Announcements

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

Apple Maps

Multi-view Stereo

Point Grey's Bumblebee XB3

Point Grey's ProFusion 25

CMU's 3D Room

Multi-view Stereo

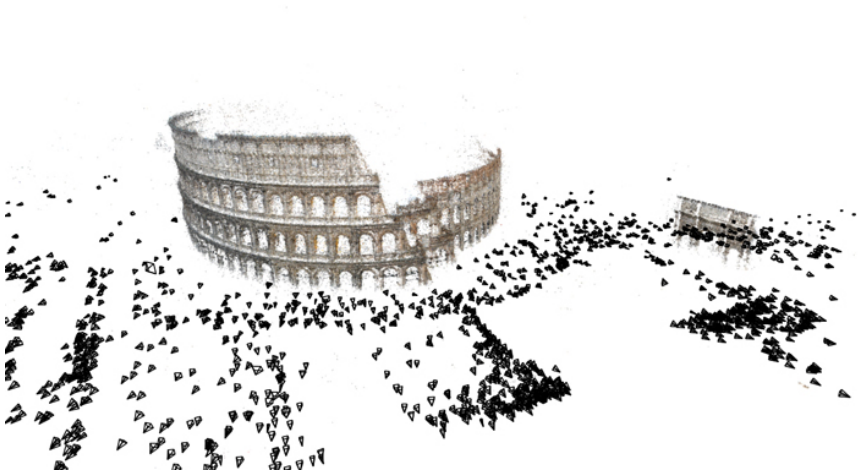

Multi-view Stereo

Input: calibrated images from several viewpoints

Output: 3D object model

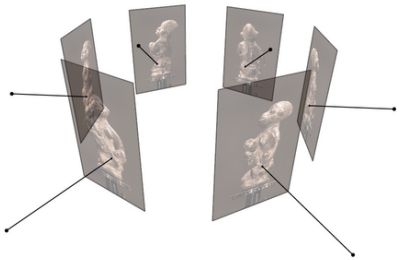

Figure by Carlos Hernandez

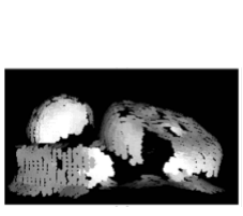

Fua
1995

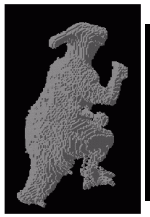

Seitz, Dyer
1997

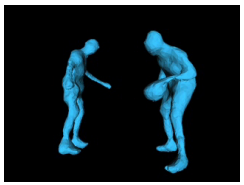

Narayanan, Rander, Kanade
1998

Faugeras, Keriven
1998

Hernandez, Schmitt
2004

Pons, Keriven, Faugeras
2005

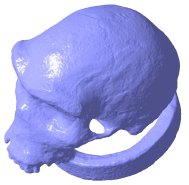

Furukawa, Ponce
2006

Goesele et al.
2007

Stereo: basic idea

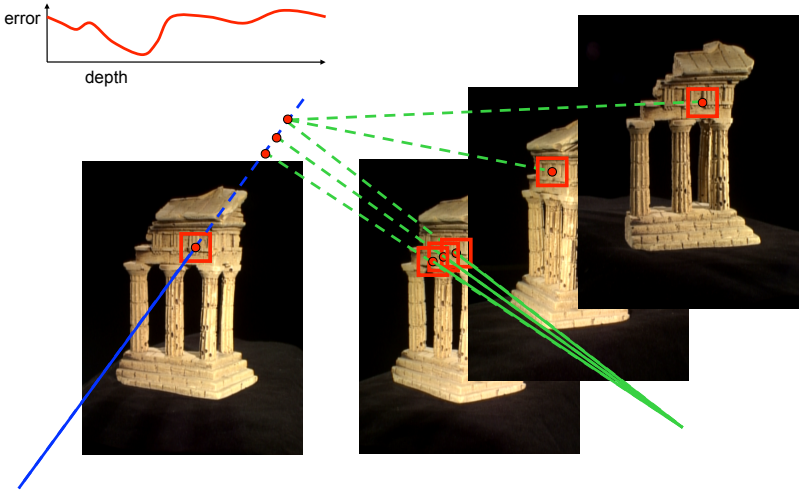

Choosing the stereo baseline

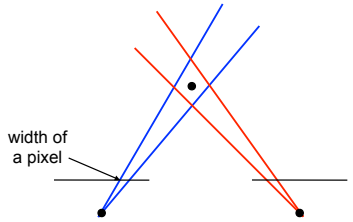

Large Baseline

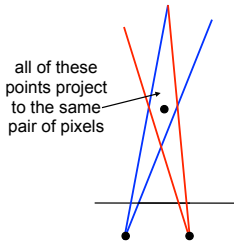

Small Baseline

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.

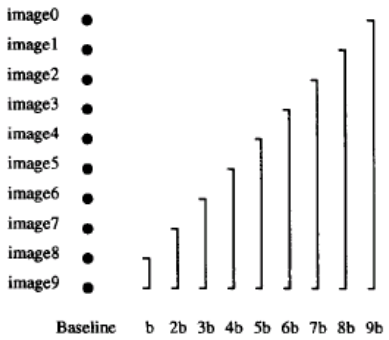

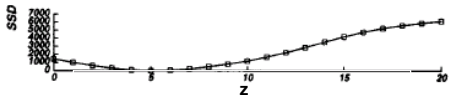

pixel matching score

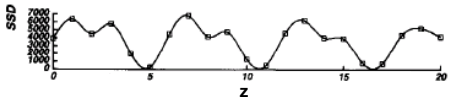

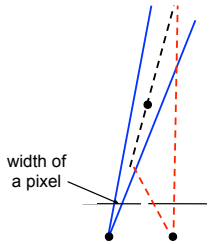

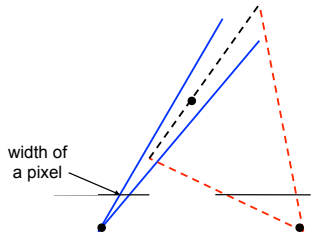

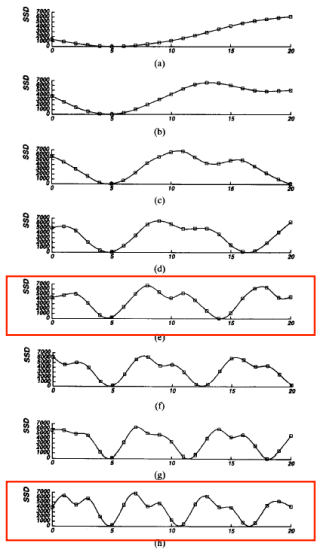

Fig. 5. SSD values versus inverse distance: (a) $B = b$; (b) $B = 2b$; (c) $B = 3b$; (d) $B = 4b$; (e) $B = 5b$; (f) $B = 6b$; (g) $B = 7b$; (h) $B = 8b$. The horizontal axis is normalized such that $8bF = 1$.

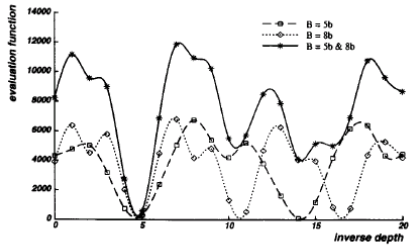

Fig. 6. Combining two stereo pairs with different baselines.

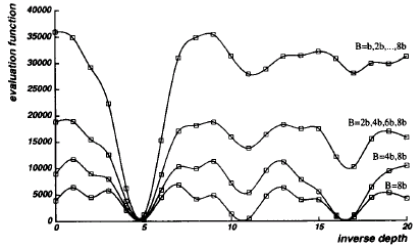

Fig. 7. Combining multiple baseline stereo pairs.

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:

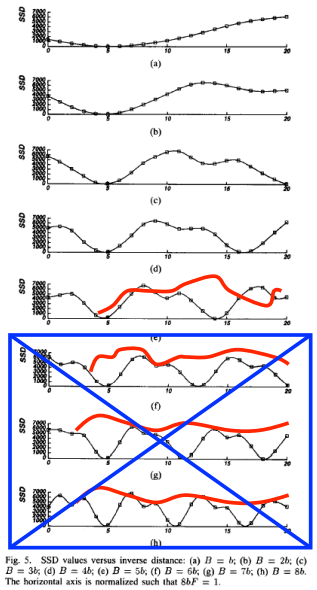

Fig. 5. SSD values versus inverse distance: (a) $B = b$; (b) $B = 2b$; (c) $B = 3b$; (d) $B = 4b$; (e) $B = 5b$; (f) $B = 6b$; (g) $B = 7b$; (h) $B = 8b$. The horizontal axis is normalized such that $8bF = 1$.

Problem: *visibility*

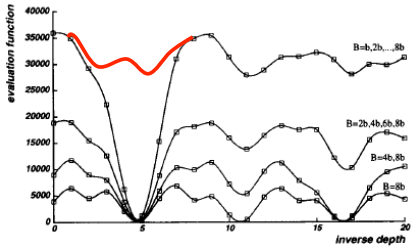

Fig. 7. Combining multiple baseline stereo pairs.

Some Solutions

- Match only nearby photos
- Ignore SSD values > threshold
- Alternative to SSD

Merging Depth Maps

vrip [Curless and Levoy 1996]

- compute weighted average of depth maps

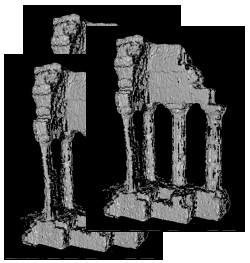

set of depth maps (one per view)

merged surface mesh

Merging depth maps

depth map 1 | depth map 2 | Union

Parametric Surfaces

Naïve combination (union) produces artifacts

Better solution: find “average” surface

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

Least squares solution

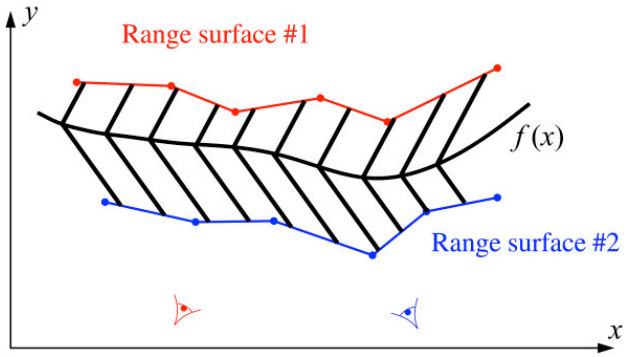

$$E(f) = \sum_{i=1}^{N} \int d_i^2(x, f)\,dx$$

VRIP [Curless & Levoy 1996]

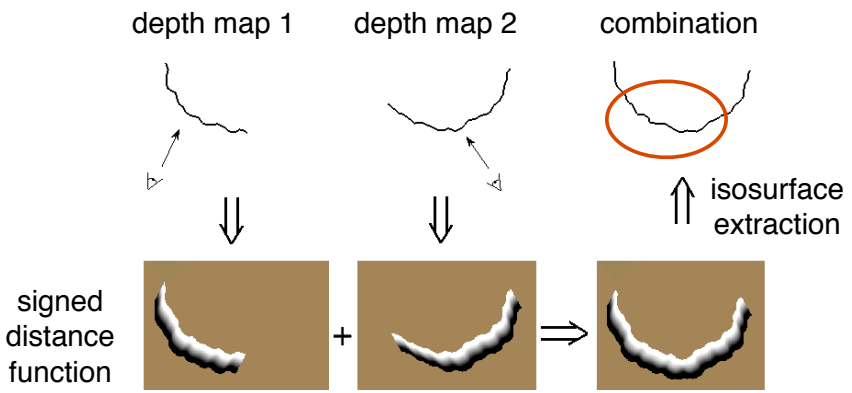

Merging Depth Maps: Temple Model

input image

317 images (hemisphere)

ground truth model

Goesele, Curless, Seitz, 2006

Michael Goesele

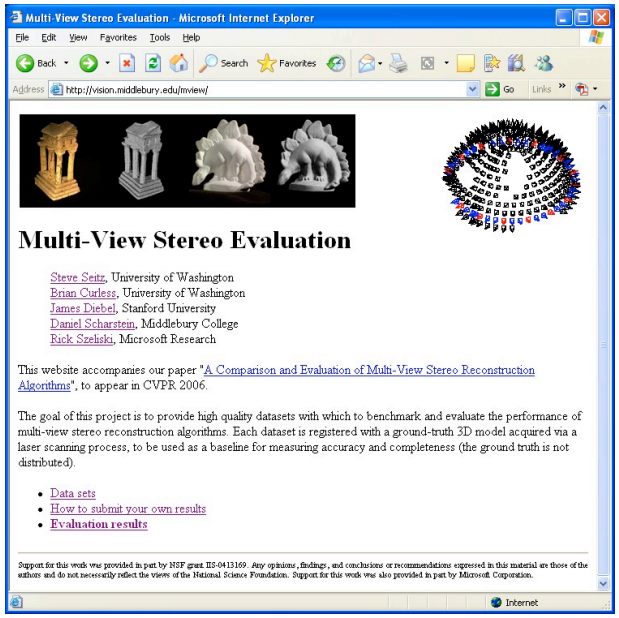

From Sparse to Dense

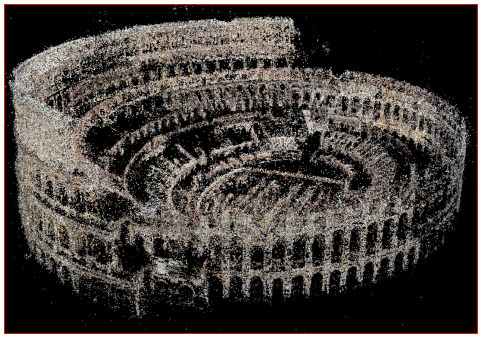

Sparse output from the SfM system

From Sparse to Dense

Furukawa, Curless, Seitz, Szeliski, CVPR 2010

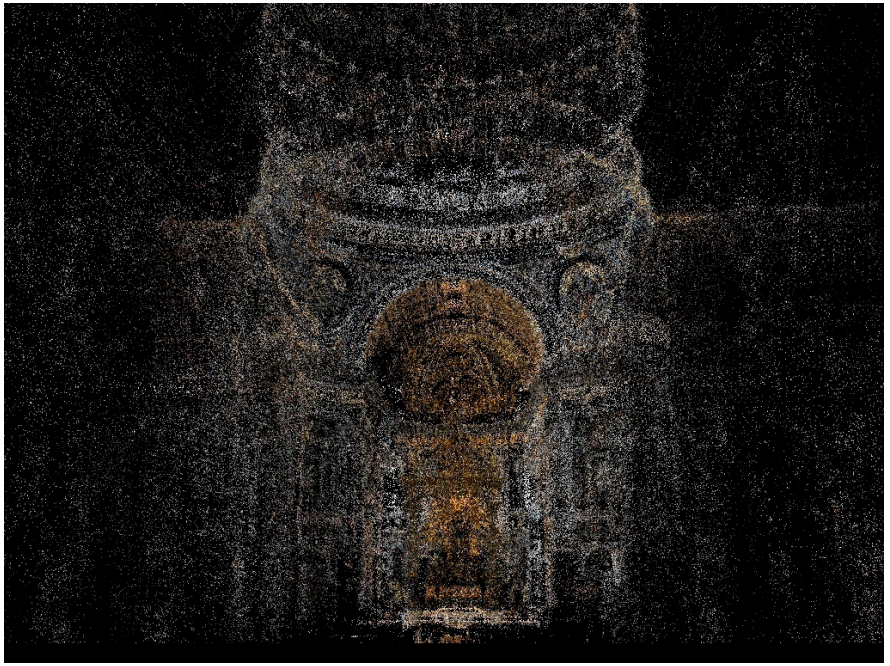

Venice Sparse Model