Announcements

- Final project proposal: **today at noon**
- Final project presentations: **next Wednesday in class**
  - > **~5 min ppt talk**
- Project 2 artifact winners...

Readings

- Seitz et al., A Comparison and Evaluation of Multi-View Stereo Reconstruction Algorithms, CVPR 2006, pp. 519-526
  - > http://vision.middlebury.edu/mview/seitz_mview_cvpr06.pdf

Multi-view Stereo

Multi-view Stereo

Input: calibrated images from several viewpoints

Output: 3D object model

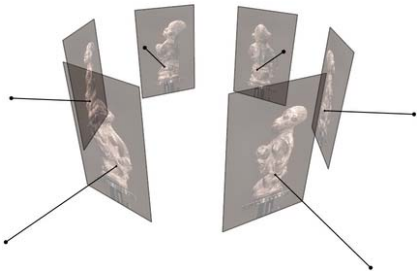

Figures by Carlos Hernandez

History

number of papers in CVPR, ECCV, and ICCV, by year

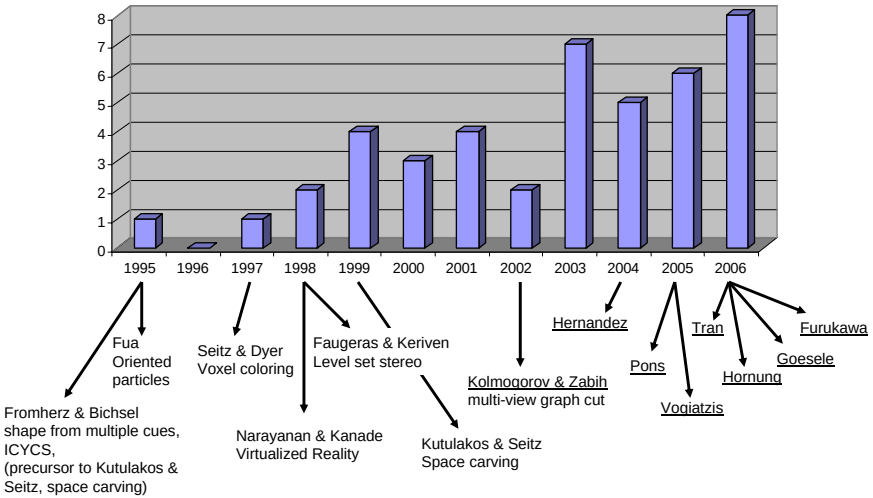

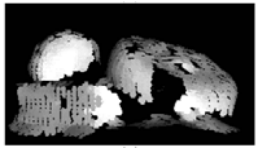

Fua
1995

Seitz, Dyer
1997

Narayanan, Rander, Kanade
1998

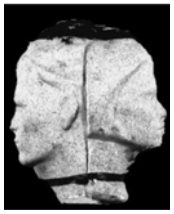

Faugeras, Keriven
1998

Hernandez, Schmitt
2004

Pons, Keriven, Faugeras
2005

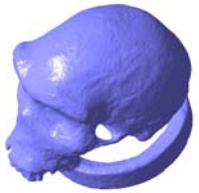

Furukawa, Ponce
2006

Goesele et al.
2007

Stereo: basic idea

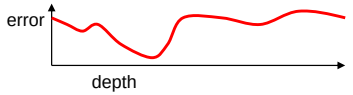

Choosing the stereo baseline

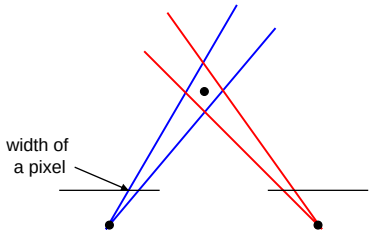

Large Baseline

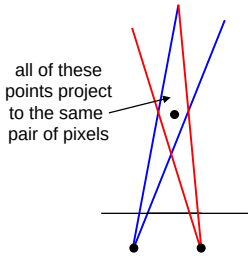

Small Baseline

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.

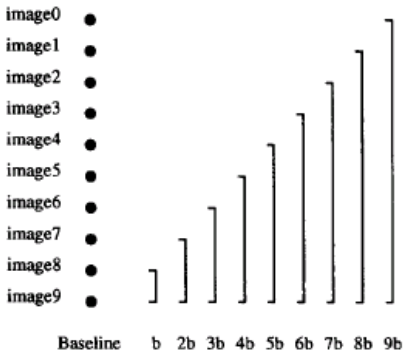

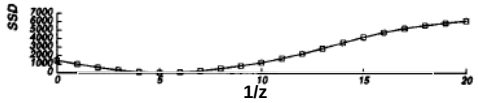

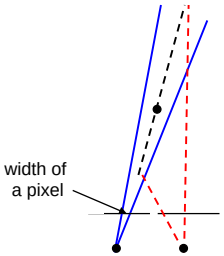

pixel matching score

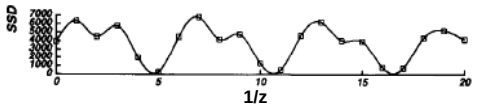

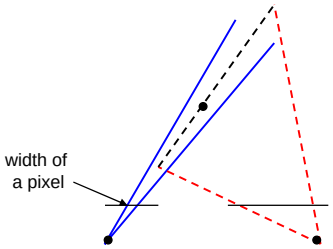

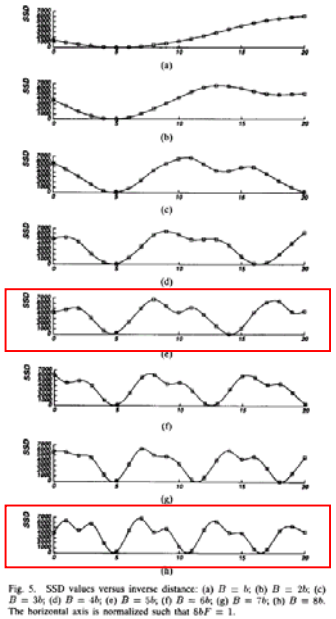

Fig. 5. SSD values versus inverse distance: (a) $B = b$; (b) $B = 2b$; (c) $B = 3b$; (d) $B = 4b$; (e) $B = 5b$; (f) $B = 6b$; (g) $B = 7b$; (h) $B = 8b$. The horizontal axis is normalized such that $bbF = 1$.

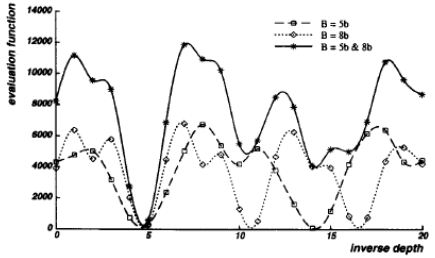

Fig. 6. Combining two stereo pairs with different baselines.

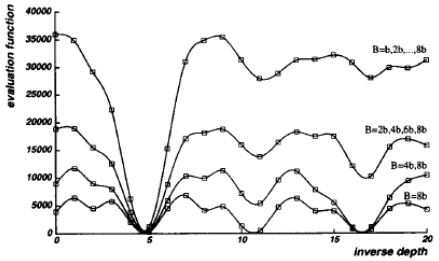

Fig. 7. Combining multiple baseline stereo pairs.

Multibaseline Stereo

Basic Approach

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

Limitations

- Only gives a depth map (not an “object model”)
- Won’t work for widely distributed views:

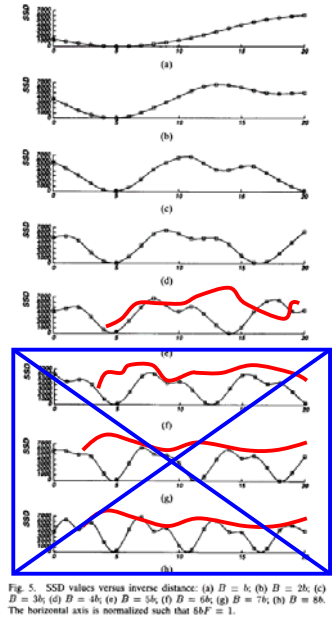

Fig. 5. SSD values versus inverse distance: (a) $B = b$; (b) $B = 2b$; (c) $B = 3b$; (d) $B = 4b$; (e) $B = 5b$; (f) $B = 6b$; (g) $B = 7b$; (h) $B = 8b$. The horizontal axis is normalized such that $bbF = 1$.

Problem: *visibility*

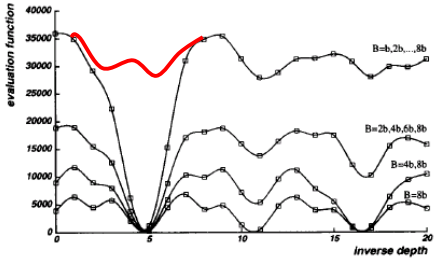

Fig. 7. Combining multiple baseline stereo pairs.

Some Solutions

- Match only nearby photos [Narayanan 98]
- Use NCC instead of SSD,
  Ignore NCC values > threshold
  [Hernandez & Schmitt 03]

Merging Depth Maps

vrip [Curless and Levoy 1996]

- compute weighted average of depth maps

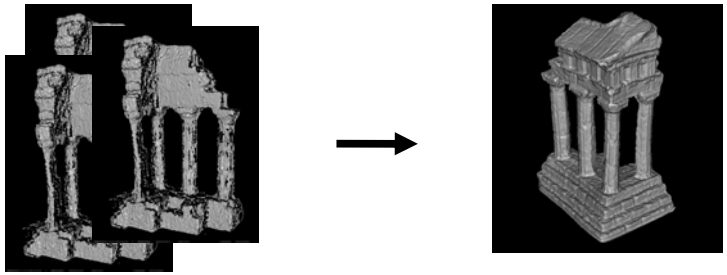

set of depth maps
(one per view)

merged surface
mesh

Merging depth maps

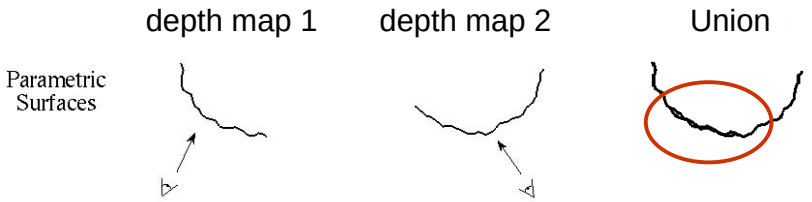

Naïve combination (union) produces artifacts

Better solution: find “average” surface

- Surface that minimizes sum (of squared) distances to the depth maps

Least squares solution

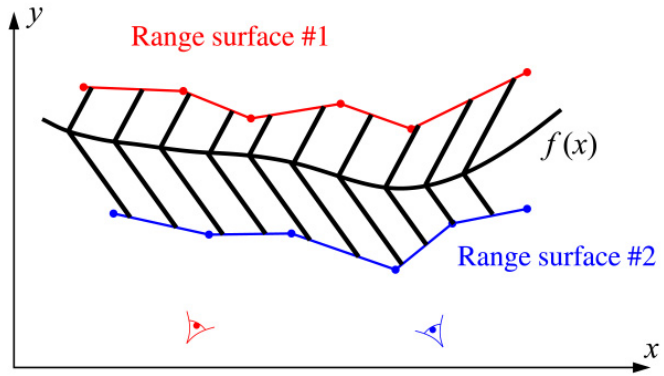

$$E(f) = \sum_{i=1}^{N} \int d_i^2(x, f)\,dx$$

VRIP [Curless & Levoy 1996]

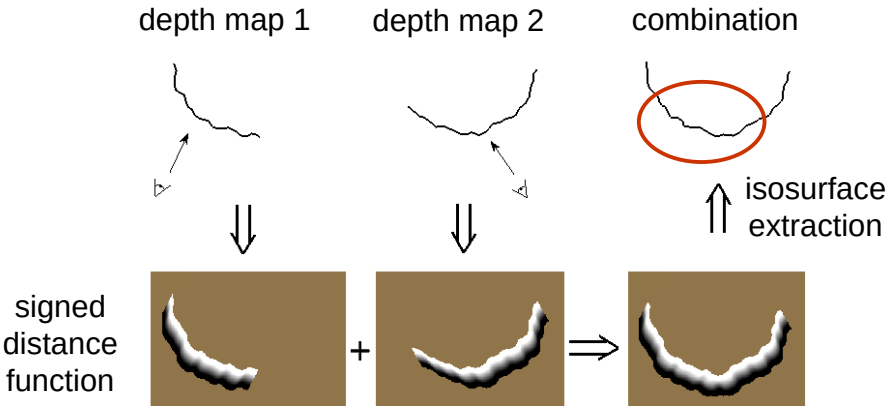

Merging Depth Maps: Temple Model

input image

317 images (hemisphere)

ground truth model

Goesele, Curless, Seitz, 2006

Michael Goesele

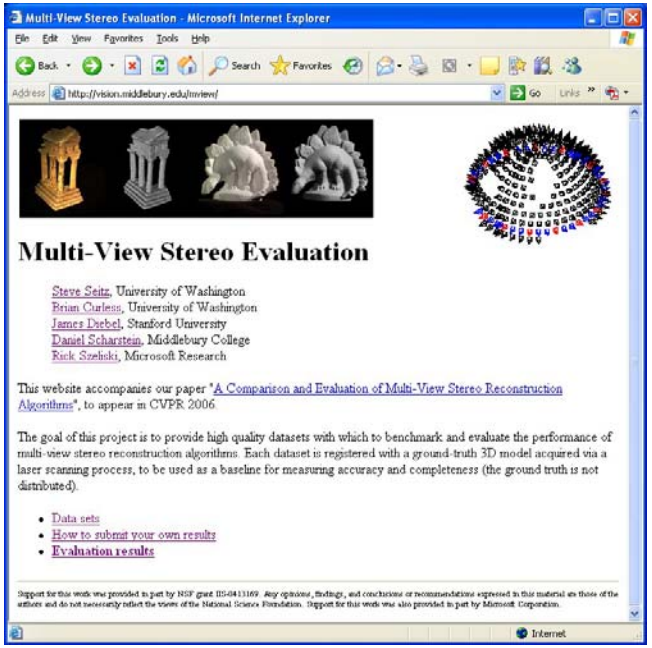

Multi-view stereo from Internet Collections

[Goesele, Snavely, Curless, Hoppe, Seitz, ICCV 2007]

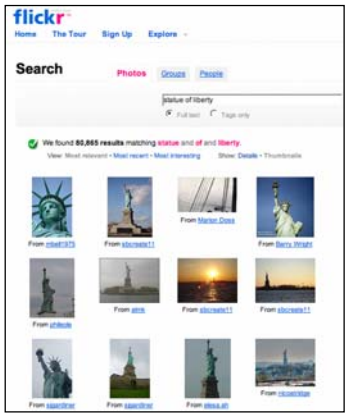

Challenges

- appearance variation

- resolution

- massive collections

82,754 results for photos matching **notre** and **dame** and **paris**.

Law of Large Image Collections

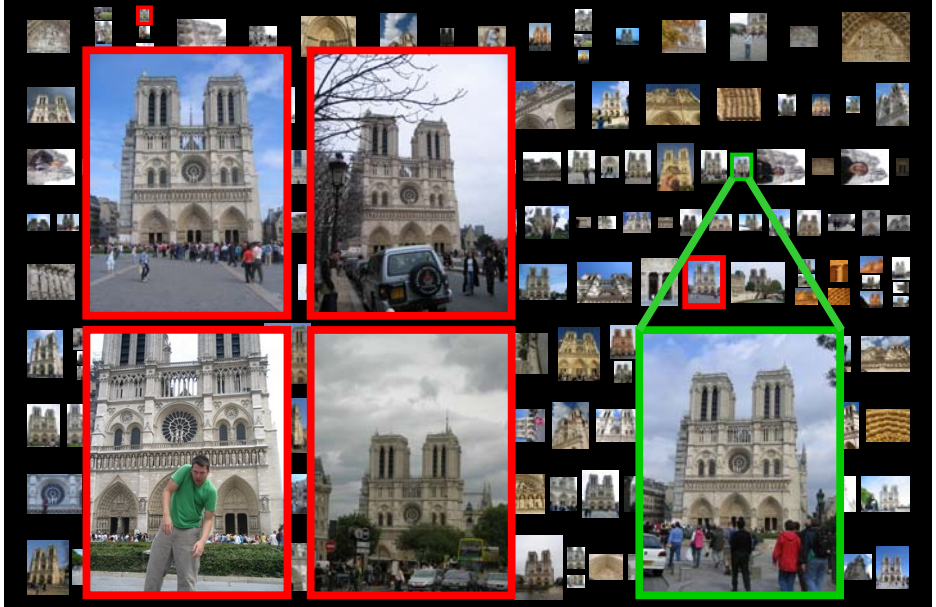

206 *Flickr* images taken by 92 photographers

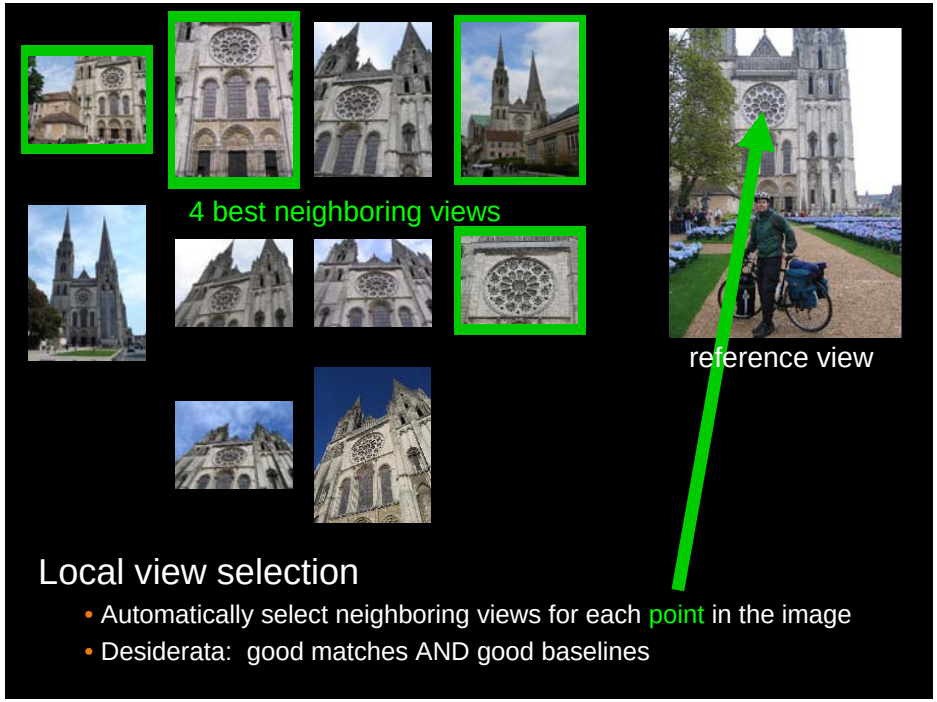

Local view selection

- Automatically select neighboring views for each point in the image
- Desiderata: good matches AND good baselines

Local view selection

- Automatically select neighboring views for each point in the image
- Desiderata: good matches AND good baselines

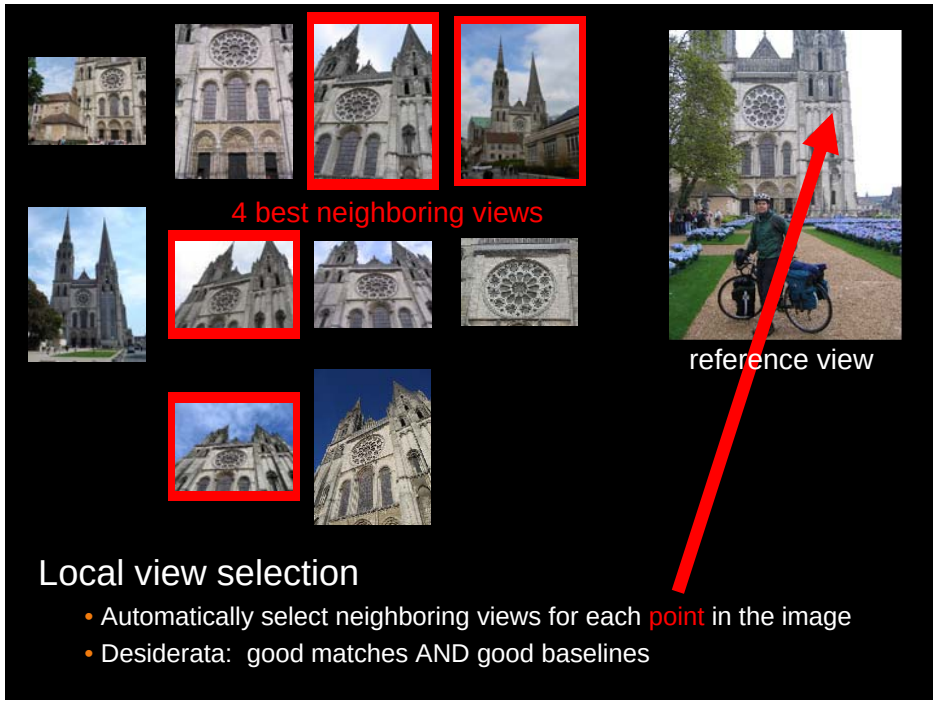

Local view selection

- Automatically select neighboring views for each point in the image
- Desiderata: good matches AND good baselines

Results




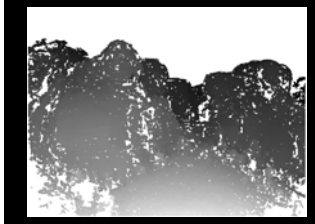
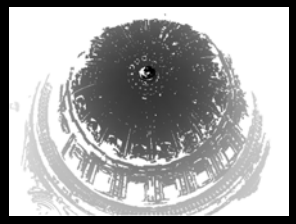
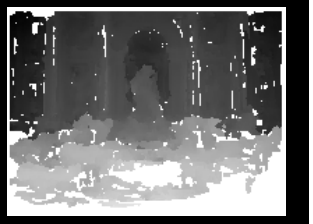

Mt. Rushmore
160 images
60 photographers

St. Peter
151 images
50 photographers

Trevi Fountain
106 images
51 photographers

Notre Dame de Paris

653 images
313 photographers

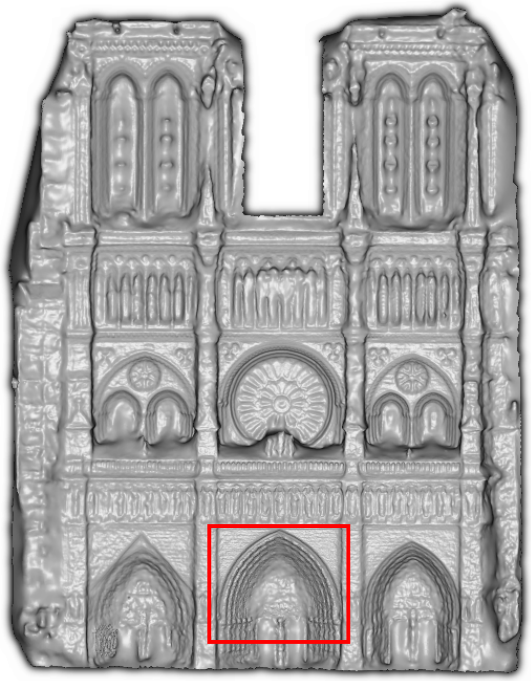

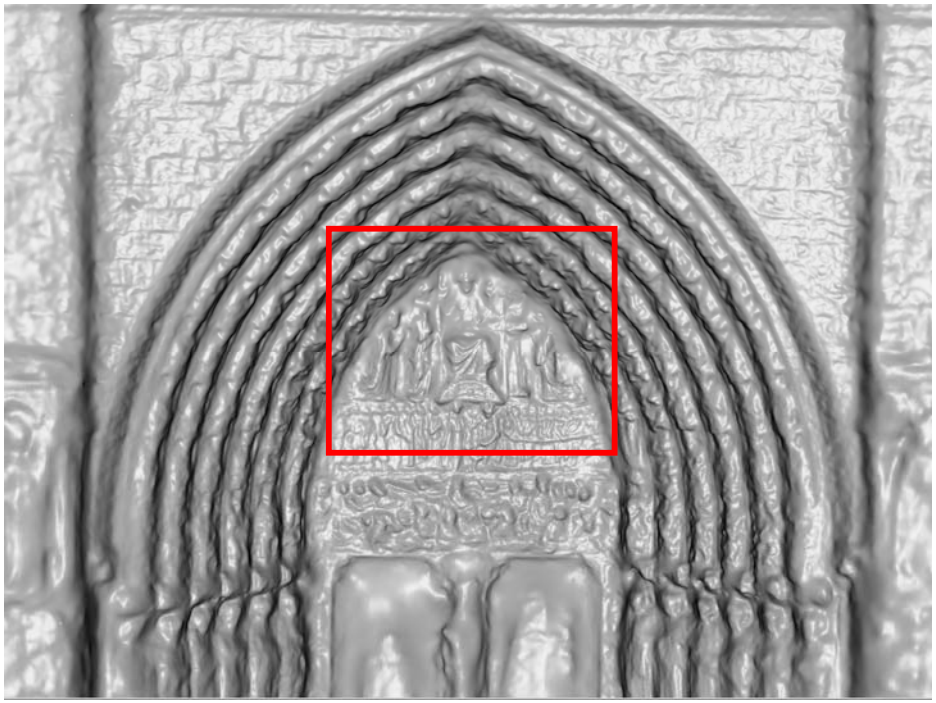

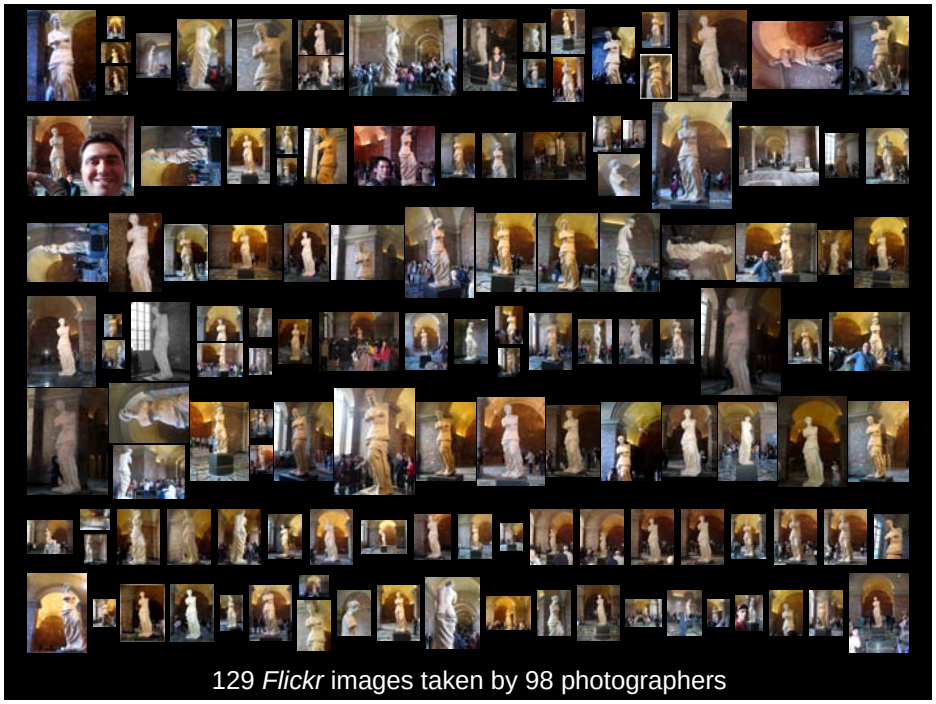

129 *Flickr* images taken by 98 photographers

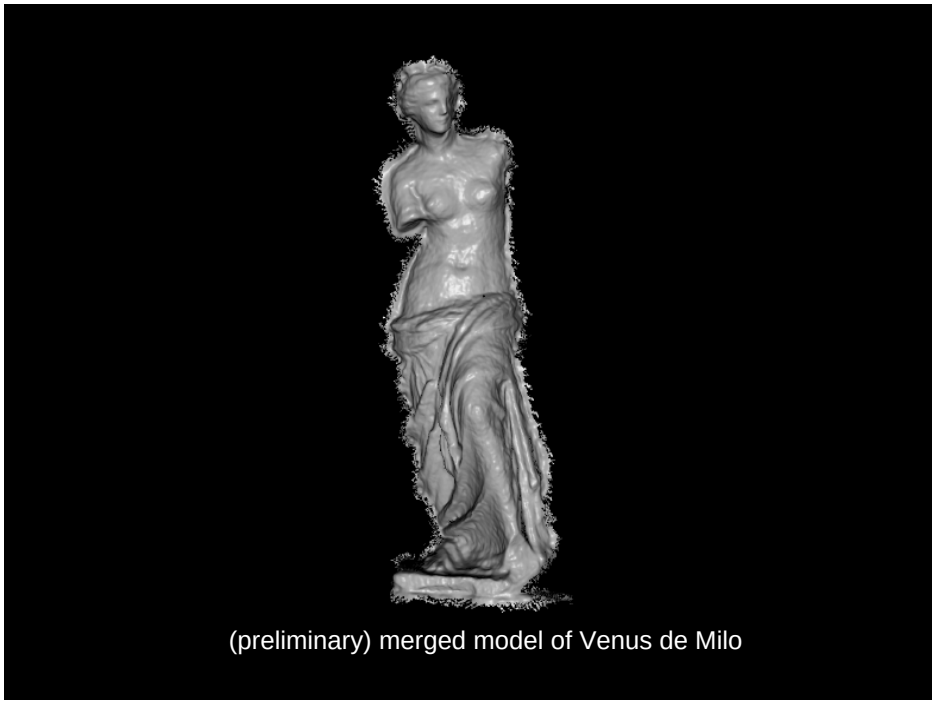

(preliminary) merged model of Venus de Milo

56 *Flickr* images taken by 8 photographers

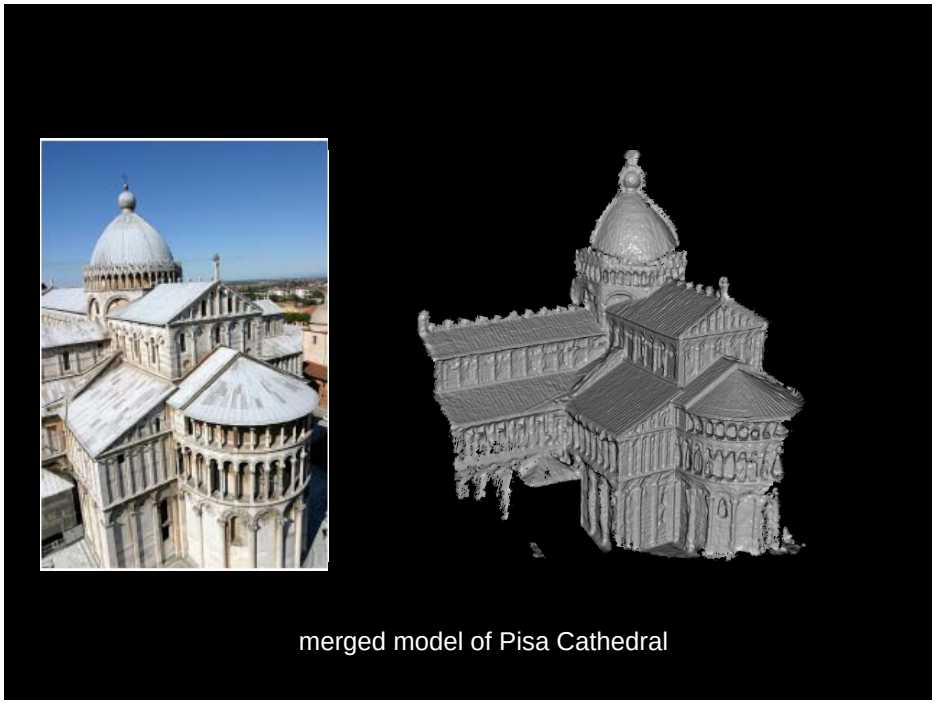

merged model of Pisa Cathedral

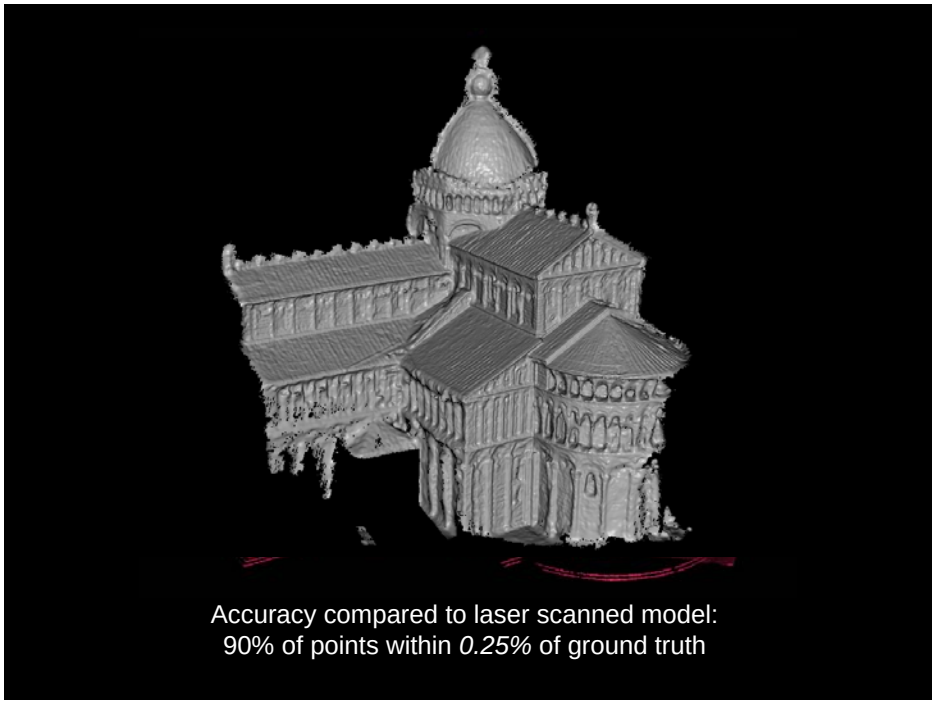

Accuracy compared to laser scanned model:
90% of points within *0.25%* of ground truth

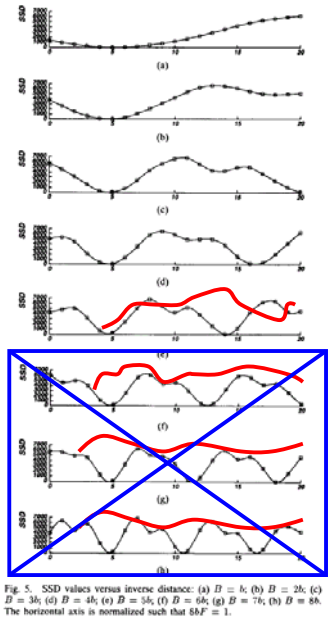

Fig. 5. SSD values versus inverse distance: (a) $B = b$; (b) $B = 2b$; (c) $B = 3b$; (d) $B = 4b$; (e) $B = 5b$; (f) $B = 6b$; (g) $B = 7b$; (h) $B = 8b$. The horizontal axis is normalized such that $bbF = 1$.

Problem: *visibility*

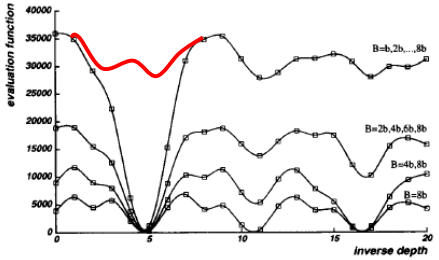

Fig. 7. Combining multiple baseline stereo pairs.

Some Solutions

- Match only nearby photos [Narayanan 98]
- Use NCC instead of SSD,
  Ignore NCC values > threshold
  [Hernandez & Schmitt 03]

The visibility problem

Which points are visible in which images?

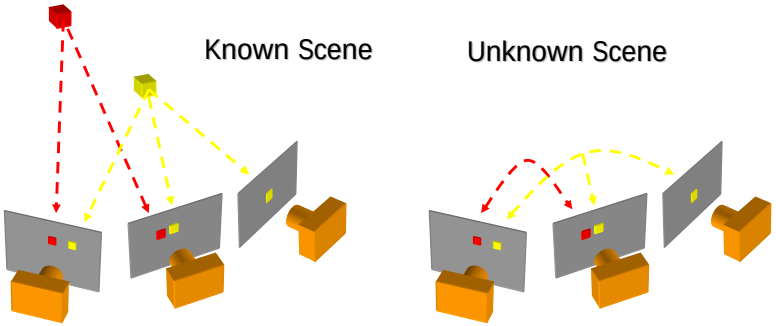

Forward Visibility *Inverse Visibility*

Volumetric stereo

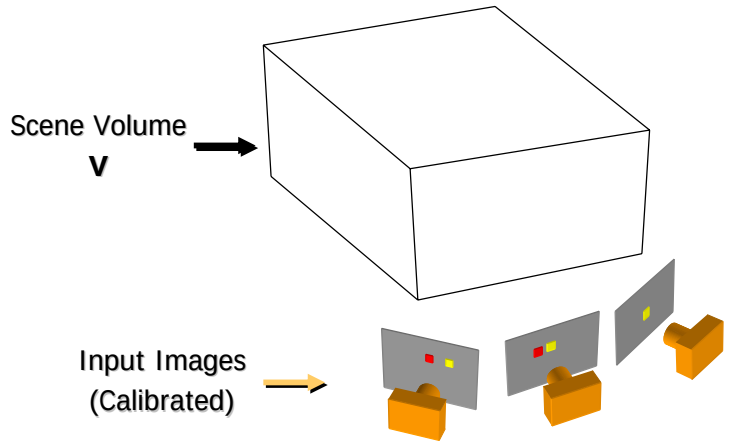

Goal: **Determine occupancy, "color" of points in V**

Discrete formulation: Voxel Coloring

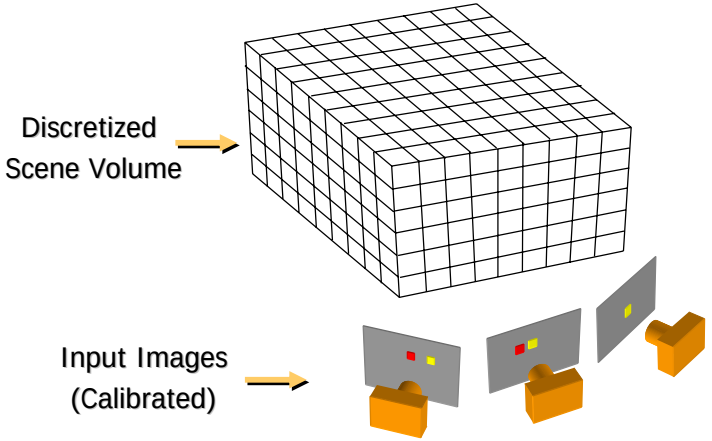

Goal: **Assign RGBA values to voxels in V**
photo-consistent **with images**

Complexity and computability

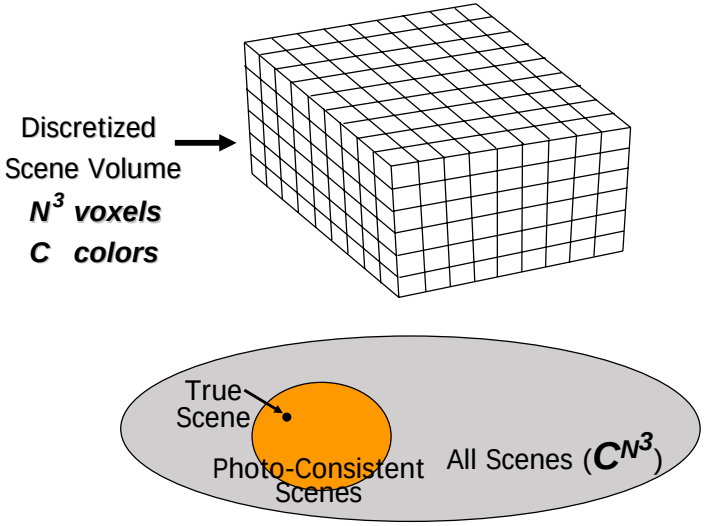

Issues

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)

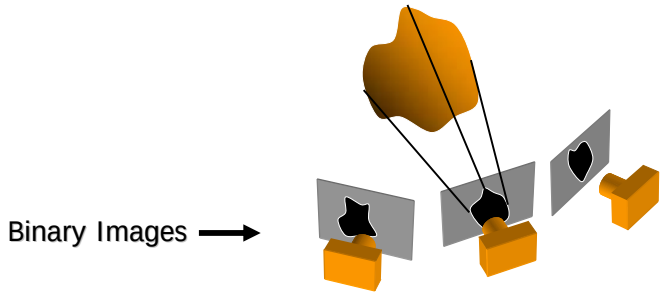

Approach:

- *Backproject* each silhouette
- Intersect backprojected volumes

Volume intersection

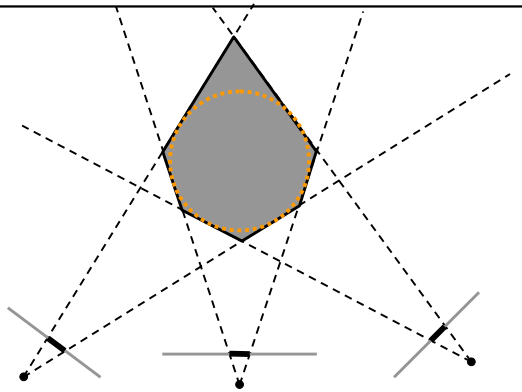

Reconstruction Contains the True Scene

- 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

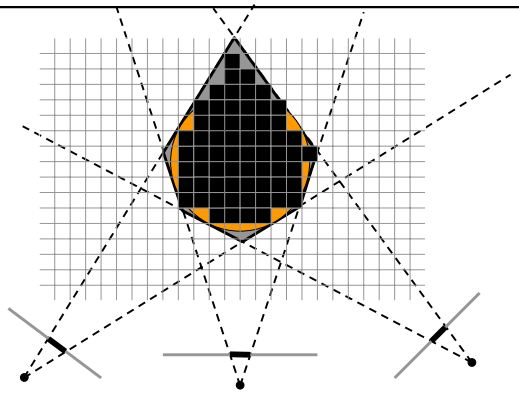

Color voxel black if on silhouette in every image

- 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/iicv99.pdf
3. General Case
   - Space carving [Kutulakos & Seitz 98]

Voxel Coloring Approach

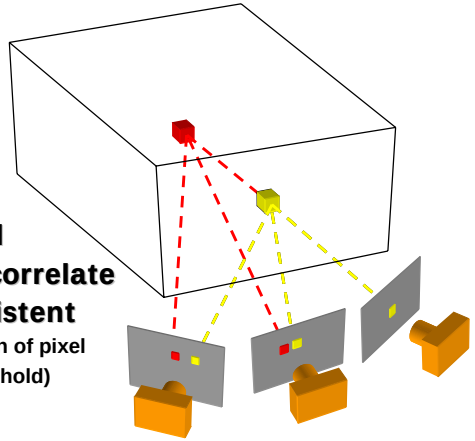

1. Choose voxel
2. Project and correlate
3. Color if consistent
(standard deviation of pixel colors below threshold)

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

Depth Ordering: visit occluders first!

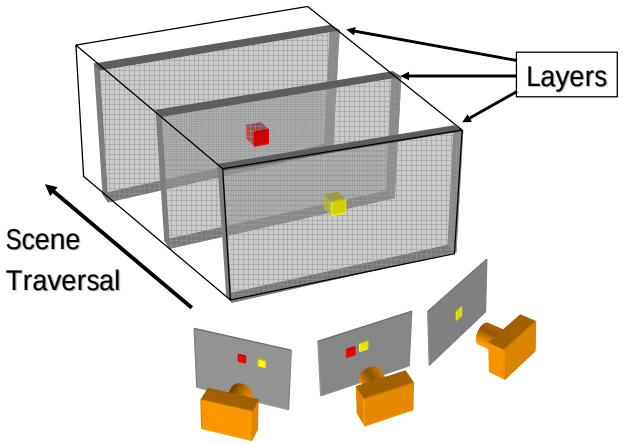

Condition: **depth order is the *same for all input views***

Calibrated Image Acquisition

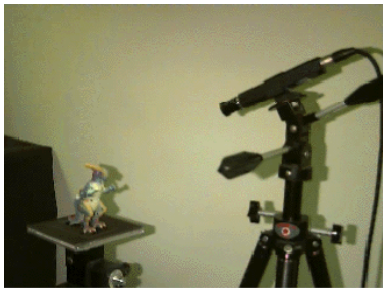

Calibrated Turntable

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

Improvements

Unconstrained camera viewpoints

- Space carving [Kutulakos & Seitz 98]

Evolving a surface

- Level sets [Faugeras & Keriven 98]
- More recent work by Pons et al.

Global optimization

- Graph cut approaches
  - > [Kolmogoriv & Zabih, ECCV 2002]
  - > [Vogiatzis et al., PAMI 2007]

Modeling shiny (and other reflective) surfaces

- e.g., Zickler et al., *Helmholtz Stereopsis*

See today's reading for an overview of the state of the art