Lecture 8
SFM & Volumetric stereo

- SFM: Self-calibration
- Volumetric stereo:
  - Space carving
  - Voxel carving

Reading:
[HZ] Chapters 19 "Auto-calibration"
[Szelisky] Chapter 7 "Structure from motion"
[Szelisky] Chapter 11 "Multi-view stereo"

Structure from motion problem

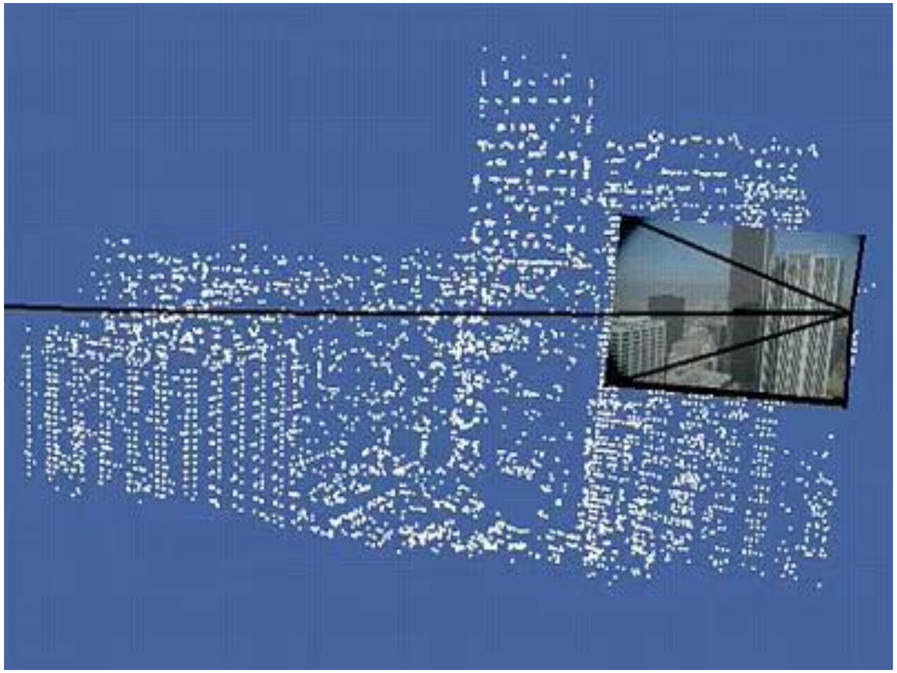

Courtesy of Oxford **Visual Geometry Group**

Structure from motion problem

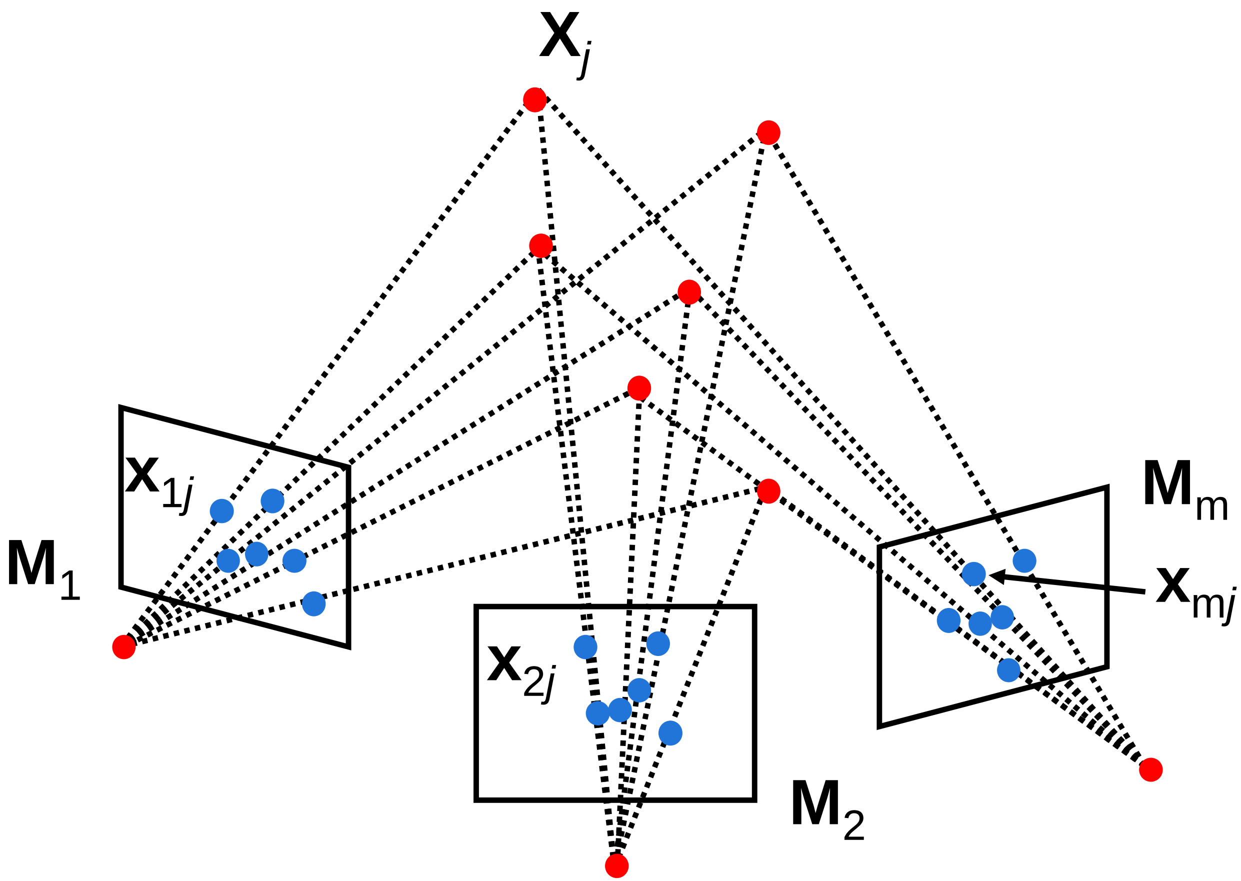

From the m×n correspondences $\mathbf{x}_{ij}$, estimate:

- m projection matrices $\mathbf{M}_i$ — motion
- n 3D points $\mathbf{X}_j$ — structure

Projective Ambiguity

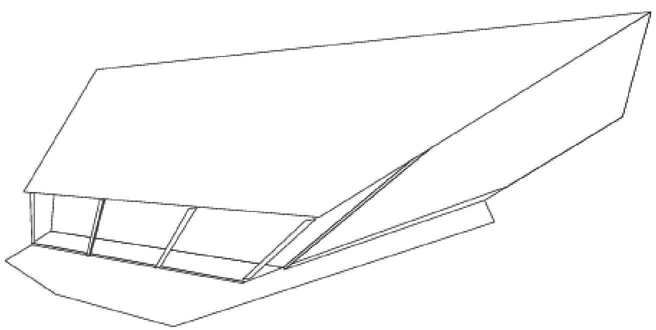

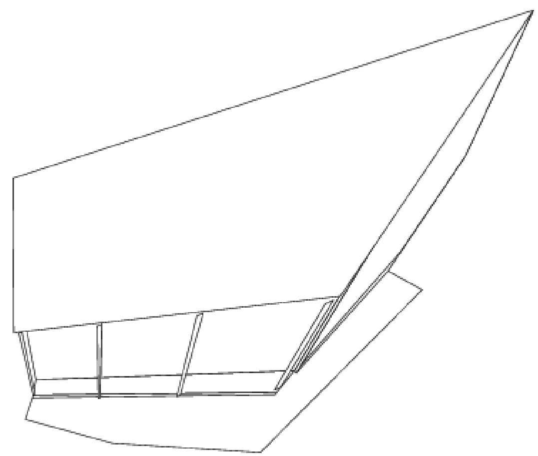

R. Hartley and A. Zisserman, Multiple View Geometry in Computer Vision, 2nd edition, 2003

Metric reconstruction (upgrade)

- The problem of recovering the metric reconstruction from the perspective one is called self-calibration

- Stratified reconstruction:
  - from perspective to affine
  - from affine to metric

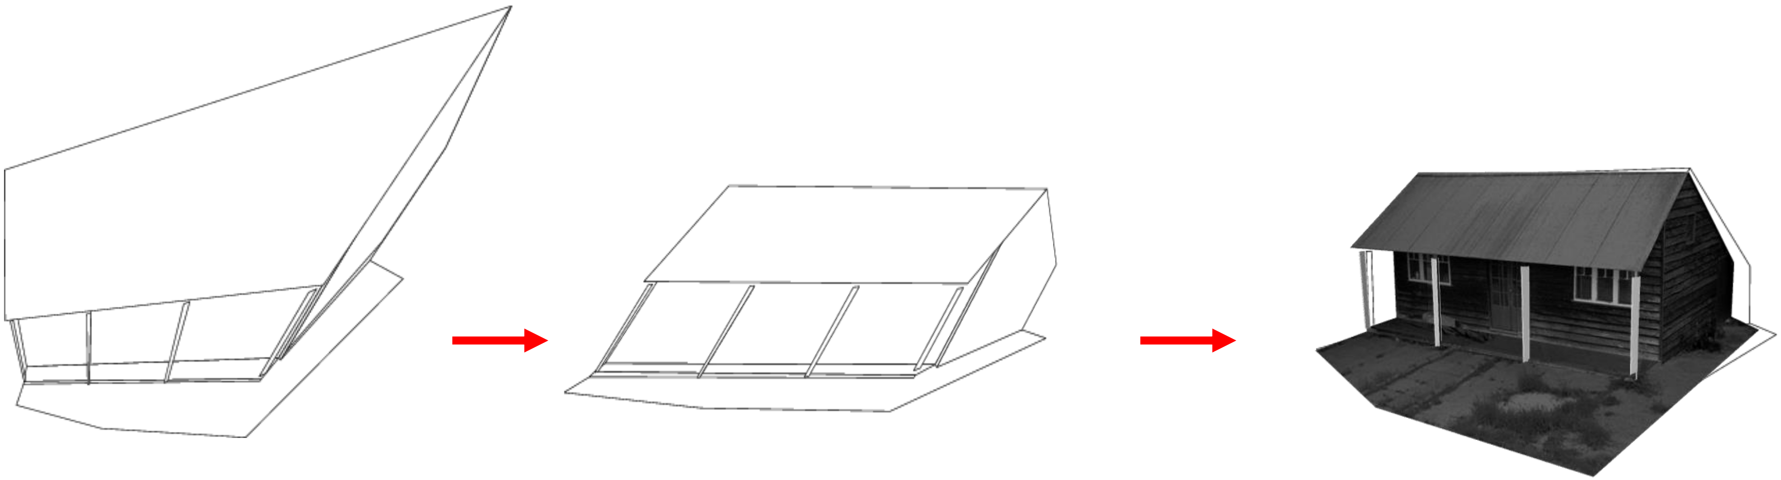

SFM problem - summary

1. Estimate structure and motion up perspective transformation
    1. Algebraic
    2. factorization method
    3. bundle adjustment

2. Convert from perspective to metric (self-calibration)

3. Bundle adjustment

** or **

1. Bundle adjustment with self-calibration constraints

Self-calibration

[HZ] Chapters 19 "Auto-calibration"

Several approaches:

- Use single-view metrology constraints (lecture 4)
- Direct approach (Kruppa Eqs) for 2 views
- Algebraic approach
- Stratified approach

Direct approach

We use the following results:

1. A relationship that maps conics across views
2. Concept of absolute conic and its relationship to K
3. The Kruppa equations

Projections of conics across views

$$x^T C_w \; x = 0$$

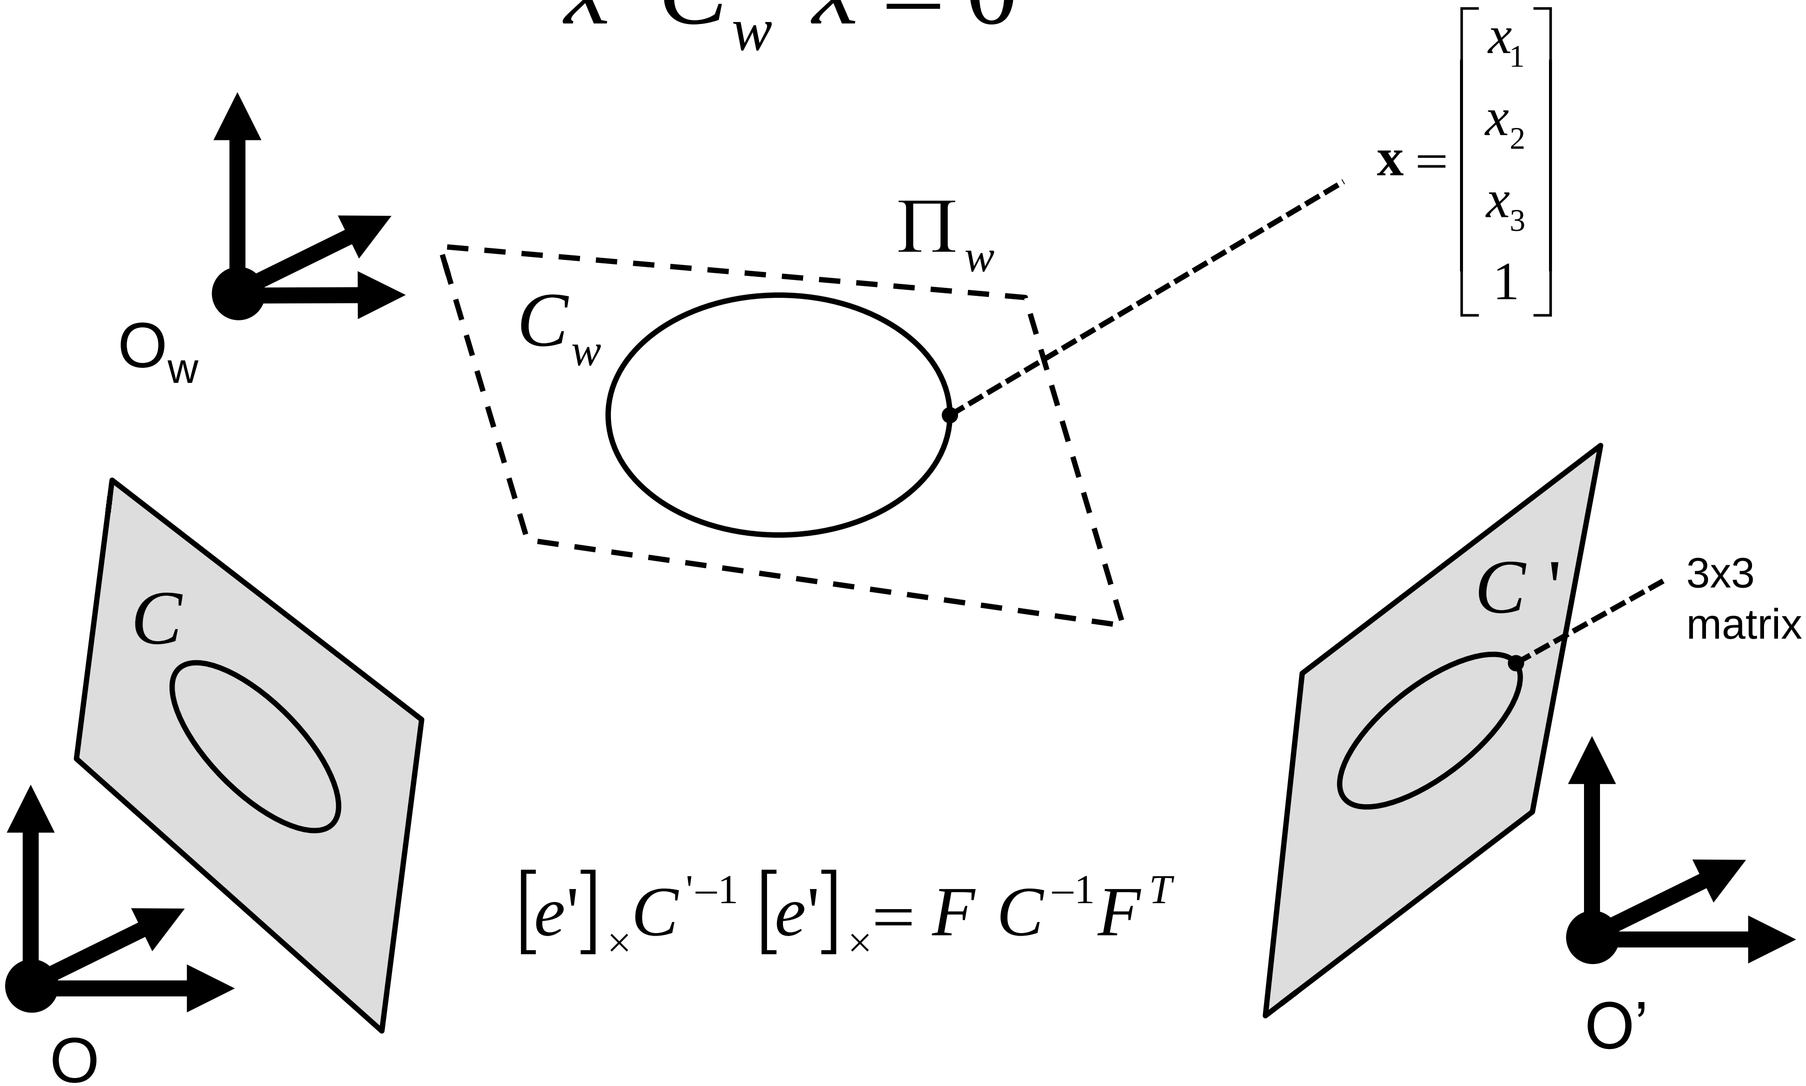

$$[e']_\times C'^{-1} [e']_\times = F \; C^{-1} F^T$$

Projection of absolute conics across views

From lecture 4, [HZ] page 210, sec. 8.5.1

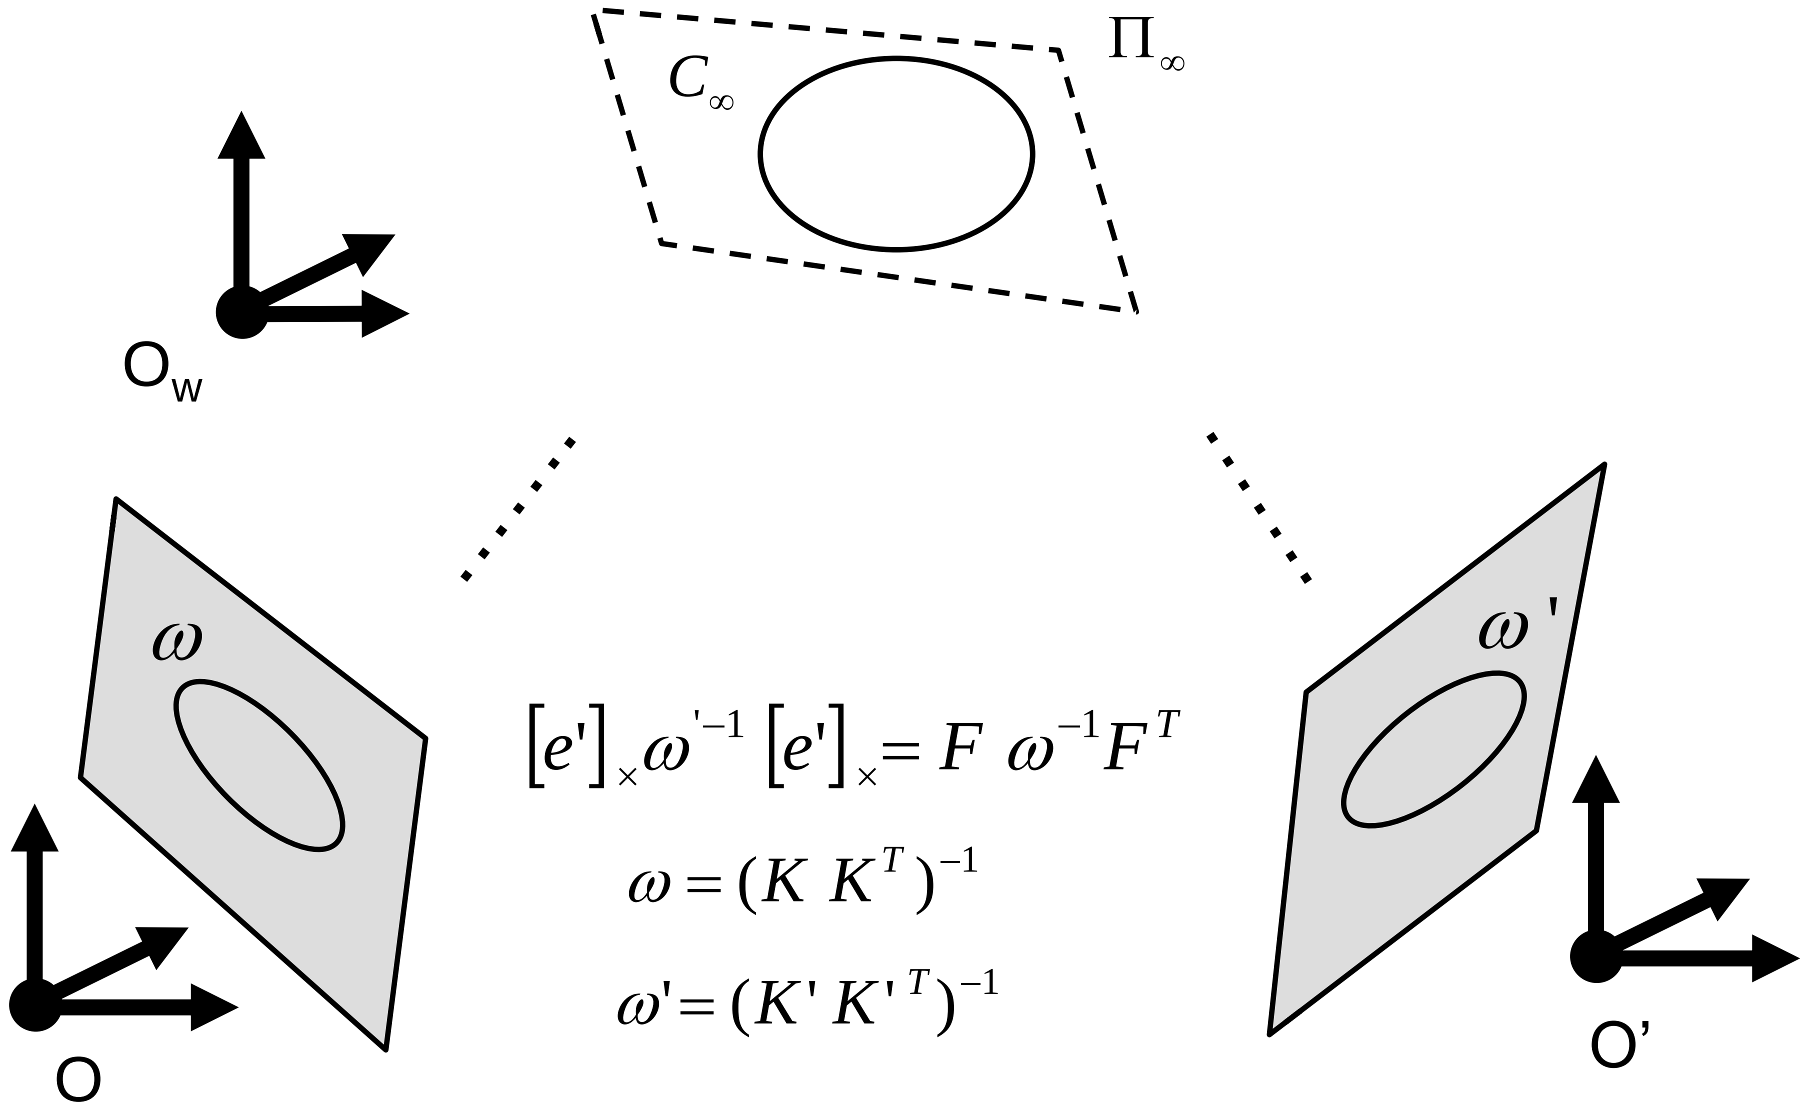

$$[e']_\times \omega'^{-1} [e']_\times = F \ \omega^{-1} F^T$$

$$\omega = (K \ K^T)^{-1}$$

$$\omega' = (K' \ K'^T)^{-1}$$

Kruppa equations

[Faugeras et al. 92]

$$\begin{pmatrix} u_2^T K' K'^T u_2 \\ -u_1^T K' K'^T u_2 \\ u_1^T K' K'^T u_1 \end{pmatrix} \times \begin{pmatrix} \sigma_1^2 v_1^T K K^T v_1 \\ \sigma_1 \sigma_2 v_1^T K K^T v_2 \\ \sigma_2^2 v_2^T K K^T v_2 \end{pmatrix} = 0$$

- Where u_i, v_i and σ_i are the columns and singular values of SVD of F

These give us two independent constraints in the elements of Ks

Kruppa equations

[Faugeras et al. 92]

$$\begin{pmatrix} u_2^T K' K'^T u_2 \\ -u_1^T K' K'^T u_2 \\ u_1^T K' K'^T u_1 \end{pmatrix} \times \begin{pmatrix} \sigma_1^2 v_1^T K K^T v_1 \\ \sigma_1 \sigma_2 v_1^T K K^T v_2 \\ \sigma_2^2 v_2^T K K^T v_2 \end{pmatrix} = 0$$

$$\frac{u_2^T K K^T u_2}{\sigma_1^2 v_1^T K K^T v_1} = \frac{-u_1^T K K^T u_2}{\sigma_1 \sigma_2 v_1^T K K^T v_2} = \frac{u_1^T K K^T u_1}{\sigma_2^2 v_2^T K K^T v_2}$$

- Special case where $K' = K = \begin{pmatrix} f & 0 & 0 \\ 0 & f & 0 \\ 0 & 0 & 1 \end{pmatrix}$

$$\alpha f^2 + \beta f + \gamma = 0 \longrightarrow f$$

Kruppa equations

[Faugeras et al. 92]

- Powerful if we want to self-calibrate 2 cameras with unknown focal length

- Limitations:
  - Work on a camera pair
  - Don’t work if R=0

$$[e']_\times \omega^{-1} [e']_\times = F\ \omega^{-1} F^T$$ becomes trivial

Since: $F = [e']_\times$

Algebraic approach

Multi-view approach

Suppose we have a projective reconstruction $\{\mathrm{M}_{\mathrm{i}}, \mathrm{X}_{\mathrm{j}}\}$

Let H be a homography such that:

- First perspective camera is canonical: $\mathrm{M}_1 = [\mathrm{I} \quad 0]$
- i^th perspective reconstruction of the camera (known): $\mathrm{M}_{\mathrm{i}} = [\mathrm{A}_{\mathrm{i}} \quad \mathrm{a}_{\mathrm{i}}]$

$$\left(\mathrm{A}_{\mathrm{i}} - \mathrm{a}_{\mathrm{i}}\mathrm{p}^{\mathrm{T}}\right)\mathrm{K}_1\mathrm{K}_1^{\mathrm{T}}\left(\mathrm{A}_{\mathrm{i}} - \mathrm{a}_{\mathrm{i}}\mathrm{p}^{\mathrm{T}}\right)^{\mathrm{T}} = \mathrm{K}_{\mathrm{i}}\ \mathrm{K}_{\mathrm{i}}^{\mathrm{T}} \qquad \text{i=2…m}$$

p is an unknown 3x1 vector

How many unknowns?
- 3 from p
- 5 x (m+1) from Ks

How many equations? 5 independent equations [per view]

$K_i \ K_i^T$ is 3x3 symmetric and defined up scale

Algebraic approach

Art of self-calibration:

use constraints on Ks to generate enough equations on the unknowns

Condition	*N. Views*
•Constant internal parameters	3
•Aspect ratio and skew known •Focal length and offset vary	4
•Aspect ratio and skew constant •Focal length and offset vary	5
•skew =0, all other parameters vary	8

Issue: the larger is the number of view,
the harder is the correspondence problem

Bundle adjustment helps!

Lecture 8

SFM & Volumetric stereo

- SFM: Self-calibration
- Volumetric stereo:
  - Space carving
  - Shadow carving
  - Voxel carving

Reading:

[HZ] Chapters 19 "Auto-calibration"
[Szelisky] Chapter 7 "Structure from motion"
[Szelisky] Chapter 11 "Multi-view stereo"

“Traditional” Stereo

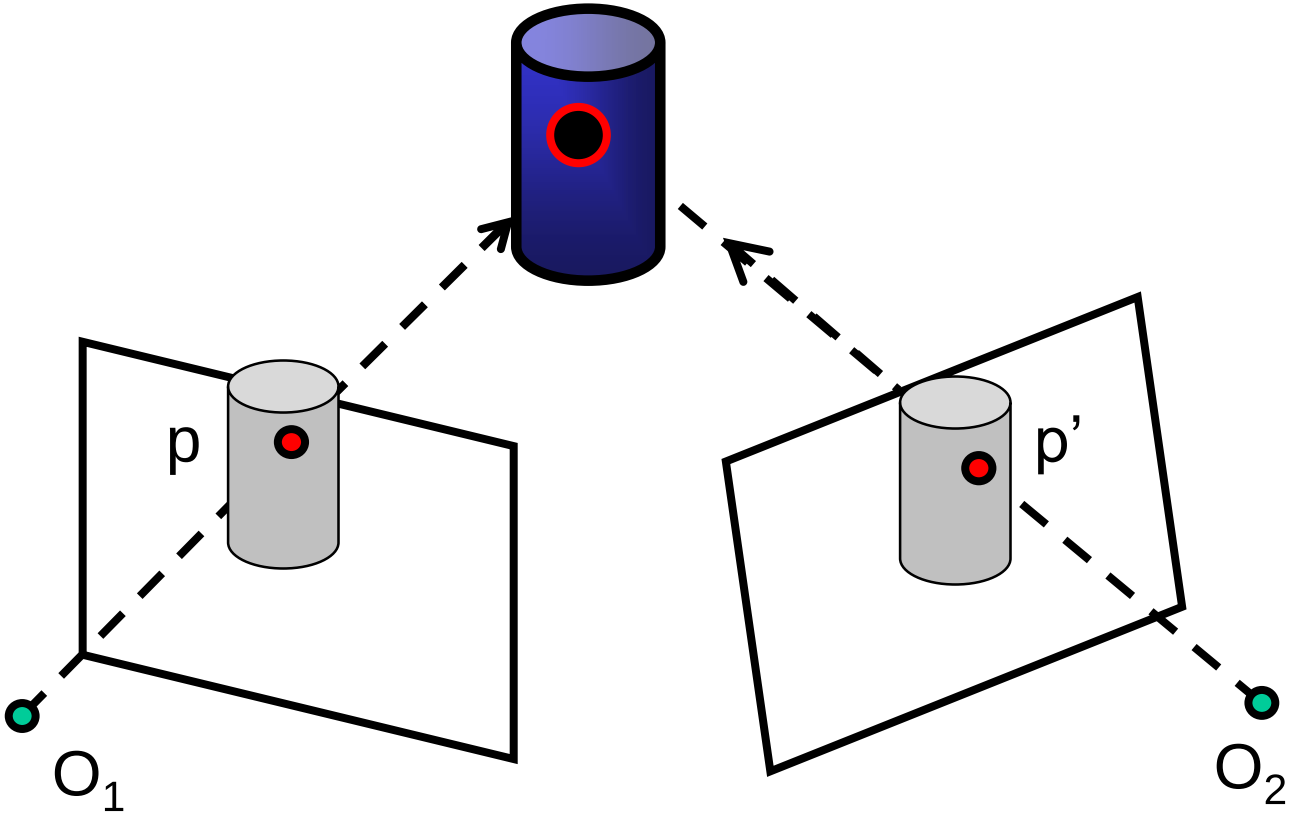

Goal: estimate the position of P given the observation of P from two view points

Assumptions: known camera parameters and position (K, R, T)

“Traditional” Stereo

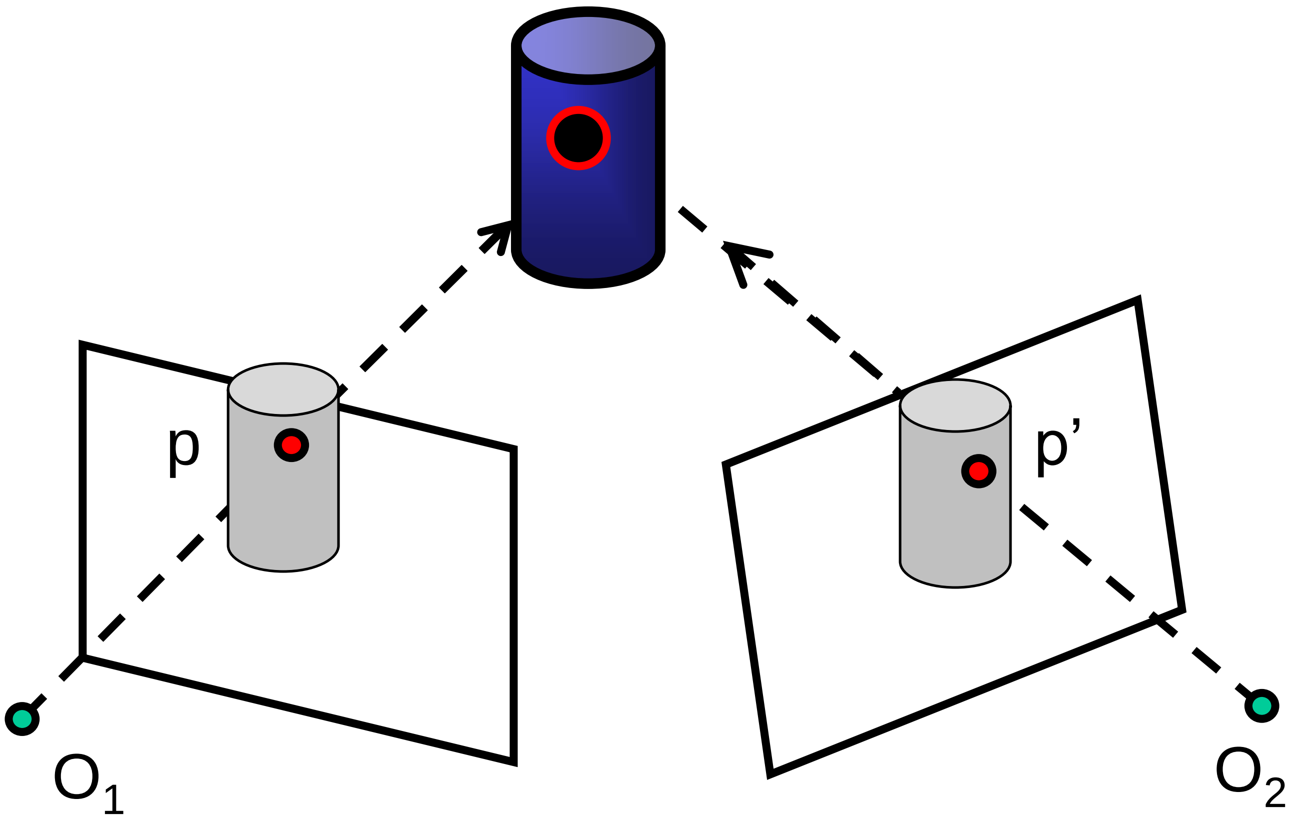

Subgoals:

1. Solve the correspondence problem
2. Use corresponding observations to triangulate

Volumetric stereo

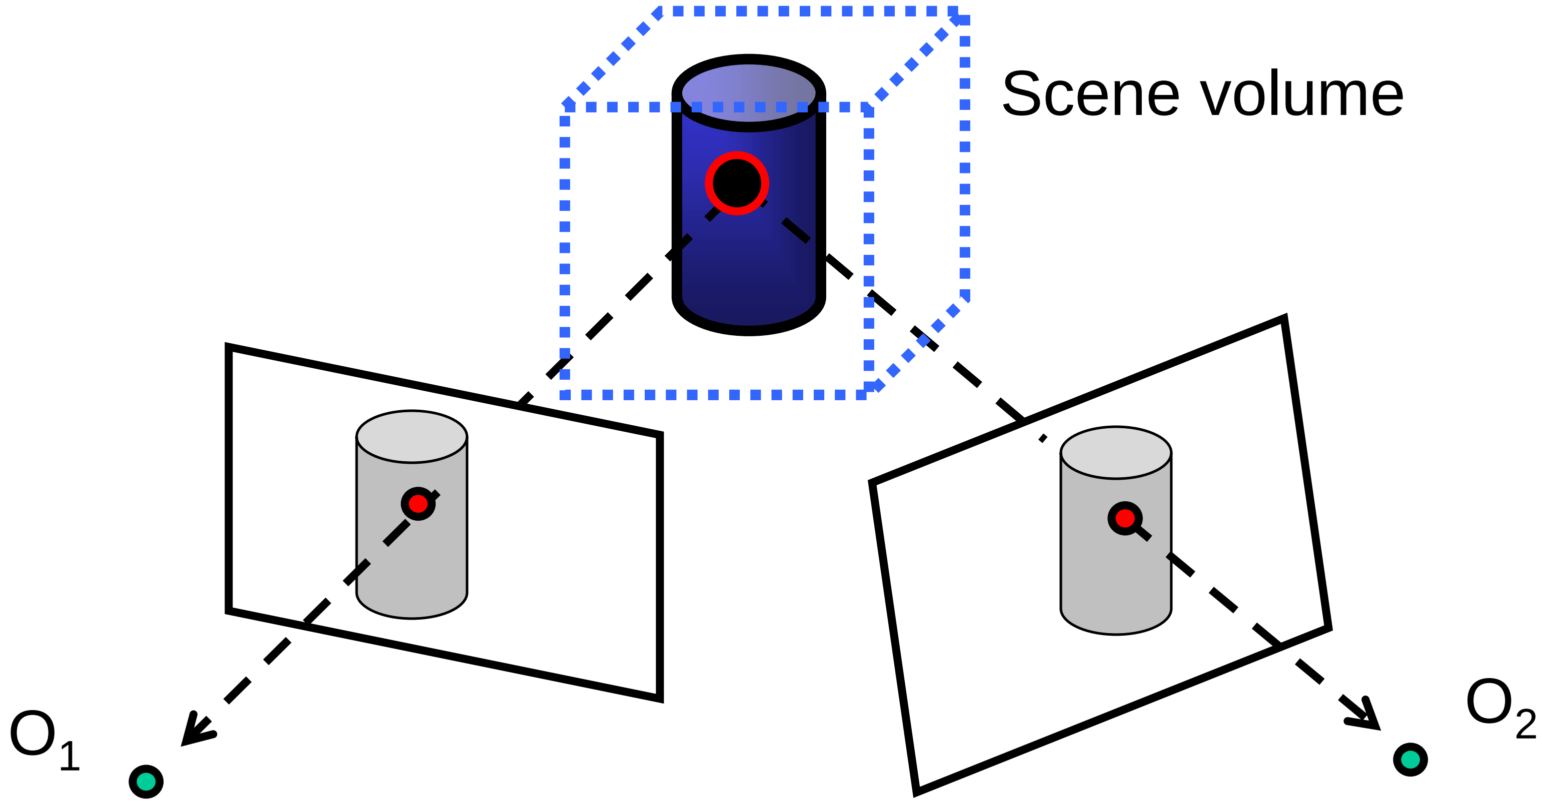

1. Hypothesis: pick up a point within the volume
2. Project this point into 2 (or more) images
3. Validation: are the observations consistent?

Assumptions: known camera parameters and position (K, R, T)

Consistency based on cues such as:

- Contours/silhouettes
- Shadows
- Colors

Volumetric Stereo

- Contours are a rich source of geometric information

Apparent Contour

[sato & cipolla]

- <u>DEFINITION</u>: projection of the locus of points on the surface which separate the visible and occluded parts on the surface

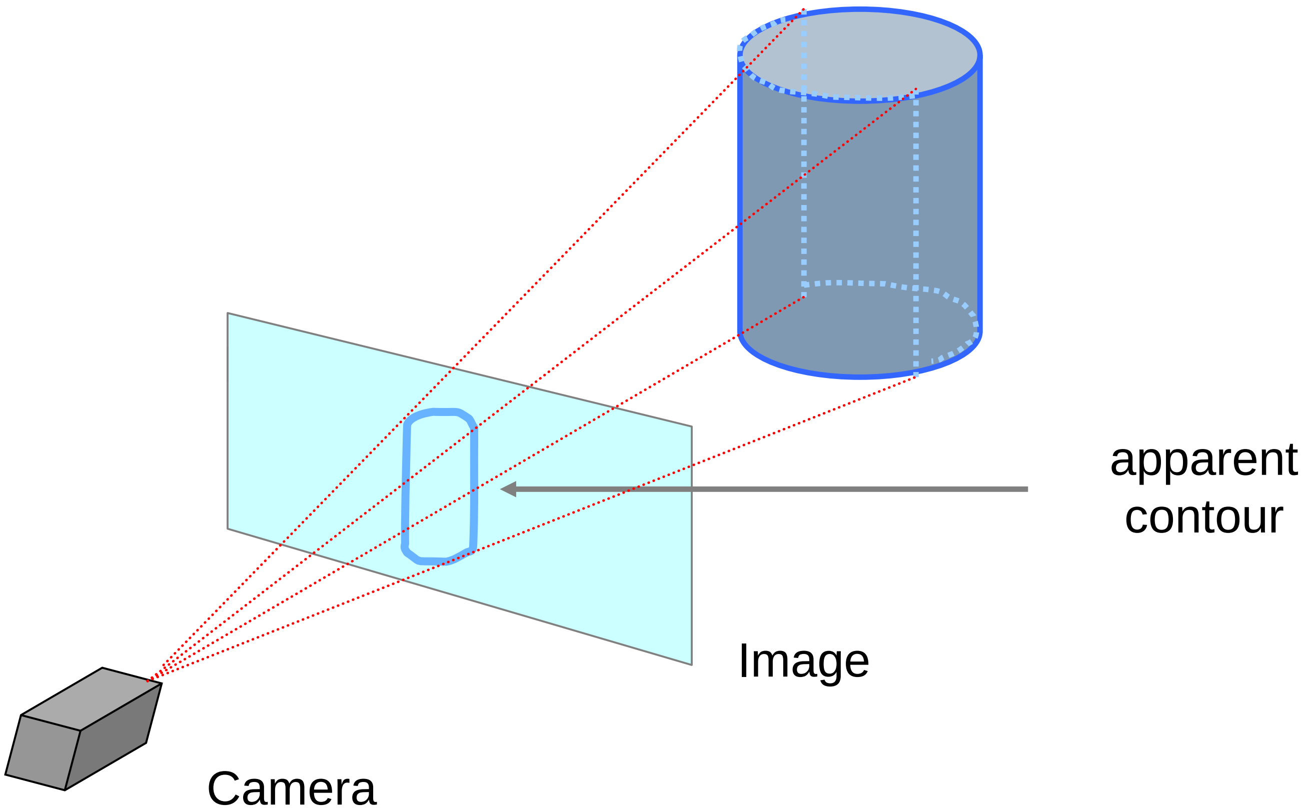

Silhouettes

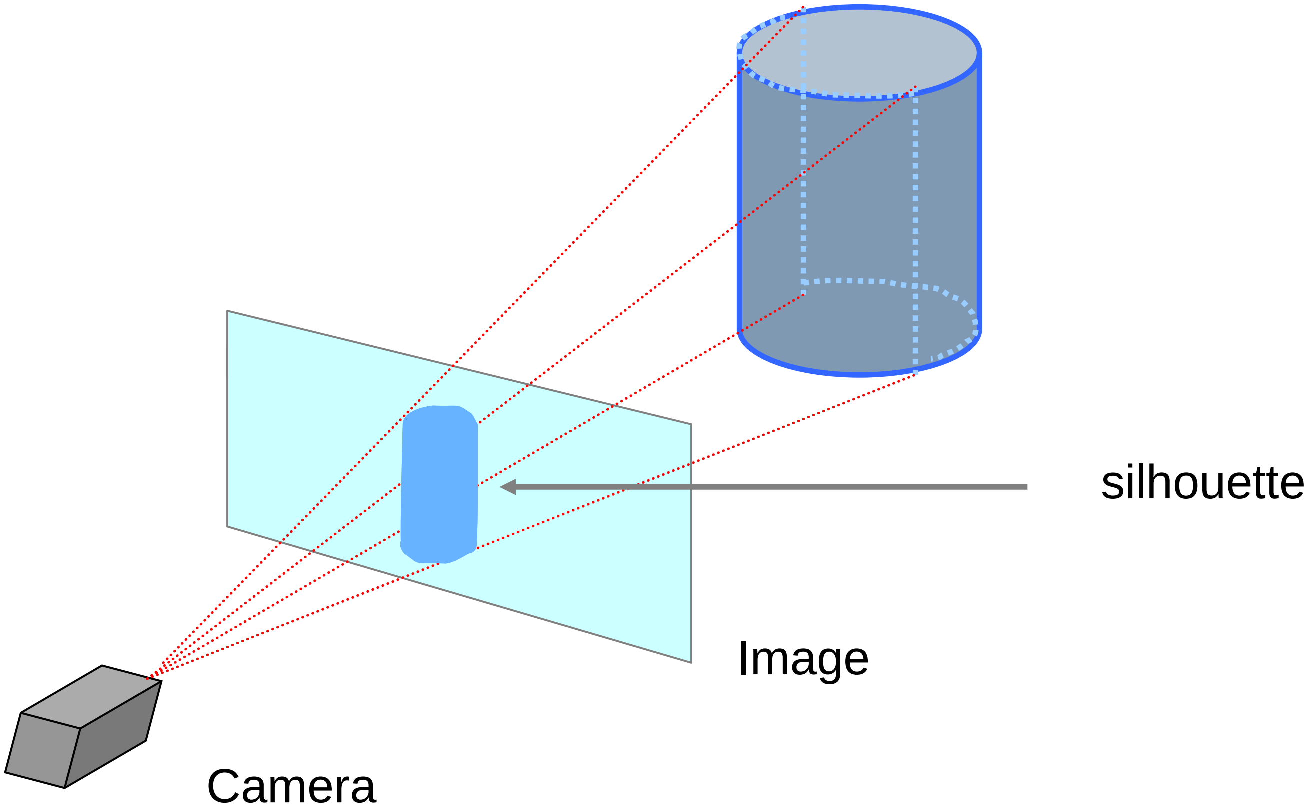

Easy to detect

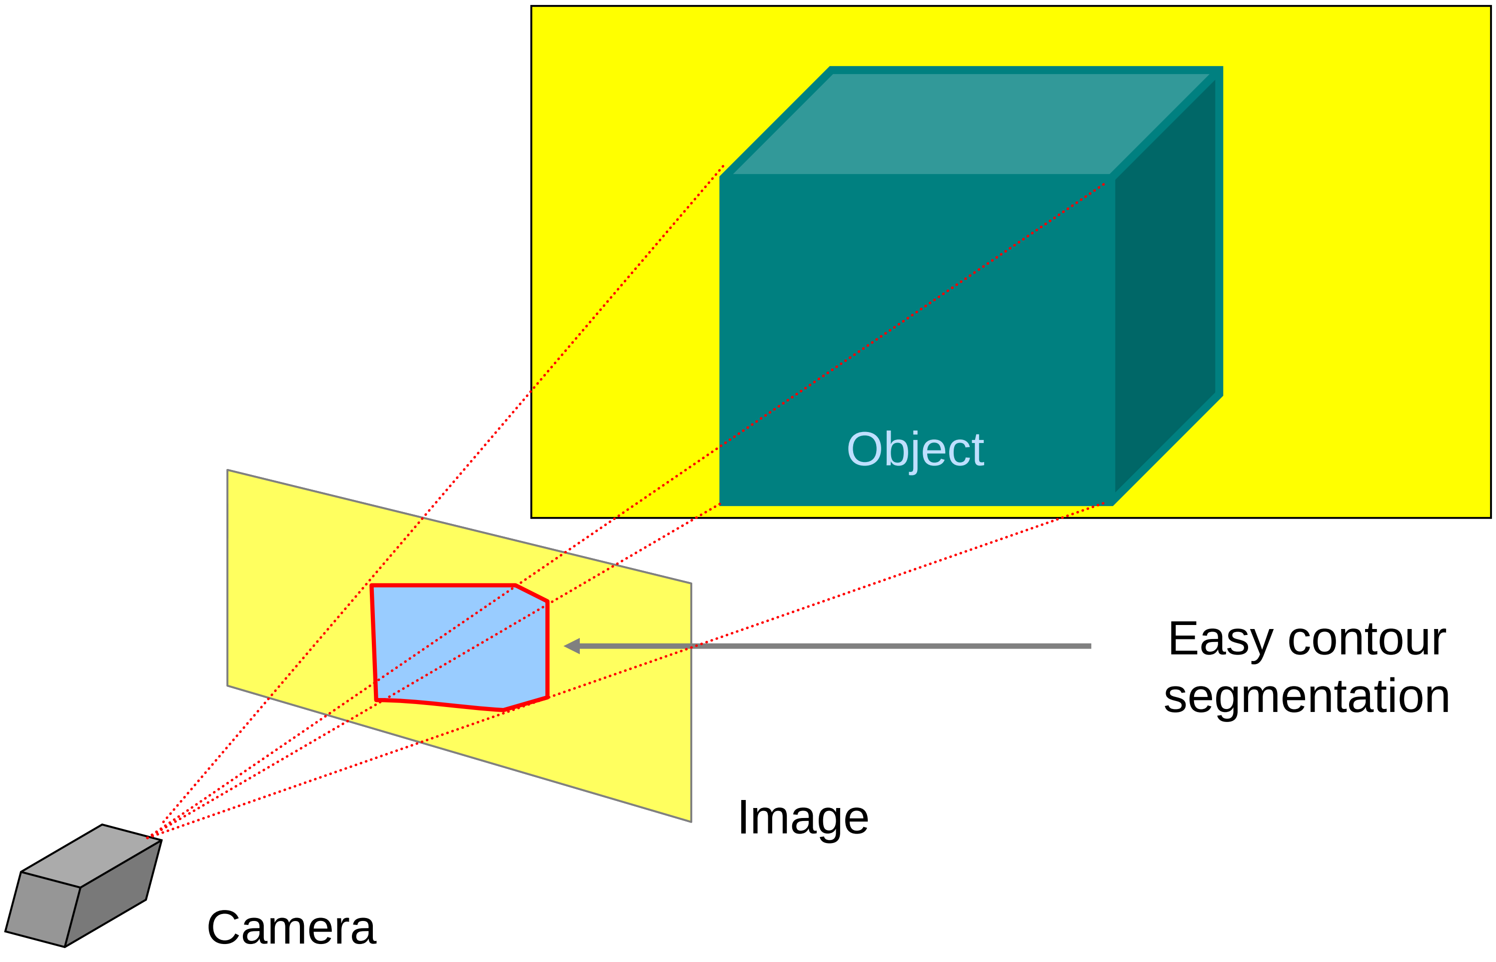

How can we use contours?

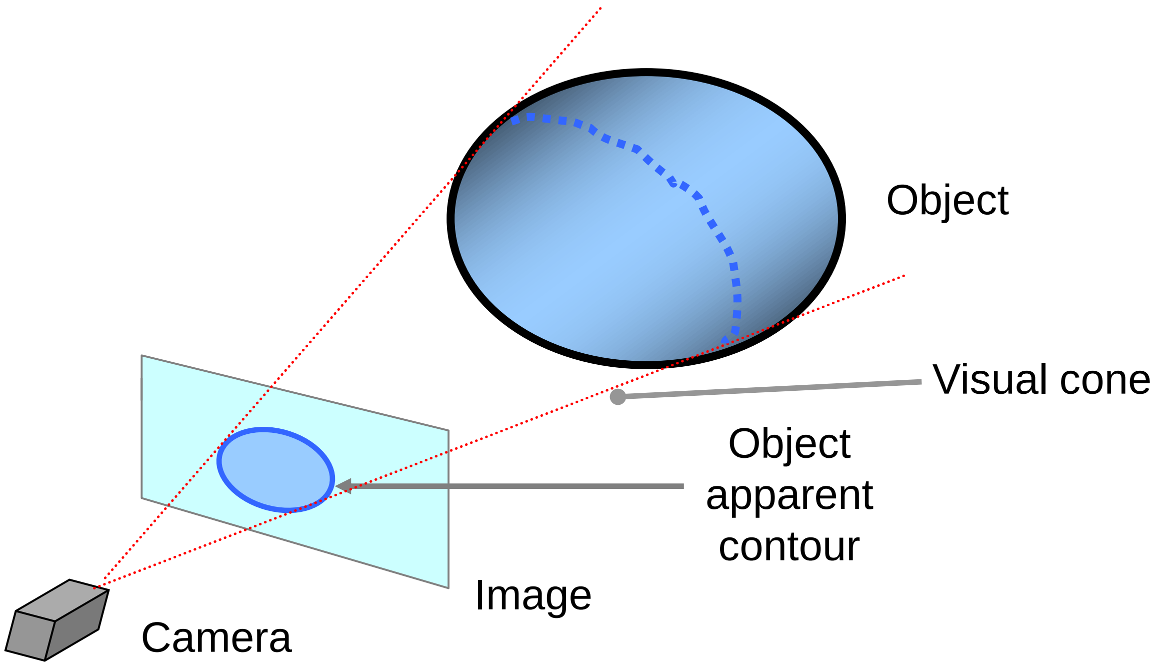

How can we use contours?

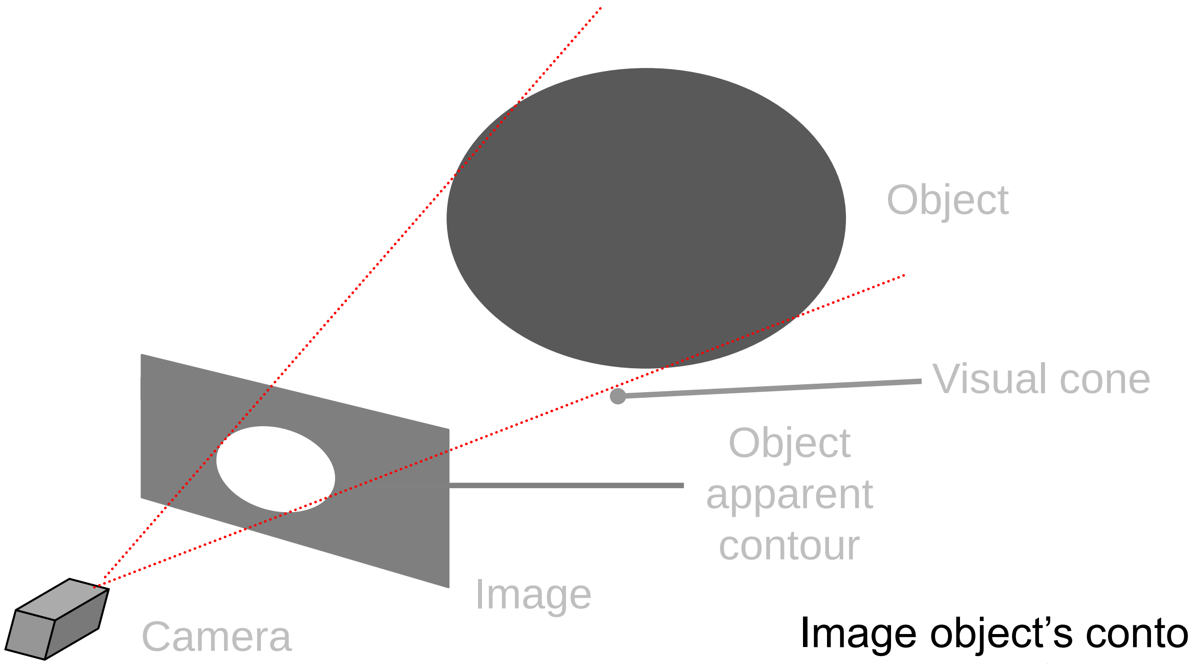

in 2D:

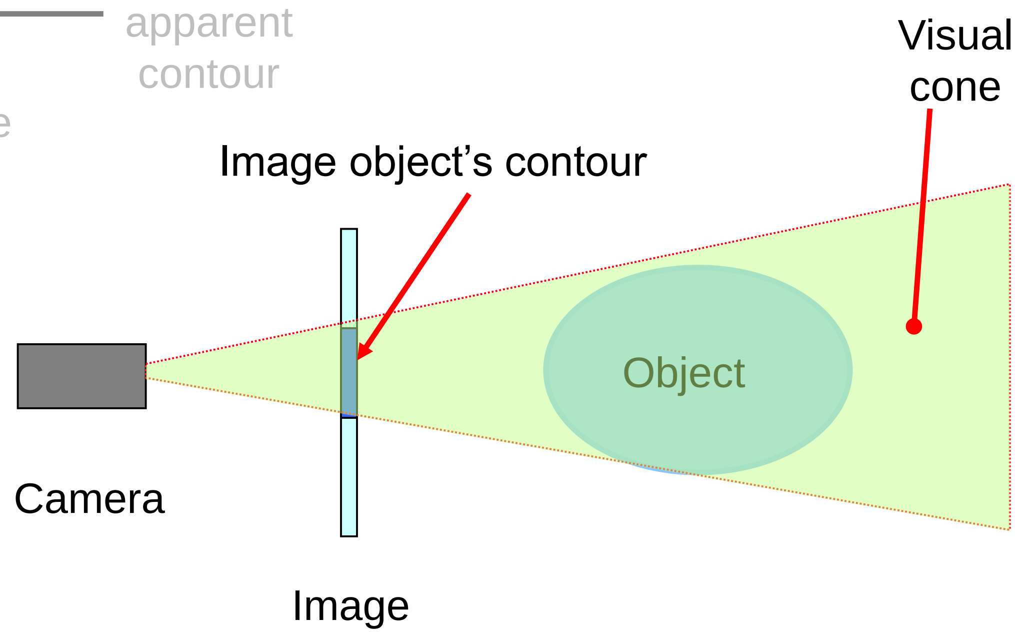

How can we use contours?

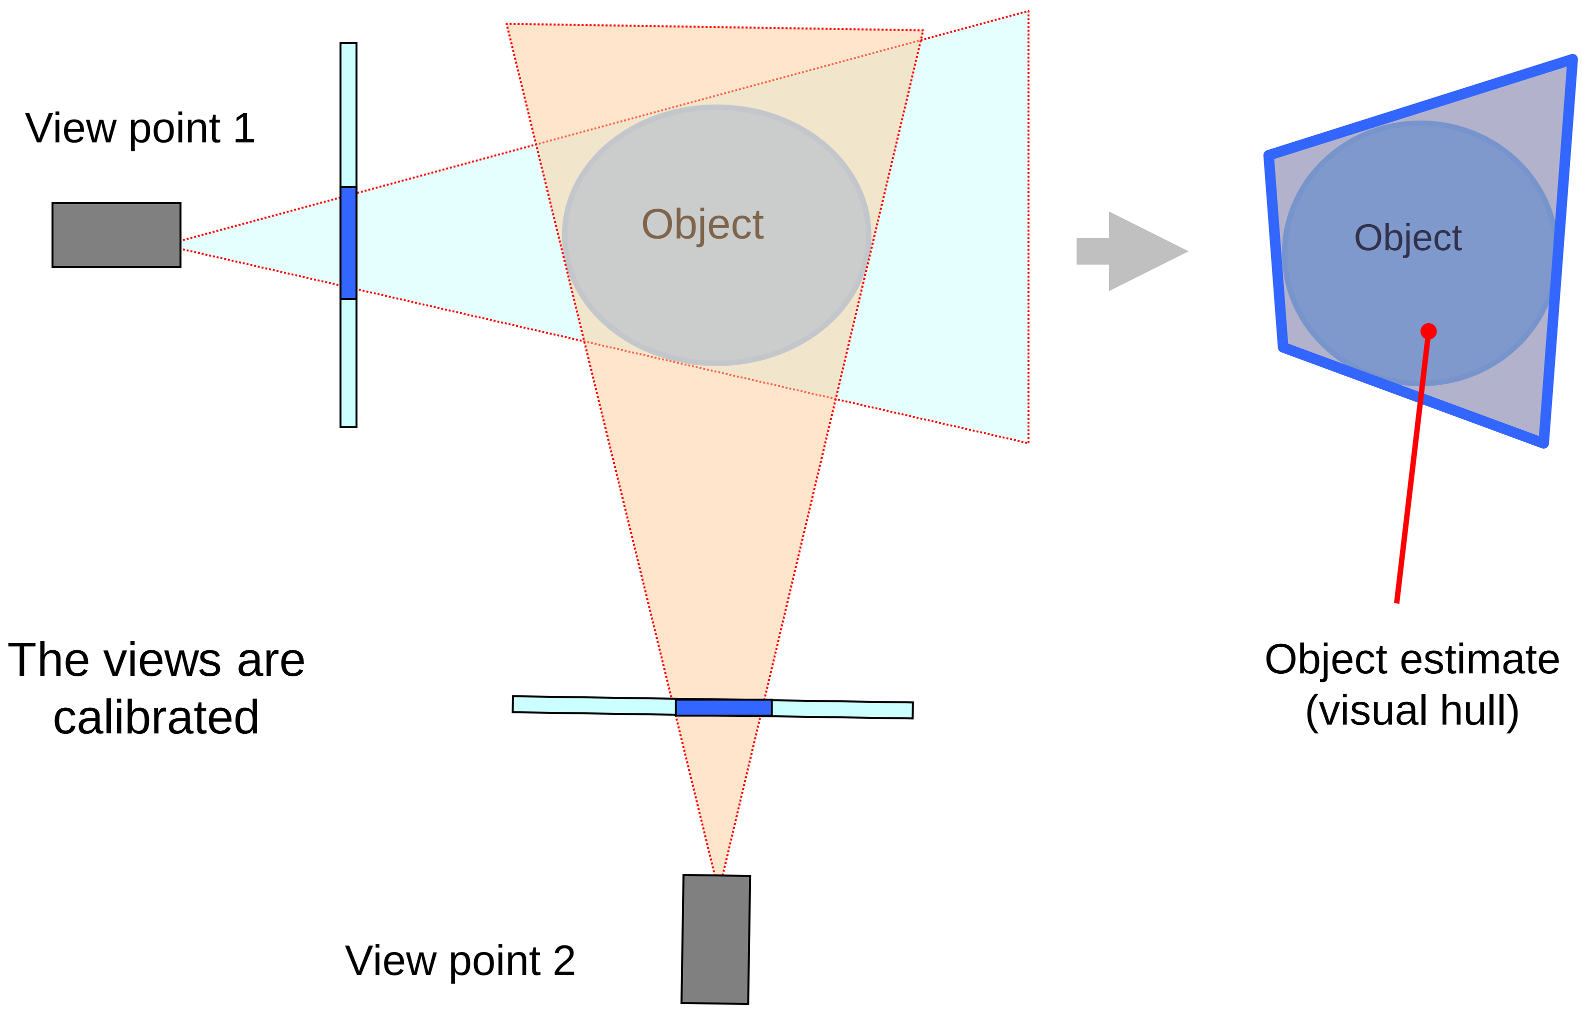

How to perform visual cones intersection?

- Decompose visual cone in polygonal surfaces
  (among others: Reed and Allen '99)

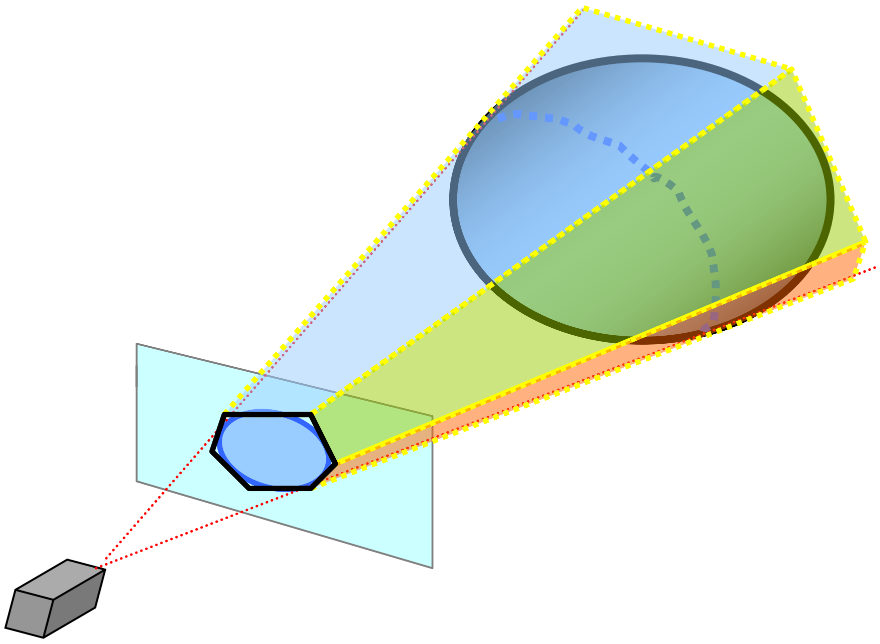

Space Carving

[Martin and Aggarwal (1983)]

- Using contours/silhouettes in volumetric stereo, also called **space carving**

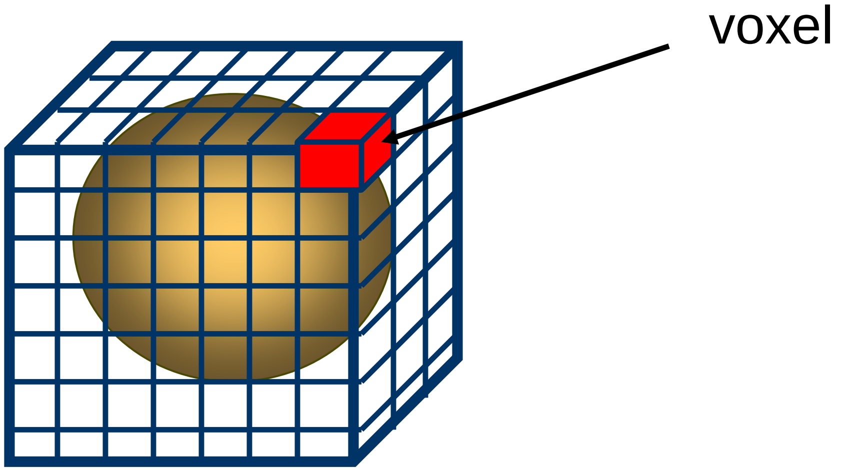

Computing Visual Hull in 2D

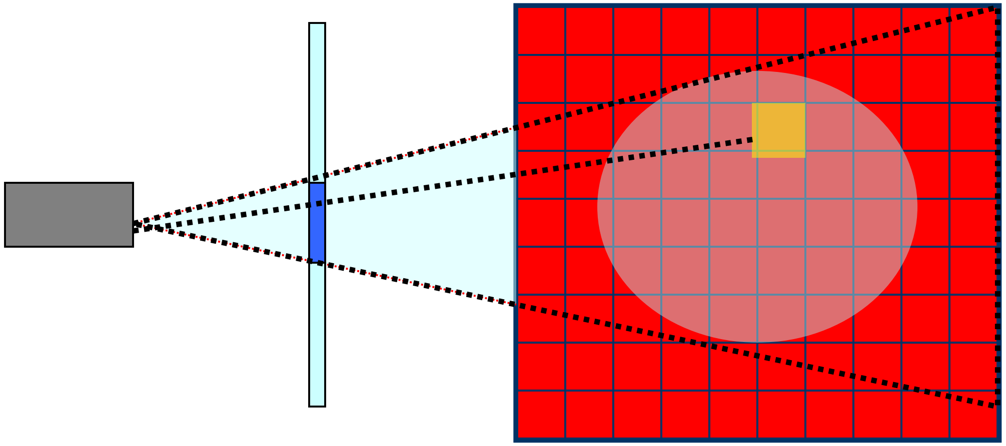

Computing Visual Hull in 2D

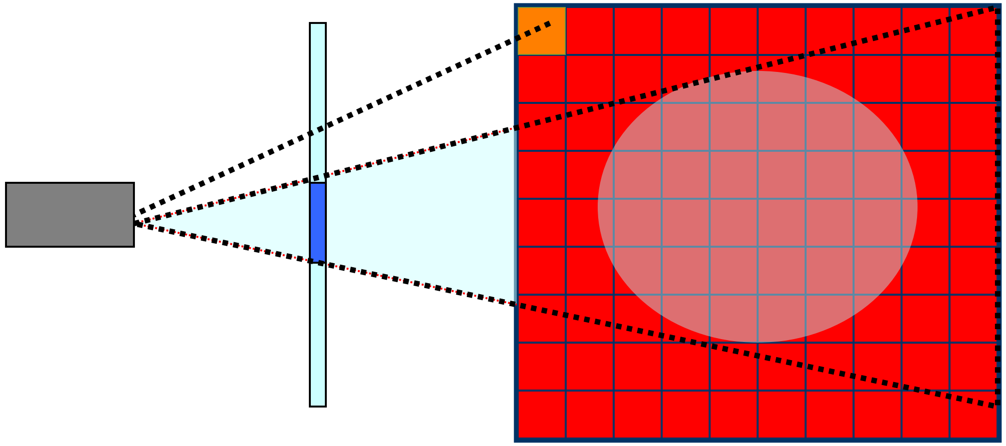

Computing Visual Hull in 2D

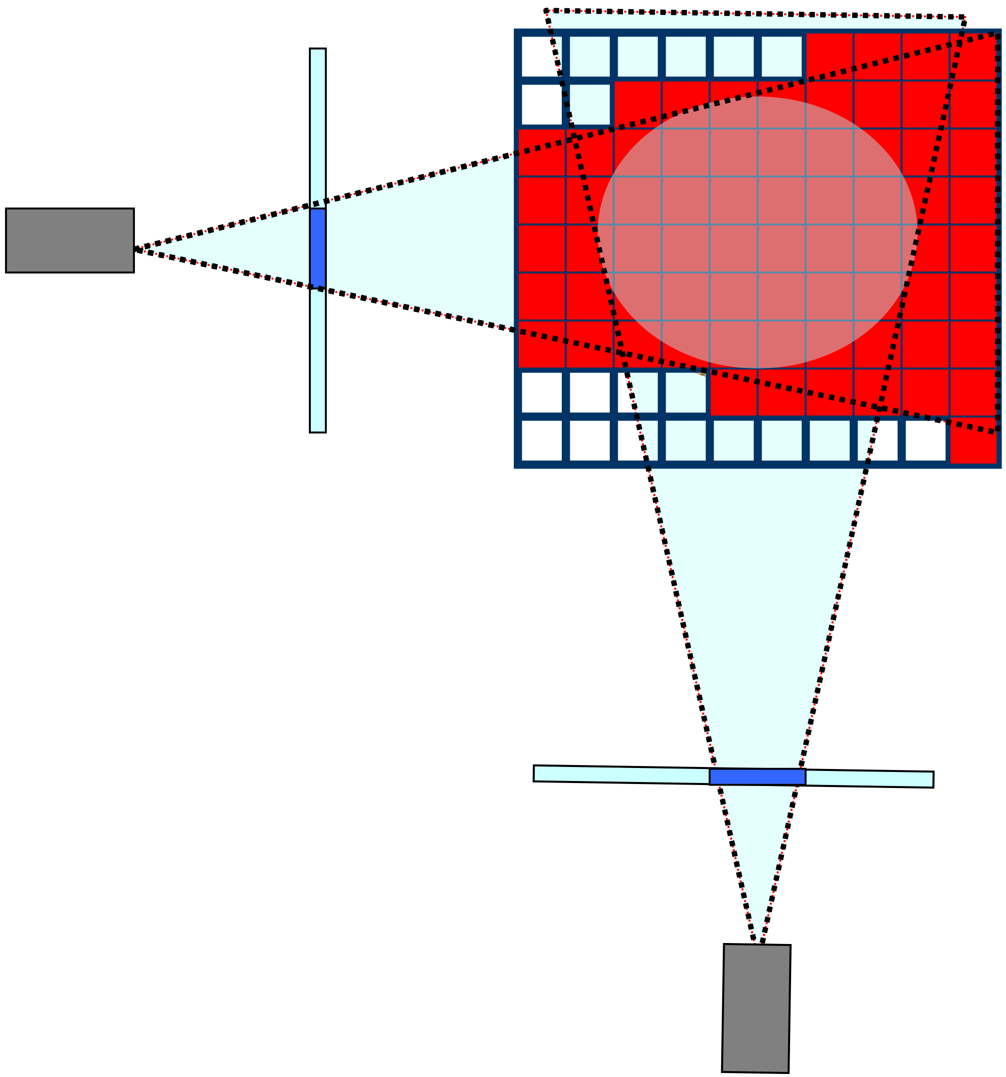

Computing Visual Hull in 2D

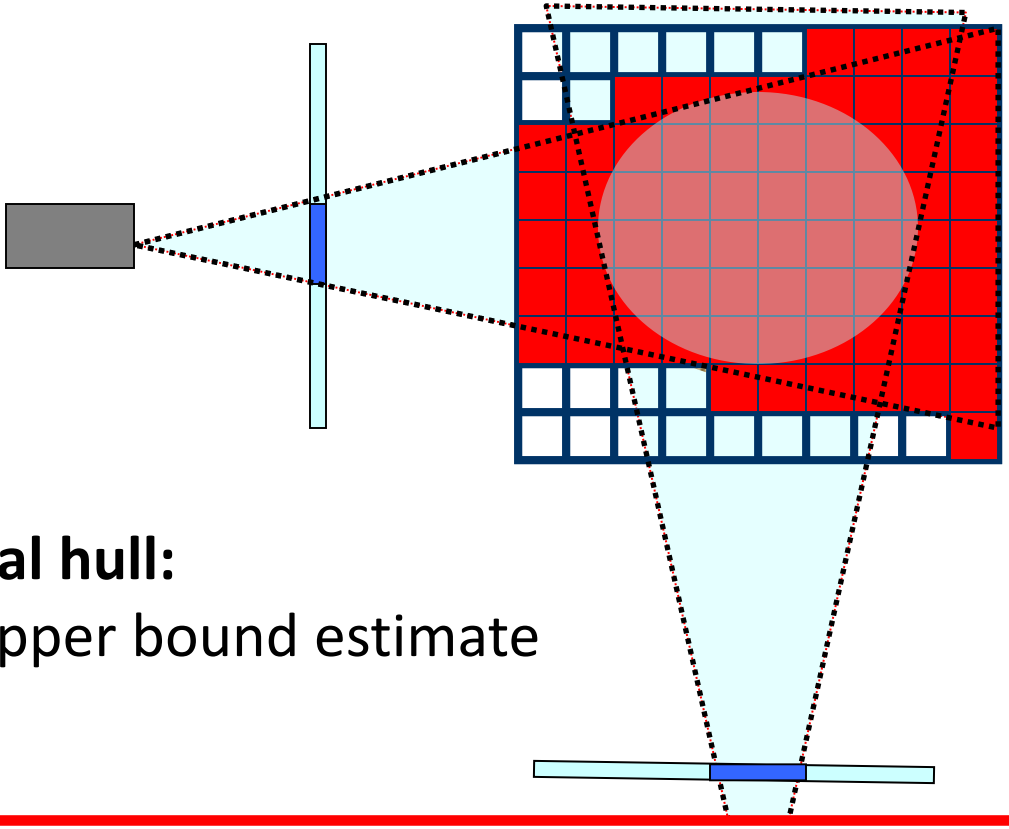

Visual hull:
an upper bound estimate

Consistency:

A voxel must be projected into a silhouette in each image

Space Carving has complexity …

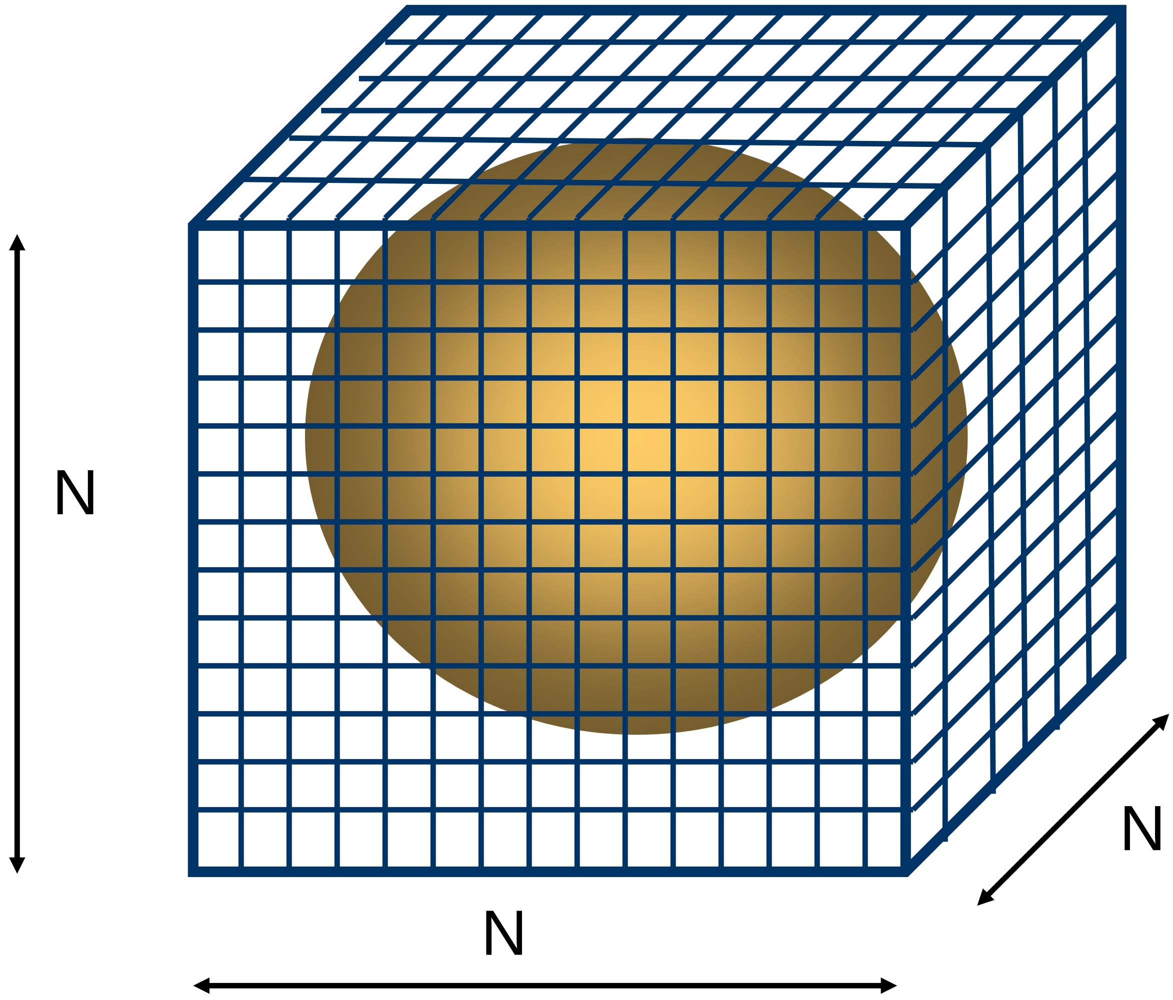

$O(N^3)$

Octrees
(Szeliski '93)

Complexity Reduction: Octrees

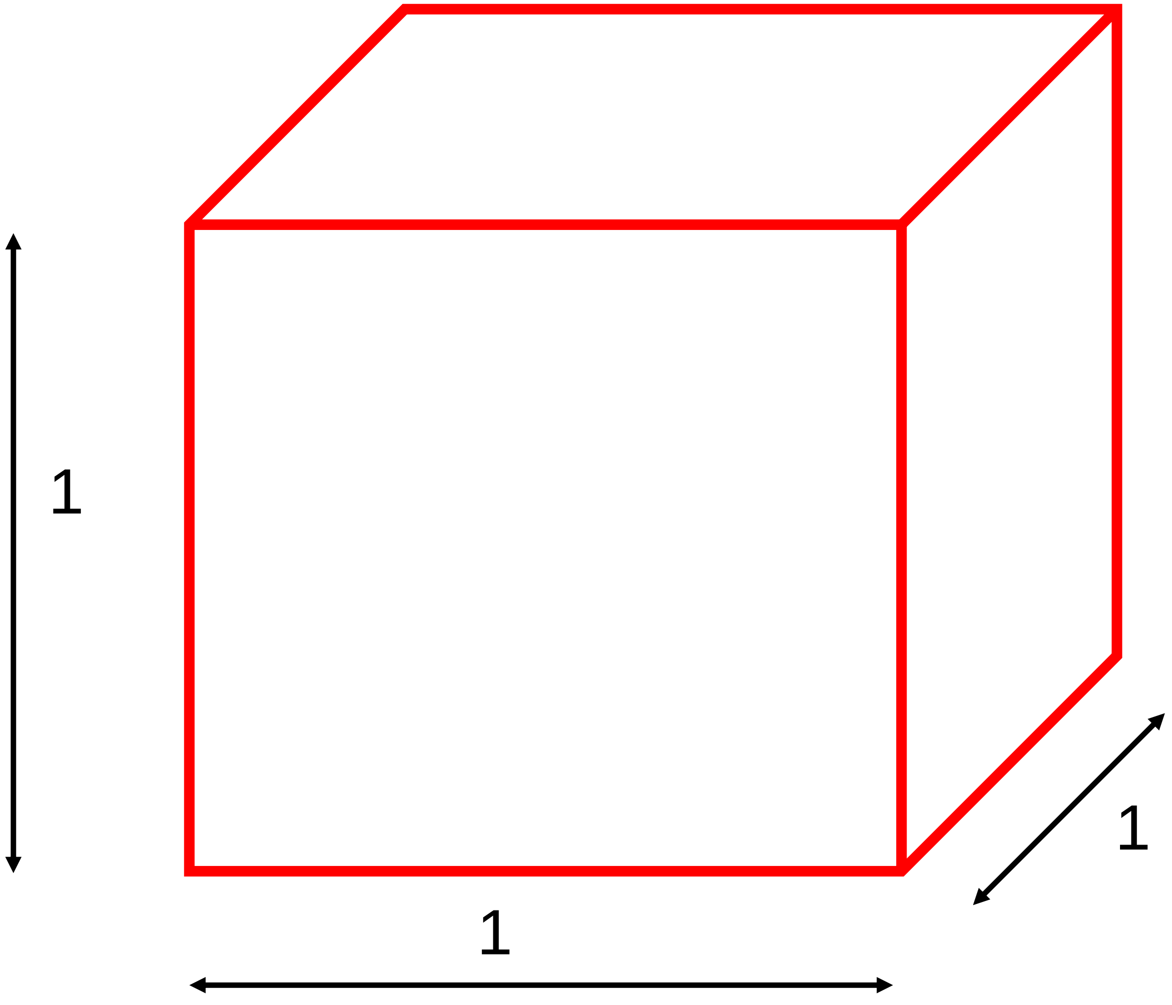

Complexity Reduction: Octrees

- Subdiving volume in voxels of progressive smaller size

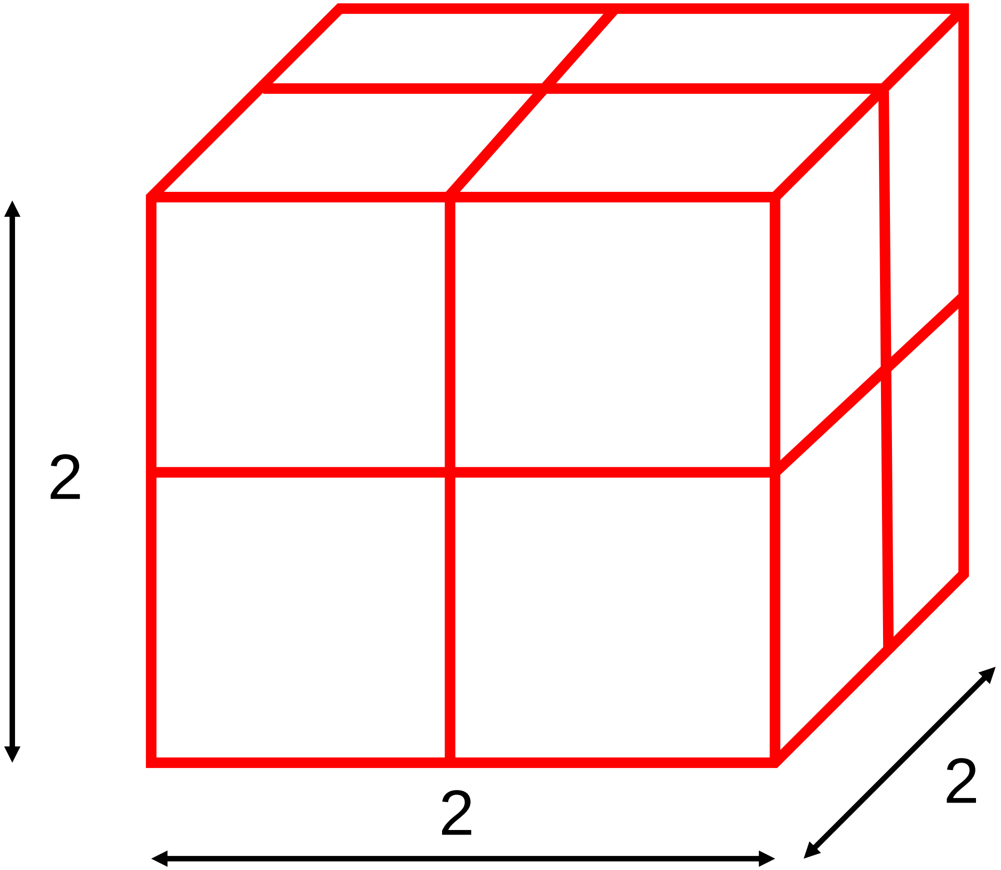

Complexity Reduction: Octrees

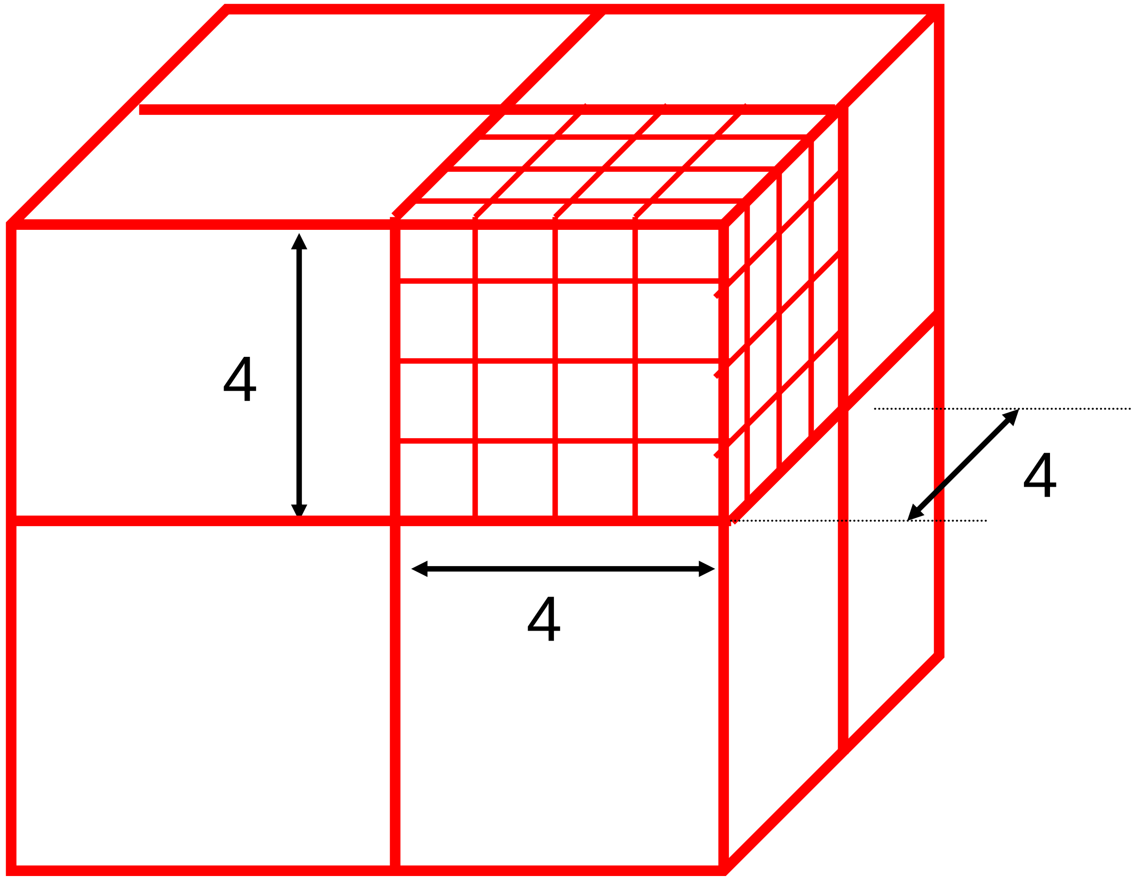

Complexity reduction: 2D example

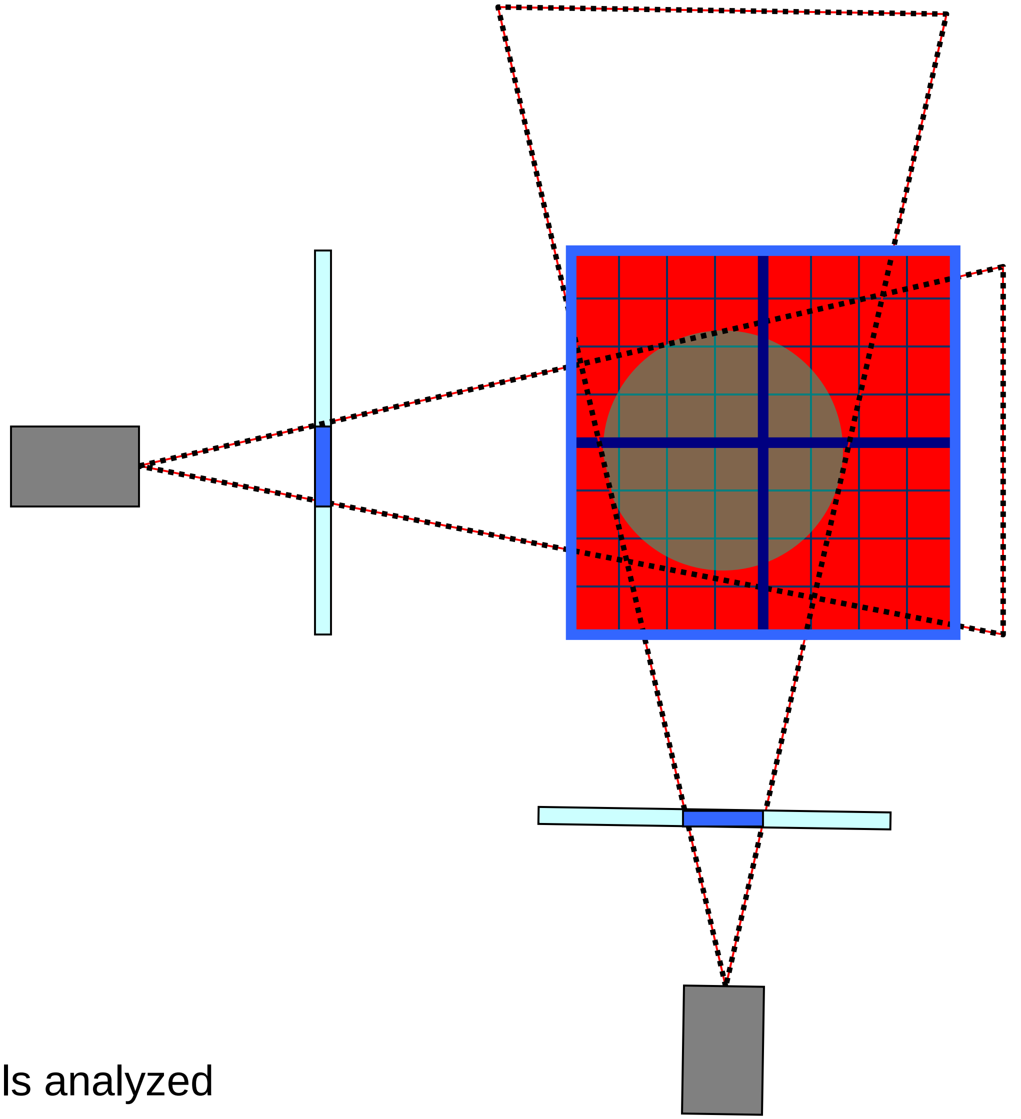

4 voxels analyzed

Complexity reduction: 2D example

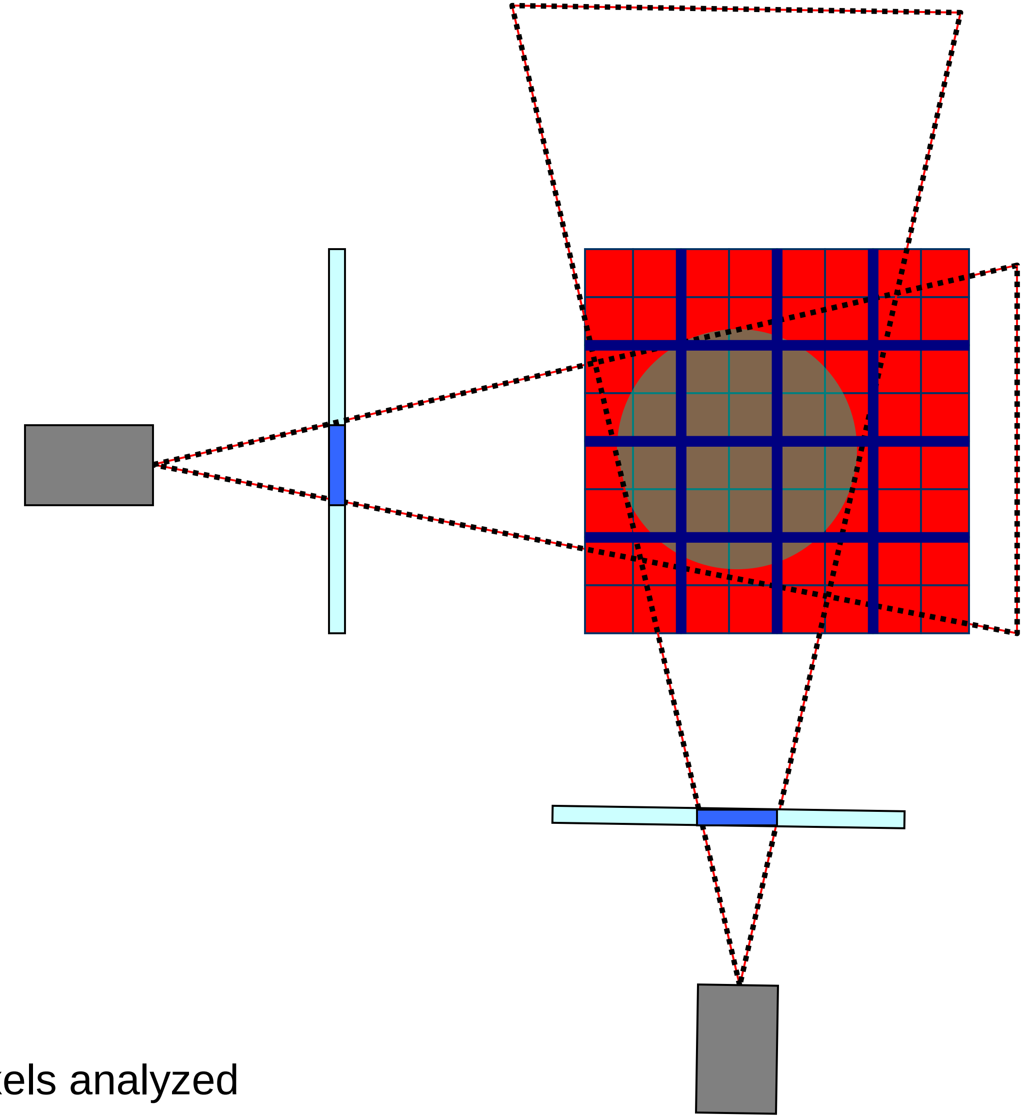

16 voxels analyzed

Complexity reduction: 2D example

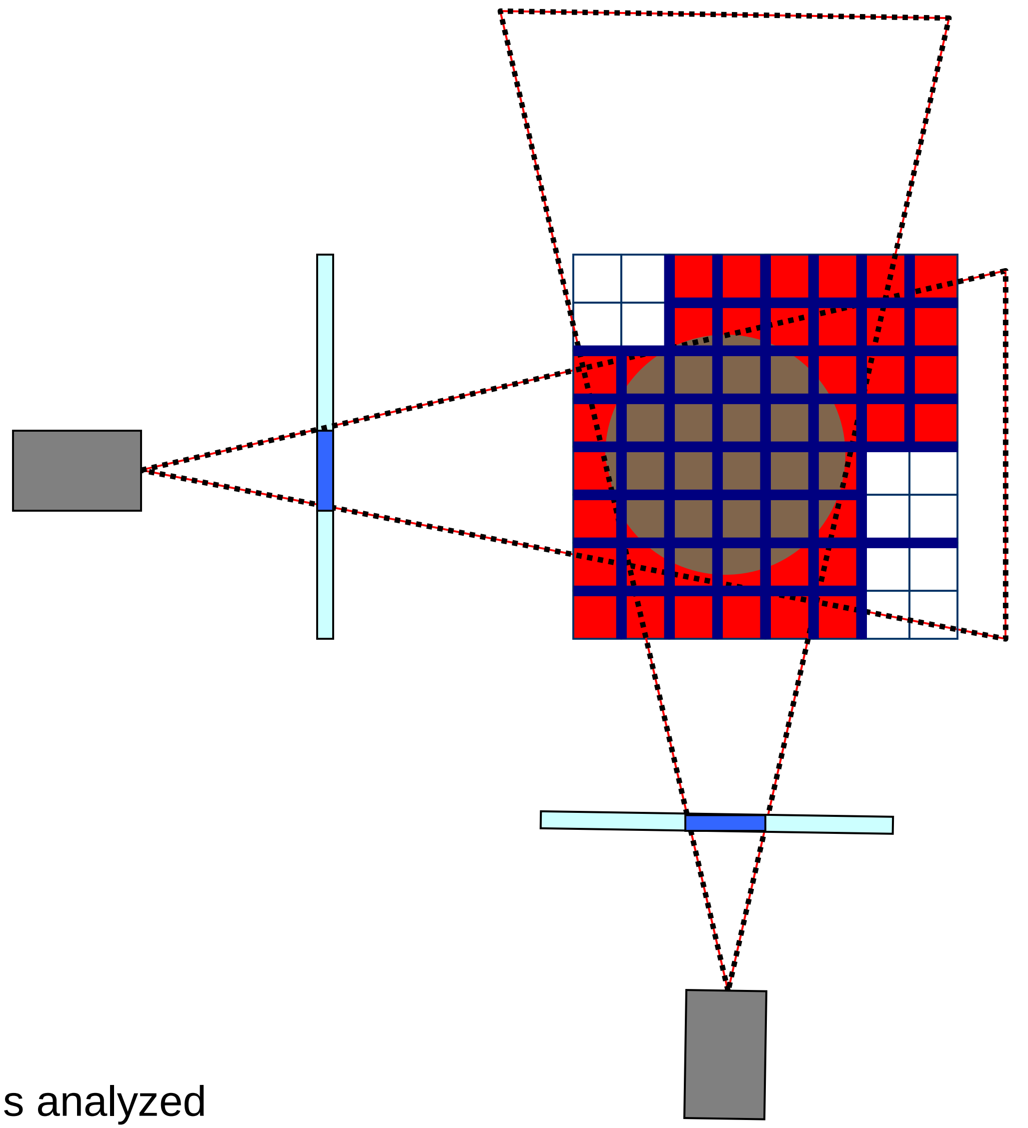

52 voxels analyzed

Complexity reduction: 2D example

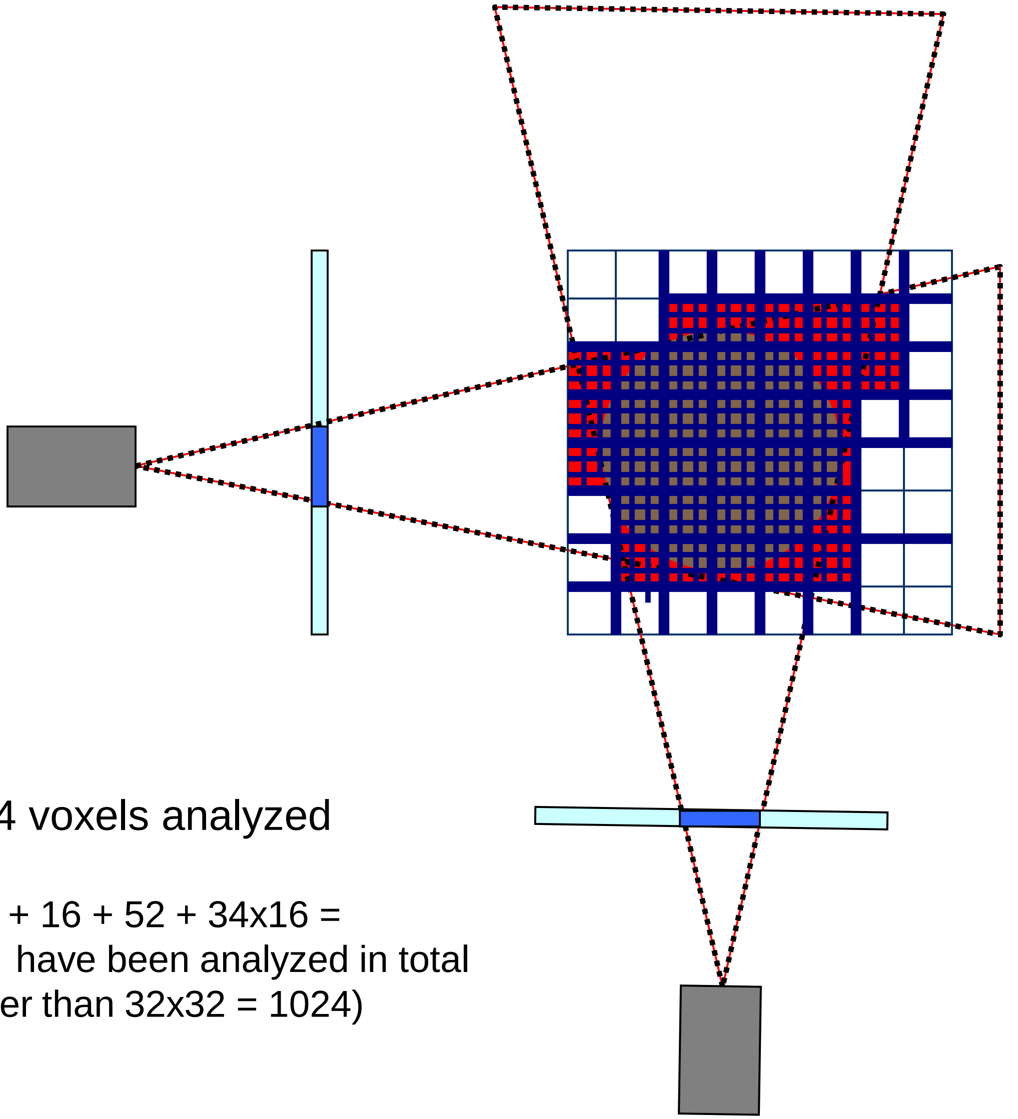

16x34 voxels analyzed

1+ 4 + 16 + 52 + 34x16 =
617 voxels have been analyzed in total
(rather than 32x32 = 1024)

Advantages of Space Carving

- Robust and simple
- No need to solve for correspondences

Limitations of Space Carving

- Accuracy function of number of views

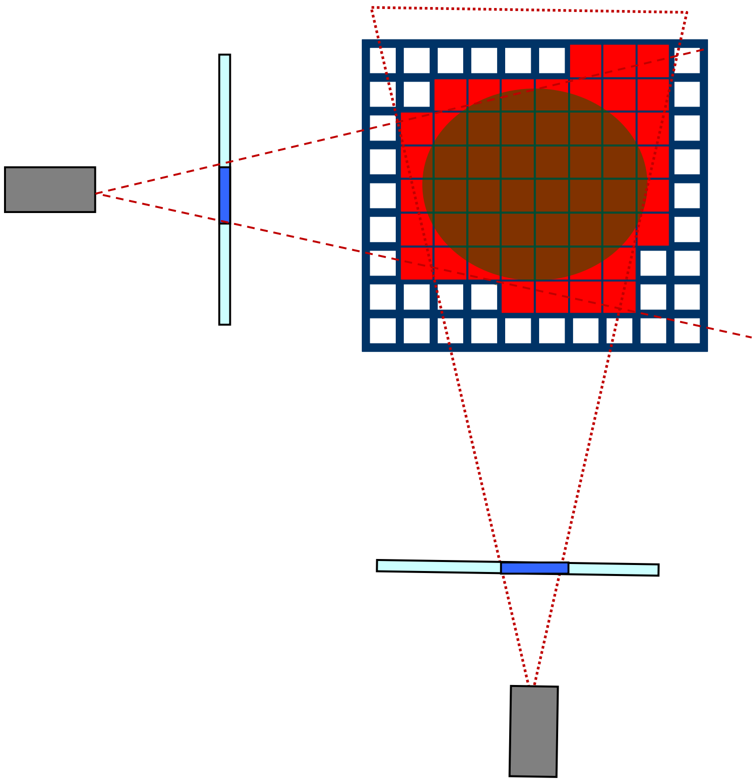

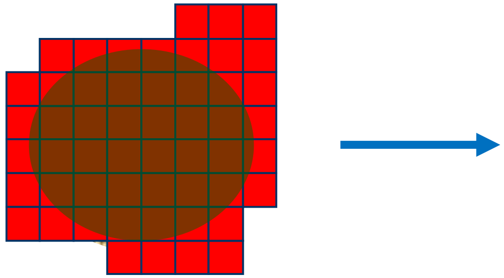

Not a good estimate

What else?

Limitations of Space Carving

- Concavities are not modeled

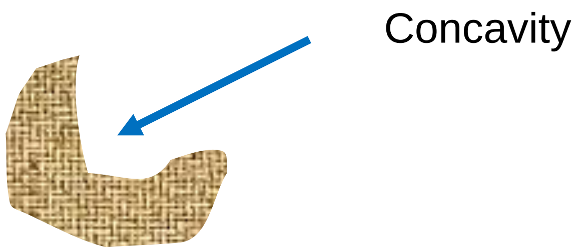

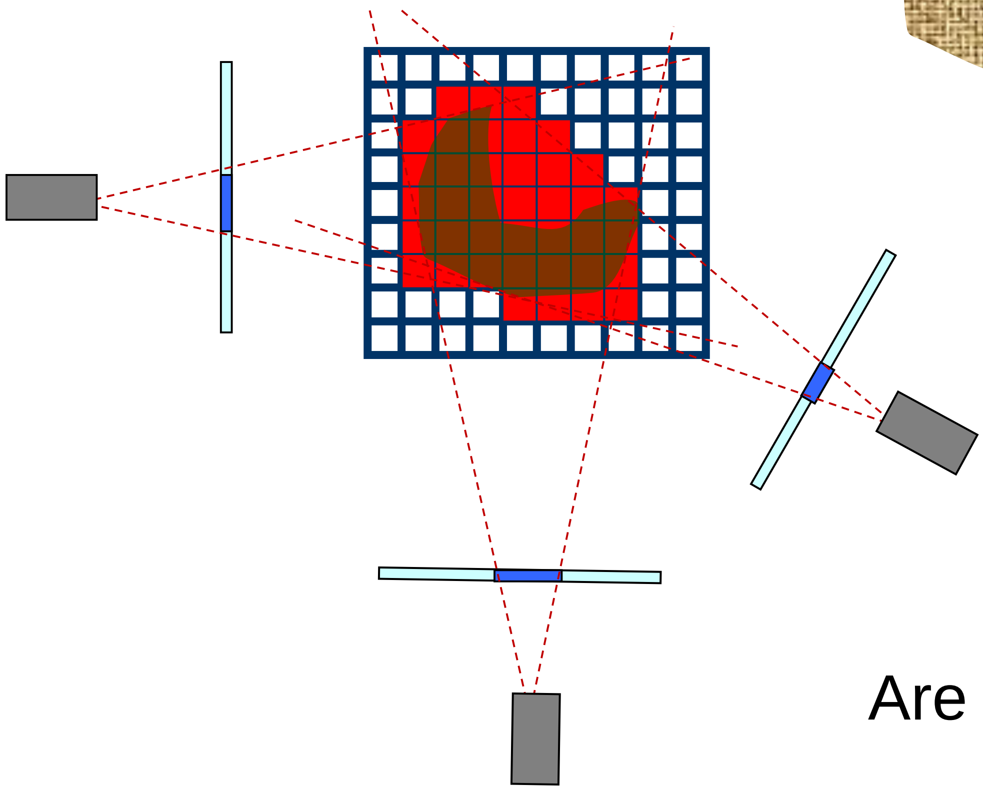

For 3D objects:
Are all types of concavities problematic?

Limitations of Space Carving

- Concavities are not modeled

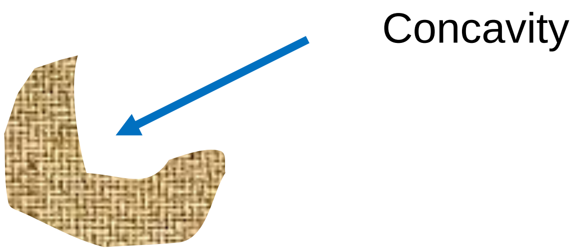

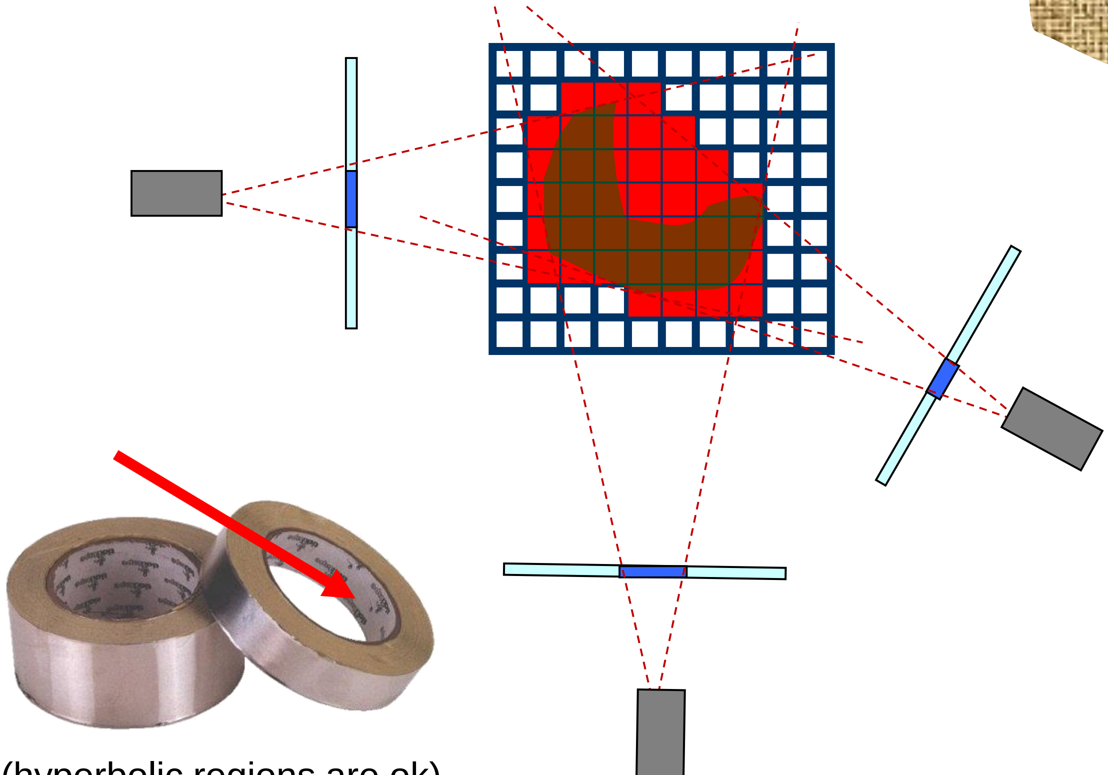

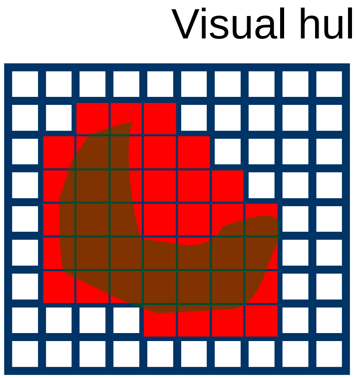

(hyperbolic regions are ok)

- Laurentini (1995)
  - Closest approximation
  - Conservative

Space Carving: A Classic Setup

Courtesy of Seitz & Dyer

Space Carving: A Classic Setup

Space Carving: Experiments

24 poses (15^{O})

voxel size = 2mm

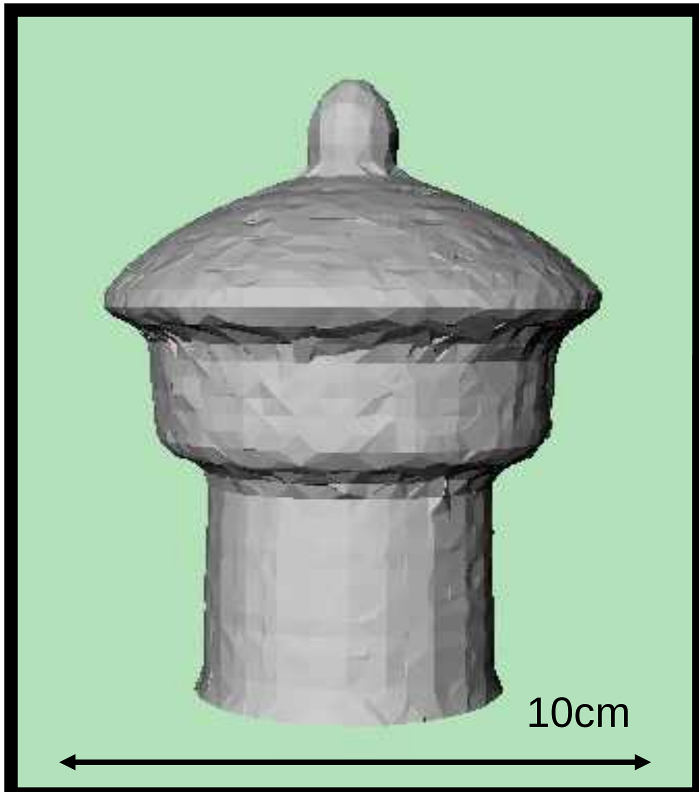

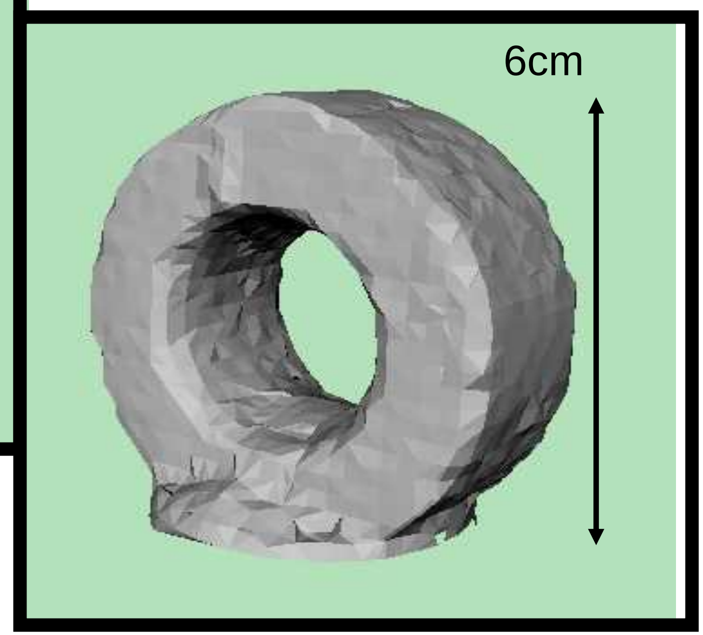

Space Carving: Experiments

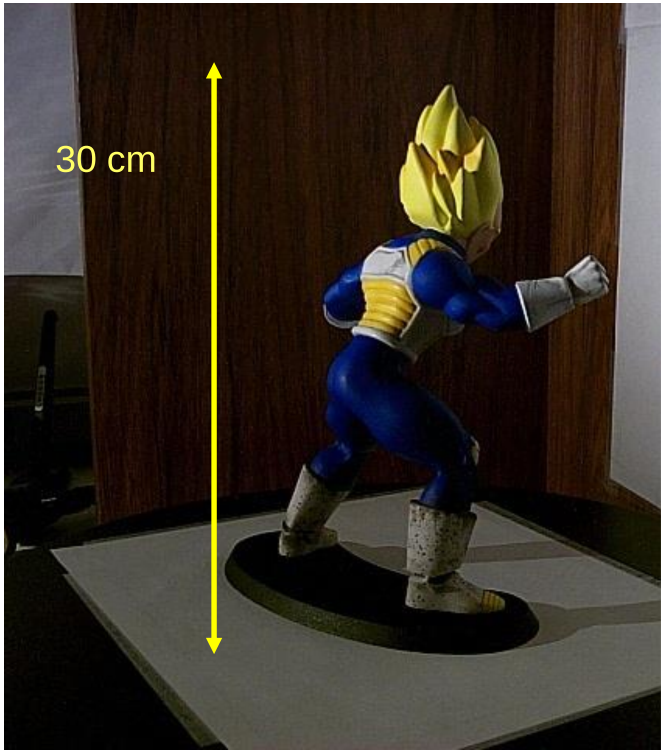

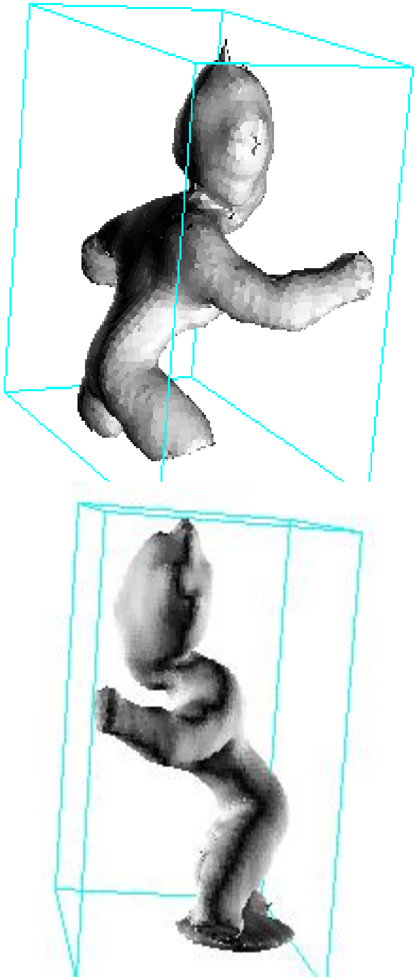

24 poses (15^{O})

voxel size = 1mm

Space Carving: Conclusions

- Robust
- Produce conservative estimates
- Concavities can be a problem
- Low-end commercial 3D scanners

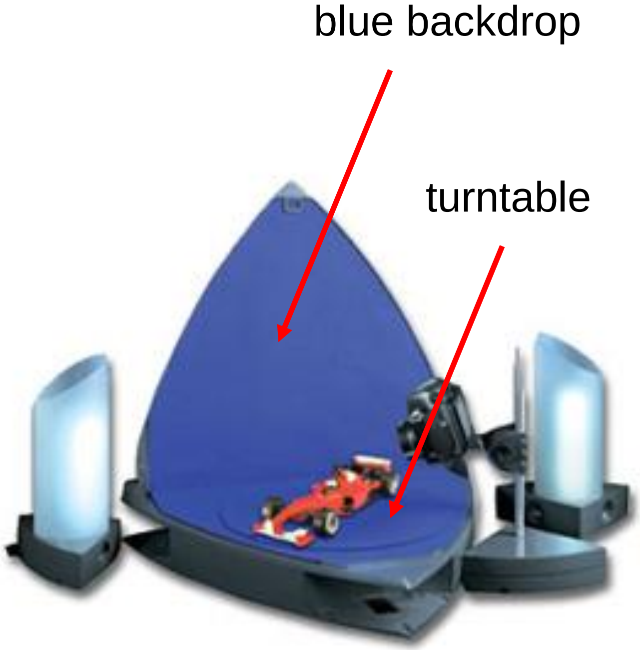

Space Carving: Conclusions

- Analyzing changes in apparent contours

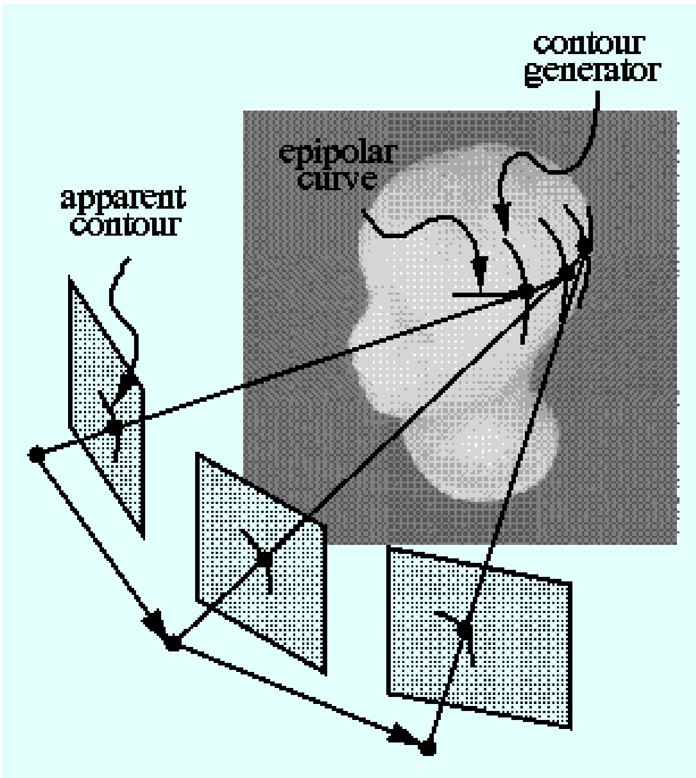

Picture from of Sato & Cipolla

- Giblin and Weiss (1987)
- Cipolla and Blake (1992)
- Vaillant and Faugeras (1992)
- Ponce ('92), Zheng('94)
- Furukawa et al. ('05…)

J. Sato and R. Cipolla. Affine reconstruction of curved surfaces from uncalibrated views of apparent contours. *PAMI*, 1999.

Lecture 8

SFM & Volumetric stereo

- SFM: Self-calibration
- Volumetric stereo:
  - Space carving
  - Shadow carving
  - Voxel carving

Reading:

[HZ] Chapters 19 "Auto-calibration"
[Szelisky] Chapter 7 "Structure from motion"
[Szelisky] Chapter 11 "Multi-view stereo"

Shape from Shadows

Volumetric Stereo

- Definition
- Shape from Contours
- Shape from Shadows
- Voxel coloring

Shape from Shadows

- Self-shadows are visual cues for shape recovery

Self-shadows indicate concavities
(no modeled by contours)

Shadow Carving: The Setup

[Savarese et al ’01]

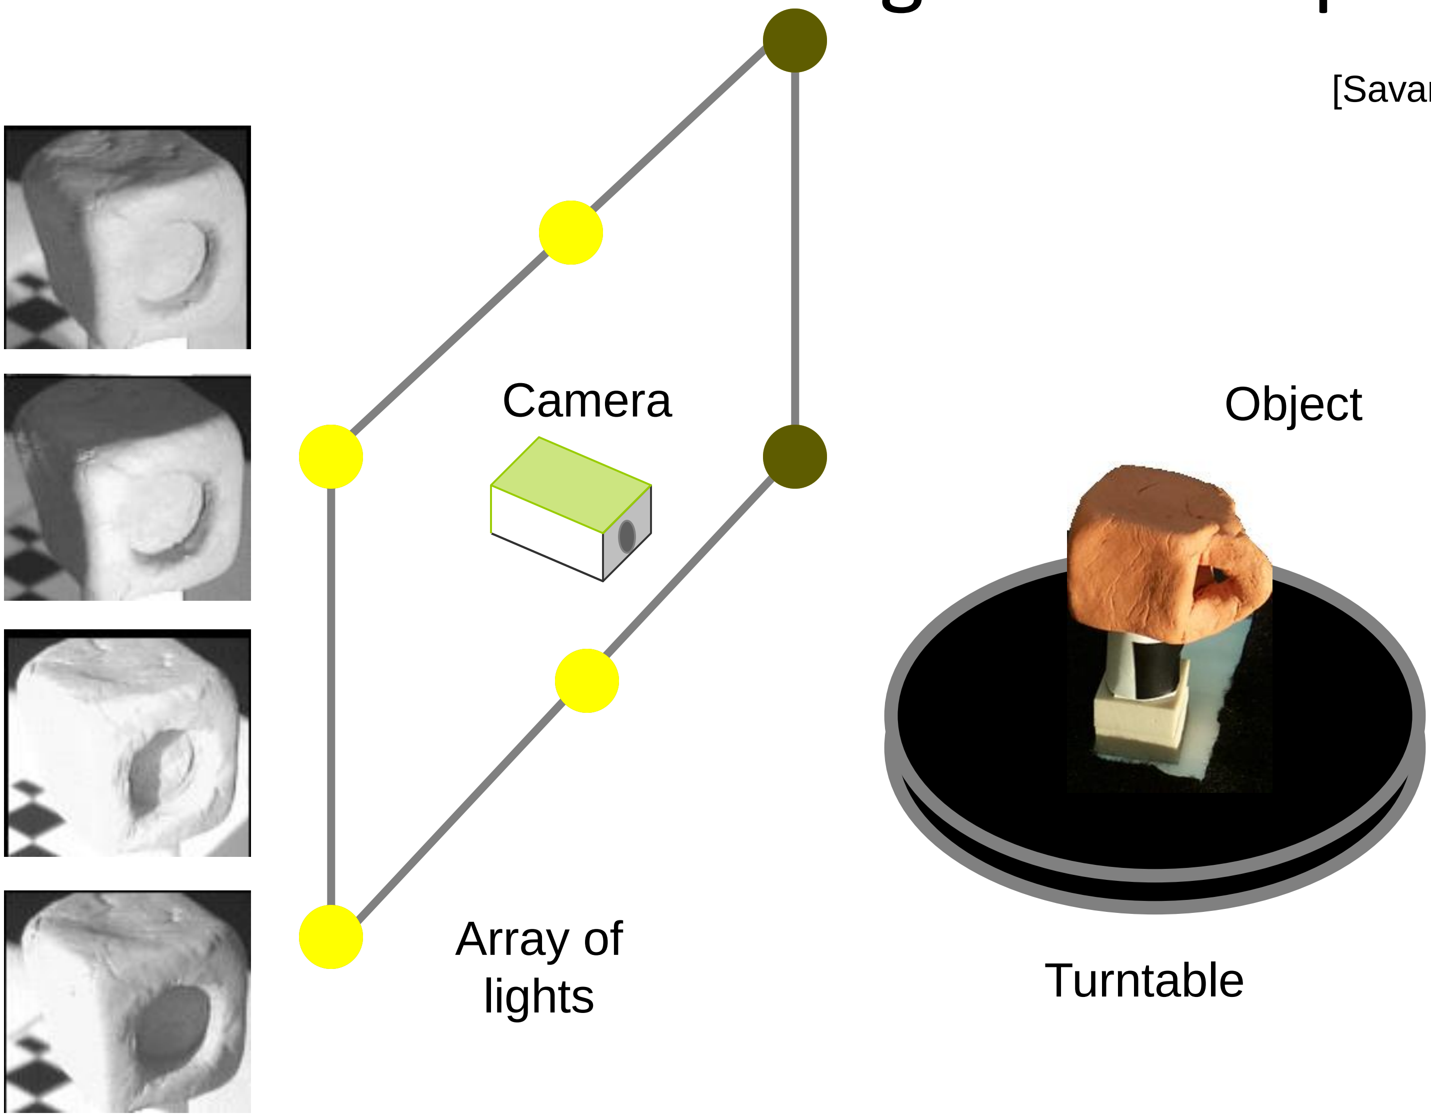

Shadow Carving: The Setup

[Savarese et al '01]

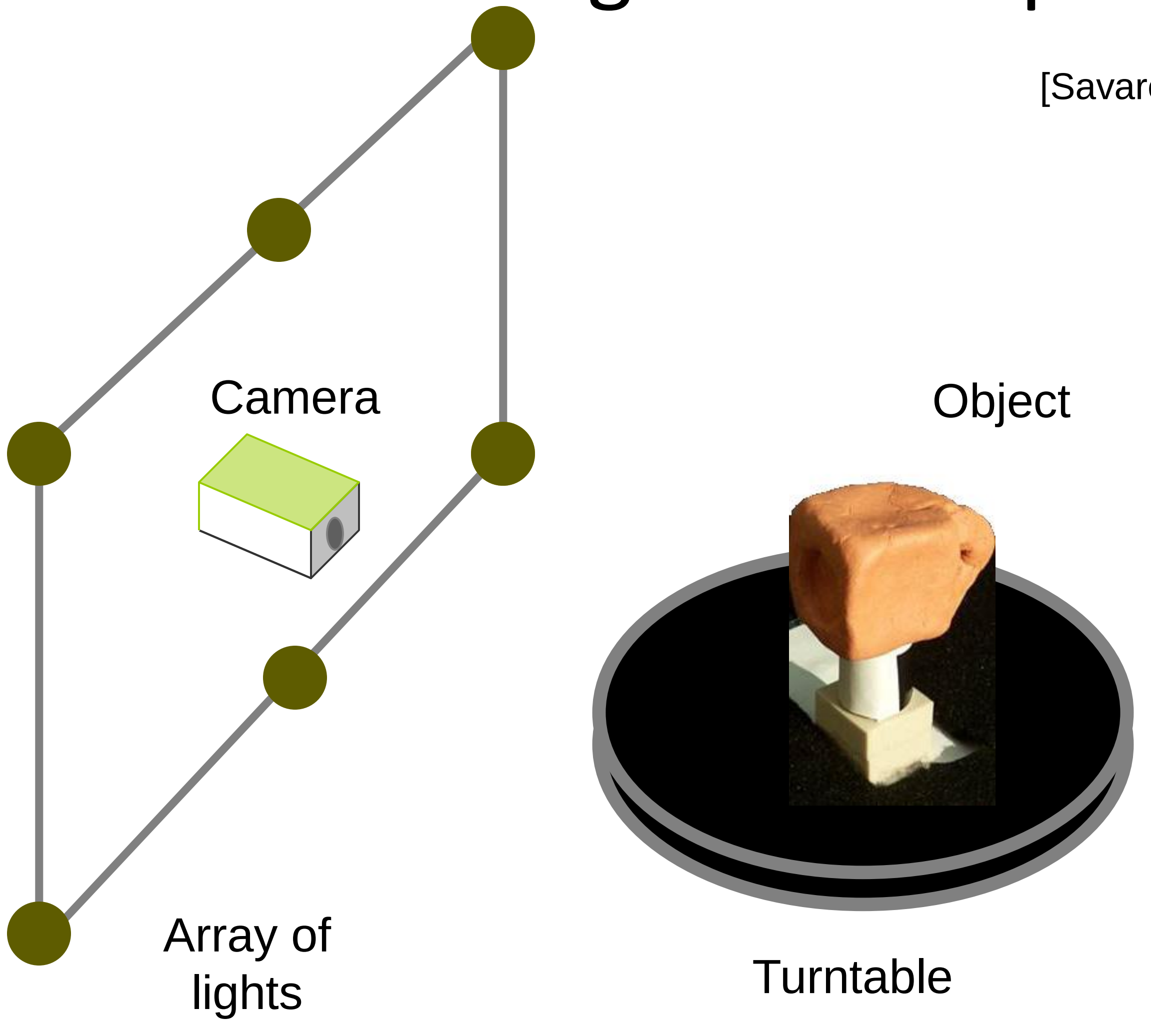

Shadow Carving

[Savarese et al '01]

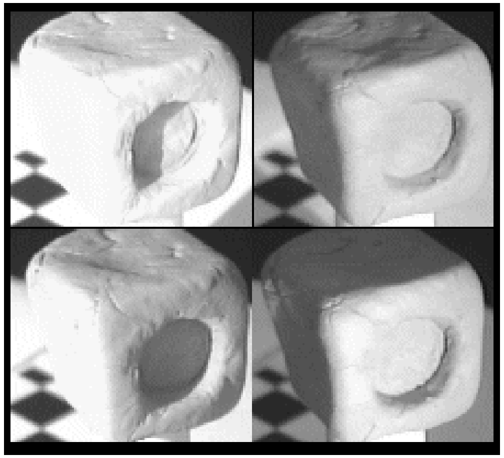

Self-shadows

+

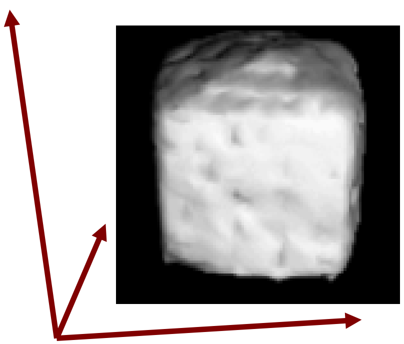

Object's upper bound

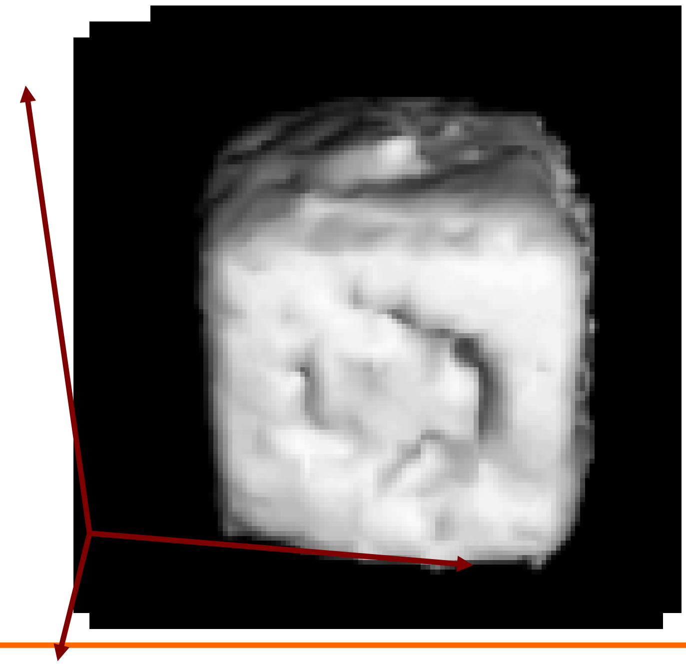

Robust with respect to shadow estimates

Object with arbitrary topology (no 2.5D terrains)

Algorithm: Step k

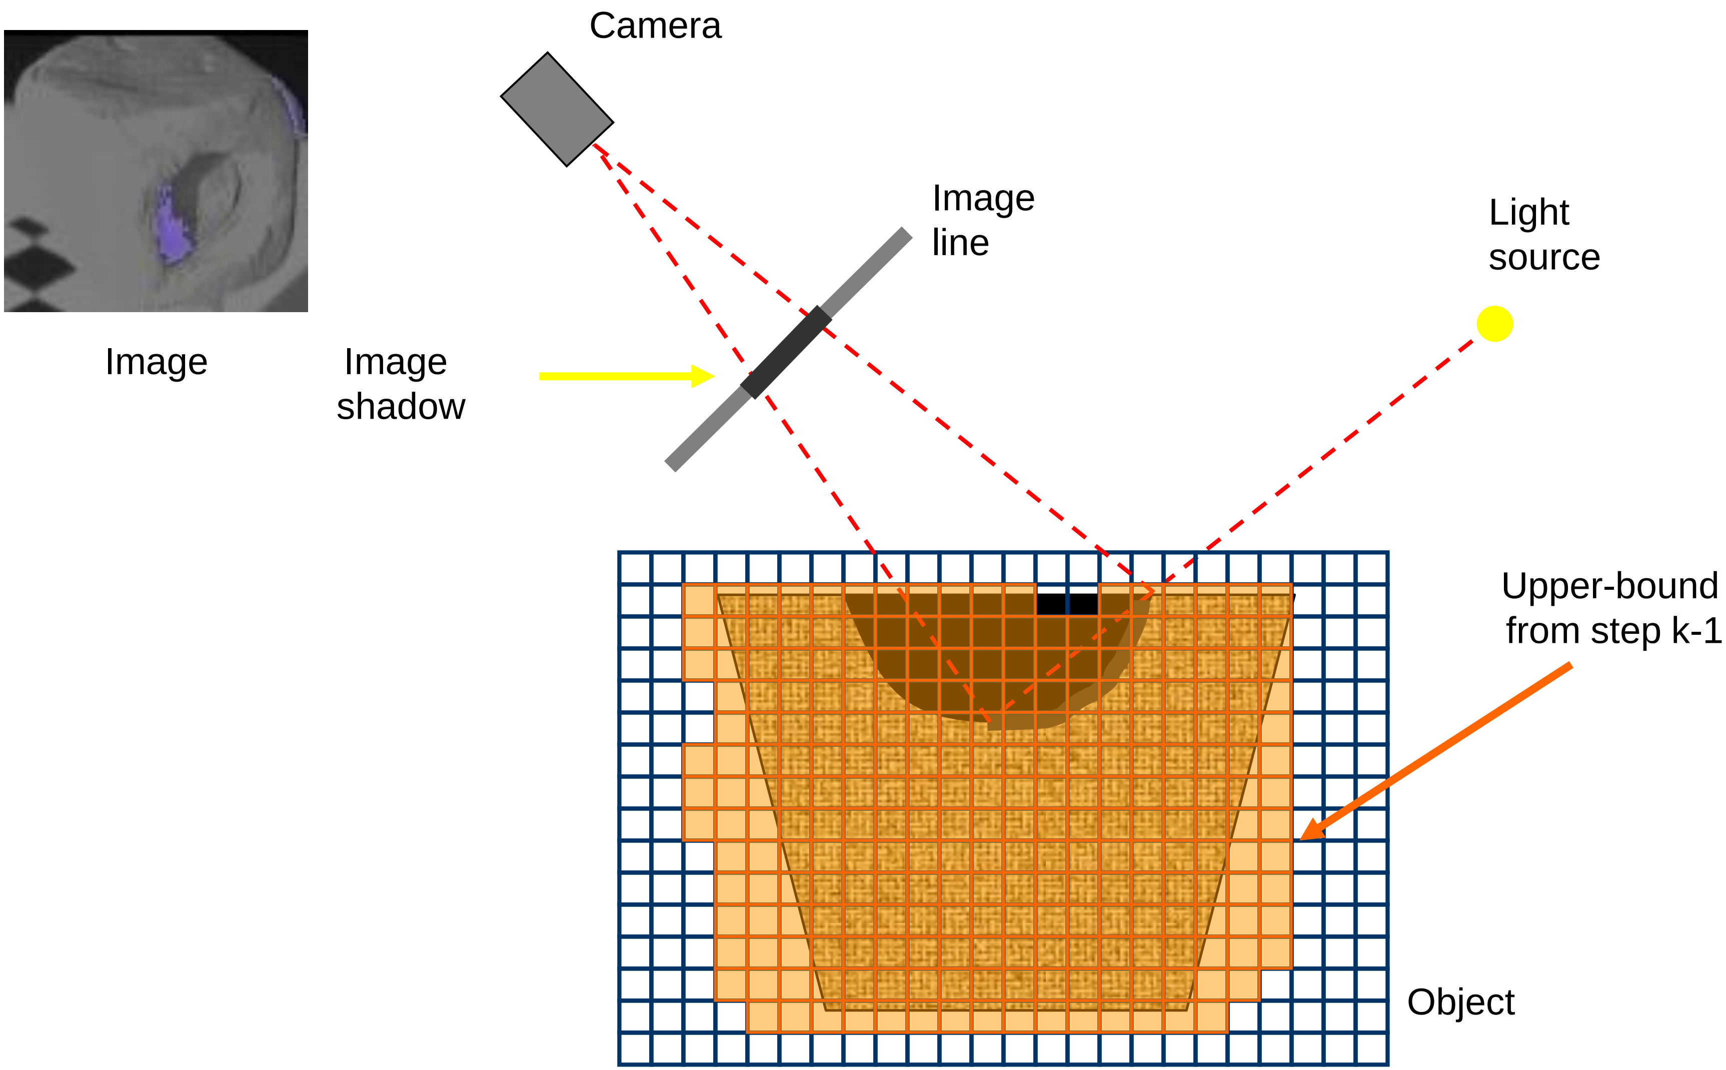

Algorithm: Step k

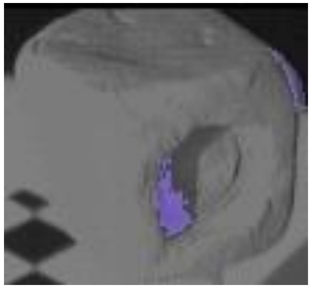

Image

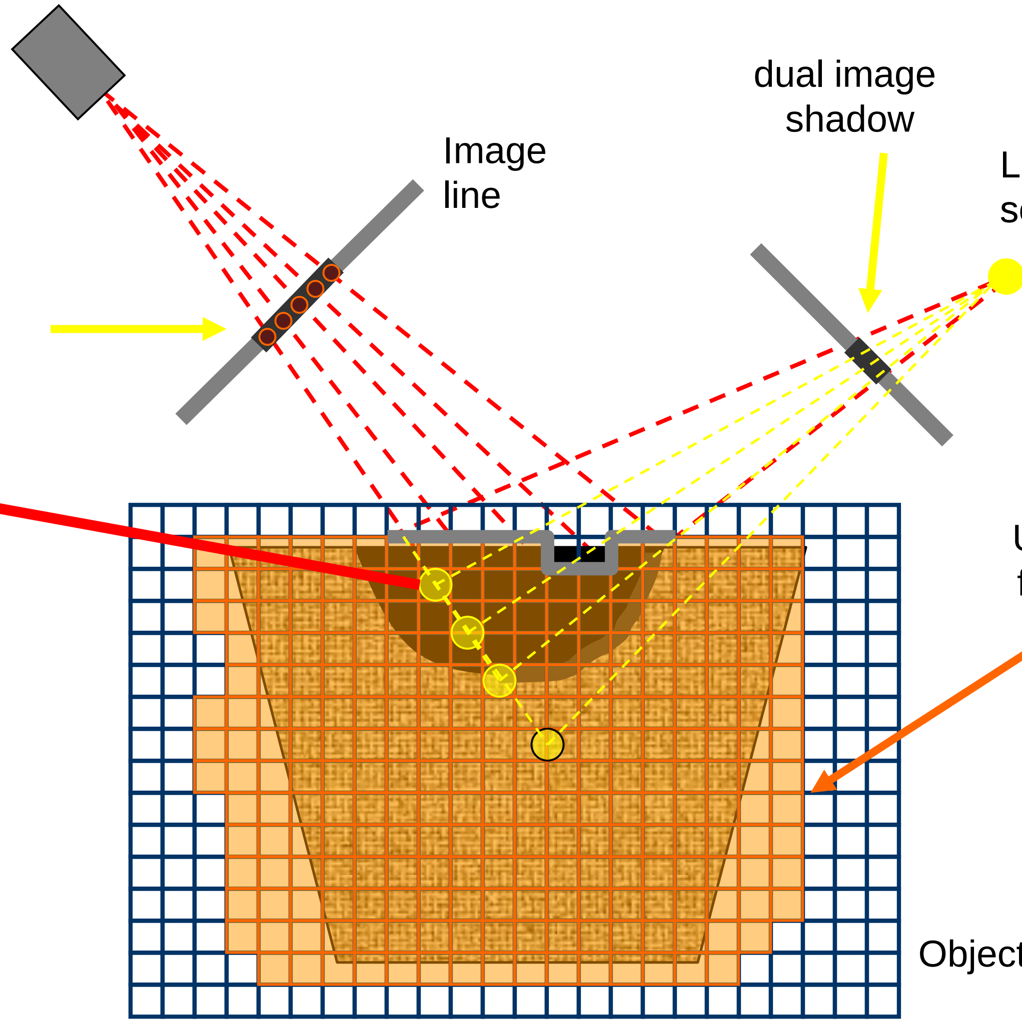

Algorithm: Step k

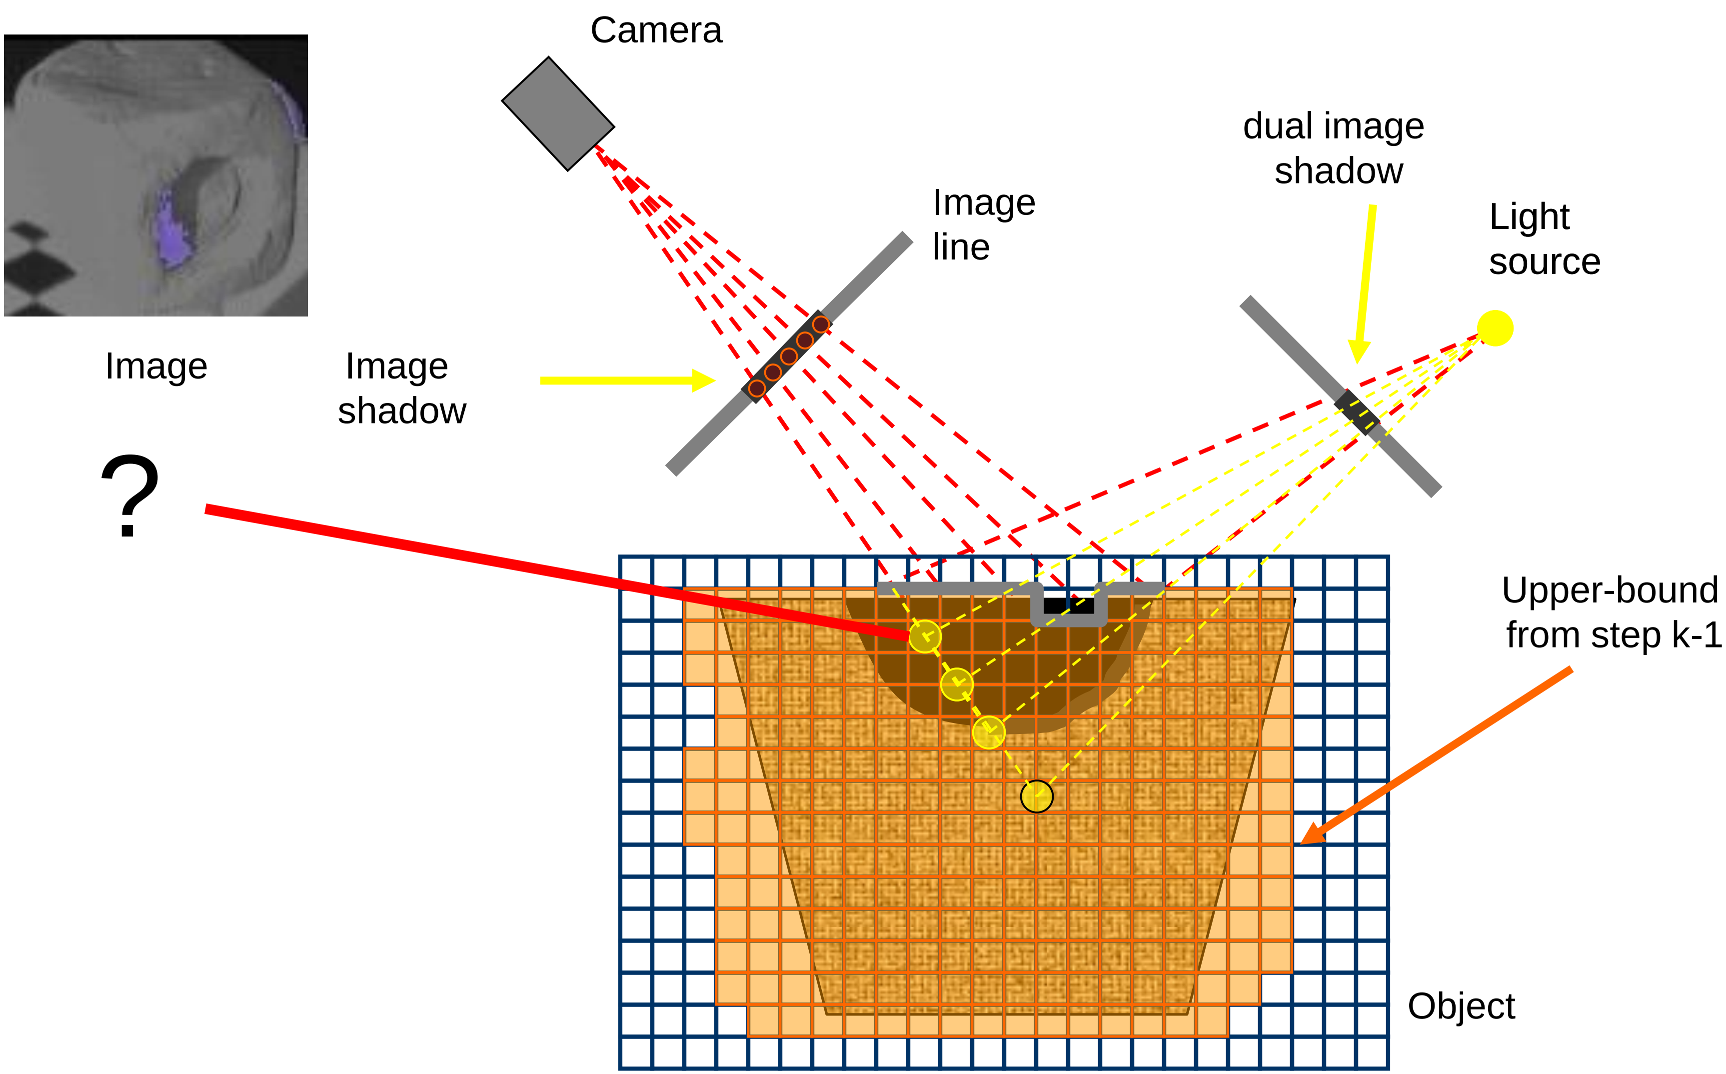

Algorithm: Step k

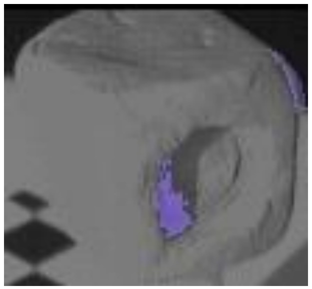

Image

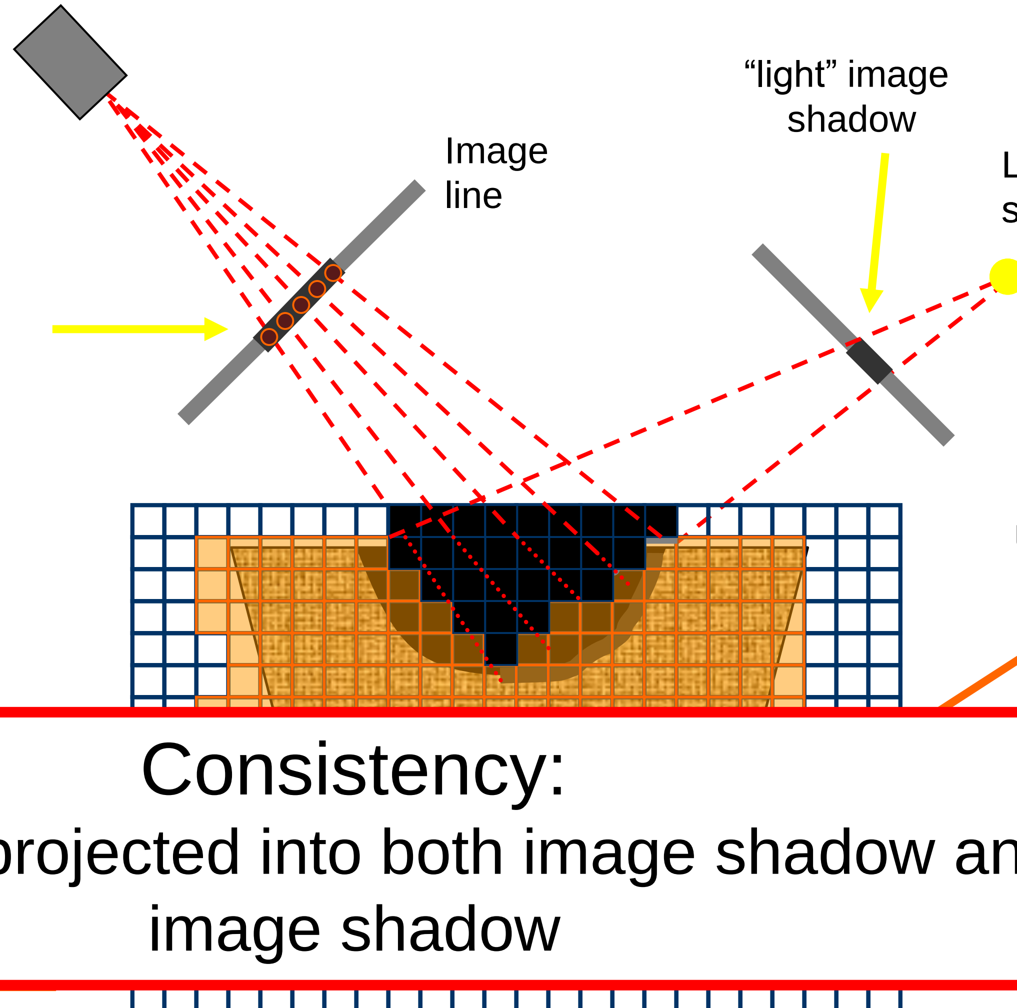

No further volume can be removed

Carving process

Consistency:

A voxel must be projected into both image shadow and dual image shadow

Algorithm: Step k

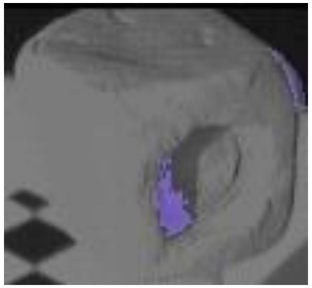

Image

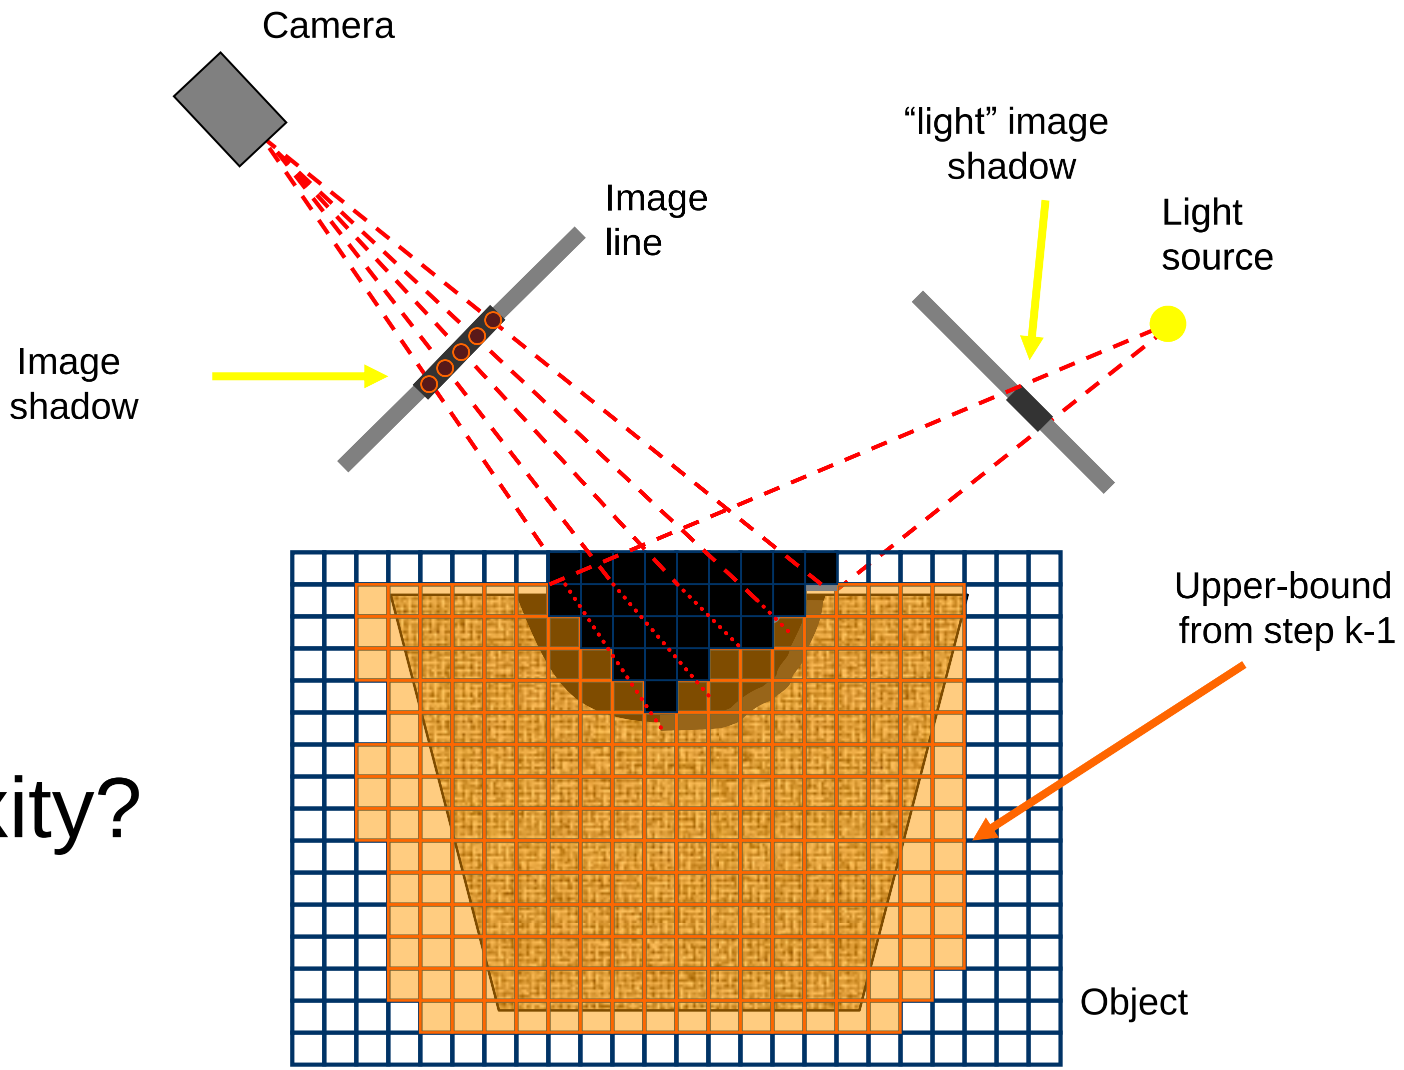

Complexity?

$O(2N^3)$

Simulating the System with 3D Studio Max

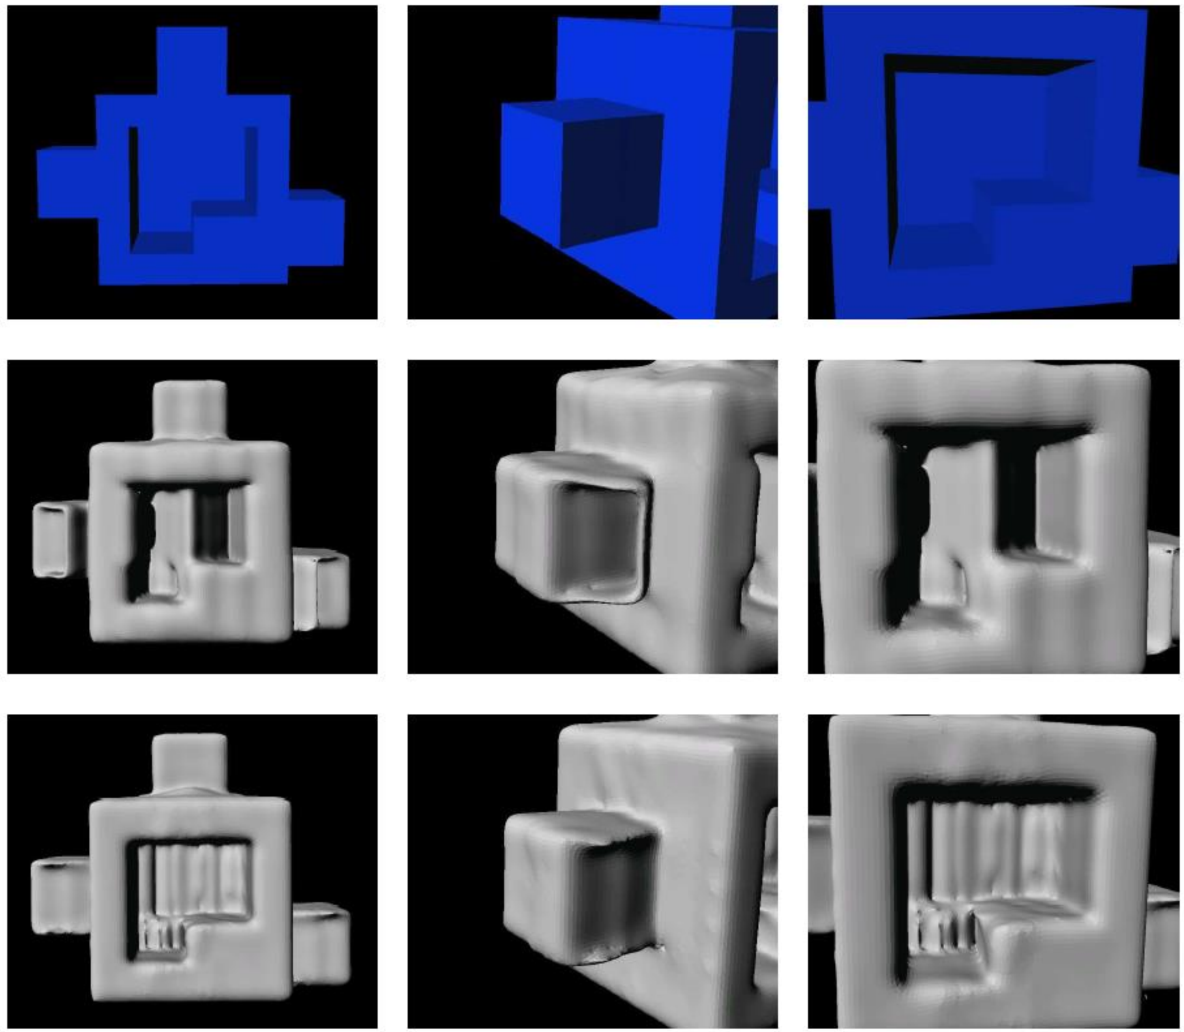

- 24 positions
- 4 lights

- 72 positions
- 8 lights

Simulating the System with 3D Studio Max

- 16 positions
- 4 lights

Simulating the System with 3D Studio Max

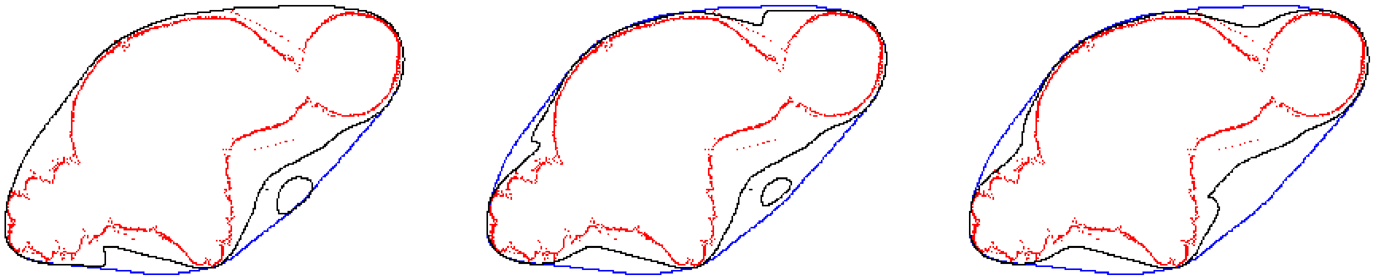

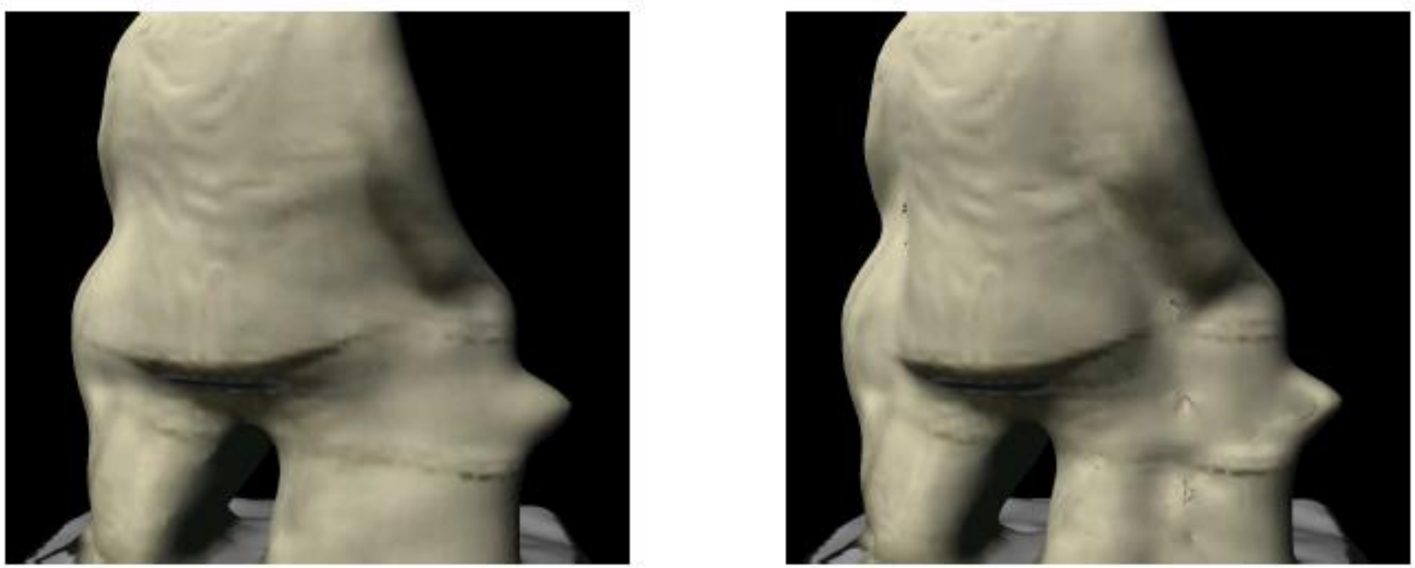

Shadow Carving: Summary

- Produces a conservative volume estimate
- Accuracy depending on view point and light source number
- Limitations with specular & low albedo regions

Lecture 8
SFM & Volumetric stereo

- SFM: Self-calibration
- Volumetric stereo:
  - Space carving
  - Shadow carving
  - Voxel carving

Reading:
[HZ] Chapters 19 "Auto-calibration"
[Szelisky] Chapter 7 "Structure from motion"
[Szelisky] Chapter 11 "Multi-view stereo"

Voxel Coloring

[Seitz & Dyer ('97)]
[R. Collins (Space Sweep, '96)]

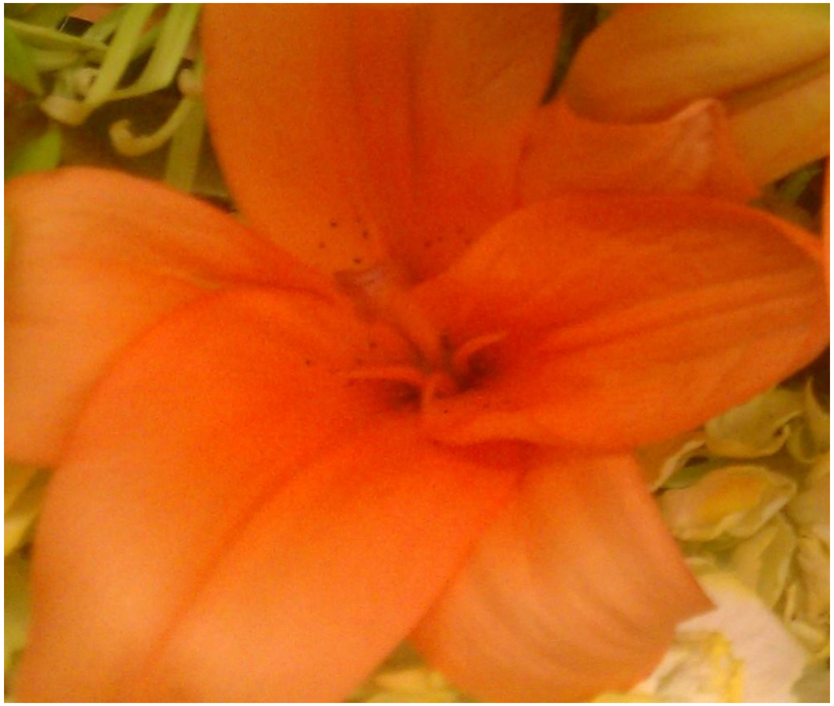

- Color/photo-consistency
- Jointly model structure and appearance

Basic Idea

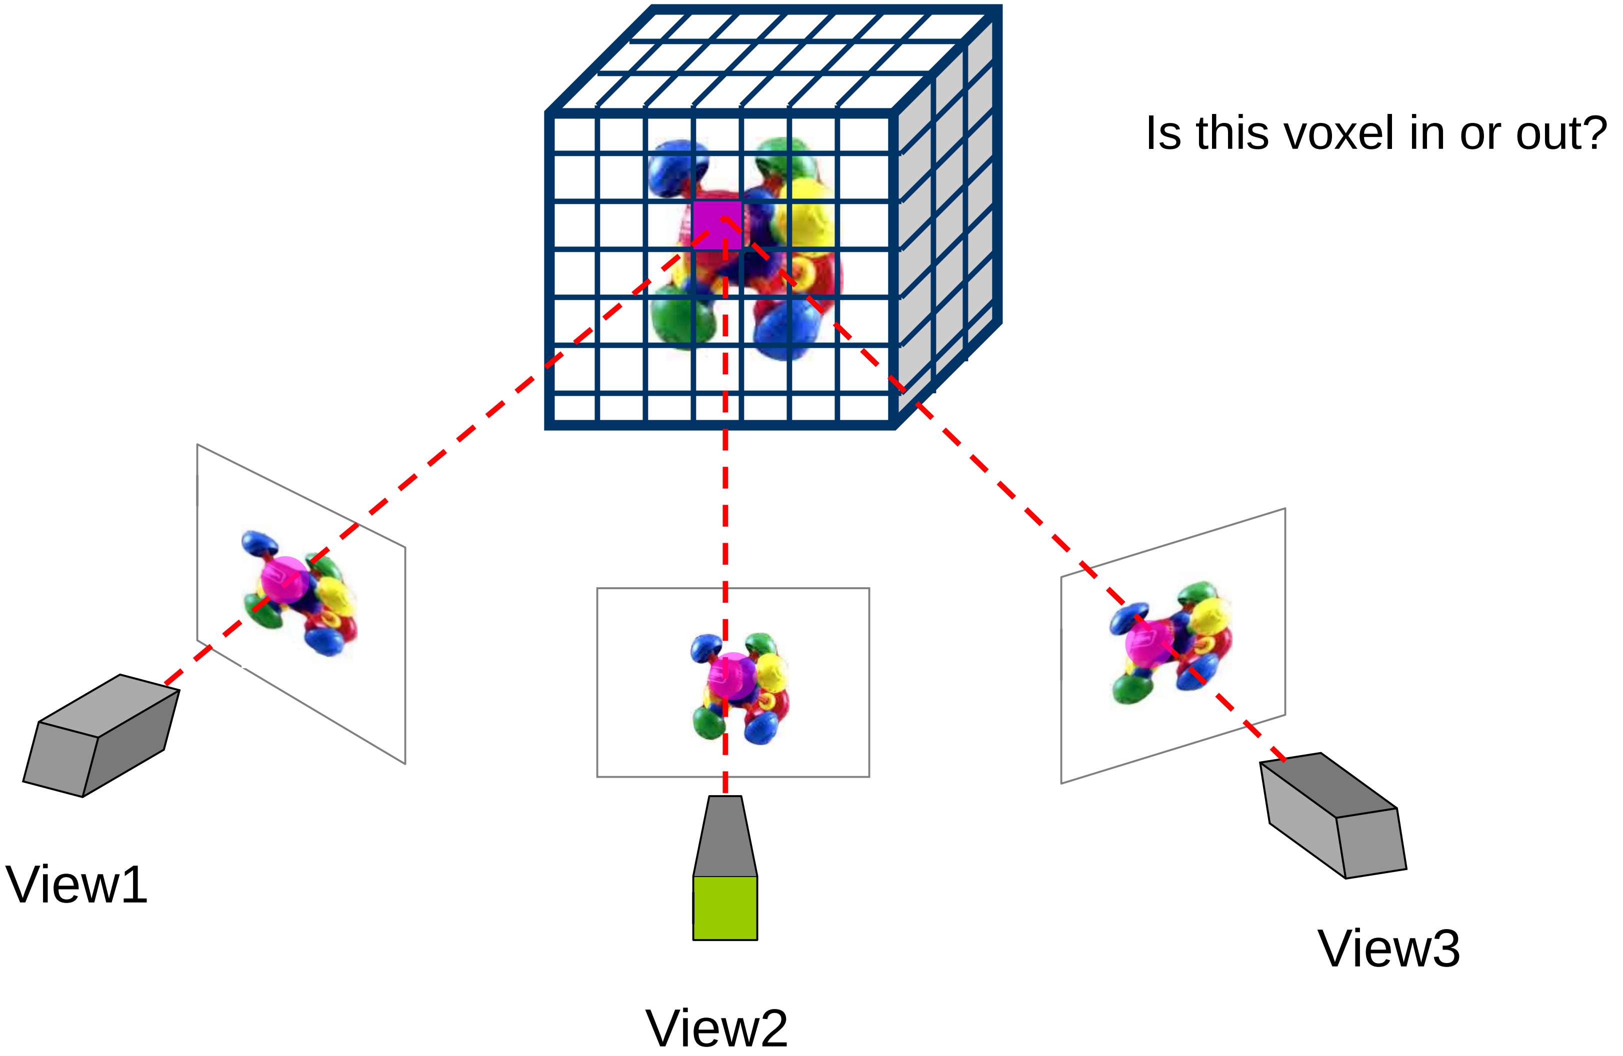

Basic Idea

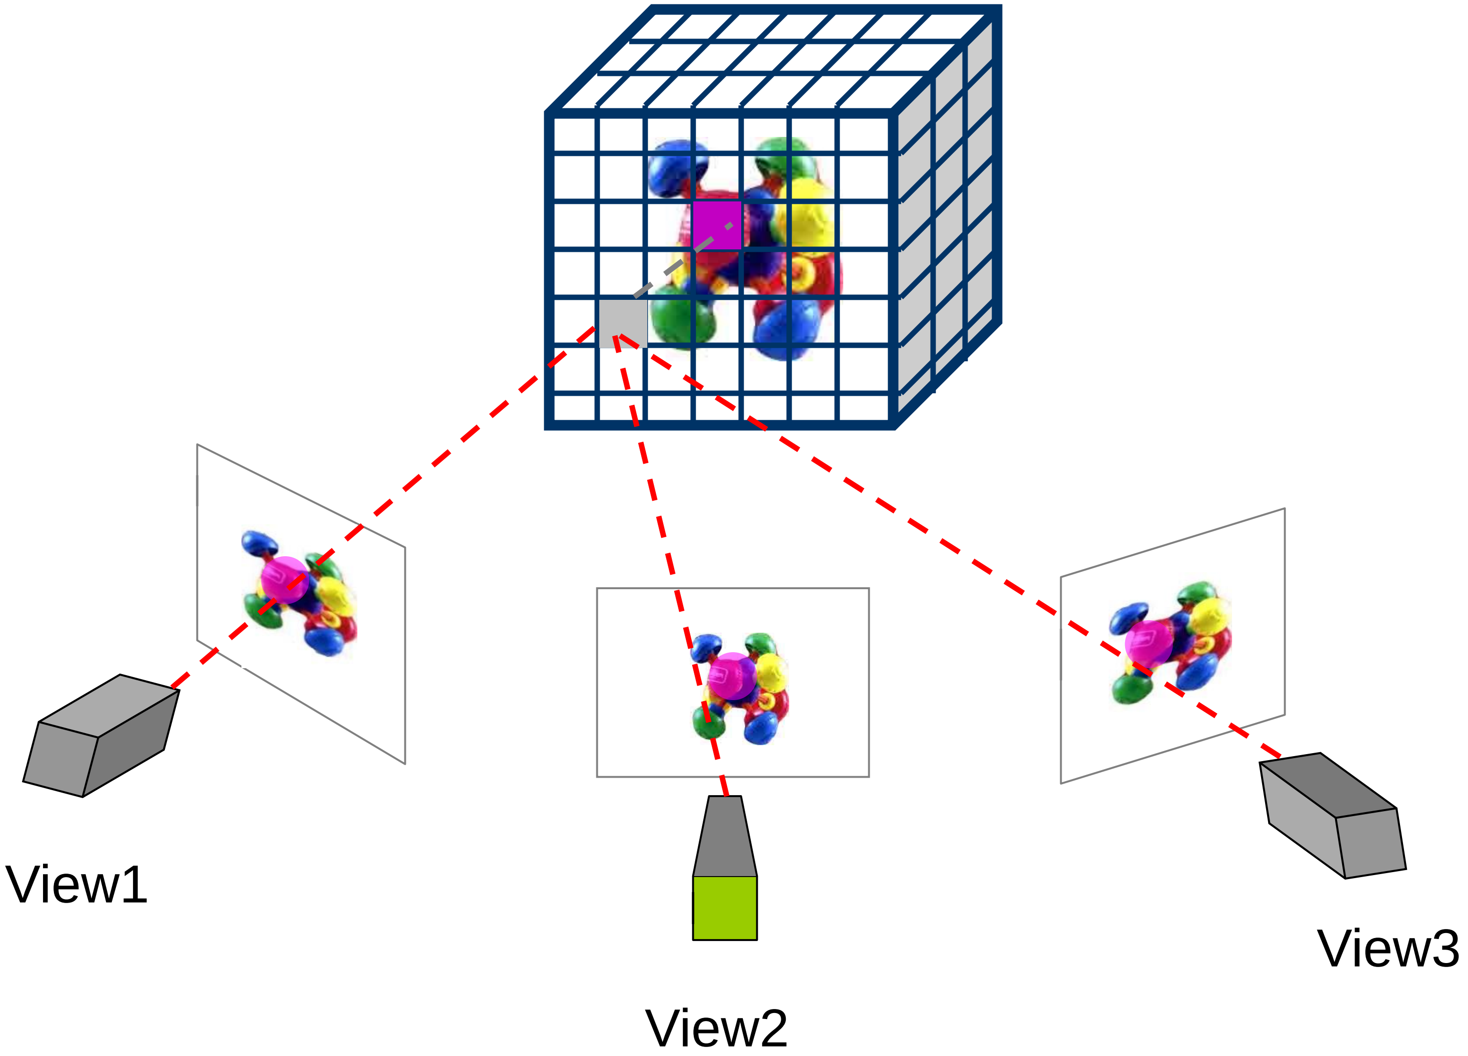

Uniqueness

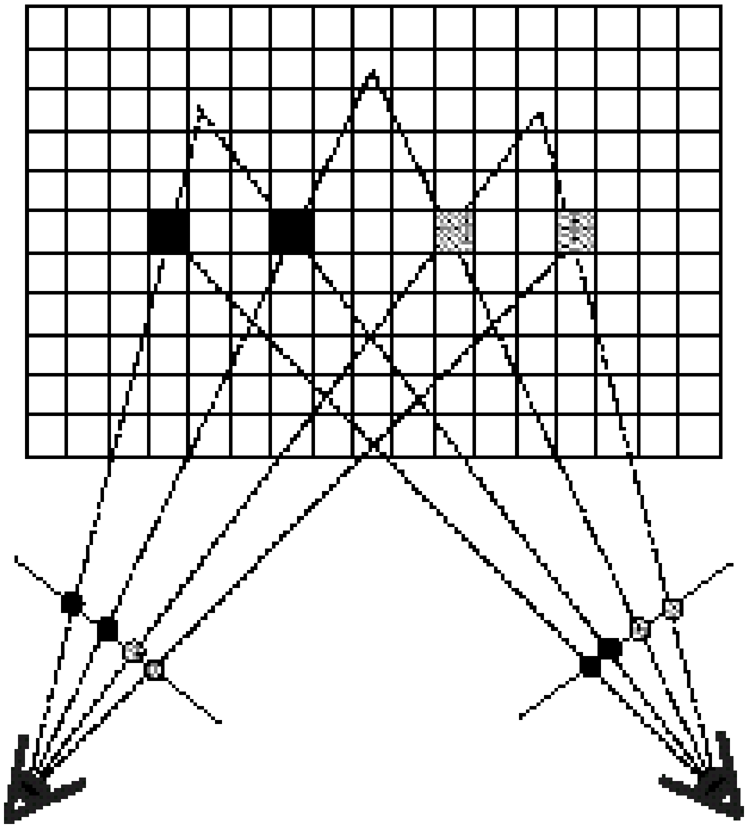

- Multiple consistent scenes

Uniqueness

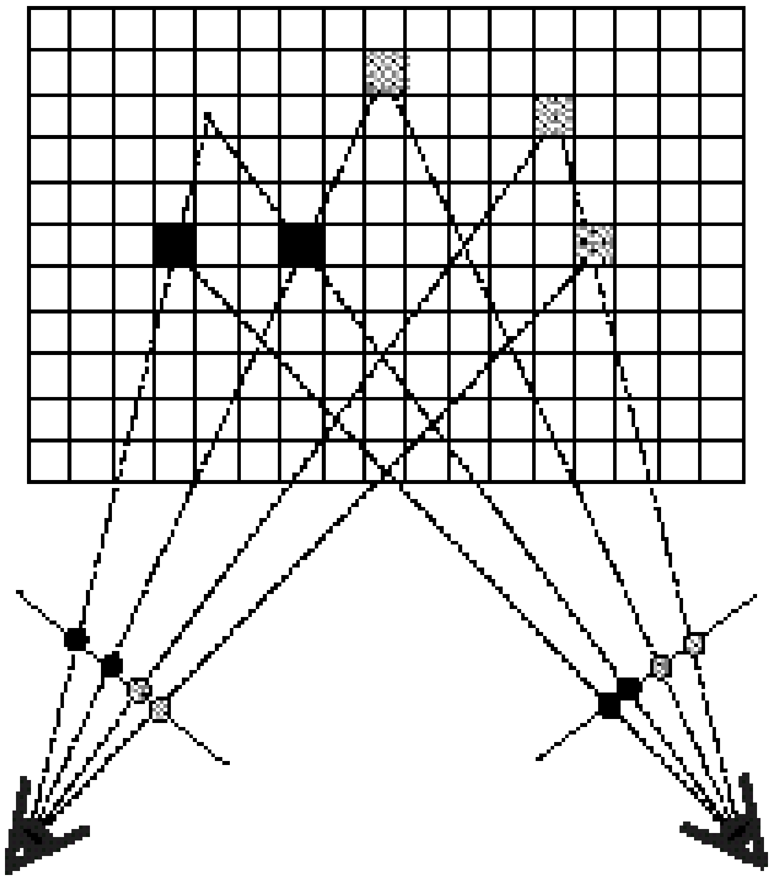

- Multiple consistent scenes

How to fix this?
Need to use a visibility constraint

The Algorithm

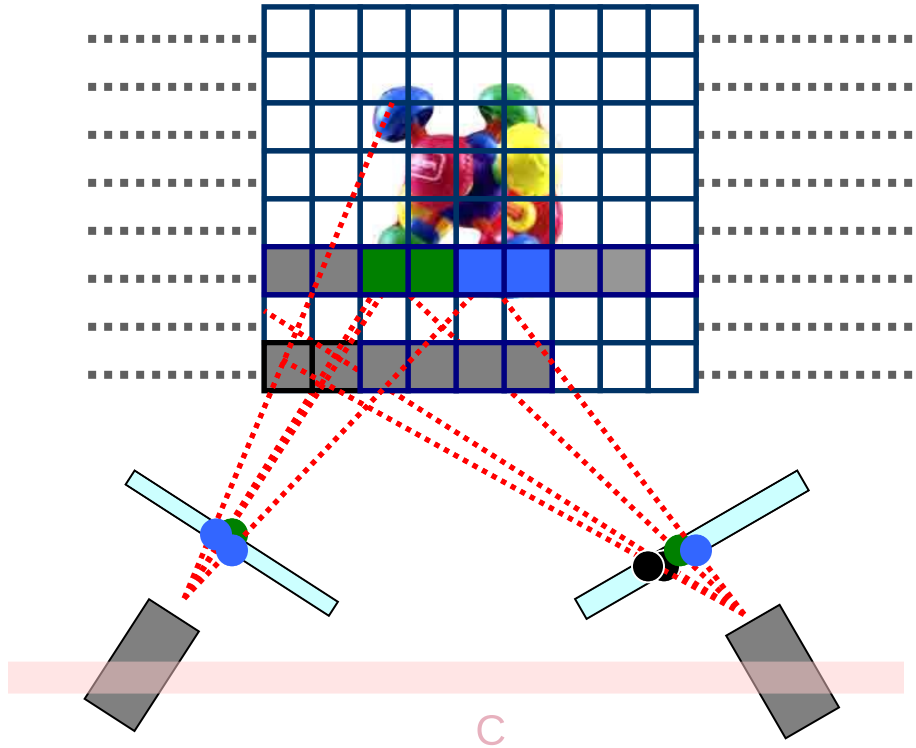

Algorithm Complexity

- Voxel coloring visits each N^3 voxels only once
- Project each voxel into L images

$\rightarrow$ $O(L\ N^3)$

NOTE: not function of the number of colors

Photoconsistency Test

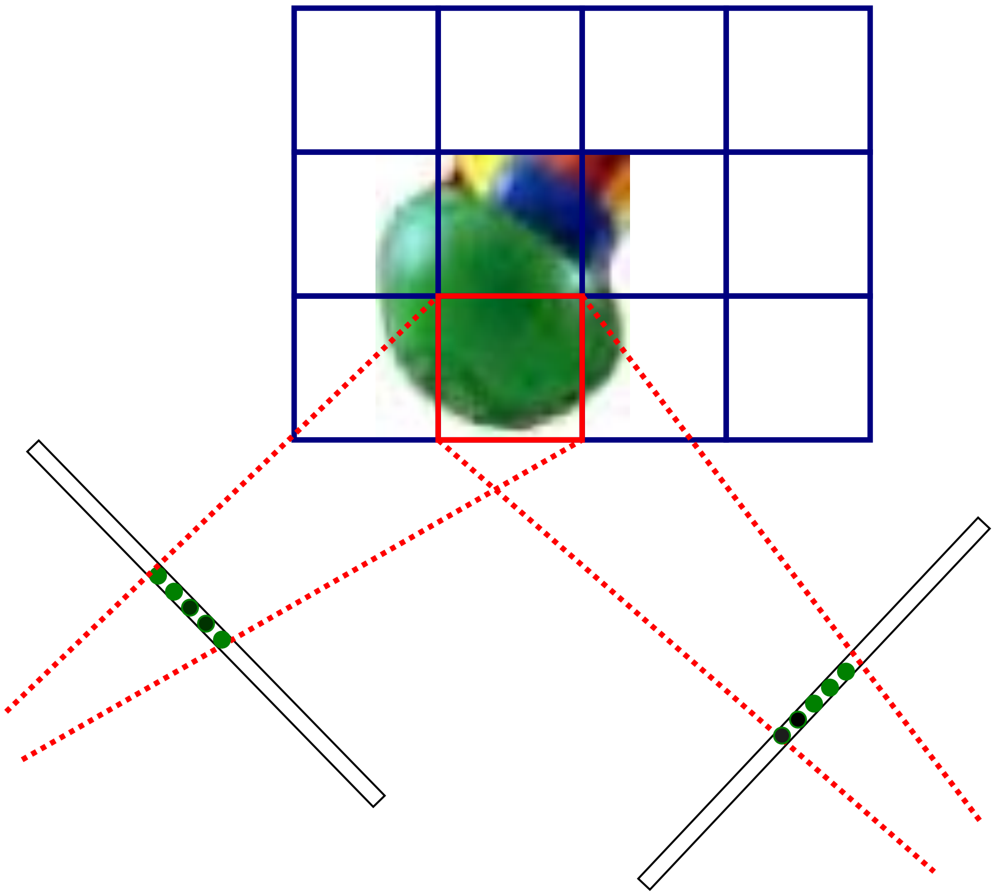

λ = corr (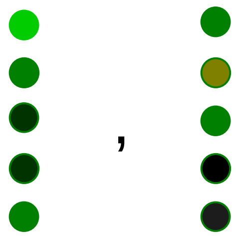 ,)

If λ > Thresh ➔ voxel consistent

A Critical Assumption: Lambertian Surfaces

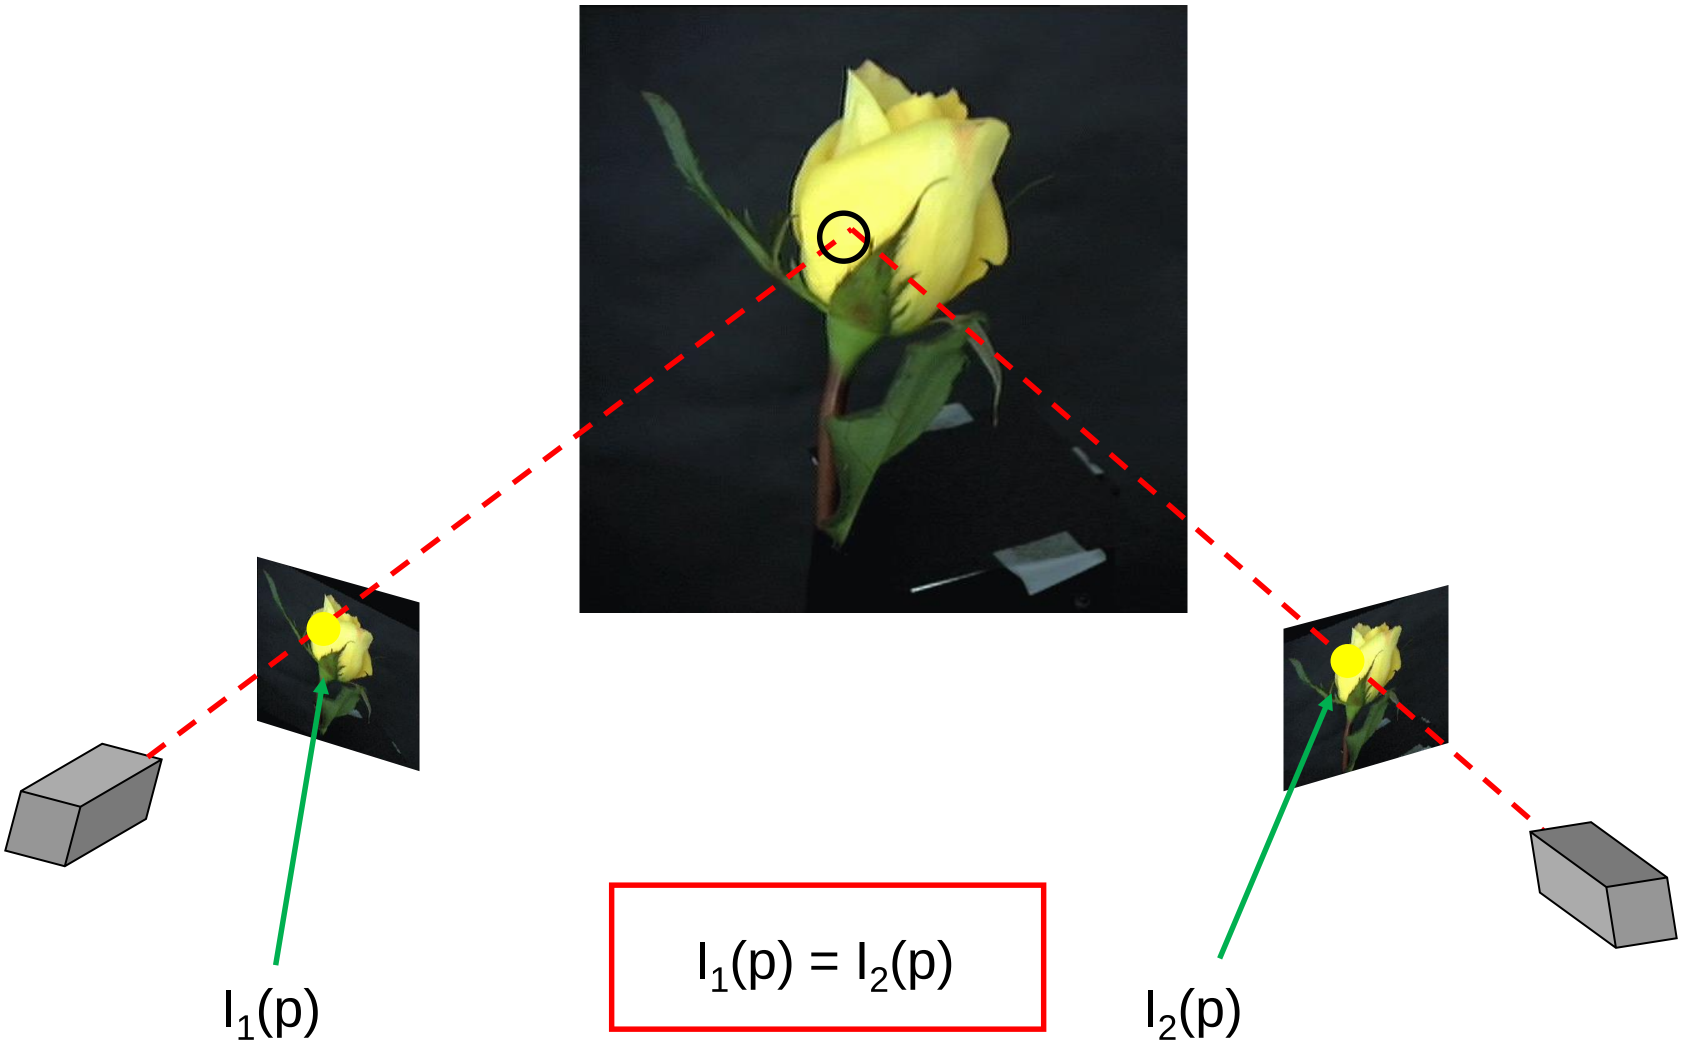

Non Lambertian Surfaces

Experimental Results

Dinosaur

72 k voxels colored
7.6 M voxels tested
7 min to compute on a 250MHz

Image source: http://www.cs.cmu.edu/~seitz/vcolor.html

Experimental Results

Flower

70 k voxels colored
7.6 M voxels tested
7 min to compute on a 250MHz

Image source: http://www.cs.cmu.edu/~seitz/vcolor.html

Experimental Results

Room + weird people

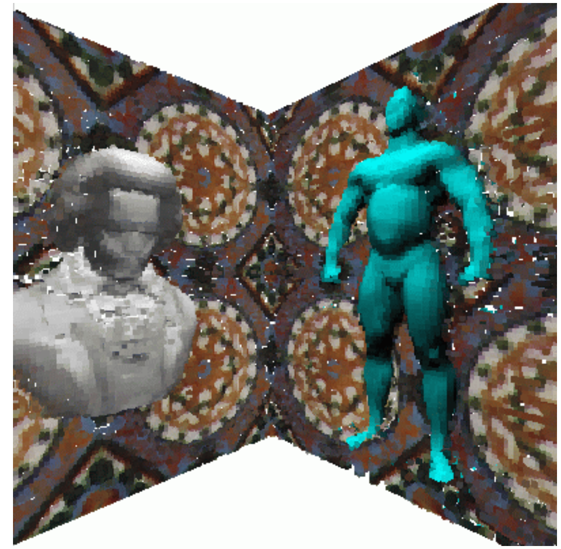

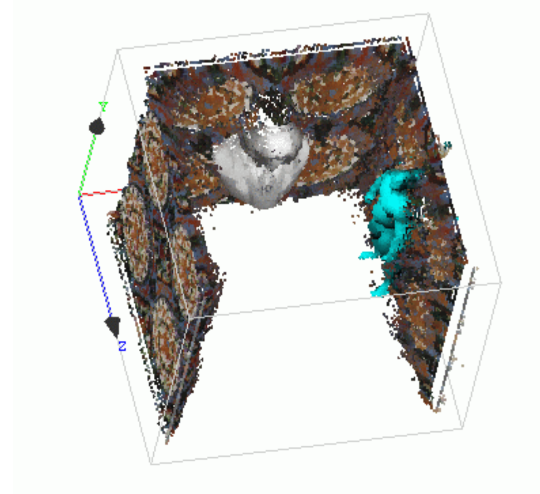

Image source: http://www.cs.cmu.edu/~seitz/vcolor.html

Voxel Coloring: Conclusions

- Good things
  - Model intrinsic scene colors and texture
  - No assumptions on scene topology

- Limitations:
  - Constrained camera positions
  - Lambertian assumption

Space Carving

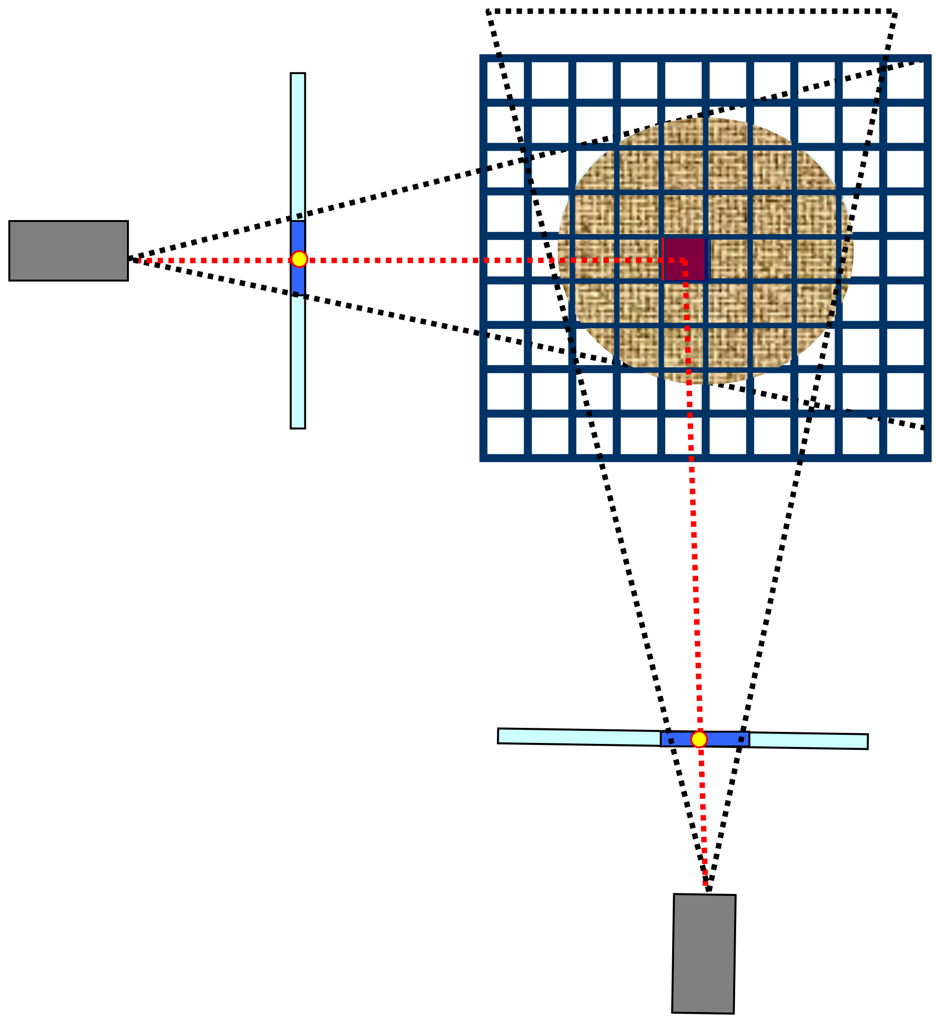

- Space carving is a binary voxel coloring
- No visibility assumption is needed

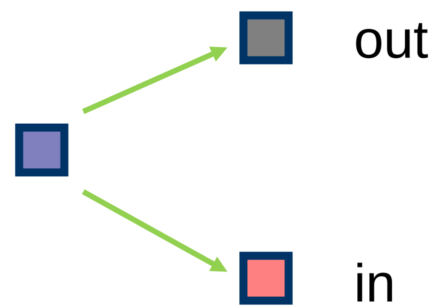

Further Contributions

- A Theory of Space Carving [Kutulakos & Seitz ’99]
  - Voxel coloring in more general framework
  - No restrictions on camera position

- Probabilistic Space Carving

[Broadhurst & Cipolla, ICCV 2001]
[Bhotika, Kutulakos et. al, ECCV 2002]

Next lecture…

Fitting and Matching