Computational Photography

Last Time?

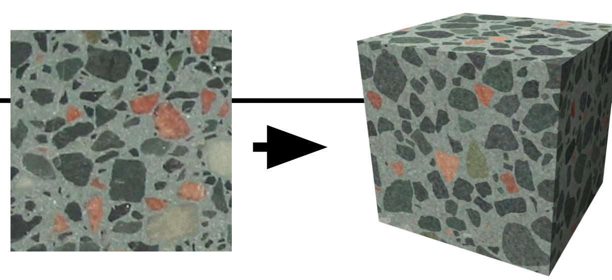

- Texture Synthesis
- Markov Model
- Image Completion
- Volumetric Texture Synthesis

"I spent an interesting evening recently with a grain of salt."

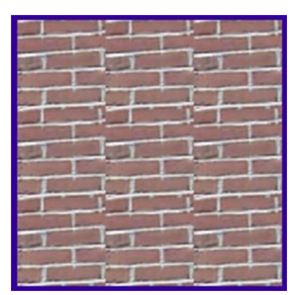

(a) 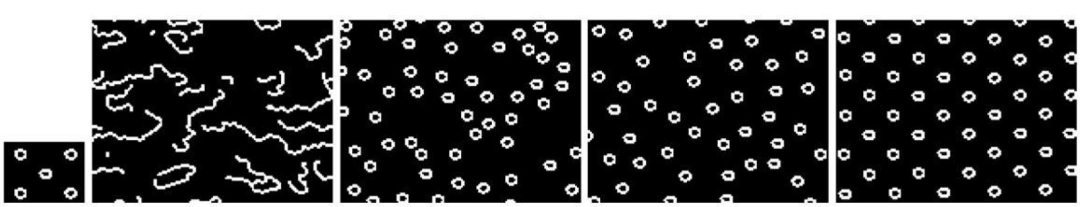

Today

- Papers for Today
- Announcements: Quiz & Final Projects
- Photography Lesson: Tilt/Shift Lenses
- Structure From Motion
- Multi-viewpoint Rendering
- Matting & Compositing
- Helmholtz Reciprocity
- Light Fields

"Flocks, Herds, and Schools: A Distributed Behavioral Model", Craig W. Reynolds, SIGGRAPH 1987

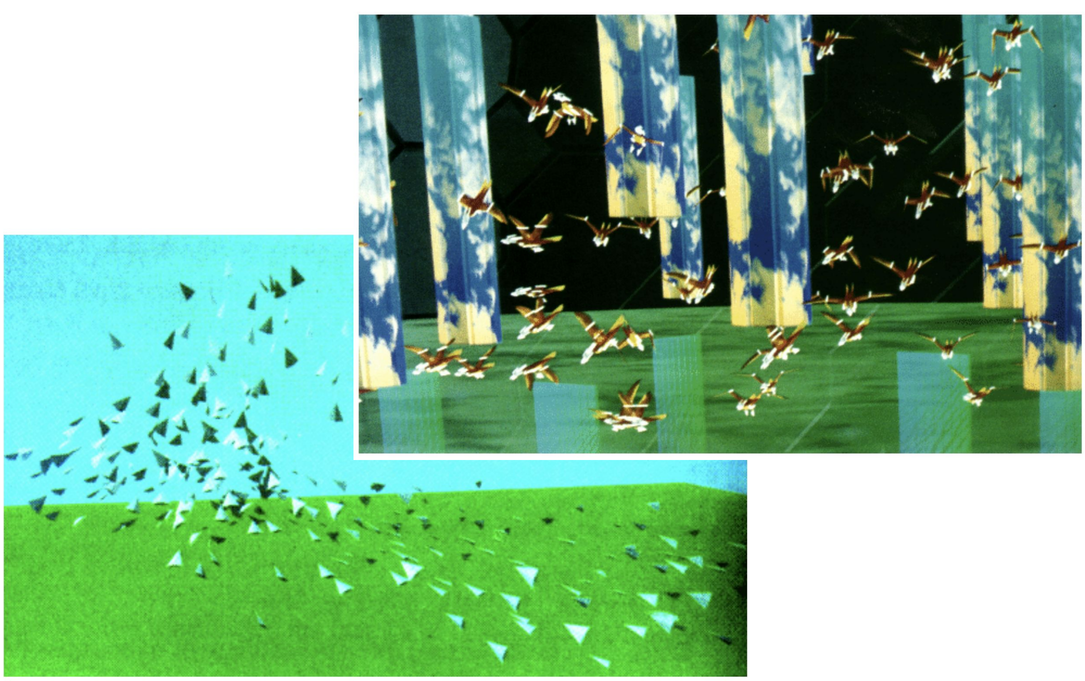

“On Demand Solid Texture Synthesis Using Deep 3D Networks”, Gutierrez, Rabin, Galerne, and Hurtut, 2018, preprint.

“Interactive Digital Photomontage”, Agarwala, Dontcheva, Agrawala, Drucker, Colburn, Curless, Salesin, & Cohen SIGGRAPH 2004

Today

- Papers for Today
- Announcements: Quiz & Final Projects
- Photography Lesson: Tilt/Shift Lenses
- Structure From Motion
- Multi-viewpoint Rendering
- Matting & Compositing
- Helmholtz Reciprocity
- Light Fields

Remaining Schedule...

- **Tuesday April 20th**:
  attendance optional office hours / work day
- **Friday, April 23rd**: Quiz 2, during class time
  - practice problems are on the calendar
- **Tuesday Apr 27th** & **Friday Apr 30th**:
  - Final Project Presentations
  - *mandatory attendance for everyone*
    ask questions & "peer grading" / feedback
- **Monday May 3rd** @ midnight:
  Final Project Reports (& source code) due
  please, please, please no late days/extensions

Final Presentation Schedule

10 min (individual)
+2 min questions/setup
20 min (team of 2)
+2 min questions/setup

Tue Apr 27th

2:30 Dan
2:42 Alex
2:54 Marcus & Michael
3:16 Jay F.
3:28 Jay I. & Lyon
3:50 Olivia
4:02 Reed
4:14 done!

Let me know ASAP if this doesn't work for you

Fri Apr 30th

2:30 Mack
2:42 Ben & Matt
3:04 Thomas
3:16 Jon & Sam C.
3:38 Jacy
3:50 Aaron & Austen
4:12 Justin
4:24 Sam O. & Sophie
4:46 Shane
4:58 done!

Final Presentation

- Summarize prior work as necessary
  - You don't need to discuss papers we covered in class
- Be technical:
  - What were the challenges?
  - How did you solve them?
- Live demo / video / lots of images (depends on project)
  - Use plenty of examples (both of success & failure)
- Teams of 2:
  - Both should present & make it clear who did what
- Use your time wisely! Practice! & time yourself!
  - *I will stop you mid-sentence if you run over*

Well-written Research Paper / Report

- Motivation/context/related work
- Accomplishments / contributions of this work
- Clear description of algorithm
  - Sufficiently-detailed to allow work to be reproduced
  - Work is theoretically sound (hacks/arbitrary constants discouraged, but must be documented)
- Results
  - well chosen examples
  - clear tables/illustrations/visualizations
- Conclusions
  - limitations of the method are clearly stated

Today

- Papers for Today
- Announcements: Quiz & Final Projects
- Photography Lesson: Tilt/Shift Lenses
- Structure From Motion
- Multi-viewpoint Rendering
- Matting & Compositing
- Helmholtz Reciprocity
- Light Fields

The Sandpit, O'Hare, 2010

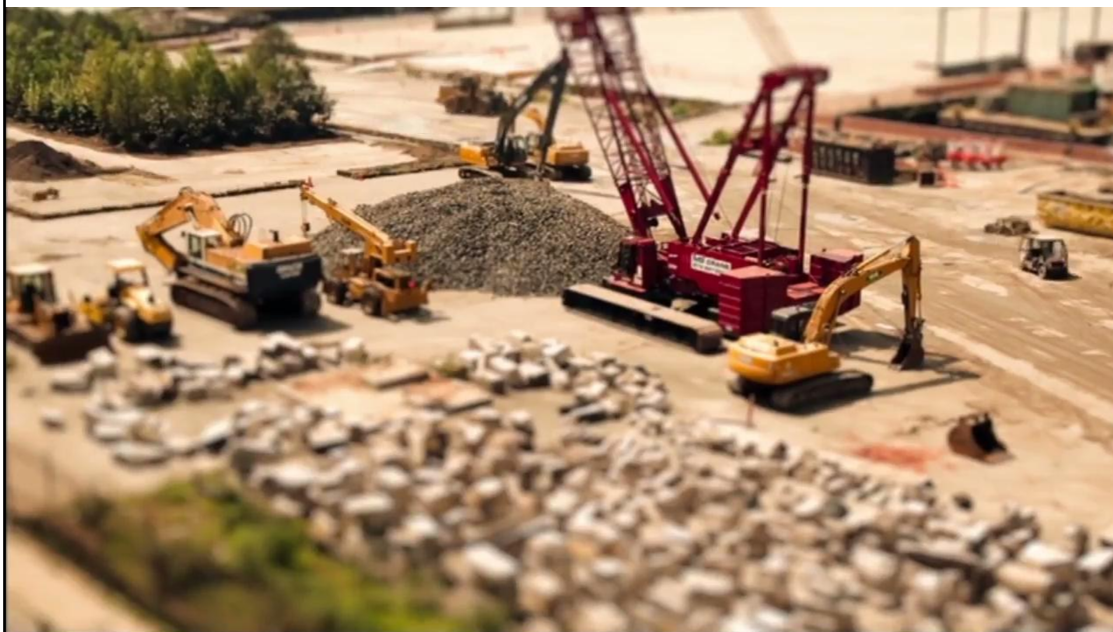

Tilt-Shift Camera Lens

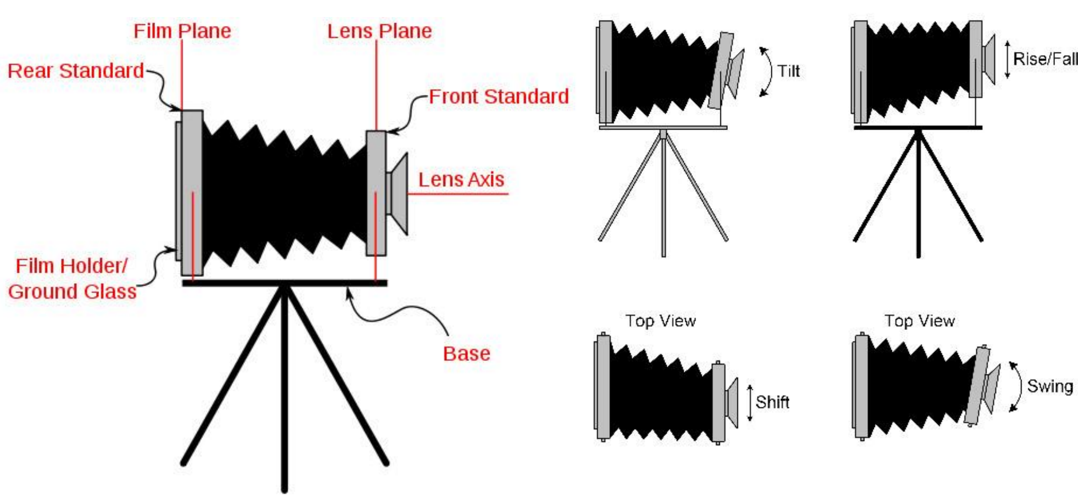

Illustrations by Chris Heald
https://photography.tutsplus.com/tutorials/an-introduction-to-large-format-photography--photo-7987

Tilt-Shift Camera Lens

https://en.wikipedia.org/wiki/Tilt%E2%80%93shift_photography#/media/File:24mm-tilt-lens.jpg

https://en.wikipedia.org/wiki/Tilt%E2%80%93shift_photography#/media/File:Nikon-35mm-left.jpg

Shift/Rise for Perspective-Control

https://www.colesclassroom.com/5-tips-to-take-architectural-photography-next-level/

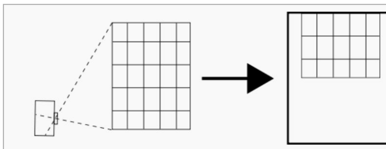

(**a**) Keeping the camera level, with an ordinary lens, captures only the bottom portion of the building.

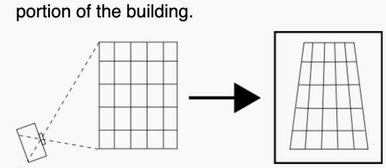

(**b**) Tilting the camera upwards results in perspective distortion.

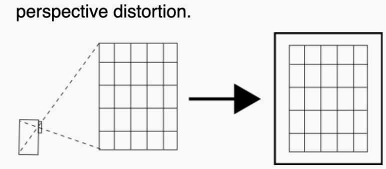

(**c**) Shifting the lens upwards results in a picture of the entire subject without perspective distortion.

https://en.wikipedia.org/wiki/Tilt%E2%80%93shift_photography

Tilt/Swing for Focus Control

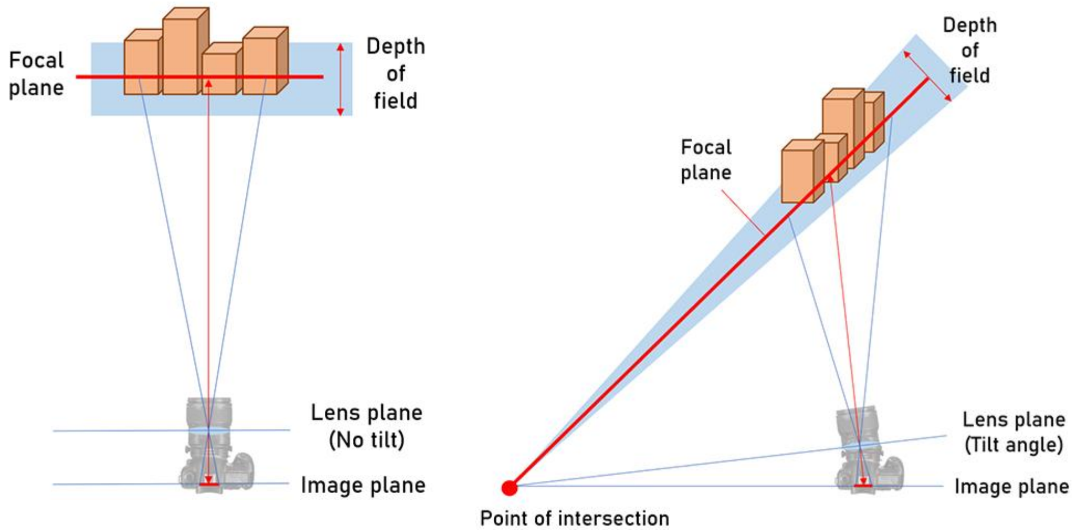

https://snapshot.canon-asia.com/india/article/en/what-you-didnt-know-about-the-tilt-function-on-tilt-shift-lenses

Tilt/Swing for Focus Control

https://upload.wikimedia.org/wikipedia/commons/e/ee/Tilt-lens_photo_of_model_train.jpg

Tilt/Swing for Focus Control

https://luminous-landscape.com/focusing-tilt-shift-lenses/

Focus in the distance

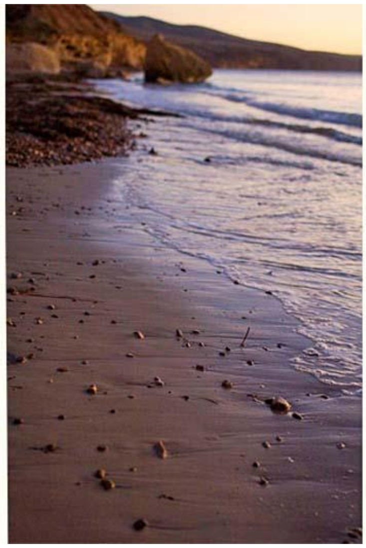

Focus on the foreground

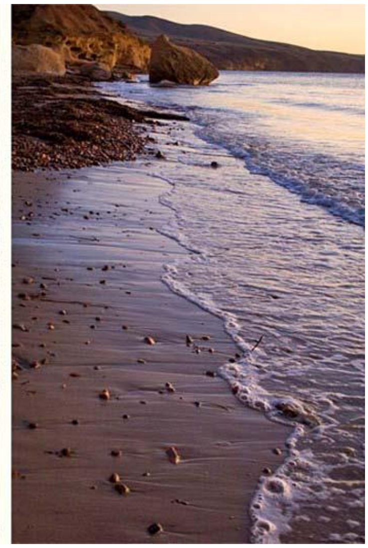

Tilted focal plan,
(most) everything in focus

Tilt/Swing for Selective Focus

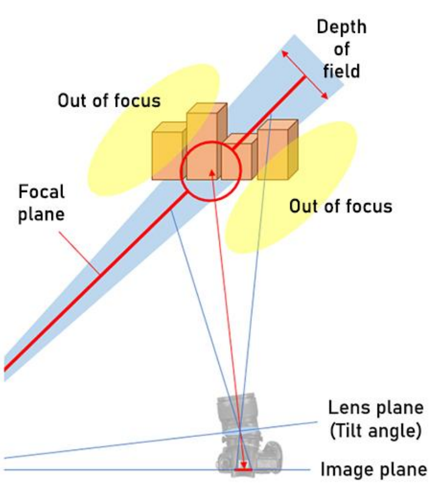

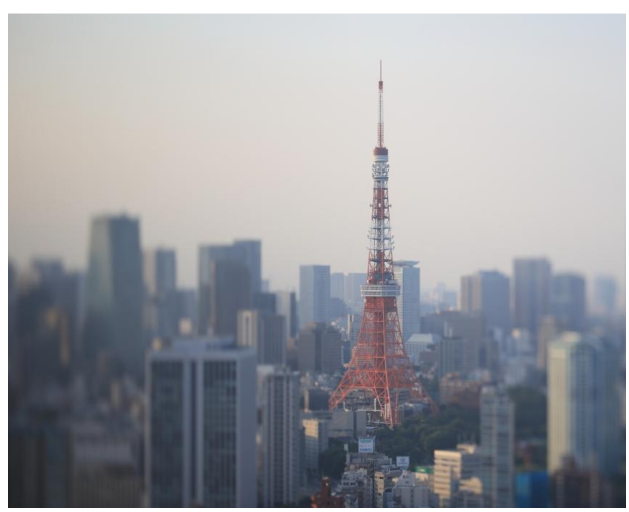

https://snapshot.canon-asia.com/india/article/en/what-you-didnt-know-about-the-tilt-function-on-tilt-shift-lenses

Tilt/Swing for Miniature Effect

Ben Thomas
http://benthomas.co/cityshrinker/wamrft5d4mk0ua9ahjxtu9cu2qtxsp

Tilt/Swing for Miniature Effect

Ben Thomas
http://benthomas.co/cityshrinker/4pgljfasa898t1z6cjnio3pilyaxzi

Tilt/Swing for Miniature Effect

How to achieve the effect?

- High-angle shot (from "above") from a distance
- With a shallow depth of field
  - Foreground and background out of focus

Why does it look like a miniature model?

- Actual miniatures are usually photographed from close up…
- Depth of field is narrower/shallower when the camera is closer to the objects

Today

- Papers for Today
- Announcements: Quiz & Final Projects
- Photography Lesson: Tilt/Shift Lenses
- Structure From Motion
- Multi-viewpoint Rendering
- Matting & Compositing
- Helmholtz Reciprocity
- Light Fields

Structure From Motion

- Input: Sequence of frames (e.g., video) of a moving object (or moving camera)
- Output: Approximate geometry of object & camera pose for each frame
- How?
  - Automatically detect features in each frame
  - Determine correspondences between features
  - Infer camera calibration & object geometry
- Humans do it all the time… but it's a really hard problem!

Photo Tourism

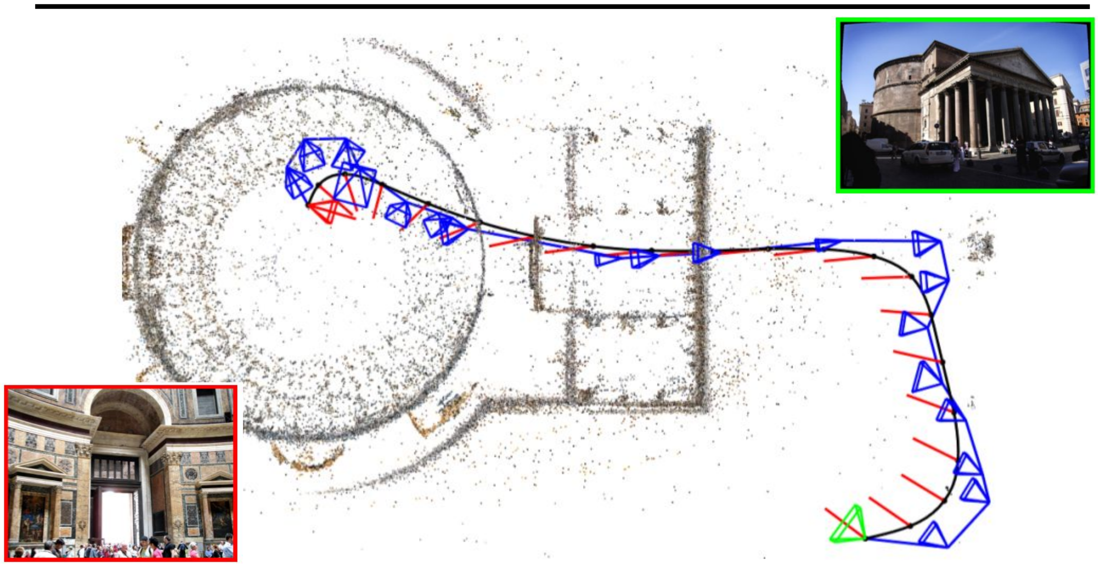

Finding Paths through the World's Photos,
Snavely, Garg, Seitz, & Szeliski, SIGGRAPH 2008
Photo tourism: Exploring photo collections in 3D,
Snavely, Seitz, & Szeliski, SIGGRAPH 2006

Reading Comments

- Finding Paths through the World's Photos
  - Computer vision + computer graphics
  - SIFT keypoints
  - What if people don't take good photos?
  - Lighting adjustment & removal/separate treatment of foreground would improve the overall quality of the results
  - Resulting video is indeed helpful for exploring / understanding a new 3D enviornment

"Image Based Tree Modeling", Tan et al., SIGGRAPH 2007

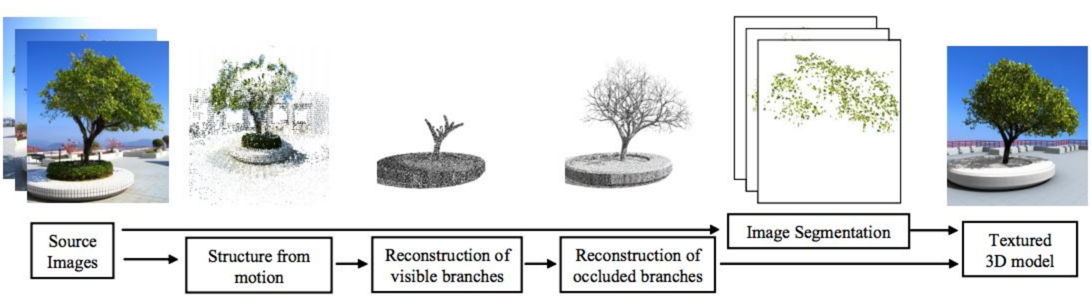

“Approximate Image-Based Tree-Modeling using Particle Flows”, Neubert et al., SIGGRAPH 2007

Figure 1: A tree is modeled using a set of input photographs. We show some examples of input and resulting 3D tree models. If image information is not available, e.g. the foliage is missing, the user is able to sketch it (right). The models approximate the input images while forming botanically plausible branching structures.

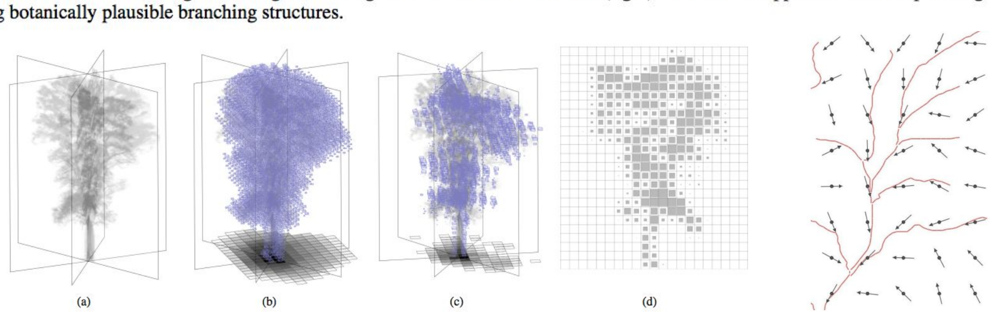

Image-Based Modeling and Photo Editing

Oh, Chen, Dorsey, & Durand, SIGGRAPH 2001

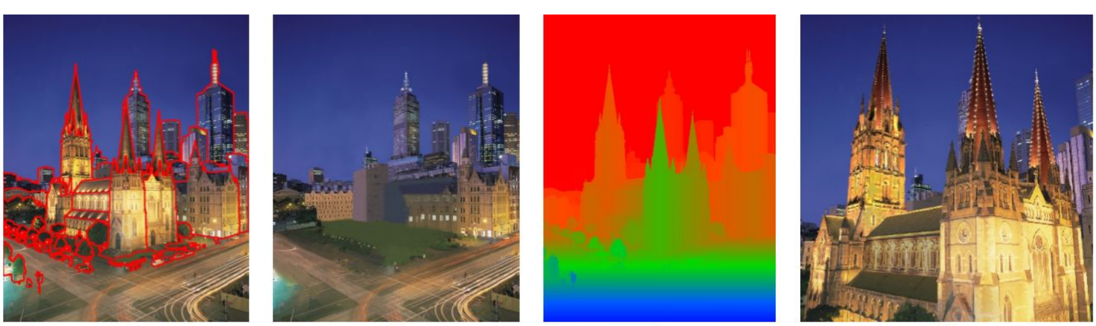

(a) (b) (c) (d)

Figure 1: St Paul’s Cathedral in Melbourne. (a) Image segmented into layers (boundaries in red). (b) Hidden parts manually clone brushed by the user. (c) False-color rendering of the depth of each pixel. (d) New viewpoint and relighting of the roof and towers.

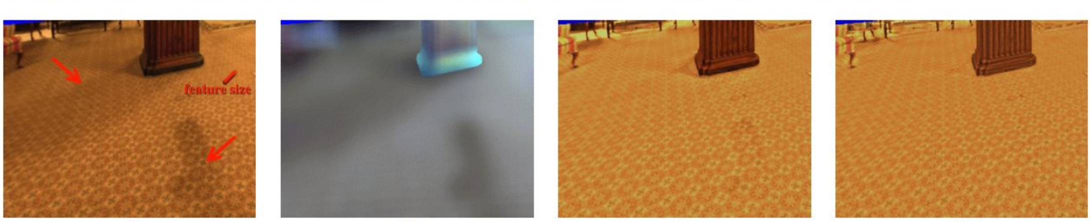

(a) (b) (c) (d)

Figure 10: Texture-illuminance decoupling. (a) Input image. (b) Initial illuminance estimation using simple Gaussian filtering. (c) Initial texture estimation, note the artifacts corresponding to shadow boundaries. (d) Texture computed using bilateral filtering.

Video Matching, Sand, Siggraph 2004

"Video Matching", Sand & Teller, SIGGRAPH 2004

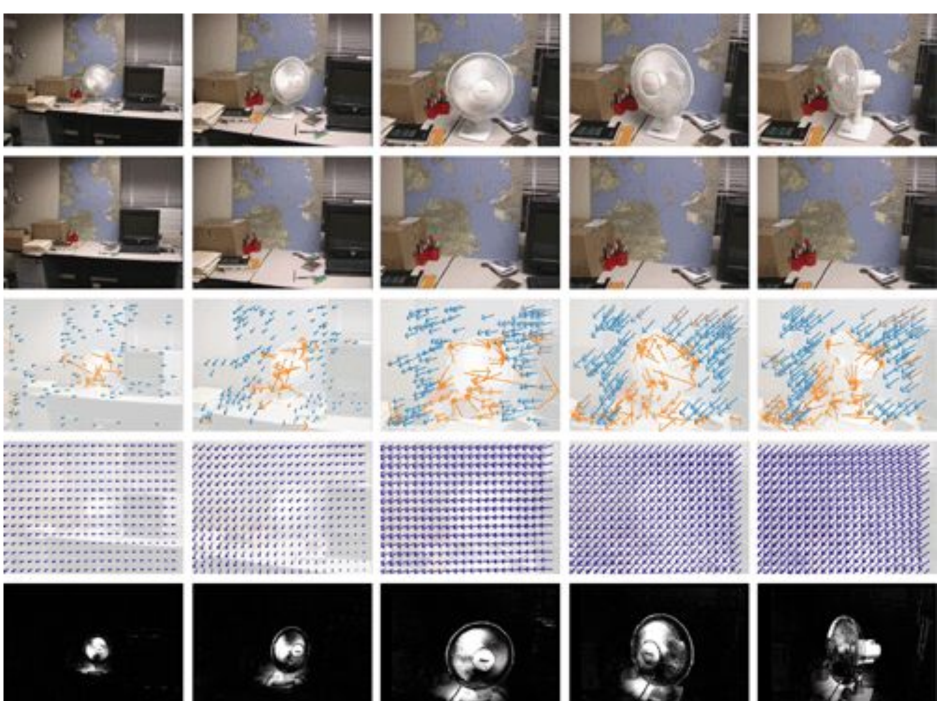

Today

- Papers for Today
- Announcements: Quiz & Final Projects
- Photography Lesson: Tilt/Shift Lenses
- Structure From Motion
- Multi-viewpoint Rendering
- Matting & Compositing
- Helmholtz Reciprocity
- Light Fields

Multi-Viewpoint Panoramas

"Photographing long scenes with multi-viewpoint panoramas", Agarwala, Agrawala, Cohen, Salesin, & Szeliski, SIGGRAPH 2006

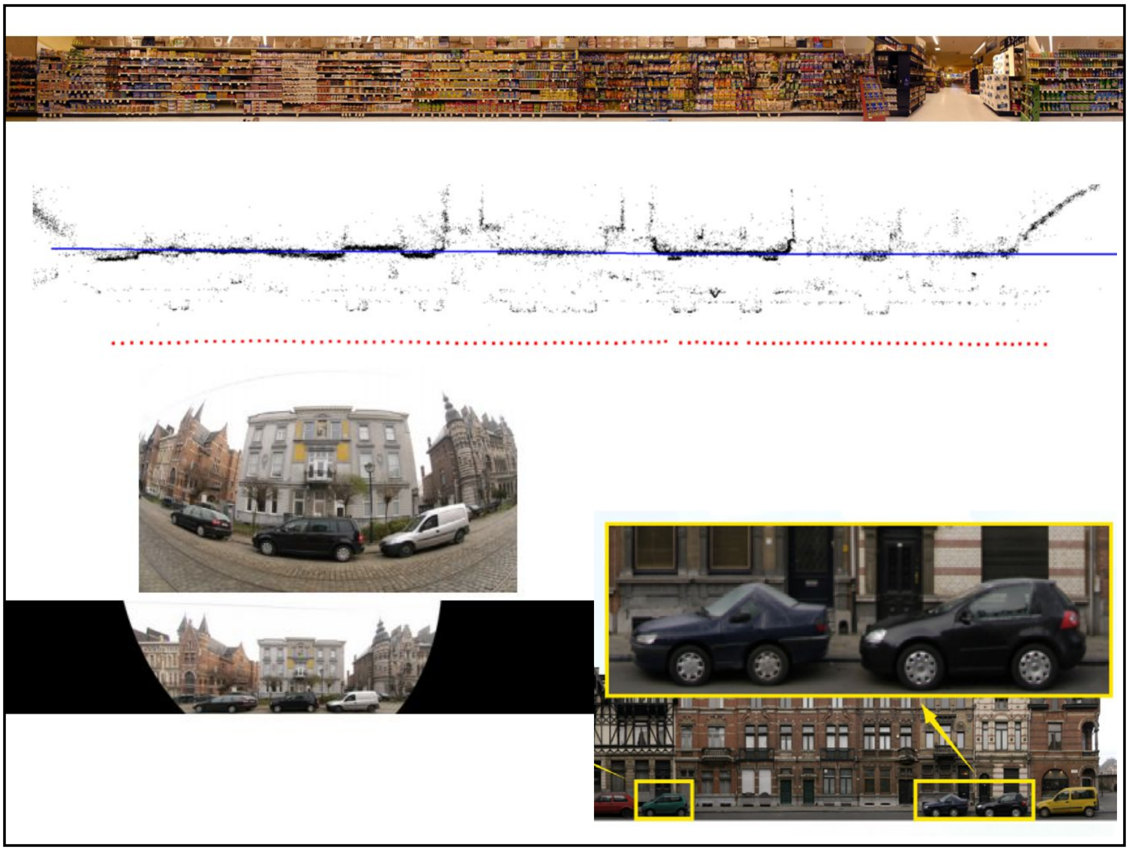

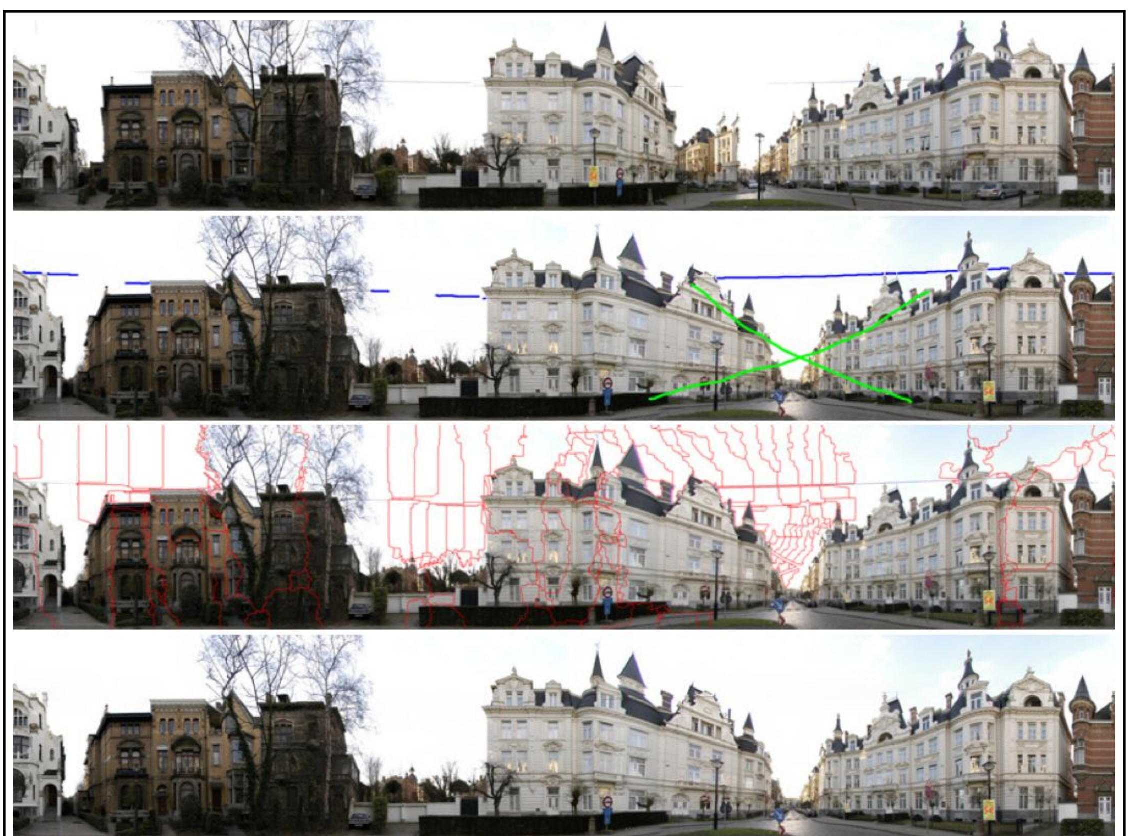

Multi-Viewpoint Panoramas

- Like many non-photorealistic rendering methods, this paper aims to mimic the style of a particular artist or style of art
- Well designed user interface:
  - Most components automated
  - User can adjust dominant plane, view selection, seams, & inpainting

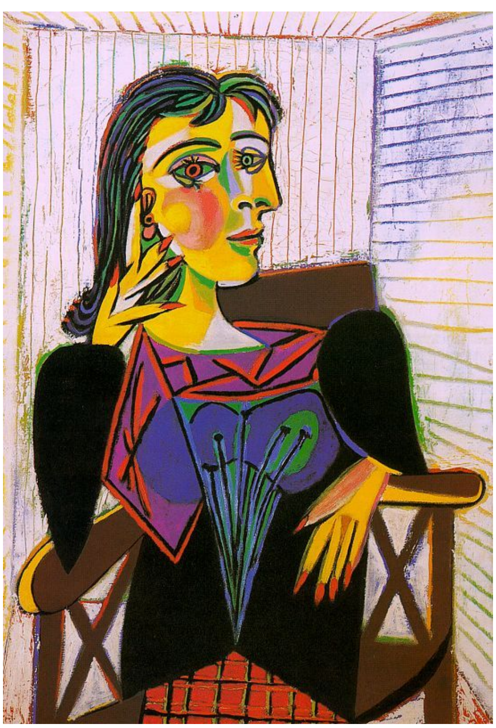

Portrait of Dora Maar
Pablo Picasso

Portrait of a Woman
Pablo Picasso

Multi-Perspective Rendering

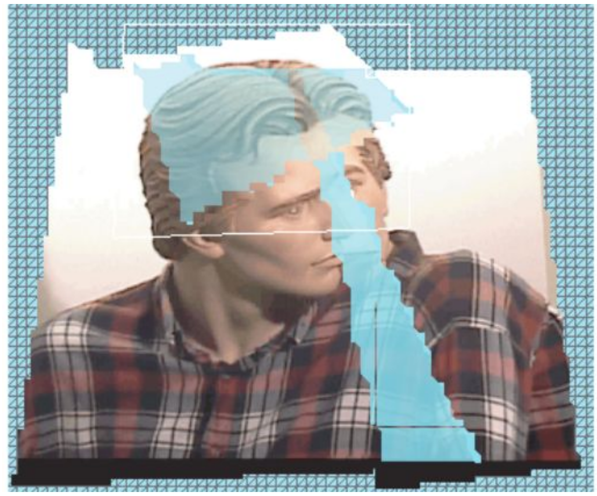

J. Yu & L. McMillan
"A Framework for Multiperspective Rendering"
Eurographics Symposium on Rendering 2004

Opening Scene from Disney's Pinocchio

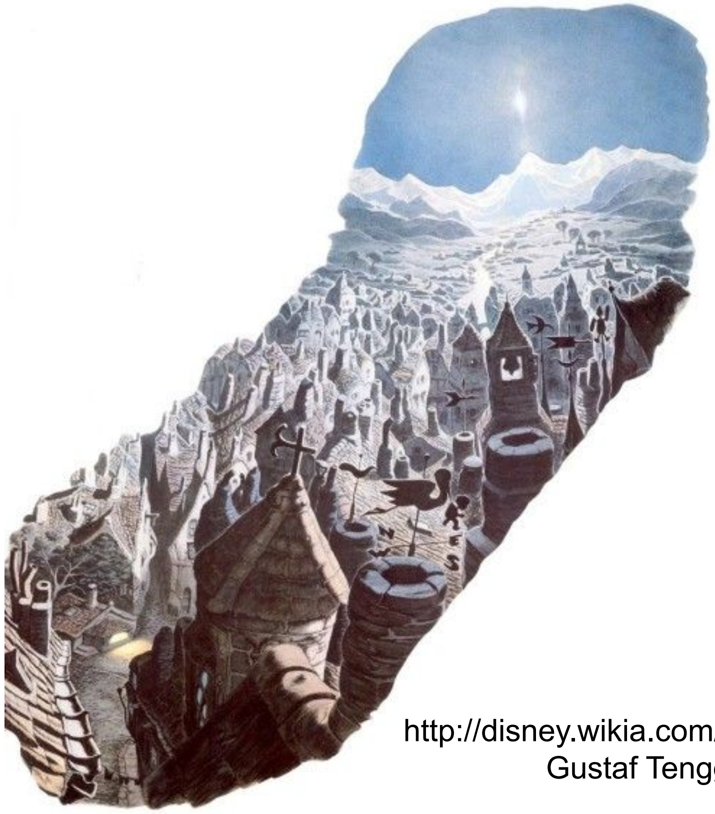

http://disney.wikia.com/wiki/Pinocchio
Gustaf Tenggren

Photo Montage

- David Hockney

http://www.hockneypictures.com/photos/photos_collages_05_large.php

Questions?

Zac Bubnick http://www.princetonol.com/groups/iad/lessons/high/cubismphoto.htm

Today

- Papers for Today
- Announcements: Quiz & Final Projects
- Photography Lesson: Tilt/Shift Lenses
- Structure From Motion
- Multi-viewpoint Rendering
- Matting & Compositing
- Helmholtz Reciprocity
- Light Fields

“Environment Matting and Compositing”
Zongker, Werner, Curless, & Salesin, SIGGRAPH 1999

“Environment Matting and Compositing”
Zongker, Werner, Curless, & Salesin, SIGGRAPH 1999

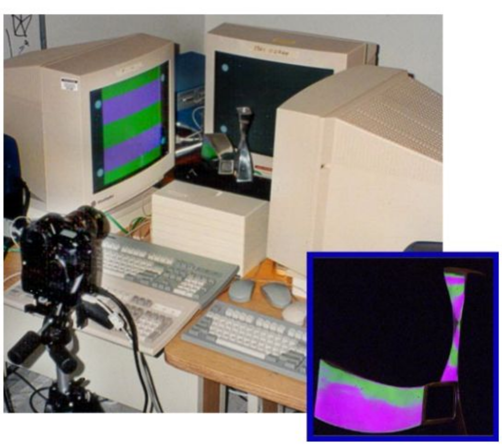

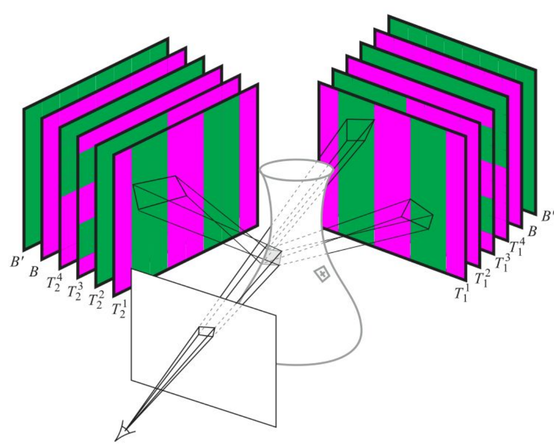

Today

- Papers for Today
- Announcements: Quiz & Final Projects
- Photography Lesson: Tilt/Shift Lenses
- Structure From Motion
- Multi-viewpoint Rendering
- Matting & Compositing
- Helmholtz Reciprocity
- Light Fields

Helmholtz Reciprocity

- BRDF is symmetric: % of light reflected from direction *i* off surface point *p* to direction *j* is the same as the % of light reflected from direction *j* off surface point *p* to direction *i*

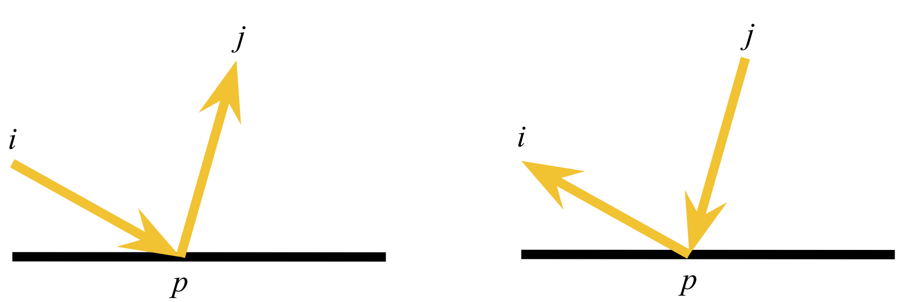

Helmholtz Reciprocity

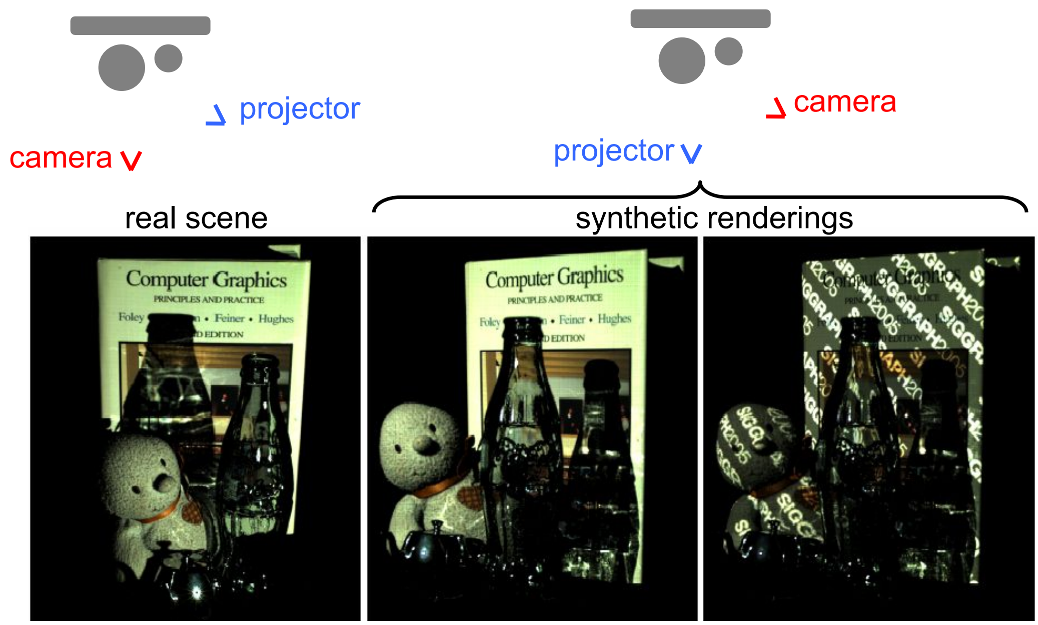

"Dual Photography", Sen, Chen, Garg, Marschner, Horowitz, Levoy, & Lensch, *SIGGRAPH 2005*

"Dual Photography", Sen, Chen, Garg, Marschner, Horowitz, Levoy, & Lensch, *SIGGRAPH 2005*

(a) (b) (c) (d)

Figure 16: Dual photography with indirect light transport. (a) A projector illuminates the front of a playing card while the camera sees only the back of the card and the diffuse page of the book. An aperture in front of the projector limits the illumination only onto the card. The card was adjusted so that its specular lobe from the projector did not land on the book. Thus, the only light that reached the camera underwent a diffuse bounce at the card and another at the book. (b) Complete camera view under room lighting. The back of the card and the page of the book are visible. It seems impossible to determine the identity of the card from this point of view simply by varying the incident illumination. To acquire the transport matrix, a 3×3 white pixel was scanned by the projector and 5742 images were acquired to produce a dual image of resolution 66×87. (c) Sample images acquired when the projector scanned the indicated points on the card. The dark level has been subtracted and the images gamma-corrected to amplify the contrast. We see that the diffuse reflection changes depending on the color of the card at the point of illumination. After acquiring the T matrix in this manner, we can reconstruct the floodlit dual image (d). It shows the playing card from the perspective of the projector being indirectly lit by the camera. No contrast enhancement has been applied. Note that the resulting image has been automatically antialiased over the area of each projector pixel.

Today

- Papers for Today
- Announcements: Quiz & Final Projects
- Photography Lesson: Tilt/Shift Lenses
- Structure From Motion
- Multi-viewpoint Rendering
- Matting & Compositing
- Helmholtz Reciprocity
- Light Fields

Light Fields

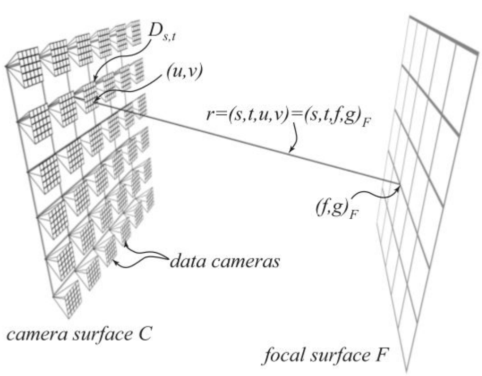

Plenoptic Modeling: An Image-Based Rendering System, McMillan & Bishop, SIGGRAPH 1995

Dynamically reparameterized light fields, Isaksen, McMillan, & Gortler, SIGGRAPH 2000

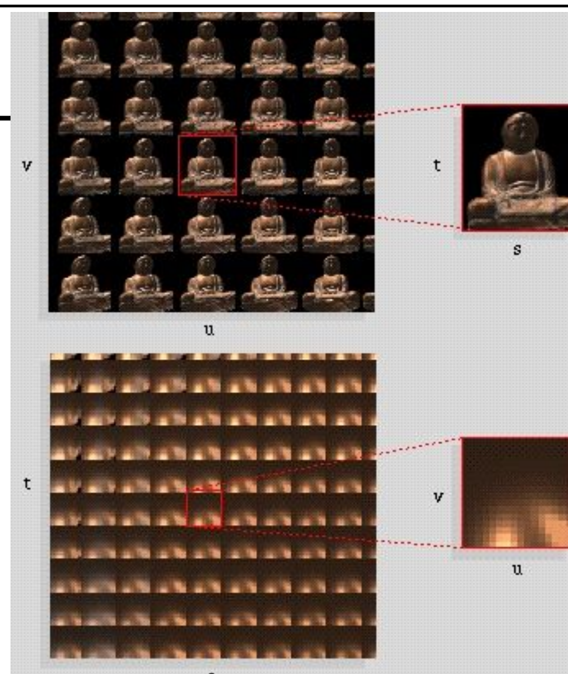

Light Field Rendering, Levoy & Hanrahan, SIGGRAPH 1996

Unstructured Lumigraph Rendering”
Buehler et al. SIGGRAPH 2001

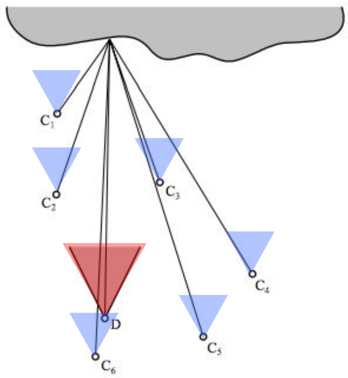

Figure 1: When available, approximate geometric information should be used to determine which source rays correspond well to a desired ray.

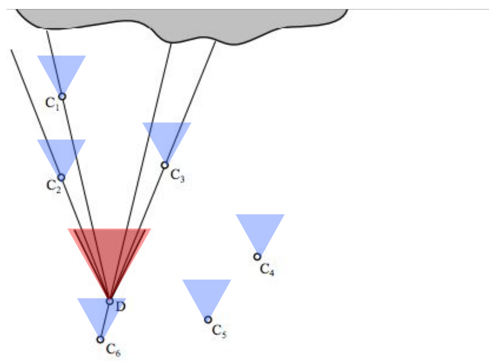

Figure 2: When a desired ray passes through a source camera center, that source camera should be emphasized most in the reconstruction.

“Unstructured Lumigraph Rendering”
Buehler et al. SIGGRAPH 2001

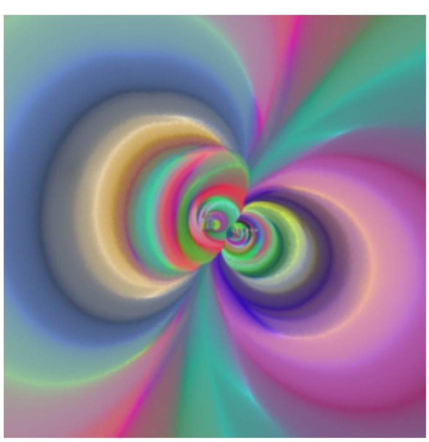

Figure 7: A visualized color blending field. Camera weights are computed at each pixel. This example is from the “hallway” dataset

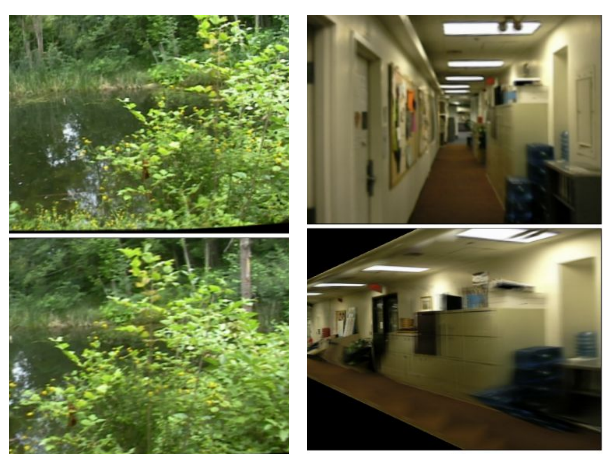

Unstructured Lumigraph, Buehler, 2001

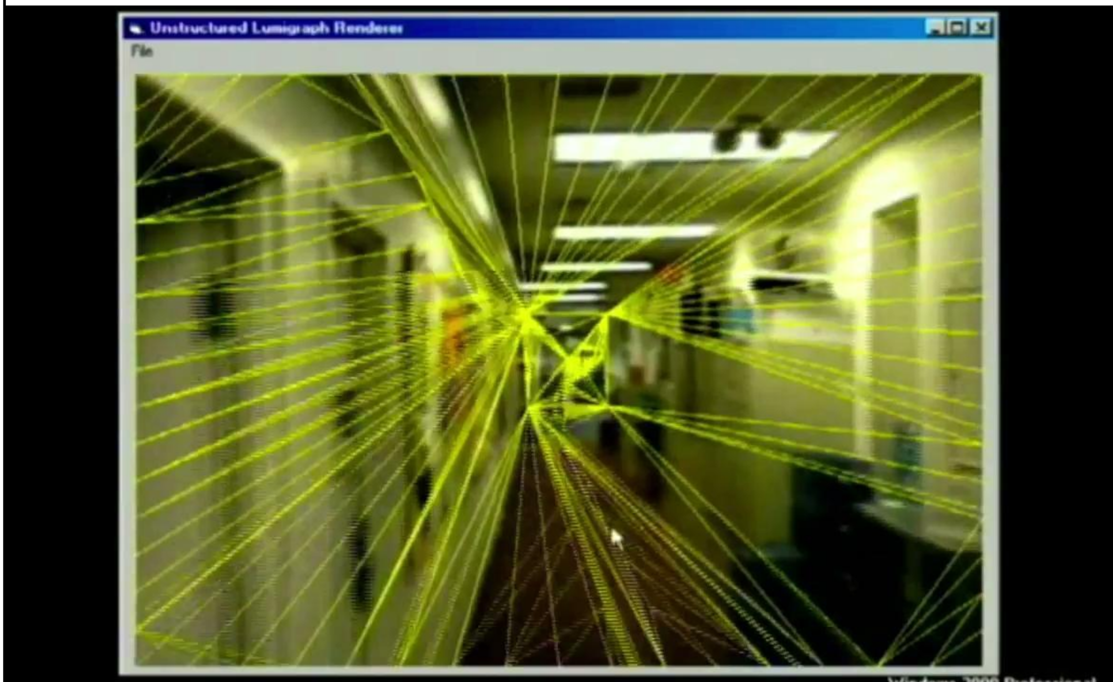

Light Field Camera

- *After* taking the photograph, we can:
  - Adjust focus
  - Change viewpoint
  - Change illumination
  - & more?

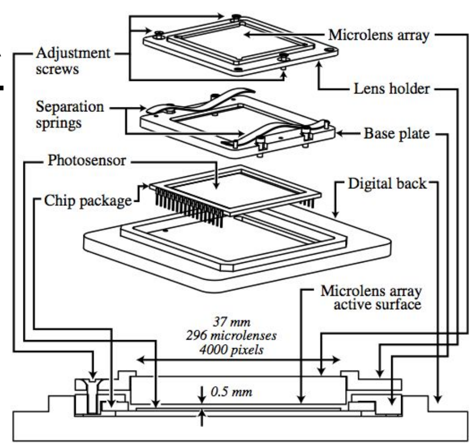

Light Field Photography with a Hand-Held Plenoptic Camera,
Ng, Levoy, Bredif, Duval, Horowitz, & Hanrahan,
Stanford Tech Report, 2005