# Handling Massive or Incomplete Information

## Plan for Today

- Tangentially Related to Privacy…
  Reconstruction from Partial Information
- Readings for Today:
  - "QSplat: A Multiresolution Point Rendering System for Large Meshes"
  - "LabelMe: online image annotation and applications"
- How to brainstorm/foster radical ideas?
- How to formally review a paper for acceptance to a conference/journal
- Homework 7 & Homework 8
- Readings for Tuesday

# Synthetic aperture confocal imaging

Marc Levoy
Billy Chen
Vaibhav Vaish
Mark Horowitz
Ian McDowall
Mark Bolas



# Synthetic aperture photography using an array of mirrors

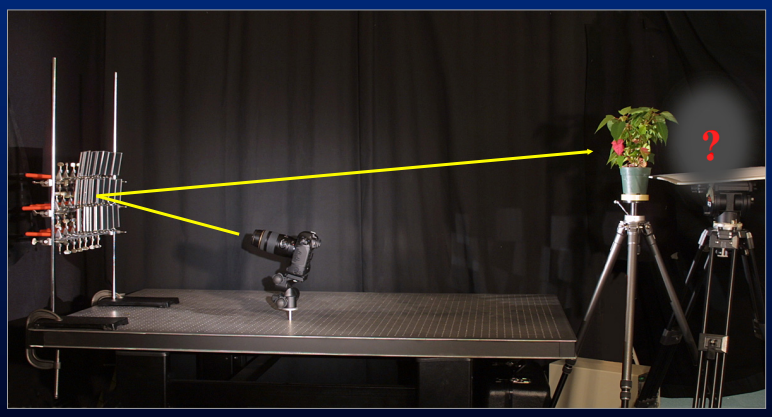


- 11-megapixel camera
- 22 planar mirrors


© 2004 Marc Levoy

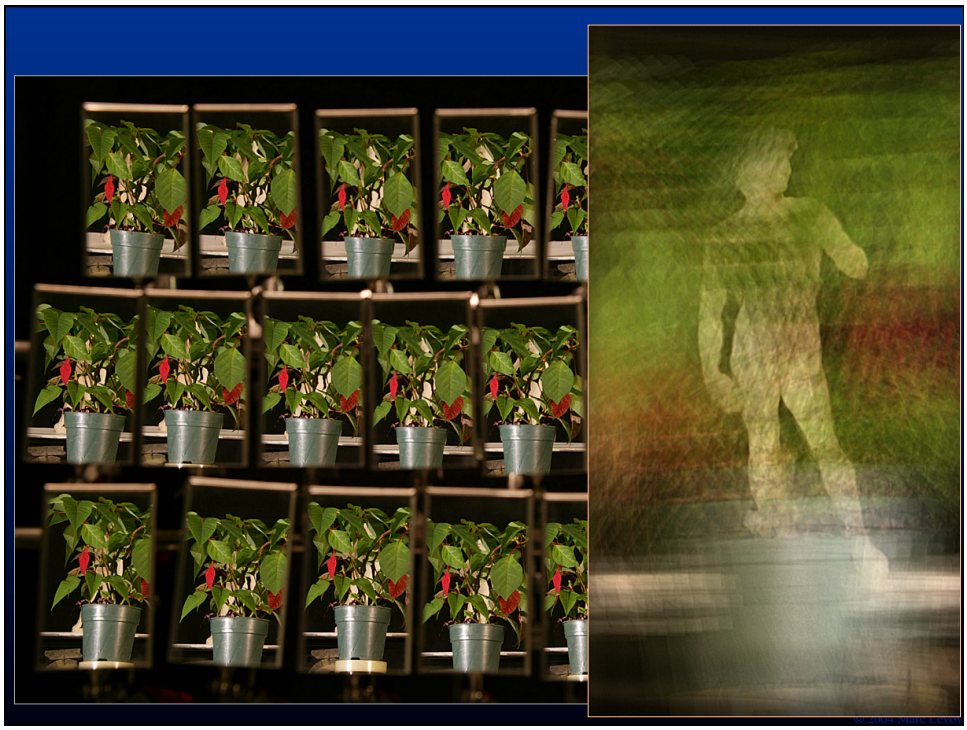
© 2004 Marc Levoy

# Confocal imaging in scattering media

- small tank
  - too short for attenuation
  - lit by internal reflections



# Stanford Multi-Camera Array

[Wilburn 2002]

- 640 × 480 pixels × 30fps × 128 cameras
- synchronized timing
- continuous video streaming
- flexible physical arrangement

# Experiments in a large water tank

- stray light limits performance
- one projector suffices if no occluders

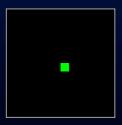



# Seeing through turbid water

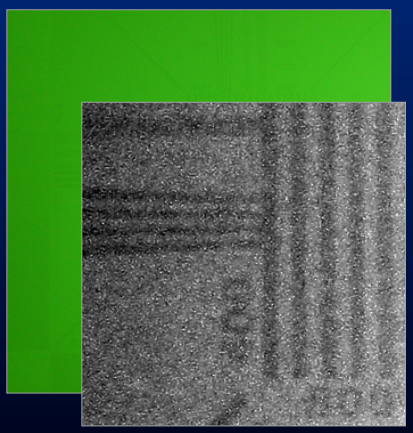

floodlit

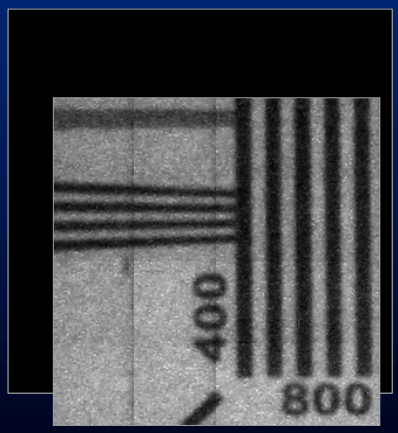


scanned tile

# Application to underwater exploration

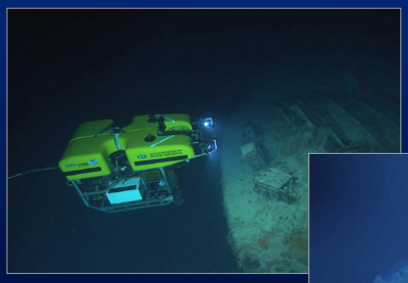

[Ballard/IFE 2004]

[Ballard/IFE 2004]



# BRDF: Bidirectional Reflectance Distribution Function

- Ratio of light coming from one direction that gets reflected in another direction
- 4D function: incoming $\theta_i, \phi_i$ outgoing $\theta_o, \phi_o$

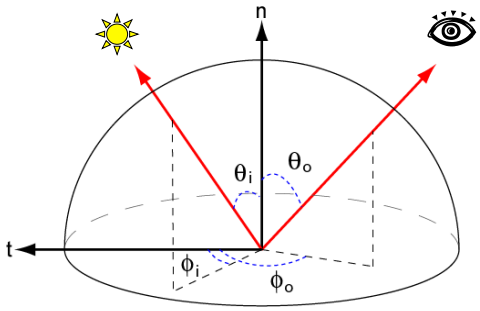


Complete material data capture: Gonioreflectometer

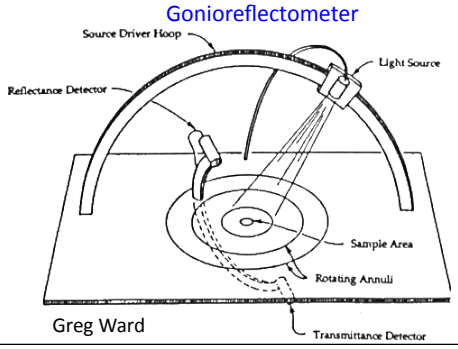


Greg Ward

# 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*

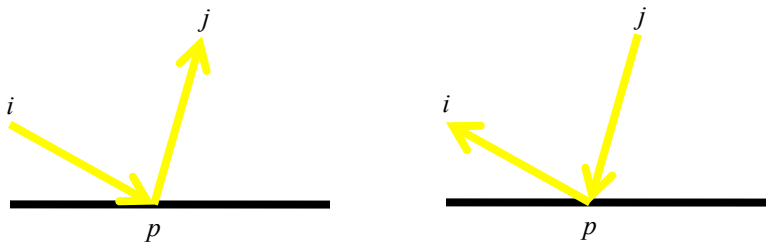


# Helmholtz Reciprocity

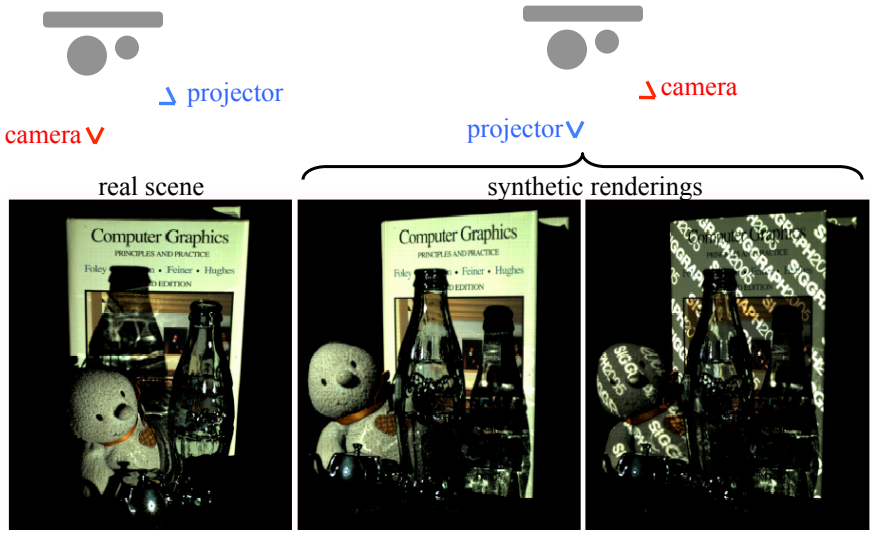


"Dual Photography", Sen, Chen, Garg, Marschner, Horowitz, Levoy, & Lensch, *SIGGRAPH 2005*

## "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 $3 \times 3$ white pixel was scanned by the projector and 5742 images were acquired to produce a dual image of resolution $66 \times 87$. (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.

## Plan for Today

- Tangentially Related to Privacy… Reconstruction from Partial Information
- Readings for Today:
  - "QSplat: A Multiresolution Point Rendering System for Large Meshes"
  - "LabelMe: online image annotation and applications"
- How to brainstorm/foster radical ideas?
- How to formally review a paper for acceptance to a conference/journal
- Homework 7 & Homework 8
- Readings for Tuesday

## Readings for Today:

- "QSplat: A Multiresolution Point Rendering System for Large Meshes", Rusinkiewicz & Levoy,
- SIGGRAPH 2000

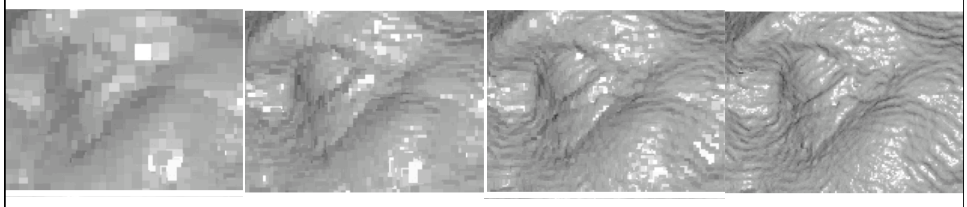

## Design Criteria & Guaranteed Interactivity

- Low powered computer
  - portable
- Interactive visualization
  - E.g., find holes in data -> suggest new scan points
- 127 million points
- Pixel size
  - Screen Space criteria
  - Visibility culling
- File layout & streaming data, pre-fetching

- Good discussion between memory & runtime
- Wanted to jump in and start coding! Explained in sufficient detail to reproduce.
  - Nice comparison of different ways to render splats
- Well-written
  - Why was related work at the end of the paper? (Where is the "right" place?)
  - Discussed the pros & cons of each decision they had to make
- Data structure description was lacking
- Arbitrary constants
- Relatively low frame rate? (5-10 fps)
- Non-traditional rendering algorithm
- Big data + interactivity + rendering
- Question about storing normals/colors

# PointShop3D

- Why deal with triangles & connectivity when the triangles are smaller than a pixel?
- "Pointshop 3D: An Interactive System for Point-Based Surface Editing" Zwicker, Pauly, Knoll, Gross, SIGGRAPH 2002

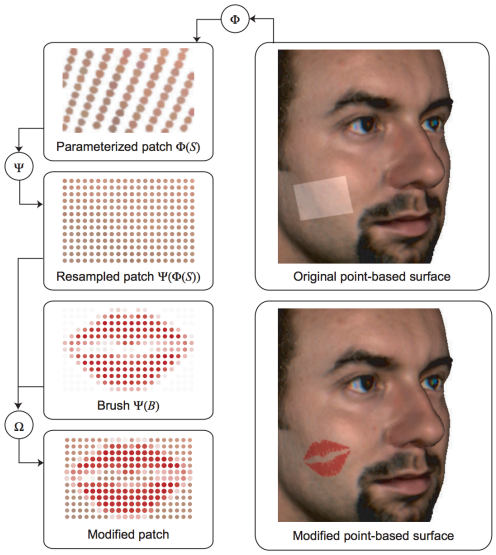


**Figure 2:** Overview of the operator framework for point-based surface editing.

# Plan for Today

- Tangentially Related to Privacy…
  Reconstruction from Partial Information
- Readings for Today:
  - "QSplat: A Multiresolution Point Rendering System for Large Meshes"
  - "LabelMe: online image annotation and applications"
- How to brainstorm/foster radical ideas?
- How to formally review a paper for acceptance to a conference/journal
- Homework 7 & Homework 8
- Readings for Tuesday

# Readings for This Week:

- "LabelMe: online image annotation and applications"
  Torralba, Russell, & Yuen, IEEE, 2010

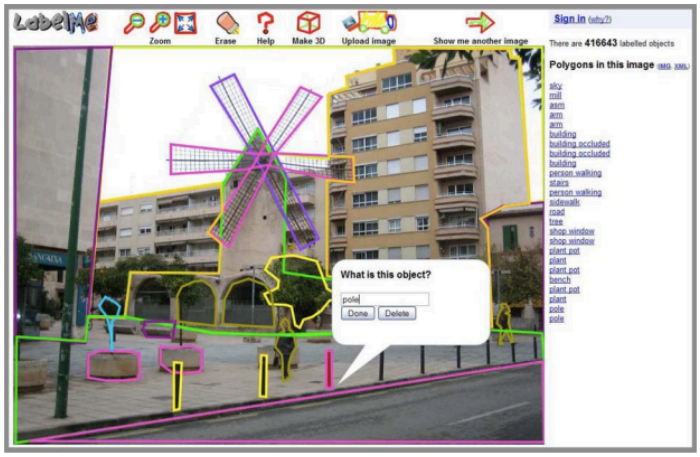

- object recognition under a variety of conditions, object class recognition rather than object instance recognition, not just canonical pose, learning about objects embedded in a scene, not just caption or tag
- number of labels, number of objects with each label still growing over time, who are the labelers?
- visualize closest neighbors to a specific image in database
- visualization of images organized by similarity, smooth transition between some types of scenes
- what types of images are present in the database? lower performance on indoor scenes (fewer examples, more variability of visual appearance)
- automatically recovered spatial relationships between objects (standing on, supported by, supported by, part of)
- Current/Future work: extrapolate scene knowledge to unlabeled portion of image, infer 3D! & extend to video

- "Crowd-sourced" data labeling
  - free & public (lots of data, but will need to deal with some junk data, if it is intentional sabotage it may skew results)
  - Data collection is "less glamorous" than designing new algorithms
  - Limited to user's understanding of the object, and their energy/laziness to completely and accurately outline & label
  - How do they check for bad data?
  - Google's game for helping with image search discontinued because of spam/junk descriptions
- Good choice of saturated colors for contrast with the natural images that are being labeled
- Images should be higher resolution?
- Construct new scenes by pulling parts from other images
  - How easy is it for a 3rd party to download & use this data? Format seems complicated.
- When you collect this much data, something (can't necessarily predict what) will be discovered. But is it and the quantity of other conclusions worth the investment?
  - Conclusions somewhat unreliable
- Security camera footage

# Image segmentation from Quick Approximate Outlining

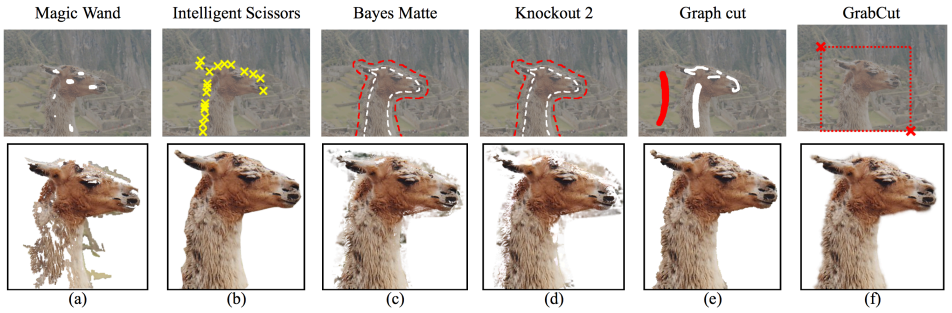


Figure 2: **Comparison of some matting and segmentation tools.** The top row shows the user interaction required to complete the segmentation or matting process: white brush/lasso (foreground), red brush/lasso (background), yellow crosses (boundary). The bottom row illustrates the resulting segmentation. GrabCut appears to outperform the other approaches both in terms of the simplicity of user input and the quality of results. Original images on the top row are displayed with reduced intensity to facilitate overlay; see fig. 1. for original. Note that our implementation of Graph Cut [Boykov and Jolly 2001] uses colour mixture models instead of grey value histograms.

"GrabCut -Interactive Foreground Extraction using Iterated Graph Cuts",
Rother, Kolmogorov, Blake, SIGGRAPH 2004

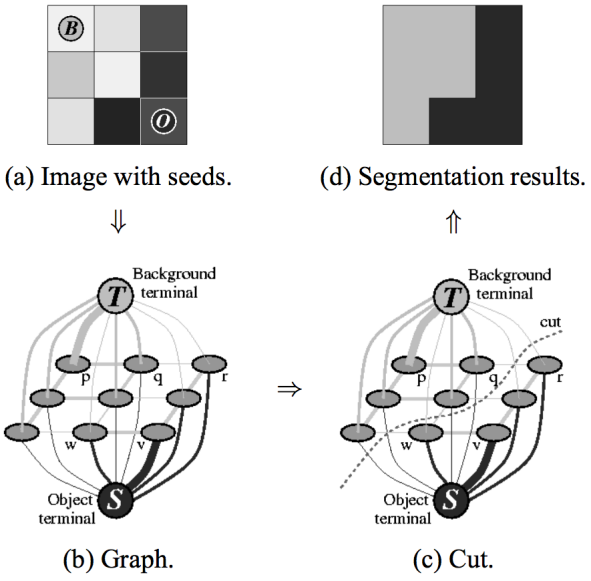


``Graph Cuts and Efficient N-D Image Segmentation''
Boykov & Funka-Lea, IJCV 2006

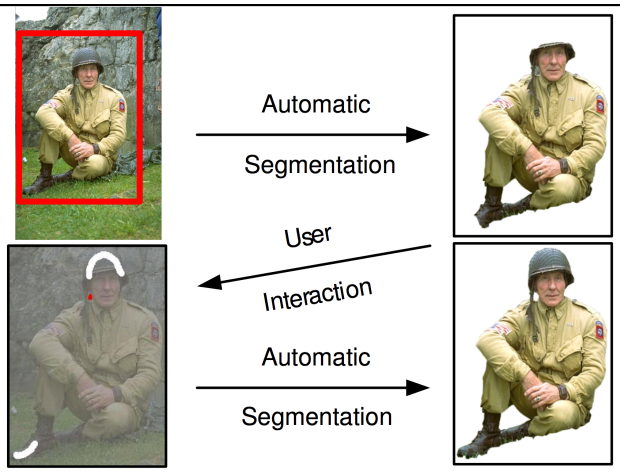


Figure 5: **User editing.** After the initial user interaction and segmentation (top row), further user edits (fig. 3) are necessary. Marking roughly with a foreground brush (white) and a background brush (red) is sufficient to obtain the desired result (bottom row).

"GrabCut -Interactive Foreground Extraction using Iterated Graph Cuts", Rother, Kolmogorov, Blake, SIGGRAPH 2004

# Plan for Today

- Tangentially Related to Privacy… Reconstruction from Partial Information
- Readings for Today:
  - "QSplat: A Multiresolution Point Rendering System for Large Meshes"
  - "LabelMe: online image annotation and applications"
- How to brainstorm/foster radical ideas?
- How to formally review a paper for acceptance to a conference/journal
- Homework 7 & Homework 8
- Readings for Tuesday

## How to encourage/foster/recognize radical new ideas?

- Be open-minded
- Brainstorming rule: generate ideas, no negativity, no early criticism/rejection
- Be a "Paper Champion", not a "Paper Killer"
- Suggest other applications/datasets
- Don't say what's wrong, say how to improve it
- Really specific comments, not general criticism
- Use "I" statements (Say "I think…" or "I feel")
  - Not "your paper is bad" but "I feel your paper could be better if you do… "

## Formal task of paper reviewer (1 of 2)

- Description: Briefly describe the paper and its contribution to computer graphics and interactive techniques. Please give your assessment of the scope and magnitude of the paper's contribution.
- Clarity of Exposition: Is the exposition clear? How could it be improved?
- Quality of References: Are the references adequate? List any additional references that are needed.
- Reproducibility: Could the work be reproduced from the information in the paper? Was any code or data submitted with the supplemental materials? If so, does it support the claims in the paper? Are all important algorithmic or system details discussed adequately in the paper?

Review form for SIGGRAPH

# Formal task of paper reviewer (2 of 2)

- Rating: Please rate this paper on a continuous scale from 1 to 5, where:
  - 1 = Definitely reject. I would protest strongly if it's accepted.
  - 2 = Probably reject. I would argue against this paper.
  - 3 = Possibly accept, but only if others champion it.
  - 4 = Probably accept. I would argue for this paper.
  - 5 = Definitely accept. I would protest strongly if it's not accepted.

  Please base your rating on the paper as it was submitted.
- Reviewer Expertise: Please rate your expertise in the subject area of the paper on a continuous scale from 1 to 3, where:
  - 1=Beginner
  - 2=Knowledgeable
  - 3=Expert.
- Explanation of Rating: Explain your rating by discussing the strengths and weaknesses of the submission, contributions, and the potential impact of the paper. Include suggestions for improvement and publication alternatives, if appropriate. Be thorough. Be fair. Be courteous. Your evaluation will be forwarded to the authors during the rebuttal period.
- Private Comments: You may enter private comments for the papers committee here. These comments will not be sent to the paper author(s). Please do not mention any other papers that are currently in review, or the names of people associated with these papers

Review form for SIGGRAPH

# Radical Inspiration for Final Project?

- Unusual interaction
  - What could you do with touch-based interaction?
  - What could you do with direct object manipulation? (rather than sliders or buttons spatially separate from your data)
- Layers/Levels/Level of Detail/Complexity of your visualization?
- Formal Evaluation
  - How would you measure, in a controlled setting, the effectiveness of your visualization?
- Maybe this will be our next crayon exercise!

# Plan for Today

- Tangentially Related to Privacy…
  Reconstruction from Partial Information
- Readings for Today:
  - "QSplat: A Multiresolution Point Rendering System for Large Meshes"
  - "LabelMe: online image annotation and applications"
- How to brainstorm/foster radical ideas?
- How to formally review a paper for acceptance to a conference/journal
- Homework 7 & Homework 8
- Readings for Tuesday

# Assignment 7: Final Project Ideas

- ½ points of a regular 1 week assignment
- Invent 2 different Final Project Ideas
  - "Who" (audience), "why" (research question), "what" (the finished visualization)
  - One technical challenge for the project. What makes it difficult? What is a potential "risk" for completion? For example:
  - Do you already have a partner?
- Make LMS post by Thursday
- Reply to 3 other students on LMS by Monday
  - Ask a detailed question about the project idea,
  - Suggest a specific dataset source,
  - Suggest a specific visualization toolkit to use for the project,
  - Suggest a reference (paper, book, URL, etc.), or
  - Pose a related research question and/or hypothesis.

## Assignment 8: Volumetric Visualization

- Learn about ParaView
- Learn about VTK
- Learn about Streamlines
- Learn about Volumetric Visualization

- If you already have volumetric data… use it!
- If you don't… procedurally generate some!

- Keep simmering those final project ideas…

## Readings for Tuesday: *(choose one)*

- "Farthest Point Seeding for Efficient Placement of Streamlines", Mebarki, Alliez, & Devillers, IEEE Visualization 2005.

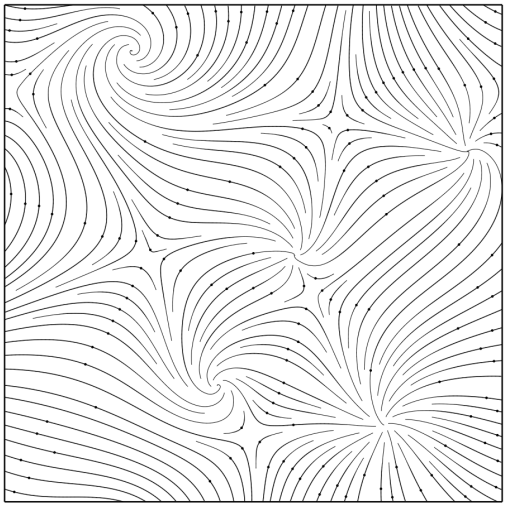

# Readings for Tuesday: *(choose one)*

- "Image Based Flow Visualization",
  Jarke J. van Wij, SIGGRAPH 2002.