Computational Photography

End of Semester

- Today is the last lecture!
- Quiz on Friday
  - Sample problems are posted on course website
- Final Project Presentations
  Tues Apr 30th, Fri May 3rd, Tues May 7th
  - Attendance mandatory (please don't be late!)
  - No laptops allowed during your classmates' presentations
  - You will be giving each other written feedback & peer grade
  - Ask good questions (participation grade)
- Presentation 10pts (peers)
- Project Report 20pts (instructor)

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
- Practice! & time yourself!
  - We have a tight schedule
  - I will stop you midsentence if you run over

Final Presentation Schedule

Tues May 1st	*Fri May 4th*	*Tues May 8th*
2:00	2:00	2:00
2:15	2:15	2:15
2:40	2:40	2:40
2:55	2:55	2:55
3:10	3:20	3:20
3:25 Optix presentation	3:35	3:35
3:50 done!	3:50 done!	3:50 done!

Total time (including setup & questions):
15 min (individual), 25 min (team of 2)

Last Time?

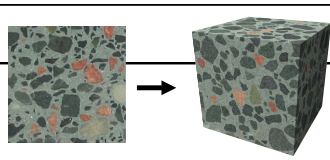

- Texture Synthesis
- Markov Model
- Image Completion
- Volumetric Texture Synthesis

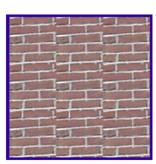

"I spent an interesting evening recently with a grain of salt."

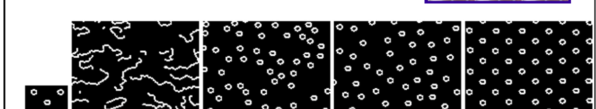

Reading from last time:

"Fragment-based image completion", Drori, Cohen-Or, Yeshurun, SIGGRAPH 2003

- Coarse to fine completion
- Confidence & traversal order
- Search for best match over different scales, rotations, & resolutions (texture frequency)
- Compositing fragments

Today

- 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

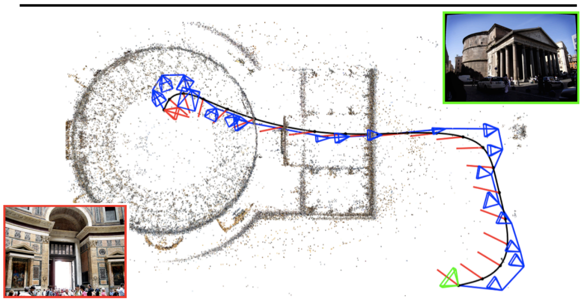

Finding Paths through the World's Photos,
Snavely, Garg, Seitz, & Szeliski, SIGGRAPH 2008
Photo tourism: Exploring photo collections in 3D,
Snavely, Seitz, & Szeliski, SIGGRAPH 2006

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

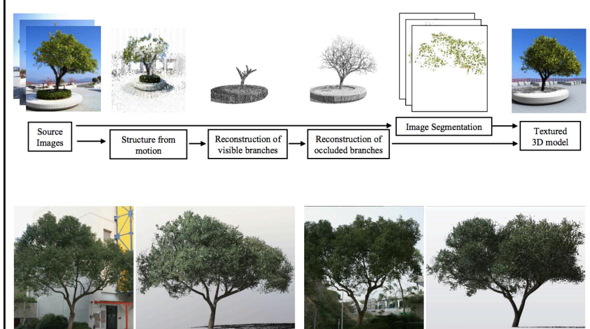

"Approximate Image-Based Tree-Modeling using Particle Flows", Neubert et al., SIGGRAPH 2007

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.

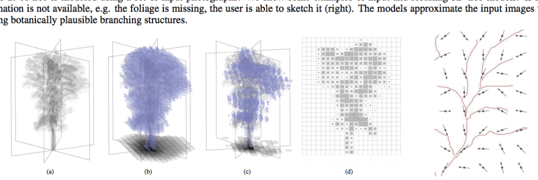

Image-Based Modeling and Photo Editing

Oh, Chen, Dorsey, & Durand, SIGGRAPH 2001

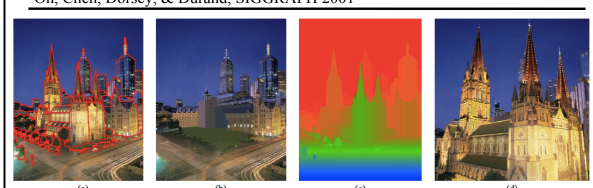

Figure 1: St Paul's Cathedral in Melbourne. (a) Image segmented into layers (boundaries in red). (b) Hidden parts manually clone brushed by the user. (c) False-color rendering of the depth of each pixel. (d) New viewpoint and relighting of the roof and towers.

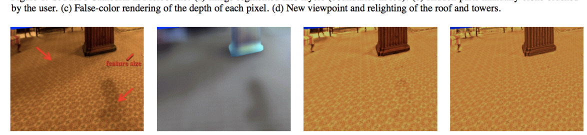

Figure 10: Texture-illuminance decoupling. (a) Input image. (b) Initial illuminance estimation using simple Gaussian filtering. (c) Initial texture estimation, note the artifacts corresponding to shadow boundaries. (d) Texture computed using bilateral filtering.

“Video Matching”,
Sand & Teller, SIGGRAPH 2004

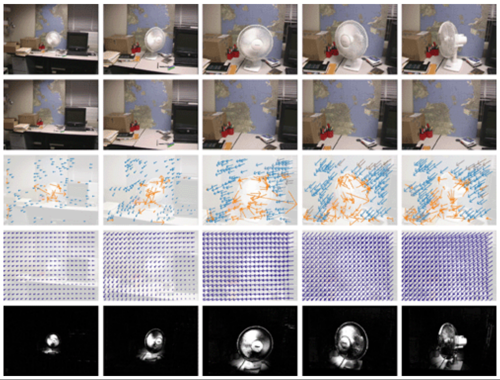

Today

- 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

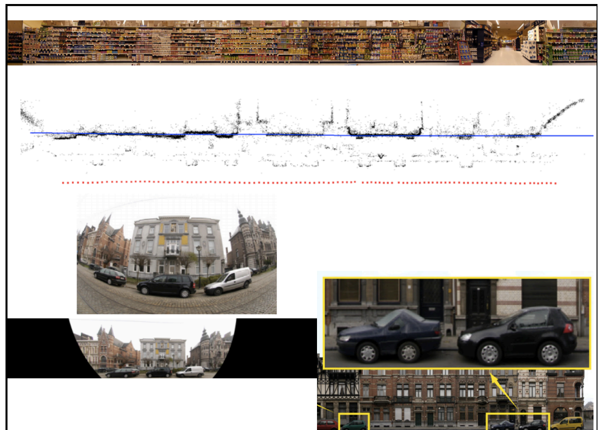

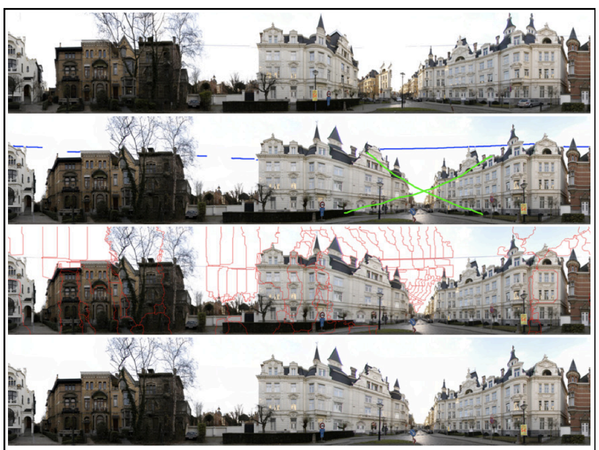

Multi-Viewpoint Panoramas

- 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

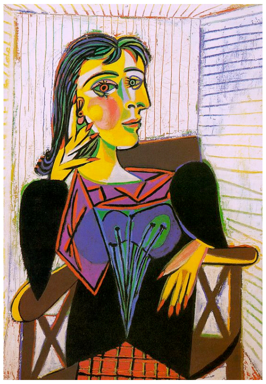

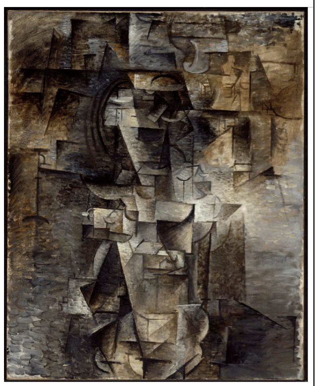

Portrait of Dora Maar
Pablo Picasso

Portrait of a Woman
Pablo Picasso

Multi-Perspective Rendering

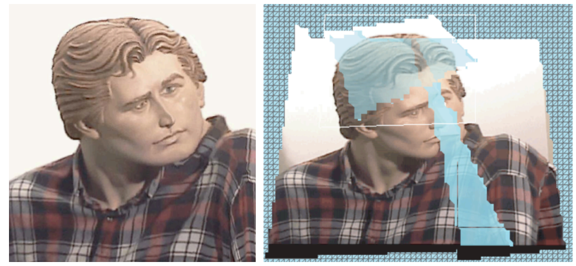

J. Yu & L. McMillan
"A Framework for Multiperspective Rendering"
Eurographics Symposium on Rendering 2004

Opening Scene from Disney's Pinocchio

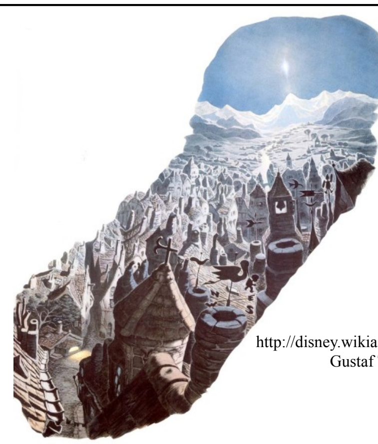

http://disney.wikia.com/wiki/Pinocchio
Gustaf Tenggren

Photo Montage

- David Hockney

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

Questions?

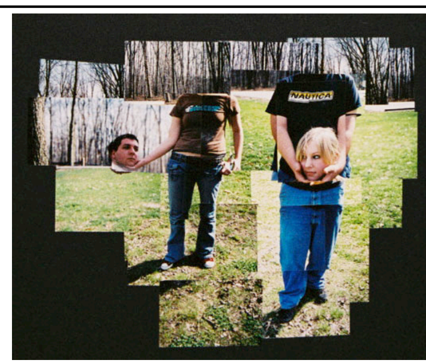

Zac Bubnick http://www.princetonol.com/groups/iad/lessons/high/cubismphoto.htm

Today

- 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

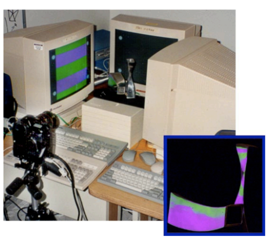

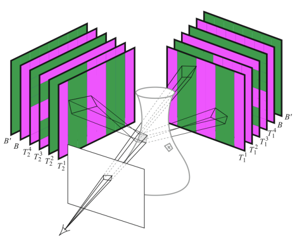

Today

- 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*

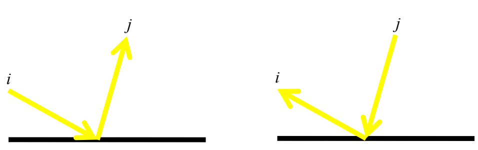

Helmholtz Reciprocity

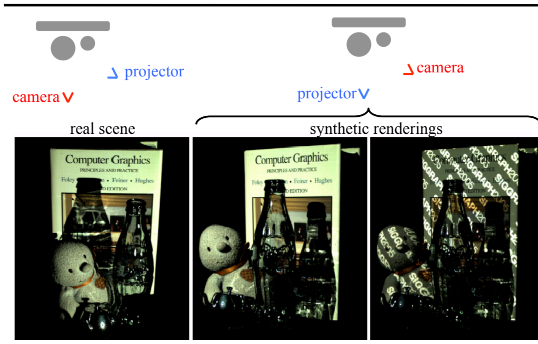

“Dual Photography”, Sen, Chen, Garg, Marschner, Horowitz, Levoy, & Lensch, *SIGGRAPH 2005*

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

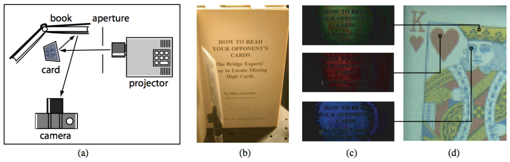

Figure 16: Dual photography with indirect light transport. (a) A projector illuminates the front of a playing card while the camera sees only the back of the card and the diffuse page of the book. An aperture in front of the projector limits the illumination only onto the card. The card was adjusted 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

- Structure From Motion
- Multi-viewpoint Rendering
- Matting & Compositing
- Helmholtz Reciprocity
- Light Fields

Light Fields

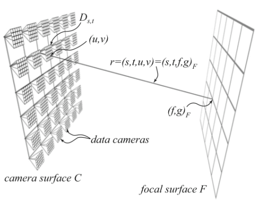

Plenoptic Modeling: An Image-Based Rendering System, McMillan & Bishop, SIGGRAPH 1995

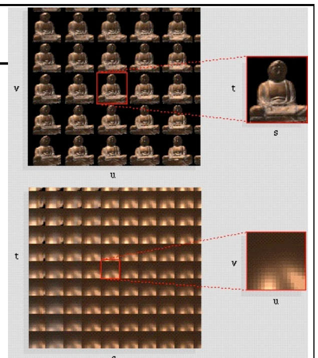

Dynamically reparameterized light fields, Isaksen, McMillan, & Gortler, SIGGRAPH 2000

Light Field Rendering, Levoy & Hanrahan, SIGGRAPH 1996

Unstructured Lumigraph Rendering" Buehler et al. SIGGRAPH 2001

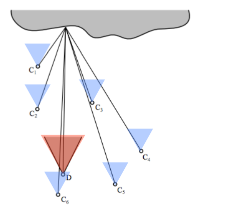

Figure 1: When available, approximate geometric information should be used to determine which source rays correspond well to a desired ray.

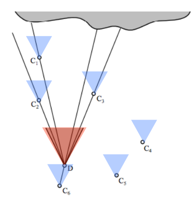

Figure 2: When a desired ray passes through a source camera center, that source camera should be emphasized most in the reconstruction.

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

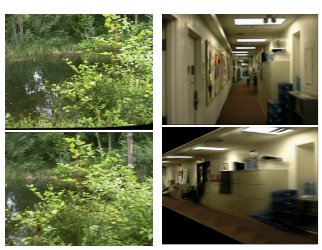

Light Field Camera

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

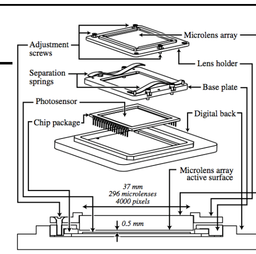

Light Field Photography with a Hand-Held Plenoptic Camera, Ng, Levoy, Bredif, Duval, Horowitz, & Hanrahan, Stanford Tech Report, 2005

"Coded Rolling Shutter Photography: Flexible Space-Time Sampling" Gu, Hitomi, Mitsunaga, & Nayar, ICCP 2010

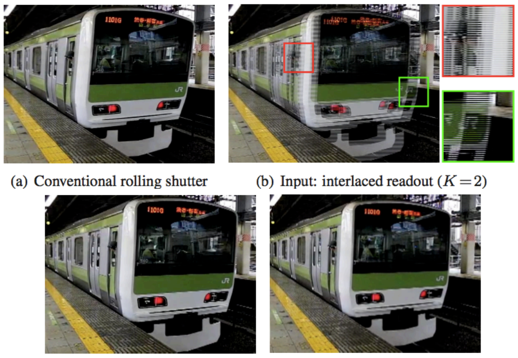

(a) Conventional rolling shutter (b) Input: interlaced readout ($K = 2$)

(c) Interpolated sub-image I_1 (d) Interpolated sub-image I_2

"Coded Rolling Shutter Photography: Flexible Space-Time Sampling" Gu, Hitomi, Mitsunaga, & Nayar, ICCP 2010

- Global Shutter vs. Rolling Shutter *plus* Coded
- Interlaced vs. Staggered
- Skew Compensation
- High Speed Photography
- Interpolation of High Resolution
- High Dynamic Range
- Adaptive Row-wise Auto Exposure
- Simulation ➔ Prototype Camera Hardware