TU WIEN

CVL Computer Vision Lab

3D Reconstruction

Simon Brenner

Digitale Transformation in der Handschriftenforschung – Computer Vision

Lecture Outline

- Introduction
  - Why 3D reconstruction?
  - Basic concepts
- Reconstruction Approaches
  - Active triangulation
  - Range scanners
  - Photogrammetry (with how-to)
  - Photometric Stereo
- Application to written heritage
  - Examples from research

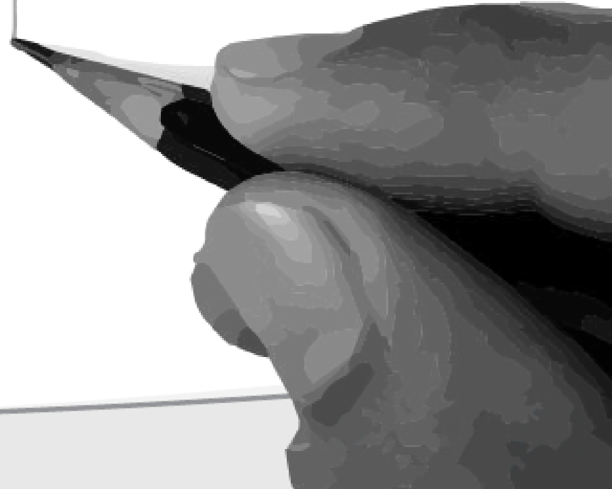

Introduction

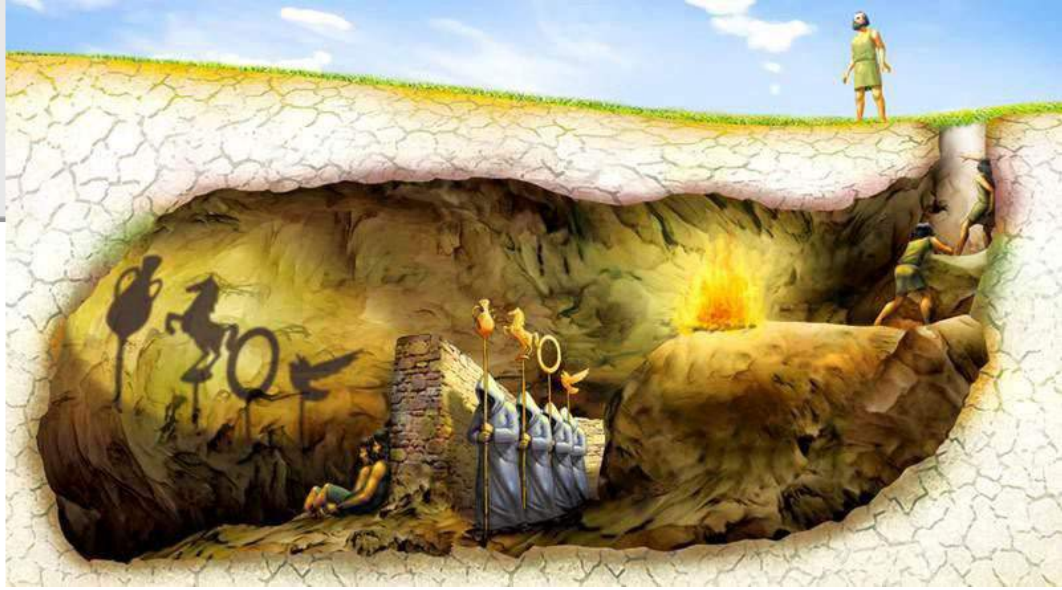

- Photographs are 2D projections of a 3D world
- Result depends on
  - Angle of view
  - Focal length / distance
  - Lighting
- Sometimes we need to know the actual shapes of things
  - Documentation
  - Measurements (lengths, volumes)
  - Simulations
  - Animations / Games

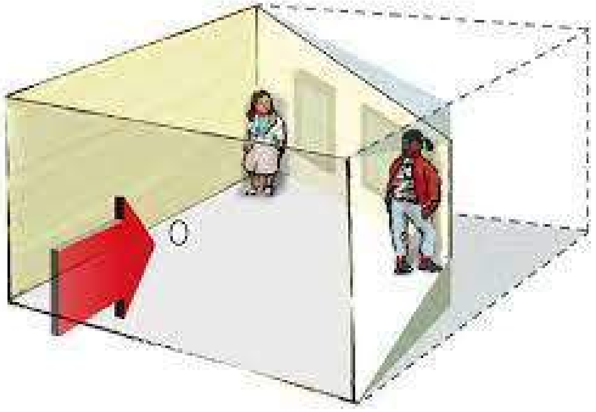

3D written heritage – some examples

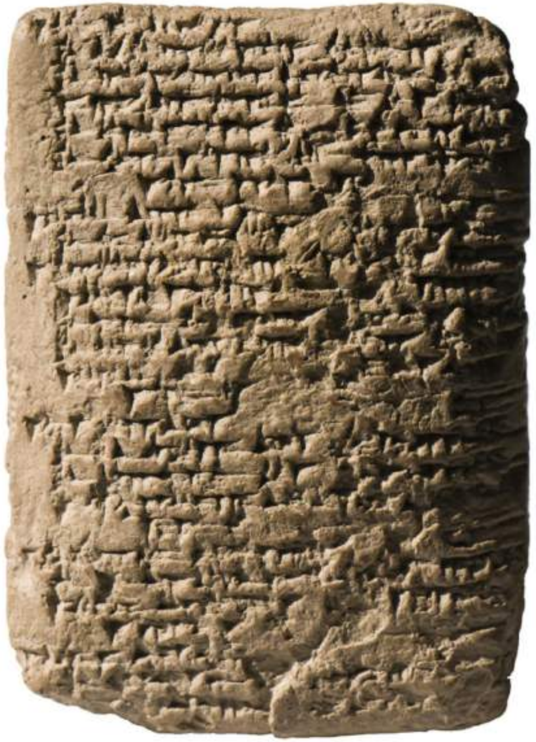

Cuneiform tablets

Etruscan bronze mirrors

Coins

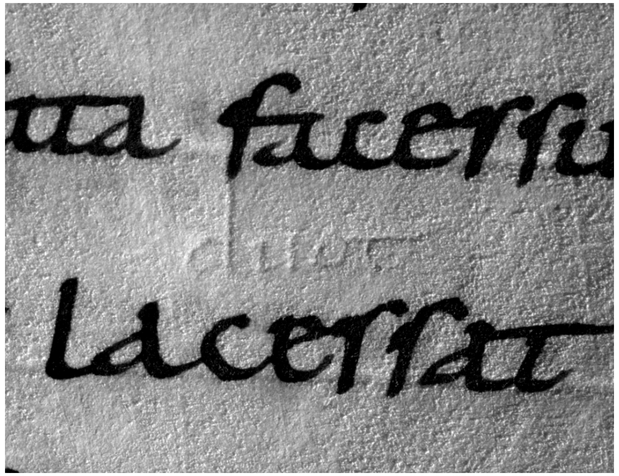

Dry-point glosses

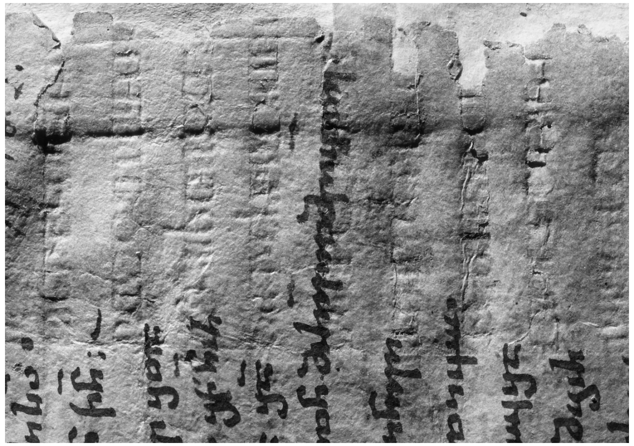

Ink corrosion

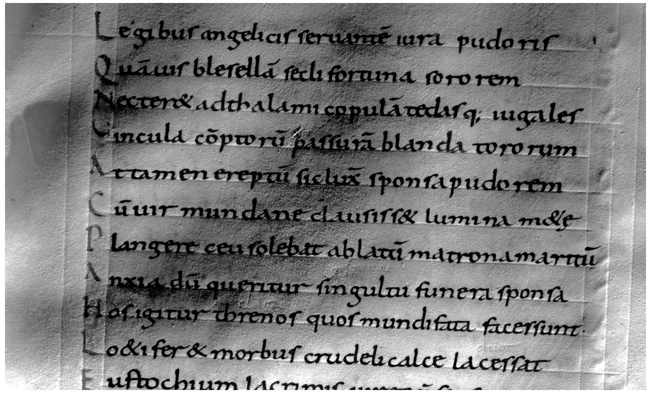

Dry-point ruling

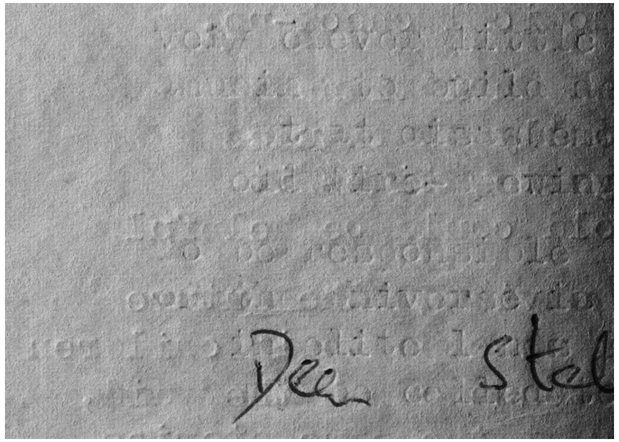

Typewriter indentations

Lecture Outline

- Introduction
  - Why 3D reconstruction?
  - Basic concepts
- Reconstruction Approaches
  - Active triangulation
  - Range scanners
  - Photogrammetry (with how-to)
  - Photometric Stereo
- Application to written heritage
  - Examples from research

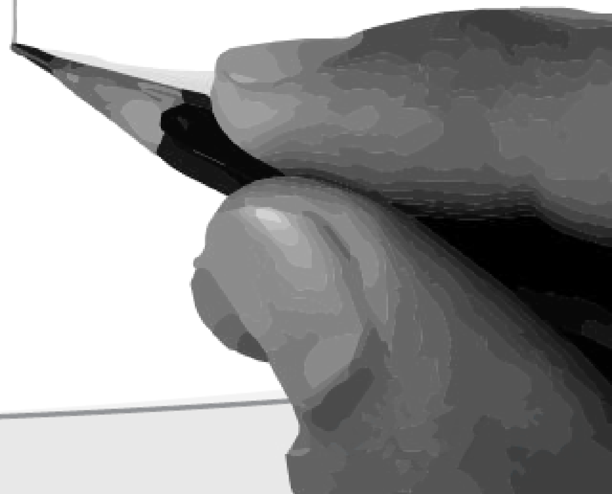

3D Representations

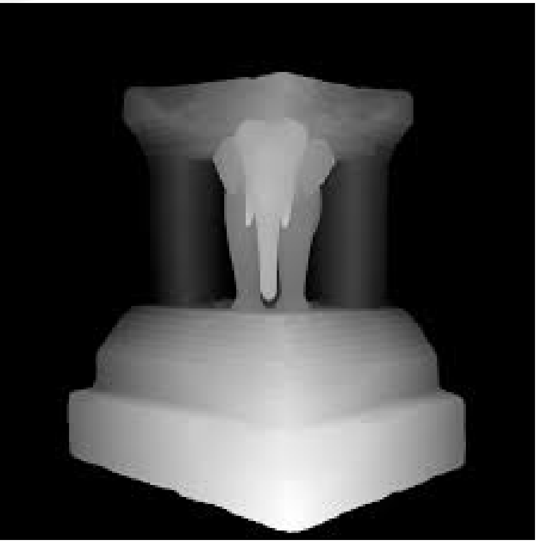

Depth map

- An image, where each pixel represents the distance to the observer
- Only one perspective → "2.5D"

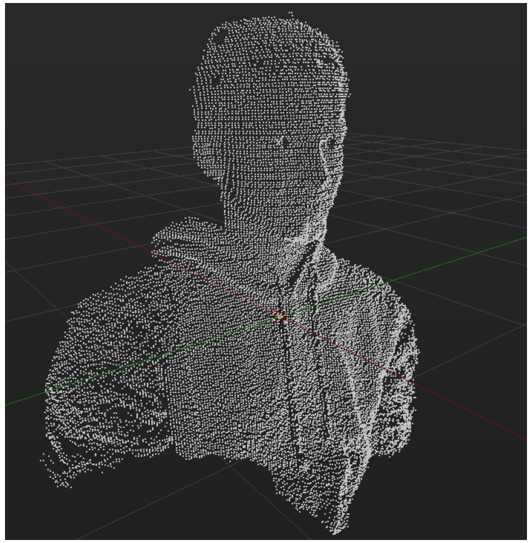

Point cloud

- Collection of points with 3D coordinates (x,y,z)
- No neighborhood information
- Raw 3D scanning data often in this form

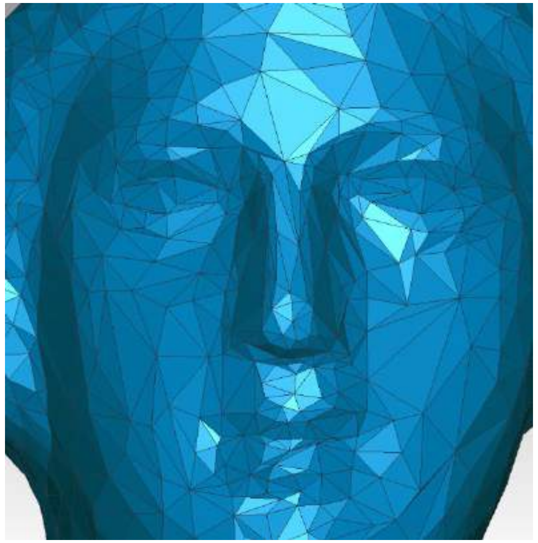

Surface model / mesh

- Vertices (points)
- Edges between vertices
- Faces between edges
- Very useful representation for visualizations, animations, games, etc.

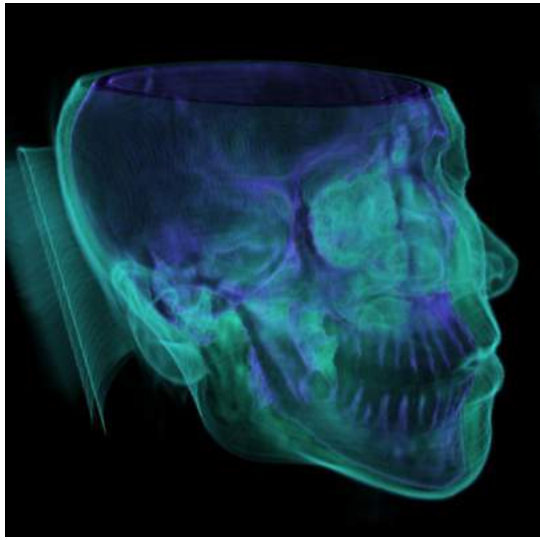

Volumetric data

- 3D space is partitioned into "Voxels", representing e.g. material density
- No explicit surfaces
- Produced by CT, MRI

Surface model / mesh

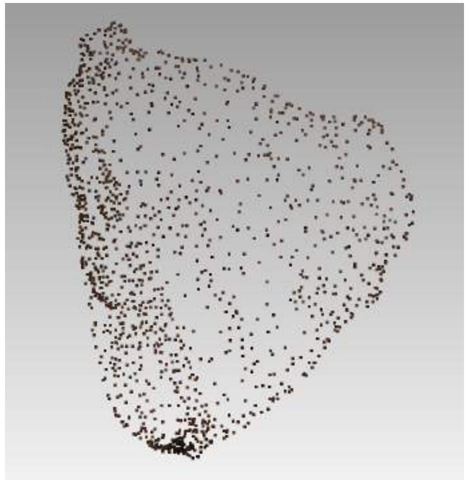

vertices

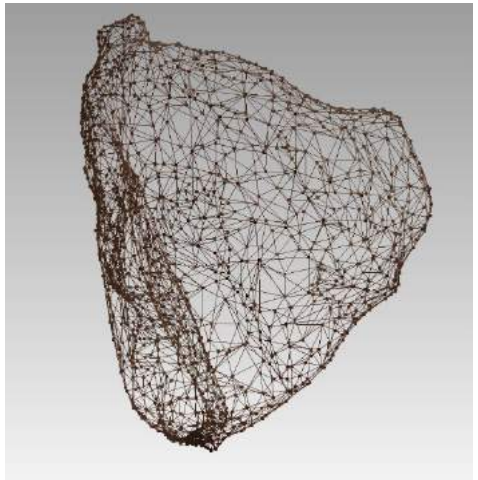

vertices + edges

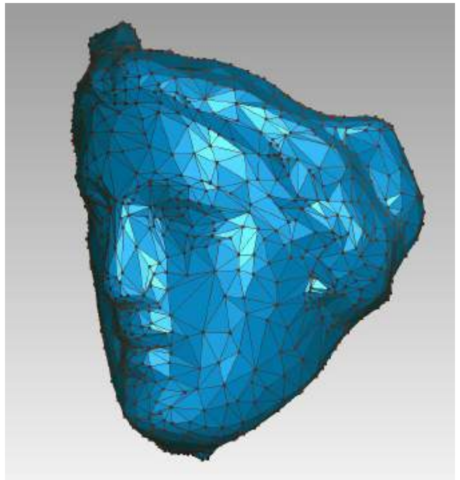

vertices + edges + faces = mesh

Surfaces are oriented (have a front and back side)

Normal vectors indicate the orientation (and are vital for rendering)

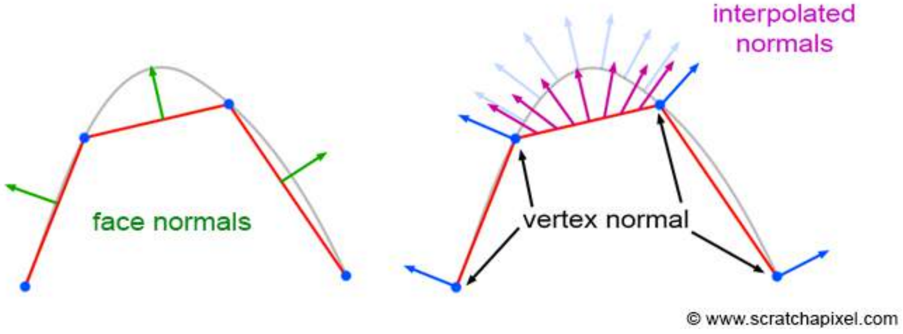

Pinhole camera model

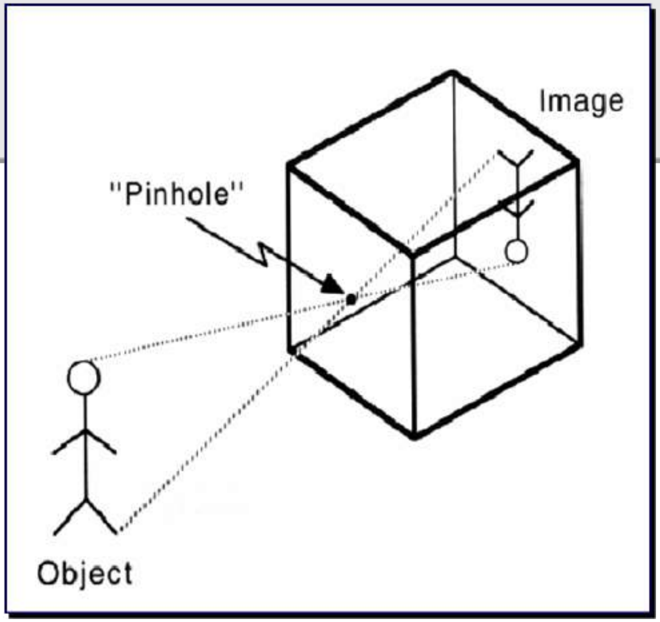

- Simplified image formation model
- Reasonable approximation for many 3D vision applications

$$x = \frac{f}{Z} X$$

$$y = \frac{f}{Z} Y$$

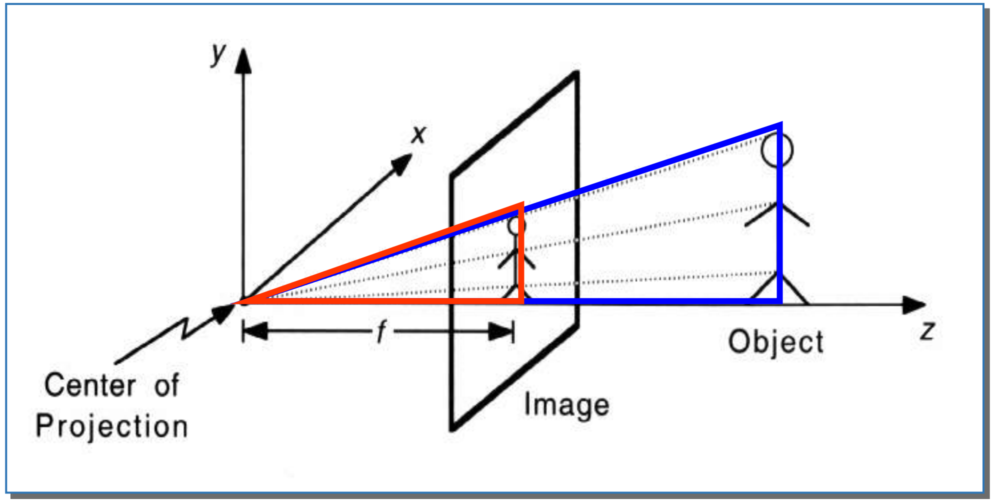

$$\frac{x}{X} = \frac{f}{Z}$$

$$\frac{y}{Y} = \frac{f}{Z}$$

(x, y)… image coordinates

(X, Y, Z)… object coordinates

Pinhole camera model

- Intrinsic Parameters
  - Focal length (f_x, f_y)
  - Principal point (p_x, p_y)
    - Intersection of principal axis with sensor

- Extrinsic Parameters
  - Rotation
  - Translation (position)

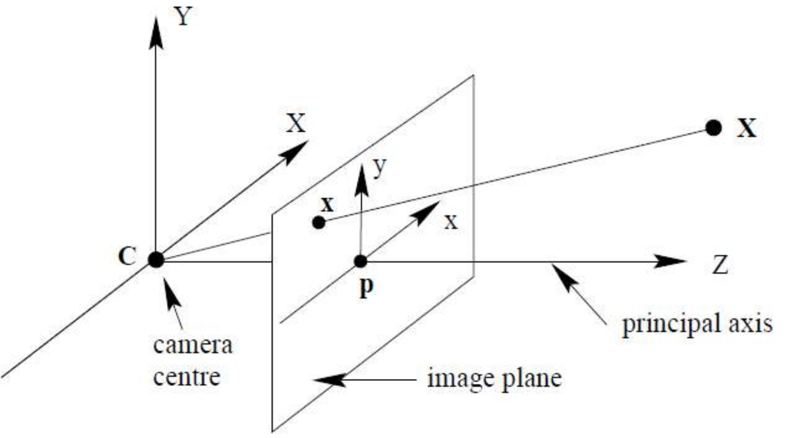

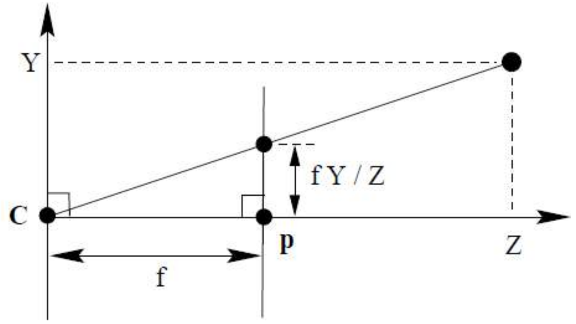

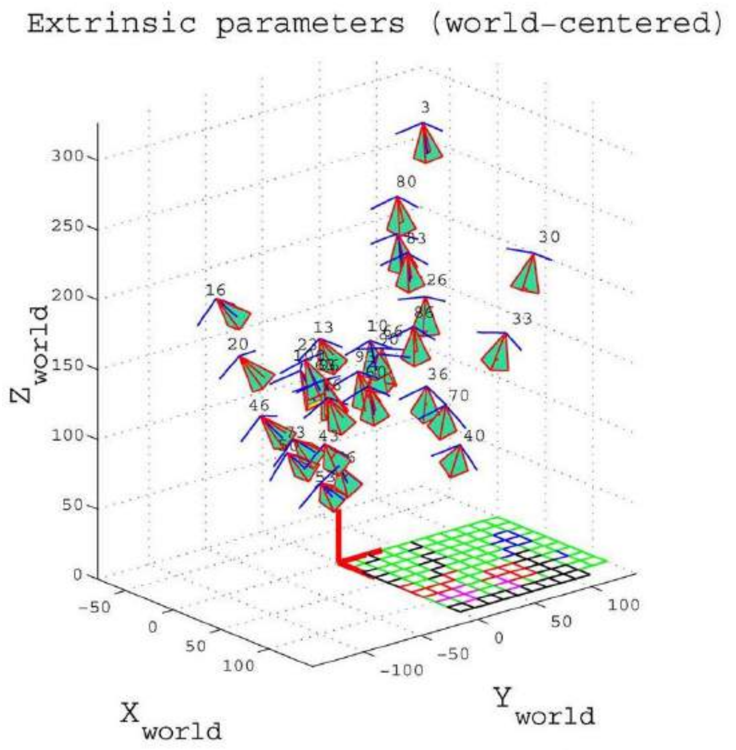

Lecture Outline

- Introduction
  - Why 3D reconstruction?
  - Basic concepts
- Reconstruction Approaches
  - Active triangulation
  - Range scanners
  - Photogrammetry (with how-to)
  - Photometric Stereo
- Application to written heritage
  - Examples from research

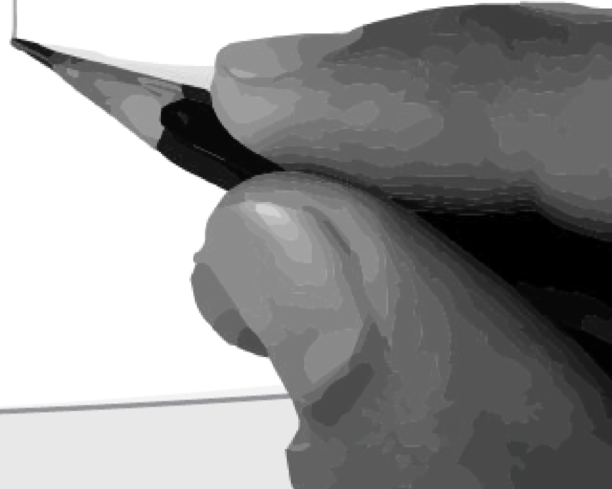

Active triangulation - basics

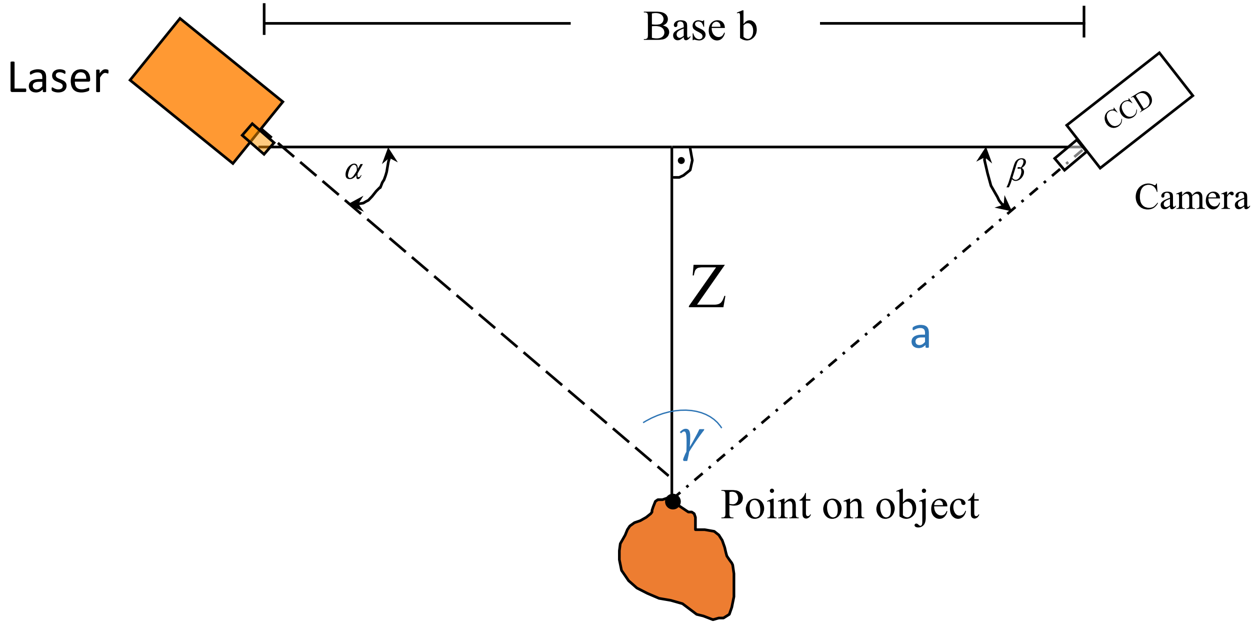

$$\frac{Z}{\sin\beta} = a$$

$$\frac{a}{\sin\alpha} = \frac{b}{\sin\gamma} \Rightarrow a = \frac{b\sin\alpha}{\sin\gamma}$$

$$\frac{Z}{\sin\beta} = \frac{b\sin\alpha}{\sin\lambda} \Rightarrow Z = \frac{b\sin\alpha\sin\beta}{\sin\lambda}$$

$$Z = \frac{b\sin\alpha\sin\beta}{\sin(180-\alpha-\beta)}$$

- Known Parameters:
  - α: Angle between base and laser beam
  - β: Angle between base and light point (seen from camera)
  - b: Distance between projector and camera (Base)
- Wanted:
  - Z: depth

Single Spot

- Laser beam hits the object
- Camera records the position (angle) of the bright spot → single surface point

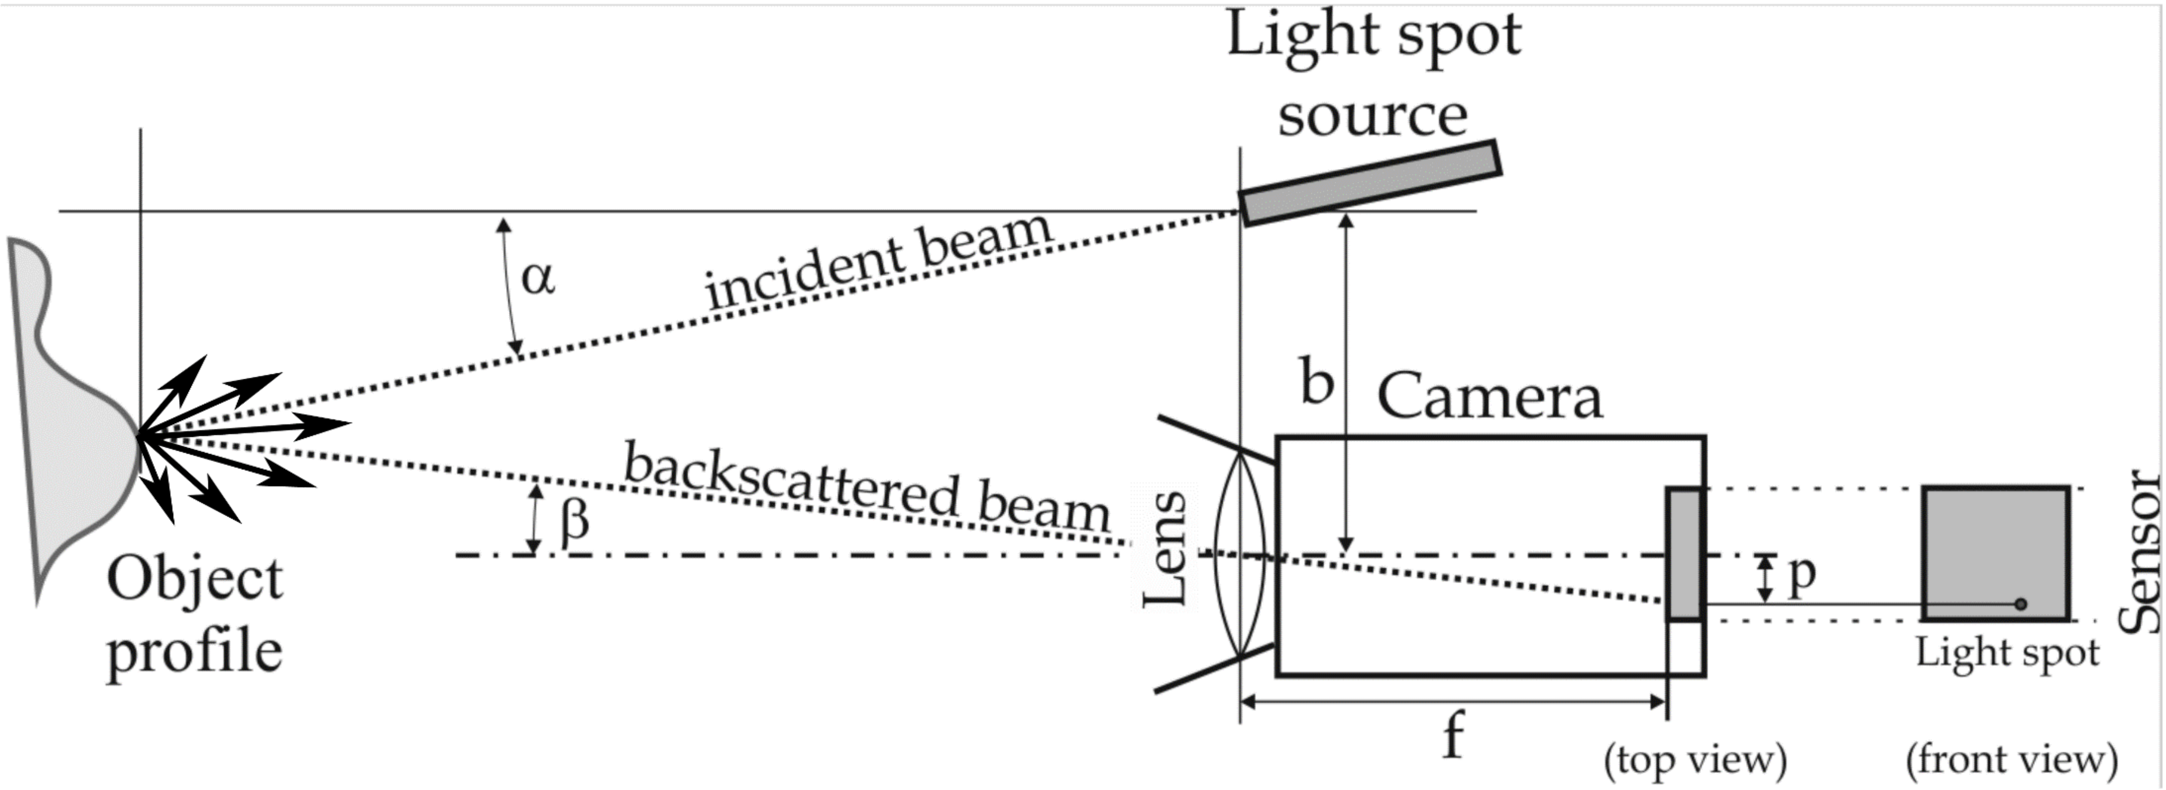

Plane of light

- Laser plane hits the object
- Camera records the bright line → single surface profile

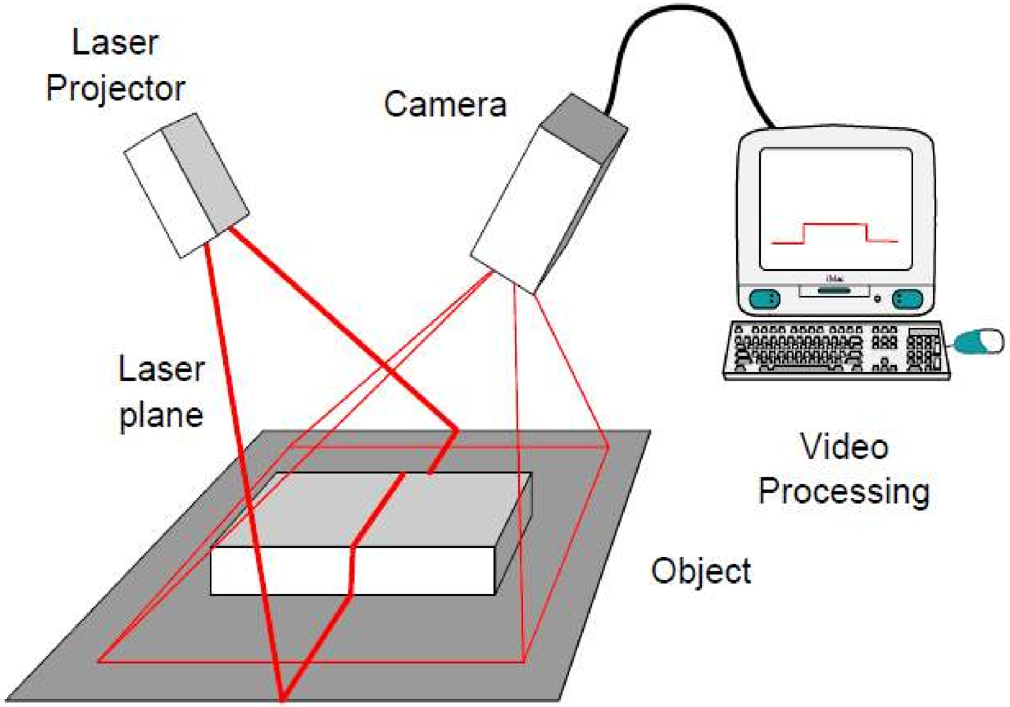

Multiple planes

- Idea: record multiple planes at once, speed up the process
- Problem: parts of strips can be occluded → discontunuities
- no un-ambiguous identification of light strips possible

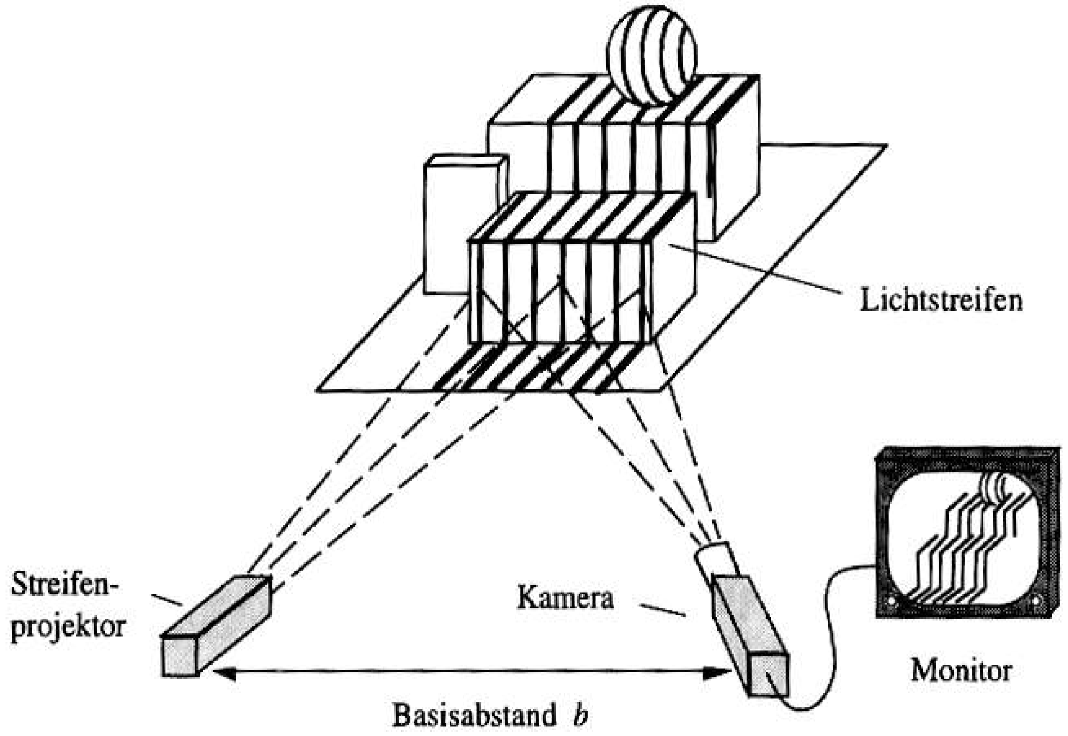

Solution: Gray-code pattern

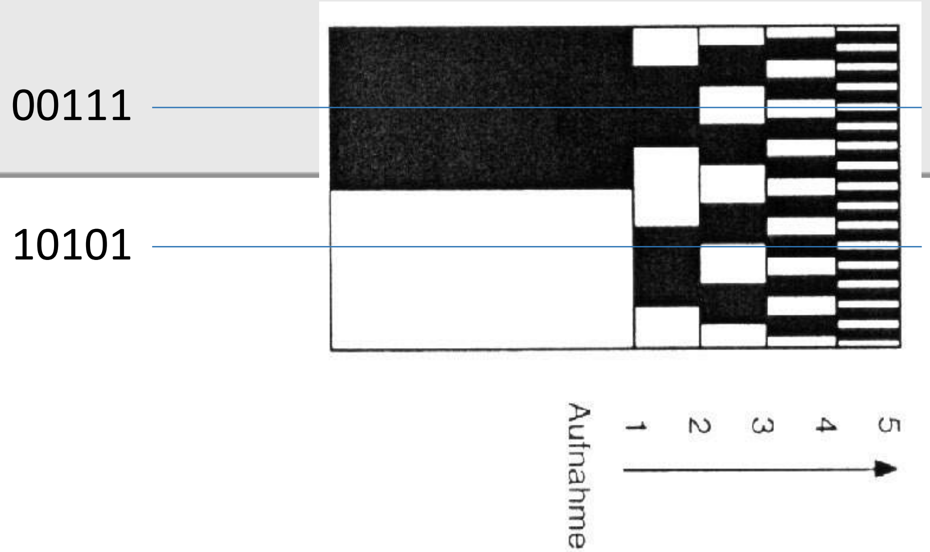

- Project a sequence of stripes
- Each pixel receives a series of bright/dark intensities → binary code
- Unique localization in the pattern possible

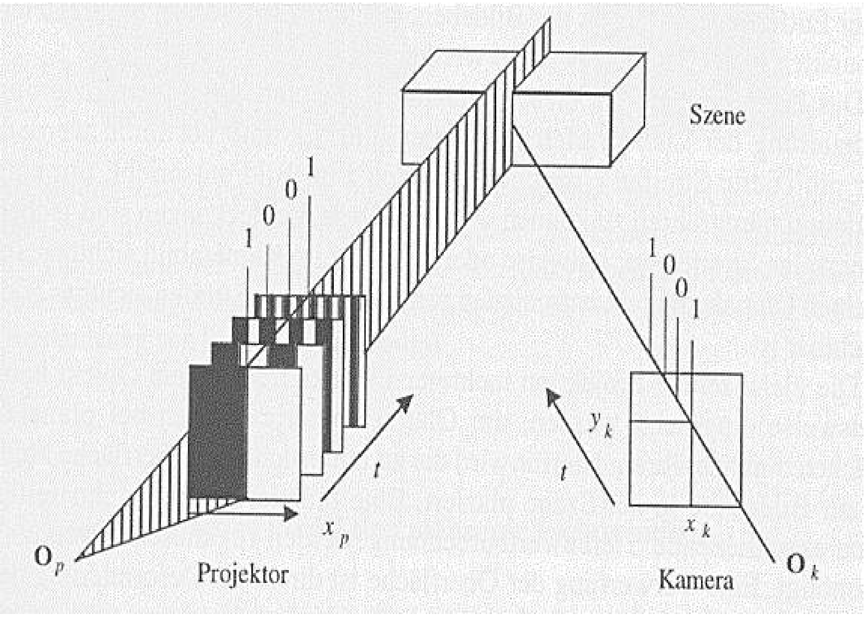

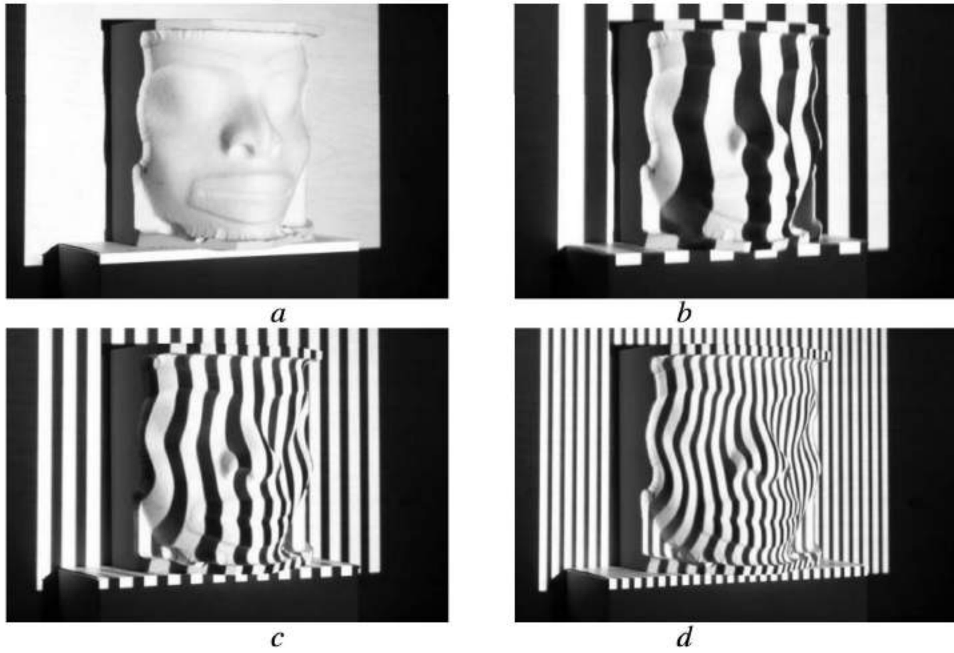

Active Triangulation – applications

- Rather for close-range applications
- Archaeological objects, industrial applications
- High accuracy (micrometers)

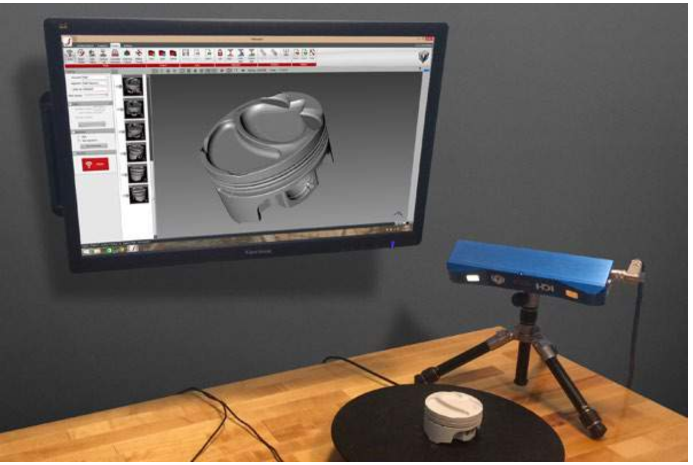

Why are long ranges problematic?

Larger ranges require larger baselines

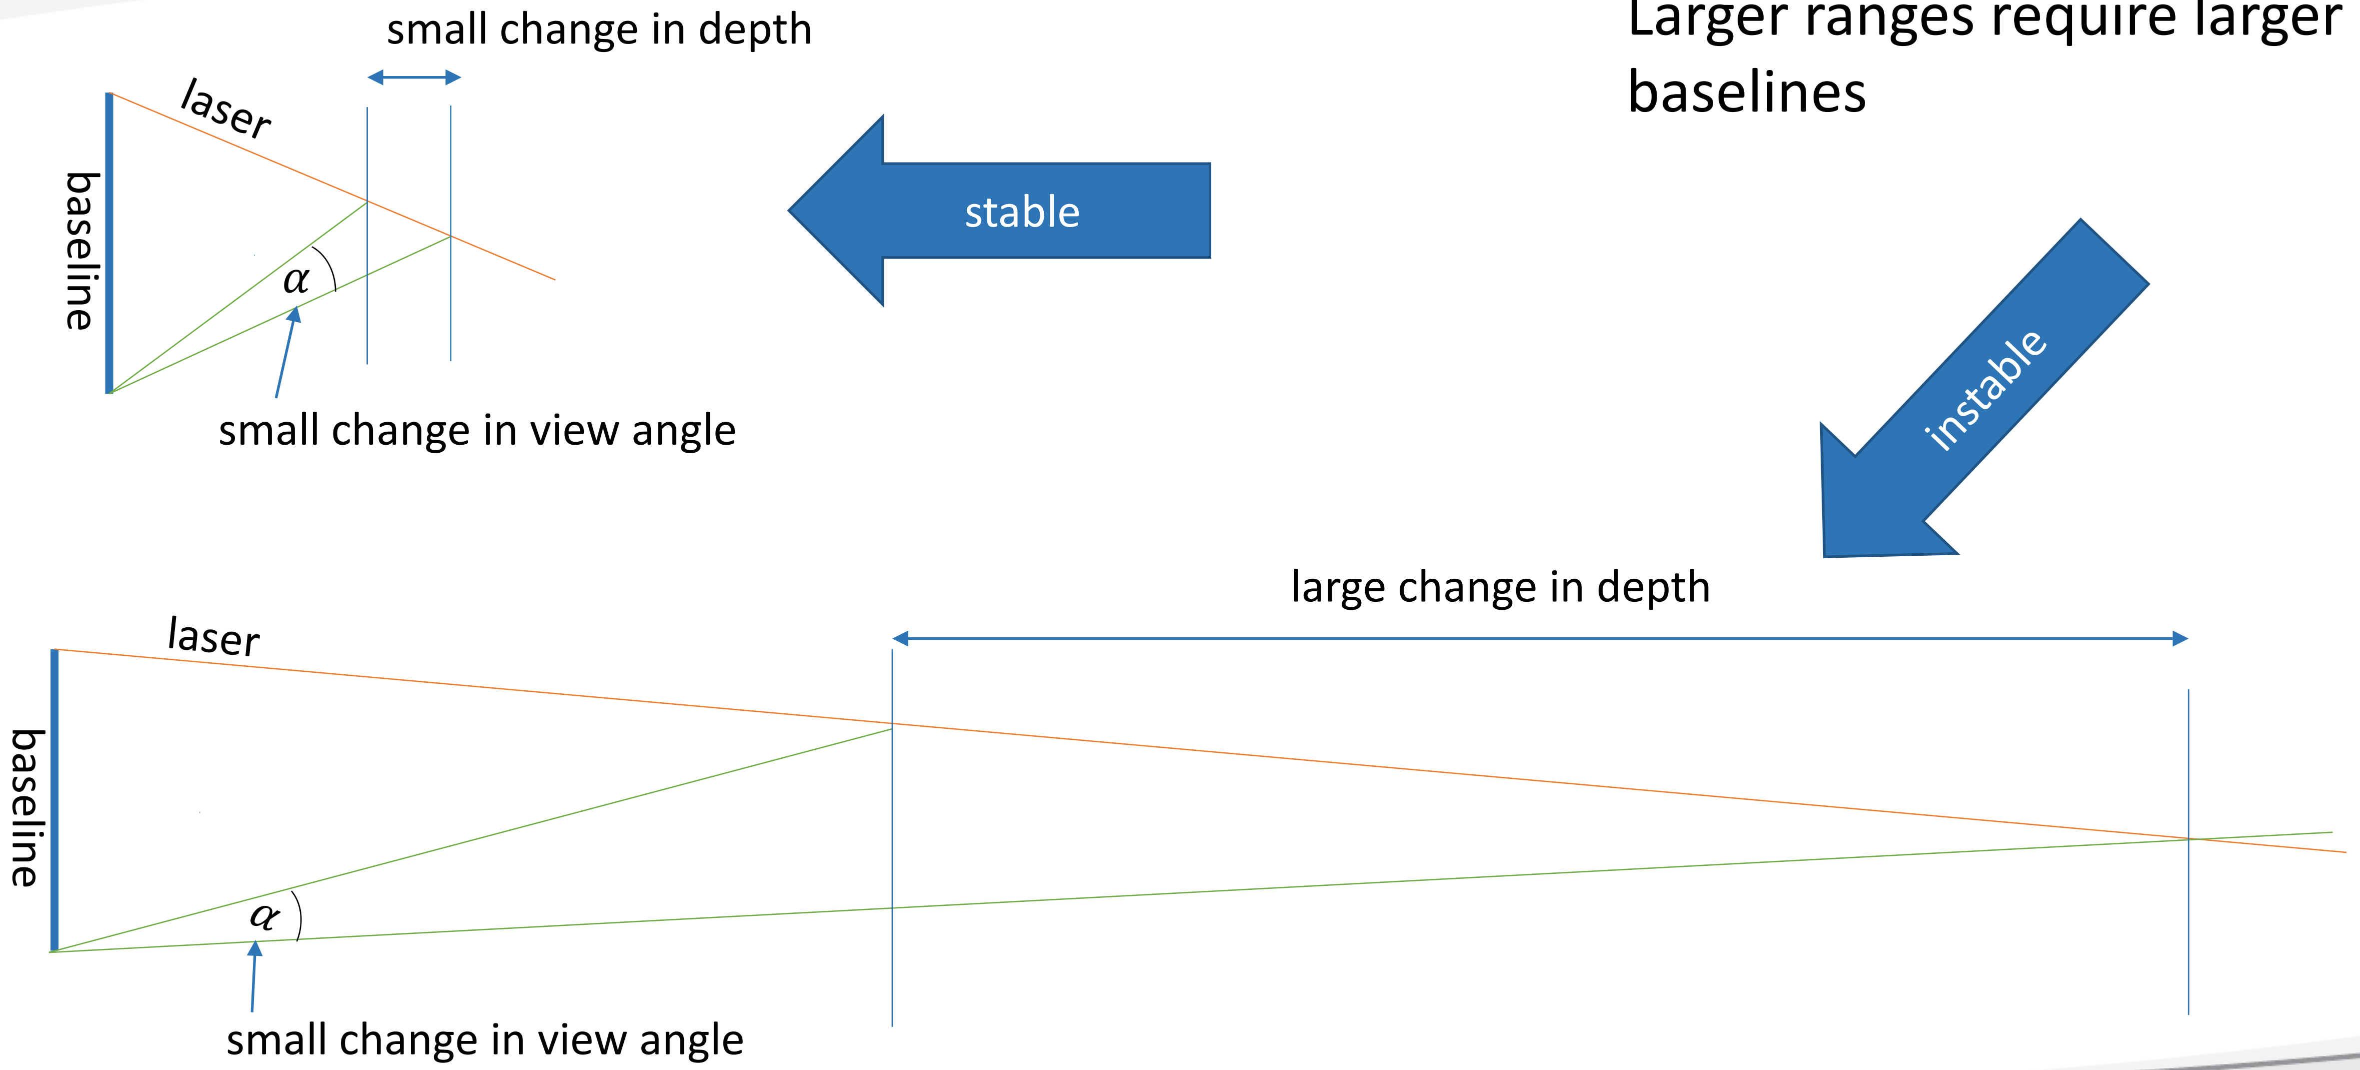

Lecture Outline

- Introduction
  - Why 3D reconstruction?
  - Basic concepts
- Reconstruction Approaches
  - Active triangulation
  - Range scanners
  - Photogrammetry (with how-to)
  - Photometric Stereo
- Application to written heritage
  - Examples from research

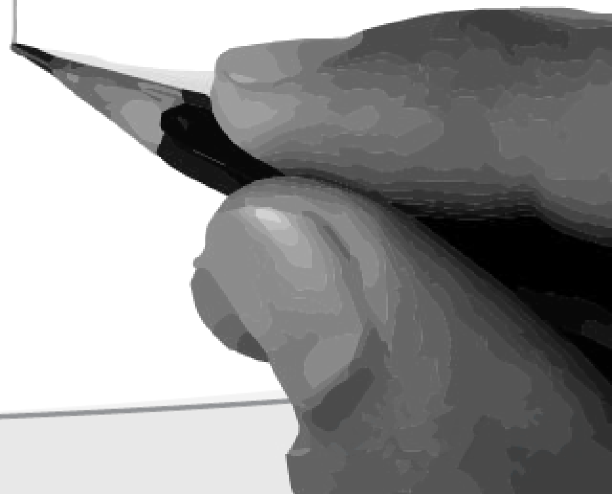

Range scanners – Time of Flight

- Shoot a light pulse at the object and measure the time until its reflection is measured
- "LIDAR" (Light Radar, Light Detection and Ranging)
- $d = \frac{c * \Delta t}{2}$
- c… speed of light
- Accuracy of time measurement is limiting factor

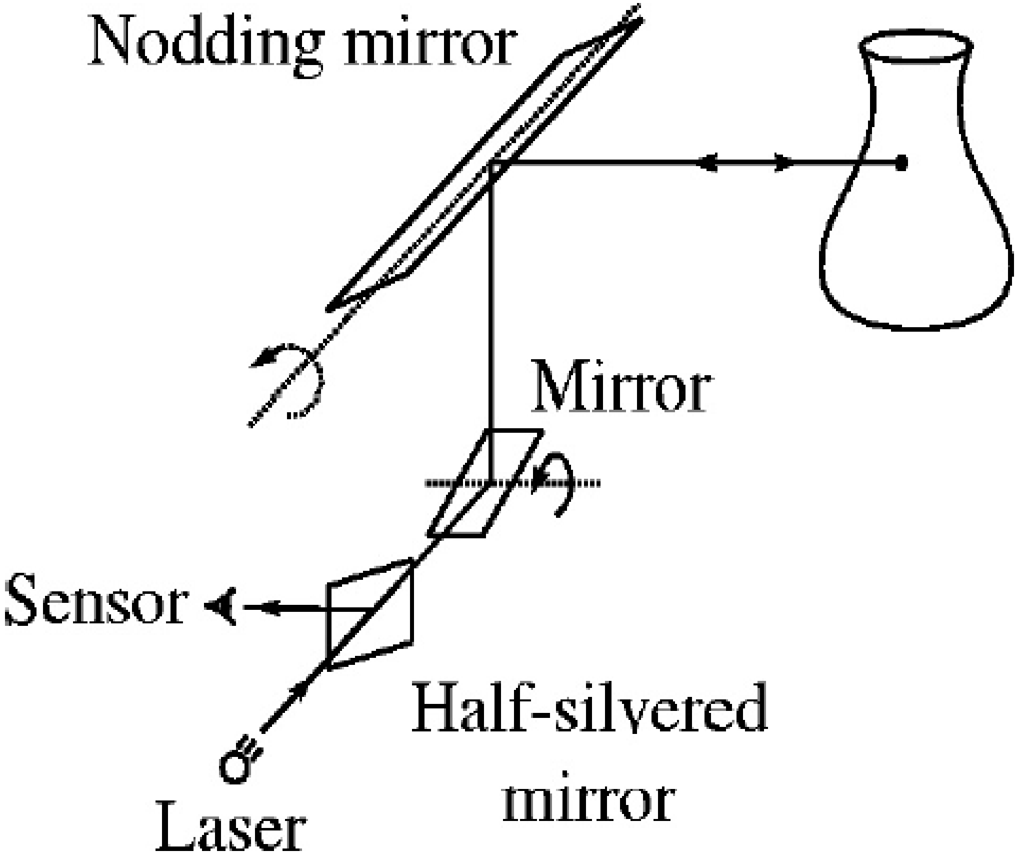

Range scanners – Time of Flight

- For long ranges
- Accuracy in mm-cm area
- Also airborne applications

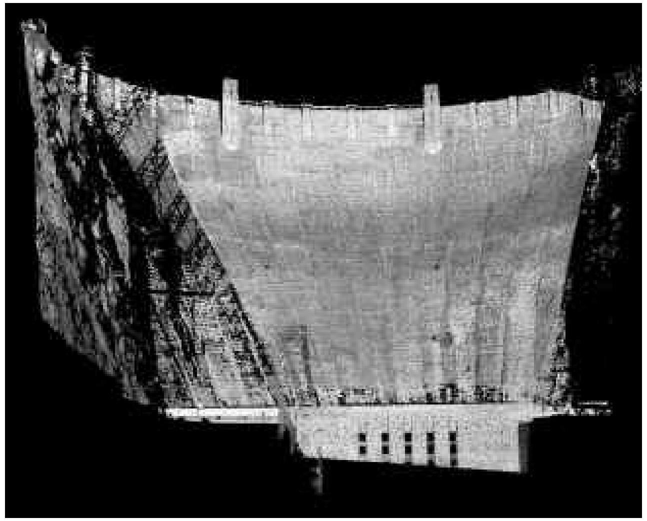

cyrax

Mensi GS100

Range scanners – phase difference

- AMCW: Amplitude-Modulated Continuous Wave
- Measure the phase difference between emitted and measured wave
- Ambiguity if distance is longer than modulation wavelength

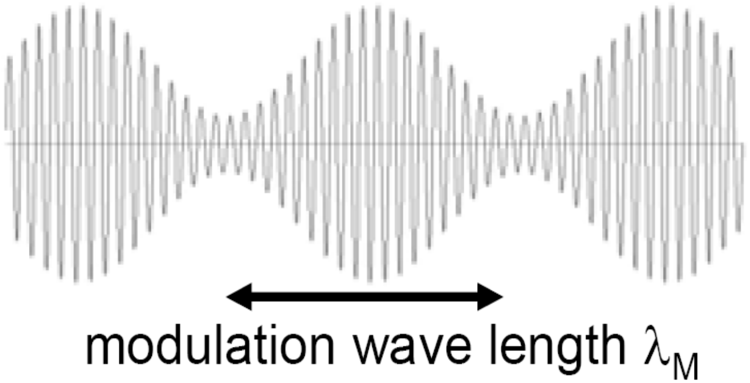

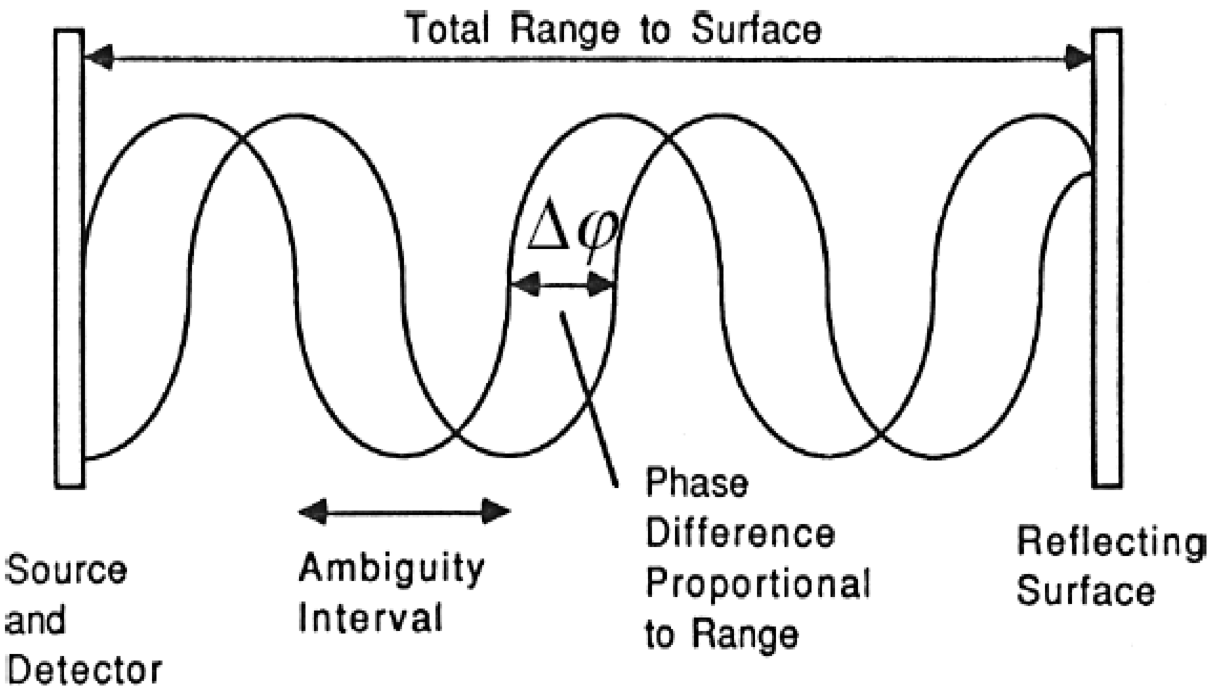

Range scanners – phase difference

- Medium distances → ambiguity interval
- Accuracy in mm area

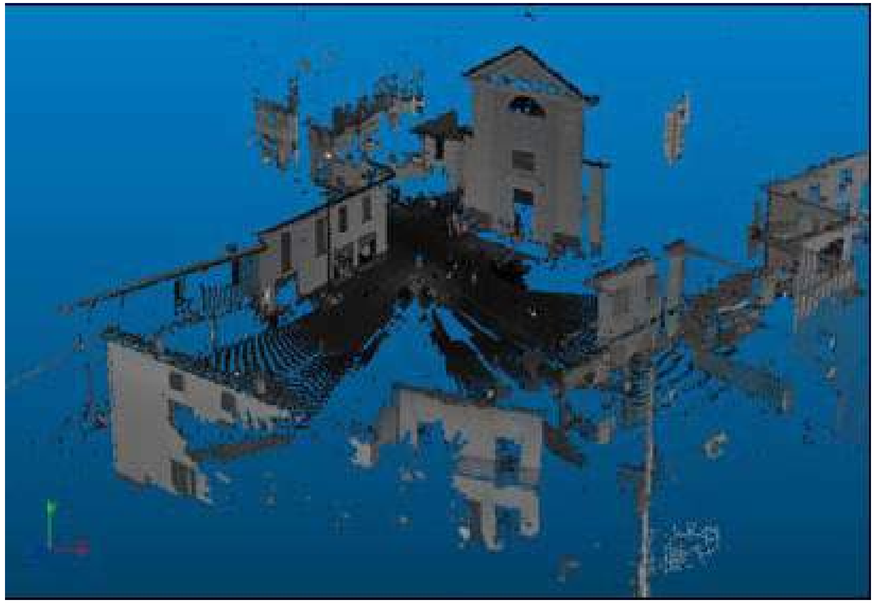

Zoller Fröhlich IMAGER 5300

Lecture Outline

- Introduction
  - Why 3D reconstruction?
  - Basic concepts
- Reconstruction Approaches
  - Active triangulation
  - Range scanners
  - Photogrammetry (with how-to)
  - Photometric Stereo
- Application to written heritage
  - Examples from research

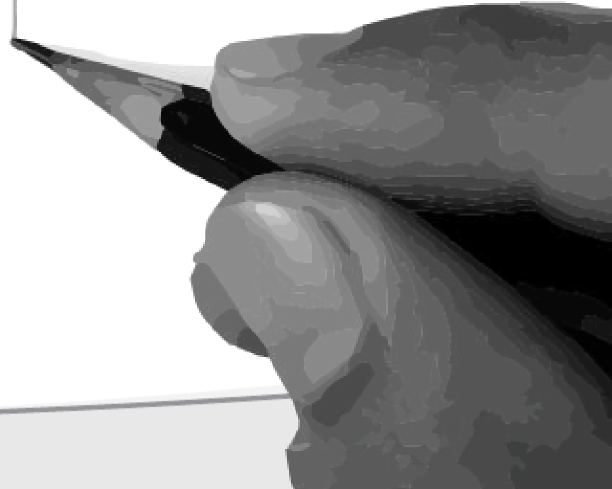

Photogrammetry

- Reconstruct the shape from a set of images, taken from different positions
- Passive method - no additional light / radiation is introduced
- Depth estimation via triangulation
- **Relies on finding corresponding points in different images**

Photogrammetry – Basic algorithm

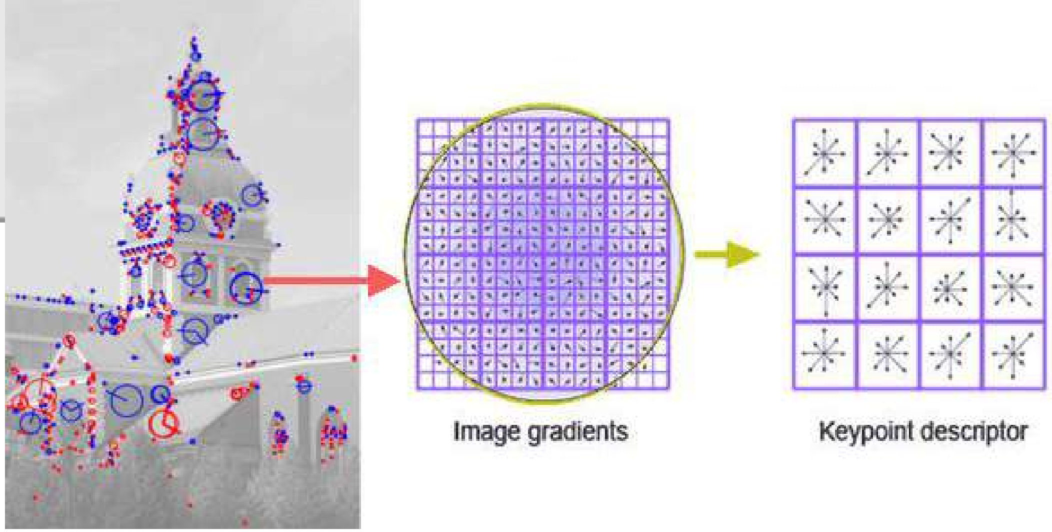

1. Find **matching points** between images
   a) Detect distinctive points in each image
   b) Compute descriptors for each point
   c) Find best matching points between images
2. Estimate camera position & orientation for each photo

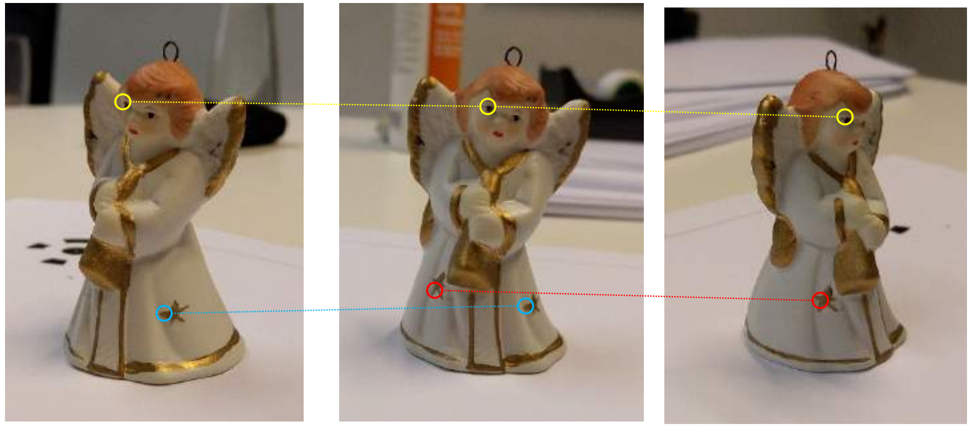

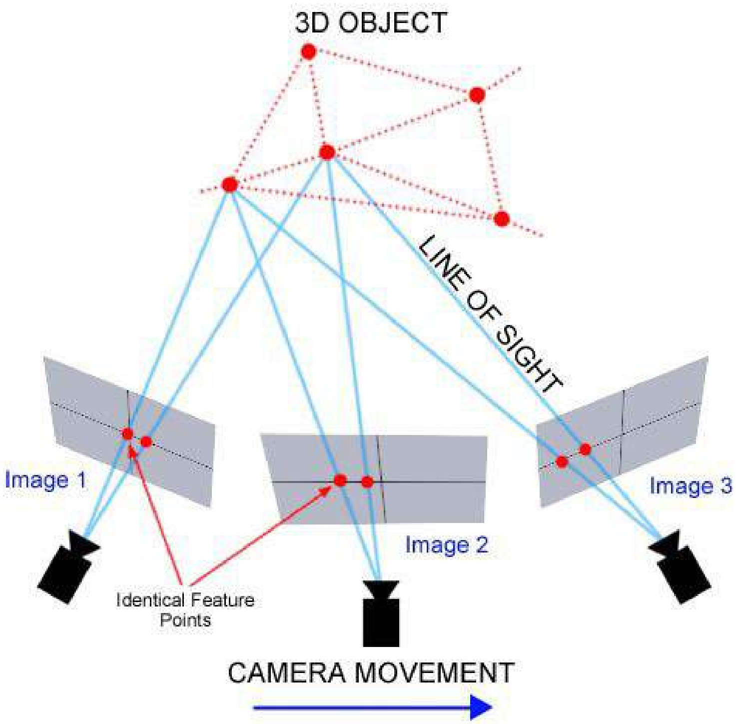

Photogrammetry – Basic algorithm

3. Dense reconstruction
    a) Knowledge of camera poses simplifies the point matching problem
    b) **Epipolar geometry**: a point in one image corresponds to a line in the other image
    c) Search for matching points only along epipolar lines
    d) Triangulate point positions

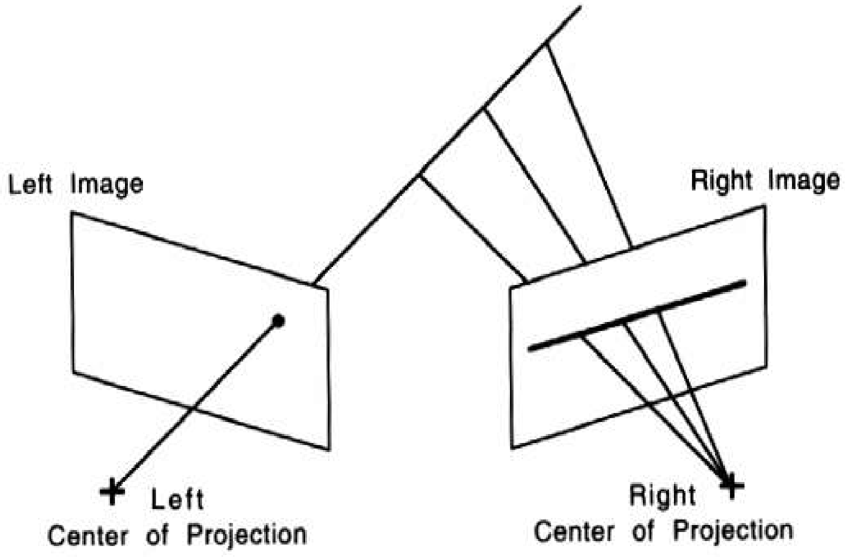

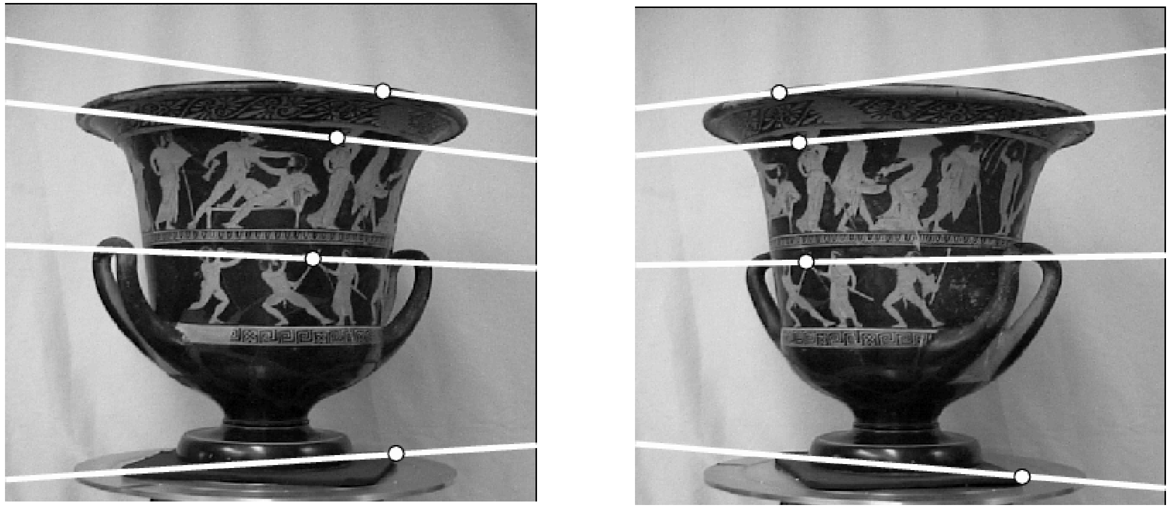

Photogrammetry – Basic algorithm

4. **Surface reconstruction:**

   Unordered 3D Points → Surface Mesh

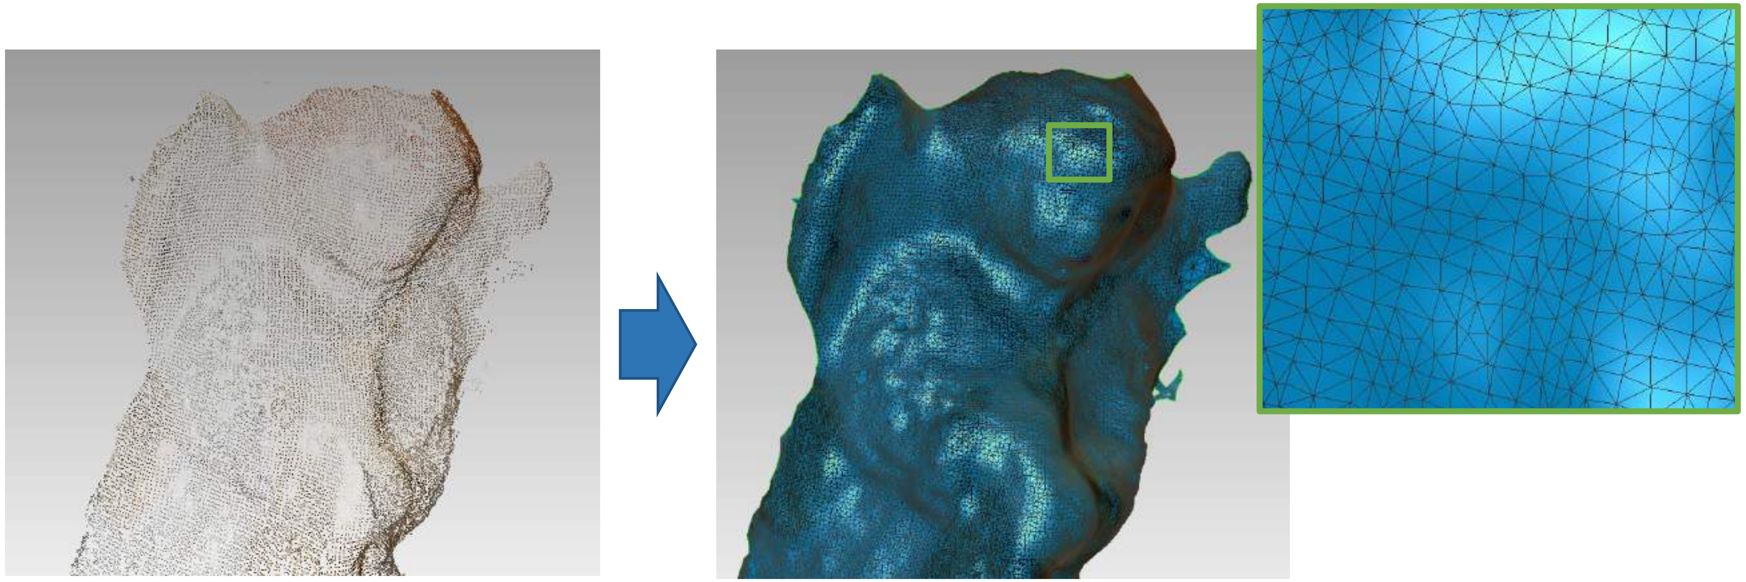

How to Photogrammetry

1. Check if the method is applicable for your object
2. Take **good** photos (most difficult and important)
3. Process them with a Photogrammetry Software, e.g.:
   - Agisoft Metashape (commercial)
   - Meshroom (open source)
4. (optional: editing)
5. Marvel at your results!

1. Applicability

Here the surface color depends on the angle of view!

No unique point matches possible!

Shape might change during acquisition!

Don‘t:

Shiny

Transparent/translucent

Smooth, uniform surfaces

Deformable

Do:

- "Natural" materials
  - Wood
  - Stone
  - Clay
  - …
- Corroded / worn materials
- Ragged, irregular surfaces

2. Take **good** pictures

- If you want nice results, invest some time & effort here!
- Different acquisition strategies for different applications

Fixed camera, rotating object

Moving camera, fixed object

Multiple synchronized cameras – nothing moves

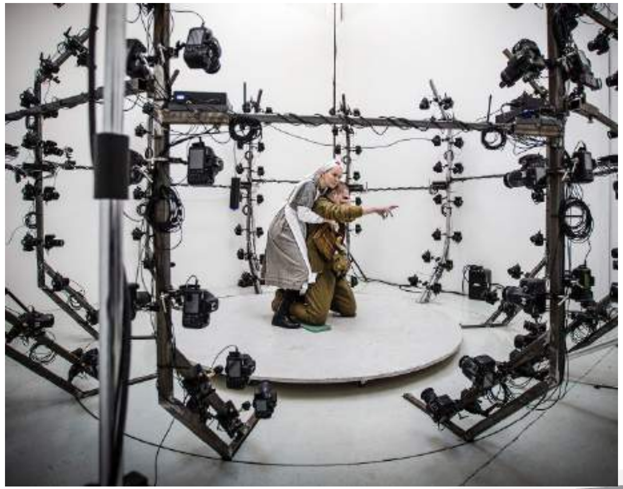

Recommended for beginners: object fixed, camera moving

- Minimal setup efforts, very versatile
- Applicable for a wide range of objects
  - Small artifacts, buildings, excavation sites, cities…
- Applicable in a wide range of environments
  - Indoors, outdoors, airborne (e.g. with drones)…




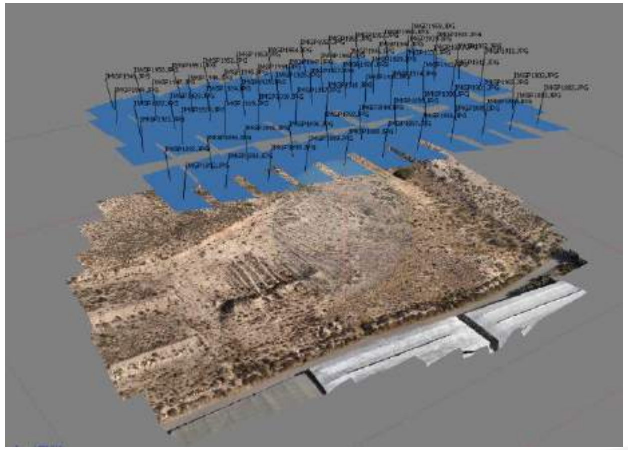

For smaller objects

- Put your object somewhere accessible (stand, stool, etc.)
- Walk around the object in ~10° steps → 1 full round = ~36 pictures (approximately – no need to measure that)
- Do this from 2-3 elevation angles
- NEVER MOVE THE OBJECT during acquisition
  - relation between object and background must not change
  - Different position → different shadows → different surface color
- Try to keep approx. the same distance
- If possible/necessary: Turn the object upside down and repeat (to capture the bottom side) - this will result in a separate model!

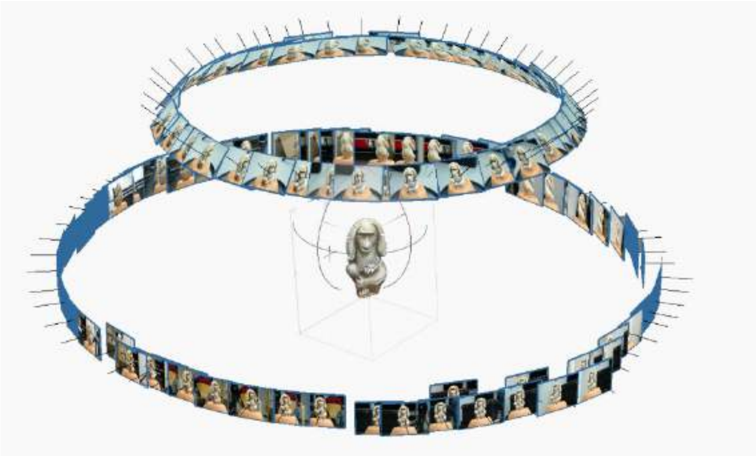

For larger objects (buildings, etc.)

- Walk parallel to the wall/object surface
- Consecutive images should overlap by min. 80% (remember: to reconstruct a point, it must be visible in min. 2 images!)
- Good idea: find out how many steps you should take between two shots
- Sometimes turn left and right to
  - increase overlap areas
  - Create larger baselines
  - Reach blind spots

$$\frac{sensorsize}{coverage} = \frac{focallength}{objectdistance}$$

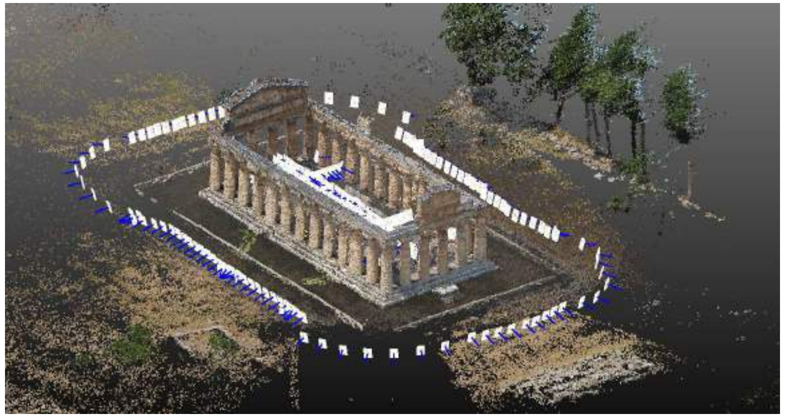

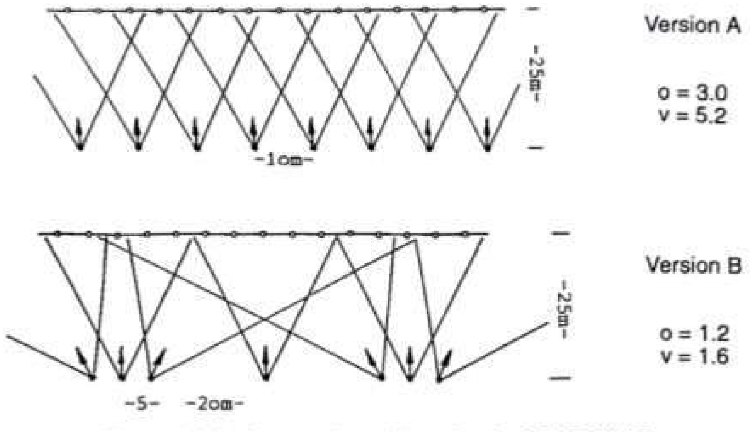

Figure 4.5-3: Camera dispositions for the LONG WALL

Taking photos – general advice

- Use a "real" camera
  - DSLR, mirrorless,...
- Shoot in manual mode, keep camera parameters constant
- Don't change the focal length (zoom)
- Care about illumination
  - Enough light!
  - Smooth, uniform
  - **Life-hack: shoot outdoors on a cloudy day!**

Use a small aperture

- Most of the object should be in focus
- Small aperture (=high f-number) → large depth of field
- This can lead to longer exposure times → Use a tripod!

Put a reference object in the scene

- Necessary to correctly scale your model
  - Grid
  - Ruler

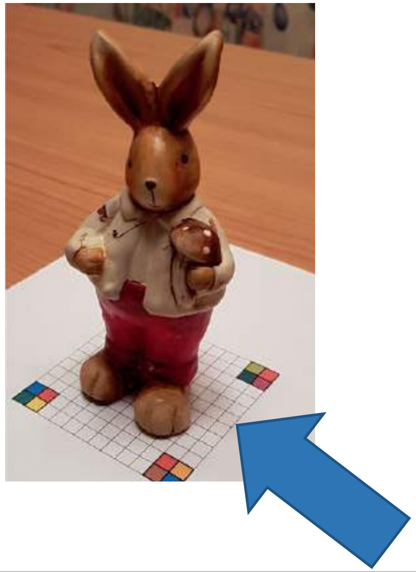

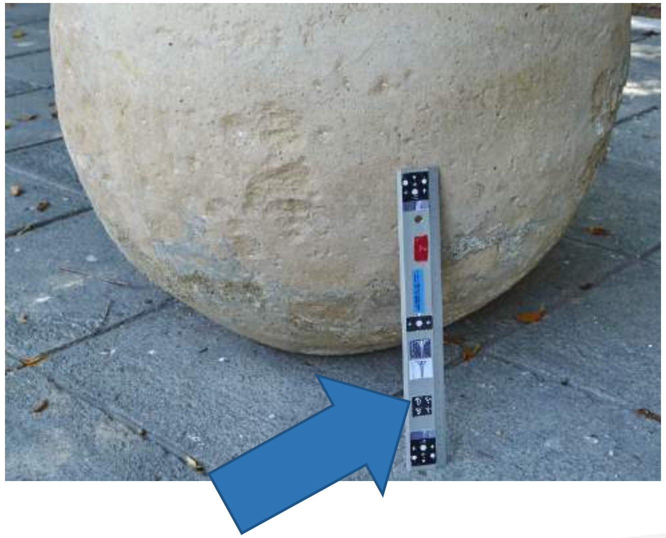

3. 3D Reconstruction

- Just use an existing software tool and follow the manual…
  - My favourites:

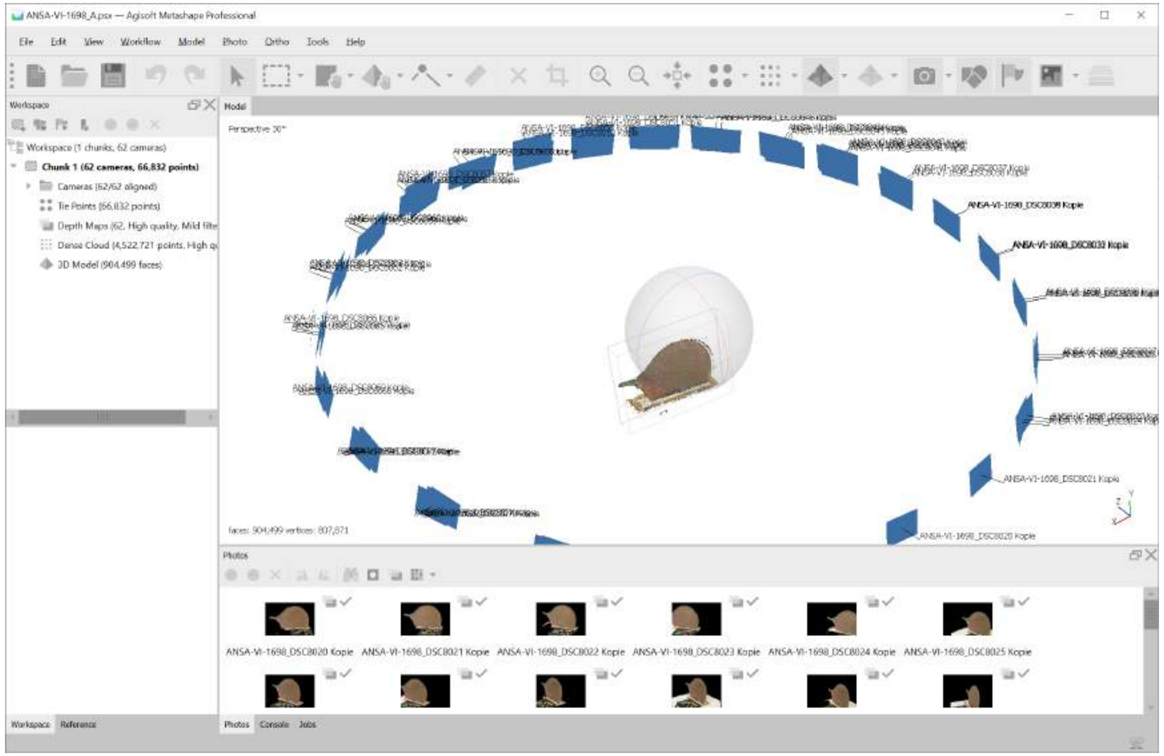

Agisoft Metashape Professional (30 days free trial available)

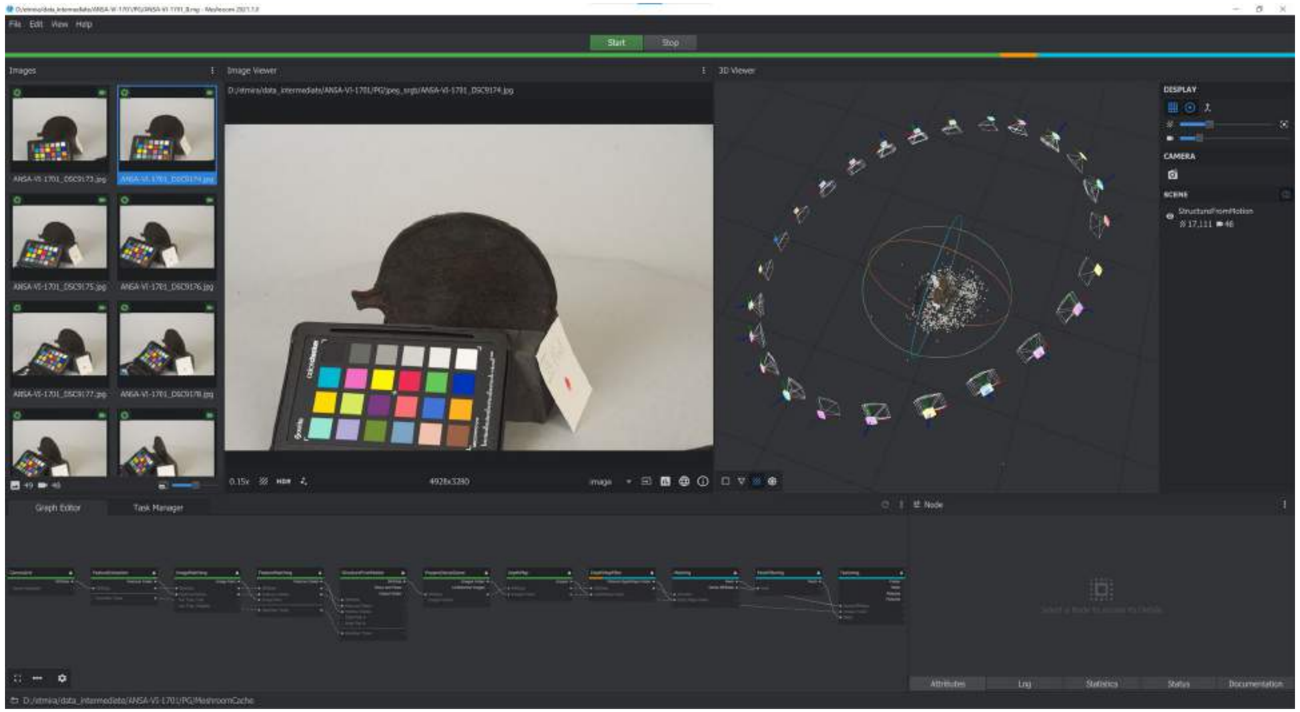

Meshroom (open source!)

4. Editing

- Close holes
- Remove unnecessary parts
- Repair broken geometry
  - Broken normal vectors
  - Non-manifolds
  - Overlaps/intersections
- Professional software: *Geomagic Wrap*
- Free options:
  - *Autodesk Meshmixer*: very user friendly, a bit limited
  - *Blender*: 3D modelling software – very powerful, steep learning curve
  - *Meshlab*: many features, a bit overwhelming for the non-scientific user

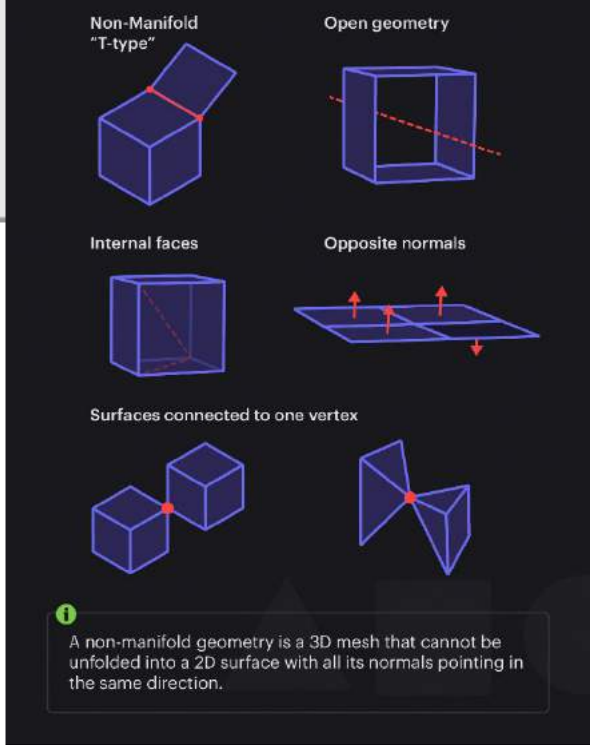

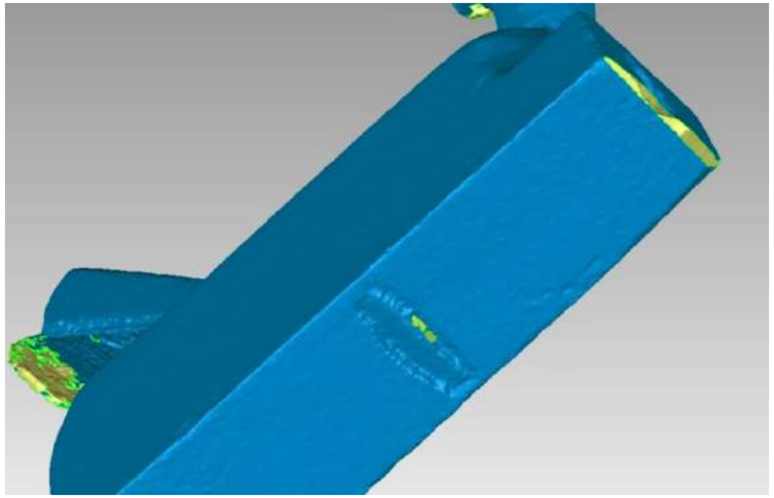

5. Marvel at your results

Lecture Outline

- Introduction
  - Why 3D reconstruction?
  - Basic concepts
- Reconstruction Approaches
  - Active triangulation
  - Range scanners
  - Photogrammetry (with how-to)
  - Photometric Stereo
- Application to written heritage
  - Examples from research

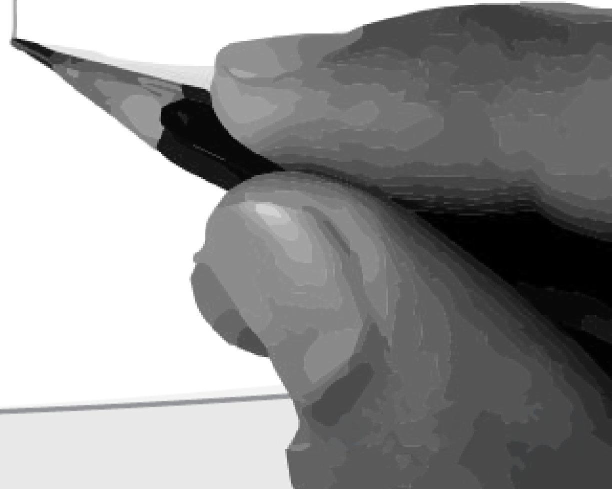

Photometric Stereo

- Reconstruct 3D surface from multiple raking light images
- Advantages (over other 3D reconstruction techniques):
  - Relatively cheap and simple acquisition
  - Spatial resolution only limited by resolution of input images
- Disadvantages
  - Susceptible to low-frequency-noise
    → distorted global shape
  - Ambiguities
- ➔ good option for quasi-flat objects with fine details
  - Such as manuscripts…

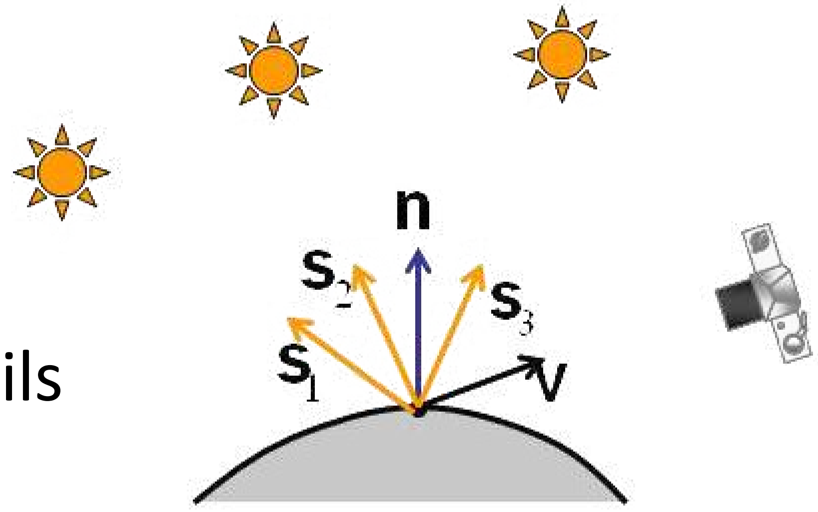

Example

Input images

Albedo map

Surface normals

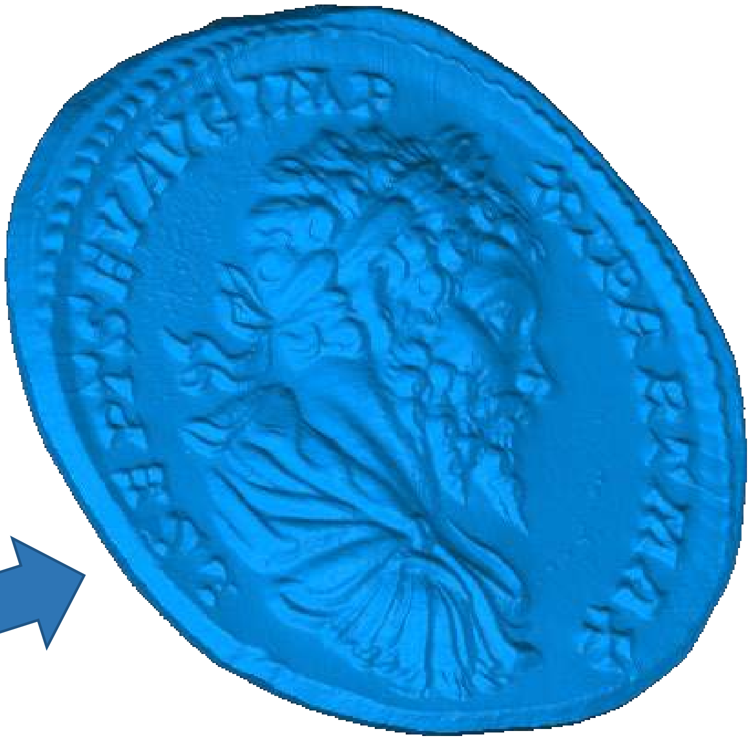

3D model

Shape from Shading vs. Photometric Stereo

- Single input image: „Shape from Shading“
  - Estimate surface inclination from brightness values
  - Reverse imaging process
  - ill-posed problem, highly ambiguous
- Multiple input images: „Photometric stereo“
  - lit from different directions
  - Unique solution possible

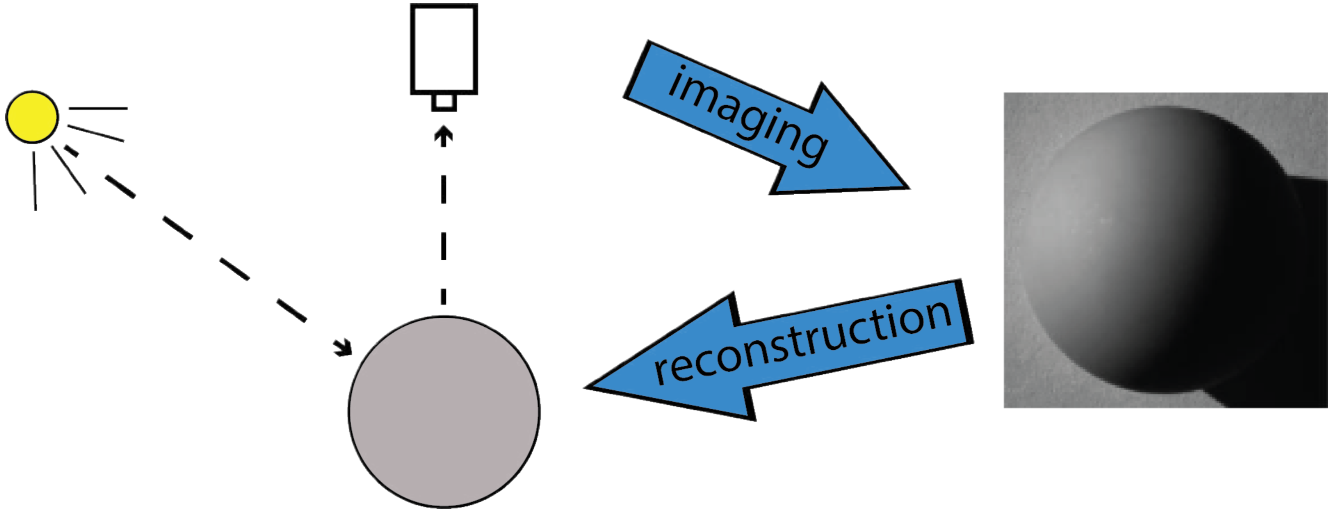

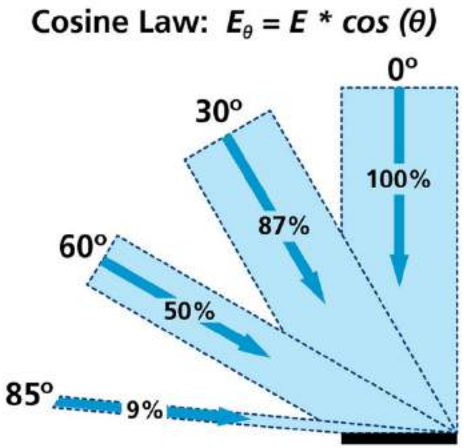

Example setup: light dome

- 54 high-power LEDs in a dome structure
- Open design for easy object handling

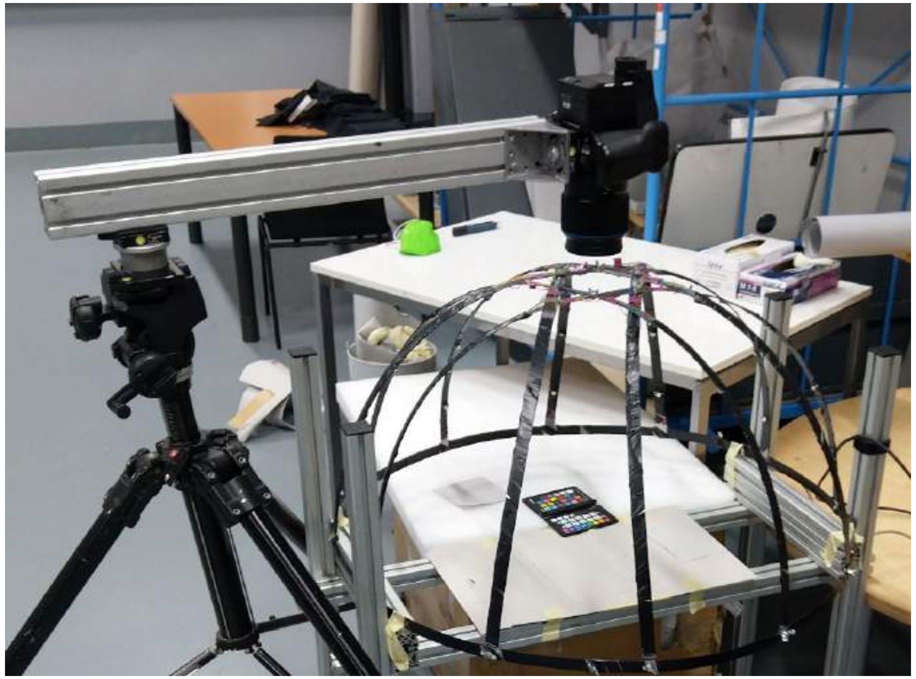

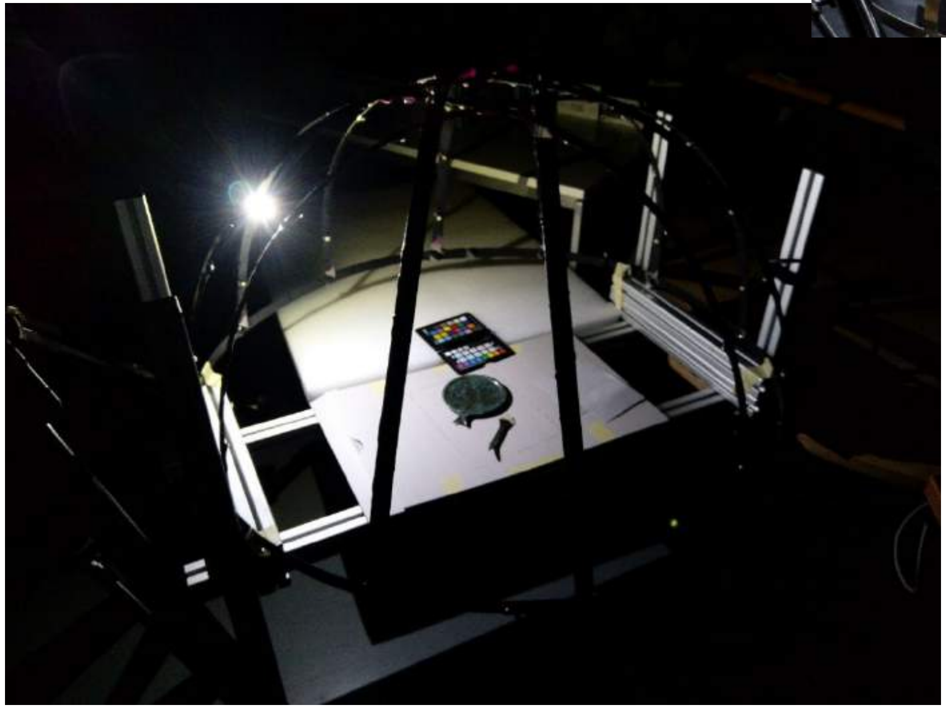

Photometric Stereo vs. Structured Light – Surface details

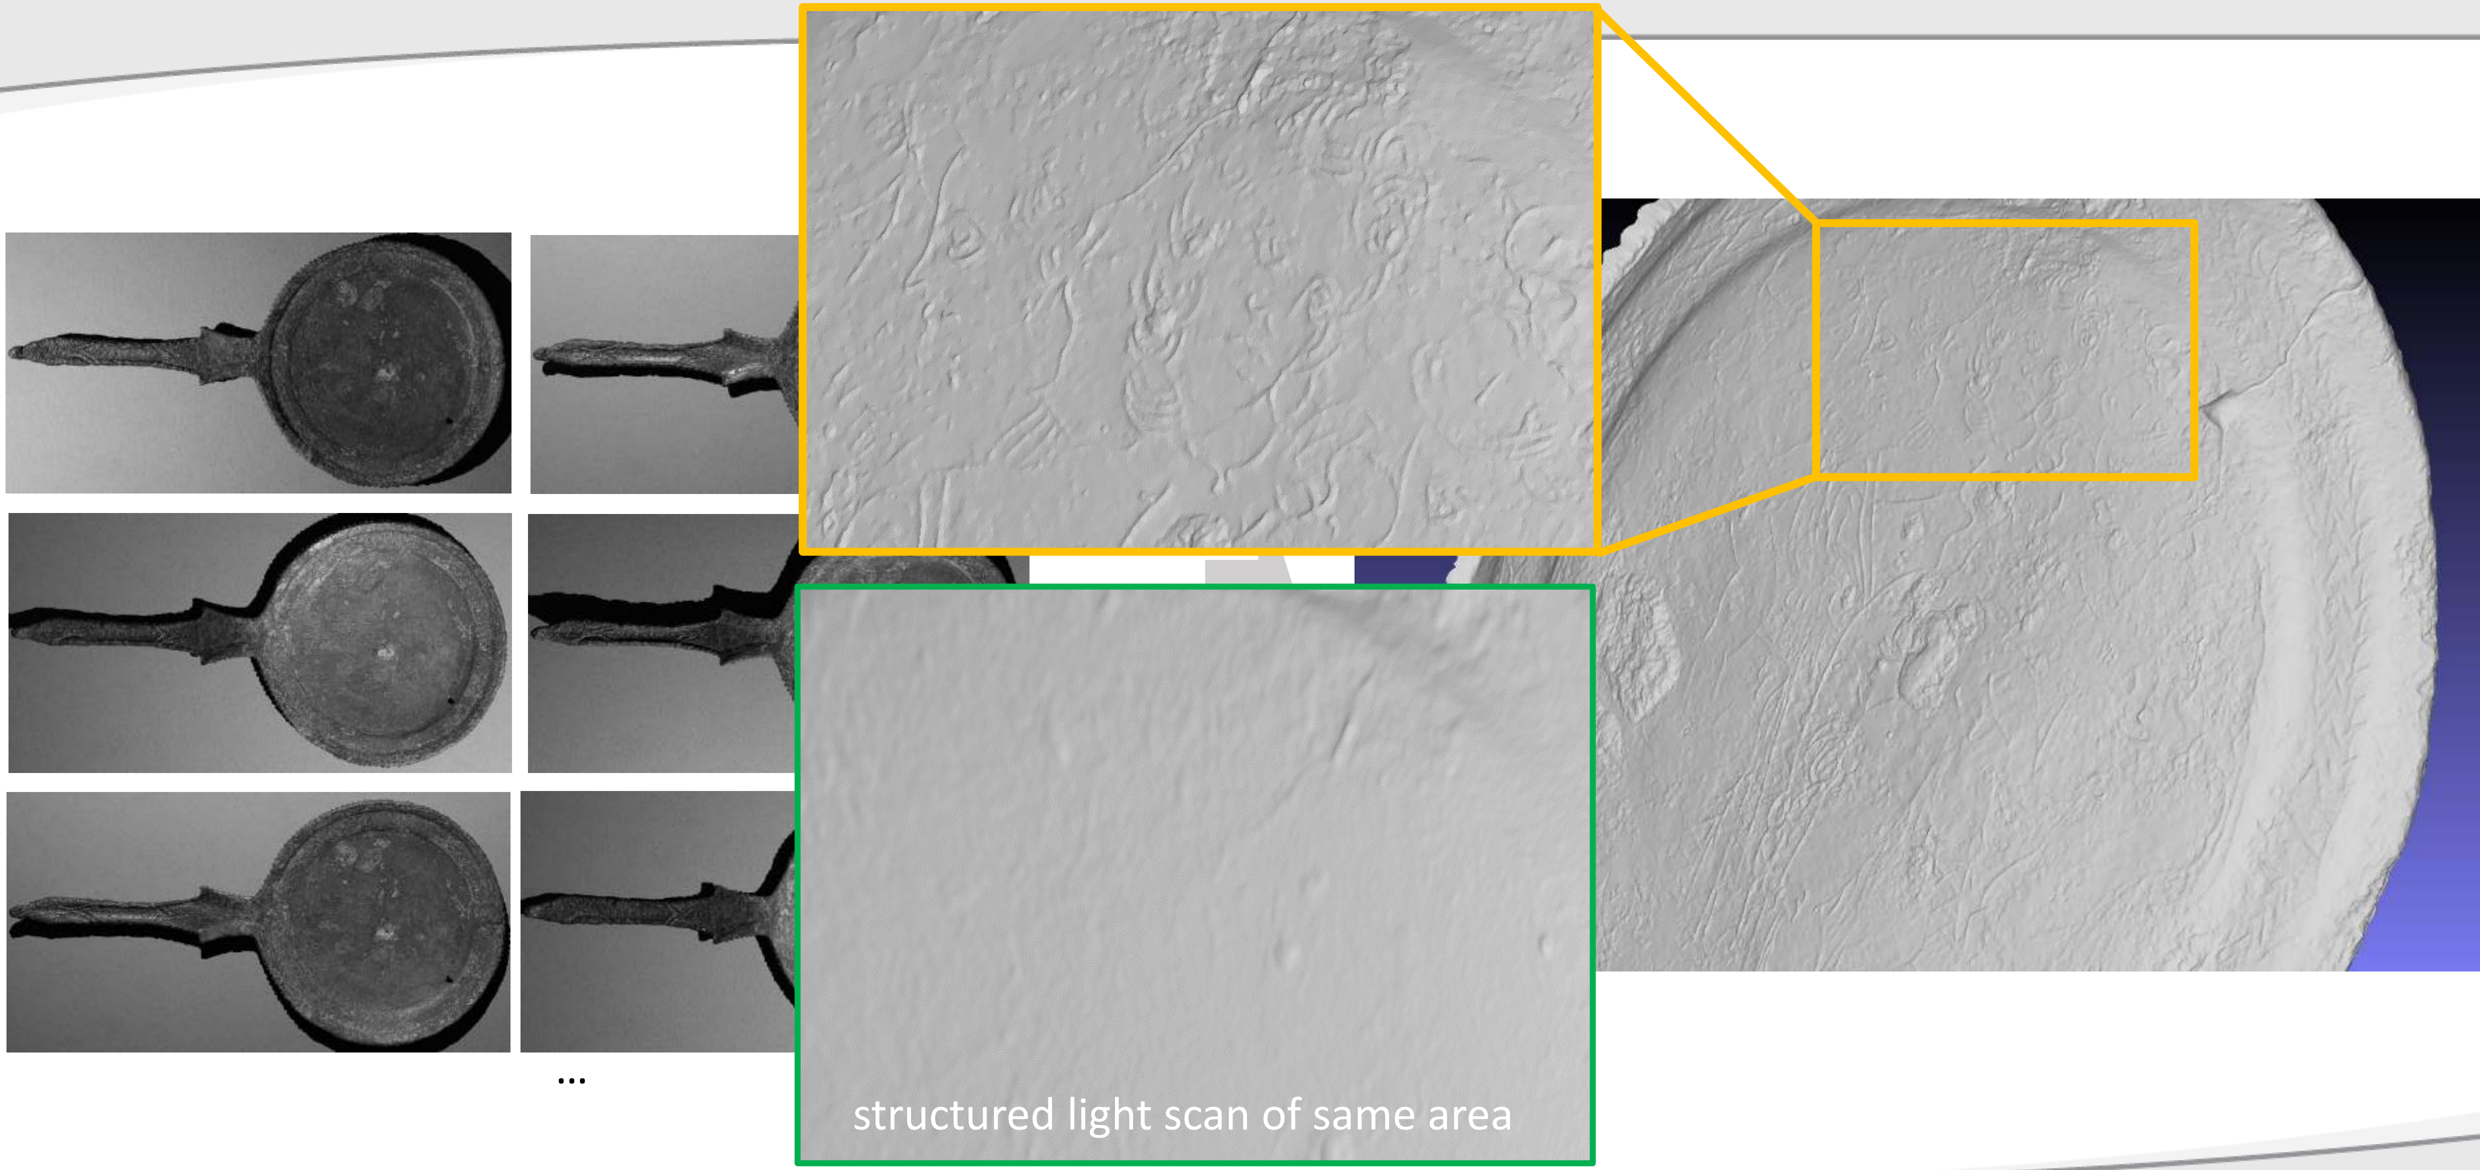

Photometric Stereo vs. Structured Light – Global shape

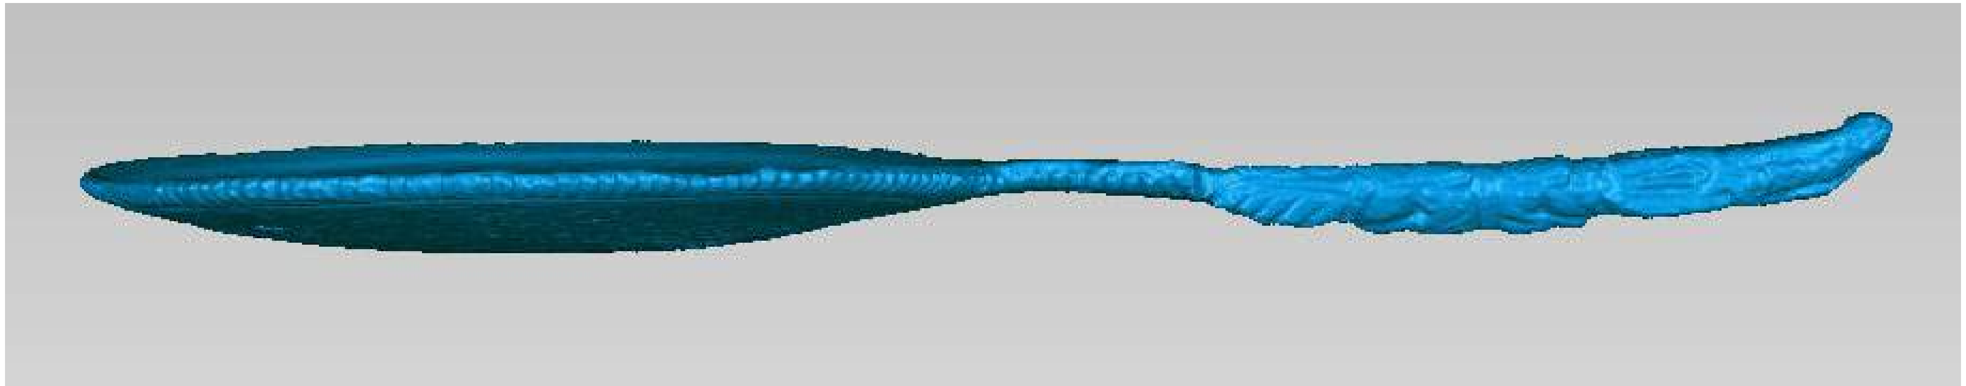

Structured Light

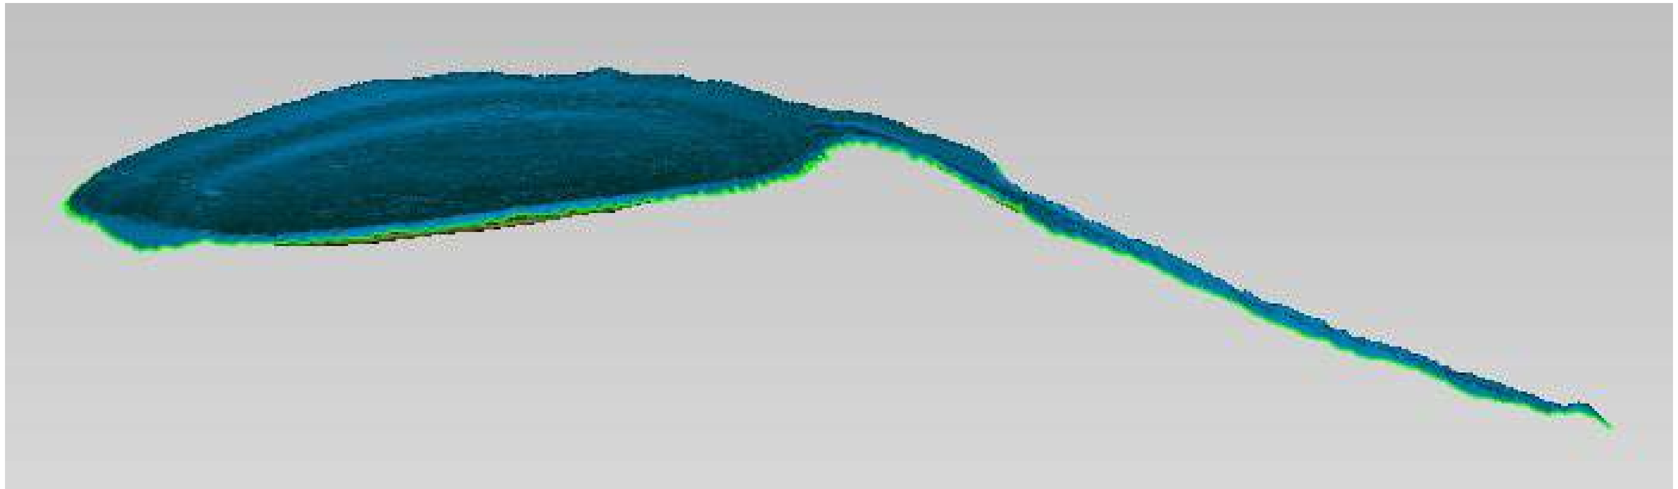

Photometric Stereo

Related: Reflectance Transformation Imaging (RTI)

- Same kind of input images as photometric stereo
- Result: reflectance model
  - Interactive re-lighting
  - Enhancement of details
  - **No real 3D information**
- Free tools available:
  - RTI Builder (creating the model from input images)
    - https://culturalheritageimaging.org/What_We_Offer/Downloads/Process/index.html
    - Needs **reflective balls** in the input images!
  - RTI Viewer (viewing the model)
    - https://culturalheritageimaging.org/What_We_Offer/Downloads/View/index.html

Highlights indicate light directions

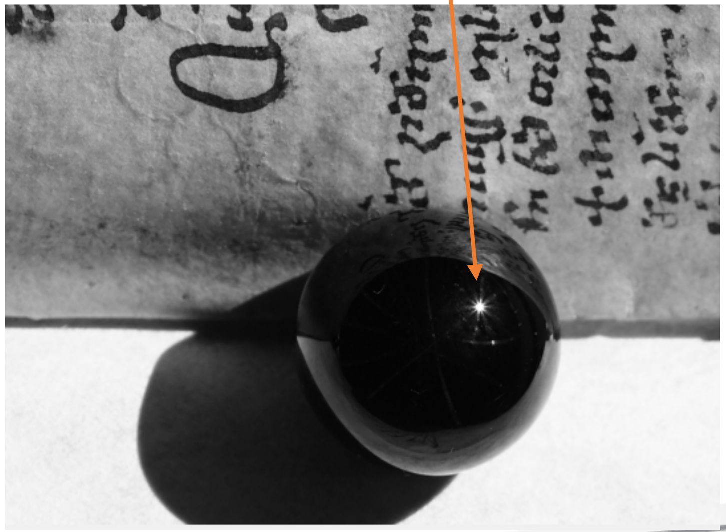

Lecture Outline

- Introduction
  - Why 3D reconstruction?
  - Basic concepts
- Reconstruction Approaches
  - Active triangulation
  - Range scanners
  - Photogrammetry (with how-to)
  - Photometric Stereo
- Application to written heritage
  - Examples from research

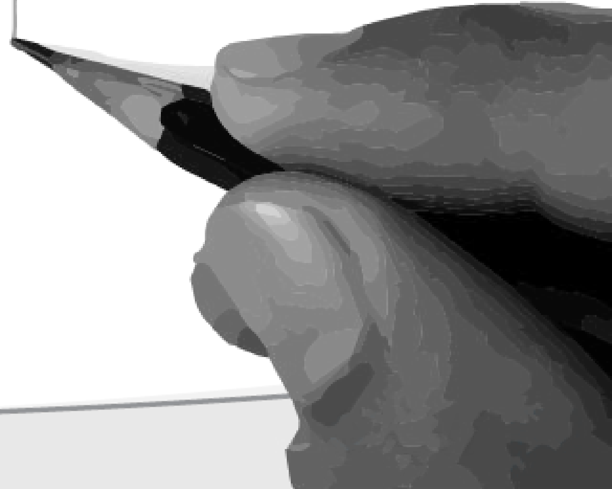

Cuneiform recognition

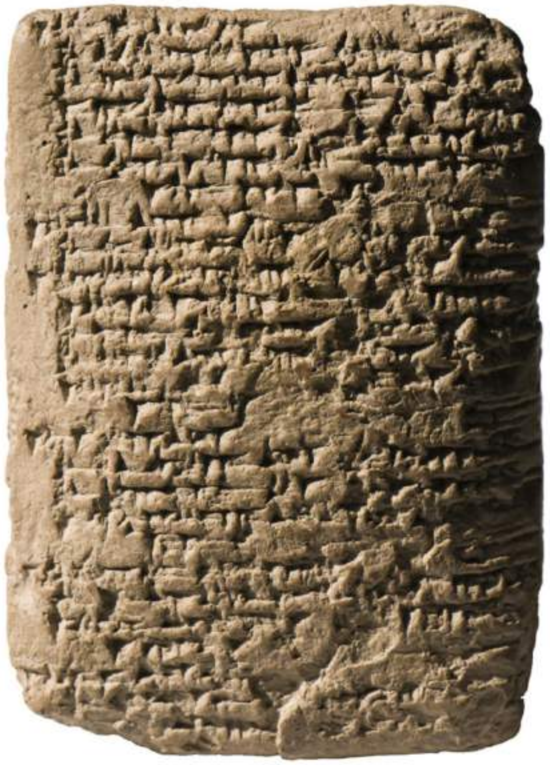

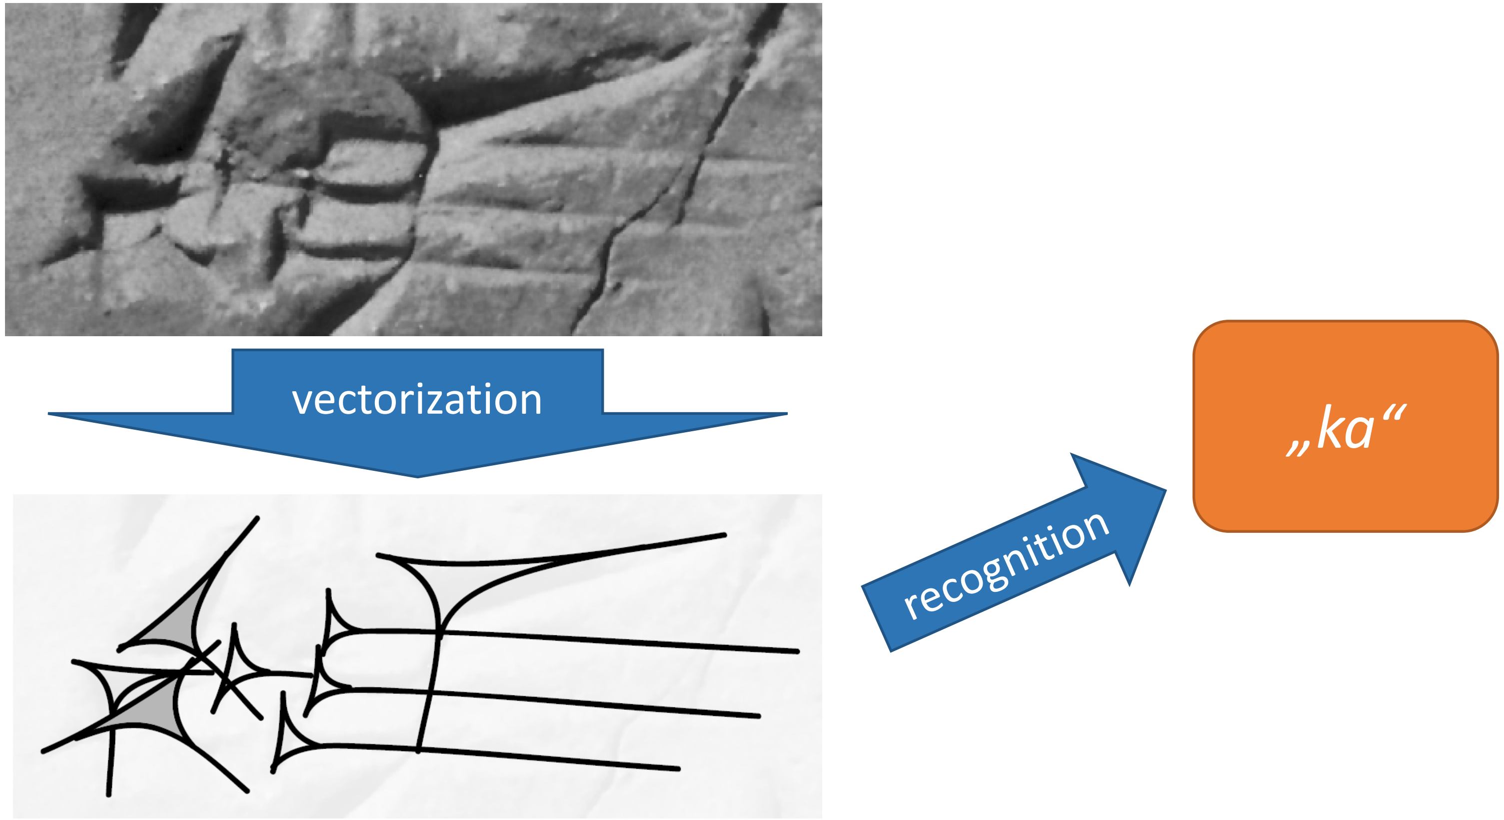

H. Mara and S. Kromker, “Vectorization of 3D-Characters by Integral Invariant Filtering of High-Resolution Triangular Meshes,” in *2013 12th International Conference on Document Analysis and Recognition*, Washington, DC, USA, Aug. 2013, pp. 62–66. doi: 10.1109/ICDAR.2013.21.

B. Bogacz, M. Gertz, and H. Mara, “Character retrieval of vectorized cuneiform script,” in *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*, Tunis, Tunisia, Aug. 2015, pp. 326–330. doi: 10.1109/ICDAR.2015.7333777.

Cuneiform vectorization

Structured light scan

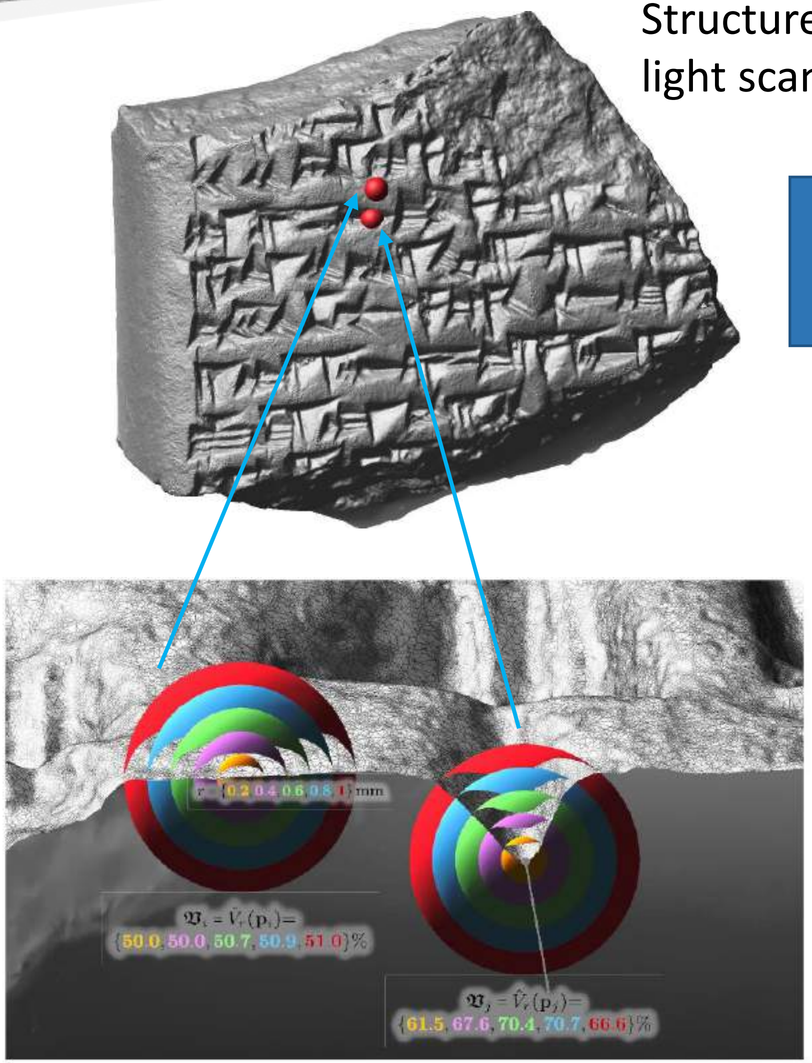

Segmentation, based on MSII

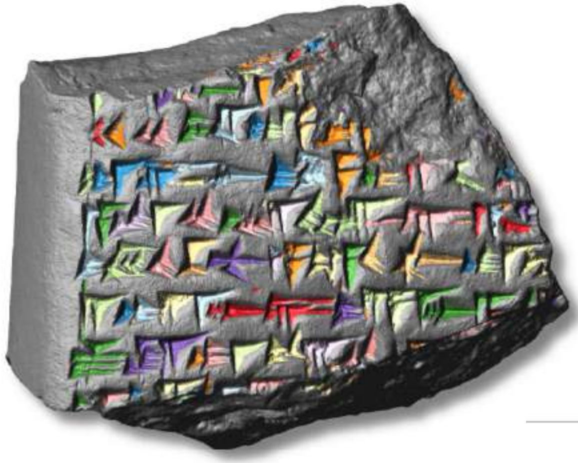

skeletonization

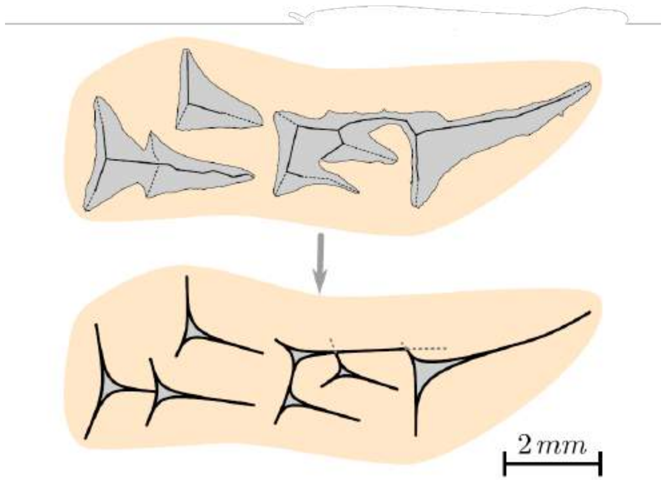

Vectorized result

Multi-Scale Integral Invariants (MSII):

- Place a sphere at each surface point
- Compute percentage of intersection
- Do that for several sphere sizes → each point is described by a vector of ratios

Cuneiform recognition

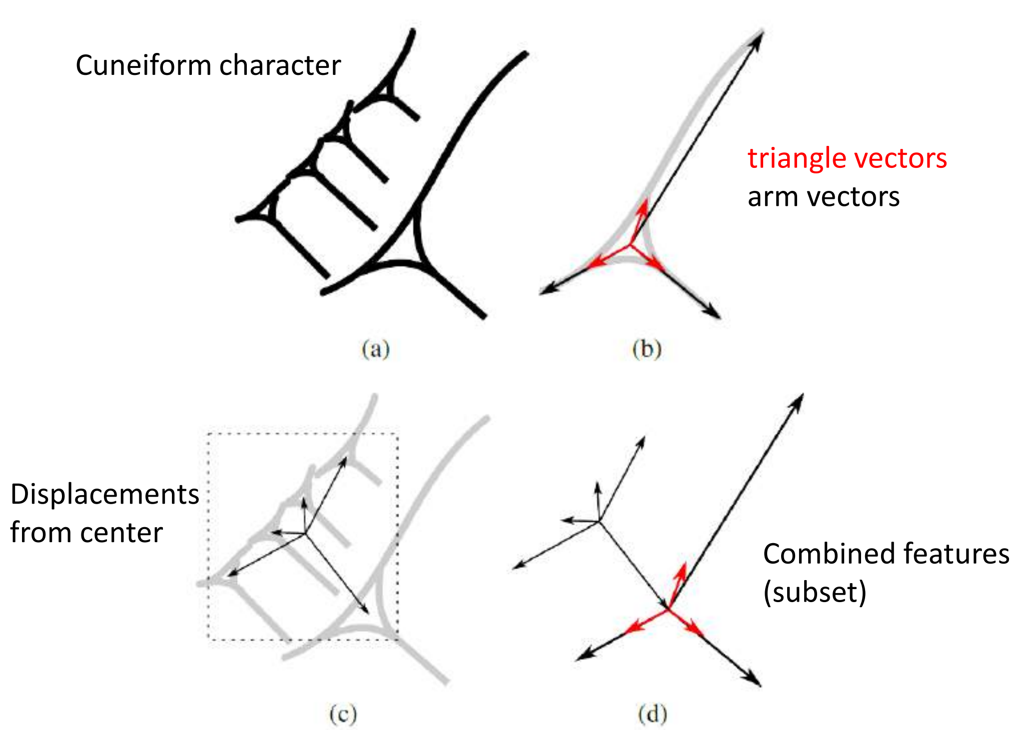

Extraction of structural features

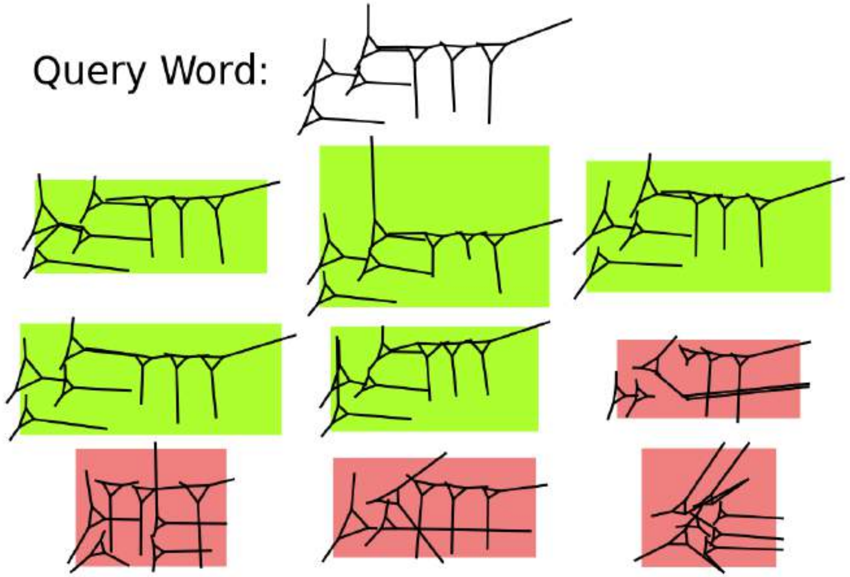

Recognition by comparison to a query word

Auden Letters: visualizing typewriter indentations

- Current project: illuminate W.H. Auden's life & work in Austra
- Based on a collection of letters, documents, etc.
- Some of them show indentations from a typewriter
- Our job: make it readable

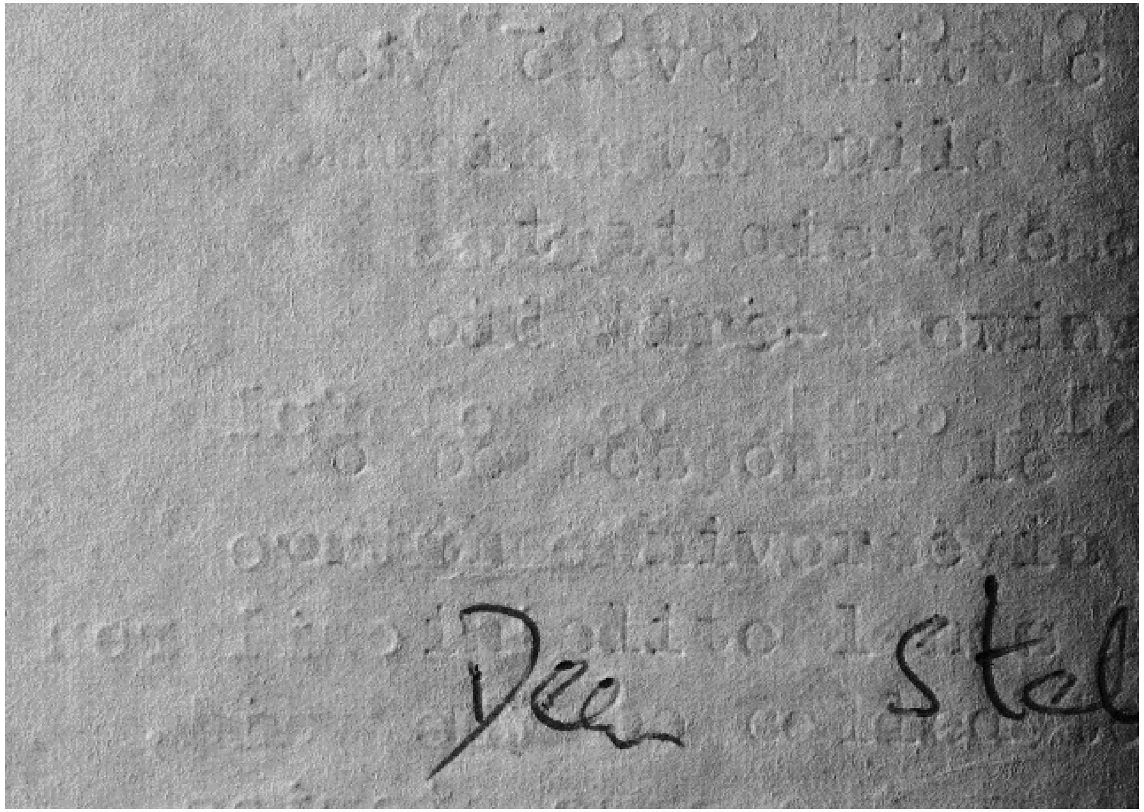

Auden Letters - Approach

1. Photometric stereo reconstruction
2. Filtering
3. Combine results in a false color image

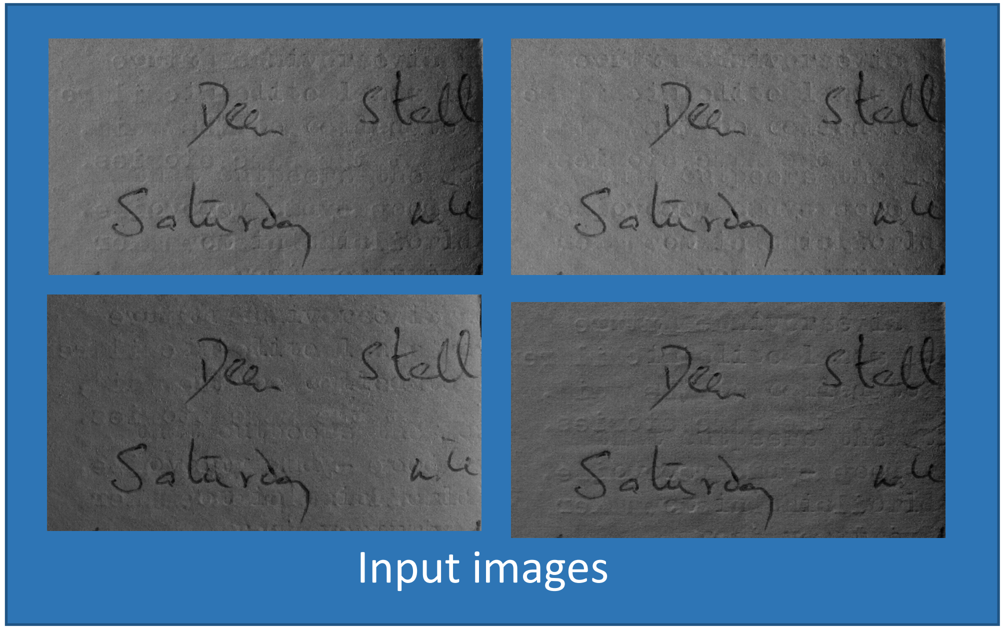

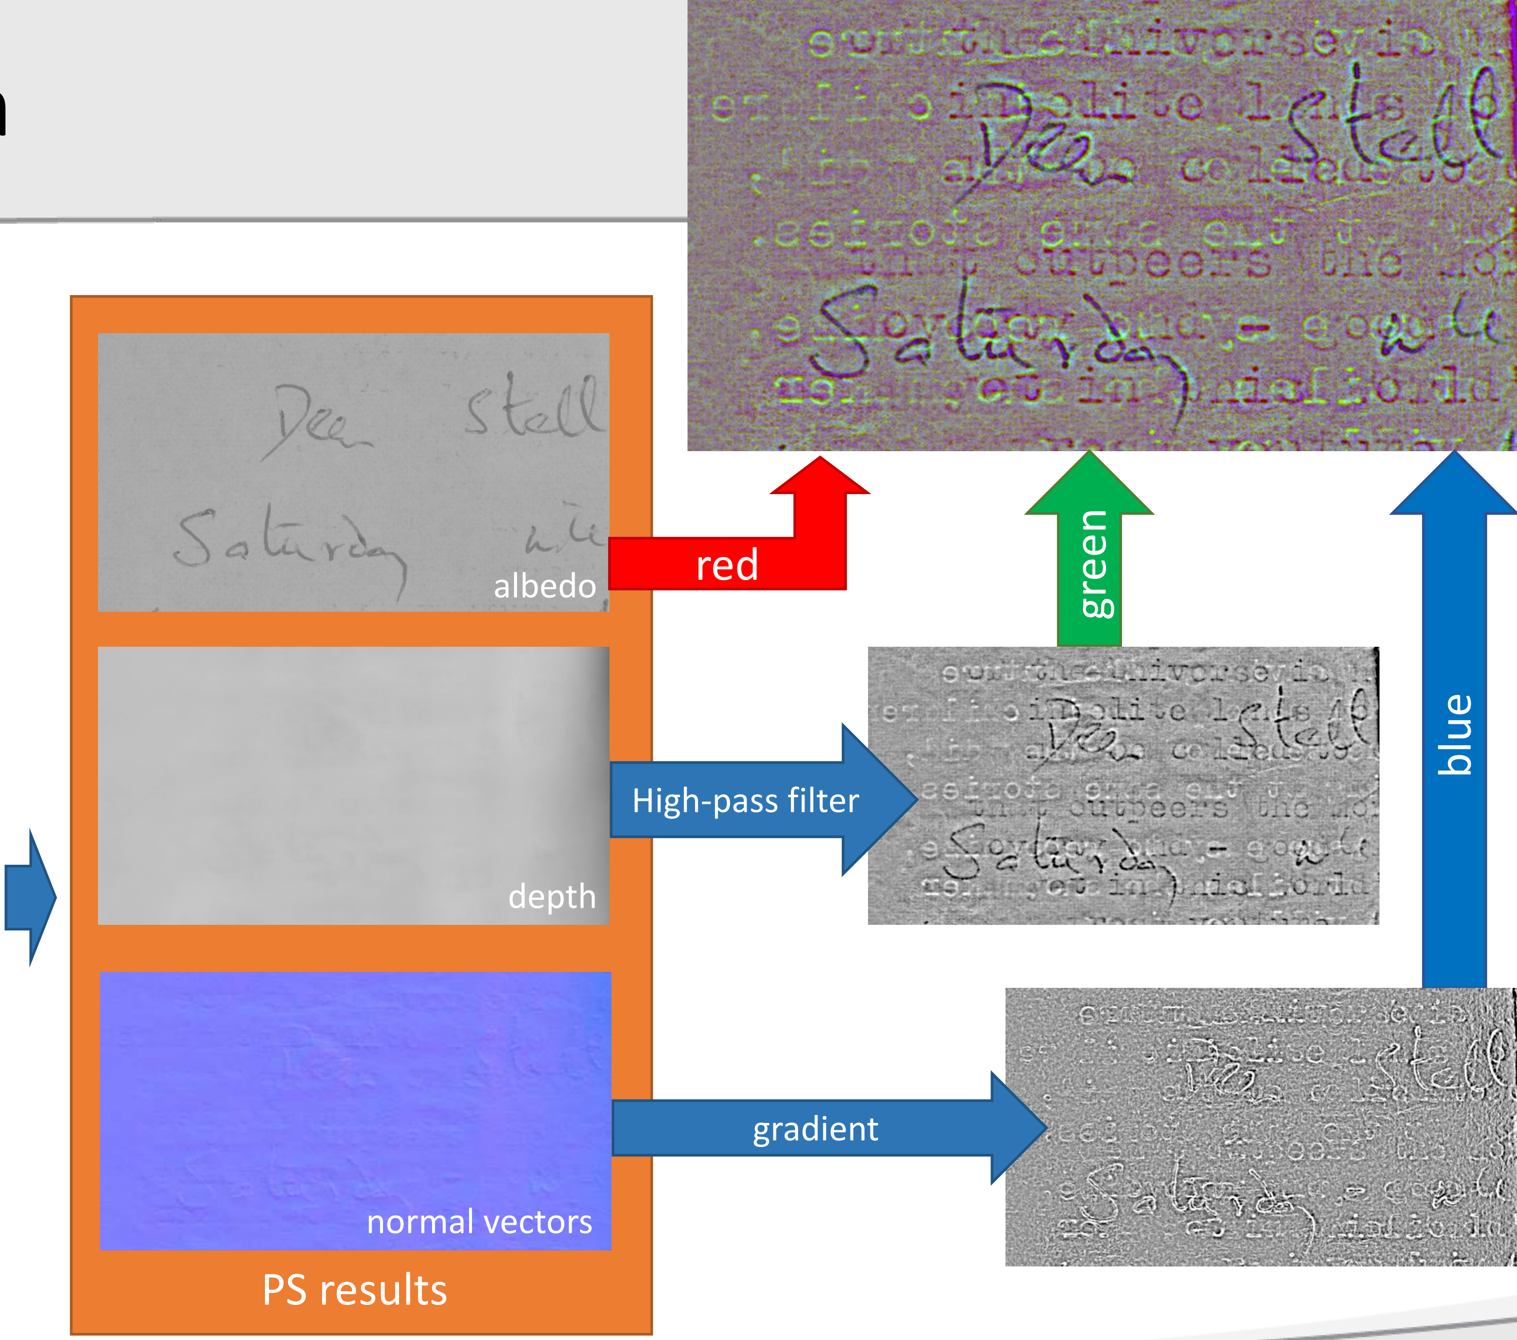

Auden Letters - Results

- Three layers of text
  - Handwritten (black)
  - Impressions from the viewer‘s direction (red)
  - Impressions from the opposite direction (green)
- An unknown intermediate version of a poem could be identified (and dated)

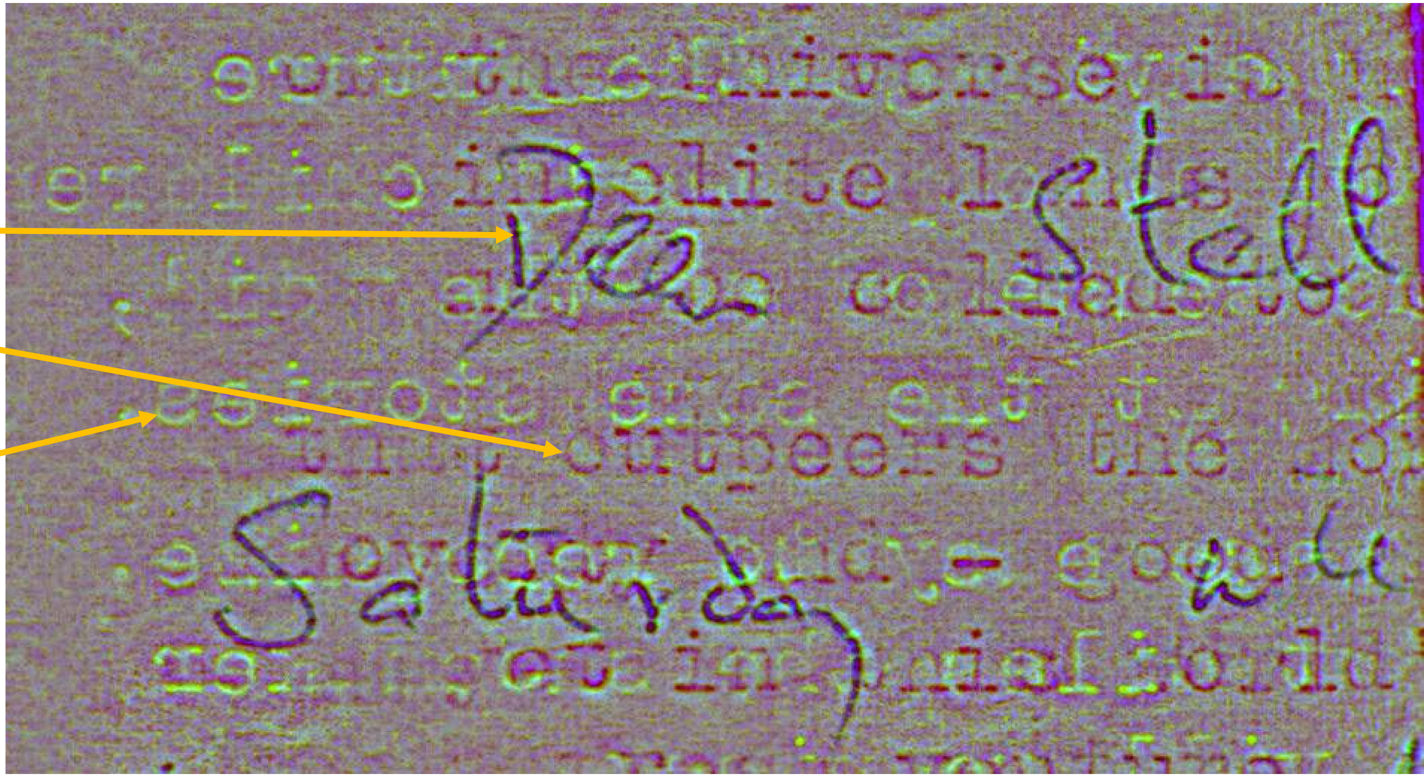

Question time