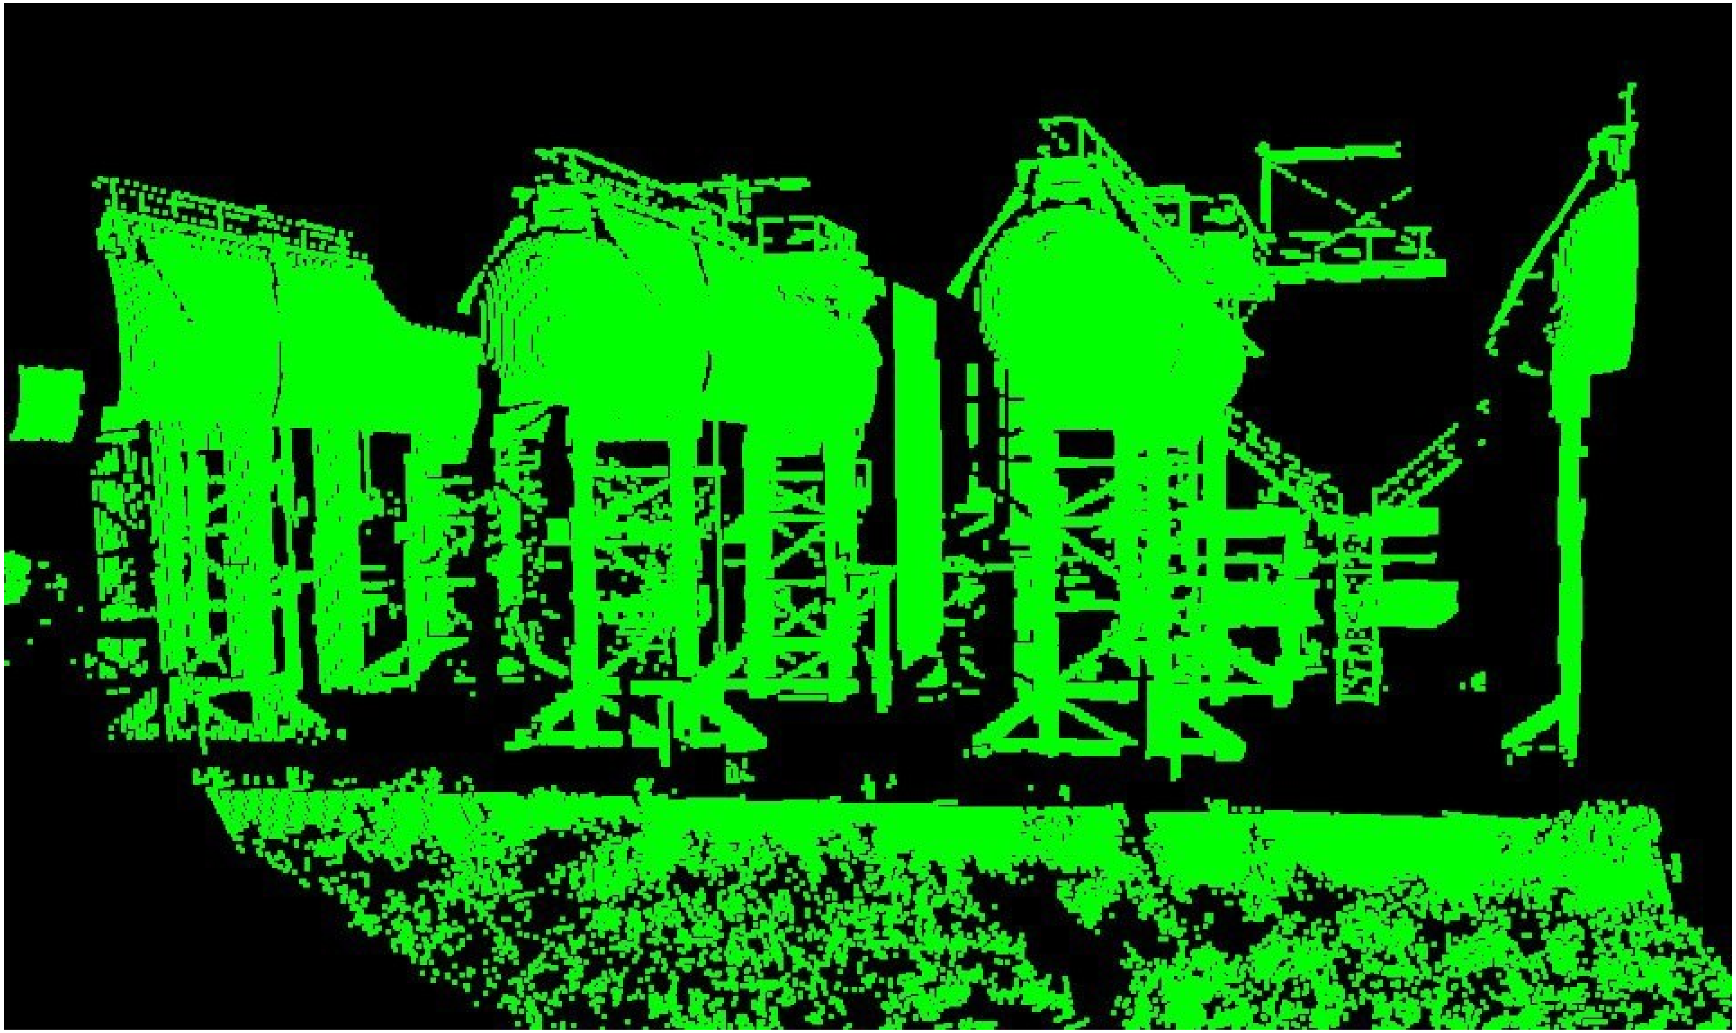

Step 3:
Model fitting in-the-field

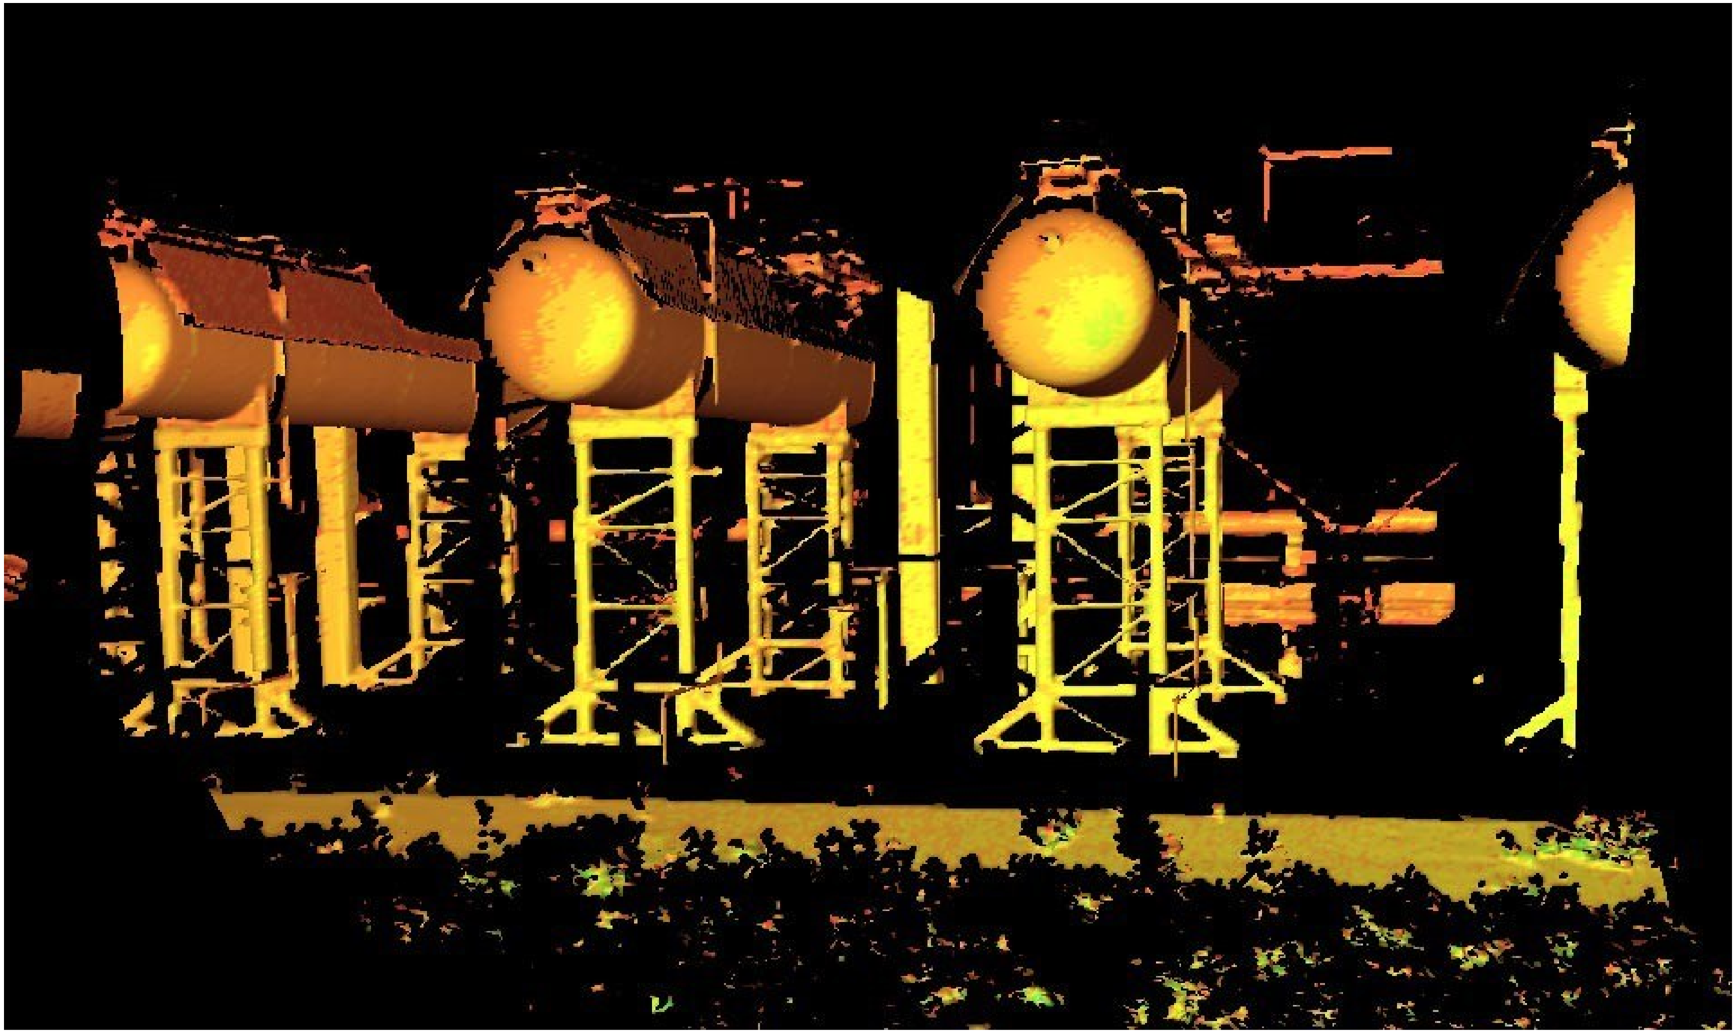

Result

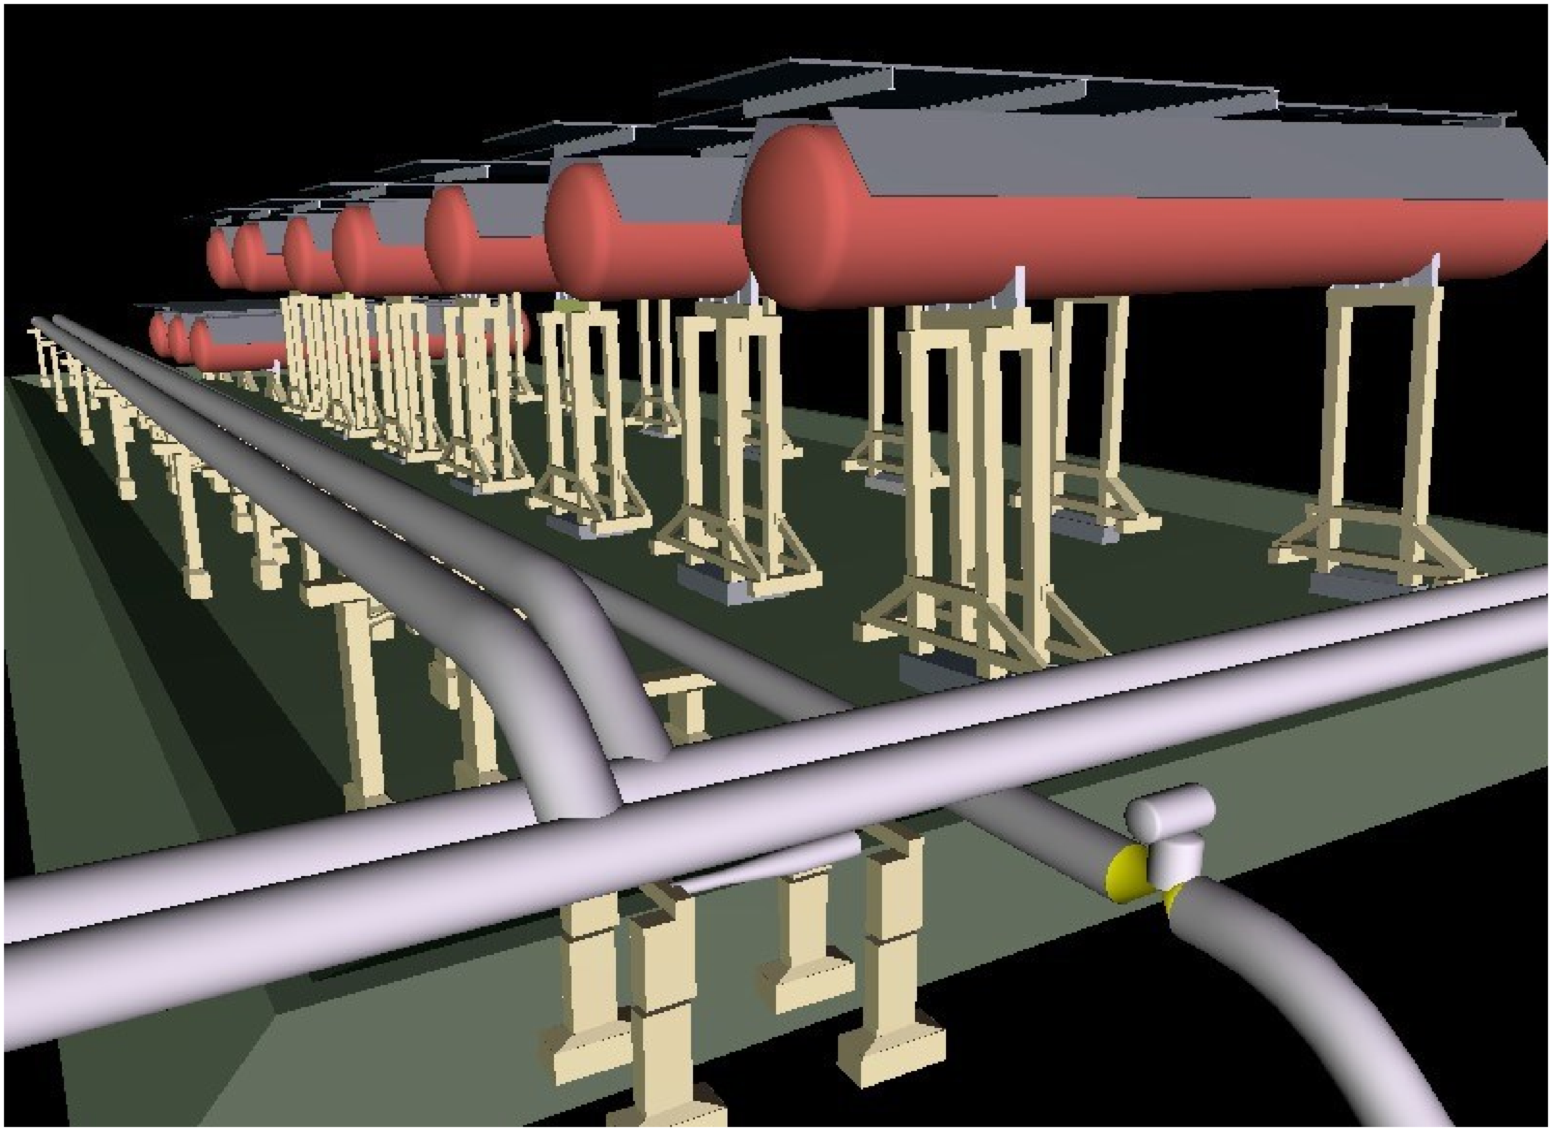

Project: As-built of Chevron hydrocarbon plant

- 400’x500’ area
- 10 vessels; 5 pumps
- 6,000 objects
- 81 scans from 30 tripod locations
- *Cyrax* field time = 50 hrs

Cost Benefits

Added Value Benefits

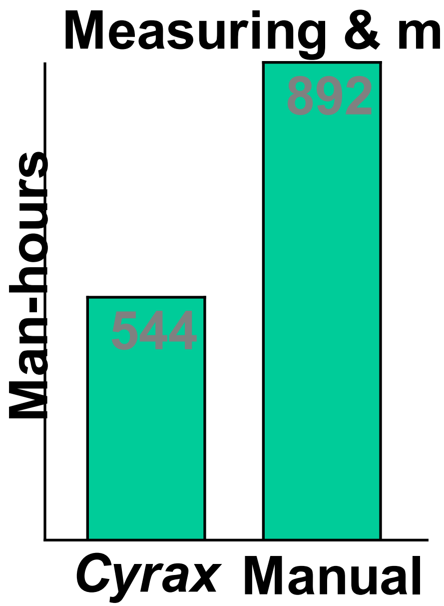

- Greater detail & no errors
- Higher accuracy
- Fewer construction errors
- 6 week schedule savings

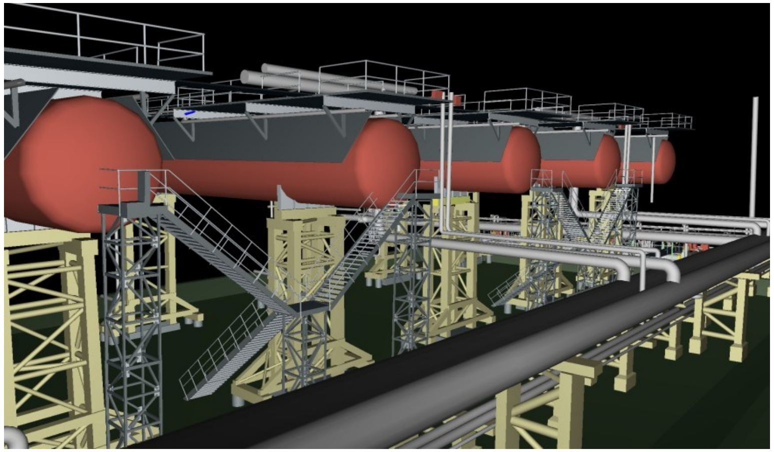

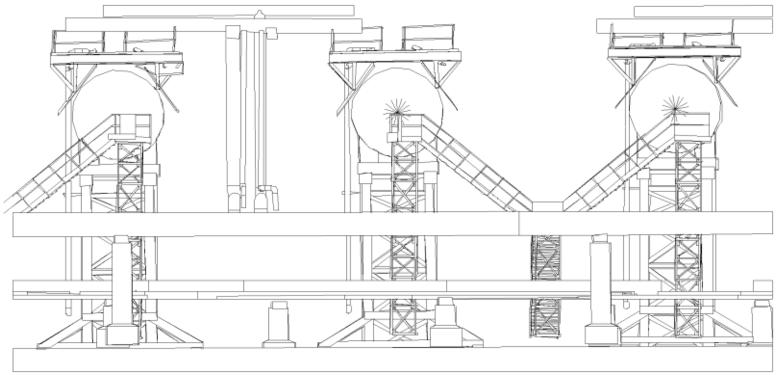

Application
Modeling movie sets

STARSHIP TROOPERS

Image courtesy of Tippett Studio

Lidar data with Riegl LMS-Z390i

courtesy of RWTH Aachen, L. Kobbelt et al.

Comparison Lidar - passive

Image courtesy of Tippett Studio

3D acquisition taxonomy

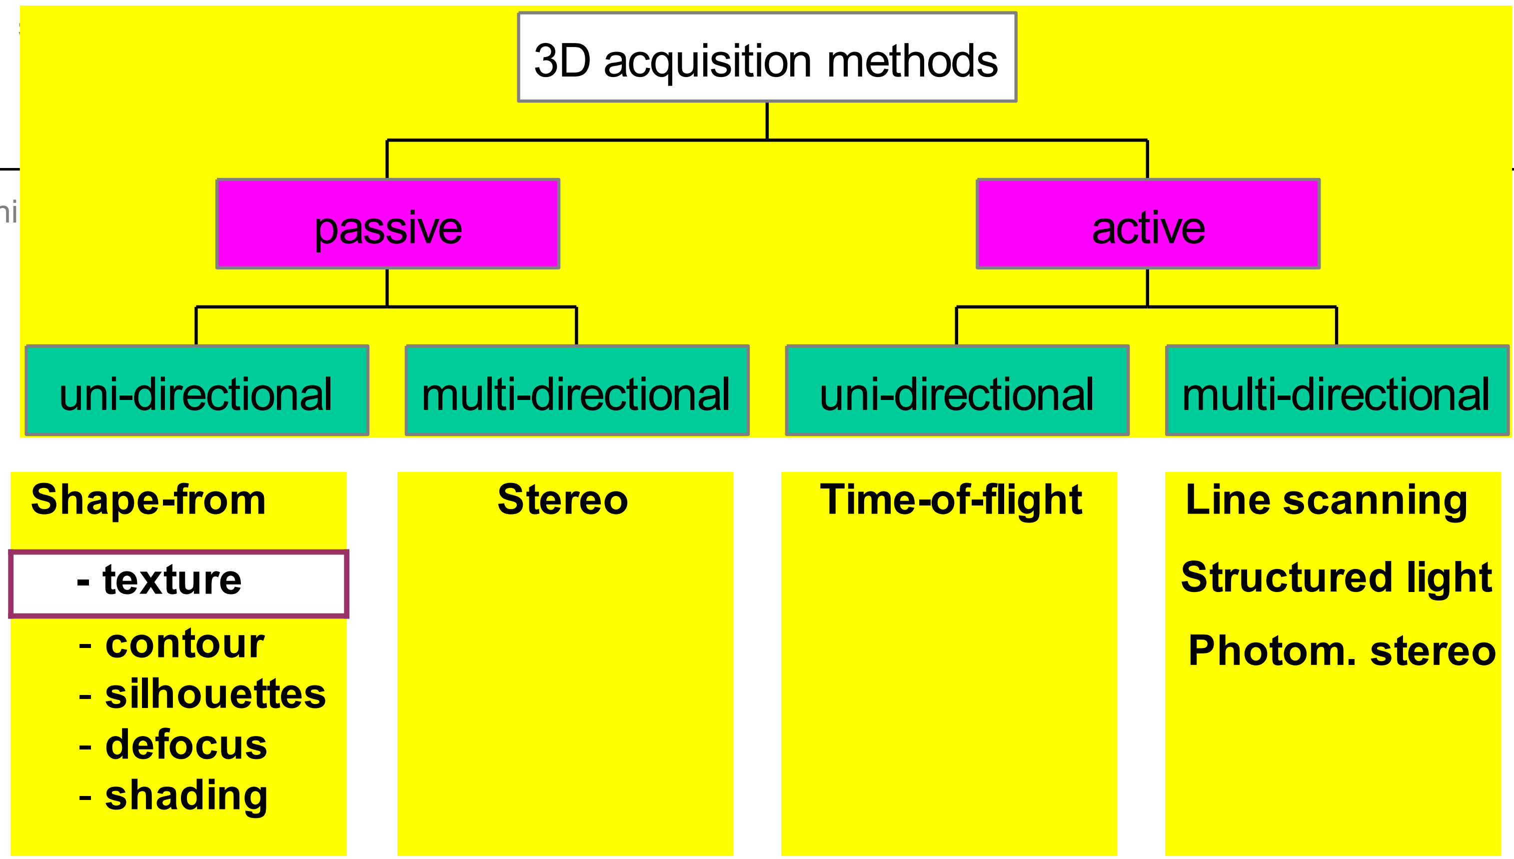

Shape-from-texture

assumes a slanted and tilted surface to have a homogeneous texture

inhomogeneity is regarded as the result of projection

e.g. anisotropy in the statistics of edge orientations

⇓

orientations deprojecting to maximally isotropic texture

è

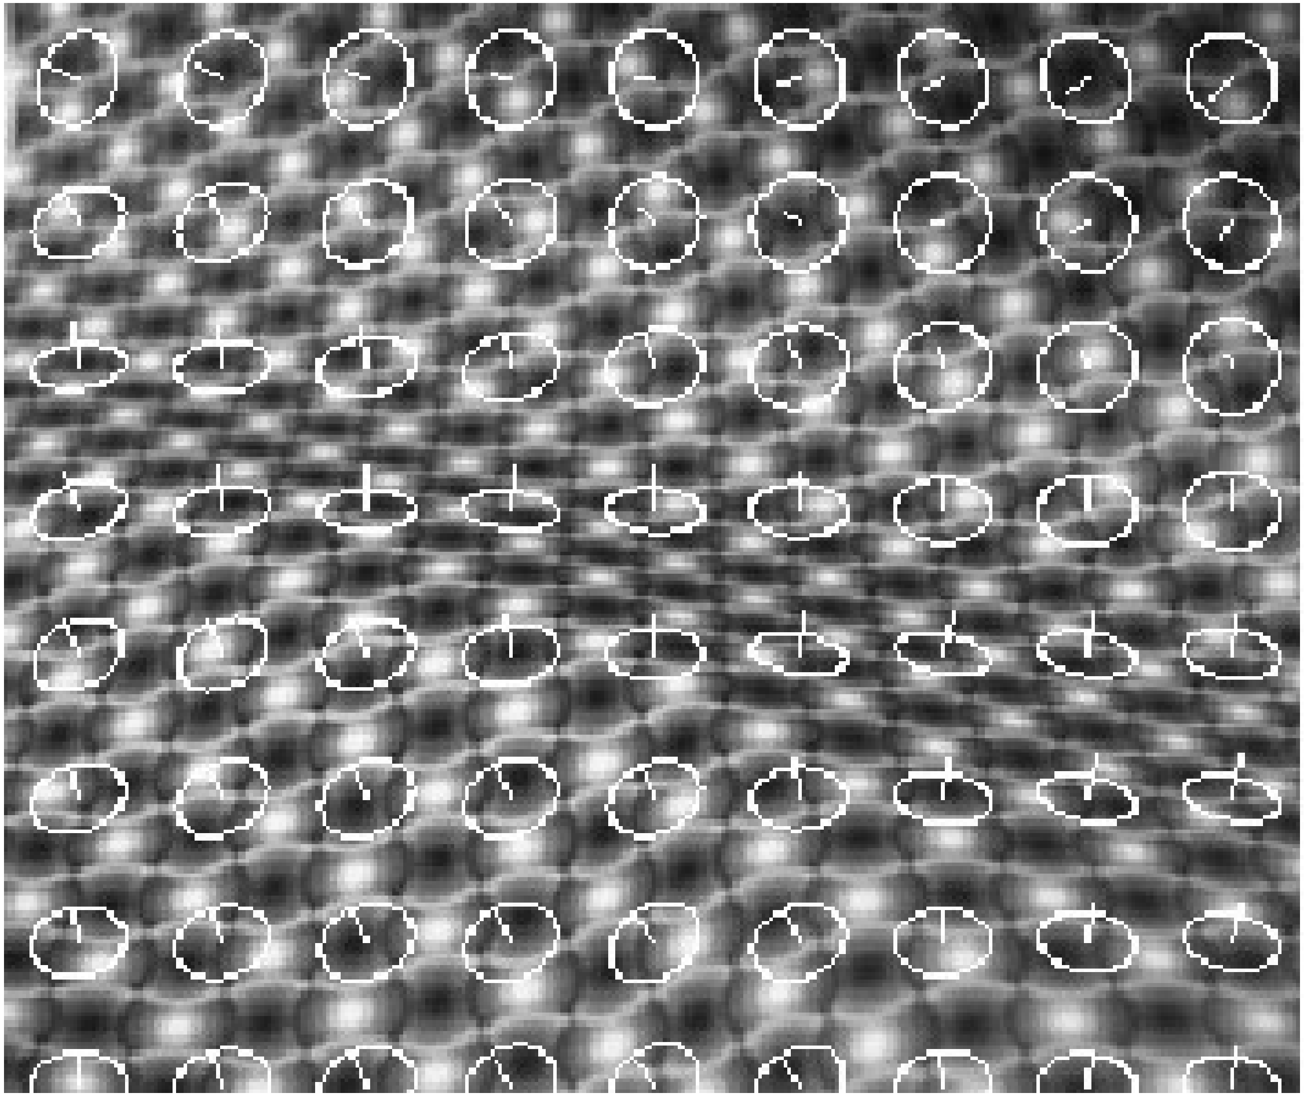

è

3D acquisition taxonomy

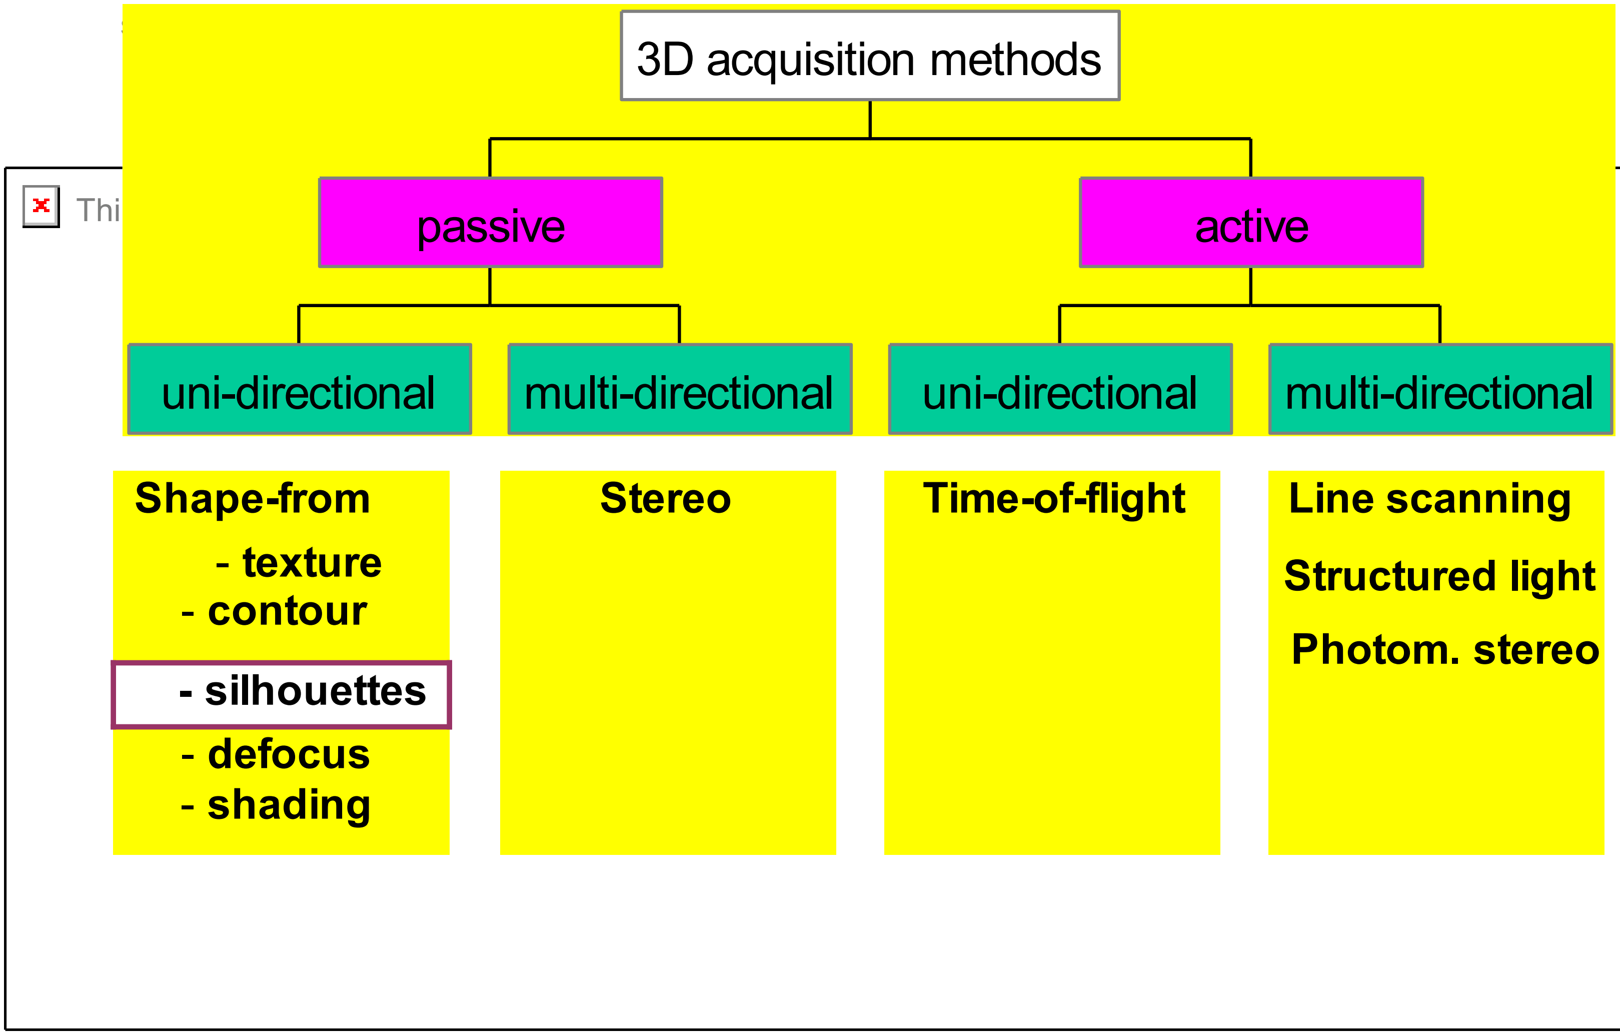

Shape-from-silhouettes

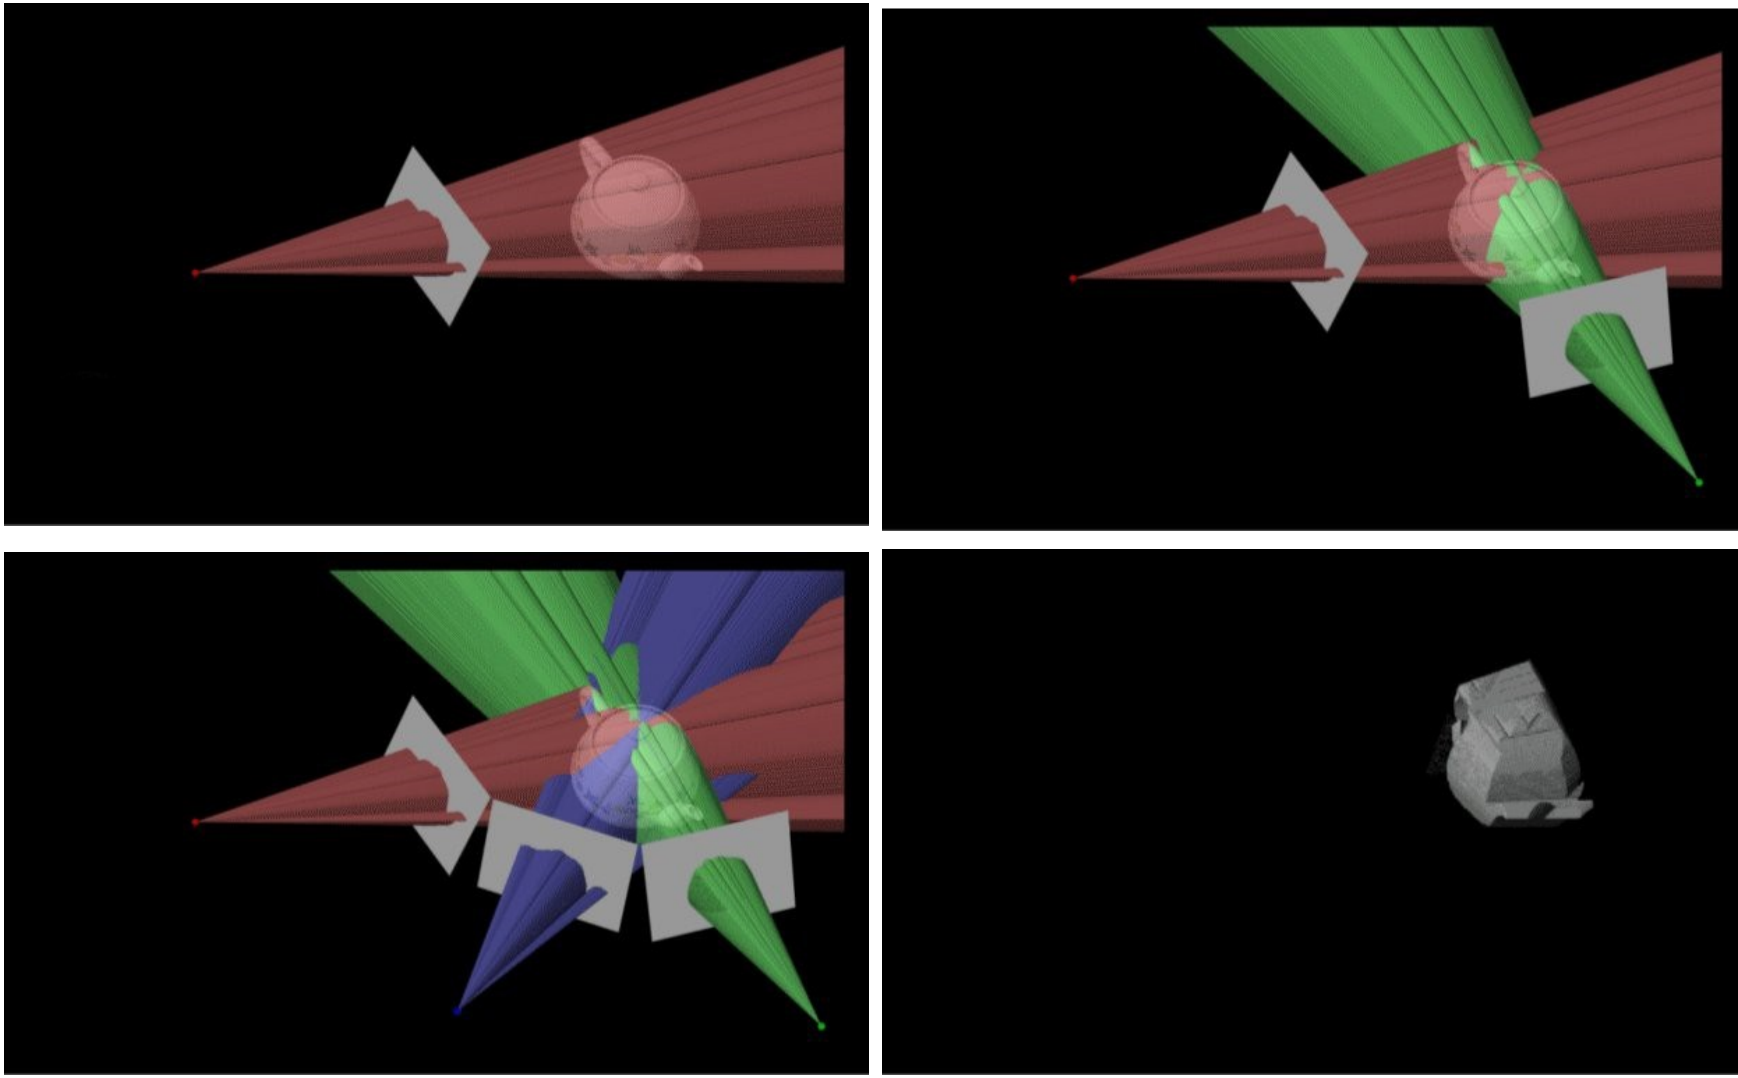

Shape from silhouettes - uncalibrated

tracking of turntable rotation
- **volumetric modeling from silhouettes**
- **triangular textured surface mesh**

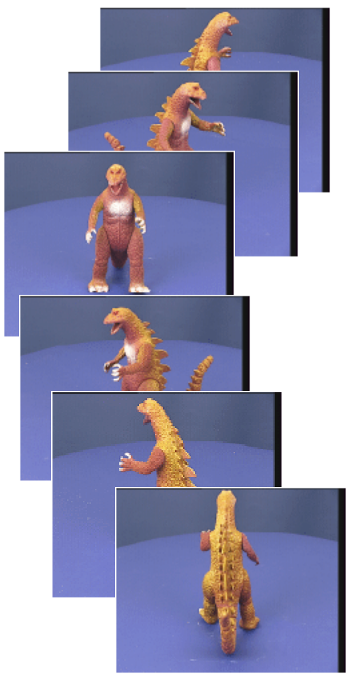

Turntable sequence

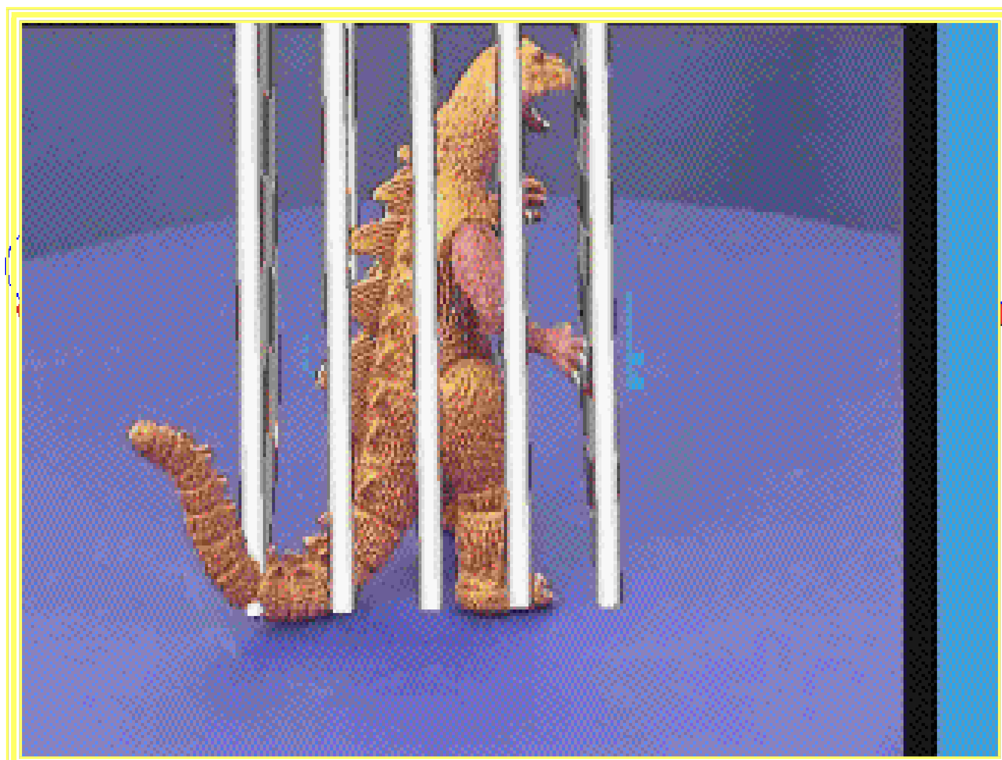

Camera tracking

VRML model

3D acquisition taxonomy

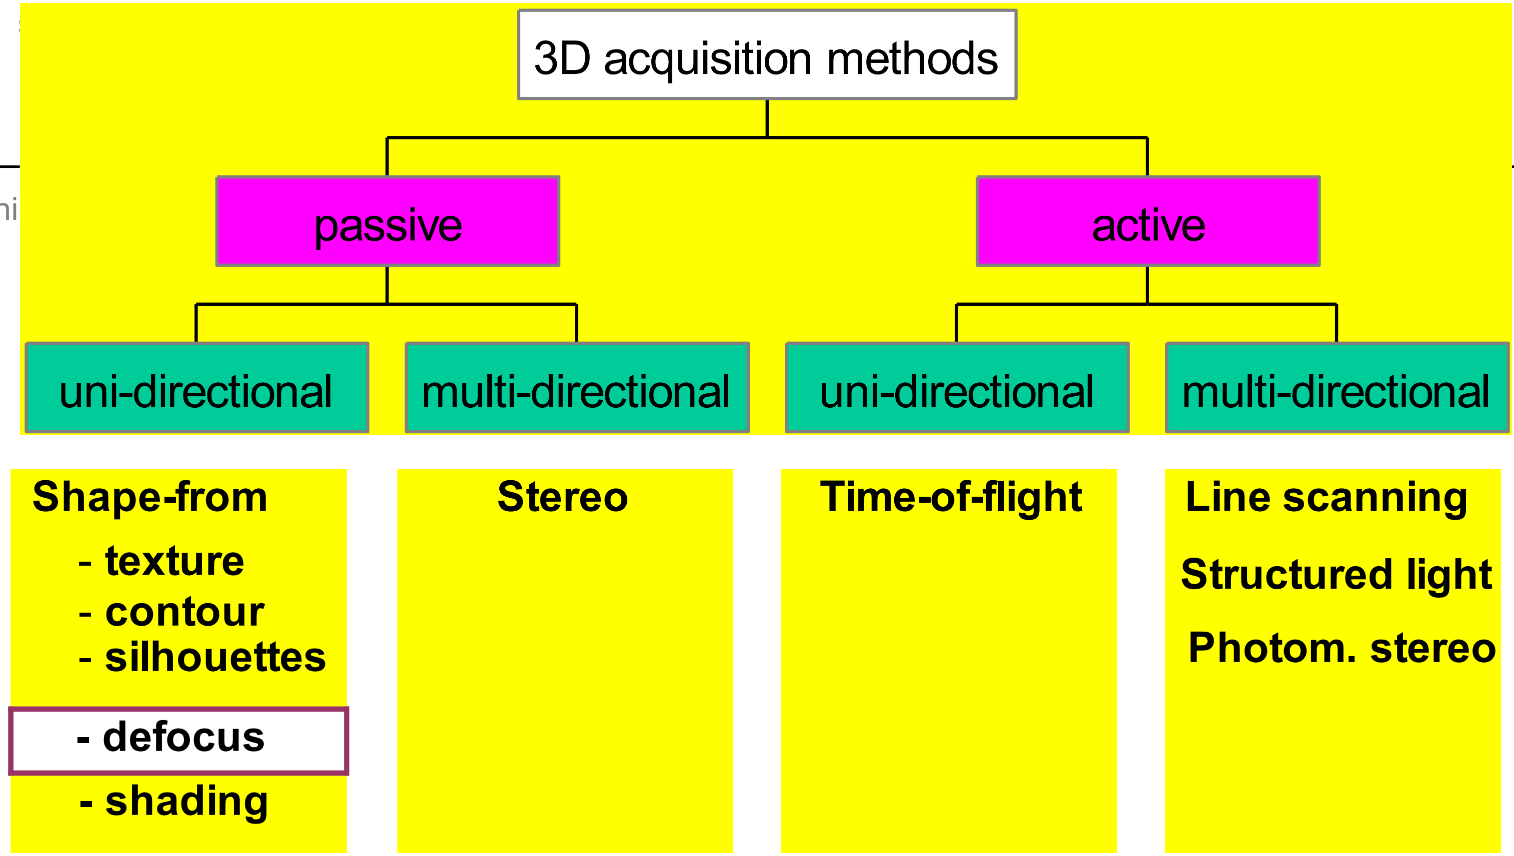

3D acquisition taxonomy

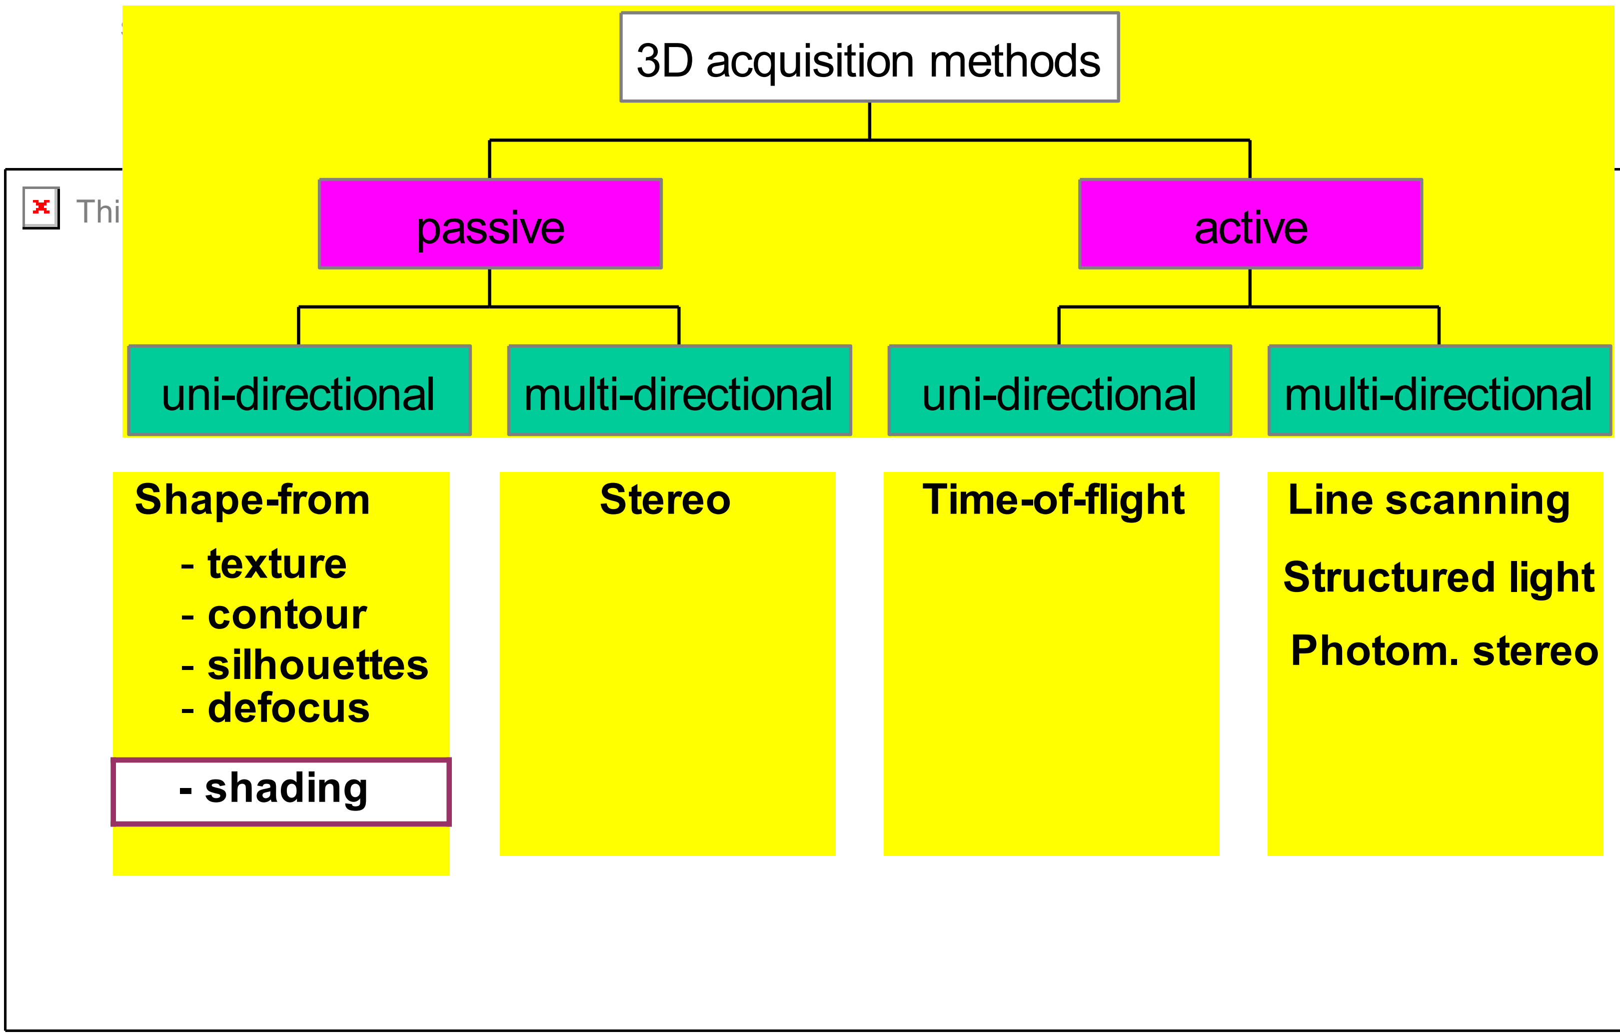

(a) (b) (c)

Fig. 1. Three images of the same scene taken with different camera settings. (a) The apple is brought into focus. (b) The grapefruit is brought int focus. (c) The background is brought into focus. When the background is brought into focus, both the grapefruit and the apple are defocused. I addition, the apple is more blurred than the grapefruit since it is farther from the background than the grapefruit.

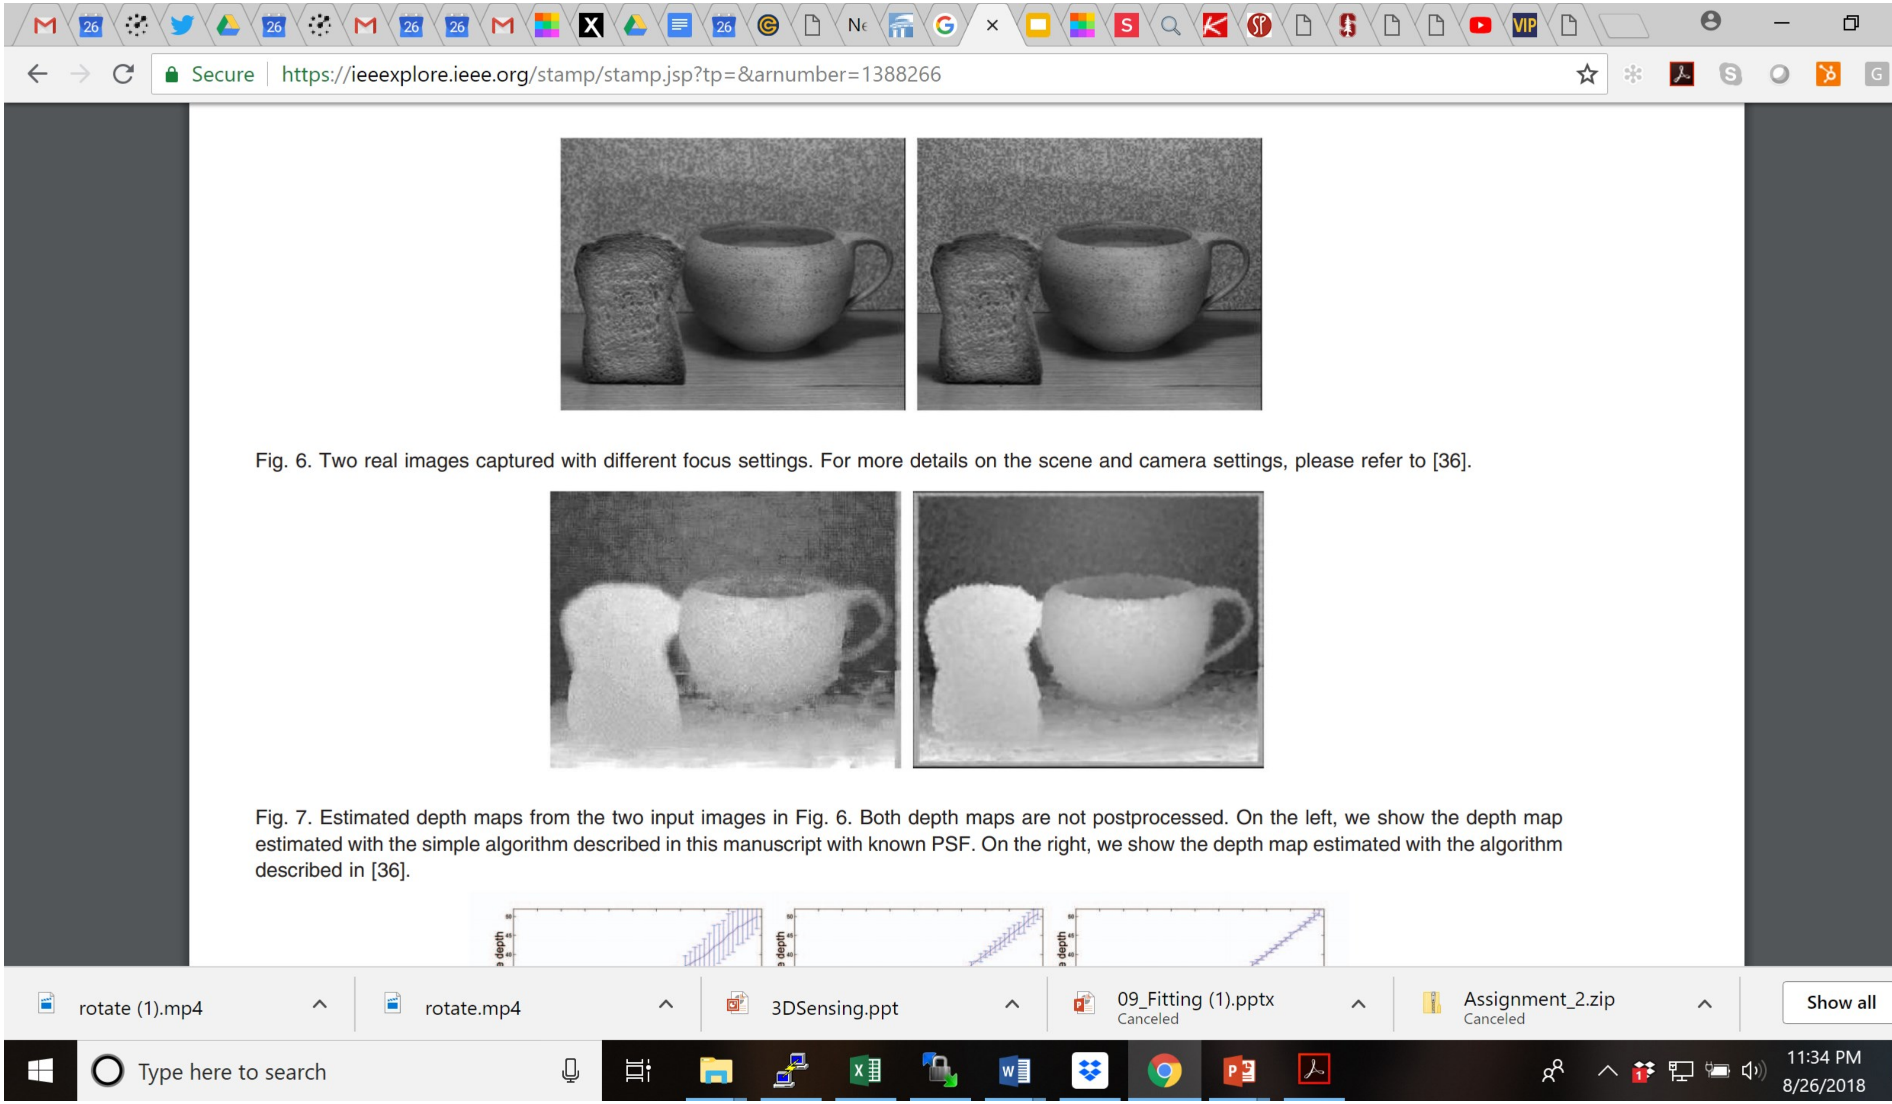

Fig. 6. Two real images captured with different focus settings. For more details on the scene and camera settings, please refer to [36].

Fig. 7. Estimated depth maps from the two input images in Fig. 6. Both depth maps are not postprocessed. On the left, we show the depth map estimated with the simple algorithm described in this manuscript with known PSF. On the right, we show the depth map estimated with the algorithm described in [36].

PAMI, Stephen Soatto , UCLA, 2005

Shape-from-shading

Uses directional lighting, often with known direction

local intensity is brought into correspondence with orientation via *reflectance maps*

orientation of an isolated patch cannot be derived uniquely

extra assumptions on surface smoothness and known normals at the rim

è

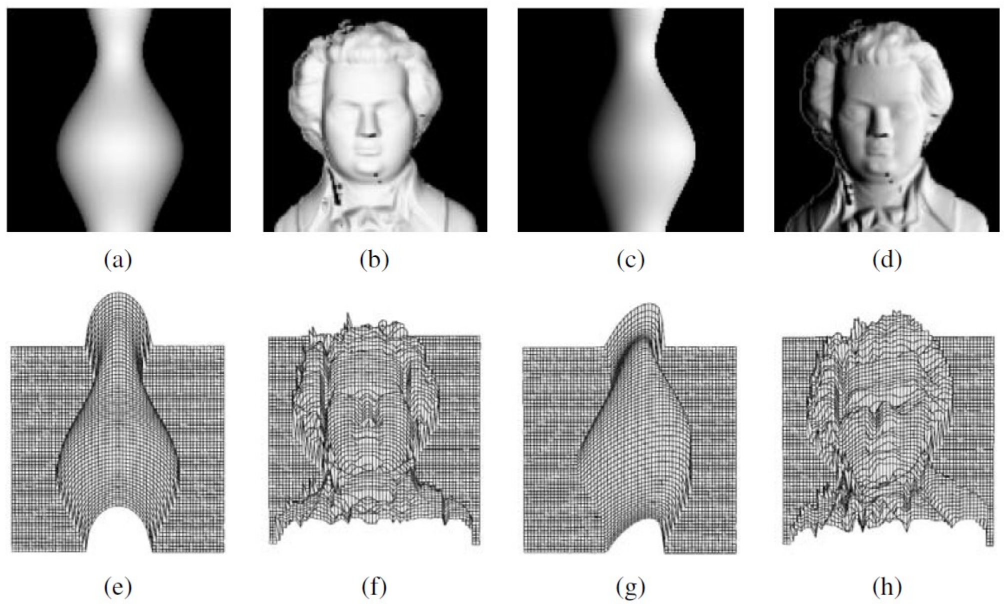

Figure 12.2 Synthetic shape from shading (Zhang, Tsai, Cryer *et al.* 1999) © 1999 IEEE: shaded images, (a–b) with light from in front $(0, 0, 1)$ and (c–d) with light the front right $(1, 0, 1)$; (e–f) corresponding shape from shading reconstructions using the technique of Tsai and Shah (1994).

3D acquisition taxonomy

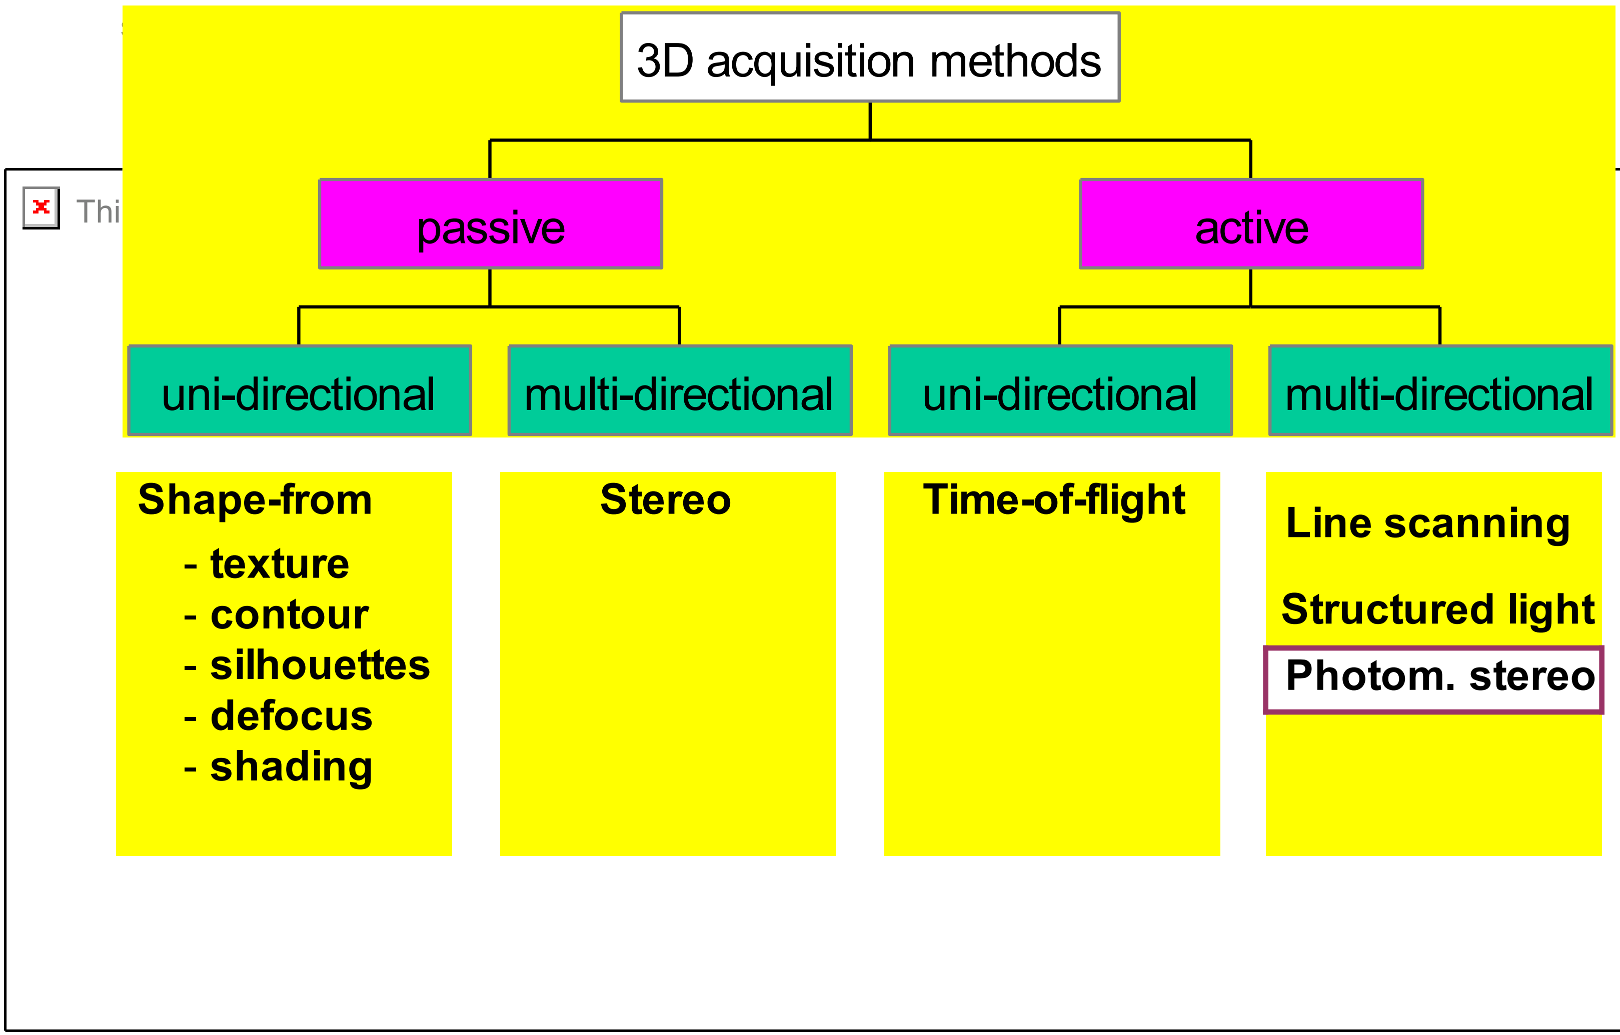

Photometric stereo

Photometric stereo is a technique in computer vision for estimating the surface normals of objects by observing that object under different lighting conditions. It is based on the fact that the amount of light reflected by a surface is dependent on the orientation of the surface in relation to the light source and the observer.[1] By measuring the amount of light reflected into a camera, the space of possible surface orientations is limited. Given enough light sources from different angles, the surface orientation may be constrained to a single orientation or even overconstrained.

è

Mini-dome for photometric stereo

Instead of working with multi-directional light applied simultaneously with the colour trick, one can also project from many directions in sequence…

.wikipedia.org/wiki/Photometric_stereo#/media/File:Photometric_stereo.png

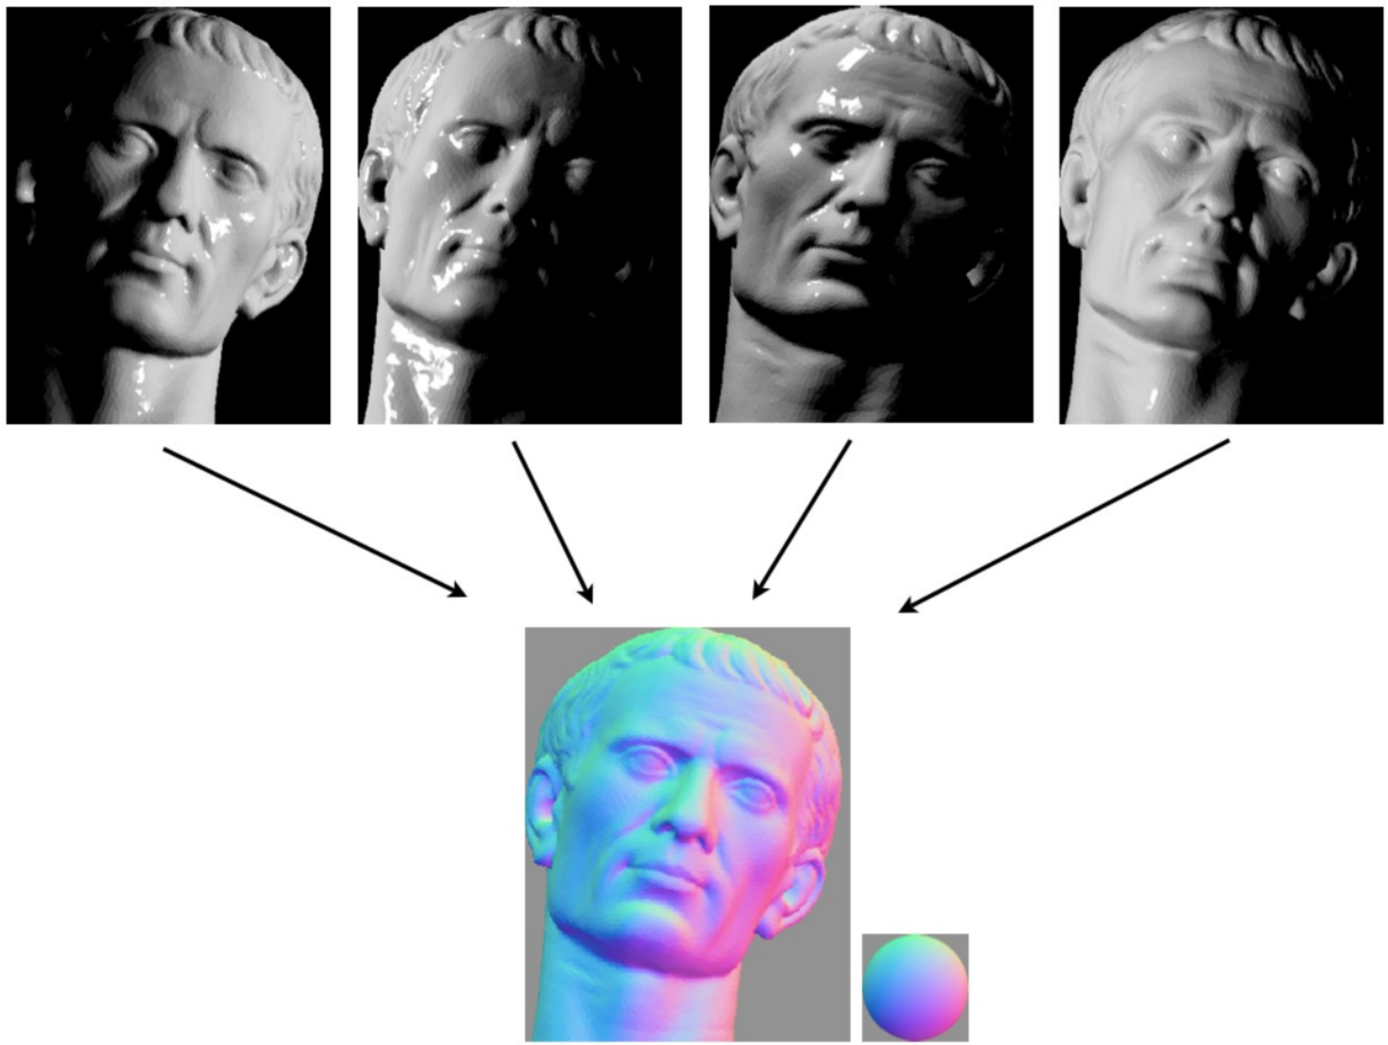

ultiple images of an object under different lighting conditions to estimate a normal direction at each pixel.

Strongest 3D cues for us are 2D...