Announcements

- Final: Thursday, December 15, 8am, here.
- Review Session, Wednesday, Dec 14, 1pm, AV Williams 4424.
- Review sheet with practice problems on-line.

Hints for Final

- Focus on core techniques/ideas: convolution, gradients, statistical modeling, 3D geometry, motion matrices and optimization methods we studied, and their use in vision.
- Of course, other topics from course may show up.

Vision is inferential



(http://www-bcs.mit.edu/people/adelson/checkershadow_illusion.html)

Boundary Detection

- Images don't come color-coded with boundaries.
 - Filtering to combine evidence.
 - Define discontinuities in 2D.
 - Snakes to combine evidence further.
 - Good continuation, comes from knowledge of shape.
 - Texture understanding as modeling random processes.
- Boundary detection involves combination of different kinds of knowledge.

Modeling + Algorithms

- Build a simple model of the world (eg., flat, uniform intensity).
- Find provably good algorithms.
- Experiment on real world.
- Update model.

Problem: Too often models are simplistic or intractable.

Stereo

- Modeling: Geometry, Photometry.
 - Perspective camera, known pose allows reconstruction.
 - Epipolar Constraint.
 - Ordering constraint
 - Match similar intensities.
 - Disparities similar in nearby pixels.
- Algorithms
 - Straightforward Reconstruction.
 - Shortest path, Graph Algorithms.

Where is Computer Vision Going?

- More Data, Faster Machines =>
- More Interaction with Other Fields.
- Fundamental Problems Remain

Optimization

- Partly push of video meant bigger optimization problems.
- Early 90s, SVD, Gradient Descent, Filtering.
- More recently, Graph Algorithms, Particle Filtering, Mean-shift, Multi-grid....
- Techniques from theoretical CS, applied math, physics, learning,







(Seitz)

State of the art method

Boykov et al., <u>Fast Approximate Energy Minimization via Graph Cuts</u>, International Conference on Computer Vision, September 1999. Ground truth



(Comaniciu and Meer)

Learning

- Recognition using supervised learning.
 - Given examples of an object
 - Use classifiers: eg, SVMs, Winnow, Boosting.
- Grouping using unsupervised learning.
 - Eg., E-M
- Probabilistic Modeling
 - Eg., Graphical models, texture,



Figure 8: Example of frontal upright face images used for training.

Viola and Jones: Real time Face Detection



Viola and Jones: Real time Face Detection

Figure 10: Output of our face detector on a number of test images from the MIT+CMU test set.

Graphics

- Common interest in modeling reflection, light, 3D shape.
- Image-Based Rendering.



input depth image novel view

Szeliski



Figure 1: 3D tracking software developed at Digital Domain was used on nearly every shot of the movie Titanic. (From Vision in Film and Special Effects *Doug Roble*, Digital Domain)



Figure 3: 2D tracking was used to apply digital makeup to Brad Pitt in Interview with a Vampire.



Image Editing: Snakes, Intelligent Scissors, Contour-based editing (Vision-Assisted Image Editing Eric N. Mortensen Brigham Young University)

Biomedical Engineering

- Segmentation
 - Identify organs to measure them.
 - Find tumors.
- Tracking
 - Is a heart beating properly? Is there dead tissue?
- Registration/Matching.
 - Positions of Tumors in Surgery.











Big Data Sets

- Image Data Bases
 - Kodak, commercial data bases w/ tens of thousands of images.
 - Internet, with millions?
- Satellite imagery (Petabytes).
 - Monitoring effects of climate change.
- Custom Data Sets



Vision at Maryland



Background Subtraction



Tracking and Activity Classification











cameras looking outward. (b) A system diagram of an Argus eye.





which is the input to the SfM reconstruction algorithm. Figure 13. Two frames of the original video sequence



Figure 14. Different views of the 3D model after texture mapping.



tom row shows sample frames from one probe video.





Figure 20. Activity-encoded particle filtering.



Figure 24. Reconstruction of a person from eight widely separated views.

Smart Thumbnails





















Many Big Problems Not Answered

- Matching when brightness changes.
- Interaction between motion and lighting
- Single-frame segmentation and multiple images.
- Representing 3D shape, and its appearance in 2D images.
- •