


CSCI 4530/6530 Advanced Computer Graphics

<https://www.cs.rpi.edu/~cutler/classes/advancedgraphics/S25/>

Lecture 23: Computational Photography





<https://imgur.com/a/Rj8tU>



Last Time?

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

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




Today


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

The Sandpit, O'Hare, 2010






Camera Obscura / Pinhole Camera








Optics: The principle of the camera obscura. Engraving, 1752.

Camera Obscura / Pinhole Camera

- “Pinhole” = tiniest aperture/opening
 - Limited light reaches the image plane
 - Requires very sensitive film / sensors and/or long exposure (stationary scene)
 - Entire scene is in focus
- Larger aperture/opening
 - Lens required to collect additional light, but bend it to land on a single point on the image plane
 - Lens geometry optimized for distance ratio object-to-lens : lens-to-image-plane
 - Only objects near this optimal distance are in focus



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-shift_photography#/media/File:24mm-tilt-lens.jpg

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




Shift/Rise for Perspective-Control




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


https://en.wikipedia.org/wiki/Tilt-shift_photography



(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.

Tilt/Swing for Focus Control



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

Focus in the distance



Focus on the foreground




**Tilted focal plane,
(most) everything in focus**



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

Tilt/Swing for Selective Focus



Tilt/Swing for Miniature Effect



Ben Thomas
[http://benthomas.co/
cityshrinker/wamrft5d
4mk0ua9ahjxtu9cu2q
txsp](http://benthomas.co/cityshrinker/wamrft5d4mk0ua9ahjxtu9cu2qtxsp)

Tilt/Swing for Miniature Effect



Ben Thomas
[http://benthomas.co/
cityshrinker/4pgljfasa
898t1z6cjni03pilyaxzi](http://benthomas.co/cityshrinker/4pgljfasa898t1z6cjni03pilyaxzi)

Tilt/Swing for Miniature Effect

- How to achieve the effect?
 - High-angle shot, from “above” & from a distance
 - Using a shallow depth of field,
so the foreground and background are 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

Questions?

Today

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

Quiz 2

- Friday April 11th
- During lecture time, 2:00-3:50 pm
- Practice problems are on the Calendar / in Course Materials
- 1 Page of notes allowed
(front & back of 1 piece of paper)
handwritten or printed

Today

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

Final Project 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 or 3:
 - All 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!***

Final Presentation Schedule

10 min (individual)

18 min (team of 2)

26 min (team of 3)

+2 min questions/setup

Tue Apr 15th

2:00 X

2:12 X&X

2:32 X

2:44 X&X

3:04 X

3:16 X&X

3:36 X

3:48 done!

Fri Apr 18th

2:00 X&X&X

2:28 X

2:40 X&X&X

3:08 X

3:20 X&X

3:40 done!

Tue Apr 22nd

2:00 X&X&X

2:28 X

2:40 X&X

3:00 X

3:12 X&X

3:32 X

3:44 done!

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

Peer Grading Final Project Presentations

- Tuesday April 15th & Friday 18th & Tuesday April 22nd:
- ***Mandatory attendance for everyone for all presentations***
- Laptops should be closed during your peers' presentations
- Time to ask questions after each presentation
- Enter "peer grading" & written feedback via Submittity
 - Bring your laptop (phone probably ok...)
- ***NOTE: We expect that students presenting the 15th & 18th may not have finished their project yet, since the report is not due until Wednesday April 23rd.***

Today

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

Final Project Report

- Due Wednesday April 23rd @ 11:59pm
- No late days may be used on the Final Project Report
- What to Submit:
 - .pdf report
 - 200x200 .png thumbnail
 - All code that you wrote
 - NOTE: Reports will be posted on course webpage, unless you opt out by emailing instructor

Final Report

Your final report should be formatted using pdf and follow a standard technical writing format (e.g., motivation, related work, algorithm/technique, results, conclusions, and bibliography). All students are encouraged to format their final report for submission to a relevant academic graphics or graphics-related conference ([SIGGRAPH](#), [Eurographics](#), etc.). The report must be a minimum of 2000 words (equivalent of 8 pages double spaced text, excluding figures & references). Your report should describe the technical details of your project; in particular, we want to know about:

- Any algorithms or data structures you implemented.
- The core features of your assignment and how you tested them.
- The challenges that you overcame (or failed to overcome). *Note: As you're working on your project, save "bloopers" images or video that show your intermediate results and how you debugged your project, and include them in your report & presentation.*
- Images/screenshots/visualizations/video showing the results of your project. Include simple and moderately complex examples.
- Any known bugs or limitations in your implementation, and potential avenues for future work.
- How long it took you to complete the assignment, and who did what.

Make sure to acknowledge (in proper bibliographic format) all references you consulted to complete the assignment and the extent of any collaboration with other students or outside sources. As with the other assignments, submit your source code, any necessary data files, and sample images or video showing the results.

The final project reports are due on the date specified on the calendar. *You may not use any late days for the report.* We will not accept any late material unless accompanied by a note from the Dean of Students office.

Please include a representative 200x200 image with your submission for the project index page (see [final projects from prior years](#)). All projects reports will be posted online by default. If you do not want your project report posted online please send the instructor an email.

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
 - with descriptive captions!
- Conclusions & Potential Future Work
 - limitations of the method are clearly stated


Questions?

Today

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

Reading for Today


- "An Approximate Image-Space Approach for Interactive Refraction", Chris Wyman, SIGGRAPH 2005



- Results are impressive
- Surprising how well it worked (almost too good to be true?)
- Technique is not perfect, but actually quite realistic when evaluated!
- Refraction is confusing and humans cannot predict or spot errors
- Short and concise and well-written
- Image-space algorithm, everything recomputed every frame (real-time!)

Reading for Today

- “Glaze: Protecting Artists from Style Mimicry by Text-to-Image Models”
Shan, Cryan, Wenger, Zheng, Hanocka, & Zhao
Proceedings of USENIX Security Symposium, 2023.



- Paper emphasizes IRB (Institutional Review Board) approval!
- Survey of experienced professional & part-time/free-lance artists
- 91% knowledgeable and worried about their art being used in AI
 - (they can spot flaws in AI art, but assume public may not be able to)
 - Primary concern: work taken without permission & compensation
- Decreased job security (especially junior positions), discourage students from studying art, diminish creativity, require limiting sharing of art online
- Previous work (cloaking to prevent facial recognition) aims to protect all information/features in the image rather than just the style.
- Cloaking artwork with a contrasting style results in mimicked artwork that is an undesirable mix of the actual artist style & cloaked style
- Impressive evaluation, using non-public artwork by volunteer professional artists & huge diverse set of styles from publicly available art/artists

- In 2 years since publication, new AI models have broken Glaze's protection (cite source?)
- Conspiracy Theory: Glaze could be stealing artists original work
- "Dead Internet Theory"
- All Computer Science students should learn & study ethics
- If AI eliminates all human artists, no new art styles will be created
- In addition to artists evaluation, would be interesting to get the public/non-artist feedback on both glazed art and the mimicry failures
- Reminded of prior work to auto-solve captchas

Reading for Today

- “Nightshade: Prompt-Specific Poisoning Attacks on Text-to-Image Generative Models”, Shan, Ding, Passananti, Wu, Zheng & Zhao, 2024



- Models trained on millions/billions of images, but only thousands of images of a specific concept - “concept sparsity”
- Do we need a ‘backdoor’ with sufficient access to poison the model?
 - Even massive data models can be broken by small amounts of corrupt data
- Is Nightshade “abuse”?
 - Opt in/opt out is being ignored by AI scrapers (cite source?)
 - Can AI complain about being ‘attacked’ by corrupted data it steals?
 - Might it harm human users who rely on accurate captions & alt text (internet accessibility) and not just protect artists from AI theft?

Today


- Photography Lesson: Tilt/Shift Lenses
- Announcements: Quiz
- Announcements: Final Project Presentations Schedule
- Announcements: Final Project Report
- Papers for 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



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 environment



“Image Based
Tree Modeling”,
Tan et al.,
SIGGRAPH 2007





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








Image-Based Modeling and Photo Editing

Oh, Chen, Dorsey, & Durand, SIGGRAPH 2001



Today

- Photography Lesson: Tilt/Shift Lenses
- Announcements: Quiz
- Announcements: Final Project Presentations Schedule
- Announcements: Final Project Report
- Papers for 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



“A Framework for Multiperspective Rendering”

J. Yu & L. McMillan, 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

- Photography Lesson: Tilt/Shift Lenses
- Announcements: Quiz
- Announcements: Final Project Presentations Schedule
- Announcements: Final Project Report
- Papers for 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



“Video Matching”,
Sand & Teller,
SIGGRAPH 2004





Matte-Based Effects

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





Today

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


Helmholtz Reciprocity

- The BRDF is symmetric!
- % of light reflected from direction i at surface point p in direction j
is the same as the
% of light reflected from direction j at surface point p in direction i





Helmholtz Reciprocity



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



real scene



synthetic renderings



“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


- Photography Lesson: Tilt/Shift Lenses
- Announcements: Quiz
- Announcements: Final Project Presentations Schedule
- Announcements: Final Project Report
- Papers for 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




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

“Unstructured Lumigraph Rendering” Buehler et al. SIGGRAPH 2001




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

Questions?
