3D Reconstruction

CS 6384 Computer Vision

Professor Yu Xiang

The University of Texas at Dallas

3D Reconstruction

- How to obtain 3D models of objects or scenes?
  - Stereo matching
  - SfM and SLAM
  - 3D scanning
  - Multi-view stereo

3D Reconstruction

Triangulation-based 3D Scanner

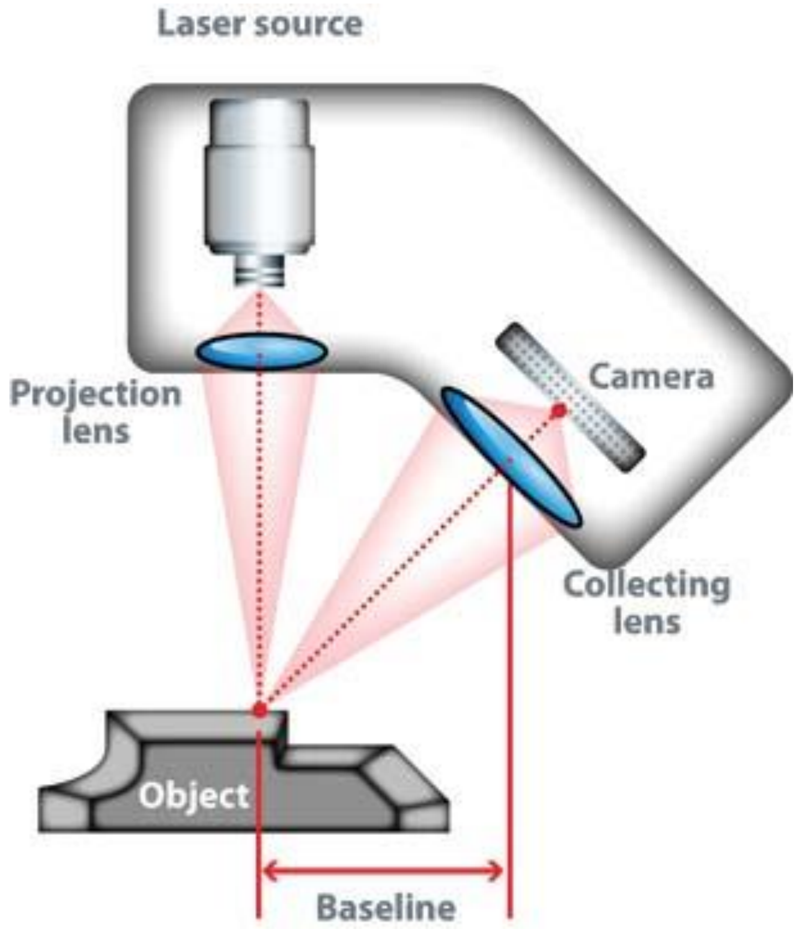

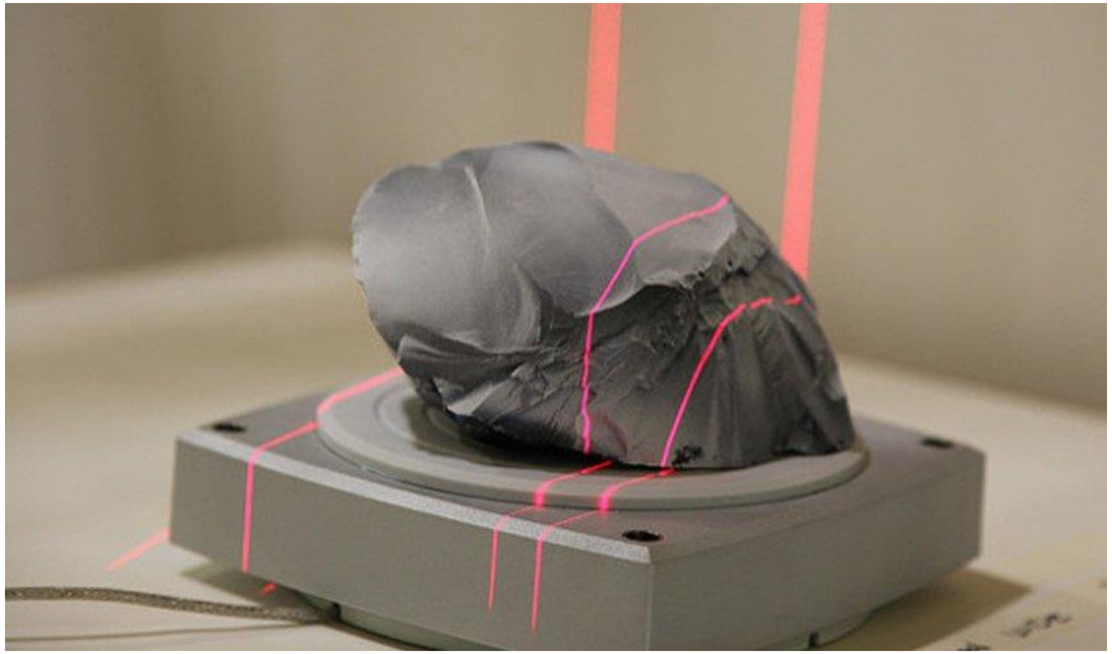

https://3dscanningservices.net/blog/need-know-3d-scanning/

Triangulation-based 3D Scanner

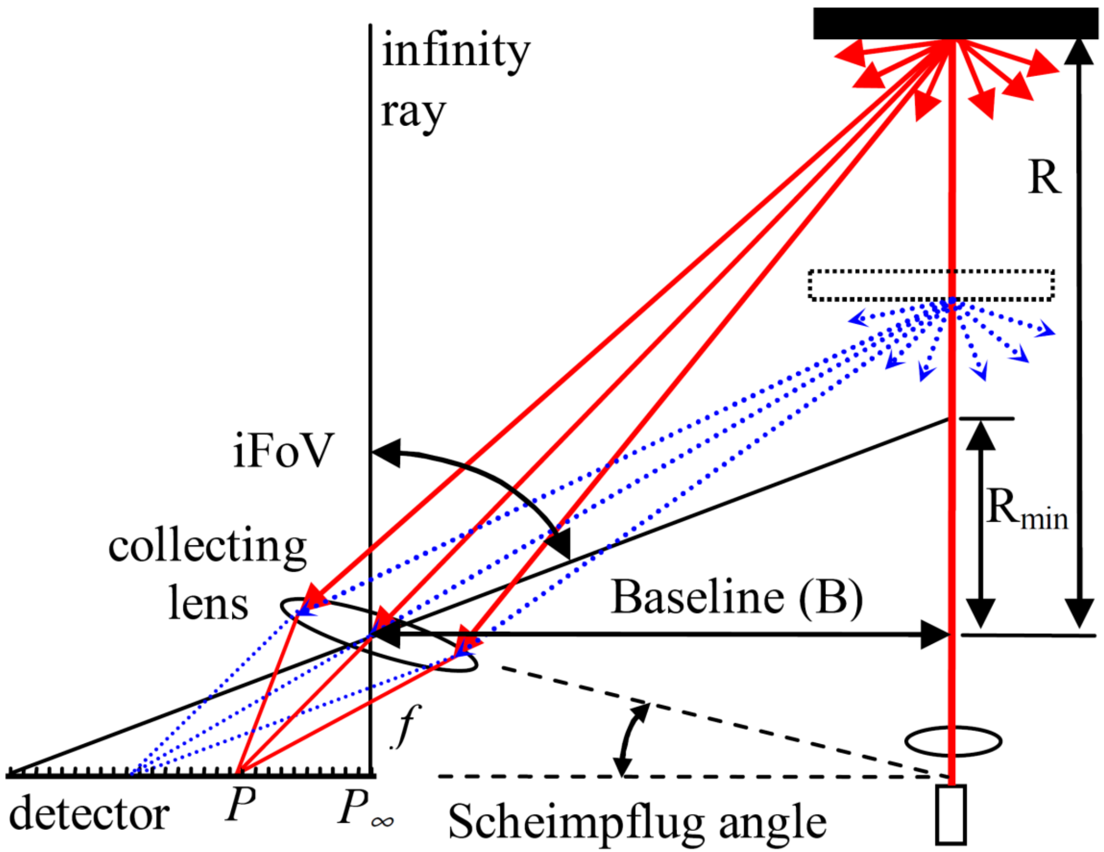

$$R = \frac{Bf}{P_\infty - P}$$

The complementary nature of triangulation and ladar technologies. Chad English, SPIE'05

Triangulation-based 3D Scanner

- Digital Michelangelo Project (1990)

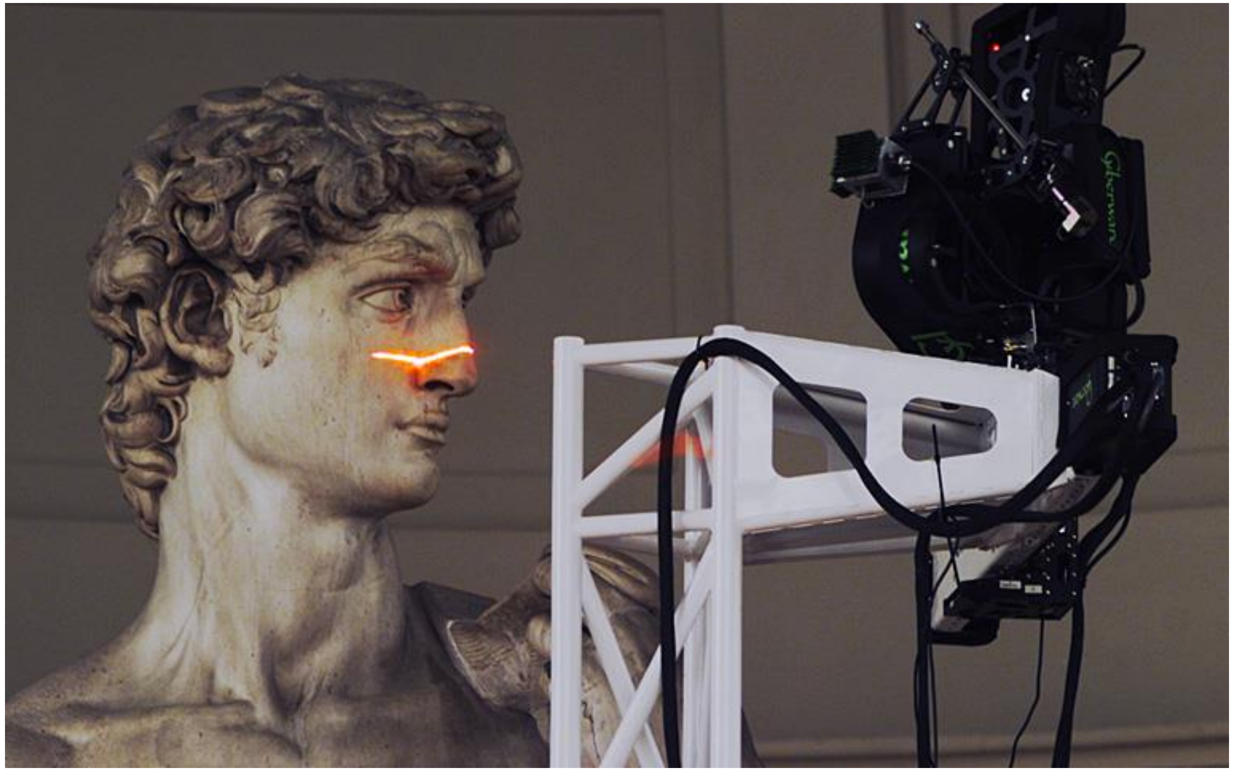

https://accademia.stanford.edu/mich/

Structured Light 3D Scanner

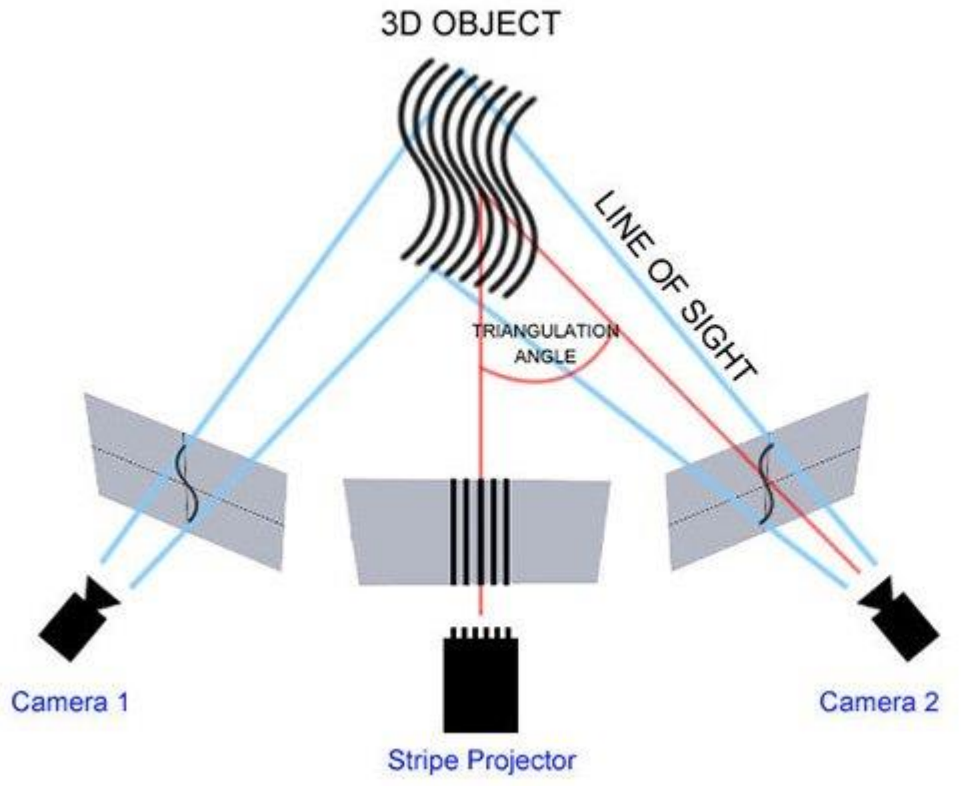

https://www.3dnatives.com/en/laser-3d-scanner-vs-structured-light-3d-scanner-080820194/

Microsoft Kinect 1

- Structured light infrared (IR)

IR stereo

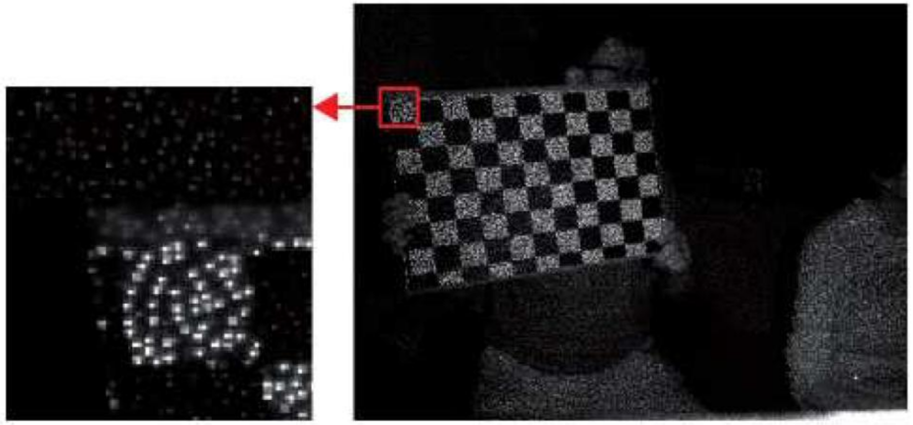

infrared (IR) speckle pattern

Time-of-flight 3D Scanner

- Long range 3D scan
  - light detection and ranging (LiDAR)

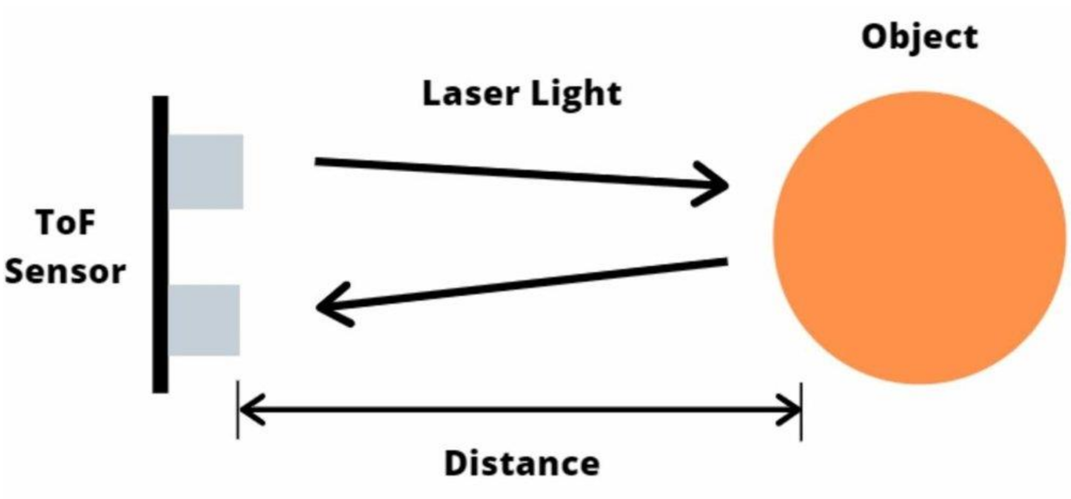

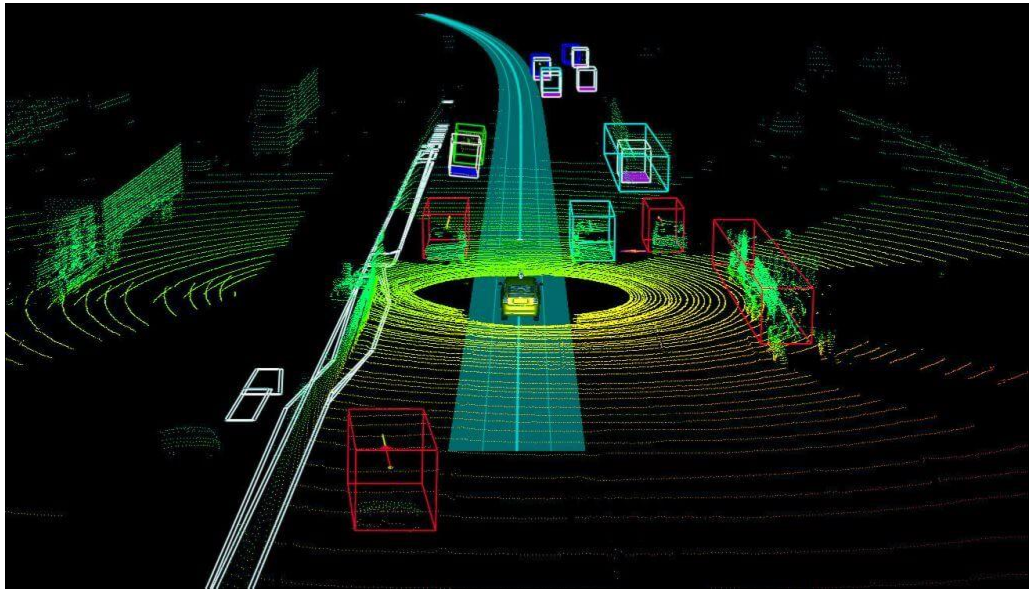

https://all3dp.com/2/tof-sensors-time-of-flight/

Microsoft Kinect 2 and Azure

- Time-of-flight infrared (IR)

The value of pixels in the clean IR reading is proportional to the amount of light returned from the scene.

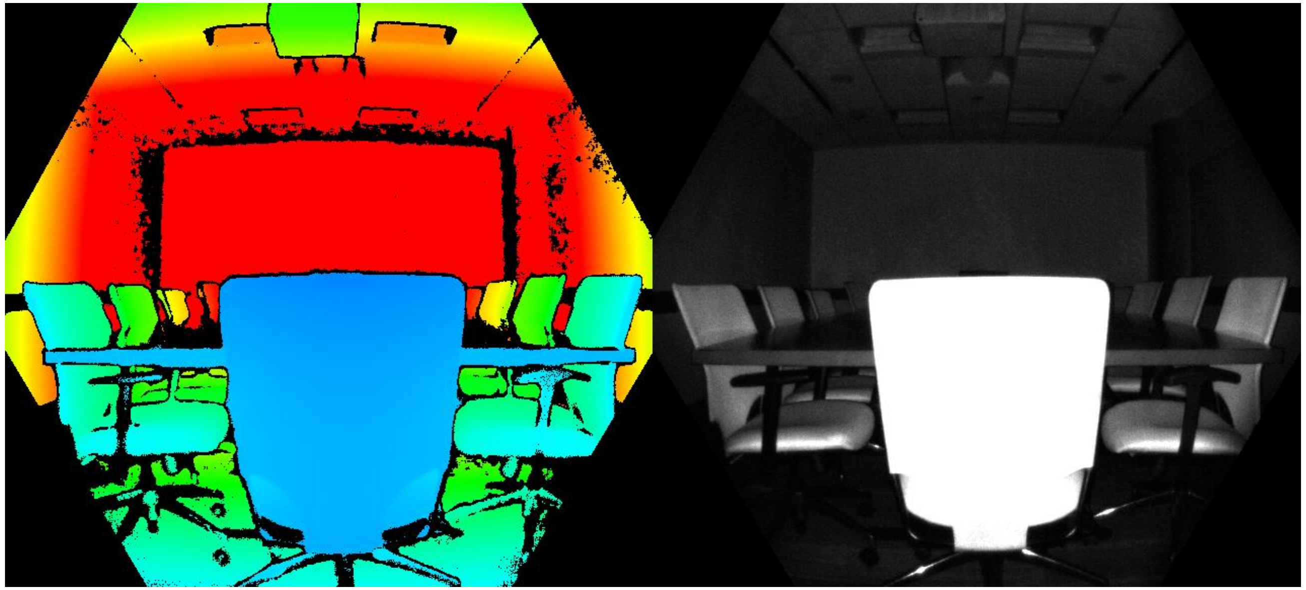

Depth Image

IR Image

https://docs.microsoft.com/en-us/azure/kinect-dk/depth-camera

Range Data Merging

- Each scan/capture generates a depth image or a point cloud
- How can we combine these data into a 3D model?
  - Alignment/registration
    - E.g., iterative closest point (ICP) algorithm
  - Merging

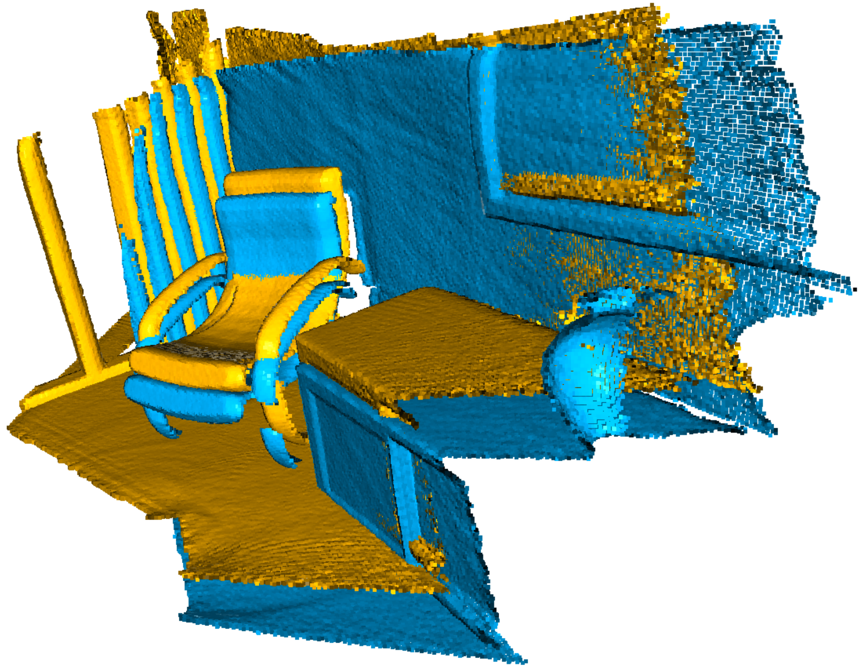

http://www.open3d.org/docs/latest/tutorial/Basic/icp_registration.html

Volumetric Integration

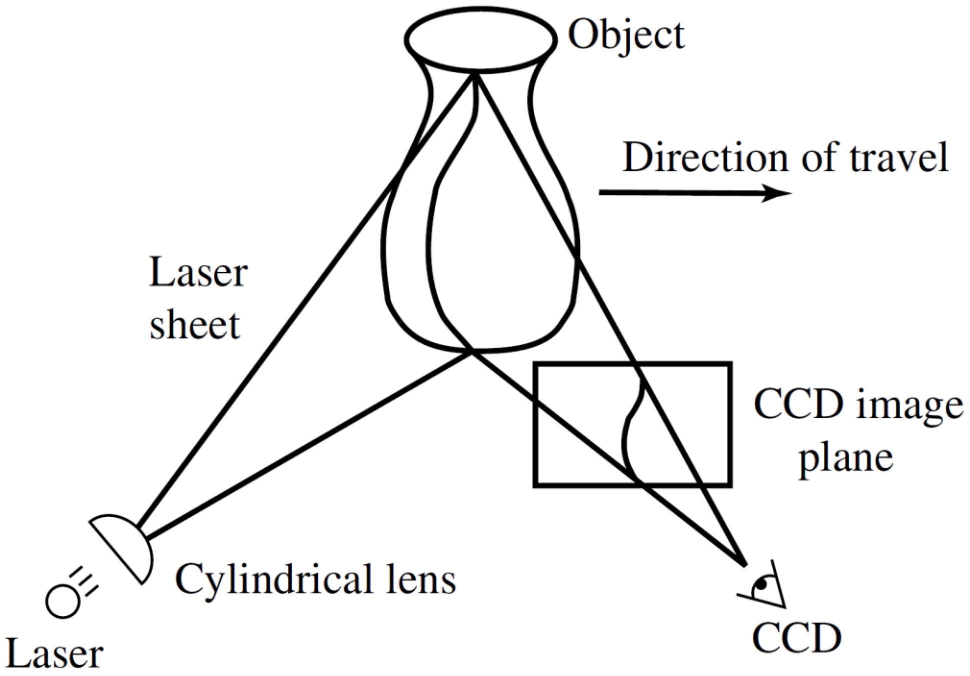

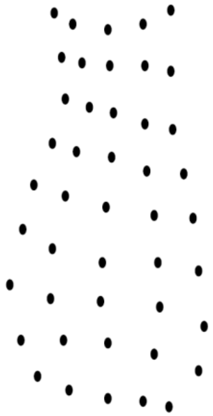

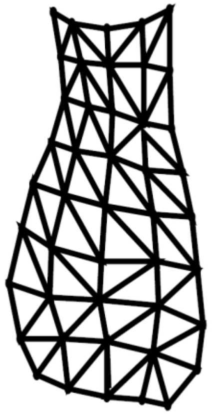

A Volumetric Method for Building Complex Models from Range Images. Curless & Levoy. SIGGRAPH'96.

Volumetric Integration

- Signed Distance Function (SDF)

$$\phi : \Omega \subseteq \mathbb{R}^3 \rightarrow \mathbb{R}$$

Signed distance to the closest object boundary

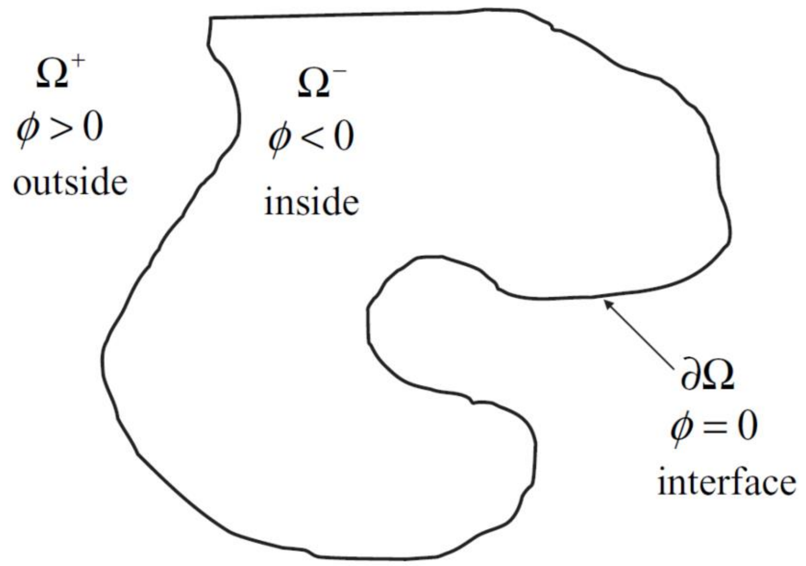

Volumetric Integration

- Signed Distance Function (SDF)

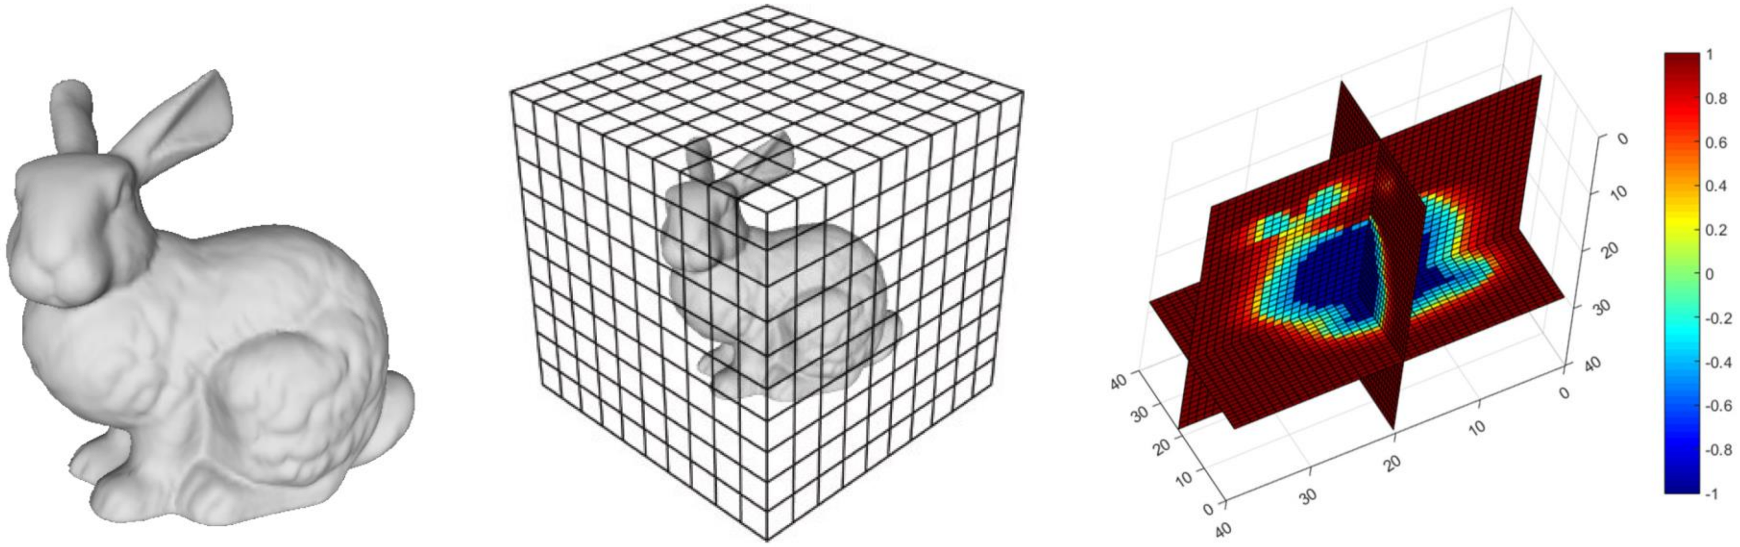

(a) Surface view. (b) Bounding volume. (c) Generated SDF.

Signed Distance Fields for Rigid and Deformable 3D Reconstruction. Miroslava Slavcheva.

Volumetric Integration

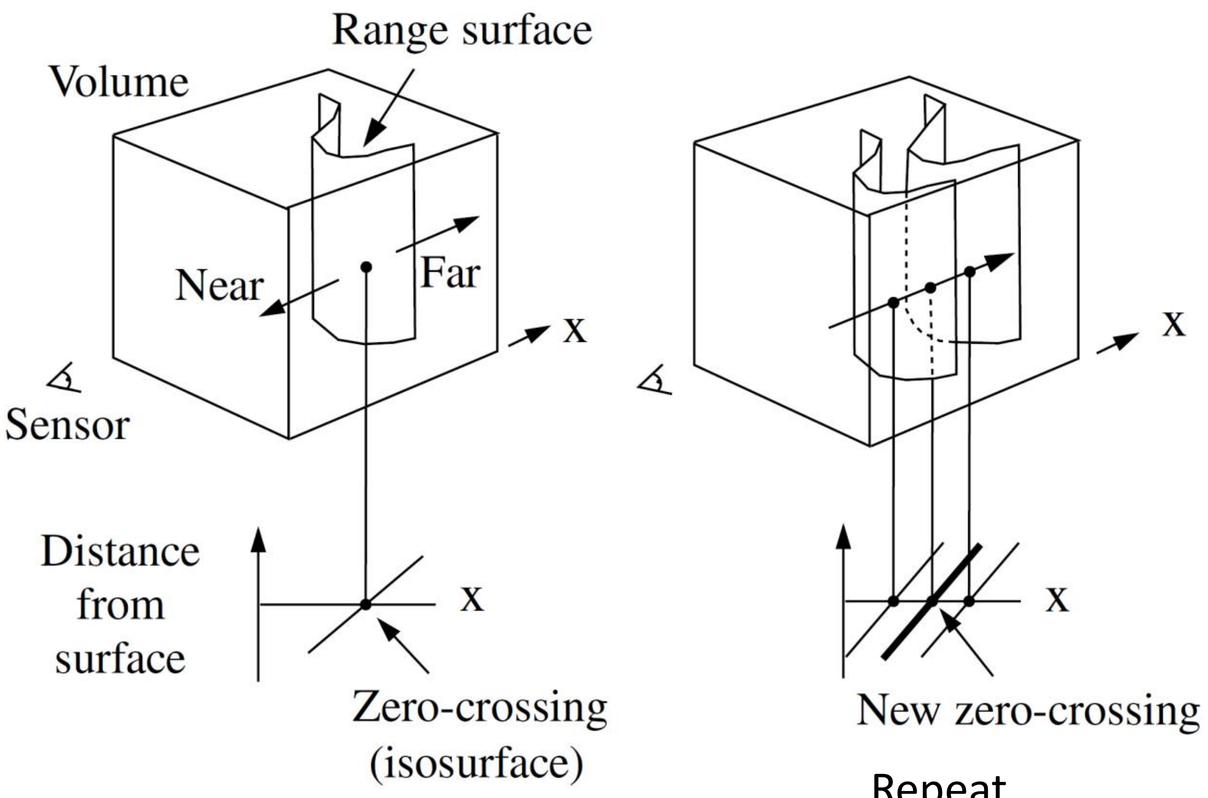

Repeat

SDF for the range image

$$D_{i+1}(\mathbf{x}) = \frac{W_i(\mathbf{x})D_i(\mathbf{x}) + w_{i+1}(\mathbf{x})d_{i+1}(\mathbf{x})}{W_i(\mathbf{x}) + w_{i+1}(\mathbf{x})}$$

$$W_{i+1}(\mathbf{x}) = W_i(\mathbf{x}) + w_{i+1}(\mathbf{x})$$

Weight function

A Volumetric Method for Building Complex Models from Range Images. Curless & Levoy. SIGGRAPH'96.

Volumetric Integration

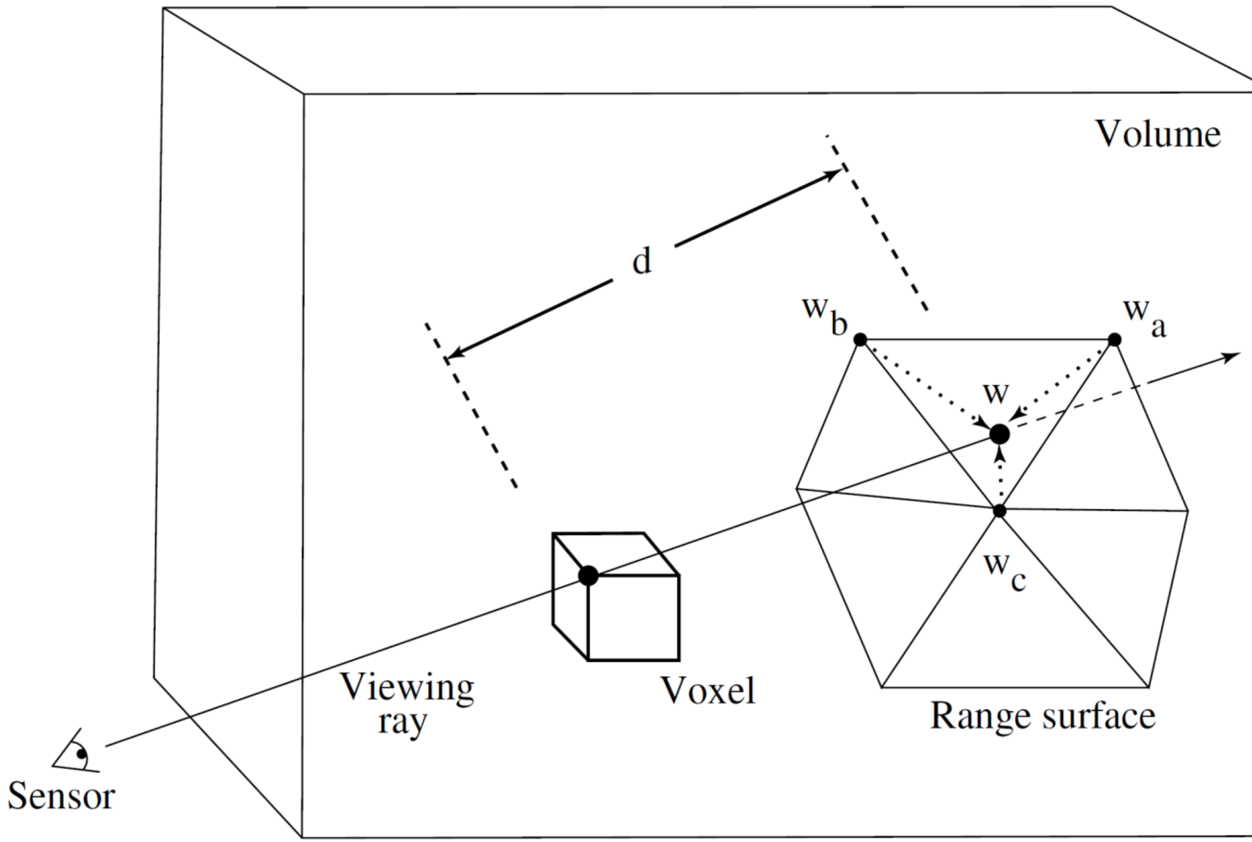

- Tessellate the range image into a triangle mesh
- The vertex weight depends on the dot product between each vertex normal and the viewing direction.
- Linearly interpolating the weights

We can fuse color (RGB) in a similar way.

A Volumetric Method for Building Complex Models from Range Images. Curless & Levoy. SIGGRAPH'96.

Volumetric Integration

Image Single scan Merged scan

A Volumetric Method for Building Complex Models from Range Images. Curless & Levoy. SIGGRAPH'96.

KinectFusion

Single scan

Rendered normal map

Rendered 3D model

Image-based 3D Reconstruction

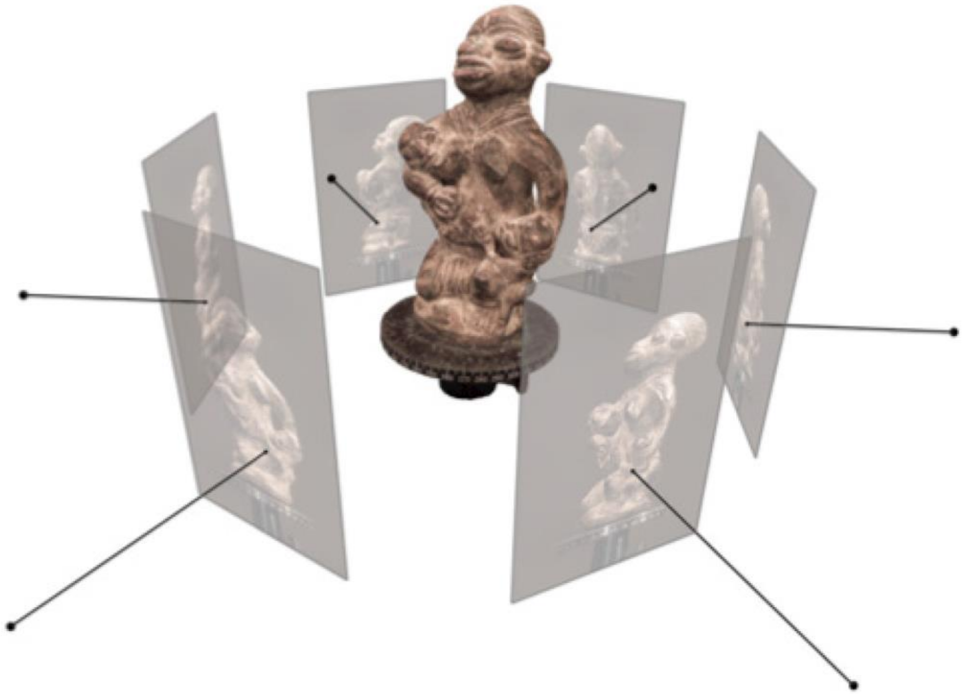

A set of images

3D model

Multi-View Stereo: A Tutorial. Yasutaka Furukawa and Carlos Hernández

Multi-view Stereo

- Image-based 3D reconstruction techniques
  - Use stereo correspondences as the main cue
  - Use more than two images

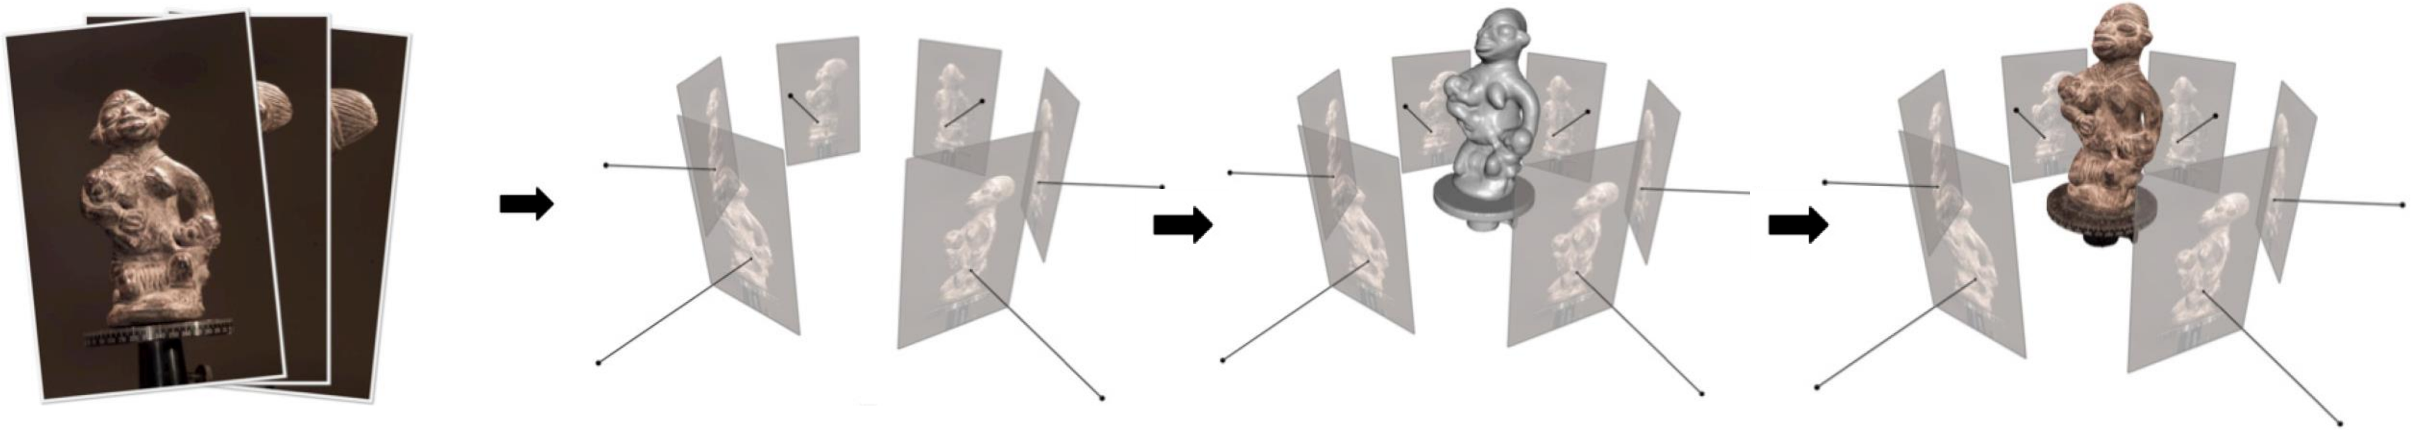

A set of images → Camera parameters → 3D geometry → Materials

Multi-View Stereo: A Tutorial. Yasutaka Furukawa and Carlos Hernández

Photo-consistency

- The projections of a 3D point to multi-view images should be consistency

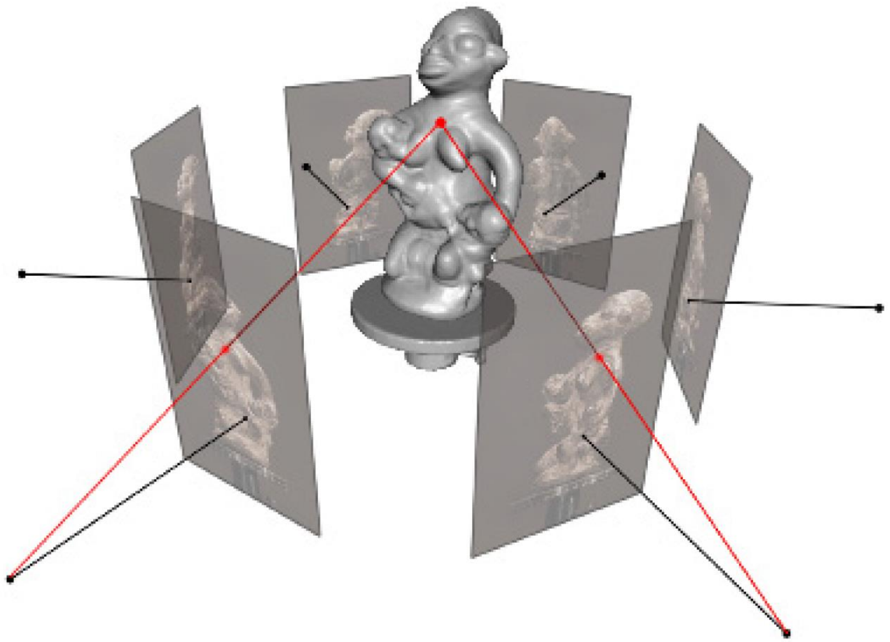

$$\mathcal{C}_{ij}(p) = \rho(I_i(\Omega(\pi_i(p))), I_j(\Omega(\pi_j(p))))$$

Similarity measure

Projection

3D point

Support domain

Photo-consistency

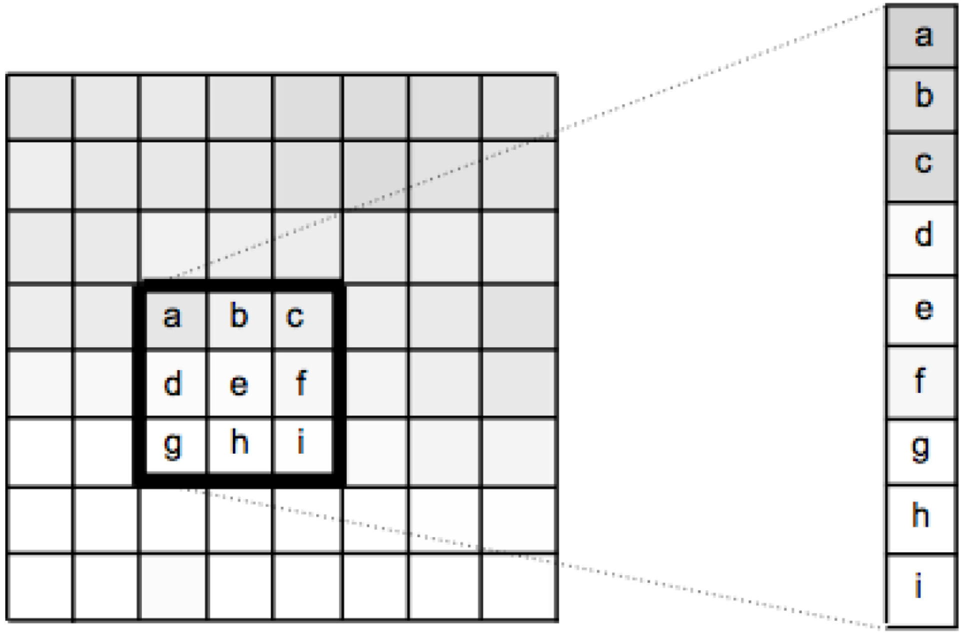

Support domain

Similarity measurement

- Normalized Cross Correlation $\sigma \equiv \sqrt{\mathrm{E}[(X-\mu)^2]}$

$$\rho_{NCC}(f,g) = \frac{(f-\bar{f})\cdot(g-\bar{g})}{\sigma_f \sigma_g} \in [-1,1]$$

- Sum of Squared Differences

$$\rho_{SSD}(f,g) = ||f-g||^2$$

- Sum of Absolute Differences

$$\rho_{SAD}(f,g) = ||f-g||_1$$

Photo-consistency

- How to use photo-consistency?
  - Estimate 3D points (geometry) that maximize photo-consistency

- The visibility problem: which points are visible in which images?

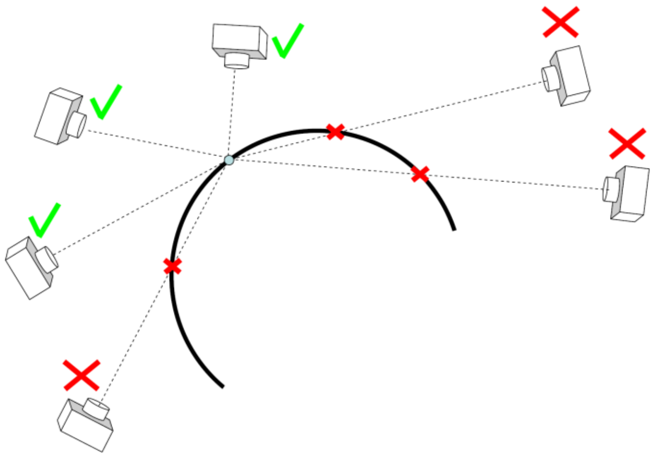

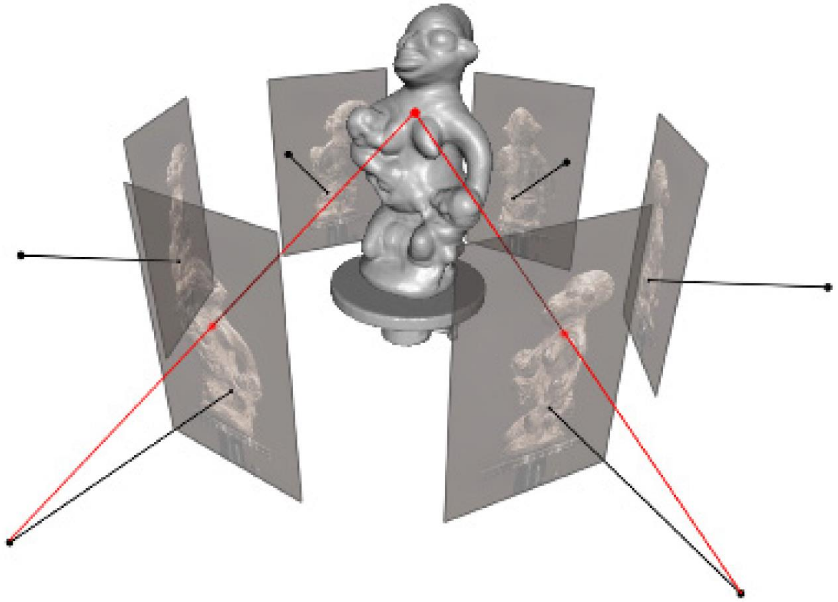

Space Carving from Silhouettes

- Silhouettes

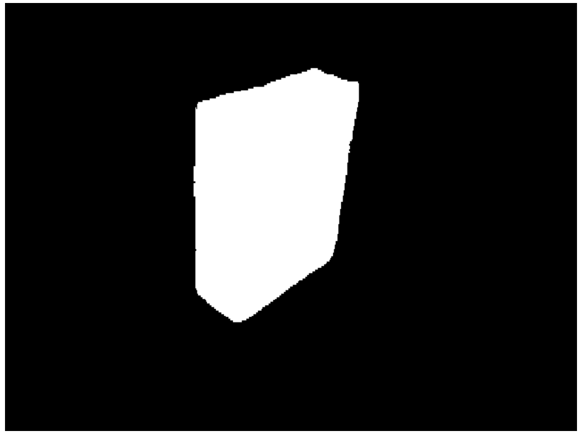


Space Carving from Silhouettes

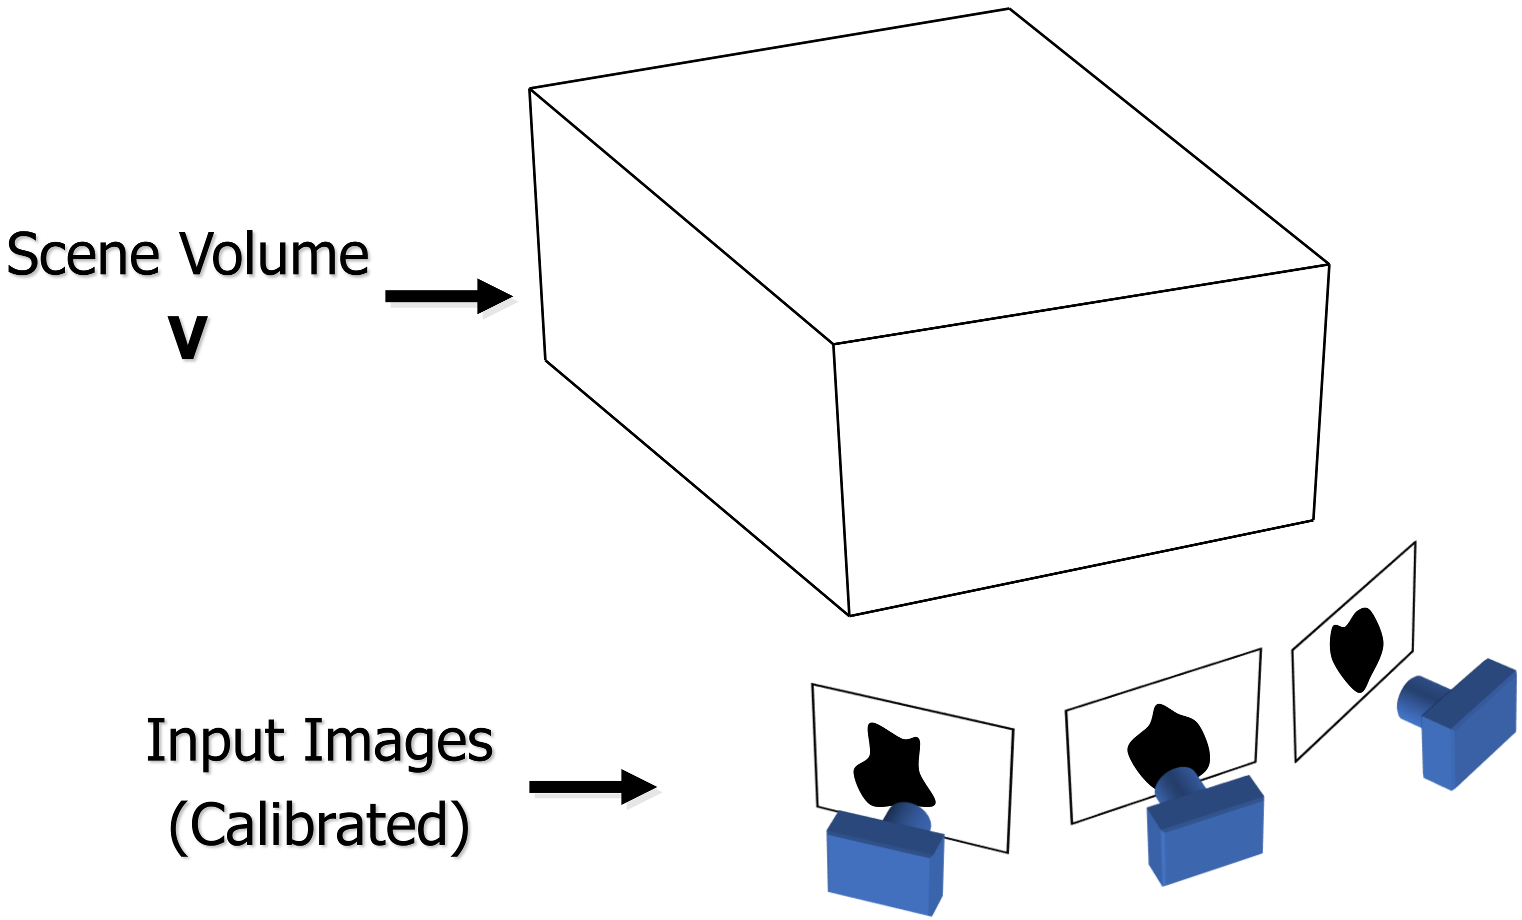

Credit: Steve Seitz

Space Carving from Silhouettes

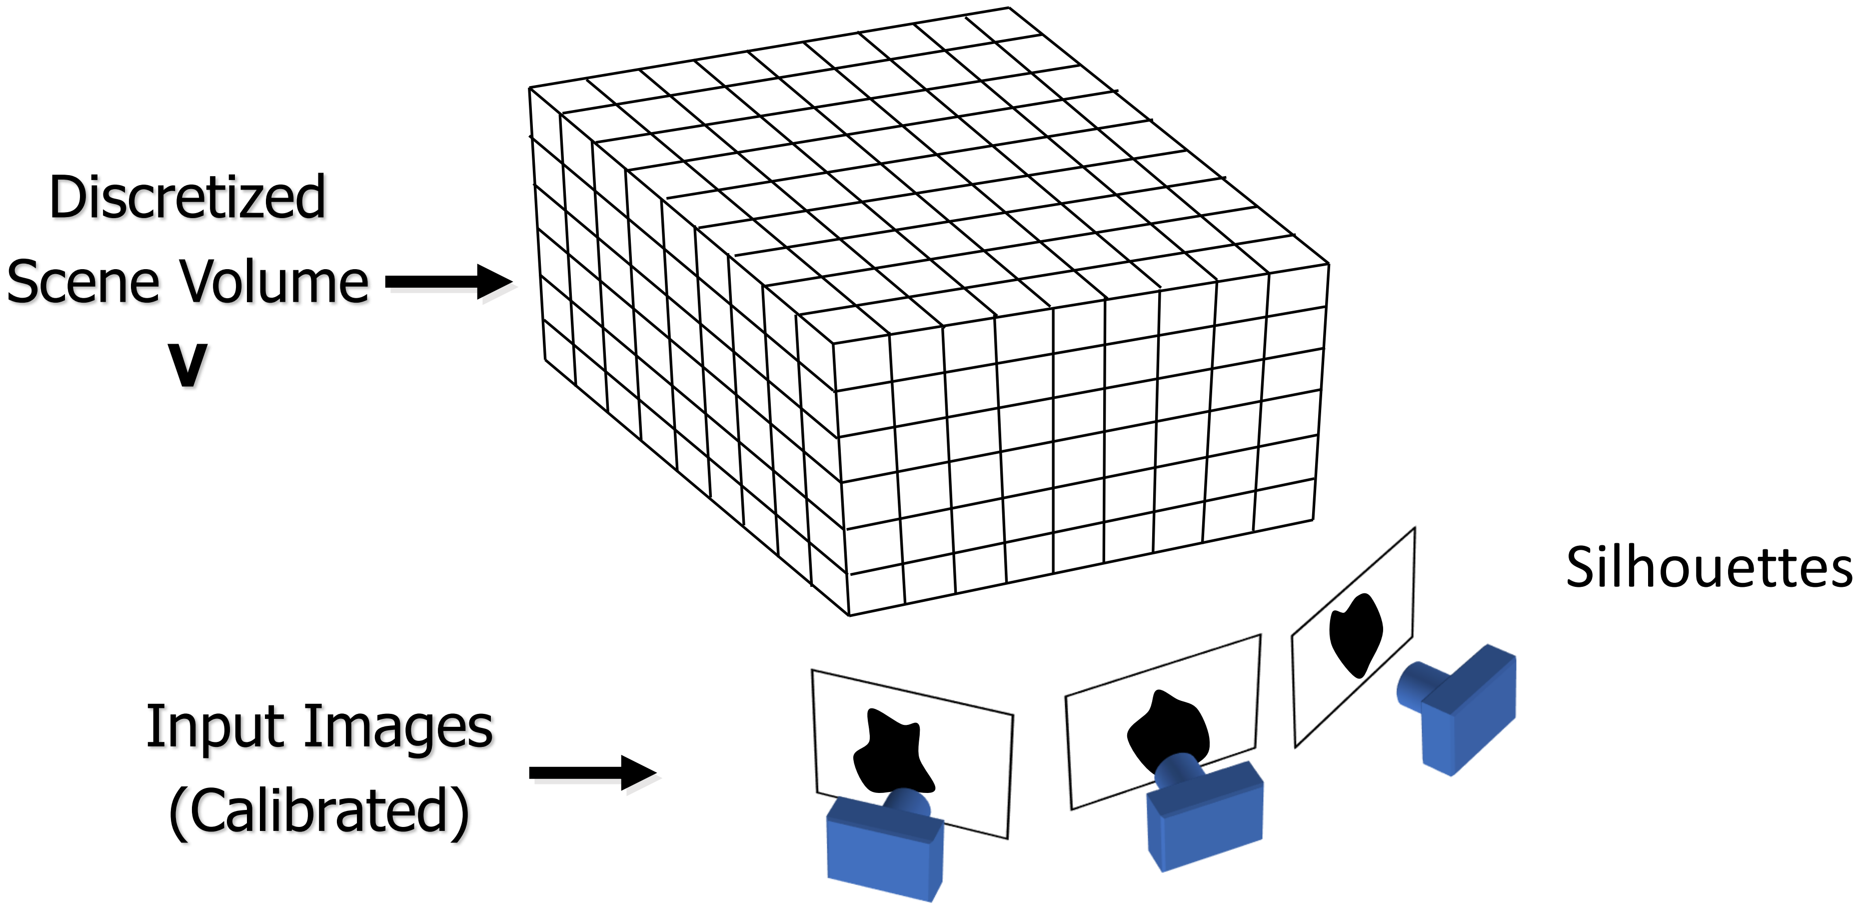

Credit: Steve Seitz

Space Carving from Silhouettes

- Color voxel black if it is projected on silhouette in every image

- Photo-consistency?

$$C_{ij}(p) = \rho(I_i(\Omega(\pi_i(p))), I_j(\Omega(\pi_j(p))))$$

- Binary comparison

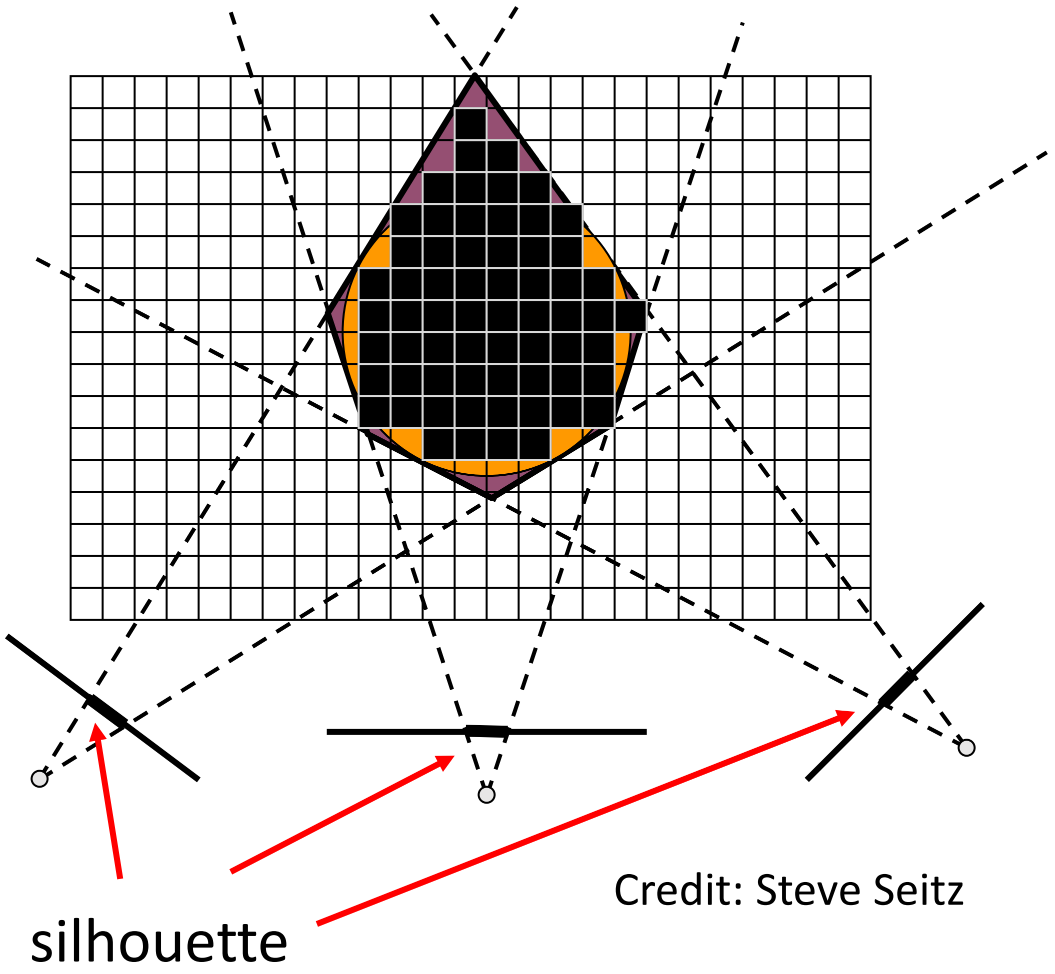

Credit: Steve Seitz

Space Carving from Silhouettes

- What if a voxel is occluded?
  - Visibility problem

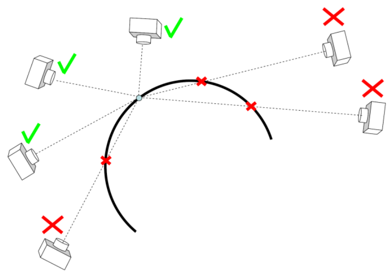

- The voxel still projects to the silhouette

- No need to check for occlusion in this case

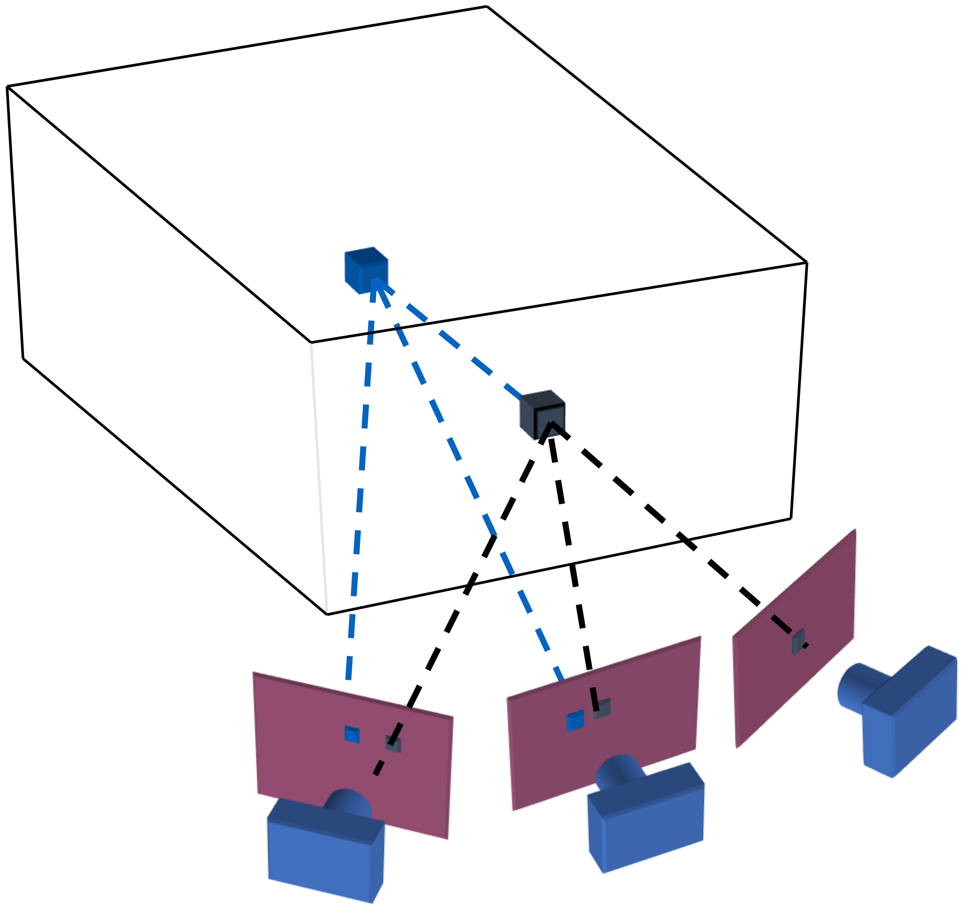

3D Reconstruction with Voxel Coloring

- Assign colors (RGB) to voxels
  1. Choose a voxel
  2. Project to each image and compute photo-consistency using colors
  3. Color the voxel if consistent

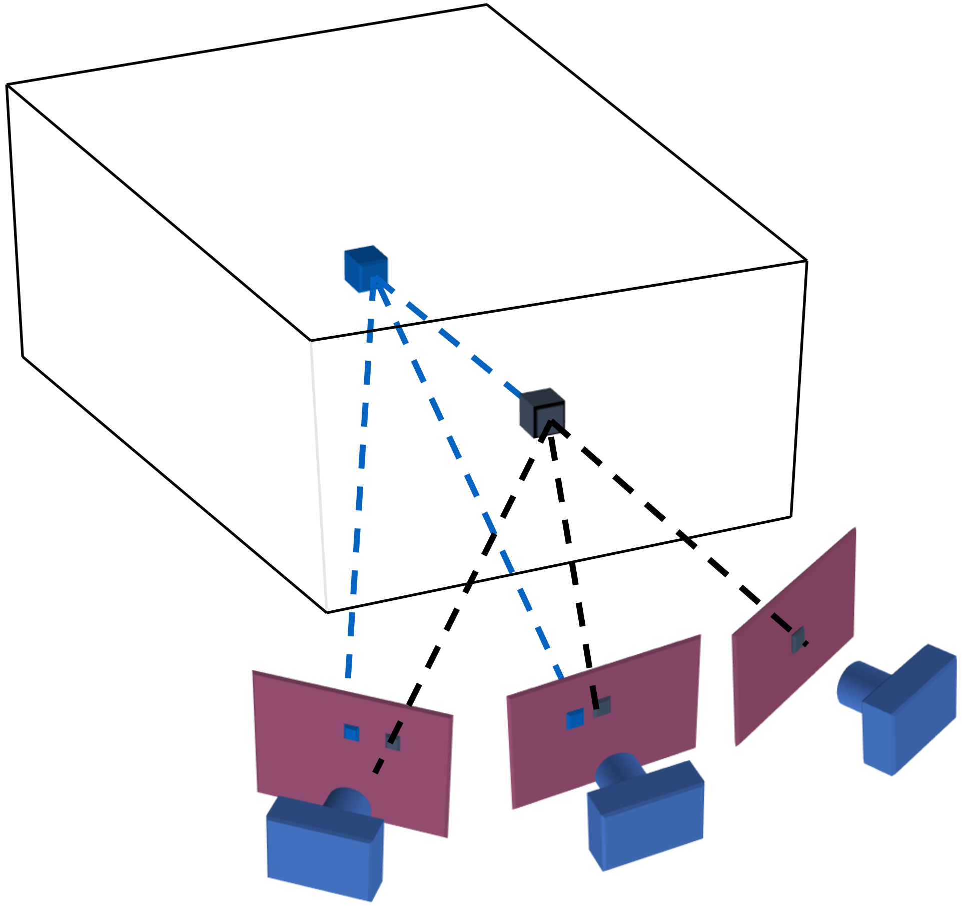

Visibility Problem: in which images is each voxel visible?

Seitz & Dyer 97

3D Reconstruction with Voxel Coloring

- Handle occlusions: visit occluders first
- Pixels will be marked if explained by visited voxels
- Only consider unmarked pixels in photo-consistency

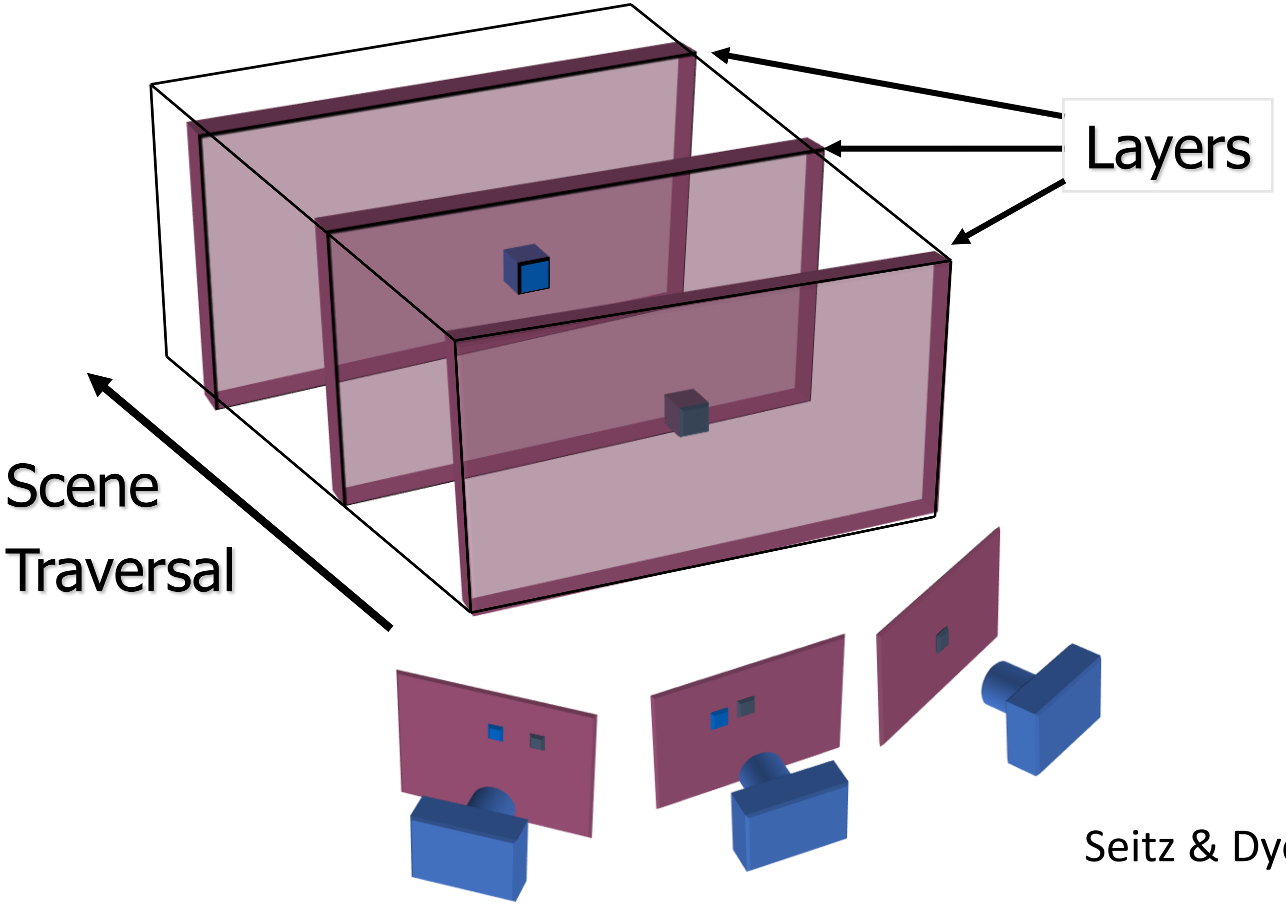

Seitz & Dyer 97

3D Reconstruction with Voxel Coloring

- Panoramic layering (inside to outside)

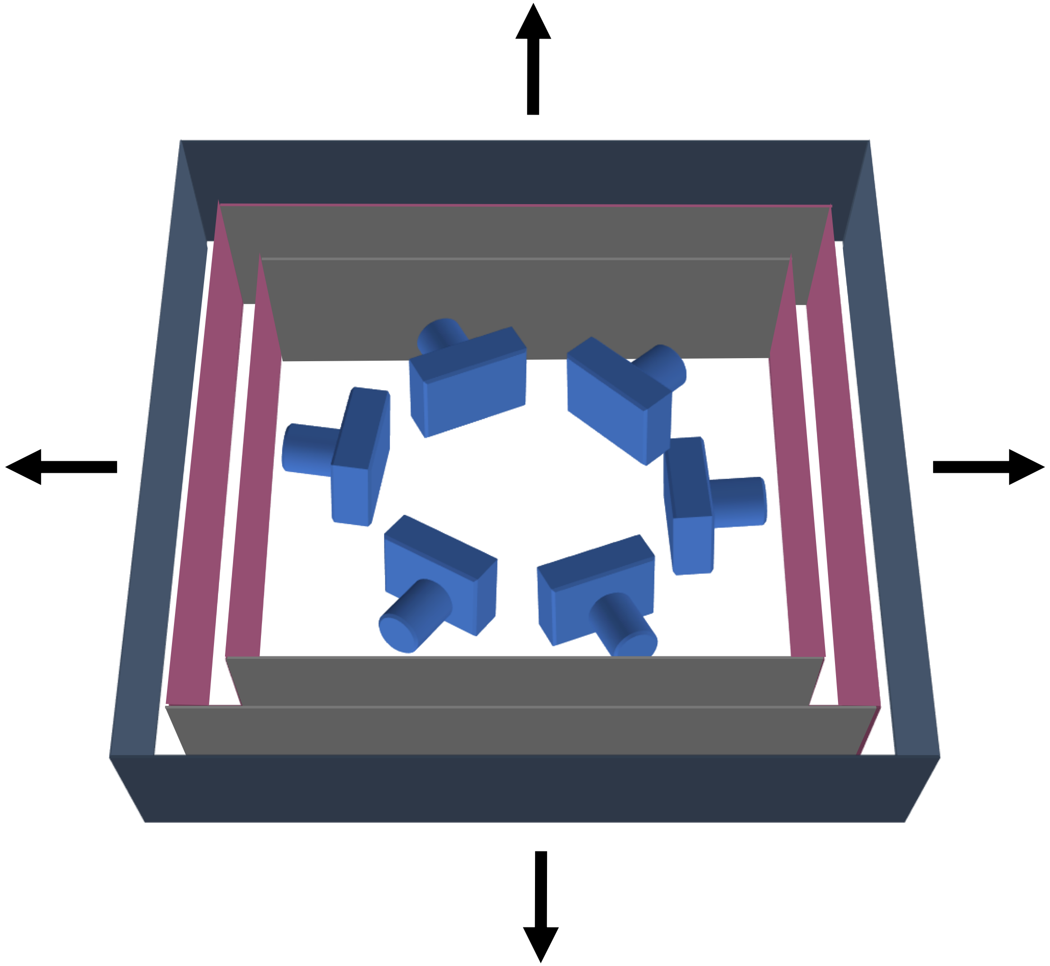

Seitz & Dyer 97

3D Reconstruction with Voxel Coloring

- *Calibrated Turntable*

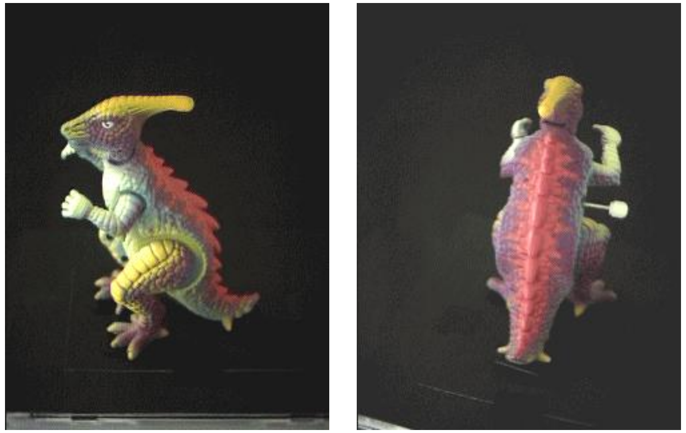

Selected Dinosaur Images

Selected Flower Images

Seitz & Dyer 97

3D Reconstruction with Voxel Coloring

Dinosaur Reconstruction
72 K voxels colored
7.6 M voxels tested
7 min. to compute
on a 250MHz SGI

Flower Reconstruction
70 K voxels colored
7.6 M voxels tested
7 min. to compute
on a 250MHz SGI

Seitz & Dyer 97

Multi-view Stereo

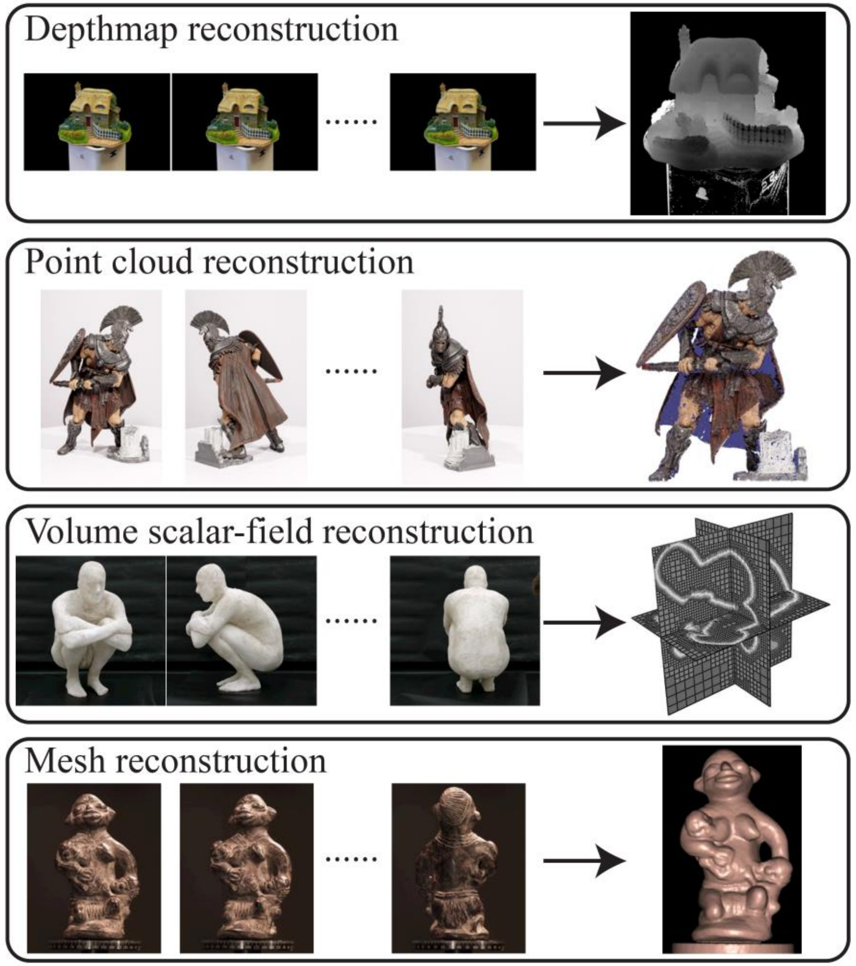

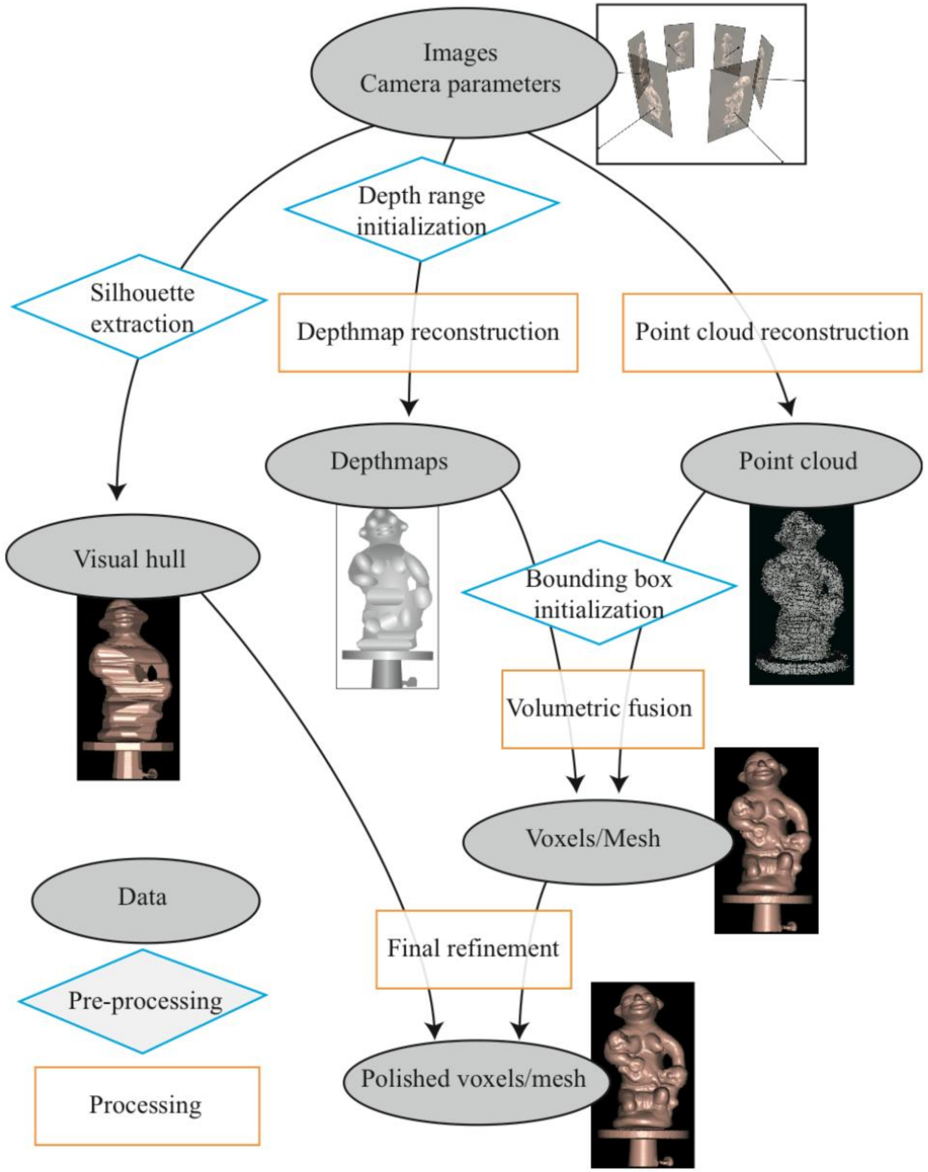

Multi-View Stereo: A Tutorial. Yasutaka Furukawa and Carlos Hernández

Further Reading

- Section 12.7, Chapter 13, Computer Vision, Richard Szeliski
- A Volumetric Method for Building Complex Models from Range Images. Curless & Levoy. SIGGRAPH'96.
- Photorealistic Scene Reconstruction by Voxel Coloring S. M. Seitz and C. R. Dyer, IJCV'99.
- Multi-View Stereo: A Tutorial. Yasutaka Furukawa and Carlos Hernández, 2015. https://carlos-hernandez.org/papers/fnt_mvs_2015.pdf