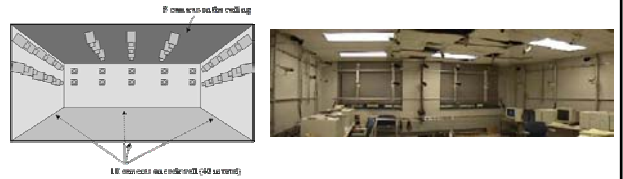## Announcements

- Project 3 grades out today, write-ups online asap
- Project 4 extension: Thurs, 11:59pm

## Multiview stereo
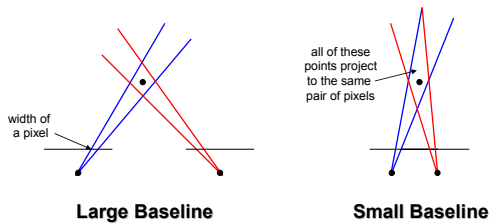


CMU's 3D Room

Readings
- S. M. Seitz and C. R. Dyer, Photorealistic Scene Reconstruction by Voxel Coloring, *International Journal of Computer Vision*, 35(2), 1999, pp. 151-173.
  - http://www.cs.washington.edu/homes/seitz/papers/ijcv99.pdf

## Choosing the stereo baseline



all of these points project to the same pair of pixels
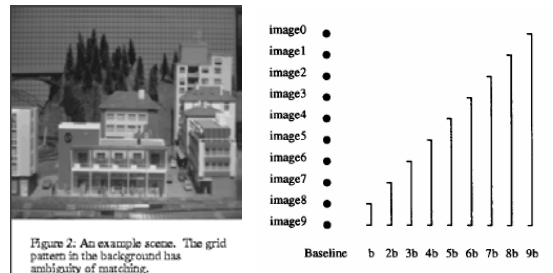
width of a pixel

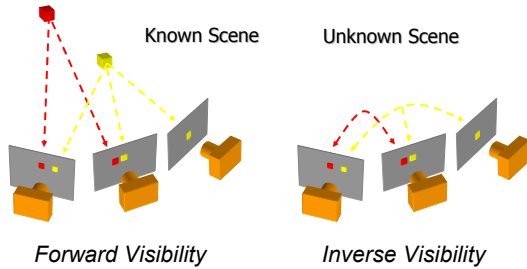**Large Baseline**          **Small Baseline**

What's the optimal baseline?
- Too small:  large depth error
- Too large:  difficult search problem

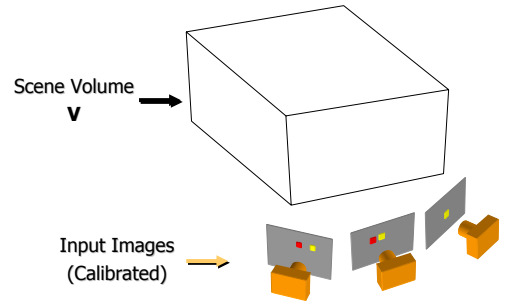## The Effect of Baseline on Depth Estimation



Figure 2: An example scene. The grid pattern in the background has ambiguity of matching.

image0
image1
image2
image3
image4
image5
image6
image7
image8
image9

Baseline    b   2b  3b  4b  5b  6b  7b  8b  9b

pixel matching score



Fig. 5. SSD values versus inverse distance: (a) $B = h$; (b) $B = 2h$; (c) $B = 3h$; (d) $B = 4h$; (e) $B = 5h$; (f) $B = 6h$; (g) $B = 7h$; (h) $B = 8h$. The horizontal axis is normalized such that $hF = 1$.

Fig. 6. Combining two stereo pairs with different baselines.

Fig. 7. Combining multiple baseline stereo pairs.

# Multibaseline Stereo

## Basic Approach
- Choose a reference view
- Use your favorite stereo algorithm BUT
  - > replace two-view SSD with SSD over all baselines

## Limitations
- Must choose a reference view (bad)
- Visibility!

## CMU's 3D Room Video



Figure 7: The CMU Video-Rate Stereo Machine Prototype System: (a) camera head; (b) processor boards.

## The visibility problem

Which points are visible in which images?



Known Scene

Unknown Scene

*Forward Visibility*

*Inverse Visibility*

## Volumetric stereo



Scene Volume
**V**

Input Images
(Calibrated)

**Goal:** **Determine occupancy, "color" of points in V**

## Discrete formulation: Voxel Coloring



Discretized
Scene Volume

Input Images
(Calibrated)
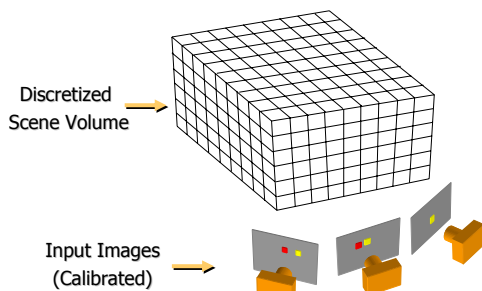
**Goal:** **Assign RGBA values to voxels in V**
***photo-consistent*** **with images**

## Complexity and computability



Discretized
Scene Volume
$N^3$ **voxels**
$C$ **colors**

True
Scene

Photo-Consistent
Scenes

All Scenes ($C^{N^3}$)
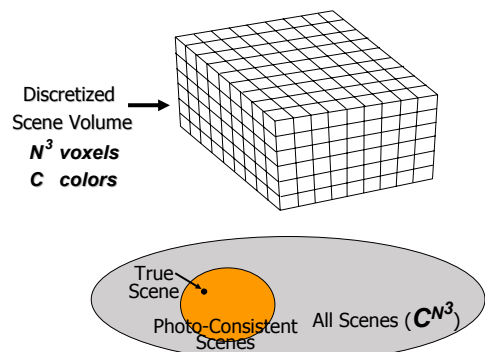
## Issues

### Theoretical Questions
- Identify class of *all* photo-consistent scenes

### Practical Questions
- How do we compute photo-consistent models?

## Voxel coloring solutions

1. **C=2 (shape from silhouettes)**
   - Volume intersection [Baumgart 1974]
     > For more info: *Rapid octree construction from image sequences.* R. Szeliski, CVGIP: Image Understanding, 58(1):23-32, July 1993. (this paper is apparently not available online) or
     > W. Matusik, C. Buehler, R. Raskar, L. McMillan, and S. J. Gortler, *Image-Based Visual Hulls*, SIGGRAPH 2000 ( pdf 1.6 MB )

2. **C unconstrained, viewpoint constraints**
   - Voxel coloring algorithm [Seitz & Dyer 97]

3. **General Case**
   - Space carving [Kutulakos & Seitz 98]

## Reconstruction from Silhouettes (C = 2)

Binary Images ⟶
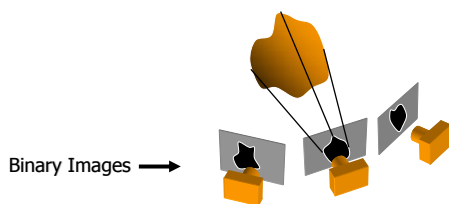
**Approach:**
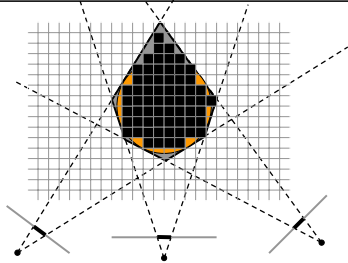- *Backproject* each silhouette
- Intersect backprojected volumes

## Volume intersection
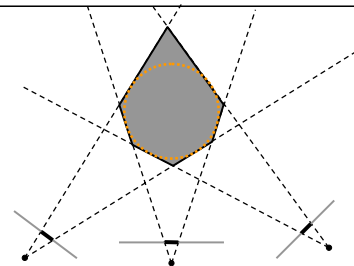
**Reconstruction Contains the True Scene**
- But is generally not the same
- In the limit (all views) get *visual hull*
  > Complement of all lines that don't intersect S

## Voxel algorithm for volume intersection



Color voxel black if on silhouette in every image
- O( ? ), for M images, $N^3$ voxels
- Don't have to search $2^{N^3}$ possible scenes!

## Properties of Volume Intersection

Pros
- Easy to implement, fast
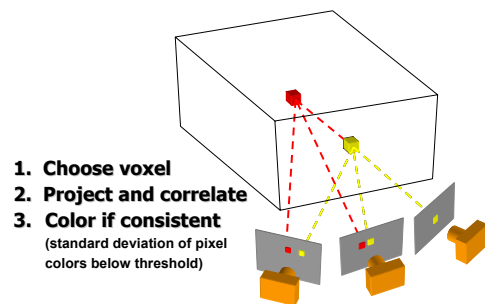- Accelerated via octrees [Szeliski 1993] or interval techniques [Matusik 2000]

Cons
- No concavities
- Reconstruction is not photo-consistent
- Requires identification of silhouettes

## Voxel Coloring Solutions

1. C=2 (silhouettes)
   - Volume intersection [Baumgart 1974]

2. C unconstrained, viewpoint constraints
   - Voxel coloring algorithm [Seitz & Dyer 97]
     > For more info: http://www.cs.washington.edu/homes/seitz/papers/ijcv99.pdf

3. General Case
   - Space carving [Kutulakos & Seitz 98]

## Voxel Coloring Approach



1. **Choose voxel**
2. **Project and correlate**
3. **Color if consistent**
   **(standard deviation of pixel colors below threshold)**

**Visibility Problem:** in which images is each voxel visible?

## Depth Ordering: visit occluders first!



Layers

Scene
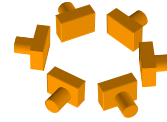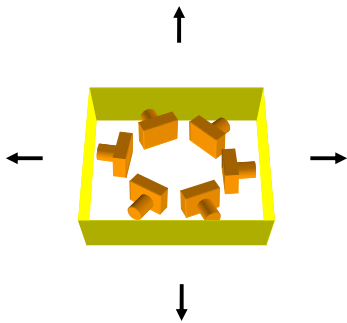Traversal
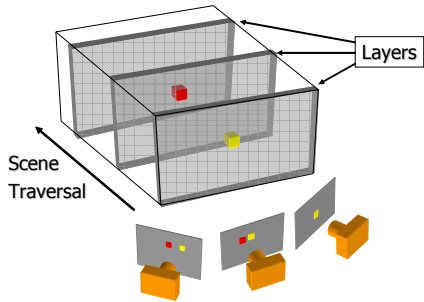
**Condition:** depth order is the *same for all input views*

## Panoramic Depth Ordering

- Cameras oriented in many different directions
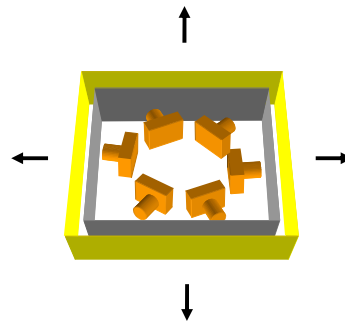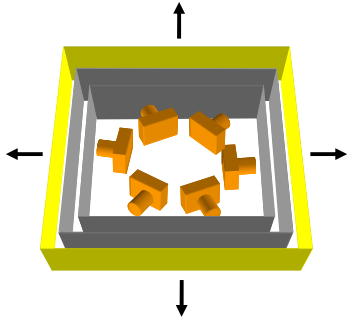- Planar depth ordering does not apply



## Panoramic Depth Ordering



**Layers radiate outwards from cameras**

## Panoramic Layering



**Layers radiate outwards from cameras**
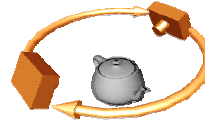
## Panoramic Layering



**Layers radiate outwards from cameras**

## Compatible Camera Configurations

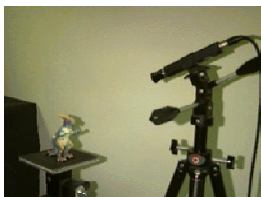Depth-Order Constraint
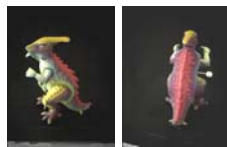- Scene outside convex hull of camera centers



*Inward-Looking*          *Outward-Looking*

## Calibrated Image Acquisition



*Calibrated Turntable*

**Selected Dinosaur Images**

**Selected Flower Images**

## Voxel Coloring Results (Video)



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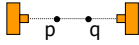**72 K voxels colored**
**7.6 M voxels tested**
**7 min. to compute**
**on a 250MHz SGI**

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

## Limitations of Depth Ordering

A view-independent depth order may not exist



Need more powerful general-case algorithms
- Unconstrained camera positions
- Unconstrained scene geometry/topology

## Voxel Coloring Solutions

1. C=2 (silhouettes)
   - Volume intersection [Baumgart 1974]

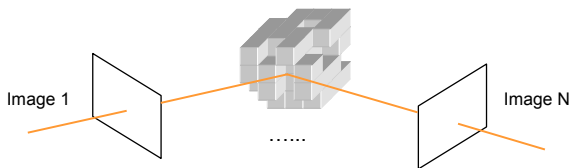2. C unconstrained, viewpoint constraints
   - Voxel coloring algorithm [Seitz & Dyer 97]

3. General Case
   - Space carving [Kutulakos & Seitz 98]
     > For more info: http://www.cs.washington.edu/homes/seitz/papers/kutu-ijcv00.pdf

## Space Carving Algorithm



Image 1    ……    Image N

Space Carving Algorithm
- Initialize to a volume V containing the true scene
- Choose a voxel on the current surface
- Project to visible input images
- Carve if not photo-consistent
- Repeat until convergence

## Which shape do you get?
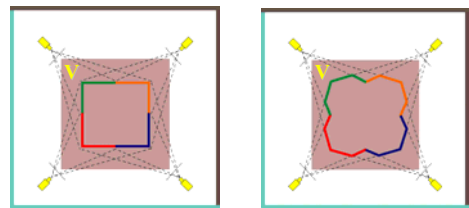


**True Scene**          **Photo Hull**

The Photo Hull *is the UNION of all photo-consistent scenes in V*
- It is a photo-consistent scene reconstruction
- Tightest possible bound on the true scene

## Space Carving Algorithm
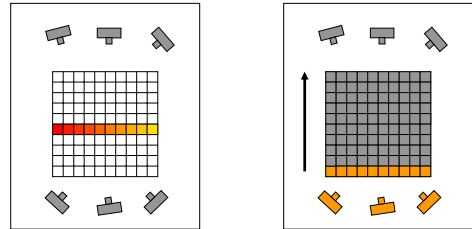
The Basic Algorithm is Unwieldy
- Complex update procedure

Alternative: Multi-Pass Plane Sweep
- Efficient, can use texture-mapping hardware
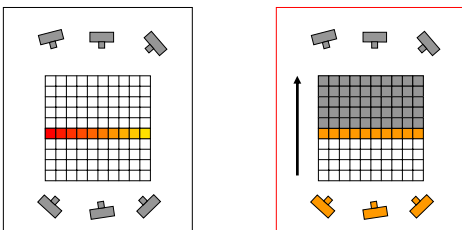- Converges quickly in practice
- Easy to implement

## Multi-Pass Plane Sweep

- Sweep plane in each of 6 principle directions
- Consider cameras on only one side of plane
- Repeat until convergence

## Multi-Pass Plane Sweep

- Sweep plane in each of 6 principle directions
- Consider cameras on only one side of plane
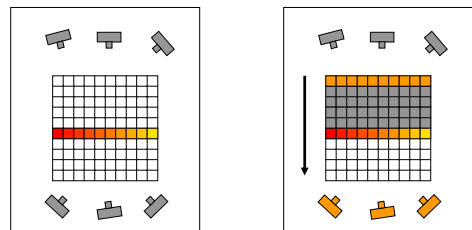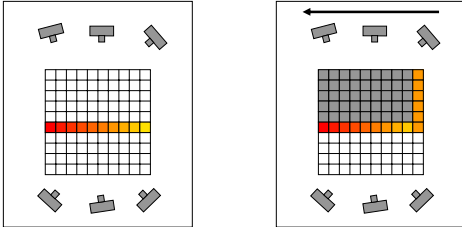- Repeat until convergence

## Multi-Pass Plane Sweep

- Sweep plane in each of 6 principle directions
- Consider cameras on only one side of plane
- Repeat until convergence

## Multi-Pass Plane Sweep

- Sweep plane in each of 6 principle directions
- Consider cameras on only one side of plane
- Repeat until convergence



## Multi-Pass Plane Sweep

- Sweep plane in each of 6 principle directions
- Consider cameras on only one side of plane
- Repeat until convergence



## Multi-Pass Plane Sweep

- Sweep plane in each of 6 principle directions
- Consider cameras on only one side of plane
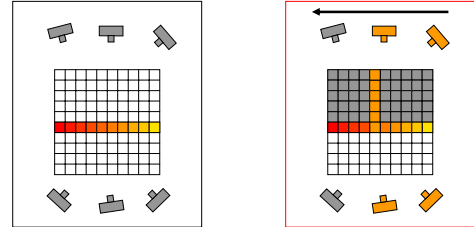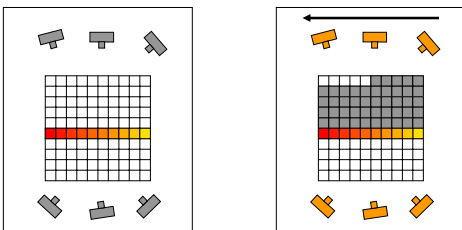- Repeat until convergence



## Space Carving Results:  African Violet



**Input Image (1 of 45)**     **Reconstruction**

**Reconstruction**     **Reconstruction**

## Space Carving Results:  Hand



**Input Image
(1 of 100)**

**Views of Reconstruction**

## Other Approaches

Level-Set Methods *[Faugeras & Keriven 1998]*
- Evolve implicit function by solving PDE's

Probabilistic Voxel Reconstruction *[DeBonet & Viola 1999],
[Broadhurst et al. 2001]*
- Solve for voxel uncertainty (also transparency)

Transparency and Matting *[Szeliski & Golland 1998]*
- Compute voxels with alpha-channel

Max Flow/Min Cut *[Roy & Cox 1998]*
- Graph theoretic formulation

Mesh-Based Stereo *[Fua & Leclerc 1995], [Zhang & Seitz 2001]*
- Mesh-based but similar consistency formulation

Virtualized Reality *[Narayan, Rander, Kanade 1998]*
- Perform stereo 3 images at a time, merge results